


IBM **@server** BladeCenter Networking Options

Step-by-step configuration
instructions for the BladeCenter
Ethernet Switch Module

Working network configurations to get
your IBM **@server** BladeCenter
solution up and running quickly

Helpful information for rapid
deployment in a Cisco or Nortel
network infrastructure



Seth Atkins
Javier Chang
Jennifer Churchill
David Green
John Parker
Matt Slavin



International Technical Support Organization

IBM @server BladeCenter Networking Options

July 2003

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

First Edition (July 2003)

This edition applies to IBM @serverBladeCenter (product number 8677-1XX).

© Copyright International Business Machines Corporation 2003. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Preface	ix
The team that wrote this Redpaper	ix
Become a published author	xi
Comments welcome	xi
 Chapter 1. Introduction to the IBM @server BladeCenter 4-Port Gb Ethernet Switch	
Module	1
1.1 Introduction to ESM - capabilities, features and functions	2
1.2 Switch management and operating concepts	6
1.2.1 Switch management and control	6
1.2.2 Switch operating concepts	6
1.3 Ports and performance features	8
1.4 Switch and network management	8
1.5 Network cables	9
1.6 Supported network standards	9
 Chapter 2. ESM management and initial configuration	11
2.1 ESM management through the management module	12
2.1.1 Establishing a physical connection to the management module	13
2.1.2 Using the management module Web interface to initially configure ESM	13
2.2 ESM management using the ESM Web interface	17
2.2.1 Basic ESM configuration examples	17
2.2.2 Configuring the switch IP address	23
2.3 ESM management using IBM Director	25
 Chapter 3. Ethernet switching fundamentals	29
3.1 Bridging basics	30
3.2 Switching basics	30
3.2.1 Full duplex versus half duplex	31
3.2.2 Half duplex flow control	31
3.2.3 Full duplex flow control	31
3.2.4 Aging time	31
3.2.5 Frame forwarding modes	32
3.2.6 Virtual LANs	32
 Chapter 4. Supported protocols and standards	35
4.1 802.1D Spanning Tree	36
4.1.1 Spanning tree sequence	37
4.1.2 Switches and network topology changes	38
4.1.3 Setting the parameters that control the spanning tree	40
4.1.4 Summary of the IEEE 802.1d Spanning Tree algorithm	40
4.2 IEEE 802.1Q VLANs	40
4.2.1 The problem with traditional VLANs	41
4.2.2 To tag or not to tag	42
4.2.3 Port VLAN identifiers (PVIDs)	44
4.2.4 VLAN identifiers (VIDs)	45

4.2.5	VLAN membership	45
4.2.6	To egress or not to egress	46
4.2.7	Ingress filtering	47
4.2.8	Summary of VLAN tagging, filtering, and forwarding	48
4.2.9	Sample VLAN configuration	49
4.2.10	Dynamic VLANs and GVRP	51
4.2.11	Priority tagging (IEEE 802.1p)	53
4.3	Link aggregation	57
4.3.1	Static link aggregation	58
4.3.2	802.3ad link aggregation	58
4.3.3	Characteristics common to both static and 802.3ad link aggregation	61
4.4	Link aggregation configuration examples	63
4.4.1	Configuring static link aggregation	66
4.4.2	Configuring IEEE 802.3ad link aggregation	68
Chapter 5.	Deploying the IBM @server BladeCenter in a Cisco environment	73
5.1	Introduction	74
5.2	Architecture summary	74
5.2.1	Datacenter networks introduction	74
5.2.2	Common Cisco components	75
5.3	Guidelines, rules and comments	76
5.3.1	Rules for attaching the eServer BladeCenter to a Cisco infrastructure	77
5.4	Preliminary information on configuration examples	82
5.5	Configuration examples	91
5.5.1	Single ESM, single link to a single Cisco switch	91
5.5.2	Single ESM, single link to two Cisco switches	98
5.5.3	Single ESM, four port LACP aggregation to a single Cisco switch	107
5.5.4	Single ESM, four port static aggregation to a single Cisco switch	121
5.5.5	Single ESM, dual port LACP aggregation to two Cisco switches	134
5.5.6	Dual ESMs, each with a single link to the same Cisco switch	148
5.5.7	Dual ESMs with a single link to two different Cisco switches	149
5.5.8	Dual ESMs each with one link to separate Cisco switches	150
5.5.9	Dual ESMs with four port LACP aggregation to different Cisco switches	160
5.5.10	Dual ESMs with two port LACP aggregation to two Cisco switches	161
5.6	Troubleshooting ESM connections to Cisco devices	178
5.6.1	Troubleshooting specifics	178
Chapter 6.	Deploying the IBM @server BladeCenter in a Nortel environment	181
6.1	Products test specifications	182
6.2	Nortel Networks feature descriptions	182
6.3	Limitations of configuration examples	184
6.4	Preliminary configuration for examples	184
6.4.1	General recommendations	184
6.4.2	Base configuration options common to all examples	185
6.4.3	Basic configuration procedures	185
6.4.4	Base configuration tasks for the ESM	187
6.4.5	Base configuration tasks for BayStack 380-24T	189
6.4.6	Base configuration tasks for Passport 8600	192
6.5	Configuration examples	194
6.5.1	Single ESM with Link Aggregation to Single BayStack 380-24T	194
6.5.2	Validation of ESM configuration	198
6.5.3	Validation of BayStack 380-24T Configuration	198
6.5.4	Single ESM with Link Aggregation to Dual BayStack 380-24T's	200

6.5.5 Validation of ESM configuration	203
6.5.6 Validation of BayStack 380-24T configuration	203
6.5.7 Dual ESMs with Two Port Aggregation to Dual BayStack 380-24Ts	206
6.5.8 Validation of ESM configuration	209
6.5.9 Validation of BayStack 380-24T configuration	209
6.5.10 Dual ESMs with Four Port Aggregation to Dual BayStack 380-24Ts	211
6.5.11 Validation of ESM configuration	215
6.5.12 Validation of BayStack 380-24T configuration	215
6.5.13 Dual ESMs with Four Port Aggregation each to Single Passport 8600	217
6.5.14 Validation of ESM configuration	221
6.5.15 Validation of Passport 8600 Configuration	221
6.5.16 Dual ESMs with Four Port SMLT to Dual Passport 8600s	223
6.5.17 Validation of ESM configuration	229
6.5.18 Validation of Passport 8600 Configuration	229
6.6 Troubleshooting ESM Connections to Nortel Networks Devices	231
Related publications	235
IBM Redbooks	235
Other publications	235
Online resources	235
How to get IBM Redbooks	236
Index	237

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 illustrates 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. 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 IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BladeCenter™
Domino™
eServer™
ibm.com®
IBM®

Lotus®
Netfinity®
Perform™
Redbooks (logo) ™
Redbooks™

@server™
xSeries®
DB2®
Tivoli®

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Nortel Networks, the Nortel Networks Globemark, Unified Networks, and How the world shares ideas are trademarks of Nortel Networks Corporation.

Cisco, Cisco Systems, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries. All other trademarks mentioned in this document are the property of their respective owners.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Preface

This IBM Redpaper will help you install, tailor and configure the new IBM @server BladeCenter and its Ethernet switch module in various network environments.

In this Redpaper, we discuss the IBM BladeCenter 4-Port Gb Ethernet Switch Module. We discuss the features and functions of the Ethernet module and how the switch is managed.

This Redpaper is designed to help you with the initial network configurations to ensure your Ethernet switch module is configured correctly and operable for you to begin immediate communications across the network. Here you will find working configurations that have been configured and tested in our labs.

Also featured in this Redpaper are several configuration examples of the Ethernet switch module being deployed in Cisco and Nortel environments.

The intent of this Redpaper is to introduce networking terms and concepts to the users and administrators of the IBM @server BladeCenter product. Our explanation of these terms and concepts are meant to give the non-networking professional an overall view of the switching and bridging environment and not substitute in-depth training for networking fundamentals. The examples illustrated in this paper are meant to provide real “working” examples of networking configurations where an @server BladeCenter is deployed, but of course many other configurations are possible.


Some knowledge of networking technology and terminology, as well as Intel-based server technology, is assumed.

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.




Seth Atkins is a Systems Engineer in the Enterprise Network Engineering department of Nortel Networks. He provides support to the Sales Engineering organization primarily for products in the BayStack and WLAN portfolios.



Javier Chang is an Advisory Network Specialist at IBM Canada. He has six years of experience in the computer networking field. He holds a Bachelor of Applied Science degree in Electrical Engineering from the University of Toronto. His areas of expertise include Ethernet switching, Internet Protocol (IP) routing and applications, virtual local area networks (VLANs), and Microsoft Windows.



Jennifer Churchill is a Network Specialist at IBM Canada. Jennifer has over five years of networking experience, primarily in the area of network support. Her area of specialization is Campus LAN technologies. She holds a Bachelor of Science degree in Computer Science, and she is a Sniffer Certified Expert.



David Green is a Staff Engineer at IBM in the United States. He has a Bachelor of Science in Information Systems from UNC-Greensboro. His areas of expertise include IBM eServer BladeCenter, Fibre Channel, SANS, networks and Microsoft Cluster Server.



John Parker is a Program Manager; he oversees projects and plans related to the IBM internal global network. He has over 25 years in the computer and networking field. John has written extensively and has published many Redbooks on campus and ATM networking.



Matt Slavin is a Systems Engineer based in Tulsa, Oklahoma; he is employed with the Strategic Alliances group at Cisco Systems. He has been in the computer and networking industry for over 25 years, operating in a number of high-level technical support capacities. Industry certifications include MCSE, MCNE, CCNA and CCIP. Current interests include infrastructure design and support, with a special focus on wireless networking and security.

Thanks to the following people for their contributions to this project:

Rufus Credle, Project XS-5010 BladeCenter and Network Options project leader
IBM Research Triangle Park

Norm Strole, @server BladeCenter Networking
IBM Research Triangle Park

Tim Dougherty, @server xSeries Offerings
IBM Somers

Rob Sauerwalt, Global Brand Manager and Team Lead - IBM @server Marketing
IBM Research Triangle Park

Ishan Sehgal, @server BladeCenter Marketing Manager, Networking
IBM Research Triangle Park

Joe Earhart, Systems Engineer
Cisco Research Triangle Park

Christophe Henrion, Emerging Technologies National Practice. VPN, CDN; Nortel, Network Appliance and Cisco Specialist
IBM Research Triangle Park

Volkert Kreuk, IGS Cisco Networking Specialist
IBM Research Triangle Park

Mark Welch, Appliance Server Testing
IBM Research Triangle Park

Chris Verne, Blade Server Development and Operations Manager
IBM Research Triangle Park

Ted Odgers, IBM Strategic Alliance
Cisco Research Triangle Park

Andrew Wray, Server System Test Software Strategist
IBM Research Triangle Park

Wilson Velez, Development Engineer
IBM Research Triangle Park

Sanjeev Sehgal, IBM Alliance
Nortel Networks

Spencer Shoemaker, Engineer
Nortel Networks Research Triangle Park

Greg Knowles, Director Alliance Solutions
Nortel Networks Research Triangle Park

Tho Nguyen, Enterprise Network Engineering
Nortel Networks

Jeanne Tucker, Tamikia Barrow, Cecilia Bardy, Margaret Ticknor
International Technical Support Organization, Raleigh Center

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

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 papers to be as helpful as possible. Send us your comments about this Redpaper or other Redbooks in one of the following ways:

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

ibm.com/redbooks

- Send your comments in an Internet note to:

redbook@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 662
P.O. Box 12195
Research Triangle Park, NC 27709-2195



Introduction to the IBM @server BladeCenter 4-Port Gb Ethernet Switch Module

This chapter provides an introduction to the IBM® BladeCenter™ 4-Port Ethernet Switch Module. It also gives a brief overview of important network concepts and the different methods for managing the switch.

1.1 Introduction to ESM - capabilities, features and functions

The IBM BladeCenter 4-Port Ethernet Switch Module is a four port gigabit switch module that can be installed in the IBM @server BladeCenter Type 8677. The @server BladeCenter supports up to two Ethernet switch modules (ESM) installed in the chassis at any one time. In the future, the @server BladeCenter will support two additional ESMs for a total of four. Currently, the necessary hardware is not available for the blades to take advantage of the additional switch modules.

The Ethernet switch modules are hot-swappable subsystems that provide Ethernet switching capabilities within a BladeCenter chassis. The primary purpose of the switch module is to provide Ethernet connectivity among the processor blades, management modules, and the external network infrastructure.

The BladeCenter chassis supports a minimum of one switch installed in one of the slots in the rear of the chassis. The blades each ship with two 1 gigabit full-duplex links that correspond to slots 1 and 2 in the rear of the chassis. While the @server BladeCenter will function with a single switch installed, for redundancy, two independent switch modules installed in module slots 1 and 2 of the chassis are required. With this configuration, each of the server blades in the BladeCenter chassis is then able to use either of its network interfaces.

The ESM has 18 ports that can be configured by the user. Ports 1 through 14 on the switch module are gigabit ports that correspond to server blades 1 through 14, respectively. In addition, as shown in Figure 1-1, each switch module has four external 10/100/1000 Mbps Ethernet ports for connection to the external network infrastructure. These ports are identified as Ext1, Ext2, Ext3, and Ext4 in the switch configuration menus and are labeled 1 through 4 (from top to bottom) on the switch module. There are also two hidden ports for dedicated use by the management modules. These ports do not appear in either the Web or telnet management interfaces to the ESM. This prevents changes to the ports that could cause access to the management modules to fail. If this happened, the switch could no longer communicate with the management modules and the management modules would no longer be able to manage the switch. When the switch is first installed in the @server BladeCenter, by default, the switch can *only* be managed through one of these interfaces. The external interfaces are disabled for security reasons. If the ports were enabled, a user who knows the default TCP/IP address and user name for the switch could access the switch and render it inoperable and unmanageable before the switch was configured.




Figure 1-1 IBM BladeCenter 4-Port Ethernet Switch Module

Essential information such as the machine type and serial number are located on the identification label on the side of the ESM. You will need this information when you register the Ethernet switch module with IBM. The media access control (MAC) address also is located on the identification label. See Figure 1-1 for the location of the identification label.

Note: The MAC address is also located on a separate label on the information panel under the external Ethernet port connectors.

The ESM can be managed via telnet or a Web interface. Both the Web and telnet interfaces can be started by accessing the switch directly or starting a session from the management module's Web interface. Each switch by default is assigned a TCP/IP address that corresponds to the module slot the switch is installed in. Table 1-1 lists the module slots and corresponding IP address. Figure 1-2 shows the locations for the slots in the BladeCenter chassis. The default addresses can be changed to addresses on your network.

Table 1-1 Default IP address listing

Bay Number	TCP/IP Address
1	192.168.70.127
2	192.168.70.128
3	192.168.70.129
4	192.168.70.130




Figure 1-2 Bay locations rear of the IBM eServer™ BladeCenter

The ESM is not currently supported in bays 3 and 4 of the BladeCenter chassis. As stated previously, the integrated Ethernet links on the blades use the switches installed in slots 1 and 2. Support for an ESM installed in bays 3 and 4 requires an Ethernet daughter card option to be installed in any blade that uses those switches. However, the daughter card option is planned for the future.

The ESM also contains an information panel with status LEDs on the rear of the switch. The status LEDs include an OK light, error LED and link and activity lights for each of the external ports on the switch. Figure 1-3 shows the locations of the LEDs on the information panel.




Figure 1-3 Information panel

OK (power-on): This green LED is located above the four external 10/100/1000 Mbps ports on the information panel. When this LED is on, it indicates that the switch module has passed the power-on self-test (POST) and is operational.

! (Ethernet switch error): This amber LED is located next to the OK (power-on) LED on the information panel. This LED indicates that the switch module has a fault. If the switch module fails the POST, this fault LED will be lit.

Ethernet link: This green link status LED is located at the top of each external 10/100/1000 Mbps port. When this LED is lit on a port, it indicates that there is a connection (or link) to a device on that port.

Ethernet activity: This green activity LED is located at the bottom of each external 10/100/1000 Mbps port. When this LED blinks on a port, it indicates that data is being received or transmitted (that is, activity is occurring) on that port.

Depending on the application, the external Ethernet interfaces can be configured to meet a variety of requirements for bandwidth or function. The BladeCenter Ethernet switch modules have been pre-configured with default parameter settings that can be used with most typical installations. However, all switch modules will need a few basic parameter settings initially, such as a TCP/IP address for management, security access and control parameters, and basic setup of the external ports for link aggregation.

This high performance Ethernet switch is ideally suited for networking environments that require superior microprocessor performance, efficient memory management, flexibility, and reliable data storage. Performance, reliability, and expansion capabilities were key considerations in the design of the ESM. These design features make it possible for you to customize the system hardware to meet your needs today, while providing flexible expansion capabilities for the future. If you have access to the World Wide Web, you can obtain

up-to-date information about the ESM and other IBM server products on the following Web site:

<http://www.ibm.com/eserver/xseries/>

User guides, drivers and firmware updates can all be found at this site.

1.2 Switch management and operating concepts

This section provides a brief overview of several of the features of the ESM and how the switch is managed. This section also covers some of the concepts of networking. The features and concepts, as well as the management of the switch, are covered in greater detail later in this Redpaper.

1.2.1 Switch management and control

The switch supports two management user interfaces. A Web browser is included on the switch and is the primary interface. It is invoked through the management module Web interface control utility. The other interface is a telnet interface that can also be invoked through the management module control utility. As previously stated, by default the external ports on the switch are disabled. Initial access to the switch must be through the management module to enable the external ports and configure the TCP/IP addresses so the switch can be accessed from the network. Once the network settings are configured, it is no longer necessary to go through the management module control interface. The switch can be accessed either through the management module control utility or directly from the network.

After a TCP/IP address has been assigned to the ESM, you can perform many different management and control tasks. These tasks fall in the following categories:

- ▶ Configuration of switch parameters
 - Switch TCP/IP address
 - Default gateway
 - General switch information: switch location, contact, system name
- ▶ Remote management setup
- ▶ Network monitoring
 - SNMP and traps
 - View port statistics
 - Monitor data traffic
- ▶ Switch maintenance

More information on these tasks and specific instructions on configuring the switch are given later in this Redpaper.

1.2.2 Switch operating concepts

This section is a brief introduction to several of the concepts necessary to understanding how the switch functions. A more in-depth explanation of each of the concepts listed follows later in this Redpaper.

Packet forwarding

The switch stores mapping information from destination MAC addresses to the destination port in a forwarding table. This information is then used to forward packets to specific ports. This capability reduces network congestion and frees the switch from having to forward traffic to all ports.

Spanning Tree Protocol

Spanning Tree Protocol provides a mechanism to block links between switches that form loops within the network. This could be caused by an incorrectly cabled network or a desire by the network designer to have redundant links between switches. If multiple links between switches are detected, Spanning Tree configures one link as the primary and blocks the secondary. In the event of the primary link failing, the secondary link is activated automatically.

VLANs

A virtual local area network (VLAN) is a logical network topology configured on the physical network layout. A VLAN can be used to combine any collection of blade servers in a chassis into a logical network segment. VLANs typically correspond to TCP/IP subnets on the network. VLANs also divide a network into broadcast domains. Any broadcast traffic generated on a VLAN will be kept on that VLAN. This keeps broadcast traffic off the network and, depending on the type of broadcast traffic, within the BladeCenter chassis. VLAN concepts, which will be explained in greater detail later in this paper, include IEEE 802.1Q VLAN, packet forwarding as it pertains to VLANs, ingress and egress ports and VLAN tags.

Tagging and untagging

Another important concept is tagging and untagging of network packets. Tagging is the basis of 802.1Q compliant VLANs. Every port on the ESM can be configured as a tagged or untagged port. Each packet that exits the switch through a tagged port has some VLAN information inserted into the header, if that packet was not previously tagged. The tag information can then be used by other 802.1Q compliant devices to make packet forwarding decisions. The VLAN information identifies the VLAN of the packet. If a port receives a packet and is not on the VLAN contained in the packet header, the switch drops the packet.

Older devices on a network may be tag-unaware and do not conform to the 802.1Q specification. The ESM uses the untagged option for any port that is connected to one of these devices. Any packet leaving an untagged port does not have any tagging information inserted into it. In addition, any tags already in a packet leaving an untagged port are stripped from the packet prior to the switch forwarding the packet on.

The first two sections provided an introduction to the BladeCenter Ethernet switch module and some of the concepts necessary for configuration and management of the switch. The next section provides more detailed technical information on the switch and what technologies it supports.

1.3 Ports and performance features

This section lists the specifications for the ports on the switch. It also lists the performance and operational features of the ESM.

- ▶ Ports
 - Four external copper ports for making 10/100/1000 Mbps connections to backbone, end stations, and servers
 - Fourteen internal full-duplex gigabit ports, one connected to each of the BladeCenter blade servers
 - Two internal full-duplex 10/100 Mbps ports for connection to the management modules. One port connects to each management module
 - The ports can be configured for autosensing and can utilize either straight-through or crossover cables for switch-to-switch connections. The switch will detect the type of cable used and set the port accordingly. However, a cross-over cable is recommended because this cable will work with all of the modes.
- ▶ Performance and operational features of the ESM
 - Transmission method: Store-and-forward.
 - Random-access memory (RAM) buffer: 8 MB.
 - Media access control (MAC) address learning: Automatic update; supports 28K MAC address.
 - Priority queues: Four priority queues per port.
 - Forwarding table age time: Maximum age: 17 to 2100 seconds. Default is 300 seconds.
 - 802.1D Spanning Tree support. Can be disabled on the entire switch or on a per-port basis.
 - 802.1Q Tagged virtual local area network (VLAN) support, including Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP).
 - Support for 256 VLANs in total, including 128 static VLANs.
 - Internet group management protocol (IGMP) snooping support per VLAN.
 - Link aggregation on four external ports for up to two static trunk groups or two link aggregation control protocol (LACP) 802.3ad link aggregation groups.

1.4 Switch and network management

The ESM supports the following network management protocols and standards. Some of these are technologies that can be used to manage and monitor the switch. Others such as Spanning Tree are technologies the switch uses to manage the network. The network technologies listed are covered in more detail later in this paper.

- ▶ Switch monitoring and management
 - Simple network management protocol (SNMP) version 1
 - Fully configurable either in-band or out-of-band control through SNMP based software.
 - Flash memory for software upgrades. This can be done through trivial file transfer protocol (TFTP) or hypertext transfer protocol (HTTP) Web interface.
 - Supports password enabled Web-based management and a telnet remote console.

- Built-in SNMP management:
 - Bridge management information base (MIB) (RFC 1493)
 - MIB-II (RFC 1213)
- ▶ Network management
 - 802.1P/Q MIB (RFC 2674)
 - Interface MIB (RFC 2233)
 - Mini-RMON MIB (RFC 1757) - four groups. The remote monitoring (RMON) specification defines the counters for the receive functions only. However, the switch provides counters for both receive and transmit functions.
 - Spanning Tree Protocol (STP) for creation of alternative backup paths and prevention of network loops.
 - TFTP support
 - Bootstrap protocol (BOOTP) support
 - Dynamic host configuration protocol (DHCP) client support

1.5 Network cables


The following cables and cable lengths are supported by the Ethernet switch module:

- ▶ 10BASE-T:
 - UTP Category 3, 4, 5 (100 meters maximum)
 - 100-ohm STP (100 meters maximum)
- ▶ 100BASE-TX:
 - UTP Category 5 (100 meters maximum)
 - EIA/TIA-568 100-ohm STP (100 meters maximum)
- ▶ 1000BASE-T:
 - UTP Category 5e (100 meters maximum)
 - EIA/TIA-568B 100-ohm STP (100 meters maximum)

1.6 Supported network standards

The following standards are supported by the Ethernet switch module. Some of these standards will be explained in greater detail later in this document.

- ▶ IEEE 802.3 10BASE-T Ethernet
- ▶ IEEE 802.3u 100BASE-TX Fast Ethernet
- ▶ IEEE 802.3z Gigabit Ethernet
- ▶ IEEE 802.1Q Tagged VLAN
- ▶ IEEE 802.1P Priority tagging
- ▶ IEEE 802.3ab 1000BASE-T
- ▶ IEEE 802.3x Full-duplex Flow Control
- ▶ ANSI/IEEE 802.3 NWay auto-negotiation



ESM management and initial configuration

This chapter provides an overview of how to manage the BladeCenter Ethernet switch module (ESM). It introduces the different ways to establish a management session to the ESM, including what type of physical cabling to use. We also document some of the initial switch configuration options in this chapter.

2.1 ESM management through the management module

The BladeCenter management module has one external 10/100 Mbps Ethernet port that we use for management. We can only manage the management module via this external port. Initially, the four external Ethernet ports on the ESM are disabled. Therefore, you must use the external Ethernet port on the management module to configure the switch.

At the time of the writing of this paper, only one management module was supported in the BladeCenter chassis. The support for a second redundant management module is planned for the future. An ESM has one internal 100 Mbps Ethernet connection to each management module. On the management module, this internal connection is labeled Eth1. The internal connection provides the Ethernet connectivity between the management module and the switch modules for management purposes. In order to have IP connectivity between the management module and the ESM, configure the management module external interface Eth0 and the internal interface Eth1 with IP addresses in the same subnet as the ESM.

You have the option of disabling the internal connection between the ESMs and the management module. However, if we disable this internal port, the management module loses its connectivity to the ESMs installed in the BladeCenter chassis, and therefore loses its ability to manage them as well.

Note: By default, the four external Ethernet ports on the ESM are disabled. Therefore, you must initially use the 10/100 Mbps Ethernet port on the management module to configure the switch.

We recommend connecting the external 10/100 Mbps Ethernet port on the management module to a network dedicated to management (shown in Figure 2-1). We also recommend accessing the Ethernet switch modules via the 10/100 Mbps port on the management module. Ideally, the management module, management station, switch modules, and/or DHCP server should connect to the same management subnet.




Figure 2-1 A typical management network

There are two primary reasons for this design:

- ▶ Separating the production network and the management network improves data security.
- ▶ If a problem occurs on the production network, the management module and switch modules remain accessible.

2.1.1 Establishing a physical connection to the management module

The only way to manage the management module is through the external 10/100 Mbps Ethernet port on the front of the module. To establish the physical connection to the management module use one of the following methods:

- ▶ Use a Category 3, 4, 5 or higher unshielded twisted pair (UTP) straight through cable to connect the Ethernet port on the management module to a switch in a network that has an accessible management station.
- ▶ Use a Category 3, 4, 5 or higher cross-over cable to connect a management station (PC, laptop, etc.) directly to the external Ethernet port of the management module.

Note: The 10/100 external Ethernet port on the management module uses a media dependent interface (MDI). This standard describes the interface for unshielded twisted pair (UTP) cable. For two devices to communicate, the transmit wires of one device must connect to the receive wires of the other device. You can accomplish this by using a cross-over cable or a port that implements the cross-over function internally (MDI-X port).

2.1.2 Using the management module Web interface to initially configure ESM

Once you establish the physical connection to the management module, configure the management station with an available IP address in the same subnet as the management module. By default, the subnet is 192.168.70.0/24.

You have two primary methods to manage the management module:

- ▶ HTTP Web interface
- ▶ IBM Director

We use the management module Web interface to demonstrate the initial switch configuration.

Note: The Web browser you use must support Java, JavaScript 1.2, and a minimum screen resolution of 800x600 with 256 colors.

Follow these steps to establish a management session with the management module, and to configure the initial recommended switch settings:

1. Open a Web browser and connect to the management module using the configured IP address. The default IP address for the management module external interface is 192.168.70.125. Please note that the default IP address for the internal interface is 192.168.70.126.
2. Enter the userid and password. The default is USERID and PASSWORD (it is case-sensitive with a zero in the place of the letter O). Click **OK**.
3. At the initial screen, click **Continue** to access the management session.

Configuring the ESM IP address

When you first install an ESM in the BladeCenter unit, you have to use the management module to configure some basic settings. For more advanced configuration features use an ESM management session. Before you can establish a session to the ESM, you must configure the ESM IP address information. Once you configure the switch module IP address, you can initiate a management session to the ESM through the management module. To configure an IP address, subnet mask, and default gateway on an ESM follow these steps:

1. From the management module menu, select **Switch Tasks -> Management**. You will see a window similar to Figure 2-2.
2. Select **Switch Module 1**.

Note: By default, the ESMs are shipped with default IP addresses. The default IP address for the ESMs in slot 1 and slot 2 are 10.90.90.91 and 10.90.90.92, respectively. The IP address configured in the management module will take precedence.

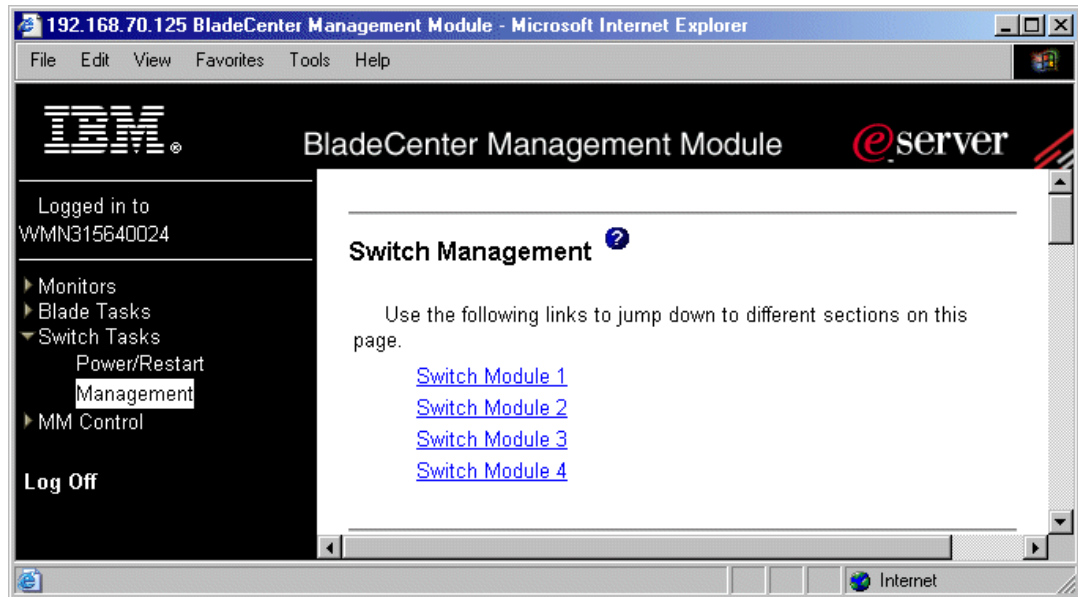


Figure 2-2 Switch management

3.As shown in Figure 2-3, complete the IP Address, Subnet Mask and Default Gateway fields.

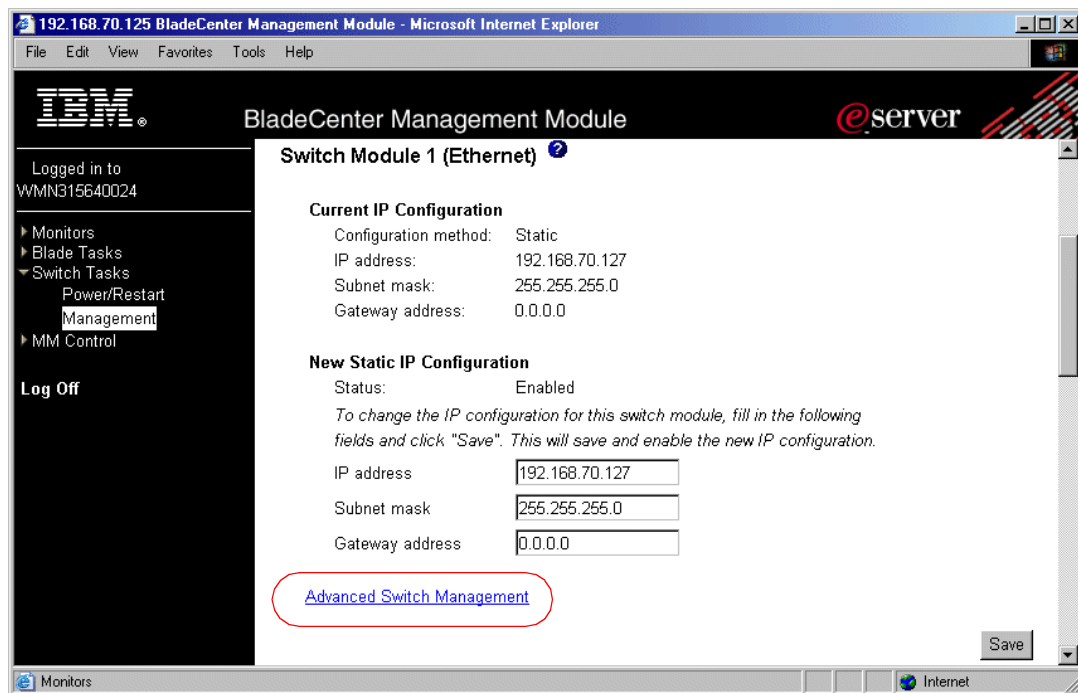


Figure 2-3 Switch module IP configuration

4. Click **Save**. The Current IP Configuration section of the panel displays the new IP address..

Note: The management module has default addresses for the switch modules; these are:

- ▶ Switch module bay 1: 192.168.70.127
- ▶ Switch module bay 2: 192.168.70.128

5. To verify the IP connectivity between the management module and the ESM, click **Advanced Switch Management** (Figure 2-3). Then click **Send Ping Requests**. Click **Cancel** to return to the previous menu.

Enabling external Ethernet ports

By default, the four external ports on the ESM are disabled. Enable the external ports on the ESM to allow Ethernet connectivity to the external infrastructure. To enable the external ports, follow these steps:

1. From the Advanced Switch Management panel, click **Advanced Setup**.
2. Select **Enabled** from the list box for the option to enable External ports (Figure 2-4).

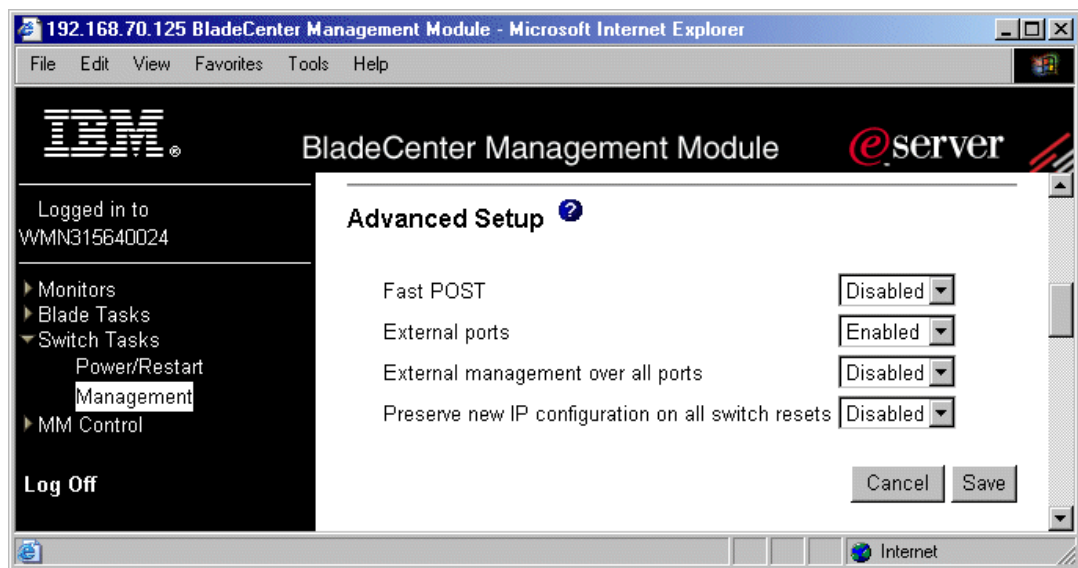


Figure 2-4 Switch Advanced Setup

3. Click **Save**.

Important: The @server BladeCenter ships with the external ESM ports disabled. You must use the management module to enable them.

Enabling external management

We recommend that you manage the ESMs through the management module; however, if you want in-band management access to the four external ports on the ESM or if you configure the ESM to send SNMP traps, then you should enable external management over all external ports. If external management is not enabled, you must establish an ESM management session through the management module.

1. From the Advanced Setup window (Figure 2-4), select **Enabled** from the External management over all ports list box.
2. Click **Save**.

Enabling option preserve IP address

In order to ensure that the ESM IP address configured in the management module is saved even after a switch factory reset, you need to enable this option.

1. From the Advanced Setup window, select **Enabled** from the Preserve new IP configuration list box.
2. Click **Save**.

Note: Repeat the ESM configuration process for switch module two, if a second switch is installed.

2.2 ESM management using the ESM Web interface

After you complete the initial configuration for the ESM through the management module, establish a management session to the ESM. Use the ESM management interface for all of the other switch configuration. The following chapters of this paper document some of the advanced configuration options.

You can manage the ESM by using any one of the following methods:

- ▶ BladeCenter management module Web interface
- ▶ IBM Director
- ▶ ESM Web interface
- ▶ ESM telnet interface

The management module Web interface is used to initially configure the ESM or to launch a Web or telnet management session to the ESM. We will now use the ESM Web interface to demonstrate the additional basic switch configuration examples.

Important: After you make changes to the ESM configuration, save the configuration to the nonvolatile random access memory (NVRAM) of the ESM. Select **Maintenance -> Save Changes**. Click **Save Configuration**.

2.2.1 Basic ESM configuration examples

This section provides examples of some of the basic recommended switch configuration. Use the following steps to establish a management session to the ESM:

1. From the management module Web interface, select **Switch Tasks -> Management**.
2. Select the appropriate ESM.
3. Click **Advanced Switch Management**.
4. Click **Start Telnet/Web Session**. You will see a window similar to Figure 2-5.

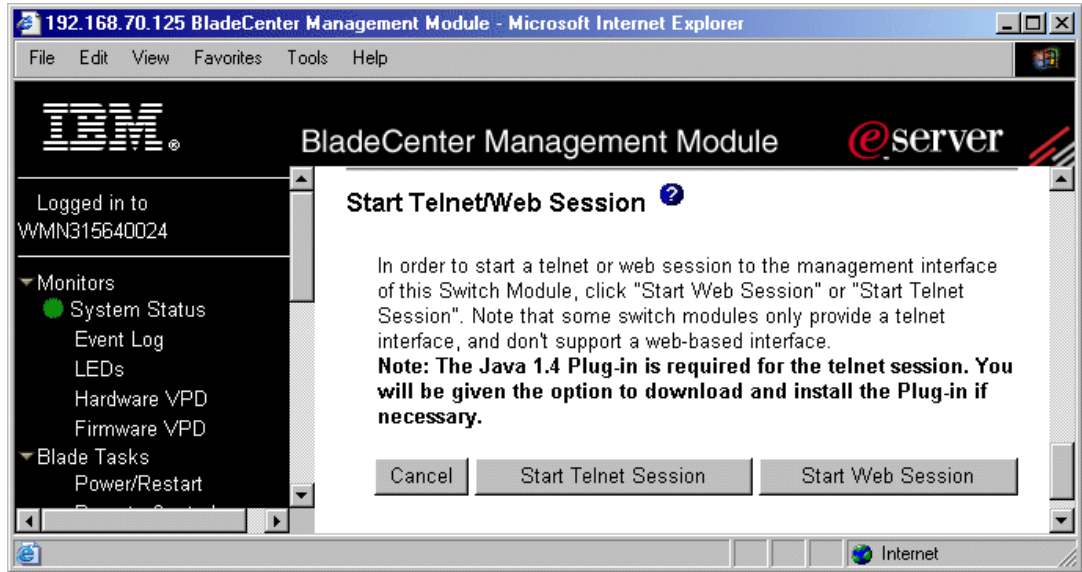


Figure 2-5 Starting a Web session

5. Click **Start Web Session**.
6. Enter the userid and password for the ESM. The default is USERID and PASSWORD (it is case-sensitive with a zero in the place of the letter O). Click **OK**.

Note: We recommend that you manage the ESM through the Ethernet port of the management module, however, if desired, you can connect directly to the ESM if you enable external management over all ports via the management module Web interface. Open a Web browser and point the browser to the IP address of the ESM to access the ESM Web interface directly. See the *ESM User's Guide* for further information.

Configuring system information

The Switch Information panel displays the MAC address of the switch, as well as the firmware and hardware versions. Use the following steps to configure the system and contact information:

1. From the ESM Web interface, select **Configuration -> Switch Information**.
2. Complete the System Name, System Location, and System Contact fields. You will see a window similar to Figure 2-6.

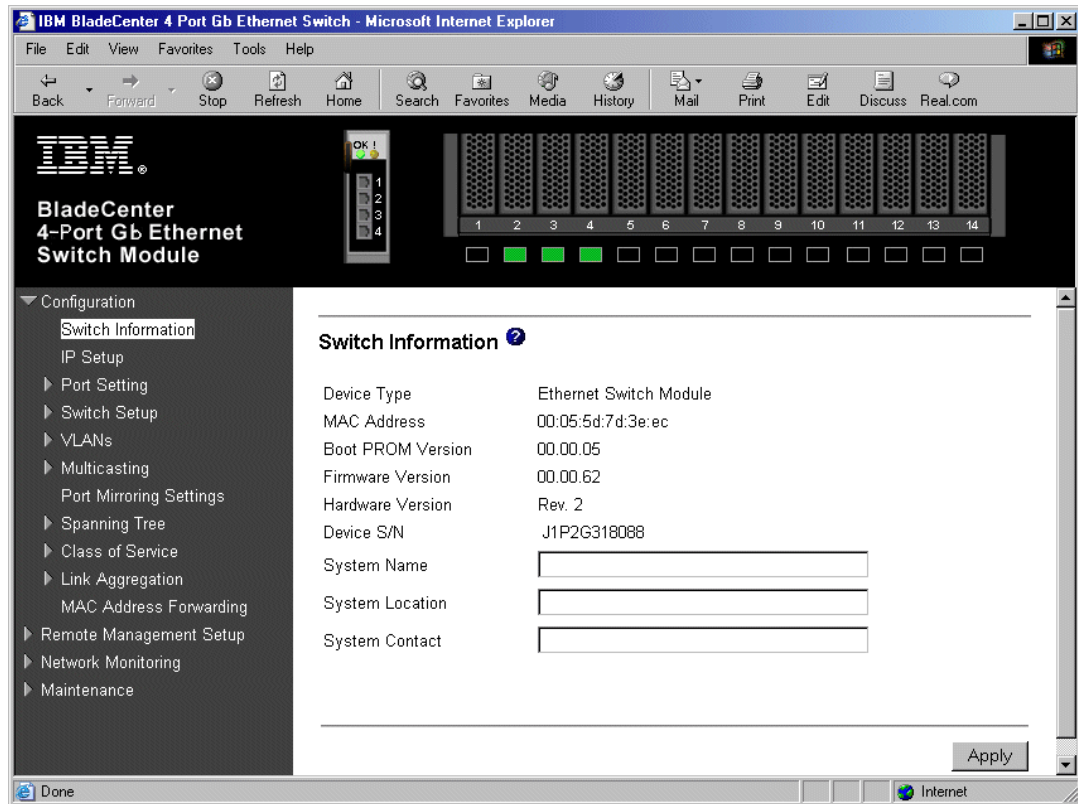


Figure 2-6 Switch information

3. Click **Apply**.

Configuring user accounts

For security reasons, we recommend that you create a new administrator account and remove the default account.

1. In Figure 2-7, select **Remote Management Setup -> Setup User Accounts**.
2. Click **Add**.

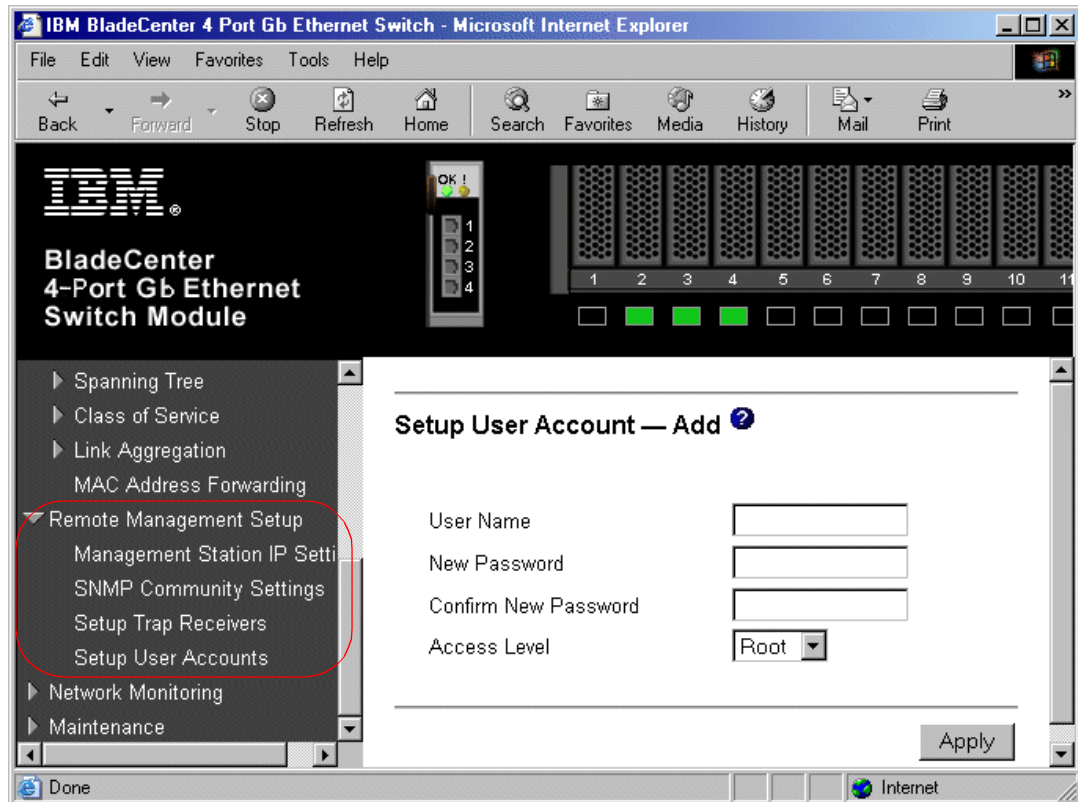


Figure 2-7 Setup user accounts

3. Complete the User Name, New Password, and Confirm New Password fields. Select **Root** from the Access Level list box. Click **Apply**.
4. To remove the default user name, select the user name **USERID** from the Setup User Accounts window. Click **Remove**. Click **OK** when prompted to delete the entry.

Use this panel to create and manage all user accounts. You can choose from three access levels: User, User+, and Root. The access level Root has full administrative access. The User option has read-only privileges, and the User+ option has the same access as User, except that it also has the ability to restart the switch.

Upgrading firmware

There are two methods to upgrade the firmware on the ESM. You can use the browser transfer or a TFTP transfer.

Important: Save any configuration changes to the NVRAM of the ESM.

Select **Maintenance -> Save Changes**. Click **Save Configuration**. If you do not complete this step before the firmware upgrade, the switch resets and it loses any unsaved changes.

Use the following steps to upgrade the firmware on the ESM via the Web browser:

1. Select **Maintenance -> Using Browser**.
2. Click **Upgrade Firmware/Configuration File**. You will see a window similar to Figure 2-8.

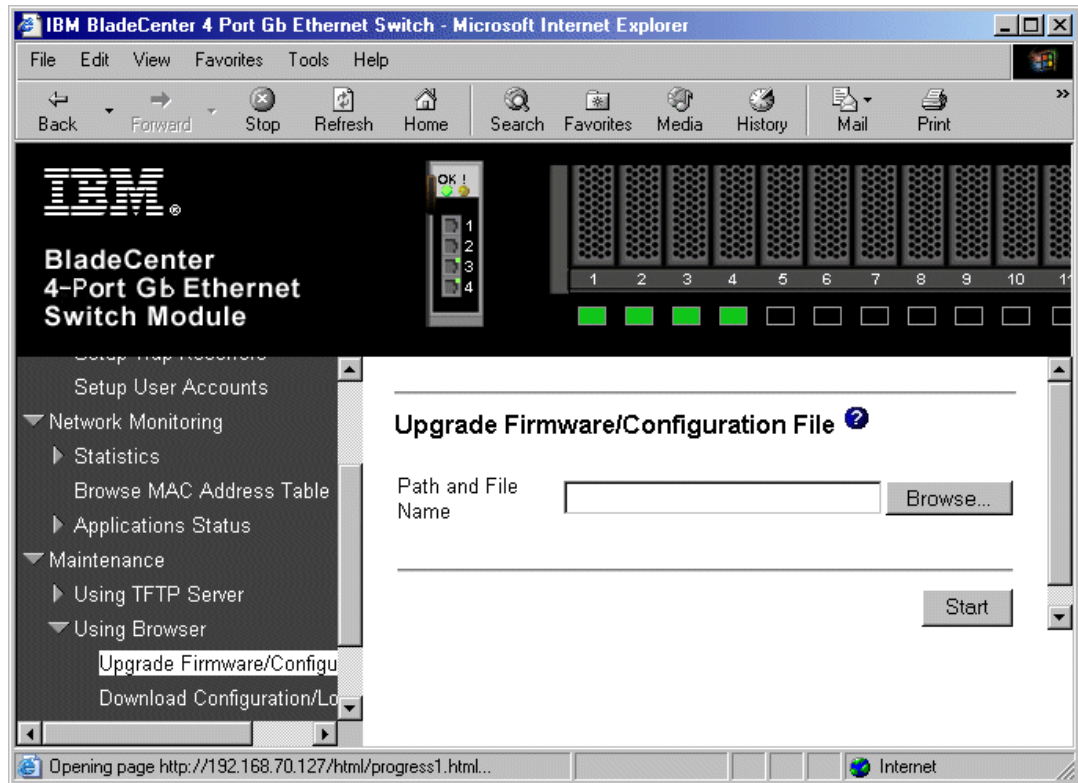


Figure 2-8 Upgrading firmware

3. Complete the Path and Filename field with the name and location of the firmware file.
4. Click **Start**. This process transfers the firmware file to the ESM, programs the NVRAM and then the switch reboots.
5. Verify the new firmware level by selecting **Configuration -> Switch Information**. This panel displays the current firmware version.

To upgrade the firmware using a TFTP, you have to have an accessible TFTP server running on the network or on your management station. You must start the TFTP server and configure it to point to the location of the firmware file.

1. Select **Maintenance -> Using TFTP Server**.
2. Click **Upgrade Firmware/Configuration File**.

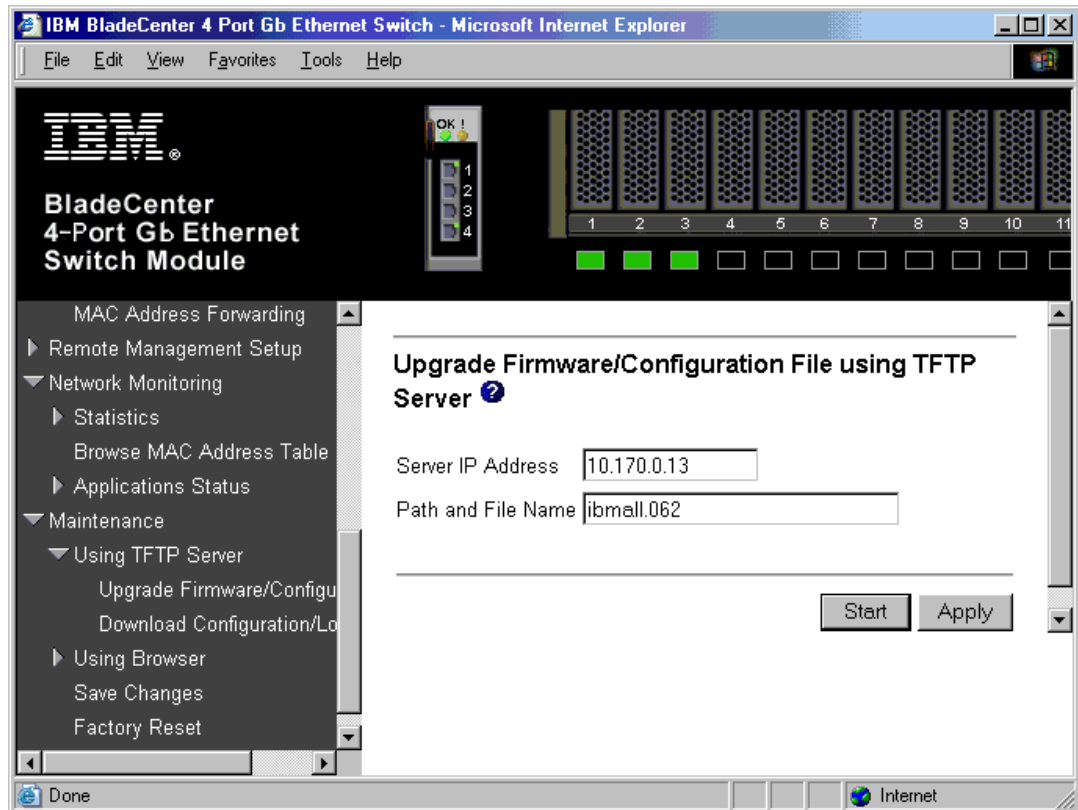


Figure 2-9 Upgrade Firmware/Configuration File using TFTP

3. Complete the TFTP Server IP address field with the IP address of the TFTP server where the firmware file is located. Complete the Filename field with the name and location of the firmware file.
4. Click **Start**. Click **OK** at the information screen. This process transfers the firmware file to the ESM, programs the NVRAM and then the switch reboots.
5. Verify the new firmware level by selecting **Configuration -> Switch Information**. This panel displays the current firmware version.

Downloading the configuration file and log file

You can download a copy of the ESM configuration and log files to your management station.

1. Select **Maintenance -> Using Browser**.
2. Click **Download Configuration/Log File**.

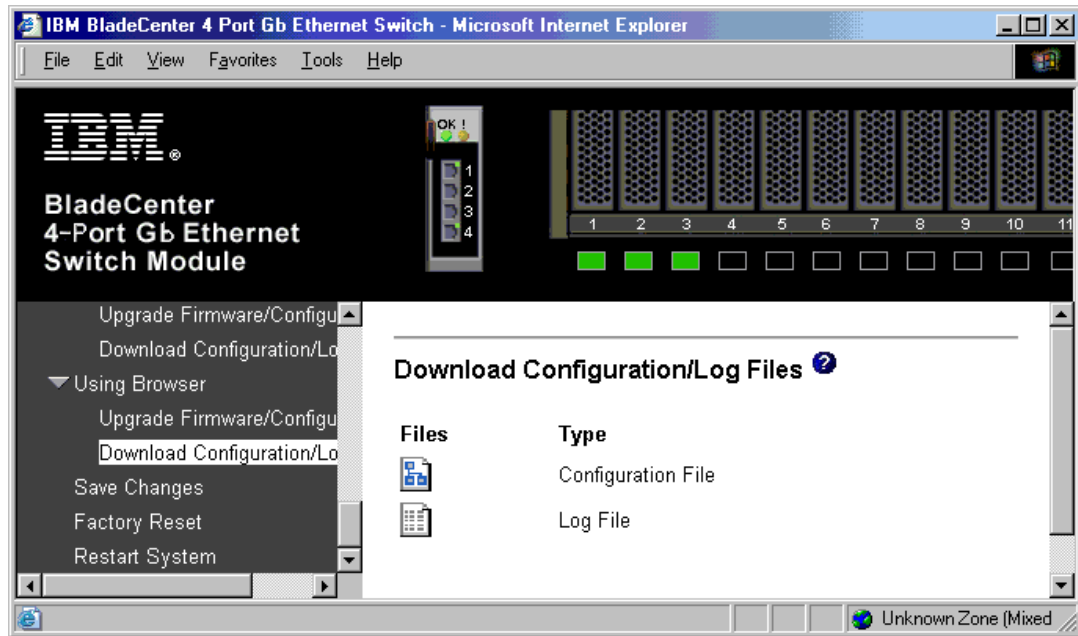


Figure 2-10 Download Configuration/Log Files

3. Click the configuration file icon under the Files heading. When the File download window appears, click **Save**.
4. Click **Close** in the Download Complete message window.

2.2.2 Configuring the switch IP address

You can change the ESM IP address using an ESM management session, but we recommend that you make the IP change in the management module as well. The primary reasons for this recommendation are as follows.

- ▶ If you connect to the ESM through the 10/100 external management module port, you lose your session and you cannot reconnect to the switch after the IP address change. The management module must update its NVRAM before it can reconnect to the switch. The management module performs this update at set intervals, so it can take up to ten minutes. In order to avoid this situation, you must also change the IP address of the switch using a management session to the management module. This immediately updates the management module NVRAM and you can reconnect to the ESM Web interface using the new IP address.
- ▶ If you change the IP address via an in-band connection to the ESM, you can reconnect to the new IP address shortly after the change. However, to ensure that the management module saves the correct switch IP address in its NVRAM, we recommend that the IP address be changed via the management module as well.
- ▶ If you change the ESM IP address via the management module Web interface first, you do not need to change it via the switch module. It is only necessary to change it in both interfaces, if the change is first made in the ESM Web interface.

The ESM has three options for IP configuration:

- ▶ To obtain an IP address from a DHCP server
- ▶ To obtain an IP address from a BOOTP server
- ▶ To manually configure the IP address

Assigning an IP address using DHCP or BOOTP

The primary reasons for using DHCP or BOOTP for assigning IP addresses to the management modules and switch modules is to avoid IP conflicts when setting up an environment with multiple BladeCenters, and to use some of the features of IBM Director. Since all of the BladeCenter management modules and switch modules use the same default IP addresses, IP conflicts can occur if you connect all devices to the same network without first assigning each device an unique IP address.

In order to use the DHCP or BOOTP option for assigning an IP address to an ESM, ensure that there is an accessible DHCP/BOOTP server on the management network. Create a DHCP reservation for each switch module. The MAC address of a device is entered into the DHCP reservation, therefore an ESM is always assigned the same reserved IP address.

To configure the ESM to obtain an IP address from a DHCP or BOOTP server:

1. Select **Configuration -> IP Setup**. In the New Switch IP Settings section of the IP Setup window, select **DHCP** or **BOOTP** from the Get IP from list box.
2. Click **Apply**. Click **OK** in the information windows.
3. You will lose the current management session. Reconnect to the ESM using the new IP address. Once you re-establish your session, save the IP address changes to NVRAM by selecting **Maintenance -> Save Changes**. Click **Save Configuration**.

Assigning a static IP address

In many production environments, many network administrators statically configure the IP addresses in their management network. By assigning static IP addresses to these critical devices, you do not need to rely on the availability of a DHCP or BOOTP server.

Note: The ESM IP address must be in the same subnet as the management module or connection via the management module will be lost.

Use the following steps to manually configure the switch IP address:

1. Click **Configuration -> IP Setup**. You will see a window similar to Figure 2-11.

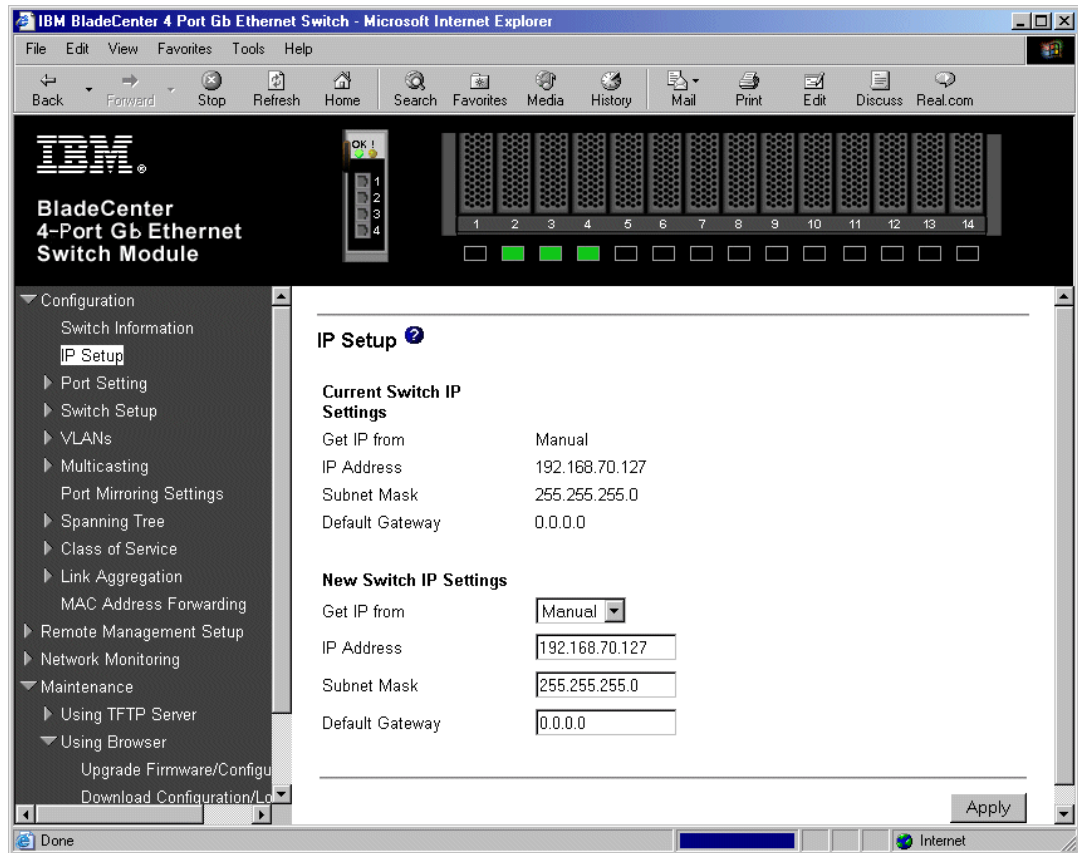


Figure 2-11 IP Setup window

2. In the New Switch IP Settings section of the IP Setup window, select **Manual** from the Get IP list box. Next, complete the IP Address, Subnet Mask, and Default Gateway fields. Click **Apply**.
3. You will lose your management session to the switch, therefore reconnect using the new IP address. If you are using a browser, simply point your browser to the new IP address. Once you re-establish your session, save the IP address changes to NVRAM by selecting **Maintenance -> Save Changes**. Click **Save Configuration**.
4. See “Configuring the ESM IP address” on page 14 for instructions on changing the ESM IP address in the management module.

2.3 ESM management using IBM Director

You can also manage the ESM by using IBM Director. After the IBM Director detects the @server BladeCenter, you can manage the chassis, blades and ESM within the BladeCenter system. You can also view the vital product data (VPD) on the switch and view current configuration information on the internal and external ports.

To manage the switch:

1. Right-click the BladeCenter chassis listed in the IBM Director.
2. Select **BladeCenter Assistant**.
3. Click **BladeCenter Management**.

4. At the bottom of the left-hand frame of the Management Processor Assistant, click **Management** to open the window shown in Figure 2-12. From here, you can reset the ESM to defaults, reboot the switch and manage the external ports.
5. Make any desired changes.
6. Click **Apply**.

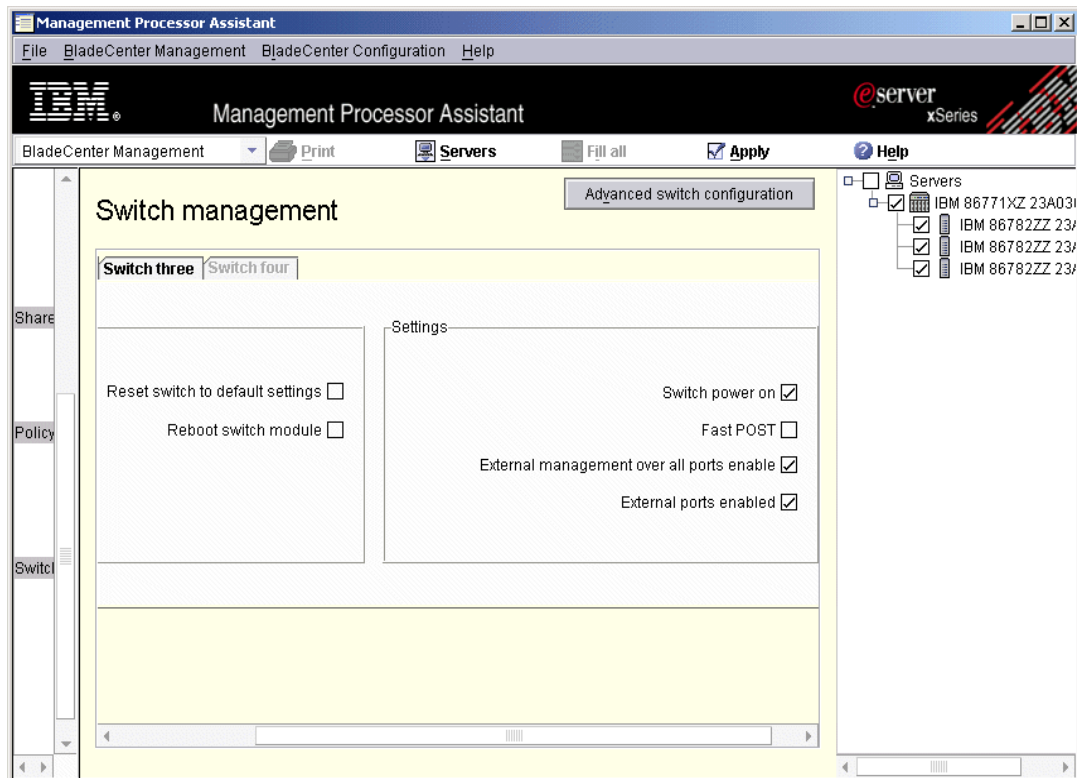


Figure 2-12 Switch management in IBM Director

Note: For full control over the switch, click the **Advanced Switch Configuration** button on the upper right of the middle frame. This will launch the Web interface to the switch described in 2.2, “ESM management using the ESM Web interface” on page 17.

To configure the switch IP address:

1. Right-click the BladeCenter chassis listed in the IBM Director.
2. Select **BladeCenter Assistant**.
3. Click **BladeCenter Management**.
4. At the bottom of the left-hand frame of the Management Processor Assistant, click **IP Configuration** to open the window shown in Figure 2-13.
5. Edit the IP address settings to the desired address.
6. Click **Apply**.

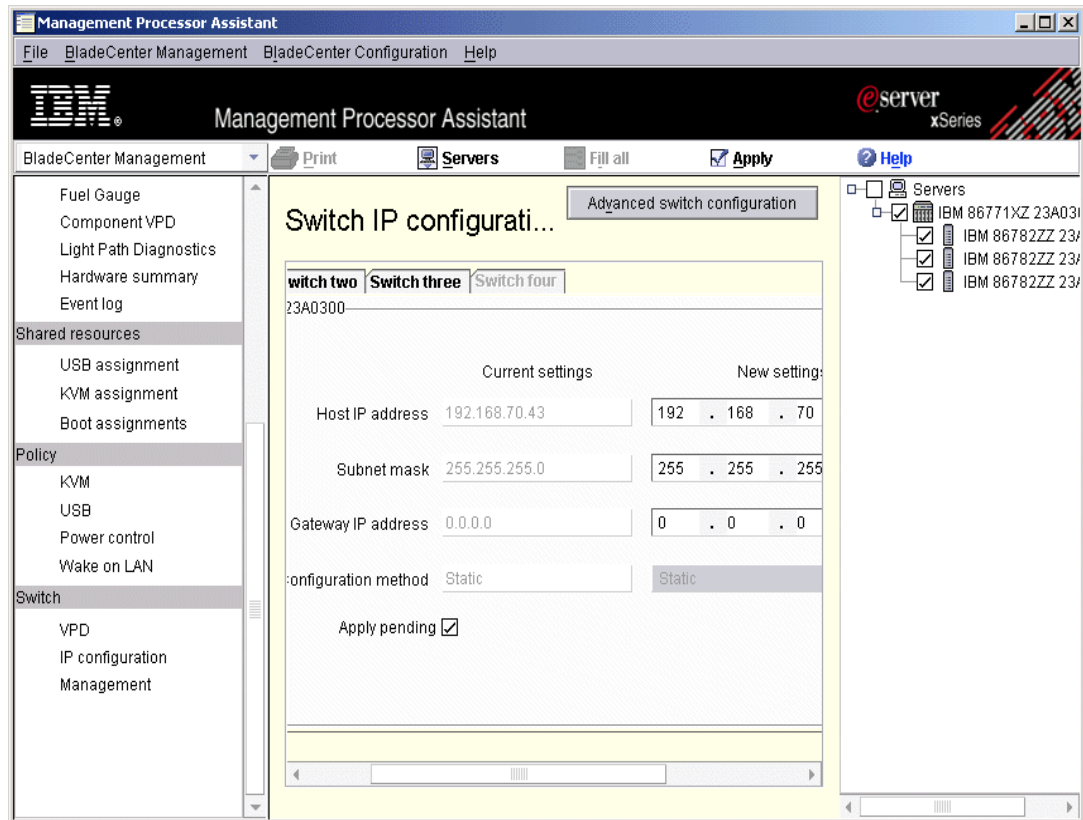


Figure 2-13 IP Setup in IBM Director

To view the VPD information on the switch:

1. Right-click the BladeCenter chassis listed in Director.
2. Select **BladeCenter Assistant**.
3. Click **BladeCenter Management**.
4. At the bottom of the left-hand frame of the Management Processor Assistant, click **VPD** to open the window shown in Figure 2-14.

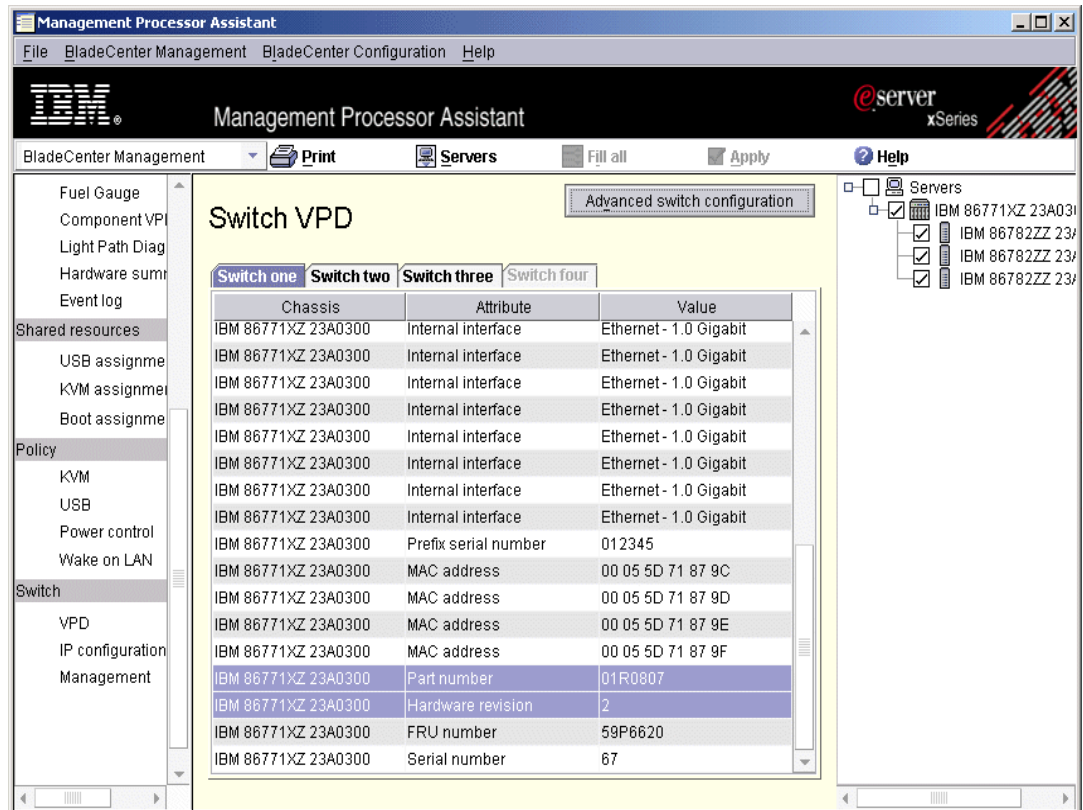


Figure 2-14 Switch VPD in Netfinity® Director



Ethernet switching fundamentals

In this chapter, we introduce the concepts of bridging and switching to learn how the Ethernet Switch Module (ESM) operates within the BladeCenter. We also explain certain switch parameters which affect the behavior of the ESM. Before reading this section, you should have a basic understanding of the Ethernet protocol and its vocabulary (for example, MAC address, frame, shared media, collision, segment, CRC error).

Important: The intent of this Redpaper is to introduce networking terms and concepts to the users and administrators of the IBM @server BladeCenter product. Our explanation of these terms and concepts are meant to give the non-networking professional an overall view of the switching and bridging environment and not substitute indepth training for networking fundamentals. The examples illustrated in this paper are meant to provide real “working” examples of networking configurations where an @server BladeCenter is deployed, but of course many other configurations are possible.

3.1 Bridging basics

In order to discuss Ethernet switching, first we need to know how bridging works. An Ethernet bridge is a two-port device that forwards traffic by analyzing the MAC addresses of frames passing through it. A bridge builds and maintains an internal *forwarding table* by inspecting the source MAC address of the frames entering its ports. The contents of the forwarding table determine whether frames entering one port should be forwarded to the other port, or discarded.

When a bridge powers on, the forwarding table is empty initially. After the first frame from the network arrives, the bridge *learns* the MAC address of the *source* machine and its location (for example, what port it is connected.) The bridge stores these two pieces of information in the table and searches for the MAC address of the *destination* machine. Because the destination machine has not transmitted any frames yet, the bridge does not know where it is. Therefore, the bridge sends the frame to the other port. The process of populating a bridge forwarding table is called *source learning*.

As other machines transmit frames, the bridge learns more source MAC addresses and their associated port location. With this information, the bridge will discard a frame (for example, not forward it to the other port) if its forwarding table indicates that the destination MAC address is on the same port as the source MAC address. By making decisions to forward or discard frames, a bridge acts as a filter to limit the amount of network traffic passed from one port to the other port.

Isolating traffic into two segments also isolates collisions in the shared media. When two or more machines in one segment transmit simultaneously, a collision occurs. After a collision event, all machines in the segment must back off (for example, stop transmitting) for a random period of time. As more devices connect to one segment of shared media, there is a greater probability of multiple collisions occurring. Any device that detects contiguous collisions must increase its backoff time exponentially.

Therefore, when a bridge splits one segment into two smaller segments, two benefits emerge. Firstly, by filtering data frames, it reduces the probability of a collision from taking place. Secondly, by isolating collision traffic, a bridge prevents collisions in one segment from disrupting communication in the other segment. In other words, Ethernet bridges separate the media into two *collision domains* to improve traffic flow.

Bridges have limitations, however. Since they typically contain two ports, bridged Ethernet networks are not very scaleable. For example, to create three collision domains, we require two bridges. Four collision domains require three bridges. Within each collision domain, multiple devices still need to share the Ethernet media. In the next section, we discuss how switches address these problems.

3.2 Switching basics

Ethernet switches are functionally equivalent to bridges. However, switches have more than two ports and contain special hardware, such as application-specific integrated circuits (ASICs), to provide much higher performance than bridges. The ability to connect to several segments makes a switch much more versatile than a bridge. We can create many small collision domains using one switch, further improving on the benefits of bridging. In fact, by connecting only one device to each switch port, we virtually eliminate the drawbacks of shared media. Every collision domain now contains only two transmitters:

1. The attached device which could be a server, workstation, printer, another switch, etc.
2. The switch itself, which forwards frames from other ports.

3.2.1 Full duplex versus half duplex

Although these new collision domains are tiny, we still consider them as shared segments because both devices cannot transmit at the same time. Otherwise, collisions occur and both parties must back off for a random time period. To eliminate the shared media altogether, we introduce the concept of a *full-duplex* Ethernet switch connection. A full-duplex connection allows both devices to transmit simultaneously with no collisions. In the BladeCenter chassis, all internal server-to-ESM connections are full-duplex, whereas the four external ports can be either full-duplex or half-duplex (for example, shared media) at either 10 or 100 Mb/s. Ports that are configured for 1000 Mb/s operate in full-duplex mode only.

3.2.2 Half duplex flow control

Full-duplex has one disadvantage compared to shared media: there is no inherent flow control mechanism. The Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocol defines how Ethernet devices should share the medium. Collisions are a natural by-product of CSMA/CD and implicitly control the flow of traffic. Recall that bridges (and half-duplex switches) create separate collision domains. Because a bridge, by design, does not forward collisions, it needs another way to limit the data flow into its ingress (inbound) port if its egress (outbound) port is congested. A bridge generates *backpressure*, which is a special repeating pattern of bits, on its ingress port to force collisions and thereby cause the transmitter to back off. Without backpressure, a bridge must discard valid frames if its buffer is full. In the BladeCenter ESM, an external port will automatically use backpressure if configured for half-duplex operation.

3.2.3 Full duplex flow control

For a full-duplex connection to have flow control, both devices must support a common protocol, such as IEEE 802.3x. In this scheme, a congested device sends a special MAC Control PAUSE frame to instruct the other machine to stop transmission for a specific period of time. The BladeCenter ESM supports 802.3x, which is automatically enabled on all internal and external full-duplex ports. Since the blade servers also support 802.3x, we know that full-duplex flow control exists on internal connections. However, any external device connected to the ESM must also support 802.3x to ensure that PAUSE frames are generated and interpreted correctly. For example, if the ESM has a full-duplex link to another switch which does not use 802.3x, then either the ESM or the external switch will discard frames during congestion.

3.2.4 Aging time

The switch forwarding table, like the one in a bridge, must contain up-to-date information about MAC addresses and their port locations. Without accurate entries, a switch might forward frames to the wrong port, to all ports, or not forward them at all. Source learning by itself is useful for adding new MAC addresses to the forwarding table only when devices transmit frames. If a machine is disconnected from the network or moved to another location, the switch needs to remove the old information. The key to this process is the *aging time*, which determines how long a MAC address is stored in the forwarding table after the device stops transmitting frames. The default value for the ESM is 300 seconds. In other words, if the switch does not detect any frames with a specific source MAC address after five minutes, it deletes the MAC address entry from the forwarding table. This will force the switch to relearn the source MAC address when the machine transmits another frame.

Increasing the aging time allows MAC addresses to stay in the table longer. This is beneficial for devices that rarely move around in the network or those that do not transmit frames within every 300 seconds. For example, some printers are very “quiet”; their primary task is to receive data for printing. The disadvantage of a long aging time is the potential for incorrect switch forwarding decisions. The switch will forward frames to the wrong port if a machine is relocated to another part of the network, but does not transmit any frames to update the switch.

Decreasing the aging time is useful when devices connect to different parts of the network on a regular basis (for example, laptops). If the switch receives a frame with a destination MAC address not stored in its forwarding table, then it *floods* the frame out all ports (except the port where the frame originated from.) This enables the relocated device to receive traffic even before it has transmitted frames. When the relocated device does transmit a frame, the switch relearns its MAC address, so flooding is not necessary. The disadvantage of a short aging time is excessive flooding. Because the forwarding table loses information too quickly, the switch frequently must send frames to all ports, similar to the operation of an Ethernet hub. This defeats the purpose of switching.

The ESM also supports static entries in the forwarding table. The aging time has no effect on static MAC addresses. In rare instances, we might want to statically configure a MAC address and map it to a port. This is primarily for security reasons, and ensures that a machine can receive traffic on the assigned port only. If someone moves the machine to another port, it cannot communicate through the switch.

3.2.5 Frame forwarding modes

Switches generally operate in one of two modes: *cut-through* or *store-and-forward*. In cut-through mode, a switch immediately forwards a frame and begins transmission after receiving the destination MAC address. This occurs while the original frame is still being received by the ingress port. The switch cannot discard the frame even if it has a CRC error, since most of the frame has already been transmitted to the destination. Cut-through operation does not work across ports that have different data rates. For example, a frame received at 10 Mb/s on one port cannot be forwarded in cut-through mode to a port at 100 Mb/s.

To forward traffic between different data rates, a switch must buffer the entire frame first. This is the store-and-forward mode of operation. Since the switch reads a complete frame before transmission, it checks for a CRC error and discards the frame if it detects an error. Store-and-forward introduces latency (delay) into the switching process, but it prevents any errors from propagating to other ports. Also, newer standards (such as IEEE 802.1Q) require that switches insert additional fields into Ethernet frames and recalculate the frame check sequence (FCS) to prevent false CRC errors. Because of these reasons, the BladeCenter ESM operates in store-and-forward mode only.

3.2.6 Virtual LANs

A virtual local area network (VLAN) is a logical association of network devices based on certain policies or rules defined in a switch. VLANs provide additional flexibility in designing networks. We are no longer constrained by the physical location of machines. To understand VLANs, we need to discuss how broadcast traffic affects a switch.

Recall that a switch reads the destination MAC address of a frame to make a forwarding decision. If the destination address is a known MAC address in the forwarding table, then the switch transmits the frame on the corresponding port. If the MAC address is not in the table, flooding occurs (see 3.2.4, “Aging time” on page 31). Frames containing a destination

address of an individual device are called *unicast* frames. In contrast, *broadcast* frames contain a special destination address that all devices recognize. Every device, upon receiving a frame with a broadcast destination address, must process the frame whether or not it was intended for that device.

When a switch receives a broadcast frame from one of its ports, it transmits a copy of the frame to all other active ports. Broadcast traffic is necessary for the operation of certain protocols and applications. Unfortunately, when too much broadcast traffic exists in the network, performance degrades. As the number of broadcast frames increases, machines must spend more CPU time to process them. To combat this problem, we create a VLAN to isolate broadcasts to a subset of ports in the switch.

Whereas switches and bridges keep Ethernet collisions within domains, VLANs isolate traffic (both unicast and broadcast) within *broadcast domains*. In the BladeCenter ESM, all ports initially belong to the same VLAN, also known as the default VLAN. Once we create a new VLAN with ports assigned to it, the ESM prevents traffic in one VLAN from interfering with traffic in the other VLAN. This is equivalent of having two physical networks with no connection between them. For example, a blade server in one VLAN cannot communicate with a blade server in another VLAN unless an external routing device connects the two VLANs together.



Supported protocols and standards

In this chapter, we discuss the supported protocols and standards for the IBM BladeCenter 4-Port Gb Ethernet Switch Module.

4.1 802.1D Spanning Tree

The spanning tree algorithm enables switches and transparent bridges to dynamically discover a loop-free network (tree) and provide a single physical path between any two stations attached to the network (spanning). Although this standard applies to both transparent bridges and ethernet switches, we will be referencing Ethernet switches in this section.

The IEEE 802.1d standard defines how the spanning tree will work in a switching environment.

If there were loops in the network topology, there would be network configurations where:

- Frames endlessly circulate within the network
- Communication is prevented because a bridge may incorrectly assert that a source and destination address are on the same side of the bridge

If there is more than one switch between any two interconnected LANs, the spanning tree algorithm will ensure that only one of them will be passing traffic between the two LAN segments, this switch will be in a forwarding state. All the other switches that are parallel to the forwarding switch will continue to participate in the spanning tree, and in the event of a failure be able to forward frames, but will not perform any frame forwarding while the root switch is active. These switches are said to be in blocking state.

Note: In the case of multiport switches, a single switch may be in forwarding state on some of its ports while it is in blocking state on the others.

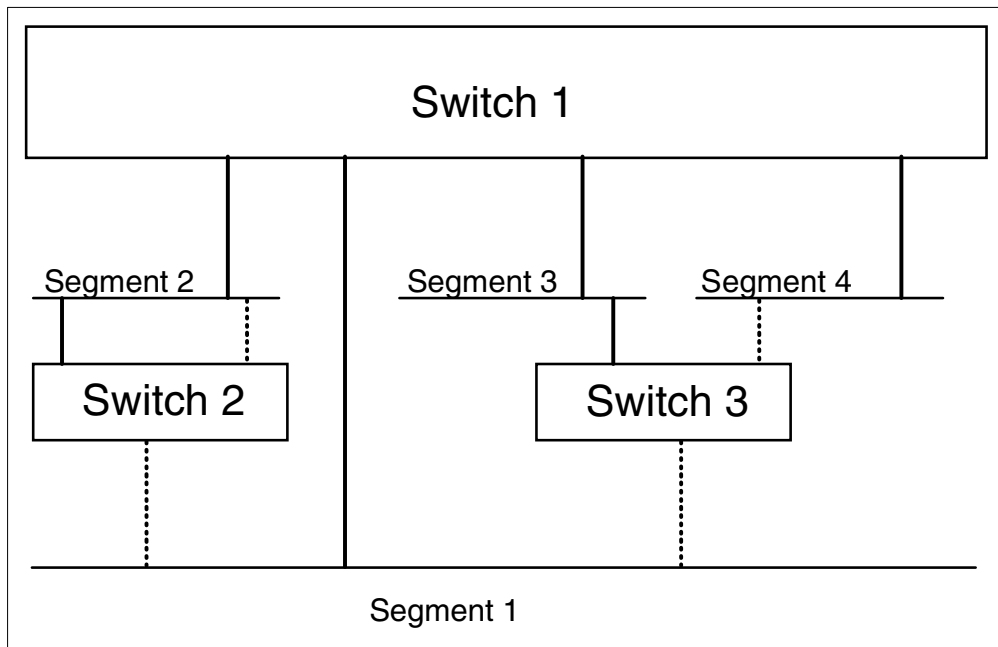


Figure 4-1 A network with spanning tree

Note: In the following explanation, the words *bridge* and *switch* are interchangeable.

The switches in Figure 4-1 above have many connections to various segments. A switch works similar to a multiport bridge. The dark lines represent physical connections that are forwarding and the broken line represents physical connections that have been blocked by the spanning tree to insure that there are no loops in the network.

4.1.1 Spanning tree sequence

To participate in the spanning tree protocol, each switch will initially assume it is the root bridge and will transmit a Hello BPDU on each of its ports.

This message will be sent every Hello time. Hello time is one of the spanning tree configuration parameters that can be specified for each switch during the switch configuration. This Hello BPDU will have the following characteristics:

1. The source address will be the address of the transmitting switch.
2. The destination address will be $\Xi'800143000000'$.
3. The source and destination SAPs will be $\Xi'42'$.
4. The Root ID field will contain the ID of the transmitting switch.
5. The Bridge ID field will contain the ID of the transmitting switch.
6. The Path Cost field will contain 0.
7. It will be sent out by the switch on all its ports.

Each Hello BPDU sent out on a switch port will be received by all the other switches that are connected to the LAN attached to that port.

BPDU's are not broadcast frames. They are sent to switches, and selectively updated and forwarded to other switches.

Each switch uses the information received in the Hello BPDU's to determine the root bridge, the designated bridges, and the designated ports within each designated bridge. To do this, each switch will continue transmitting a Hello BPDU on each of its ports until it receives a better Hello BPDU than the one it is transmitting on that port.

The better Hello BPDU will be determined based on the following information contained in the Hello BPDU (listed in order of their significance):

1. The lowest root ID
2. The lowest path cost to the switch
3. The lowest transmitting switch ID
4. The lowest port ID

As soon as a switch receives such a Hello BPDU on a port, it will stop transmitting any further Hello BPDU's on that port and will use the information received in the better Hello BPDU to transmit a new Hello BPDU on all its other ports. The new Hello BPDU will have the following characteristics:

1. The source address will be the address of this switch.
2. The destination address will be $\Xi'800143000000'$.
3. The source and destination SAP will be $\Xi'42'$.
4. The Root ID field will contain the root ID received in the better Hello BPDU.
5. The Bridge ID field will contain the ID of this switch.

6. The Path Cost field will be the sum of the path cost received in the better Hello BPDU plus the path cost defined for the switch port on which the better Hello BPDU was received.
7. The new Hello BPDU will be sent out by the switch on all its ports except the port on which the better Hello BPDU was received.

This process will be repeated by all the switches until:

1. There is one switch (root bridge) remaining that is still transmitting its original Hello BPDU.
2. One switch (designated bridge) on each LAN is transmitting the Hello BPDU based on the Hello BPDU received from the root bridge.

On the designated switch, the port on which the best Hello BPDU is received is the root port and all the ports to which the Hello BPDUs are transmitted are the designated ports.

Note: There may be some ports on the designated bridge, over which the switch will not be transmitting Hello BPDUs due to the fact that the received BPDUs on those ports are better than the one this switch would be able to transmit (but they are not better than the Hello BPDU received on its root port). Once the root and designated bridges have been elected, the root ports and the designated ports will be put in forwarding state and all the other ports will be put in blocking state.

4.1.2 Switches and network topology changes

Switches using the spanning tree algorithm automatically adjust to changes in the network topology to ensure that a loop-free network is maintained. A change in the network topology can occur in the following circumstances:

1. When switches enter or leave the network
2. When spanning tree parameters change, causing switch ports to change state or causing a change in the choice of the root bridge

The result of any of the above changes is that:

1. The spanning tree has to be reconfigured using a Topology Change Notification protocol
2. The filtering database must be updated

Filtering database update

A switch builds a filtering database for each of its ports by listening to the frames exchanged on the LAN attached to that port. This database contains the addresses of stations attached to that LAN segment and are used to forward/discard frames across the switch. When the network topology changes due to the switch addition, removal or reconfiguration, it is important that the new switches can update their filtering database quickly enough to cope with these changes in a manner that:

1. Ensures that the stations can continue to communicate with each other through the bridges.
2. Ensures the performance of the network is not affected due to the bridges forwarding the frames incorrectly and flooding the network.

To ensure the above, an aging timer is used by the bridges to delete entries within the filtering database that have not been used recently.

This timer should be able to cope with changes that happen as a result of stations physically moving from one LAN to another, as well as changes happening as a result of a bridge

addition/removal (spanning tree reconfiguration). The latter will normally result in a group of stations logically moving from one LAN to another.

In general, to cope with the changes occurring due to the station moves, a longer aging timer is required than the one required to cope with the spanning tree reconfiguration. Therefore, the standard defines two timer values for the aging timer:

1. A longer timer value is to be used in coping with normal changes due to station additions, removals or timeouts. This is a user-configurable parameter and is referred to as aging time.
2. A shorter timer value is to be used when the bridge is in a state of topology change. The forward delay timer of the root bridge is used for this purpose.

Note: The forward delay timer is specified for each bridge during its configuration, but all the bridges will use the value defined in the current root bridge.

Topology change notification

Spanning tree topology change is detected by a bridge whenever:

- ▶ A port enters the forwarding state
- ▶ A port leaves the forwarding state
- ▶ A new bridge becomes the root bridge

When a topology change occurs, the following actions will be performed:

1. The bridge detecting the change issues a Topology Change Notification (TCN) BPDU. This frame will be sent on the root port to the destination address $\Xi'800143000000'$.
2. The designated bridge on this port will acknowledge this frame by sending back a Hello BPDU with Topology Change Acknowledgment (TCA) set to 1.
3. The designated bridge will issue, on its root port, its own TCN BPDU.
4. This process repeats until a TCN BPDU reaches the root bridge.
5. The root bridge will start transmitting a Hello BPDU with the TCN set to 1 for a period equal to the sum of forward delay time and maximum age time.

The bridges that receive the Hello BPDU with TCN set to 1 will start using the shorter aging timer (forward delay) to age out filtering database entries. The forward delay timer will be used as the aging timer until a Hello BPDU with TCN set to 0 is received.

4.1.3 Setting the parameters that control the spanning tree

Table 4-1 shows the configurable spanning tree parameters that are defined as part of the standard for transparent bridging.

Table 4-1 Spanning tree parameters

Parameter	Meaning	Default
Bridge Max Age	Maximum age of received BPDU	20 seconds
Bridge Hello Time	Time interval between configuration BPDUs	2 seconds
Bridge Forward Delay	Time spent in Listening state, time spent in Learning state, short aging timer	15 seconds
Bridge Priority	Priority portion of bridge identifier	32768
Path Cost	Cost for entering this port	1000/LAN_speed (Mbps)
Port Priority	Priority portion of port identifier	128

4.1.4 Summary of the IEEE 802.1d Spanning Tree algorithm

From the previous discussion, it is clear that the key parameters governing the topology of the spanning tree are the bridge priority and the path cost. The port priority is unlikely to have any effect on the spanning topology.

In most circumstances, adjusting the bridge priority and using the IEEE defaults for all other parameters should provide acceptable control over the active topology. One approach could be to:

1. Choose three values of bridge priority: a low value, a medium value and the IEEE default.
2. Assign the low value to the switch to be the normal root bridge. This switch becomes the center of the network.
3. Assign the medium value to any other switch that may become the root. This allows for failure of the root bridge.
4. Allow all other switch in the network to use the IEEE default.

The result of the spanning tree algorithm for switches and transparent bridges is a loop-free network in which the endstations require no knowledge of the network topology to be able to communicate with the other stations through one or more switches. However, another result is a network in which there is no load balancing over switches and in case of parallel switches, all but one of the switches will be idle (blocking state).

4.2 IEEE 802.1Q VLANs

Virtual local area networks (VLANs) allow us to create logical groups within a switch in order to control broadcast traffic. Devices attached to switch ports in one VLAN can only communicate within that VLAN, unless a router connects the VLAN to another network. VLANs in a switch are functionally equivalent to isolated, bridged Ethernet networks. We have the convenience and flexibility to assign a port to any VLAN, whereas in traditional Ethernet networks, we must physically move a device or cable. In the following sections, we discuss the need for IEEE 802.1Q VLANs and how to implement them in the ESM.

Note: VLANs can be defined based on other criteria as well, such as MAC address or protocol type, but these are beyond the scope of this discussion and are not supported by the ESM.

4.2.1 The problem with traditional VLANs

The VLANs described thus far assume that all machines connect to a single switch. If we expand our network by connecting an additional switch, it cannot extend or enlarge the VLANs from the first switch. For example, in Figure 4-2, we connect port 1 on Switch B to port 6 on Switch A. Both ports belong in VLAN 3, so traffic in that VLAN flows freely between switches. However, VLAN 1 and VLAN 2 are isolated in each switch. Devices in VLAN 1 or VLAN 2 of Switch A cannot communicate with devices in VLAN 1 or VLAN 2 of Switch B.




Figure 4-2 Both VLANs 1 and 2 are isolated in each switch. Only VLAN 3 spans both switches

Alternatively, if we do not create any VLANs in Switch B (that is, all ports are in the default VLAN), then all devices attached to Switch B would be grouped into VLAN 3 in the first switch (Figure 4-3). Again, Switch B cannot access VLAN 1 and VLAN 2 in Switch A. In both cases, we can only connect one VLAN in Switch A to one VLAN in Switch B.




Figure 4-3 Switch B can communicate with Switch A in VLAN 3 only

To solve this problem, we require a mechanism to satisfy the following requirements:

1. Allow traffic from multiple VLANs in one switch to cross a single link to another switch.
2. Identify the VLAN that each frame belongs in *before* it crosses the link so the remote switch knows where to forward the traffic.

Figure 4-4 shows these two ideas at work. Port 6 on Switch A and port 1 on Switch B are no longer in VLAN 3. Instead, they are ports designated for carrying traffic to and from different VLANs. The frames that travel across the link have VLAN identification so they can be forwarded to the appropriate ports by the remote switch.




Figure 4-4 Transporting three VLANs across a single link

IEEE 802.1Q is a standard protocol that enables VLANs to span multiple Ethernet switches. Prior to 802.1Q, manufacturers used proprietary protocols to forward VLAN traffic between their switches. Unfortunately, this meant that we could not mix different brands of switches and still maintain VLAN consistency across the network. Also, with 802.1Q, we still have the flexibility to connect with non-802.1Q (legacy) switches.

4.2.2 To tag or not to tag

One of the key concepts of 802.1Q is the tag header. An Ethernet frame that contains a tag header immediately following the source MAC address field is called a *tagged* frame (see Figure 4-5 on page 43). Switches (such as the ESM, and also some network interface cards) that support 802.1Q read the contents of the tag to determine in which VLAN the frame belongs. Likewise, an *untagged* frame is an Ethernet frame that does not contain a tag header. These frames do not contain a VLAN identifier. For each VLAN, switch ports are either in tagging mode or untagging mode, depending on the type of device connected.

If we set a switch port to tagging mode for a particular VLAN, then every frame from that VLAN will be transmitted with a tag header. If we set a switch port to untagging mode for a particular VLAN, then every frame from that VLAN will be transmitted without a tag header. The tagging/untagging parameter only affects frames transmitted by the switch. It does not affect frames received by the switch. In other words, the tagging/untagging function affects egress (outgoing) traffic only.

Important: The Tag check box in the ESM's Edit 802.1Q VLANs menu determines if frames from a particular VLAN should be transmitted from that port with a tag header (checked) or without a tag header (unchecked). This option has no effect on ingress (incoming) frames.

With the tag header, switches are able to preserve VLAN traffic as it travels from the source to the destination. VLAN-tagged frames are never modified by switches. Once a tag is created at the source, all 802.1Q switches preserve the tag to ensure proper VLAN isolation. A switch removes the tag only when the frame must be sent to a device that only accepts untagged Ethernet frames (for example, PCs, servers, printers, and non-802.1Q switches). In general, we enable tagging on both ends of a link that support 802.1Q, thereby creating a VLAN *trunk* connection.¹ The trunk can carry as few as one VLAN, or as many as 256 VLANs (in the ESM).

¹ Although the word *trunk* has more than one meaning in networking technology (for example, a trunk is also known as a link aggregate or channel), we define it here as any connection that carries 802.1Q VLAN-tagged traffic.




Figure 4-5 Difference between a tagged and an untagged Ethernet frame

Example: Creating a new VLAN²

1. Click **Configuration** -> **VLANs** -> **Edit 802.1Q VLANs**.
2. Click **Add**.
3. In the VLAN ID (VID) field, type 10.
4. In the VLAN Name field, type VLAN 10.
5. Click **Apply**. Refer to Figure 4-6.

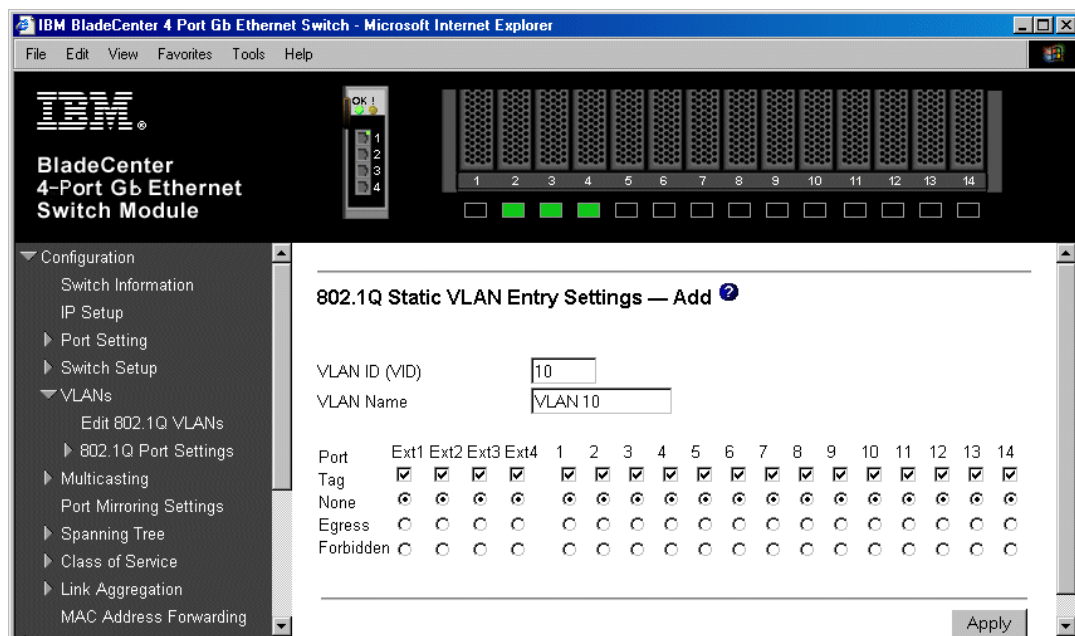


Figure 4-6 Creating a new VLAN

Example: Setting ports to untagging mode

1. Click **Configuration** -> **VLANs** -> **Edit 802.1Q VLANs**.
2. Select **VLAN ID 10**.

² All examples assume that we use the Web interface for managing the ESM.

3. Click **Edit**.
4. Under Port 2, click **Tag**. The box should now be unchecked. Repeat for Ports 3 and 4.
5. Click **Apply**. Refer to Figure 4-7.

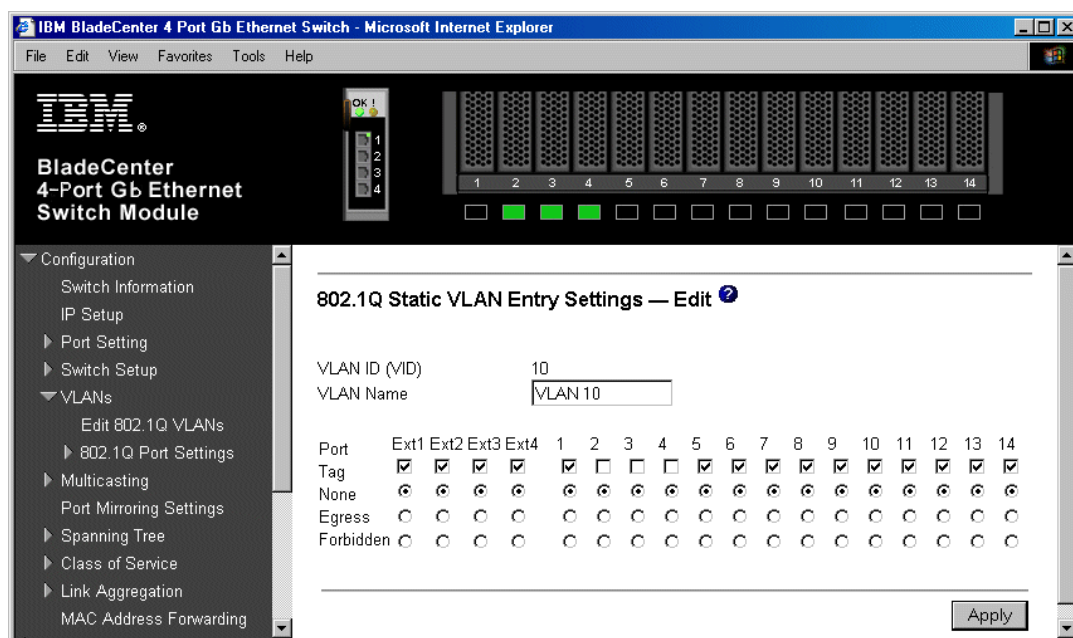


Figure 4-7 Settings ports to untagging mode

4.2.3 Port VLAN identifiers (PVIDs)

As stated in the previous section, when we enable tagging on a switch port, the switch preserves the 802.1Q VLAN tag as frames exit the port. Conversely, when we disable tagging, the switch removes the VLAN tag. A tag is never created at a port during egress. This implies that tag headers are created somewhere else in the process. Indeed, when a normal, untagged Ethernet frame arrives at a switch port, the switch must insert a tag into the frame to identify in which VLAN the frame belongs. In the ESM, each internal and external port has one (and only one) Port VLAN Identifier (PVID) associated with it. The PVID determines the VLAN identifier (VID) that is inserted into the tag of an originally untagged frame.

The PVID comes into play only for untagged frames as they enter a port. The PVID has no bearing on egress frames. By using the PVID to create the VID in the 802.1Q tag header, a switch guarantees that all frames are tagged as they enter its ports. In other words, the PVID is an ingress parameter for a port. In contrast, tagging/untagging is an egress parameter. These two concepts are critical to understanding how 802.1Q VLANs work, and how to configure the BladeCenter ESM. Although not obvious at first, PVIDs are always involved in the operation of the ESM. Even in factory-default state, all ESM ports are assigned a PVID of 1, meaning that all untagged frames are tagged with a VID of 1 at ingress. Also, by default, all ESM ports are set to untagging (that is, the Tag box is unchecked), meaning that all frames are untagged at egress.

Creating VLANs using only the PVID

When we change the PVID of a port to a value other than 1, the ESM checks if a VLAN already exists with a matching VID. If the VLAN does not exist, the ESM automatically creates a new VLAN with a matching VID (that is, if we select a PVID of 20, then a VLAN with a VID of

20 is created). In addition, the ESM sets the port to untagging mode and makes it a member of the VLAN.

Example: Assigning PVIDs

1. Click **Configuration -> VLANs -> 802.1Q Port Settings -> Port VLAN ID (PVID)**.
2. For Bays 2, 3, and 4, enter 10 in the PVID field.
3. Click **Apply**. Refer to Figure 4-8.

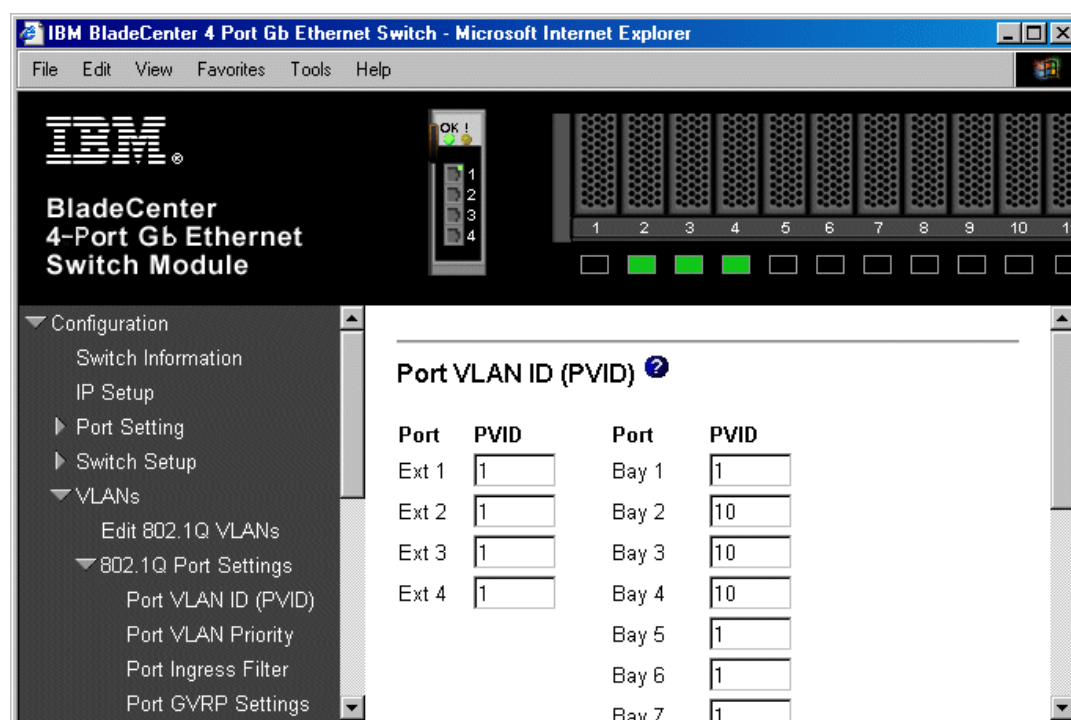


Figure 4-8 Assigning PVIDs

4.2.4 VLAN identifiers (VIDs)

Although they appear similar, VLAN Identifiers (VIDs) are not the same as Port VLAN Identifiers (PVIDs). Whereas one PVID is fixed to each port (that is, all untagged ingress traffic is associated with one VLAN), a port can potentially transmit and receive tagged frames with different VIDs. Once a frame is tagged, we no longer refer to a PVID. Instead, we refer to a tagged frame's VID. Any VLAN in an 802.1Q network is also referenced by its VID, a number which ranges from 1 to 4094. Therefore, a tagged frame with a specific VID is associated with the VLAN having the same VID. The distinction between VIDs and PVIDs allows us to configure 802.1Q VLANs in a flexible manner.

4.2.5 VLAN membership

A port, using its PVID, automatically becomes a member of that VLAN. By default, all ESM ports are members of VLAN 1, since all ports have a PVID of 1. After a switch makes a forwarding decision based on the frame's destination MAC address, it checks the VID in the tag to see if the destination port is a member of that VLAN. If the destination port is a member of the VLAN, then the frame is transmitted. If the destination port is not a member of the VLAN, the switch discards the tagged frame regardless if the intended destination MAC address is located on that port. This process, known as VLAN filtering, is applied to every

frame before it exits the port. Notice that this explanation of VLAN membership does not include ports that receive incoming traffic. In fact, VLAN filtering only applies to egress traffic. A switch does not apply VLAN filtering to ingress frames under normal circumstances.³

4.2.6 To egress or not to egress

An interesting feature of 802.1Q is that a port can be a member of more than one VLAN. In the previous sections, we discussed how assigning a PVID to a port automatically makes that port a member of the VLAN. This method of issuing VLAN membership is simple and straightforward. A second method, which allows ports to be members of multiple VLANs, is to create a new VLAN from scratch (that is, a VLAN with no member ports yet; refer to “Example: Creating a new VLAN” on page 43). We start by specifying a VLAN number, known as the VID (2 to 4094 in the ESM), and then choose the ports to be members of the VLAN.

Recall that the VLAN filtering process involves the comparison of a frame’s VID and the destination port’s membership in that VLAN. This filter is applied to a destination port before it can transmit the frame. Therefore, when we select ports to be members of the VLAN, we also refer to this as selecting *egress* ports, because VLAN membership permits frame transmission. Similarly, when we remove a port’s VLAN membership, we also refer to this as removing an *egress* port’s VLAN membership.

If we select the same port to be a member of multiple VLANs then the switch’s VLAN filter will allow frames with VIDs that match those VLANs to egress that port. If the VIDs do not match egress port membership, then the frames are discarded.

Tip: In the ESM, setting a port to Egress in the Edit 802.1Q VLANs menu makes that switch port a member of the VLAN. Setting the port to None removes it from the VLAN.

Example: Assigning VLAN membership to ports without using PVIDs

1. Click **Configuration -> VLANs -> Edit 802.1Q VLANs**.
2. Select VLAN ID 10 and click **Edit**.
3. Under Port Ext1, select **Egress**. Repeat this step for Ports Ext2, Ext3, and Ext4.
4. Click **Apply**. Refer to Figure 4-9.

³ If ingress filtering is enabled on a port, then a switch will discard received tagged frames if the port is not a member of the VLAN specified by the VID of the tagged frame.

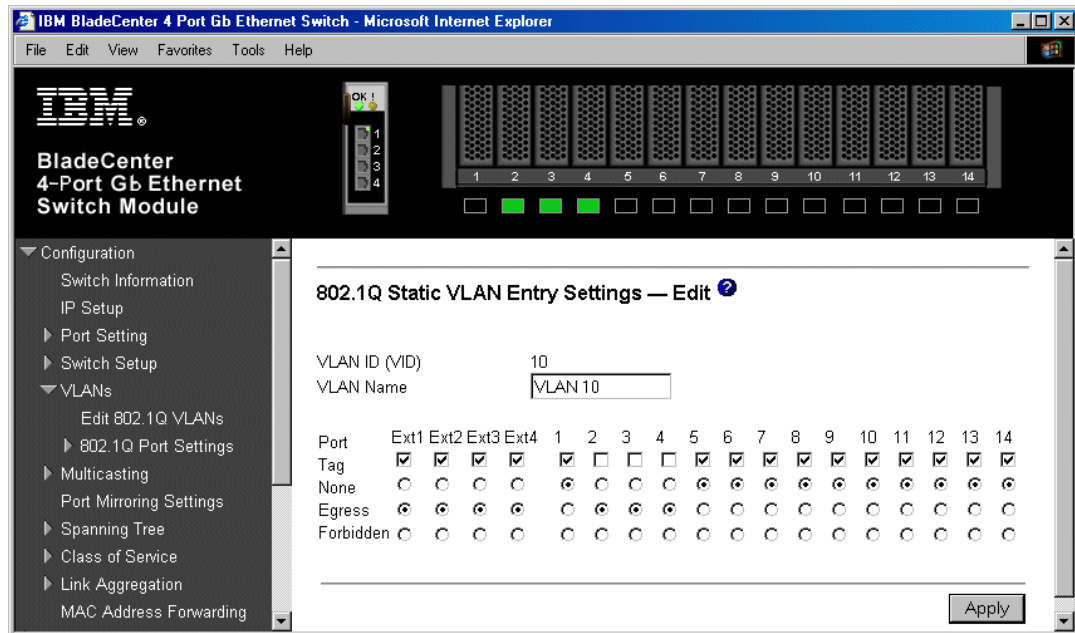


Figure 4-9 Assigning VLAN membership without using PVIDs

4.2.7 Ingress filtering

So far, we have explained how VLAN filtering affects a switch's egress traffic. In 802.1Q, a feature known as *ingress filtering* triggers a switch to look at a port's inbound traffic as well. By default, ingress filtering is disabled for every port in the ESM. When enabled, a switch discards any ingress, tagged frames if the port is not a member of the VLAN specified in the frame's VID. If the VID matches the VLAN membership of that port, then the tagged frame is processed as usual. If we enable ingress filtering and the port receives an untagged frame, then the switch uses the PVID to tag the frame and process it as usual.

Note: Port ingress filtering has *no effect* on untagged traffic.

By performing ingress filtering, a switch prevents unnecessary frame forwarding and processing later on, where the frame might be discarded anyway. There are some situations where ingress filtering is beneficial. For example, we would enable ingress filtering on a busy 802.1Q VLAN trunk port if we know that all end stations attached to that switch do not need to communicate with specific VLANs. The switch avoids the effort of flooding unknown destination or broadcast VLAN-tagged frames to all other active ports, and then discarding them by VLAN filtering at the egress ports. It is much easier (and bandwidth-efficient) to discard one unwanted tagged frame at ingress than many copies of that tagged frame at egress.

Example: Enabling ingress filtering

1. Click **Configuration -> VLANs -> 802.1Q Port Settings -> Port Ingress Filter**.
2. Enable the ingress filter for ports Ext1, Ext2, Ext3, and Ext4.
3. Disable ingress filtering on all other ports.⁴
4. Refer to Figure 4-10 on page 48.

⁴ Do not enable ingress filtering on internal ESM ports since the blade servers generate untagged frames only.

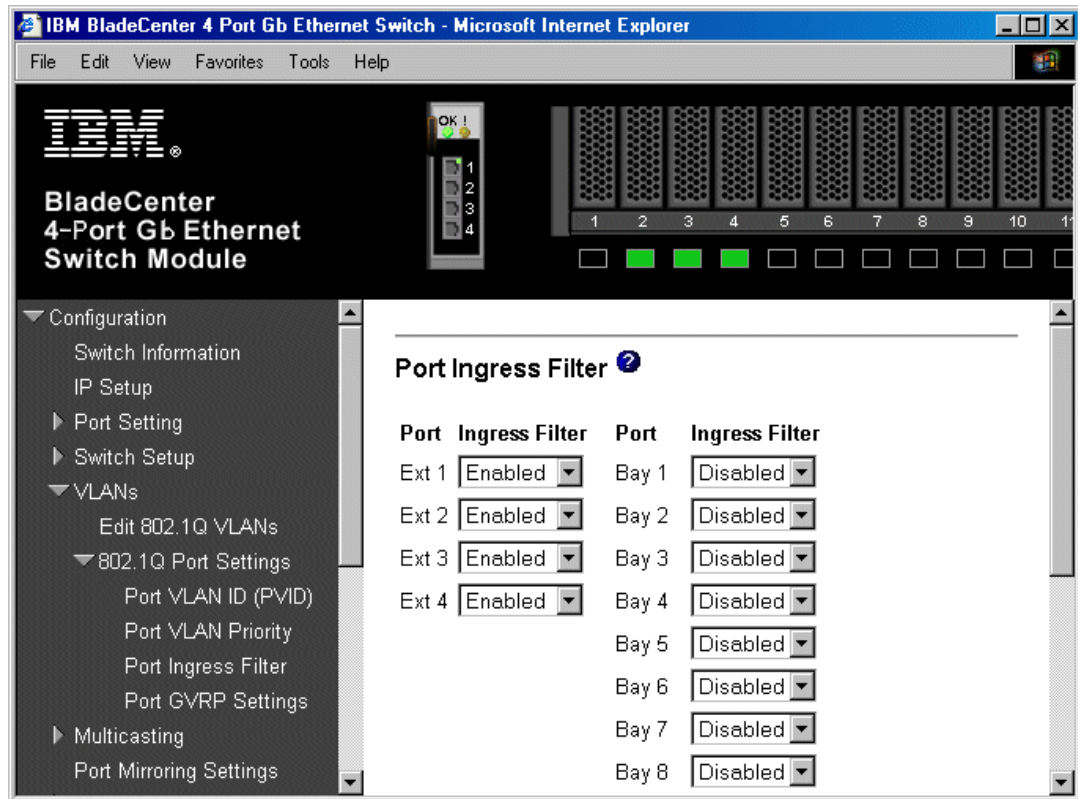


Figure 4-10 Enabling the port ingress filters

4.2.8 Summary of VLAN tagging, filtering, and forwarding

To help us integrate the main ideas of 802.1Q, we present a flowchart of the frame forwarding process in Figure 4-11. The flowchart assumes that the ingress port receives a unicast frame with a known destination MAC address. Therefore, no flooding occurs within the VLAN.

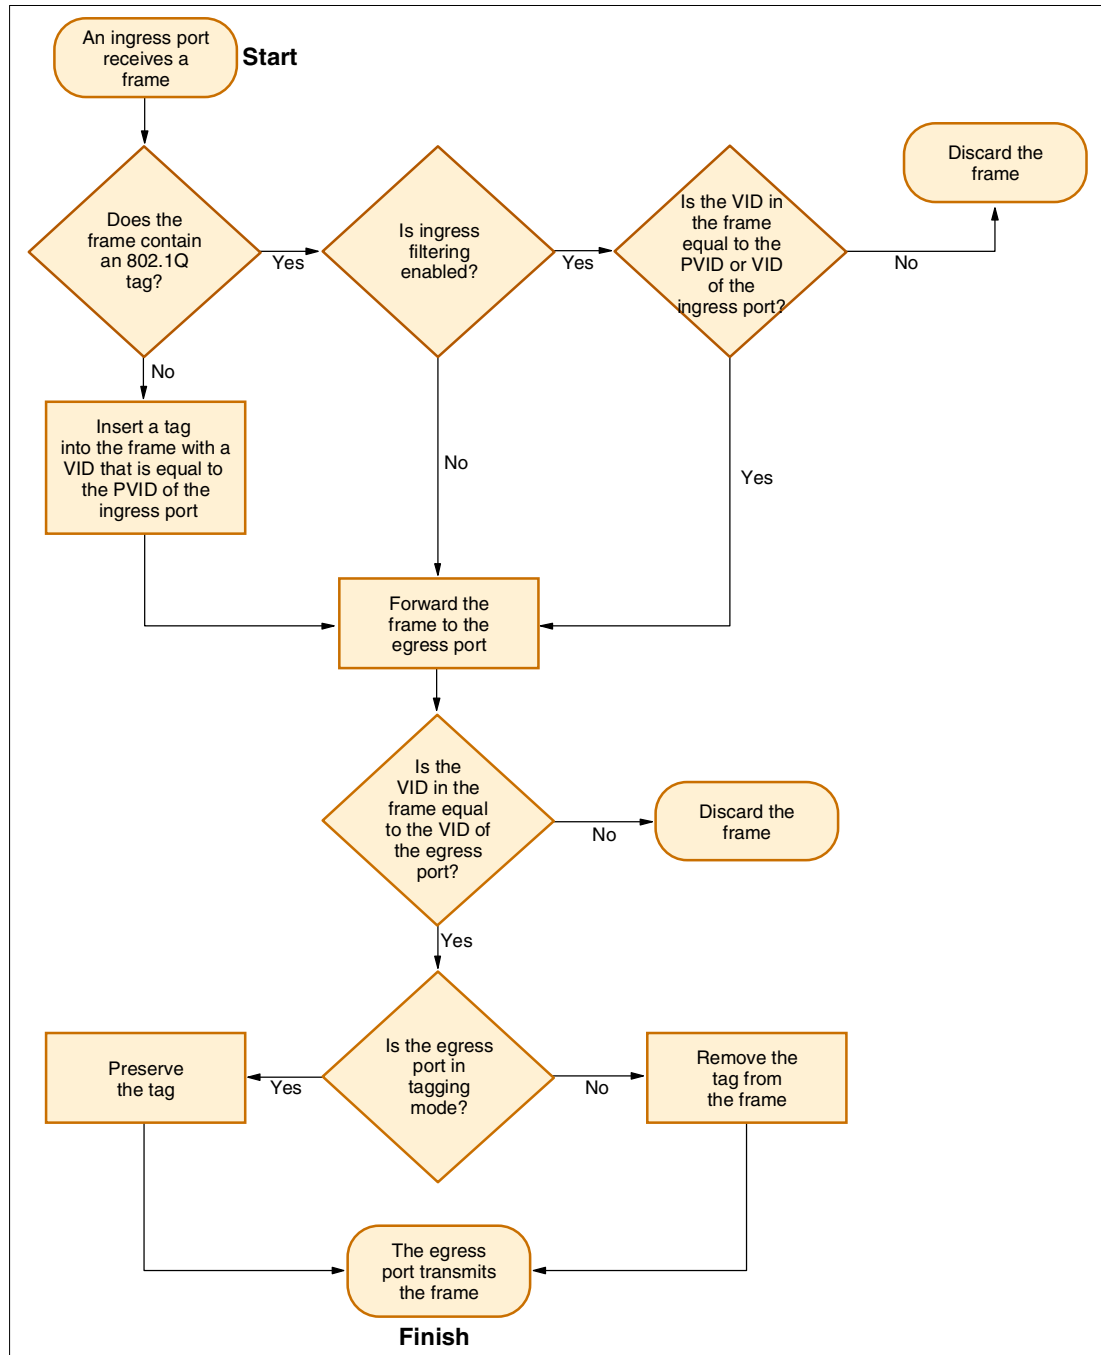


Figure 4-11 802.1Q frame filtering and forwarding

4.2.9 Sample VLAN configuration

In Figure 4-12 on page 51, we show a sample VLAN configuration that demonstrates the use of PVIDs and VIDs to control the flow of traffic within the BladeCenter chassis. The table below the illustration provides detailed information about each VLAN using a layout similar to the ESM's web interface menu. Here are some important observations (and reminders) to help us interpret the configuration:

- ▶ Always set internal ESM ports 1 to 14 to untagging mode, unless a blade server is enabled for 802.1Q VLAN tagging to accept traffic from multiple VLANs.
- ▶ Any port that is not a member of the VLAN (indicated by a "-") is always forced into tagging mode by the ESM management interface.
- ▶ A blade server transmits data into the VLAN specified in its PVID.
- ▶ A port is a member of the VLANs specified in its PVID *and* its VIDs. For example, port 14 is a member of all four VLANs.
- ▶ All external ESM ports are not members of VLAN 4. Therefore, VLAN 4 traffic is isolated within the BladeCenter chassis.
- ▶ All external ESM ports are members of VLANs 1, 2, and 3 and having tagging enabled. These ports connect to other switches that also have tagging enabled.
- ▶ All external ESM ports have a PVID of 1. Any untagged traffic entering the ESM through these ports will assigned to VLAN 1.
- ▶ Only blade server 14 can transmit data into VLAN 4 because only port 14 has a PVID of 4, but all blade servers can receive data from VLAN 4.
- ▶ Blade server 14 is unique in that it can communicate with all other servers. All other blade servers (in ports 1-13) can communicate only within their respective VLANs, and with blade server 14.
- ▶ This configuration allows VLANs to overlap yet maintains traffic isolation between selected devices. For example, although blade servers 7, 8, and 9 are egress members of VLAN 4, they can only transmit into VLAN 3. Therefore, VLAN 3 and VLAN 4 together create a bidirectional association with blade server 14.

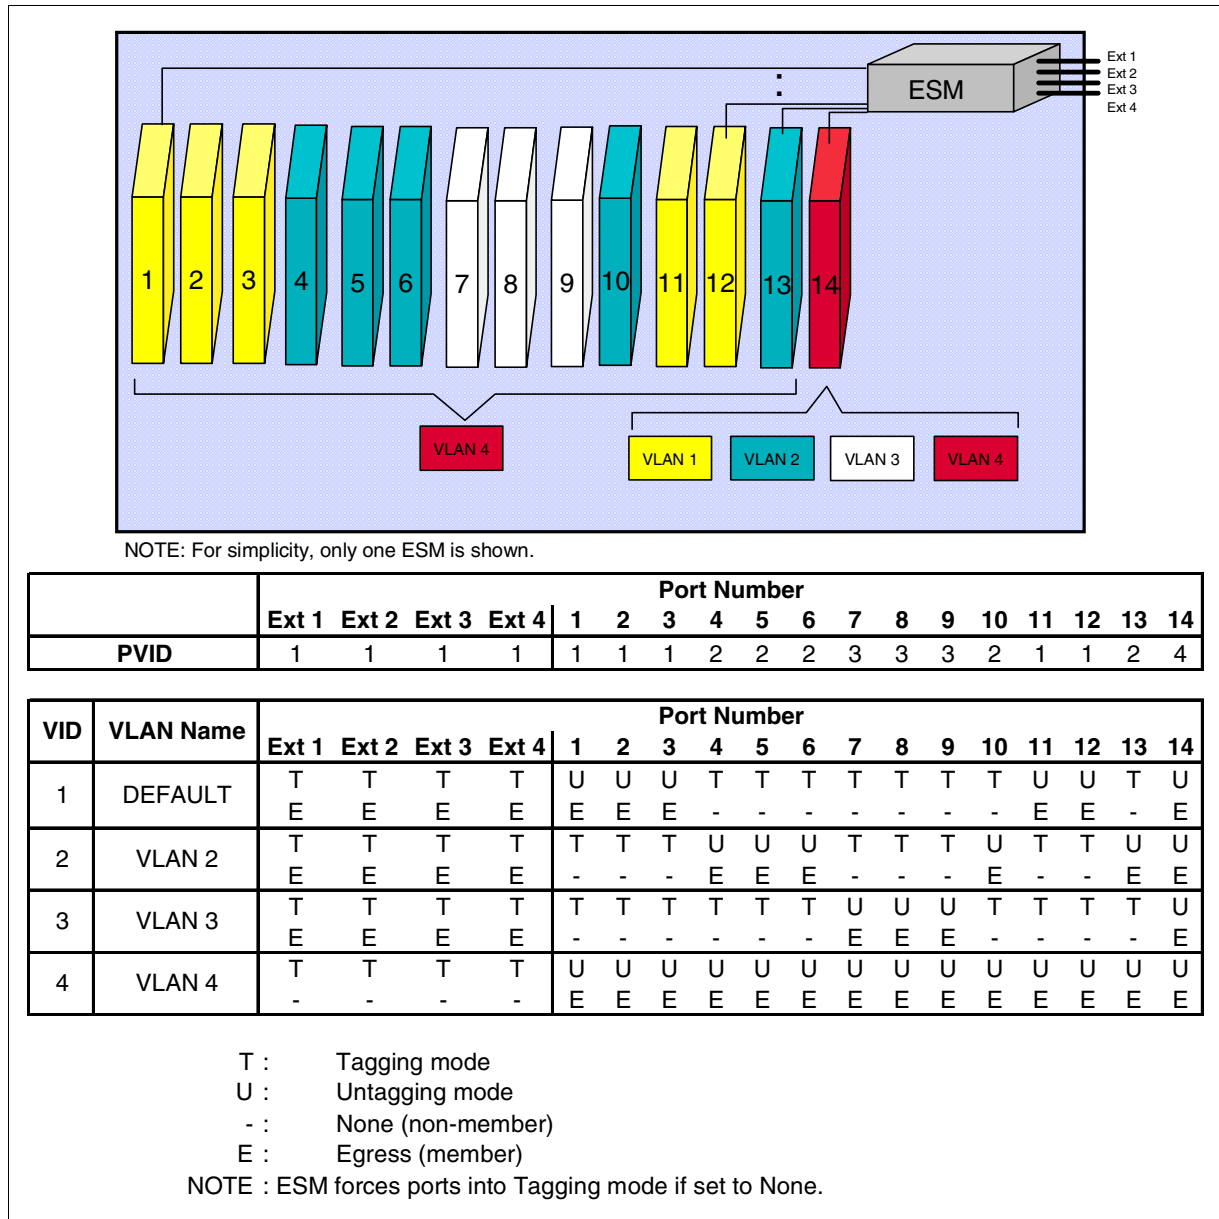


Figure 4-12 Sample VLAN configuration using both PVIDs and VIDs

4.2.10 Dynamic VLANs and GVRP

The IEEE 802.1Q standard takes VLANs one step further by supporting *dynamic VLANs*. All of the topics covered up to this point assume that we manage 802.1Q VLANs from each switch. This is fine for a small network of 802.1Q-compliant switches. Unfortunately, as we add more switches to the network, configuring 802.1Q VLANs in each switch becomes a laborious task. IEEE 802.1Q uses a protocol called GVRP (GARP VLAN Registration Protocol), which is an application of GARP (Generic Attribute Registration Protocol), to enable switches to communicate VLAN information amongst themselves. A GVRP-enabled switch advertises its VLANs to other switches and learns about VLANs from other switches.

When we manually create a VLAN in a switch, we are actually creating a *static* VLAN. A dynamic VLAN, on the other hand, is automatically created by a switch when it learns about the VLAN using GVRP. Static VLANs still remain after a switch is powered off or rebooted, but

dynamic VLANs disappear. Because of the temporary nature of dynamic VLANs, we take advantage of GVRP in certain instances only.

For example, core switches that connect multiple 802.1Q trunks in a redundant manner would benefit from GVRP for the following reasons:

1. Core switches typically have no end user connections, so they do not require static VLANs.
2. Some of the redundant links are not active due to the spanning tree protocol. It makes sense to let the switches dynamically modify their VLAN configuration based on the current network topology, rather than manually altering VLAN memberships whenever ports change from blocking to forwarding or vice versa.
3. GVRP greatly simplifies the process of creating a new VLAN that spans from edge switch to edge switch. Instead of setting up the same VLAN in every switch, we just need to create a static VLAN in the edge switches. Then, the edge switches advertise the VLAN to its GVRP-enabled neighbors. Eventually, all of the core switches learn about the new VLAN.

With this knowledge of dynamic and static VLANs, we now see the differences between the three types of port membership in an 802.1Q VLAN. In the ESM, they are labeled as follows under the Edit 802.1Q VLANs menu:

1. None - the port is not a static VLAN member, but can dynamically become a VLAN member if GVRP allows it.
2. Egress - the port is a static VLAN member.
3. Forbidden - the port can neither be a static nor a dynamic VLAN member.

Tip: If ingress filtering is enabled to block traffic from certain VLANs, we should set the port to the Forbidden state for those VLANs. Otherwise, the port becomes a dynamic member for any VLANs learned from GVRP, which defeats the purpose of the ingress filter.

Example: Enabling GVRP

1. Select **Configuration -> Switch Setup -> Switch Settings**.
2. Ensure that Switch GVRP is set to Enabled. Click **Apply**.
3. Select **Configuration -> VLANs -> 802.1Q Port Settings -> Port GVRP Settings**.
4. Enable GVRP on ports Ext 1, Ext 2, Ext 3, and Ext 4.
5. Disable GVRP on all other ports.⁵ Click **Apply**. Refer to Figure 4-13 on page 53.

⁵ Do not enable GVRP on internal ESM ports because blade servers do not require this protocol.

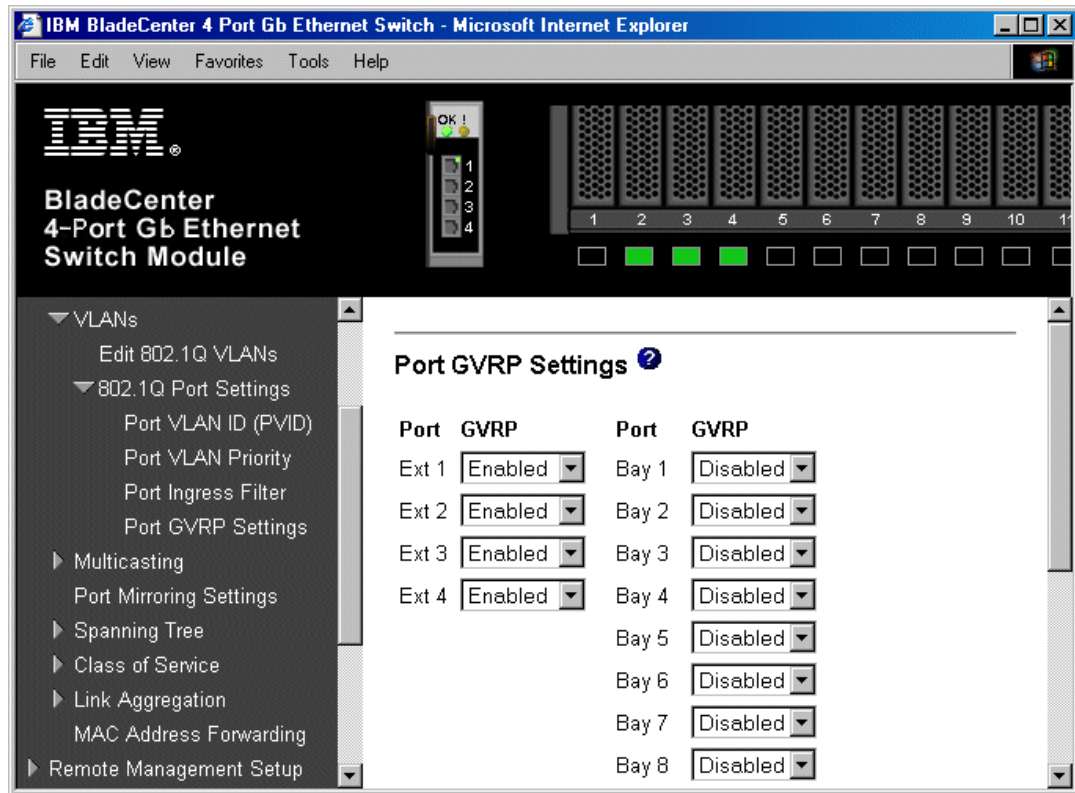


Figure 4-13 Enabling GVRP at the port level

4.2.11 Priority tagging (IEEE 802.1p)

Another feature of 802.1Q is the option to insert user priority information into the tag header (see Figure 4-5 on page 43). IEEE 802.1p⁶ (now incorporated into 802.1D) defines how switches should support the classification and transmission of time-sensitive data. Since 802.1Q also defines how switches communicate traffic priority using tags, priority tagging is sometimes referred to as “802.1Q/p”. We use priority tagging to implement application service quality in Ethernet, a technology that normally transfers data on a best-effort basis.

Port priority

Priority-tagged frames, by definition, have a VID of 0 (null). These frames are generated by end stations with network adapters capable of priority tagging. When a priority-tagged frame enters a switch port, the switch uses the PVID to insert a non-zero VID into the frame (refer to 4.2.3, “Port VLAN identifiers (PVIDs)” on page 44). This frame now becomes a VLAN-tagged frame. For untagged frames, the switch uses the *port priority* to assign a user priority to the frame, much like it uses the PVID to assign a VID to the frame. By default, every port in the ESM has a priority of 0. To change this value, select **Configuration -> VLANs -> 802.1Q Port Settings -> Port VLAN Priority**.

Note: It is possible to assign priority based on other criteria as well, such as destination MAC address or IP address, but these are beyond the scope of this discussion and are not supported by the ESM.

⁶ IEEE nomenclature is the reason for using a lowercase p in 802.1p; it is not a standard on its own. 802.1Q is a full standard, and therefore contains an uppercase Q.

Mapping user priority into queues

The priority of tagged frames ranges from 0 (default) to 7, which is the highest priority. A switch forwards frames into different output queues based on the user priority. All 802.1p-compliant switches support more than one queue (or traffic class) per port. However, not all switches have eight queues, so some user priorities must be grouped together. In the ESM, each port has four output queues for prioritizing traffic based on a *weighted round robin* algorithm. By default, the eight 802.1p user priorities map into the four queues as shown in Figure 4-14 on page 54. We can change this mapping to satisfy the requirements of different applications.

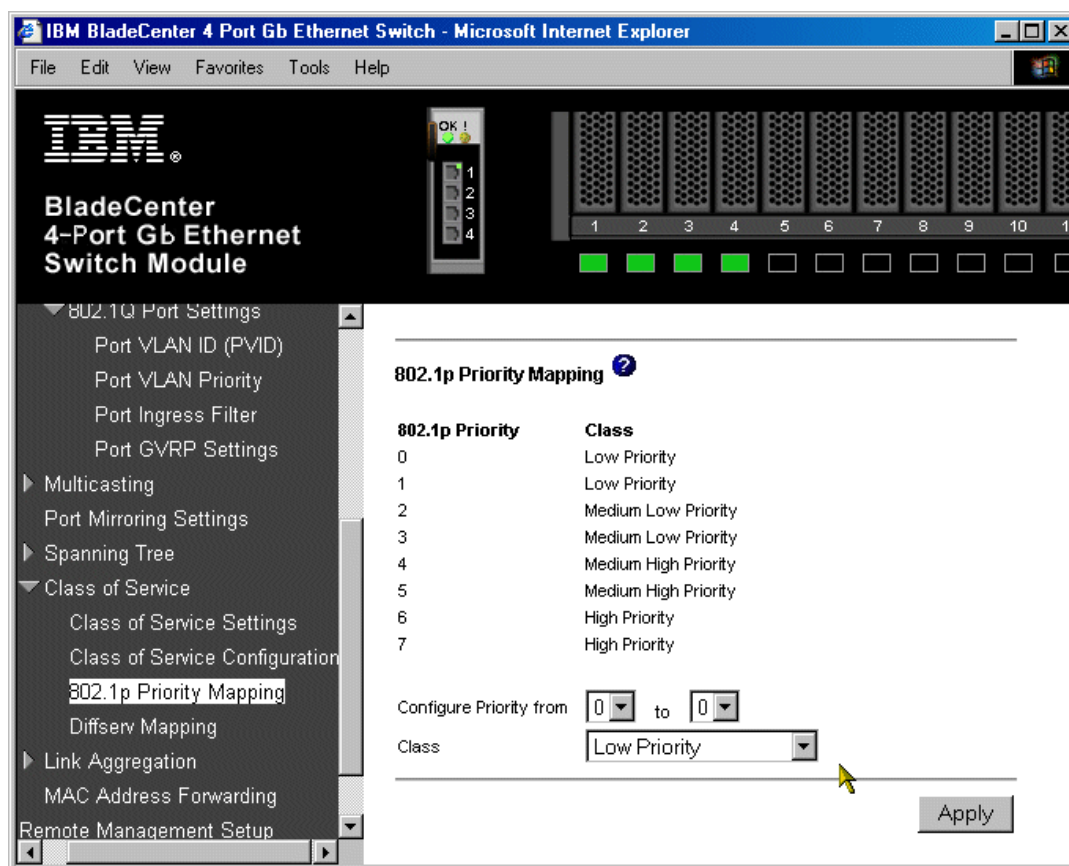


Figure 4-14 802.1p priority mapping

Weighted round robin (WRR)

The weighted round robin algorithm is disabled by default. In this case, the ESM uses normal round robin to transmit one frame from each queue in a cyclic manner. Since there are four queues, this roughly divides the port's data rate by four. For example, if we have a 1-Gb/s port, then each queue would be allocated 250 Mb/s on average. A queue still transmits each frame at 1 Gb/s, but because it must take turns with three other queues, its effective rate is reduced to 250 Mb/s. Keep in mind that this is only an approximation since all frames are not equal in size. Also, we assume that every queue has frames to transmit at all times. If all frames are classified into one traffic class (and therefore into one queue), then round robin distribution is meaningless.

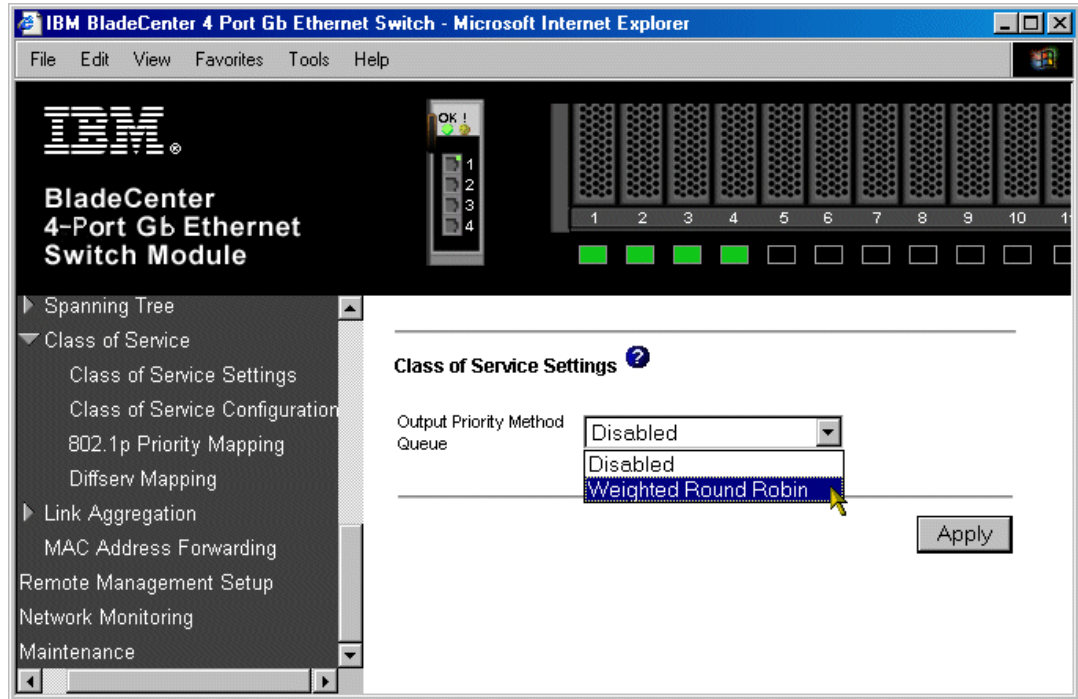


Figure 4-15 Enabling Weighted Round Robin

To activate output prioritization in the ESM, we must enable the weighted round robin (WRR) algorithm (see Figure 4-15). Once we enable WRR, the ESM uses a weight allocation illustrated in Figure 4-16 on page 56. A weight ranges from a value of 1 to 15. The only restriction is that the weight of any queue must be greater than or equal to the weight of the queue below it. Each weight represents the number of frames that the port transmits from the queue before advancing to the next queue. We can approximate the performance of prioritized traffic with this formula:

$$\text{Effective Data Rate} = \frac{\text{Queue Weight}}{\text{Total Weight}} \times \text{Port Data Rate}$$

The total weight is the sum of the individual weights of each queue. Applying the default weights of 1,4,10, and 15 to correspond with the Low, Medium Low, Medium High, and High Priority queues respectively, we obtain these effective data rates for a connection at 1 Gb/s:

- ▶ Low Priority: 33 Mb/s
- ▶ Medium Low Priority: 133 Mb/s
- ▶ Medium High Priority: 333 Mb/s
- ▶ High Priority: 500 Mb/s

Normal round robin is essentially the same as weighted round robin, but with all weights equal to 1. Figure 4-17 on page 57 provides a visual illustration of a weighted round robin mechanism at work.

In addition to assigning weights for frame transmission, we can also specify the maximum latency value for each queue. This parameter determines how long a frame stays in the queue. If a frame has been in the queue for the maximum latency period, the ESM will transmit the frame regardless of the weighted round robin algorithm. By default, the maximum latency value for every queue is 400 ms. This ensures that all frames will be sent within 400 ms of entering the queues, assuming that the port is able to transmit (for example, flow control is not stopping transmission).

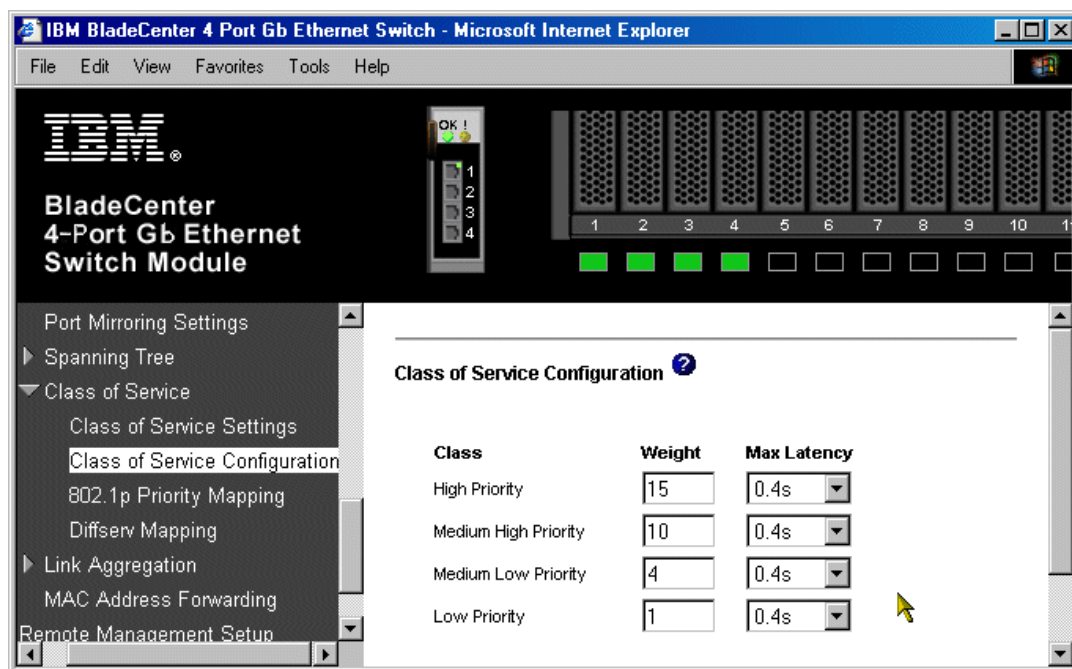


Figure 4-16 Priority queue weights and maximum latencies

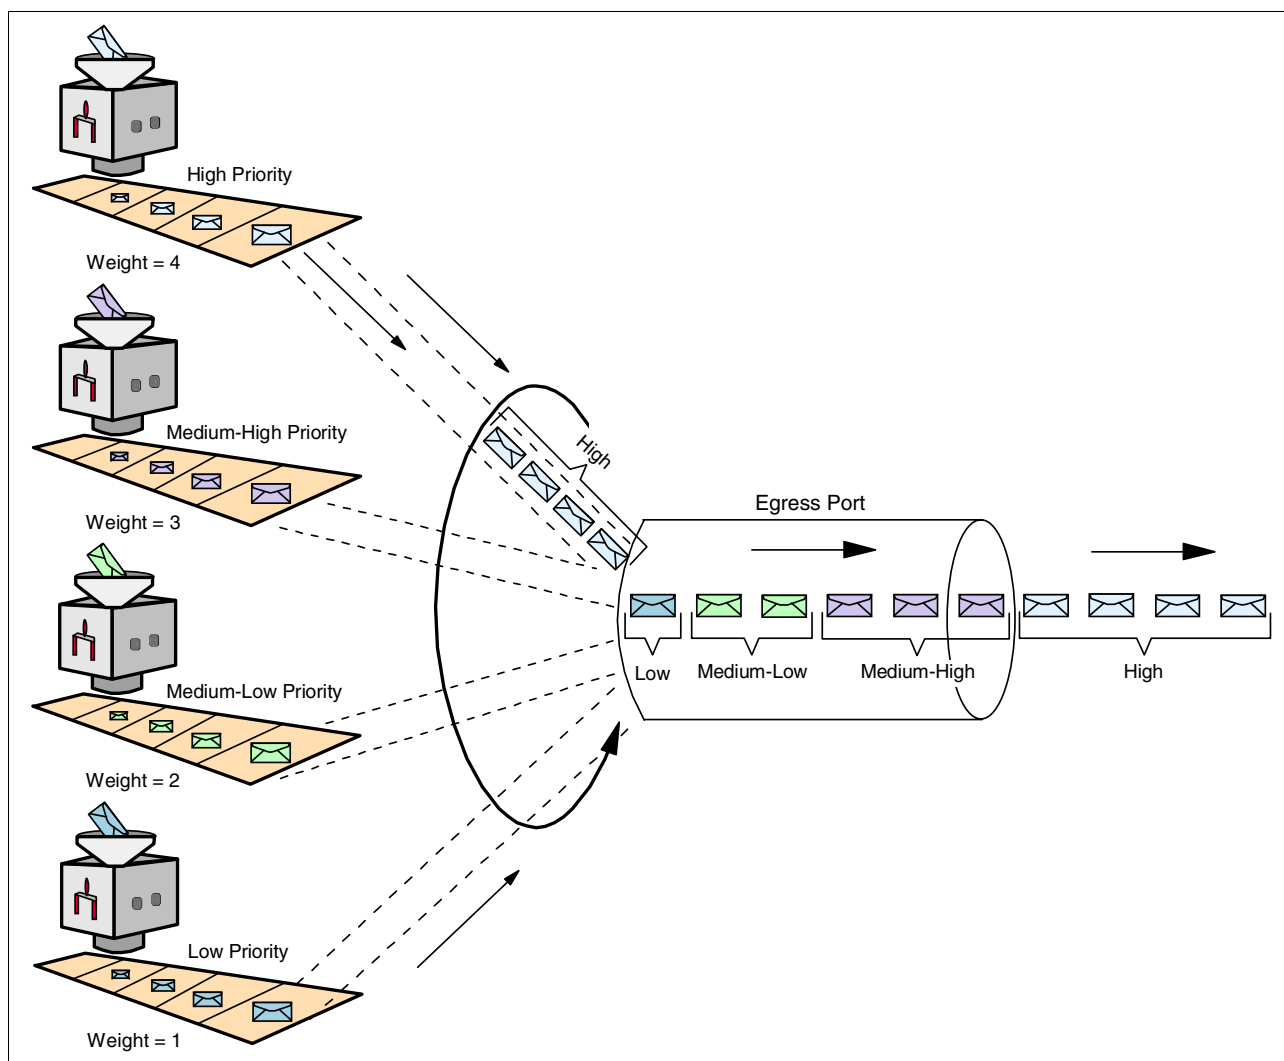


Figure 4-17 Applying a weighted round robin algorithm to frame transmission

4.3 Link aggregation

This section describes link aggregation and how to implement it using the BladeCenter ESM. In this section, we include some examples of link aggregation configuration.

As illustrated in Figure 4-18, link aggregation combines physical links into one point-to-point logical link for increased bandwidth and higher fault tolerance. To clients it appears as if it is a single link, however the traffic is load balanced across the physical links. If one link fails, the traffic is redirected across the remaining links with no noticeable impact to the clients. Link aggregation is primarily configured between switch-to-switch or switch-to-server connections.




Figure 4-18 Link aggregation

Most vendors originally implemented a non-standard mode of link aggregation. The vendors often use the term *trunk* or *channel* to describe this mode. In most cases, these implementations are interoperable. However, to avoid any problems that may exist with the non-standard methods the IEEE introduced the 802.3ad standard (now a part of 802.3-2002). In this document, we refer to the non-standard method as static link aggregation, and the standard mode as 802.3ad link aggregation. The ESM supports both of these modes of link aggregation on the four external ports.

4.3.1 Static link aggregation

Static link aggregation is a non-standard method of link aggregation. In this mode, the two devices being connected do not exchange protocol information. The ports on each device are statically assigned to a group. The ports are recognized as belonging to a static link aggregation immediately upon being added to the group. However, the devices on both sides of the connection must be configured for static mode in order for the aggregation to operate properly.

The ESM management session refers to static link aggregation as port trunking.⁷ Many switch vendors have their own implementation of static link aggregation, however without standardization some incompatibilities may exist. We recommend you use the static link aggregation option only when the other device does not support IEEE 802.3ad.

4.3.2 802.3ad link aggregation

IEEE 802.3ad introduces standardization, as well as the ability for two systems⁸ to dynamically determine which ports to aggregate. The systems use the link aggregation control protocol (LACP) to exchange aggregation capabilities and to attempt to form link aggregation groups (LAGs). A link aggregation group is one or more ports that have a point-to-point connection to another system, and that have the same operational key value. If we enable LACP the systems will attempt to dynamically assign links to a link aggregation group (LAG). Both sides of the connection must be enabled for 802.3ad. For example, if we configure one switch for 802.3ad link aggregation, but not the other switch the links are treated as individual links.

⁷ The term trunking has many different definitions in networking, however in this section we use it to describe static link aggregation.

⁸ The term system is used to refer to any device that supports 802.3ad.

Link aggregation control protocol

The 802.3ad link aggregation control protocol uses link aggregation protocol data units (LACPDU) to communicate information and states between two systems. The LACPDU structure is a basic IEEE 802.3 frame. It is not a tagged frame. There are two LACP modes that a system can use, active or passive. If a system is in LACP active mode then it transmits LACPDUs to the partner system regardless of the partner's mode. On the other hand, a passive LACP system only sends LACPDUs if its partner system is active.

When LACP is initially enabled between two systems, the systems will begin to send LACPDUs to attempt to agree on aggregation. LACP allows the systems to exchange their link aggregation capabilities. Once they determine that they are capable of link aggregation, the systems attempt to create the link aggregation group, and add links as members of the group.

Identifiers

Each system that participates in a link aggregation group must have a unique system identifier so that the link aggregation control can determine which set of links connect to the same system. The system identifier (ID) is an unique identifier that is comprised of the system priority and the MAC address of the system. LACP uses the system ID to differentiate the devices involved in the link aggregation.

Once LACP can differentiate the systems by using the system IDs, it needs a method of determining which links are capable of aggregation. Two keys are used in 802.3ad link aggregation, the administrative key and the operational key. These keys have local significance to the system they are assigned. The administrative key is a configured value that allows you to assign ports to a link aggregation group. Changing the administrative keys for a group of ports to the same value means that the ports have the potential to aggregate. If ports do aggregate, then the operational key is changed by the system so that all the ports in the link aggregation group have the same operational key. Only the system is capable of changing the operational key values. This key specifies the currently active key value that LACP is using to identify the link aggregation group. The administrative and operational keys do not necessarily have to be the same value. The important point is that all ports in a link aggregation group will have the same administrative key, as well as the same operational key.

In order for the a link aggregation group to be created, the systems must agree on the link aggregation group identifier. The system ID and operational key are used to create the link aggregation group identifier (LAG ID). The link aggregation group identifier is used to identify links belonging to a particular LAG.




Figure 4-19 Link aggregation group identifier

For example, as shown in Figure 4-19, suppose one switch has a system ID of A⁹ and another switch has a system ID of B. The operational key values of 15 and 17 are assigned to ports 1 and 2 of switches A and B, respectively. In this example, the link aggregation group identifier would be [(A,15), (B,17)] where A and B are the system IDs and 15 and 17 are the operational key values associated with the respective ports. Links with the same link aggregation group identifier are selected for membership to the LAG.

If the links are treated as individual links then the port identifier is also used for further identification. The port identifier is a combination of port priority and port number. In this example, let's assume that port 1 of both switches have a port ID of 30, and port 2 of both switches have a port ID of 40. The link aggregation group identifiers for links 1 and 2 would be represented by [(A, 15, 30), (B, 17, 30)] and [(A, 15, 40), (B,17,40)], respectively.

Some of the 802.3ad link aggregation components

Once the systems agree on the link aggregation group identifier, the links are assigned to the aggregation group. At this point, the distributor and collector functions are enabled, and the clients can communicate through the link aggregation.

The frame distribution function accepts frames from clients and distributes them to the appropriate port in the link aggregation group. It uses a frame distribution algorithm to determine which port to use to transmit the frame. The 802.3ad standard does not impose a particular frame distribution algorithm, however the standard does state that the algorithm must follow two rules:

- ▶ Maintain the order of frames between two clients.
- ▶ Do not duplicate frames.

The frame collector function passes the frames received from the individual ports in the link aggregation group to the client. The distributor ensures the frames are sent in the correct order, therefore the collector passes frames to the client in the order it receives them.

The collection of the distributor, collector and other functions used for a link aggregation group is called the aggregator. The aggregator is the interface between the client and the link aggregation group. Each port is assigned an aggregator. When multiple links exist in a link aggregation group, the aggregator of the lowest numbered port in the aggregation becomes the aggregator for the entire group. The aggregator is bound to one or more ports, however it is assigned an unique MAC address. Therefore, clients treat the link aggregation as if it is a single link. This MAC address could be the MAC address of one of the ports in the LAG or another unique address.

After the link aggregation is operational, the systems continue to periodically exchange LACPDU's to monitor and control the link aggregation. LACP monitors the validity of the group of links in the aggregation. It determines if there are new ports to aggregate, and it removes a port if its membership is no longer valid. It removes the entire group if there are no remaining member ports.

The LACPDU carries information, such as system priority, system ID, operational key, port priority, as well as whether the system is active or passive. LACP depends upon communicating link aggregation states and information, rather than commands. Each system knows what action to take based upon the information it receives in the LACPDU.

If there is a change to the status of the aggregation, the link aggregation control protocol allows quick redirection of traffic across the remaining links by exchanging LACPDU's. When there is a need to transmit something, such as state changes the system transmits

⁹ The values used in these examples are simply used to illustrate the creation of the link aggregation identifier, and they are not a true representation of the system identifier, operational key, and port identifier.

LACPDUs. Therefore, if a link fails the system notifies its partner using LACPDUs and the traffic is redirected across the remaining links.

Dynamic key changes

In systems where there is a limit on the number of links allowed to participate in a link aggregation group, there needs to be a method of determining which ports join the LAG. For example, suppose two systems only support a 2 port link aggregation, but that four links are connected between the two systems and all of the ports have a key value of 1. In this situation, operational keys may be dynamically modified to differentiate the link aggregation groups. These dynamic key changes will only be made by the system with the highest priority (the lowest system priority value). The highest priority system assigns ports to the link aggregation starting from the highest priority port and working its way down. The highest priority system will assign the two additional ports a different operational key value. In systems where standby links are supported, the two additional links will become standby links. Therefore, if one of the links fail, one of the standby links will be added to the link aggregation group.

4.3.3 Characteristics common to both static and 802.3ad link aggregation

This section describes some of the concepts that are common to both static link aggregation and 802.3ad link aggregation.

The main goals of link aggregation¹⁰ are:

- ▶ Increased data rate - The link aggregation group provides a maximum combined data rate of all of the member ports.
- ▶ Increased availability - The redirection of traffic across the remaining member ports if a link failure occurs.
- ▶ Load balancing - The distribution of the frames across multiple links.
- ▶ Fast re-configuration - Automatically redirects the frames if a link fails so that the client experiences no noticeable impact.

Distribution methods

When transmitting frames across a link aggregation, the system uses a distribution algorithm to determine which physical links the traffic goes through. The distribution algorithm specifies how the system load balances the outbound traffic across the ports in the link aggregation. The partner device on the other side of the connection determines the inbound traffic balancing.

The distributor function must operate in a manner that ensures the frames belonging to a given conversation between two MAC clients use the same physical link. Depending on the algorithm, the distributor sends frames belonging to separate conversations over different links.

Choosing from the different distribution methods

The optimal distribution method depends on the network topology. Use the configuration option that can best divide the traffic in the network. The option that provides the most variety leads to the best balancing of traffic over the ports in the link aggregation group. The distribution methods are not based on traffic load, but rather they are address based. They determine which physical port to transmit traffic to by using MAC addresses or IP addresses.

¹⁰ In this section, the term link aggregation refers to both static and 802.3ad link aggregation.

We use the following examples to demonstrate that the network topology must be taken into consideration when choosing the best distribution method.




Figure 4-20 Many to one example

The example illustrated in Figure 4-20 represents an environment where many servers typically send traffic to one particular device, such as a router. In this case, the servers need to communicate to other servers or clients in a different subnet. Therefore, the packets will have to be forwarded to a router. The router will then forward the traffic on to the appropriate destination. If the distribution method being used by switch 1 is the destination MAC address distribution method, then switch 1 distributes the frames across the ports based on the destination host's MAC address. The packets going to the same destination (router's MAC address) will travel over the same physical link, and therefore this method does not lead to the optimal utilization of all the available links. In this type of environment, using the source or destination IP method may lead to greater distribution of traffic across all of the ports.




Figure 4-21 Many to many example

The example in Figure 4-21 illustrates an environment where many stations are communicating directly with many other stations. Therefore, the frames being forwarded across the link aggregation will have many different source MAC addresses, and they will be destined to many different destination MAC addresses. In this environment, using the source or destination MAC and/or the source or destination IP may lead to the best distribution of traffic across all of the links.

Requirements and characteristics of link aggregation

When configuring link aggregation, you must comply with these requirements:

- ▶ To avoid causing a loop in the network, configure and enable link aggregation before making any physical cable connections between the two devices. If converting the ports back to individual links, disconnect all cables before disabling link aggregation to prevent any loops.
- ▶ Configure both sides of the connection for the same link aggregation type.
- ▶ Use the same speed and full duplex for all ports in the link aggregation.
- ▶ All ports must belong to the same VLANs.
- ▶ Each port can only belong to one link aggregation group, either static or 802.3ad.
- ▶ Link aggregation supports point-to-point aggregated links only.

Some of the characteristics of link aggregation are as follows:

- ▶ You can have non-continuous ports within a link aggregation group, therefore you can use non-sequential port numbers, such as Ext 1, Ext 2, and Ext 4.
- ▶ Typically, the lowest numbered port becomes the master port. When you configure this port, it configures the entire group of ports. The spanning tree uses this port to calculate the port cost and the state of the group.
- ▶ The switch treats the group as one when moving ports to/from a VLAN.
- ▶ The spanning tree protocol treats the entire group of ports as one logical port.

4.4 Link aggregation configuration examples

The ESM supports the configuration of a maximum of two link aggregation groups. It supports the bundling of two, three, or four ports in a link aggregation. Therefore, when using the ESM the maximum data rate for a link aggregation is 4 Gbps. By default, the ESM does not have link aggregation enabled, and therefore each port is treated as an individual link.

We use the ESM Web interface to demonstrate both static and 802.3ad link aggregation configuration in this section. The two examples show 2-port link aggregation between two ESMs using a port speed of 1000 Mbps. These examples assume that the ports belong to the same VLANs and that a session has been established to the ESM Web interface.

Both static link aggregation and 802.3ad link aggregation have the following configuration requirements:

- ▶ Configure the ports to the same speed and to full duplex.
- ▶ Select the distribution method. The distribution method you choose on the ESM determines the outbound traffic distribution. Configure the distribution method on the other side of the connection to balance the traffic in both directions. The ESM uses distribution methods based on MAC addresses or IP addresses. The ESM supports the use of the following distribution methods:
 - For non-IP packets:
 - Source MAC address
 - Destination MAC address
 - Source and destination MAC addresses

- For IP packets:
 - Source MAC address
 - Destination MAC address
 - Source and destination MAC addresses
 - Source IP address
 - Destination IP address
 - Source and destination IP address

Remember, do not connect the cables between the two devices until after you complete the configuration. Complete the following steps for each switch.

Configuring port speed and duplex

Configure the ports for the same speed and full duplex mode. Although you can use auto-negotiation, we recommend that you use fixed port configuration.

1. Select **Configuration -> Port Settings -> Configure Ports**.
2. Click the first port (**Ext 1**) in the picture of the ESM on the top of the window. You will see a window similar to Figure 4-22.

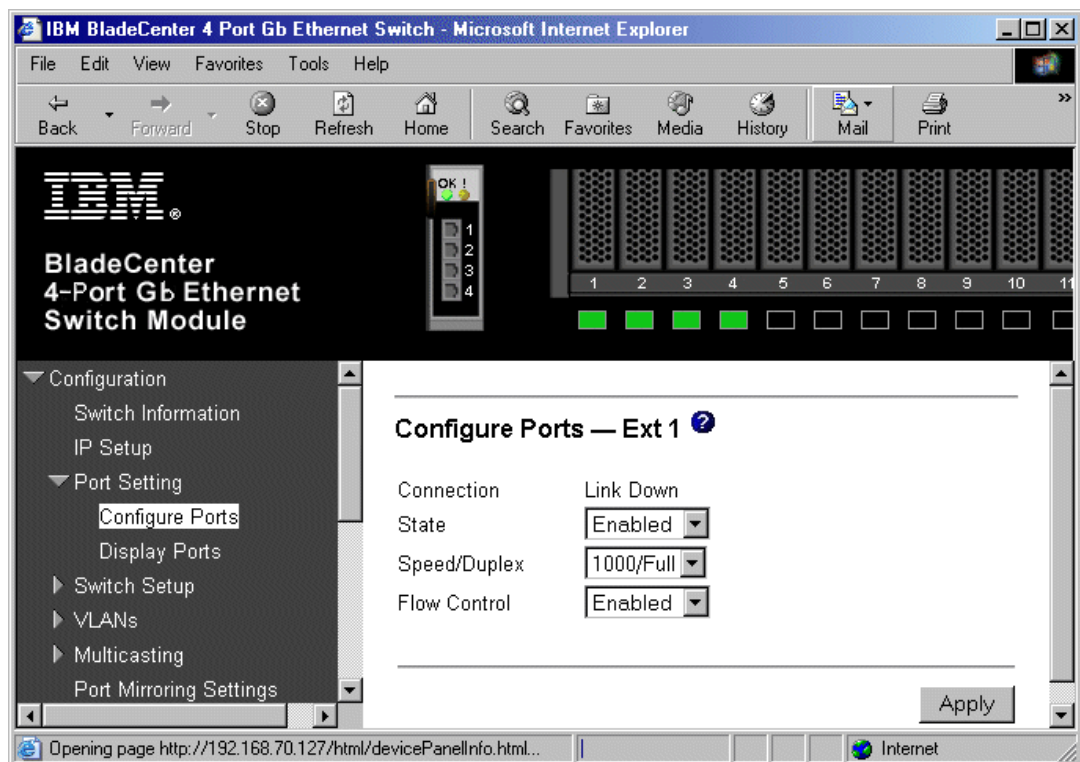


Figure 4-22 Configure external ports

3. Select **1000 Mbps/Full Duplex** from the Speed/Duplex list box and **Enabled** from the State list box. Click **Apply**.
4. Repeat steps 2 and 3 for the Ext 2 port.
5. To verify the port settings, select **Configuration -> Port Setting -> Display Ports**. You will see an example of ports 1 and 2 configured as 1000 Mbps full duplex similar to Figure 4-23.

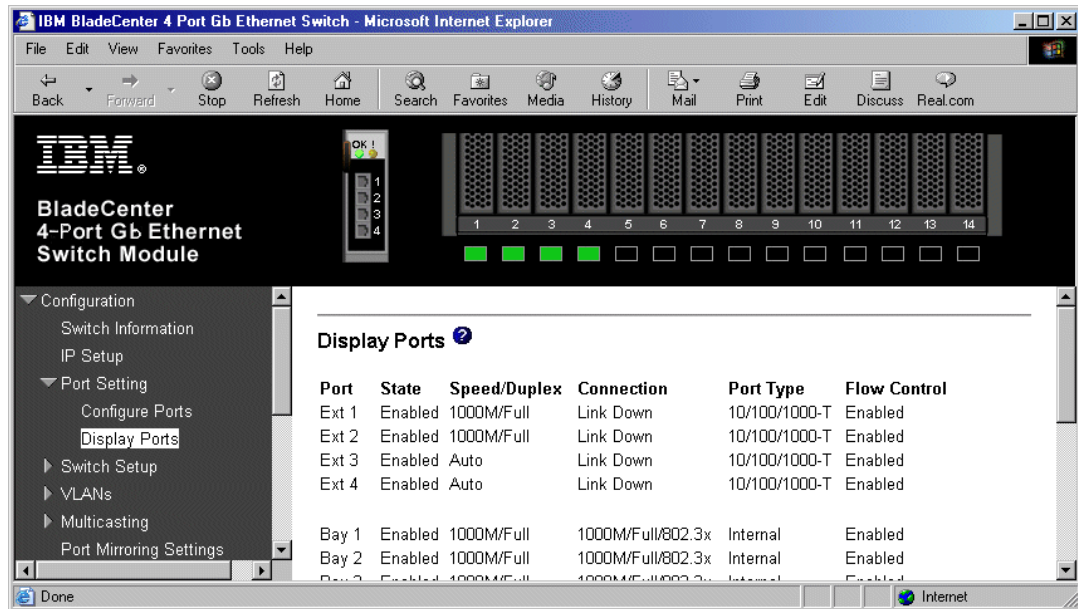


Figure 4-23 Display external port configuration

Selecting the distribution method

To transmit frames across a link aggregation, the switch uses a distribution algorithm to determine which physical links the traffic goes through. Both static link aggregation and 802.3ad link aggregation require a distribution method that specifies the outbound load distribution algorithm. In order to effectively utilize the bandwidth available for the link aggregation group, select the distribution method that best divides the traffic across all of the ports in the group.

To configure the distribution method for the ESM, use the following steps:

1. Select **Configuration -> Link Aggregation -> Link Aggregation Settings**.
2. Select the Distribution Method from both the Non-IP Packet list box and from the IP Packet list box. Refer to Figure 4-24.
 - By default, the ESM uses the source MAC address distribution method for both the non-IP packets and the IP packets. Using the source MAC address method means that the switch distributes the frames across the ports in the link aggregation based on the source MAC address, therefore all frames coming from the same source MAC address travel over the same physical link. Use this method in an environment where many stations are sending traffic to many stations on the other side of the link aggregation connection. For this example, we are using the source MAC address for both Non-IP packets and IP packets.

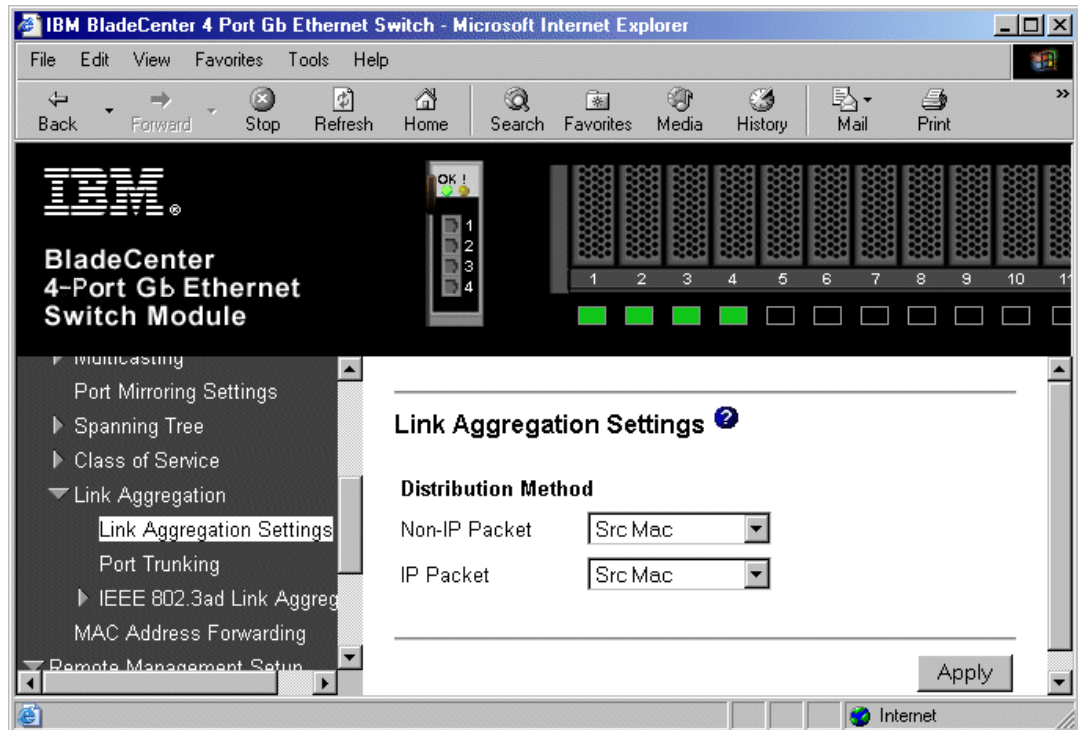


Figure 4-24 Select the distribution method

3. Click **Apply**.

4.4.1 Configuring static link aggregation

To configure static link aggregation in the ESM, you must specify the ports that belong to the group and enable the static aggregation. The ESM supports a maximum of two trunking groups, and they are identified as group ID 1 and 2. Configure both sides of the connection, using the following steps:

1. Select **Configuration -> Link Aggregation -> Port Trunking**. Refer to Figure 4-25 to see a similar port trunking configuration.

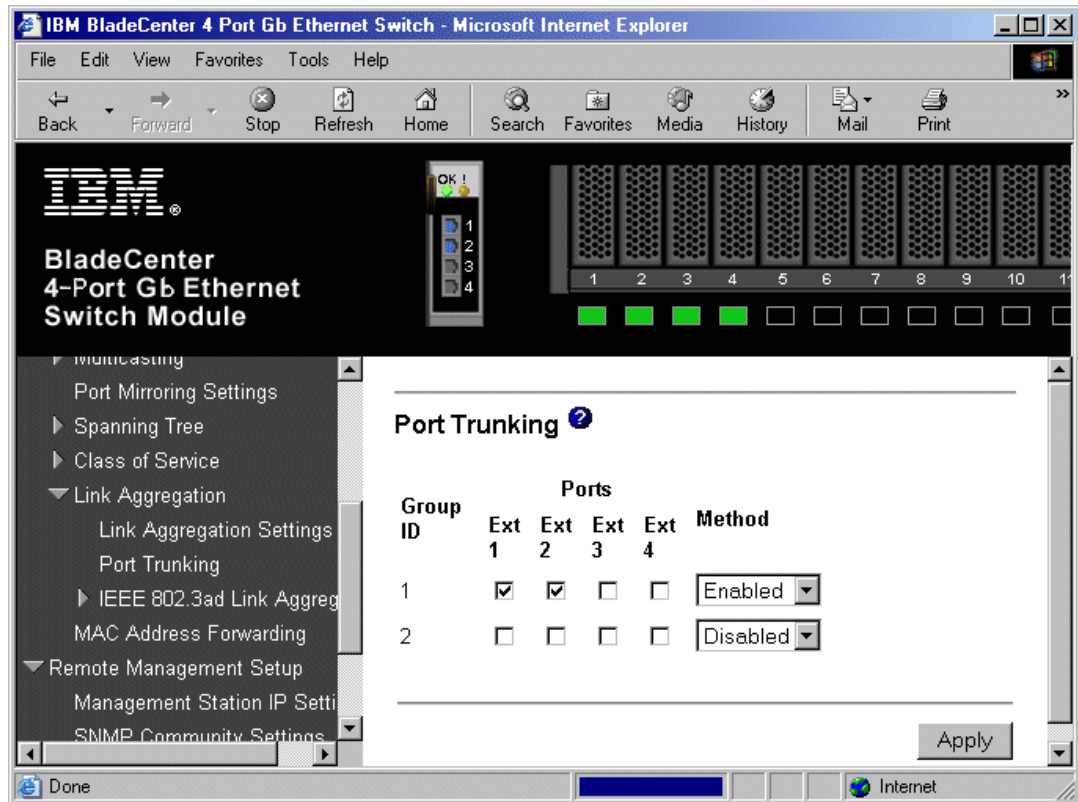


Figure 4-25 Configure static link aggregation

2. Select the ports that you want to configure to belong to the group. Select **Enabled** from the Method list box. Click **Apply**. The port static link aggregation panel should look like Figure 4-25.
3. Configure the switch on the other side of the connection. Repeat steps 2 and 3 on the other switch.
4. Connect the cables between the two switches. Since we have fixed the port speed, use cross-over cables. Verify that the trunk ports are blue in the picture of the ESM at the top of the panel.
5. To verify the status of the spanning tree, use the Spanning Tree Port Settings panel. Click **Configuration -> Spanning Tree -> STP Port Settings**. You will see a window similar to Figure 4-26. Both switches should show forwarding status for both ports.

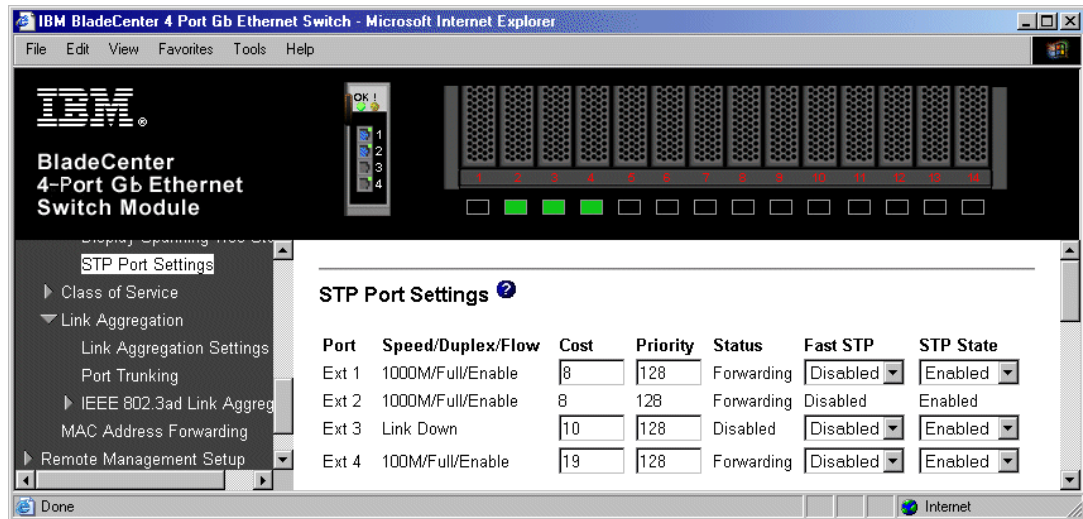


Figure 4-26 Display the spanning tree port status

4.4.2 Configuring IEEE 802.3ad link aggregation

If using 802.3ad link aggregation, both ends of the connection must support 802.3ad. Use the following steps to configure 802.3ad link aggregation:

1. Select **Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Setup IEEE 802.3ad Link Aggregation**. You will see the Setup IEEE 802.3ad Link Aggregation panel similar to Figure 4-27.

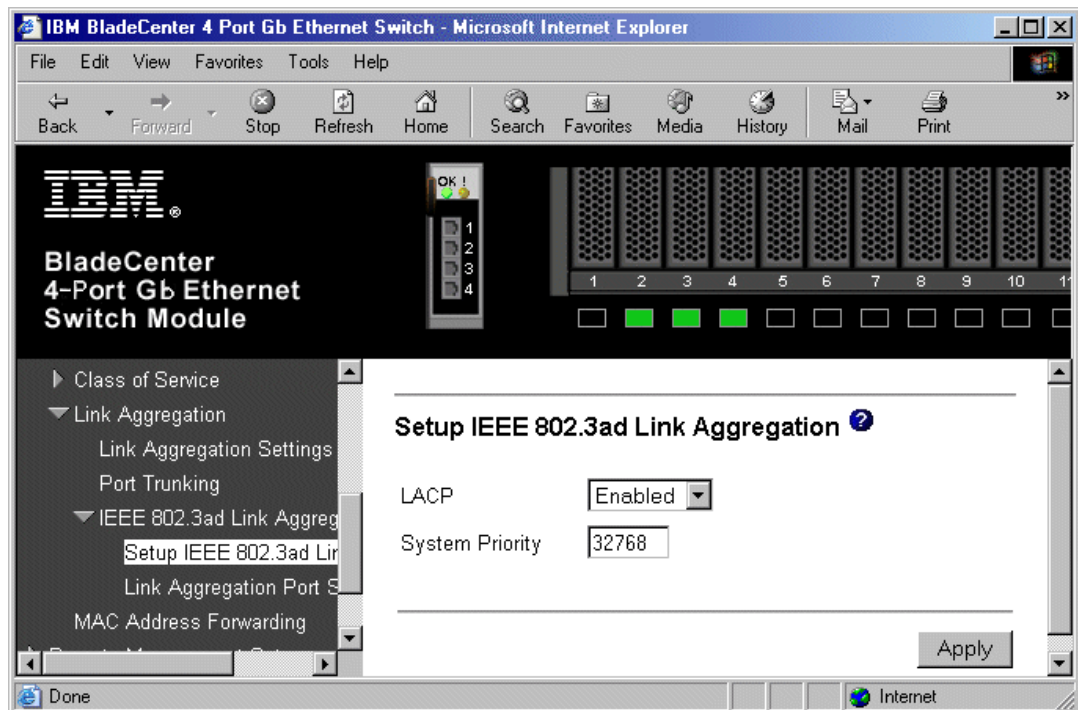


Figure 4-27 Setup IEEE 802.3ad link aggregation

2. Select **Enabled** from the LACP list box. By default, the ESM acts as an active LACP switch when LACP is enabled.

3. Enter the desired number in the System Priority field. The default system priority is 32768. This value can range from 1-65535. This value, along with the system MAC address is used to create the system ID. The lower the value the higher priority a system will have. Enter a low numeric value if you want the switch to have a higher priority. In most cases, it can be left as the default value. Click **Apply**.
4. Select **Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Link Aggregation Port Settings**. This will bring you to the Link Aggregation Port Settings panel. Refer to Figure 4-28.
5. Choose the ports Ext 1 to Ext 2 from the Configure external port from and to list box.
6. Enter the port priority in the Priority field. The default for this field is 128. This value can range from 1-255. The switch uses this value to create the port identifier. The port identifier is created by combining the port priority and the port number. The lower value means the port has a higher precedence over other enabled link aggregation ports in the group. In most cases, you can leave this value as the default.
7. Complete the Administrative Key field. This value can range from 1-255. Assign the same admin key to all of the ports. The ports with the same key value have the potential to aggregate.
8. Select **Enabled** from the Mode list box. This enables the ports to join the LAG.

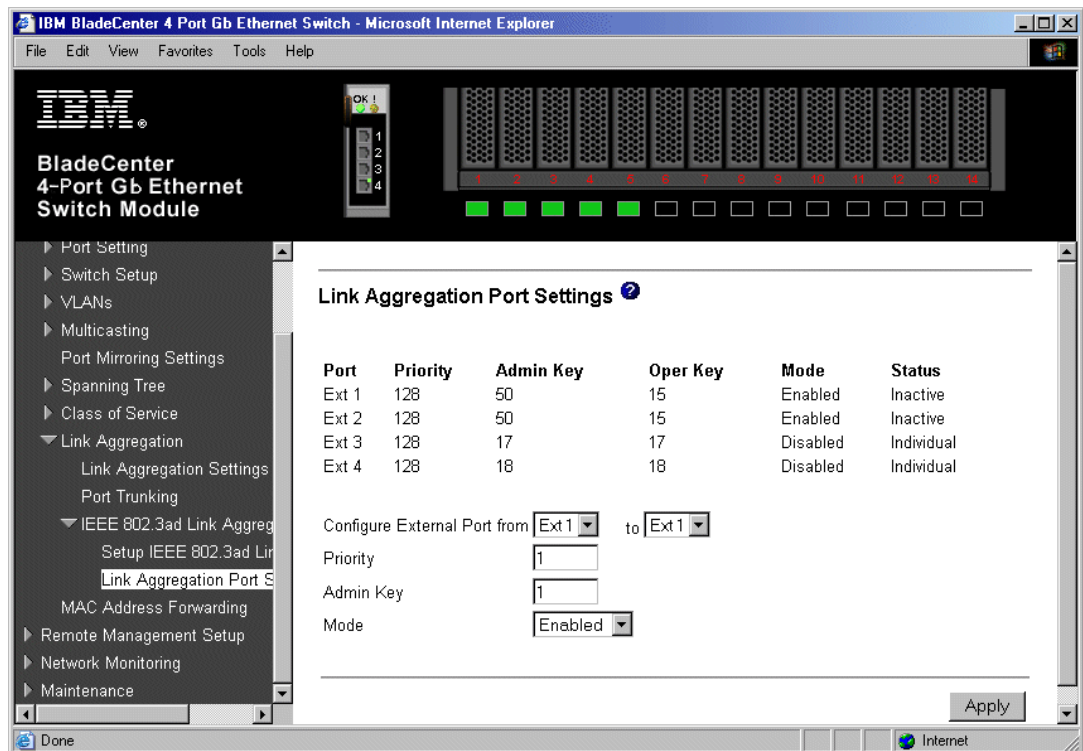


Figure 4-28 Display link aggregation port settings

9. Click **Apply**.
10. Repeat steps 1 on page 68 through 9 for the other switch.
11. Connect the cables between the two switches. Since the port speed is fixed, you must use cross-over cables.
12. To view the current state of the link aggregation group look at the status column of the Link Aggregation Port Settings panel. The ports that are part of the link aggregation group will

have a status of Active. See Figure 4-29. Otherwise, the switch displays the status as Individual for single links or Inactive for ports that can not participate in a LAG.

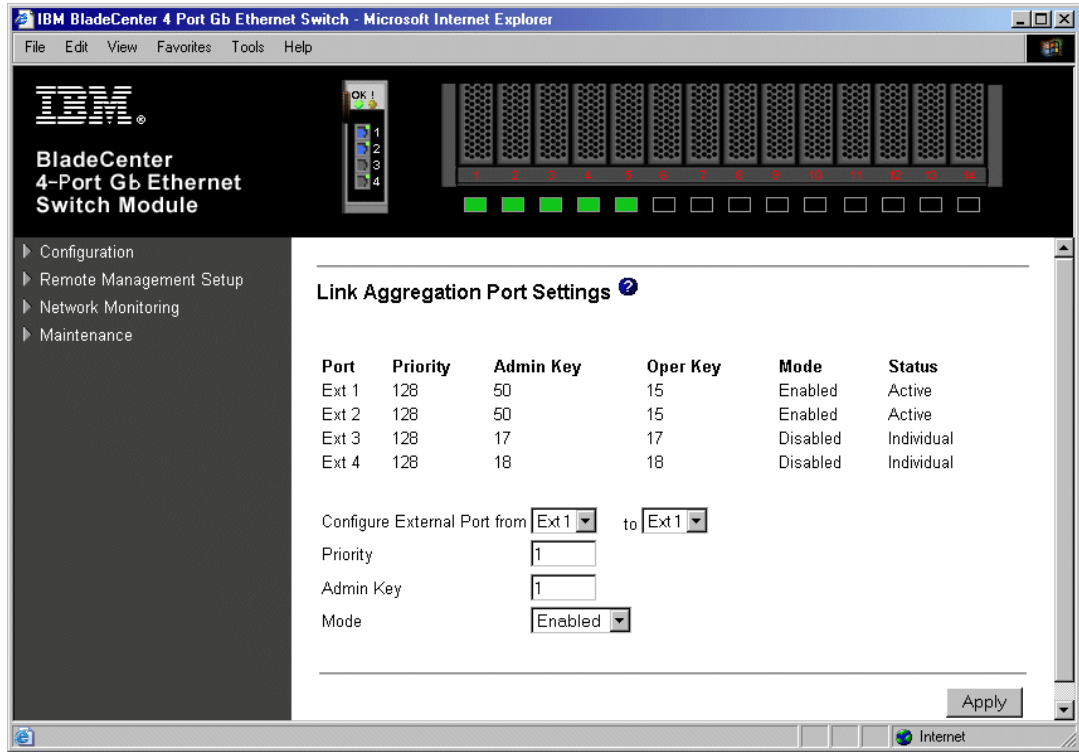


Figure 4-29 Verify the status of 802.3ad link aggregation

13. We also recommend that you verify the status of spanning tree. Always enable spanning tree when connecting multiple links between two devices. Click **Configuration -> Spanning Tree -> STP Port Settings**. This displays the STP Port Settings panel similar to Figure 4-30. Check to ensure that the status column for the ports in the link aggregation group show forwarding. Spanning tree treats all the ports as one logical link. Notice that you can only configure the Ext 1 port. This is the port that is now used to configure all settings for the link aggregation group.

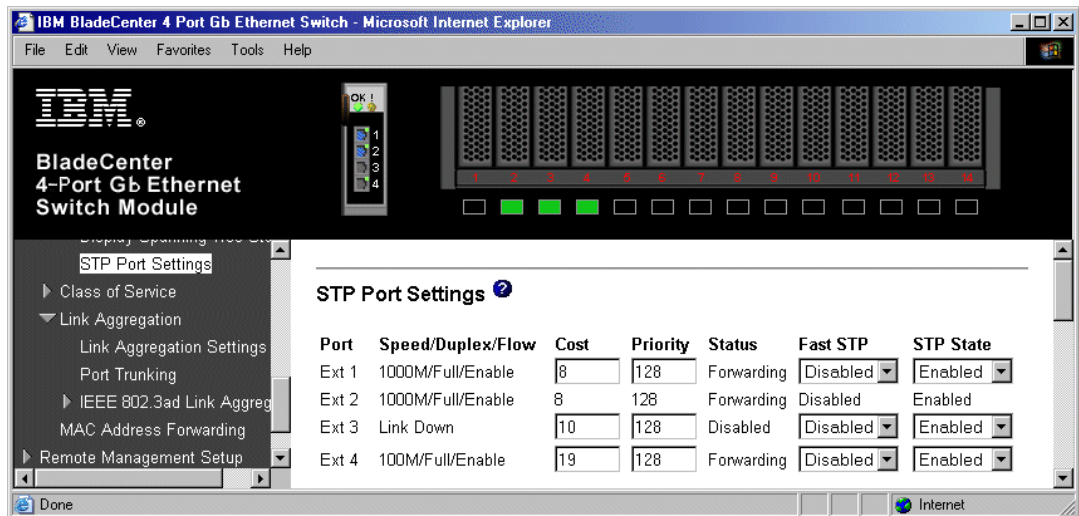


Figure 4-30 Verify spanning tree status

14. To view the status of the link aggregator, click **Network Monitoring -> Application Status -> Link Aggregator**. Refer to Figure 4-31. Ensure that Aggregating is displayed in the status column for the link aggregator.

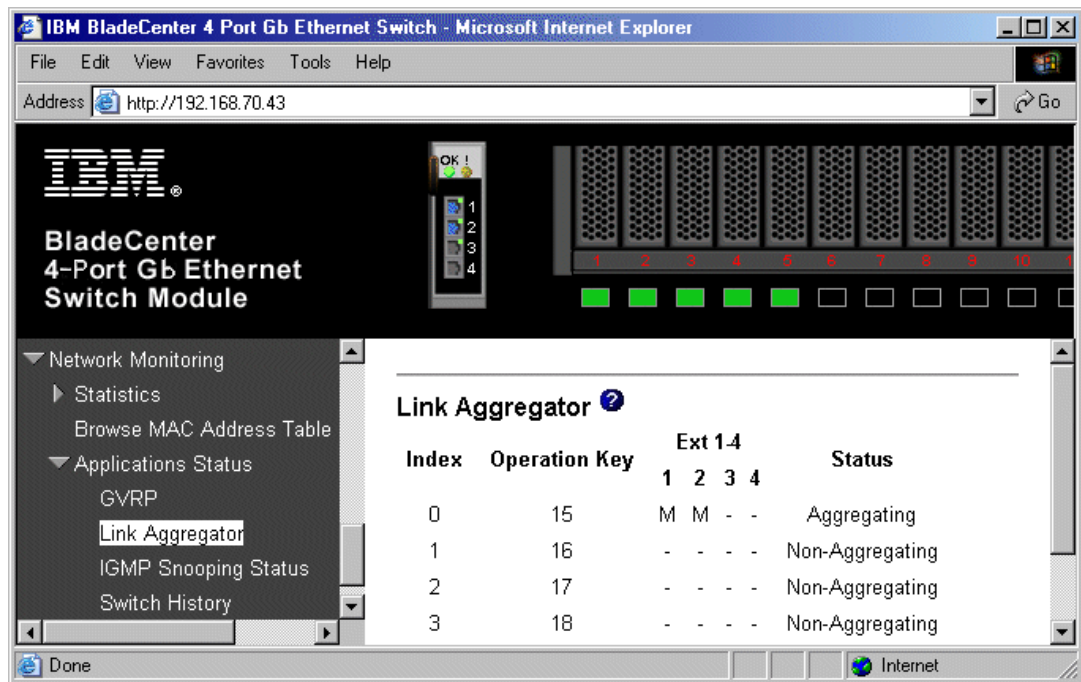


Figure 4-31 View the status of the link aggregator

15. When you want to check how traffic is being distributed, select **Network Monitoring-> Statistics -> Port Utilization**. Refer to Figure 4-32.

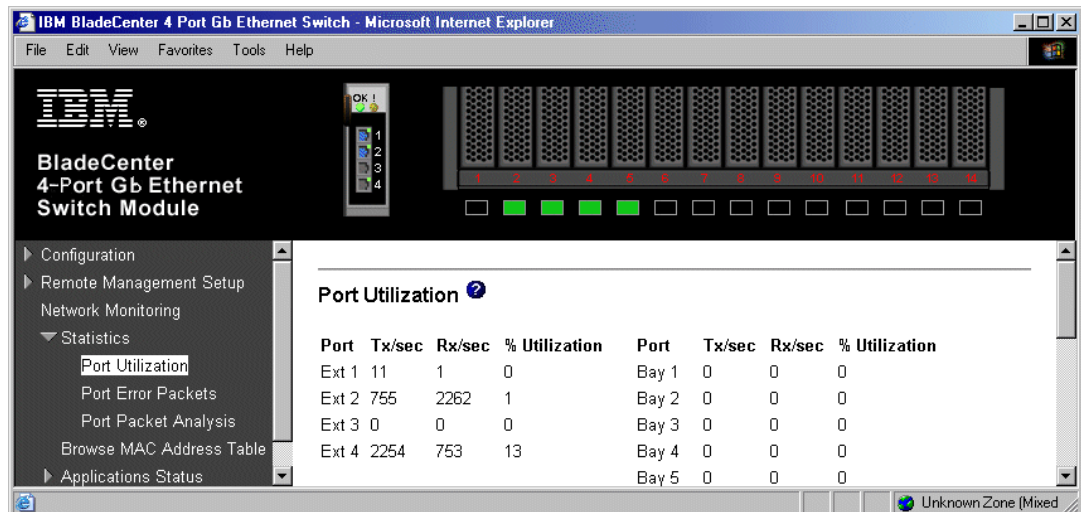


Figure 4-32 Check port utilization



Deploying the IBM @server BladeCenter in a Cisco environment

This chapter provides information on deploying the BladeCenter Ethernet Switch Module (ESM) in an infrastructure based on Cisco Systems switches. Specific ESM and Cisco configurations are covered for a number of possible scenarios.

5.1 Introduction

The current version of the ESM, known as the *4-Port Gb Ethernet Switch Module* within the IBM BladeCenter, is a fully functional, standards-based, layer 2 switching device. It provides the necessary features and functionality to attach the BladeCenter components to virtually any standards-based network infrastructure, including those based on Cisco Systems equipment.

While the ESM and Cisco share a common set of specifications with which they are compatible (this was covered in Chapter 1, “Introduction to the IBM ^ BladeCenter 4-Port Gb Ethernet Switch Module” on page 1), it is important to note that there are some features and standards that are fairly commonly deployed in Cisco infrastructures that are not found in the ESM. A partial list of Cisco enhancements and standards that are commonly encountered in a Cisco infrastructure, but that are not supported by the current ESM, follows:

- ▶ BackboneFast - To aid in the rapid convergence of layer 2 networks
- ▶ UplinkFast - To aid in the rapid convergence of layer 2 networks
- ▶ UDLD - UniDirectional Link Detection, to reduce the possibility of STP loops
- ▶ CDP - Cisco Discovery Protocol (to aid in management and troubleshooting)
- ▶ ISL - VLAN Trunking (Cisco Proprietary)
- ▶ PAgP - Port Aggregation Protocol (Cisco EtherChannel)
- ▶ VTP - VLAN Trunking Protocol (similar to GVRP)
- ▶ DTP - Dynamic Trunking Protocol (to auto-negotiate trunk type and state)
- ▶ RADIUS - Centralized administrative control of access to the switch
- ▶ TACACS+ - Centralized administrative control of access to the switch
- ▶ VMPS - VLAN Management Policy Server (only on some Cisco switches)
- ▶ PVST+ - Per VLAN Spanning Tree
- ▶ 802.1w - Rapid Reconfiguration Spanning Tree (enhancement to 802.1D)
- ▶ 802.1s - Multiple Spanning Trees (enhancement to 802.1Q)

5.2 Architecture summary

Full discussions of various network architectures are beyond the scope of this document. That said, the following section contains some introductory material on this subject.

5.2.1 Datacenter networks introduction

While what constitutes a datacenter and its associated architecture is open to discussion, it is reasonably expected that many BladeCenter deployments will occur in some form of datacenter environment. It can also be safely assumed that datacenter networks are fairly unique environments from other parts of the enterprise network, and can be more demanding in their requirements.

High availability, performance and security tend to be far more critical, with elements of firewalls, content caching/load balancing and security playing a much larger role than in other areas. With this stated, this document does not attempt to discuss all of the various elements that may constitute a datacenter, but rather focuses only on attaching the BladeCenter to Cisco switches.

Note: This document should not be considered a substitute for an Architectural Reference document, and the examples provided should not be employed without first understanding the specific needs of your particular environment.

For those seeking more details and information on Cisco's definition of datacenter architectures, along with recommendations, please visit:

http://www.cisco.com/en/US/partner/netsol/ns110/ns53/ns224/networking_solutions_packages_list.html

Note: A CCO (Cisco Connection Online) user ID and password are necessary to access this information. A CCO ID can be obtained by going to the following page and registering:

<http://tools.cisco.com/RPF/register/register.do>

5.2.2 Common Cisco components

As previously mentioned, the datacenter is a rather unique environment, with demands that can far exceed those of other areas of the enterprise network. With this in mind, it is recommended that certain highly robust and scalable designs and platforms be used in the datacenter.

With regard to design, Cisco traditionally proposes a three level architecture when referring to enterprise environments. This architecture promotes both robustness and scalability and is made up of an access level, a distribution level and a core level. When working specifically with datacenter architectures, a common current approach is more often a two level design:

- ▶ An access level (to attach servers). Sometimes referred to as the front-end layer, this is primarily an L2 (OSI) network.
- ▶ An aggregation level (to tie the components of the datacenter to each other and into the rest of the enterprise network via the enterprise "core"). In the aforementioned three level design, this aggregation level would be referred to as the distribution level. One could also look at this as being the "core" level of a datacenter, where elements of OSI L2 and L3 come together.

With a two level design such as this in mind, the BladeCenter can be attached to this type of architecture at one of two points, as follows.

- ▶ Attaching the ESM's of the BladeCenter to the access layer switches. This in essence adds a third layer to the two layer architecture.
- ▶ Attaching the ESM's of the BladeCenter directly to the aggregation layer switches. This maintains the two layer approach, with the ESM's essentially becoming the access layer.

As both of these scenarios have their advantages, architecture selection is based on specific customer requirements. When ESMs are attached to the access layer, the network architecture provides more flexibility in handling unexpected growth. When ESMs directly interface to the aggregation layer, the network benefits from a reduced number of network levels.

Attaching to the access layer tends to be strictly an OSI L2 connection, while attaching to the aggregation level may be at L2 or L3. While direct L3 connectivity is beyond the scope of this document, and all configuration examples given assume L2, it is important to consider the pros and cons of attaching to the network at L2 or L3. An example of this might be that connecting at L2 can lead to complications with the different implementations of STP employed by the ESM and the Cisco equipment, while connecting at L3 can lead to inefficient use of IP subnet addresses and limit growth potential.

With regard to platforms, the platform of choice at the aggregation level is the Cisco 6500 series of switches. This modular platform offers high performance and availability, and operation in multiple OSI layers, as well as investment protection. Beyond switching and routing, the 6500 also currently offers various service modules (such as Firewall, Content services and Intrusion Detection) that allow this platform to serve many roles in the datacenter, as well as the role of the aggregation element.

Our testing was performed utilizing the @server BladeCenter ESM and the Cisco switches, 3550 and 4003. Providing our readers with configurations utilizing the Cisco 6500 would have been ideal but the switch was unavailable during this project.

Note: As mentioned above, it is not the intent of this chapter to promote the use of the 3550 or the 4003 for datacenter deployment.

5.3 Guidelines, rules and comments

Before discussing specifics of some of the combinations of configurations available in this environment, it is necessary to discuss some basics of configuration and operation as used during the creation of this document.

All configurations were tested with the following equipment and code revisions:

- ▶ IBM BladeCenter 8677
 - Four single CPU Blade servers
 - One with Windows 2000 server with SP3
 - Three with Red Hat Linux 7.2
 - One Management module
 - Two Ethernet Switch Modules (ESM)
 - ESM (FRU 59P6620):
 - Boot PROM version: 00.00.04
 - Firmware version: 00.00.62
 - Hardware version: Rev 2
- ▶ Cisco 3550-24
 - Two ports GBIC based uplinks and 24 ports 10/100BaseT
 - IOS Version 12.1(13)EA1
- ▶ Cisco 4003
 - WS-X4012 Supervisor 1
 - CatOS Version 7.5(1)
 - ROM Version 6.1(4)
 - WS-X4232-GB-RJ 2 ports Gig uplinks and 32 ports 10/100BaseT

All configurations and testing were performed on clean systems. In the case of the ESM, the modules were restored to factory default via the GUI interface option (click **Maintenance -> Factory Reset**) prior to any configuration. In the case of the Cisco 3550, a combination of **write erase** and **delete flash:vlan.dat** commands were issued, followed by a **reload**. For the 4003, a **clear config all** command was issued, followed by a **reset** command.

Important: Performing the operations above will result in all configuration data being lost on the mentioned devices, which in turn *will* lead to network downtime if performed on production systems. The commands are presented here only to indicate preparation performed prior to commencement of lab testing.

If working in a production network, be sure to understand the consequences of any commands issued. Failure to completely understand the operation of commands can lead to network-down conditions.

Also note that available features and command syntax can be different with different versions of code. This document was prepared utilizing the features and syntax from the aforementioned revisions of code, and as such, might vary from other revisions. For complete and current lists of available features and commands for these products, please visit the IBM or Cisco Web sites.

5.3.1 Rules for attaching the eServer BladeCenter to a Cisco infrastructure

This section contains information on things to consider when attaching the BladeCenter ESM to a Cisco infrastructure. It is highly recommended that you review this entire section prior to any initial configuration changes.

Guidelines and comments - General

These are some comments and recommendations that are general in nature, and not part of a specific technology covered elsewhere.

Cable selection

Selection of the cable type (cross-over or straight-through) to use between the ESM and a Cisco switch can vary, depending if you are using auto-negotiation or not on the ESM. When a port on the ESM is configured for auto-negotiation, it will automatically try to determine the best speed, duplex and MDI/MDI-X configuration. The auto MDI/MDI-X usually permits any cable (straight-through or cross-over) to be used, as the ESM can adjust for either cable. When you change the ESM from Auto to a fixed speed/duplex, this also disables the auto MDI/MDI-X feature. When this happens, a cross-over cable *must* be used to connect an ESM to a Cisco switch.

Based on this, it is recommended that you always use a cross-over cable between the ESM and a Cisco switch, to ensure that no matter how the speed/duplex is configured, the link will continue to operate.

Speed/duplex selection

The decision to allow a port to auto negotiate its speed and duplex, or to force it to a set value, is a subject of frequent debate. Testing in the lab has shown that the ESM can correctly negotiate the link when attaching to Cisco switches. With that said, conventional wisdom indicates to always hard-code critical links within a network. On that note, it is recommended that you hard-code the speed and duplex between the ESM and the Cisco switch to the fastest combination supported by both parties. In the case of the examples presented in this chapter, this will always be set to 100/full, which is the maximum link speed of the Cisco switches in the lab while preparing this chapter.

Important: In production environments, it is strongly recommended that 1000BaseT connections be used (available on all Cisco platforms suitable for datacenter environments) to ensure the best possible throughput.

Use of the term trunk

There are a number of terms and acronyms used in the industry that have proven to be the source of much confusion. One such term is the word *trunk* or *trunking*. This term has been used to describe a number of technologies, most commonly the act of bundling links together to increase performance and reliability, and the act of carrying multiple VLANs on a single connection.

Unless otherwise stated, in this chapter, the following definitions apply:

- ▶ *Trunk* or *trunking* - The act of carrying multiple VLANs on a single connection (the connection might be a single link, or a group of links aggregated together to form a Link Aggregation Group). The IEEE specification for VLAN trunking is 802.1.Q.
- ▶ *Aggregation* or *link aggregation* - The act of bundling multiple physical links in to one logical link, for the purposes of increasing throughput and/or offering increased reliability. Link Aggregation is often referred to as EtherChannel in the Cisco world. The IEEE specification for link aggregation is 802.3ad (now part of 802.3 2002).

Note that there are some places in the Web interface to the ESM where what is defined as aggregation above is shown in the menus as *trunking*. In those cases, it will be necessary to display the menu option as it exists on the page (for example Port Trunking) but this still refers to aggregation as defined in this document.

Use of the term Native VLAN

The term Native VLAN is used throughout this chapter to describe a single designated untagged VLAN in an 802.1Q trunk. The 802.1Q specifications does not define this term, but the concept of untagged VLANs on a trunk is defined within the specification. Cisco has adopted this term to describe a VLAN that provides, among other things, backward compatibility with a device that might not understand 802.1Q such that at least some communications can take place across this link. In Cisco networks, the Native VLAN is most often VLAN 1, but can be changed to any VLAN. Any device connecting to a Cisco switch via an 802.1Q trunk, should define at least one untagged VLAN, and it must match the untagged VLAN on the Cisco switch (for example, they both define the untagged VLAN as VLAN 1). Note that both the ESM and all Cisco switches default to VLAN 1 as being the untagged VLAN. More comments about the operation and configuration of the Native VLAN can be found elsewhere within this chapter.

Guidelines and comments - VLANs and trunking

On the Cisco external switch, always configure the trunks connecting to an ESM as 802.1Q. The BladeCenter does not support Cisco ISL (Inter-Switch Link).

On the Cisco external switch, always configure the connections connecting to an ESM as a *trunk* port, rather than an *access* port. This can help to prevent unintentional Spanning Tree loops.

On the Cisco external switch, always use the **nonegotiate** option with connections going to an ESM, to ensure a trunk can be established. The BladeCenter does not support Cisco DTP (Dynamic Trunking Protocol) to determine trunk type.

GVRP is a protocol that allows dynamic administration of VLANs across a layer 2 network. At this time, it is not commonly found in Cisco infrastructure deployments, which instead more commonly make use of a Cisco protocol (VTP) for dynamic administration of VLANs. Based on the lack of broad use of GVRP in Cisco networks, and the fact that GVRP is not supported on some Cisco switches, this document only discusses non-dynamic ways to administer VLAN's between the ESM and a Cisco network.

It is recommended that the Native 802.1Q VLAN be VLAN 1 (the default for both the ESM and all Cisco switches). While you can make any VLAN the Native VLAN, it can become confusing when mixing Mono Spanning Tree and PVST+, and the end result can be unexpected and undesired operation of the network brought on by Spanning Tree loops.

Regardless of what number is assigned to the Native VLAN, it must be the same on both sides of the trunk (most commonly both sides call the Native VLAN, VLAN 1). Having different Native VLANs on either side of the trunks can lead to unexpected and undesired operation of the network brought on by Spanning Tree loops.

Make sure that the trunk is *carrying* the native VLAN (usually VLAN 1). There is some control available to prevent a given trunk from carrying a given VLAN. If using this feature, do not block the Native VLAN. By default, both the ESM and Cisco switches carry VLAN 1 on every trunk.

Guidelines and comments - Link Aggregation (EtherChannel)

Link Aggregation Control Protocol (LACP) is the preferred method for performing link aggregation between the ESM and Cisco devices. If the Cisco switch being used does not support LACP, then one has the choice of upgrading the switch to newer code that supports LACP, or to use static aggregation on both sides.

It is always advisable to check the release notes of the revision of Cisco IOS and CatOS being utilized to ensure that the feature set you require is present.

Links that are part of an aggregation group *must* have the same characteristics (speed, duplex, have the same trunk settings, carrying the same VLANs). Having links with different characteristics will result in unexpected issues, including aggregations failing to form between the ESM and Cisco switches. For troubleshooting LACP issues, see “Troubleshooting: LACP Aggregation link will not form” on page 179.

Guidelines and comments - Spanning Tree

It is recommended that you always leave STP enabled on both the BladeCenter ESM and the attaching Cisco switches. Disabling STP can lead to loops in the network, which can lead to an unexpected disruption of traffic. At a minimum, the Native VLAN (usually VLAN 1) on the Cisco external switch and the ESM should have STP enabled.

On the Cisco external switch, do not use a link configured as *Access* to connect to a Mono Spanning Tree switch (such as the ESM). Always use a *Trunk* link. Forcing the link to be a trunk link can reduce the likelihood of Spanning Tree loops.

The ESM has support for a feature referred to as *Fast STP*, which can be enabled on an individual link or an aggregated link (in whole). This option, while not discussed in the provided examples, can be used to reduce the time it takes the ESM to begin forwarding packets after a link is brought up or down on the ESM. With default configs, a link using STP will usually take approximately 30 seconds to begin forwarding after first being brought up. This is the result of STP going through its various states (15 seconds in the listening state, 15 seconds in the learning state) to ensure a loop in the network does not exist. *Fast STP* causes the link to go straight from the listening state to the forwarding state (bypassing the learning state). In most cases, this reduces the time it takes to begin forwarding packets from 30 seconds to 15 seconds.

In most cases, you will not want the ESM to become the root switch for the network, as this will more than likely result in sub-optimal flow in the layer 2 network. To prevent this from happening, one should consciously select a non-ESM switch to be root, and set this switch to a lower switch priority than the other switches in the network (good planning also dictates pre-selecting a backup root, but that is beyond the scope of this document). Cisco has several

commands for forcing the election of the root bridge (all involve setting the priority lower than other switches in the network).

When deploying in an environment where an ESM will be directly attaching to multiple Cisco switches in a common layer 2 cloud, one has the choice of letting the ESM and the connecting switches decide which redundant links to block, or manually controlling what links get blocked. It is recommended that you set the path costs on the links between the ESM and the Cisco switch, such that one can be assured that, if all links are operational, a specific set of links will go in to blocking.

The choice of path cost, and thus path selection, in your production environment is up to your network administrator, and will depend on factors such as the network architecture, the location of the root switch, the distance of the root switch to the ESMs and various port cost settings in between. The important thing your network administrator needs to keep in mind is that the ESM is a Mono Spanning Tree device while Cisco switches generally use PVST+.

The following provides an example of what this can mean.

Network description: as shown in Figure 5-1, a single ESM with dual aggregated links, each aggregated link going to separate Cisco switches, each Cisco switch joined to each other via a single link (simulating a layer 2 network beyond the switches).

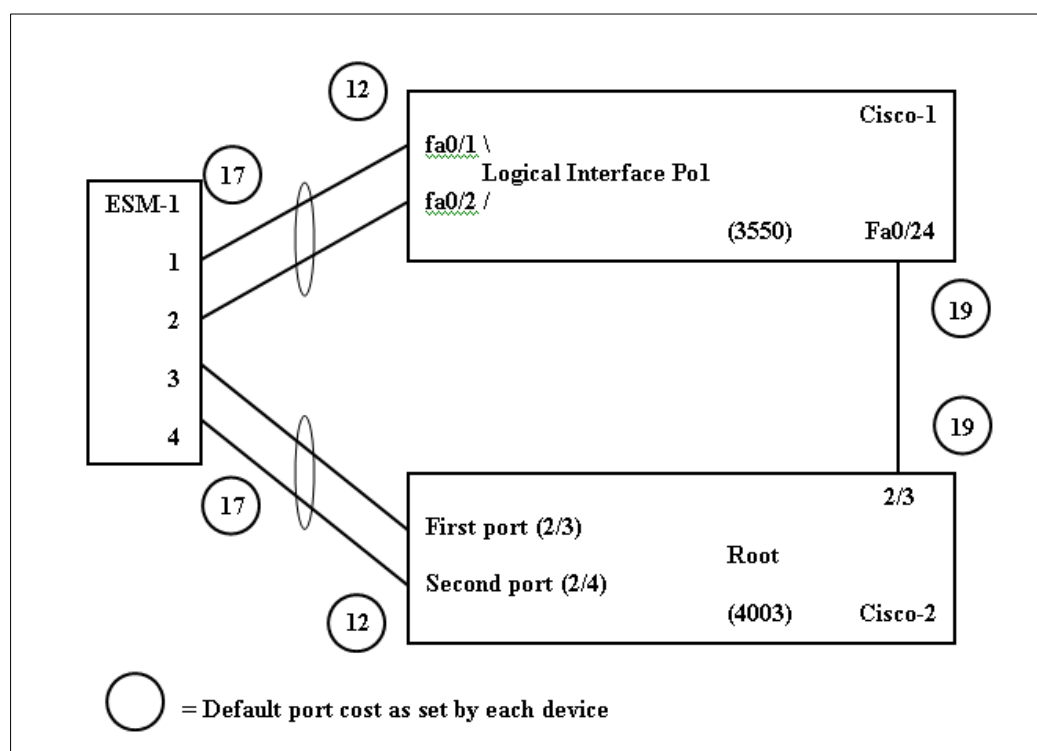


Figure 5-1 Default STP port cost for this example.

In the examples presented in this chapter, where multiple Cisco switches are used to provide redundancy, the Cisco-2 switch is always forced to be the root (chosen for consistence of results for the examples, and not because it is a good design). Having the root directly attached to the ESM is not necessarily recommended in redundant configurations, as flow patterns can become less than obvious (see example below). To simplify decision making, for examples in this chapter that make use of redundant links, the path costs on all the Cisco links to the ESMs are set to 100. While this does not necessarily result in perfect load balancing, it does result in predictable blocking and forwarding conditions (Cisco-1's links to

the ESM(s) will go in to blocking and all traffic from the ESM (or both ESMs if there are more than one) will go through Cisco-2). In the event of a failure of Cisco-2, STA will switch traffic over to Cisco-1 (albeit with the normal Spanning Tree delays).

Based on the default port cost assigned and the location of the root switch (Cisco-2) one might think Spanning Tree would block the link between Cisco-1 and the ESM and any traffic from Cisco-1 to Cisco-2 would flow over the direct link between the two Cisco switches.

While this would be the case if all devices were running either PVST+ or Mono Spanning Tree, it is not the case in this mixed environment.

What will happen in this case is that, traffic on VLAN 1, destined from Cisco-1 to Cisco-2 will follow the predicted path (through the direct link). Traffic on other VLANs, however, would not take this path, but rather would travel through the ESM, as it would appear to be the lower cost path. The following command executed on Cisco-1 demonstrates this result:

```
Cisco-1-IOS#sh span root (cost = cost to root in this command)
```

Vlan	Root ID	Root Cost	Hello Time	Max Age	Fwd Dly	Root Port
VLAN0001	8192 00b0.649c.4c00	19	2	20	15	Fa0/24
VLAN0005	8192 00b0.649c.4c04	12	2	20	15	Po1
VLAN0010	8192 00b0.649c.4c09	12	2	20	15	Po1

Notice that the path to the root for VLAN0001 is via Fa0/24 (the direct link) while the path to the root for VLAN 5 and 10 is via Po1 (the logical interface of the aggregation link toward the ESM). This split traffic is the direct result of joining a PVST+ switch to a Mono Spanning Tree switch, and is probably not the desired result.

What is happening in this example is that the SPT for VLAN 1 is taking in to account the ESM link costs, 17+12=29 (as SPT is common on VLAN 1 for both PVST+ and Mono Spanning Tree devices) and choosing to forward over the direct connection between Cisco-1 and Cisco-2 (lower cost of 19). The other VLANs see the link from Cisco-1 through the ESM to Cisco-2 as a straight connection through to Cisco-2, cost 12 (as BPDU's for VLANs other than VLAN 1 are essentially tunneled through the ESM as multicast packets for that VLAN), and are choosing to forward traffic through the ESM.

It should be noted that as long as the root is not one of the directly attached switches to the ESM, this flow will be unlikely to occur (as other path costs will come in to play that will more then likely offset this undesired behavior). Even so, it helps to be aware of the possibility of this occurring, as only traffic destined for the BladeCenter should be forwarded to the ESM to ensure maximum throughput is available for the blade servers.

For this example, a simple solution is to set the port cost on the po1 interface of Cisco-1 such that it becomes an undesirable path (anything higher than 19 will work) for all VLANs. The following demonstrates this change:

```
Cisco-1-IOS#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Cisco-1-IOS(config)#int po1
Cisco-1-IOS(config-if)#spanning-tree cost 100
Cisco-1-IOS(config-if)#end
Cisco-1-IOS#sh span root
```

Vlan	Root ID	Root Cost	Hello Time	Max Age	Fwd Dly	Root Port
VLAN0001	8192 00b0.649c.4c00	19	2	20	15	Fa0/24
VLAN0005	8192 00b0.649c.4c04	19	2	20	15	Fa0/24
VLAN0010	8192 00b0.649c.4c09	19	2	20	15	Fa0/24

Here it can be seen that, after the change to the port cost, all VLANs are now choosing the direct link as the path to the root.

5.4 Preliminary information on configuration examples

Before we discuss specific configuration examples, it is necessary to discuss some of the bases for all configurations in this chapter.

Some comments on the examples offered

Within the examples provided are syntax and screen captures of commands to complete the desired task. It is possible that a production switch might have configuration commands already in place that conflict with these commands. It is the responsibility of the person configuring the Cisco switch and the ESM to fully comprehend any changes and their resultant consequences. Failure to fully understand the commands can lead to network-down conditions.

Examples showing the use of link aggregation are always shown using an even number of sequential links. Both the BladeCenter and most Cisco products support using an odd number of links in the aggregation, as well as non-sequential links.

The examples provided assume a layer 2 network exists, and one is attempting to connect the BladeCenter to this layer 2 network. Where appropriate, comments on ports being blocked via STP, will be included. If the network behind the initial switches is a layer 3 network (routed) the comments on STP blocking will probably be incorrect, depending on your specific environment. It is important for the administrator to understand these situations.

The examples provided do not go into network architecture design, rather they only cover the specifics of interfacing the BladeCenter in to a Cisco infrastructure with certain characteristics. It is assumed that the administrator understands the need for and ramifications of a proper network design (see 5.2, “Architecture summary” on page 74 for a brief introduction to datacenter architectures).

Examples that show use of link aggregation are only in reference to layer 2 link aggregation. Some Cisco switches support layer 3 aggregation, but their description and use are not discussed in these examples.

Options exist to control the way traffic is load balanced over any aggregated links. This section assumes that default load balancing is in use.

General rules and comments for configuring the ESM

While there are many possible ways to create a desired configuration, care must always be taken to understand the consequences of any such configuration.

The ESM supports configuration via an HTTP Web interface, a telnet interface and the IBM Director software. This chapter only discusses the HTTP Web based tool, accessible by pointing your browser to the IP address of the ESM to be configured, and logging in. The default user credentials (ID and password) are as follows:

User ID: *USERID* (All Caps)

Password: *PASSWORD* (All caps, numeric 0 used in place of alpha O)

Recommended sequence of configuration

The following are the basic steps followed in the production of most of the examples given in this document:

1. Shut down or un-cable the links to be configured (Table 5-1 on page 84).
2. Configure the IBM ESM.
 - a. Configure desired speed and duplex of links.
 - b. Configure any desired VLANs and VLAN trunking options.
 - c. Configure any desired aggregation links.
 - d. Save the configuration to NVRAM.
3. Configure the Cisco equipment.
 - e. Configure desired speed and duplex of links.
 - f. Configure any desired VLANs and VLAN trunking options.
 - g. Configure any desired aggregation links.
 - h. Depending on the Cisco switch, save the config to NVRAM.
4. Re-enable or re-cable the links that had been disabled in step 1 (Table 5-2 on page 84).
5. Confirm the desired operation of the configuration.

Base configuration options common to all examples

The following are some configuration options established that are common to all of the examples. These are only for demonstration purposes in the examples, and more than likely *will not* be duplicated in your particular environment.

All configurations will have three VLANs configured: VLAN 1, VLAN 5 and VLAN 10.

All configurations assume that VLAN 1 is the native VLAN (Native VLAN is untagged).

All configurations assume that all VLANs will be carried on all trunks. It is possible to limit which VLANs are carried on a given trunk, but this is not presented in this chapter.

All configurations presented in this section force one of the Cisco switches to be the Spanning Tree root for all VLANs. There is a high probability that any existing network will already have a desired switch configured as the root. It is very important that you understand the proper selection of the root bridge, and that the ESM not be allowed to become the root bridge. Allowing the ESM to become the root bridge can result in sub-optimal data flow within the layer 2 network.

The following blade servers internal to the BladeCenter will be placed in the following VLANs during the configuration stage of each example:

- BladeServer 1: VLAN 1
- BladeServer 2: VLAN 5
- BladeServer 3: VLAN 10
- BladeServer 4: VLAN 10

Summary of disconnect procedure, to be performed for each example

When performing initial configurations or making changes to existing configurations that might have an impact on Spanning Tree (such as changing link aggregation), it is recommended that you leave connections un-cabled, or shut down, prior to making the configuration changes. This will reduce the likelihood of any temporary Spanning Tree loops and possible

network-down conditions that might result in the process of adding or changing configurations.

Table 5-1 shows three basic options to disable the connection. Choose the one most suited to your situation. For example, if you will not be physically at the equipment while you are performing the configuration, physically disconnecting the cables is not your best option.

Table 5-1 Pre configuration step - disable the links being configured

Description and comments	Action
Option 1 - Disable the ESM interface. Repeat for any other Ext interfaces involved in the configuration.	Perform the following from the Web interface: <ul style="list-style-type: none"> ► Click Configuration -> Port Settings -> Configure Ports. ► Select Ext1 interface by clicking the Ext1 connector at the top of the screen. ► Change State to Disable. ► Click Apply.
Option 2a - Disable the Cisco interface (CatOS). Repeat for any other desired interfaces.	Perform the following from the enable mode: To disable a single port: set port disable 2/3 To disable a range of ports from 2/3 to 2/6: set port disable 2/3-6
Option 2b - Disable the Cisco interface (IOS). Repeat for any other desired interfaces.	Perform the following from the enable mode: To disable a single port: config t interface fa0/1 shutdown end To disable a range of ports from fa0/1 to fa0/4: config t interface range fa0/1 - 4 shutdown end
Option 3 - Pull the cable(s).	Unplug the cable(s) from either the ESM or the Cisco external switch of the link(s) as necessary.

Summary of reconnect procedure, to be performed for each example

Table 5-2 includes the steps performed once the configuration of both sides of the connection are complete. It should be the reverse of whatever procedure was used in Table 5-1.

Table 5-2 Post configuration step - reconnecting the devices

Description and comments	Action
Option 1 - Re-enable the ESM interface. Repeat for any other Ext interfaces involved in the configuration.	Perform the following from the Web interface: <ul style="list-style-type: none"> ► Click Configuration -> Port Settings -> Configure Ports. ► Select Ext1 interface by clicking the Ext1 connector at the top of the screen. ► Change State to Enable. ► Click Apply.
Option 2a - Enable the Cisco interface (CatOS). Repeat for any other desired interfaces.	<ul style="list-style-type: none"> ► Perform the following from the enable mode: <ul style="list-style-type: none"> – To enable a single port: set port enable 2/3 – To enable a range of ports from 2/3 to 2/6: set port enable 2/3-6

Description and comments	Action
Option 2b - Enable the Cisco interface (IOS). Repeat for any other desired interfaces.	<ul style="list-style-type: none"> ► Perform the following from the enable mode: <ul style="list-style-type: none"> – To enable a single port: <pre>config t interface fa0/1 no shutdown end</pre> – To enable a range of ports from fa0/1 to fa0/4: <pre>config t interface range fa0/1 - 4 no shutdown end</pre>
Option 3 - Plug the cable(s) in to the port(s).	<ul style="list-style-type: none"> ► Plug in the cable(s) from either the ESM or the Cisco external switch of the link(s) as necessary.

ESM base configuration

As already noted, during the creation of this chapter, each device was wiped out and an initial base configuration was installed prior to the procedures given in each example.

For the ESMs there should be no changes necessary after the Factory Reset. One comment that needs to be mentioned here is that if this were a brand new @server BladeCenter, you would need to connect to the BladeCenter's Management Module (default 192.168.70.125) and Enable the external interfaces of the ESMs (default is Disabled) at least once. This can be done by logging in to the Management Module (ID= *USERID* and Password = *PASSWORD*) and clicking **Switch Tasks -> Management -> Switch Module X** (where **X** is the ESM to manage), click **Advanced Switch Management -> Advanced Setup** and change External ports to Enable (reference Figure 5-2).

Advanced Setup

Fast POST	Disabled
External ports	Enabled
External management over all ports	Enabled
Preserve new IP configuration on all switch resets	Enabled

Figure 5-2 Enabling the external ports for the first time via the Management Module

The following set of screen captures show some of the default configurations (set by Factory Default), as seen by being logged in to the ESM in question.

For this first view, reference Figure 5-3 on page 86, Click the top connector (**Ext1**) in the picture of the ESM at the top of the page.

Configure Ports — Ext 1 ?

Connection

State

Speed/Duplex

Flow Control

Link Down

Enabled

Auto

Enabled

Figure 5-3 ESM showing Ext1 set to enable (should be set the same for other Ext ports)

As in Figure 5-4, click **Configuration -> VLANs -> 802.1Q Port Settings -> Port VLAN ID (PVID)**.

Port VLAN ID (PVID) ?

Port	PVID	Port	PVID
Ext 1	1	Bay 1	1
Ext 2	1	Bay 2	1
Ext 3	1	Bay 3	1
Ext 4	1	Bay 4	1
		Bay 5	1
		Bay 6	1
		Bay 7	1
		Bay 8	1
		Bay 9	1
		Bay 10	1
		Bay 11	1
		Bay 12	1
		Bay 13	1
		Bay 14	1

Figure 5-4 ESM showing default PVID settings for each connection in the ESM

As in Figure 5-5 on page 87, click **Configuration -> VLANs -> Edit 802.1Q VLANs**, select **VLAN 1**, click **Edit**.

802.1Q Static VLAN Entry Settings — Edit ?

VLAN ID (VID)

VLAN Name

Port	Ext1	Ext2	Ext3	Ext4	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Tag	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
None	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Egress	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
Forbidden	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure 5-5 ESM showing default settings for VLAN 1

As in Figure 5-6, click **Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Link Aggregation Port Settings**.

Link Aggregation Port Settings ?

Port	Priority	Admin Key	Oper Key	Mode	Status
Ext 1	128	15	15	Disabled	Individual
Ext 2	128	16	16	Disabled	Individual
Ext 3	128	17	17	Disabled	Individual
Ext 4	128	18	18	Disabled	Individual

Configure External Port from to

Priority

Admin Key

Mode

Figure 5-6 ESM showing default settings for Link Aggregation LACP

As in Figure 5-7 on page 88, click **Configuration -> Link Aggregation -> Port Trunking**.

Port Trunking

Group ID	Ports				Method
	Ext 1	Ext 2	Ext 3	Ext 4	
1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled
2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled

Figure 5-7 ESM showing default settings for static link aggregation

Cat 3550 (IOS based switch) base configuration

As with the ESM, each switch in each example is wiped out and a base configuration installed, prior to the commencement of each example. In the case of IOS based switches, after the switch is reset to the factory default, the following configuration was applied (to simulate portions of a pre-existing Cisco network). As already mentioned, this is only for the examples shown here, and more than likely will vary from your production network.

Several things to note about this base configuration:

- ▶ Port fa0/20 has been configured as a test point to test access to VLAN 1.
- ▶ Port fa/15 has been configured as a test point to test access to VLAN 5.
- ▶ Port fa0/10 has been configured as a test point to test access to VLAN 10.
- ▶ Port fa0/1 through fa0/4 will be configured as desired in each section (this will include setting up the trunk link to the ESM, as well as any desired aggregation).
- ▶ In examples with more than one Cisco switch, port 0/24 on switch Cisco-1 will be connected to port 0/24 on Cisco-2 to simulate the layer 2 cloud behind the two Cisco switches (this will help to show any Spanning Tree blocking that might occur with a given configuration).
- ▶ In the examples with more than one Cisco switch, the switch named Cisco-2 will be forced to become the root switch (for consistence purposes only). This can be done by running the **spanning-tree vlan X priority YYYYY** command such that YYYYY is lower on Cisco-2 than it is on Cisco-1 for each of the three sample VLANs.
- ▶ Telnet and Enable passwords are all set to cisco (lowercase).
- ▶ The configuration example below is for an IOS based switch acting as Cisco-1 (VLAN 1 management IP address set to 192.168.70.200).
- ▶ For a Cisco-2 configuration of an IOS based switch, the VLAN 1 management IP address would be set to 192.168.70.201, and the name of the switch would be changed accordingly (Cisco-2-IOS).

```
!
hostname Cisco-1-IOS
!
enable secret 5 $1$QHTt$SvAEWBZAtQUvWj.uzH07Y1
enable password Cisco
!
ip subnet-zero
no ip domain-lookup
!
vtp domain IBM
vtp mode transparent
```

```

vlan 5
  name VLAN5
!
vlan 10
  name VLAN10
!
spanning-tree mode pvst
spanning-tree extend system-id
spanning-tree vlan 1 priority 24576
spanning-tree vlan 5 priority 24576
spanning-tree vlan 10 priority 24576
!
!
!
!
!
interface FastEthernet0/10
  description VLAN 10 Test point
  switchport access vlan 10
  switchport mode access
  no ip address
  spanning-tree portfast
!
interface FastEthernet0/15
  description VLAN 5 Test point
  switchport access vlan 5
  switchport mode access
  no ip address
  spanning-tree portfast
!
interface FastEthernet0/20
  description VLAN 1 Test point
  switchport mode access
  no ip address
  spanning-tree portfast
!
interface FastEthernet0/24
  description Trunk port to rest of layer 2 network
  switchport trunk encapsulation dot1q
  switchport mode trunk
  switchport nonegotiate
!
!
interface Vlan1
  ip address 192.168.70.200 255.255.255.0
  no shut
!
ip classless
!
!
!
line con 0
  exec-timeout 0 0
line vty 0 4
  exec-timeout 0 0
  password cisco
  login
!
end

```

Cat 4003 (CatOS based switch) base configuration

As with the ESM and IOS based switches, the CatOS based switch was wiped out and a base configuration applied, prior to each example. After the switch was reset to the factory default, the following configuration was applied to the CatOS based switch to simulate portions of a pre-existing Cisco network. As already mentioned, this is only for the examples shown here, and more than likely will vary from your production network.

Several things to note about this base configuration:

- ▶ Module 2 is a 4232-GB-RJ, with two ports of GBIC based GigE and 34 ports of 10/100BaseT.
- ▶ Port 2/20 has been configured as a test point to test access to VLAN 1.
- ▶ Port 2/15 has been configured as a test point to test access to VLAN 5.
- ▶ Port 2/10 has been configured as a test point to test access to VLAN 10.
- ▶ Port 2/3 through 2/6 will be configured as desired in each section. This will include setting up the trunk link to the ESM, as well as any desired aggregation.
- ▶ In examples with more than one Cisco switch, port 2/34 on switch Cisco-1 will be connected to port 2/34 on Cisco-2 to simulate the layer 2 cloud behind the two Cisco switches (this will help to show any Spanning Tree blocking that might occur with a given configuration).
- ▶ In the examples with more than one Cisco switch, the switch named Cisco-2 will be forced to become the root switch (for consistence purposes only). This can be done by running the **set spantree priority YYYYY 1** command such that YYYYY is lower on Cisco-2 than it is on Cisco-1 for each of the three sample VLANs.
- ▶ Telnet and Enable passwords are all set to `cisco` (lowercase).
- ▶ The configuration example below is for a CatOS based switch acting as Cisco-1 (VLAN 1 management IP address set to 192.168.70.202).
- ▶ For a Cisco-2 configuration of a CatOS based switch, the VLAN 1 management IP address would be set to 192.168.70.203, and the name of the switch would be changed accordingly (Cisco-2-CatOS).

```
#version 7.5(1)
!
!
#system web interface version(s)
set password $2$4CZy$7im5fQA35UD.7hR5us/U0.
set enablepass $2$NMK/$VnHGG4PHSDFPFG28d7IpP0
set logout 0
!
#system
set system name Cisco-1-CatOS
!
!
#vtp
set vtp domain IBM
set vtp mode transparent
set vlan 5,10
!
#ip
# For Cisco-1 deployment
set interface sc0 1 192.168.70.202/255.255.255.0 192.168.70.255
!
set interface sl0 down
set interface me1 down
!
```



```

#vlan <VlanId>
set spantree priority 8192 1
set spantree priority 8192 5
set spantree priority 8192 10
!
#module 2 : 34-port 10/100/1000 Ethernet
set vlan 5 2/15
!
set vlan 10 2/10
!
!
!
set trunk 2/34 nonegotiate dot1q
set spantree portfast 2/10,2/15,2/20 enable

```

5.5 Configuration examples

This section contains the actual configuration examples for installing and verifying the @server BladeCenter in a layer 2 Cisco infrastructure environment.

Note: All examples given are based on layer 2 network connections. The possibility exists to connect the layer 2 interfaces of the ESM's directly to routed interfaces (layer 3) of Cisco switches. While this option is available, and certainly desirable in certain environments, its discussion is beyond the scope of this document and is thus not included here.

Note: While some of the examples provided show a number of ways to achieve layer 2 redundancy, they do not take into account some of the possible issues that may be encountered with Spanning Tree in complex layer 2 networks. With this in mind, the examples should only be viewed as a possibility of what can be done, and not what is necessarily desirable in any specific environment. While Spanning Tree can be a valuable tool to control loops, it should not be used as a substitute for good designs that minimize possible loops but still offer the desired redundancy.

5.5.1 Single ESM, single link to a single Cisco switch

In this example (Figure 5-8 on page 92), we discuss a very basic configuration that includes a single ESM, a single Cisco switch and a single link between the two. This configuration offers minimal performance and redundancy and might be used for initial installation and testing of a @server BladeCenter, or in an environment that does not require maximum performance or redundancy.




Figure 5-8 Single ESM, Single Cisco Switch, Single link

Step 1: Take down the link(s)

It is always advisable to disable the link(s) prior to making any configuration changes. Please see Table 5-1 on page 84 for procedures.

Step 2: Configure the IBM side of the link

This section walks you through the sequence of actions required to configure the ESM for this example (reference Table 5-3).

The following assumptions have been made for this example:

- ▶ The user is already logged in to the ESM with *root* level access.
- ▶ Port EXT1 on the ESM in Switch Module Bay 1 is being used as the link between IBM and Cisco.
- ▶ Commands are being performed in the sequence shown.
- ▶ The ESM is starting from a default config per the Example on page 85.
- ▶ Cisco switches being used are 10/100 based and we will be setting the ESM port to 100Mbps full duplex. This means a cross-over cable *must* be used for the link between the ESM and the Cisco switch.
 - If using switch ports that support 1000Mbps, the procedure to set speed will be slightly different from what is shown in the example.

Table 5-3 Configuring the ESM

Description and Comments	Actions to perform via Web interface to ESM
Step 2.1- <i>Configure speed and duplex.</i> As already noted, it will be necessary to use a cross-over cable on the link between the ESM and the Cisco switch, as tuning off auto-negotiation also turns off the auto selection of MDI/MDI-X.	<ul style="list-style-type: none"> ▶ Click the top connector (Ext1) in the picture of the ESM at the top of the page. ▶ Change Speed/Duplex to 100/Full. ▶ Click Apply.
Step 2.2 - <i>Configure PVIDs</i> This places the desired blade server ports in to the desired VLANs. If the VLAN does not exist, it will be created automatically.	<ul style="list-style-type: none"> ▶ Click Configuration -> VLANs -> 802.1Q Port Settings -> Port VLAN ID (PVID). ▶ For Bay 2, change the PVID to 5. ▶ For Bays 3 and 4, change the PVID to 10. ▶ All other PVIDs should be set for 1. ▶ Click Apply.

Description and Comments	Actions to perform via Web interface to ESM
<p>Step 2.3 - <i>Configure 802.1Q trunking</i> This will allow Ext1 to carry traffic for VLAN 5 and VLAN 10 (already carrying VLAN 1 as an initial default).</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> VLANs -> Edit 802.1Q VLANs ▶ Select VLAN 5 and click Edit. ▶ Set Interface Ext1 to Egress and make sure Tag box is checked. ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked. ▶ Click Apply. ▶ Click Configuration -> VLANs -> Edit 802.1Q VLANs. ▶ Select VLAN 10 and click Edit. ▶ Set interface Ext1 to Egress and make sure the Tag box is checked. ▶ Make sure Ports 3 and 4 are set for Egress and the Tag box is unchecked for these two ports. ▶ Click Apply.
<p>Step 2.4 - <i>Save ESM config to NVRAM</i> Failure to perform this step will result in all changes to the ESM being lost if the @server BladeCenter is powered off or the ESM is otherwise restarted.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> Maintenance -> Save Changes. ▶ Click Save Configuration. ▶ Click Ok when complete.

Step 3: Configuring the Cisco switch

The following assumptions have been made for this example (Table 5-4):

- ▶ VLANs 1, 5 and 10 already exist in the Cisco configuration and will be part of the 802.1Q trunk to the ESM.
- ▶ The user is already logged in to the switch and the switch is in enable mode.
- ▶ The lowest available compatible port is being used:
- ▶ For the CatOS switch being used in this example, port 2/3 is being used.
- ▶ For the IOS switch being used in this example, port fa0/1 is being used.
- ▶ Commands are being performed in the sequence shown.
- ▶ The switch is starting from a default config per the Example on page 88 (IOS) or the Example on page 90 (CatOS).
- ▶ Cisco switch being used is 10/100 based and we will be setting the port to 100Mb full duplex.
 - If using switch ports that support 1000Mbps, the command to set speed will be different from what is shown in the example.

Table 5-4 Configuring the Cisco equipment

Description and Comments	CatOS based switch	IOS based switch
<p>Step 3.1 <i>Configure speed and duplex.</i></p>	<pre>set port speed 2/3 100 set port duplex 2/3 full</pre>	<pre>config t int fa0/1 speed 100 duplex full</pre>
<p>Step 3.2 <i>Configure 802.1Q trunking</i> Forces link to become an 802.1Q VLAN trunk.</p>	<pre>set trunk 2/3 nonegotiate dot1q</pre>	<pre>switchport trunk encap dot1q switchport mode trunk switchport nonegotiate end</pre>
<p>Step 3.3 <i>Save config to NVRAM.</i> Only necessary on IOS based switches.</p>	<p>(does not apply)</p>	<pre>write mem</pre>

Step 4: Reconnecting the devices

This is the final step to bring the connection in to full operation. This will be the reverse of whatever procedure was used in Step 1. See Table 5-2 on page 84 for details on how to reestablish the links.

Step 5: Verifying the configuration

This section provides options for verifying the correct and desired operation.

Verifying correct operation on the BladeCenter ESM

Verify the configurations on the ESM look as follows:

As in Figure 5-9, verify port state by clicking **Configuration -> Port Setting -> Display Ports**, and reviewing the status of **Ext1** (State should be *Enabled*, Speed/Duplex should be *100/Full*, and Connection should show *100M/Full/802.3x*).

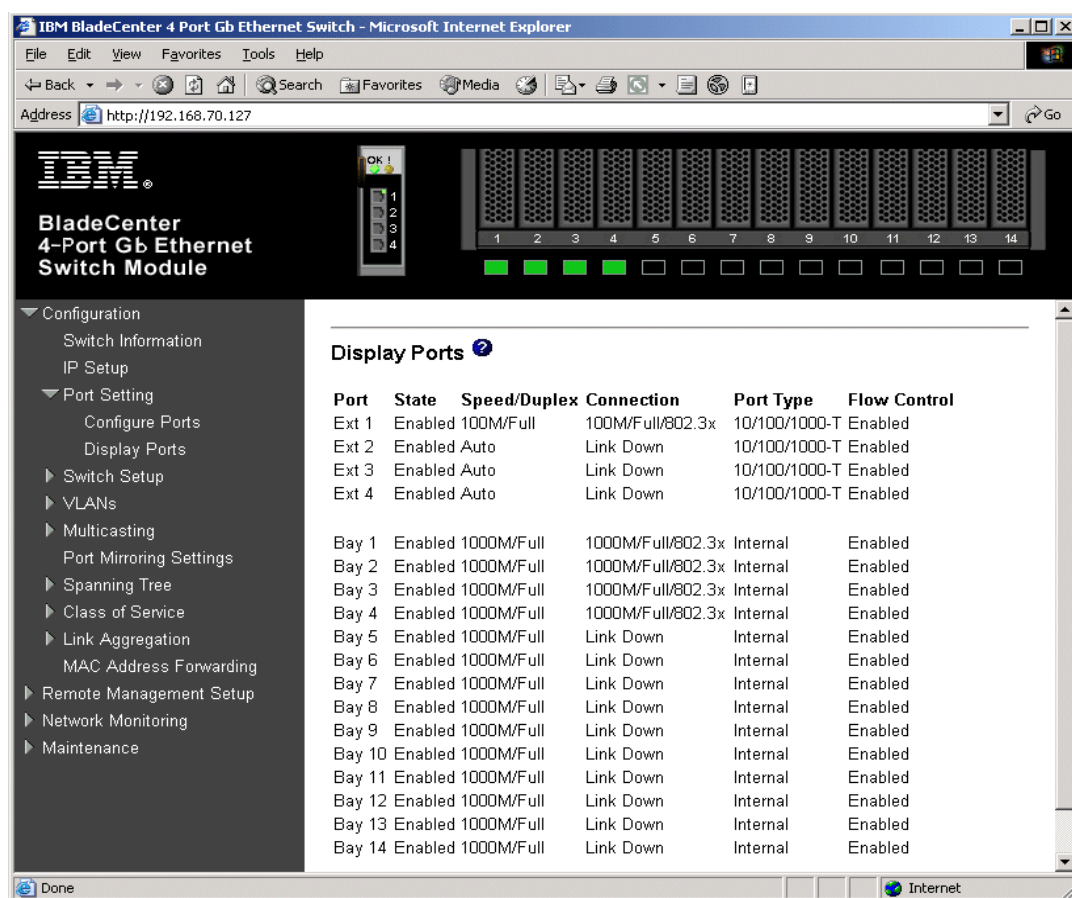


Figure 5-9 Verify ports are operational

As in Figure 5-10 on page 95, verify VLAN 1 configurations clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 1** and clicking **Edit**. Make sure the ports not on VLAN 1 (Ports 2, 3 and 4) are set to None and their associated **Tag** box is checked

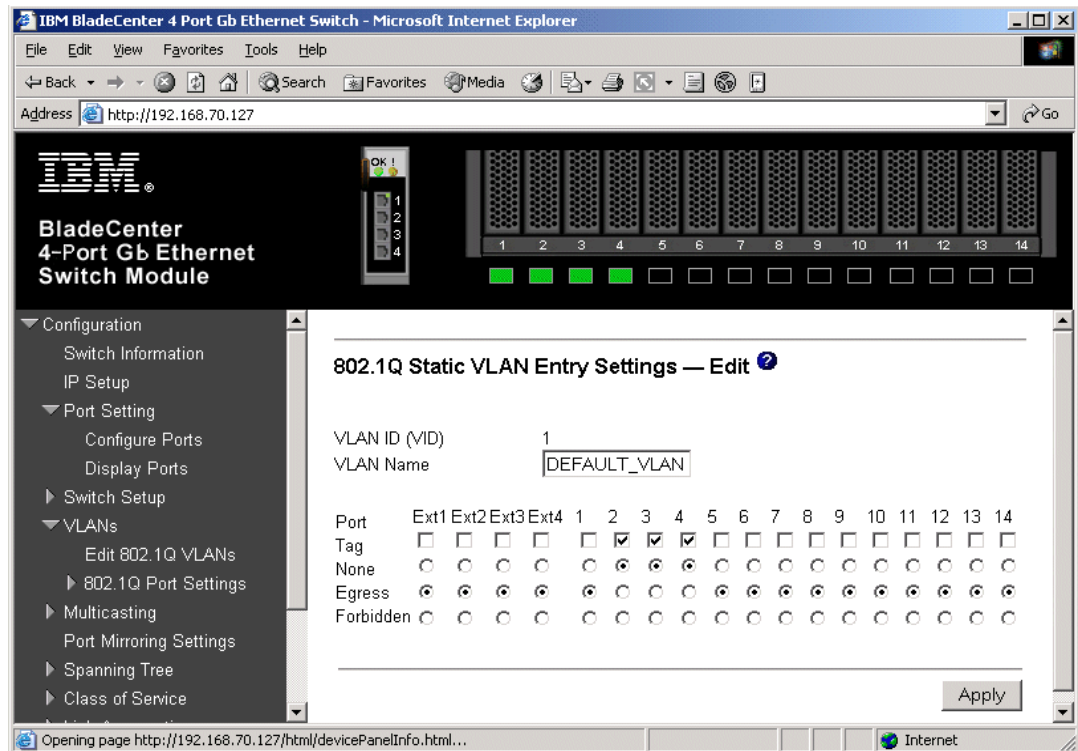


Figure 5-10 .Verifying VLAN 1

As in Figure 5-11 on page 96, verify VLAN 5 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 5** and clicking **Edit**. Make sure **Port 2** is set to **Egress** and that the **Tag** box is not checked. Make sure the Ext1 interface is set for **Egress** (allows VLAN 5 traffic to pass through Ext1) and that the box for **Tag** is checked.

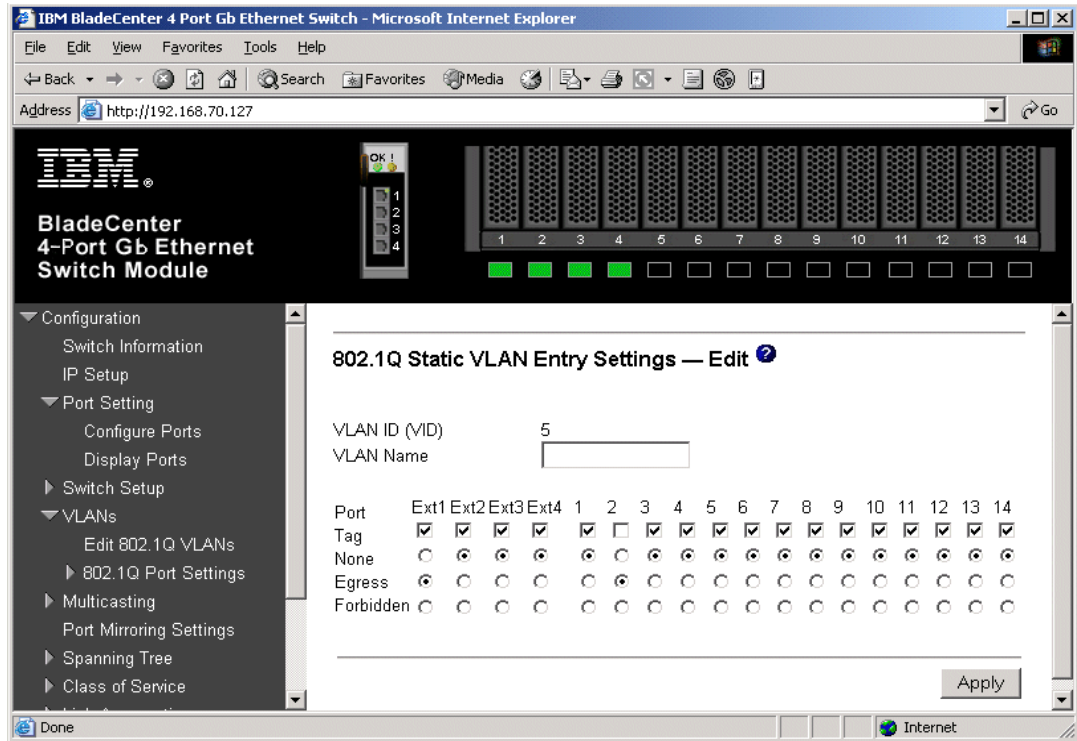


Figure 5-11 Verify VLAN 5

As in Figure 5-12 on page 97, verify VLAN 10 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 10** and clicking **Edit**. Make sure ports 3 and 4 are set to Egress and that the Tag box is not checked for either port. Make sure the Ext1 port is set for Egress (allows VLAN 10 traffic to pass through Ext1) and that the box for **Tag** is checked.

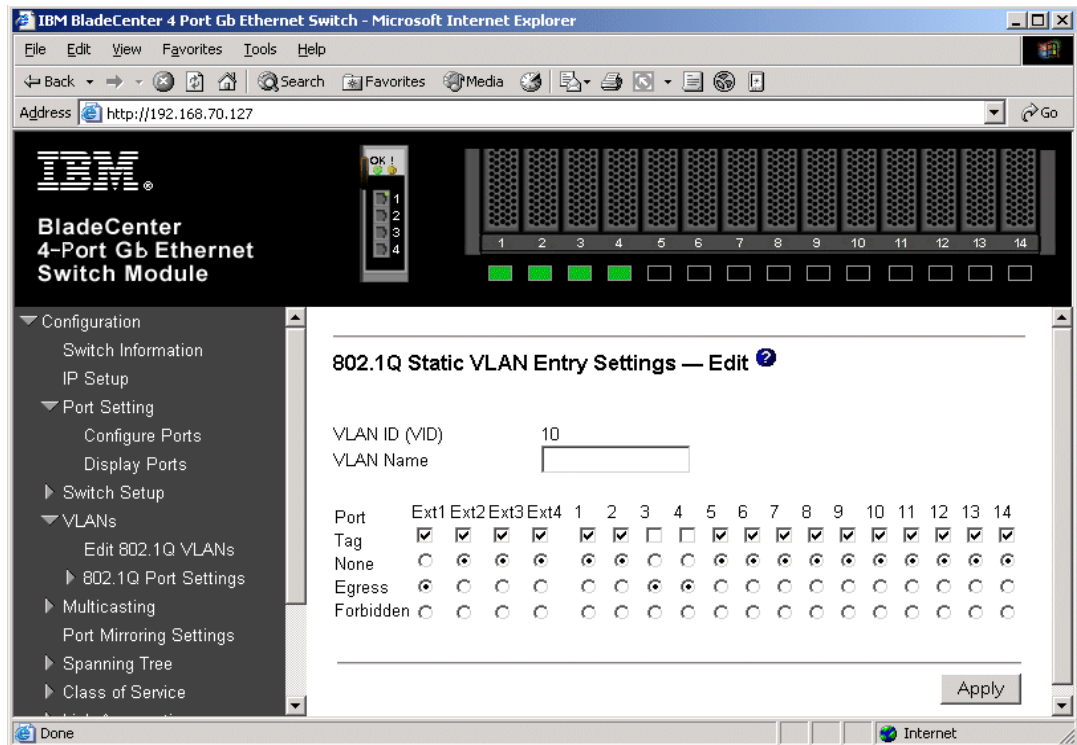


Figure 5-12 .Verify VLAN 10

As in Figure 5-13, verify Spanning Tree is forwarding traffic for the port by clicking **Configuration -> Spanning Tree -> STP Port Settings**. Check to see that the Status for Ext1 is forwarding.

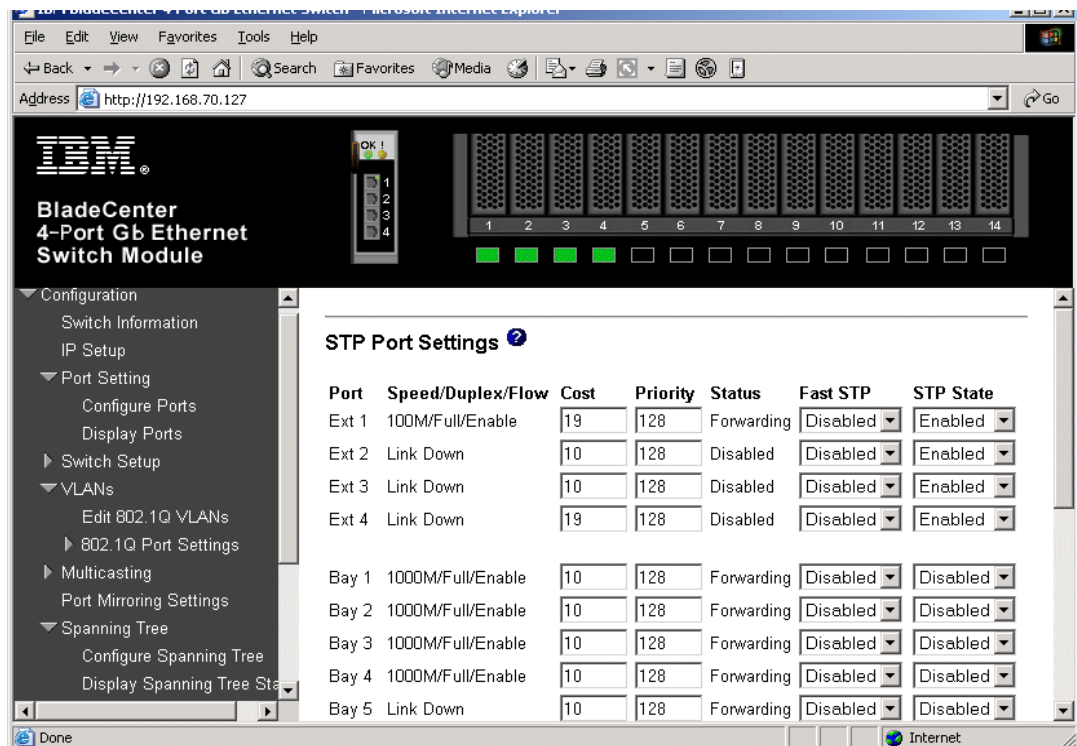


Figure 5-13 Verifying STP is forwarding to Ext1

Verifying correct operation on the Cisco external switch

The following section (reference Table 5-5) includes some commands one can use to verify the desired configuration and operation of the Cisco equipment.

Table 5-5 Verifying the configuration and operation of the Cisco external switch of the connection

Description and Comments	CatOS based switch	IOS based switch
Review running config for desired statements.	show config ▶ Review for the following: <ul style="list-style-type: none"> – set port speed 2/3 100 – set port duplex 2/3 full – set trunk 2/3 nonegotiate dot1q 1-1005,1025-4094 	show run ▶ Review for the following on int fa0/1: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport nonegotiate – duplex full – speed 100
Show speed and duplex.	show port status 2/3 ▶ Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100 	show int fa0/1 status ▶ Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100
Show trunking status. Link must be up before a trunk will come up.	show port trunk 2/3 ▶ Should show the following: <ul style="list-style-type: none"> – Mode = nonegotiate – Encapsulation = dot1q – Status = Trunking – Native VLAN = 1 	show int fa0/1 trunk ▶ Should show the following: <ul style="list-style-type: none"> – Mode = on – Encapsulation = 802.1q – Status = Trunking – Native VLAN = 1
Ping the ESM. Where x.x.x.x is the IP address of the ESM (must be in same VLAN as subnet being pinged).	ping x.x.x.x ▶ Based on the Cat4K being at 192.168.70.202 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.	ping x.x.x.x ▶ Based on the Cat 3550 being at 192.168.70.200 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.
Attach a device to VLAN 5 and 10 on the Cisco switch and verify a ping to one of the blade servers in that same VLAN works as desired.	▶ For VLAN 5, attach a device to port 2/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work) ▶ For VLAN 10, attach a device to port 2/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).	▶ For VLAN 5, attach a device to port fa0/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work). ▶ For VLAN 10, attach a device to port fa0/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).

5.5.2 Single ESM, single link to two Cisco switches

This example (see Figure 5-14 on page 99) shows a little more robust configuration, with a single ESM connecting to two different Cisco switches. It still offers minimal performance but with increased redundancy should an uplink or one of the external switches fail. It does not offer any redundancy in the event of an ESM failure.

Based on our specific configuration for this example (layer 2 network, Cisco-2 is root and port cost on both Cisco links toward the ESM set to 100), Spanning Tree will block the connection between Cisco-1 and the ESM, at the Cisco external switch until such time as the link from

the ESM to Cisco-2 (or Cisco-2 itself) goes down. The choice of root and port cost settings in this example were only made for this example, and may be a poor choice in a production network (certainly placing the root switch up against the ESM in a datacenter environment would not be very common). It is very important that any time an ESM is connected in a redundant fashion, the location of the root switch and various port costs in the network be accounted for to ensure correct data flow (we need to prevent non-BladeCenter traffic from flowing through the ESM).




Figure 5-14 Single ESM with single links to two different Cisco switches

Step 1: Take down the link(s)

It is always advisable to disable the link(s) prior to making any configuration changes. Please see Table 5-1 on page 84 for procedures.

Step 2: Configure the IBM side of the link

This section walks you through the sequence of actions required to configure the ESM for this example (reference Table 5-6 on page 100).

The following assumptions have been made for this example:

- ▶ The user is already logged in to the ESM with *root* level access.
- ▶ Port EXT1 on the ESM in Switch Module Bay 1 is being used as the link between the @server BladeCenter and Cisco-1.
- ▶ Port EXT2 on the ESM in Switch Module Bay 1 is being used as the link between the @server BladeCenter and Cisco-2.
- ▶ Commands are being performed in the sequence shown.
- ▶ The ESM is starting from a default config per the Example on page 85.
- ▶ Cisco switches being used are 10/100 based and we will be setting the ESM port to 100Mbps full duplex. This means a cross-over cable *must* be used for the link between the ESM and the Cisco switch.
 - If using switch ports that support 1000Mbps, the procedure to set speed will be slightly different from what is shown in the example.

Table 5-6 Configuring the ESM

Description and Comments	Actions to perform via Web interface to ESM
<p>Step 2.1- <i>Configure speed and duplex.</i> As already noted, it will be necessary to use a cross-over cable on the link between the ESM and the Cisco switch, as tuning off auto-negotiation also turns off the auto selection of MDI/MDI-X.</p>	<ul style="list-style-type: none"> ▶ Click the top connector (Ext1) in the picture of the ESM at the top of the page. ▶ Change Speed/Duplex to 100/Full. ▶ Click Apply. ▶ Repeat for Ext2 interface.
<p>Step 2.2 - <i>Configure PVIDs</i> This places the desired blade server ports in to the desired VLANs. If the VLAN does not exist, it will be created automatically.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID). ▶ For Bay 2, change the PVID to 5. ▶ For Bays 3 and 4, change the PVID to 10. ▶ All other PVIDs should be set for 1. ▶ Click Apply.
<p>Step 2.3 - <i>Configure 802.1Q trunking</i> This will allow Ext1 and Ext2 to carry traffic for VLAN 5 and VLAN 10 (already carrying VLAN 1 as an initial default).</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->Edit 802.1Q VLANs. ▶ Select VLAN 5 and click Edit. ▶ Set Interface Ext1 to Egress and make sure Tag box is checked. ▶ Set Interface Ext2 to Egress and make sure Tag box is checked. ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked. ▶ Click Apply. ▶ Click Configuration->VLANs->Edit 802.1Q VLANs. ▶ Select VLAN 10 and click Edit. ▶ Set interface Ext1 to Egress and make sure Tag box is checked. ▶ Set Interface Ext2 to Egress and make sure Tag box is checked. ▶ Make sure Ports 3 and 4 are set for Egress and the Tag box is unchecked for these two ports. ▶ Click Apply.
<p>Step 2.4 - <i>Save ESM config to NVRAM</i> Failure to perform this step will result in all changes to the ESM being lost if the @server BladeCenter is powered off or the ESM is otherwise restarted.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> Maintenance -> Save Changes. ▶ Click Save Configuration. ▶ Click Ok when complete.

Step 3: Configuring the Cisco switches

The following assumptions have been made for this example (reference Table 5-7 on page 101):

- ▶ VLANs 1, 5 and 10 already exist in the Cisco configuration and will be part of the 802.1Q trunk to the ESM.
- ▶ The user is already logged in to the switch and the switch is in enable mode.
- ▶ The lowest available compatible port is being used:
- ▶ For the CatOS switch being used in this example, port 2/3 is being used.
- ▶ For the IOS switch being used in this example, port fa0/1 is being used.
- ▶ Commands are being performed in the sequence shown.
- ▶ The switch is starting from a default config per the Example on page 88 (IOS) or the Example on page 90 (CatOS).
- ▶ Cisco switch being used is 10/100 based and we will be setting the port to 100Mb full duplex.
 - If using switch ports that support 1000Mbps, the command to set speed will be different from what is shown in the example.

Table 5-7 Configuring the Cisco equipment

Description and Comments	CatOS based switch	IOS based switch
Step 3.1 - <i>Configure speed and duplex.</i> Need to perform on both Cisco-1 and Cisco-2.	set port speed 2/3 100 set port duplex 2/3 full	config t int fa0/1 speed 100 duplex full
Step 3.2 - <i>Configure 802.1Q trunking.</i> Need to perform on both Cisco-1 and Cisco-2. Forces link to become an 802.1Q VLAN trunk.	set trunk 2/3 nonegotiate dot1q	switchport trunk encap dot1q switchport mode trunk switchport nonegotiate
Step 3.3 <i>Configure Spanning Tree port cost.</i> Need to perform on both Cisco-1 and Cisco-2. Setting the port cost higher than default helps to prevent VLAN traffic between Cisco-1 and Cisco-2, other than VLAN 1, from being switched through the ESM. This results in a more predictable flow. For more information, review the section titled: "Guidelines and comments - Spanning Tree" on page 79	set spantree portcost 2/3 100 ► Note that for this particular design, the port costs on all links is 19. In this situation, normal traffic flow from Cisco-1 to Cisco-2 would all be through 2/34. The port cost being set here is to get in to the habit of forcing port cost when using redundant connections from an ESM in to the production network.	spanning-tree cost 100 end ► Note that for this particular design, the port costs on all links is 19. In this situation, normal traffic flow from Cisco-1 to Cisco-2 would all be through fa0/24. The port cost being set here is to get in to the habit of forcing port cost when using redundant connections from an ESM in to the production network.
Step 3.4 <i>Save config to NVRAM.</i> Only necessary on IOS based switches.	<i>(does not apply)</i>	write mem

Step 4: Reconnecting the devices

This is the final step to bring the connection in to full operation. This will be the reverse of whatever procedure was used in Step 1. See Table 5-2 on page 84 for details on how to reestablish the links.

Step 5: Verifying the configuration

This section provides options for verifying the correct and desired operation.

Certain elements of this verification depend on the existence of a layer 2 network behind the two Cisco switches. This was simulated for this example by tying the two switches together with an 802.1Q trunk link (see the configuration Example on page 88 for IOS based switches or the configuration Example on page 90 for CatOS based switches, for how this link was configured).

Verifying correct operation on the BladeCenter ESM

Verify the configurations on the ESM look as follows.

As in Figure 5-15 on page 102, verify port state by clicking **Configuration -> Port Setting -> Display Ports**, and reviewing the status of Ext1 and Ext2 (State should be *Enabled*, Speed/Duplex should be *100/Full*, and Connection should show *100M/Full/802.3x*).




Figure 5-15 Verify ports are operational

As in Figure 5-16 on page 103, verify VLAN 1 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 1** and clicking **Edit**. Make sure the ports not on VLAN 1 (Ports 2, 3 and 4) are set to None and their associated **Tag** box is checked.

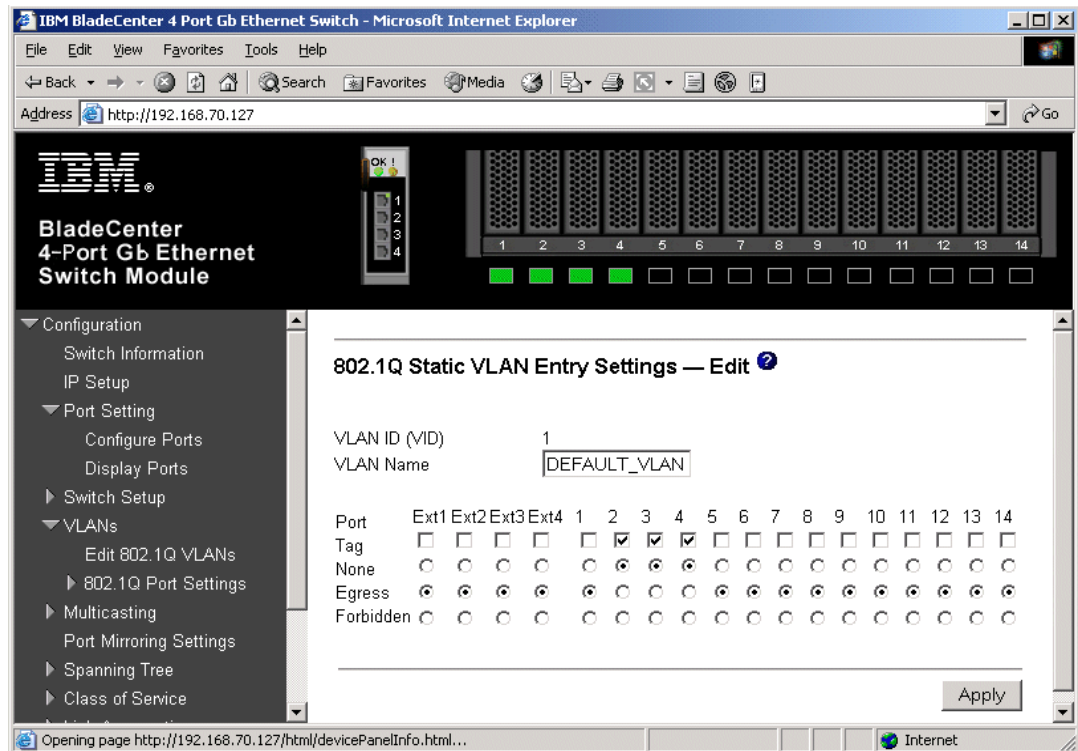


Figure 5-16 Verifying VLAN 1

As in Figure 5-17 on page 104, verify VLAN 5 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting VLAN 5 and clicking **Edit**. Make sure port 2 is set to Egress and that the Tag box is not checked. Make sure the Ext1 and Ext2 interfaces are set for Egress (allows VLAN 5 traffic to pass through these Ext interfaces) and that their associated **Tag** boxes are checked.

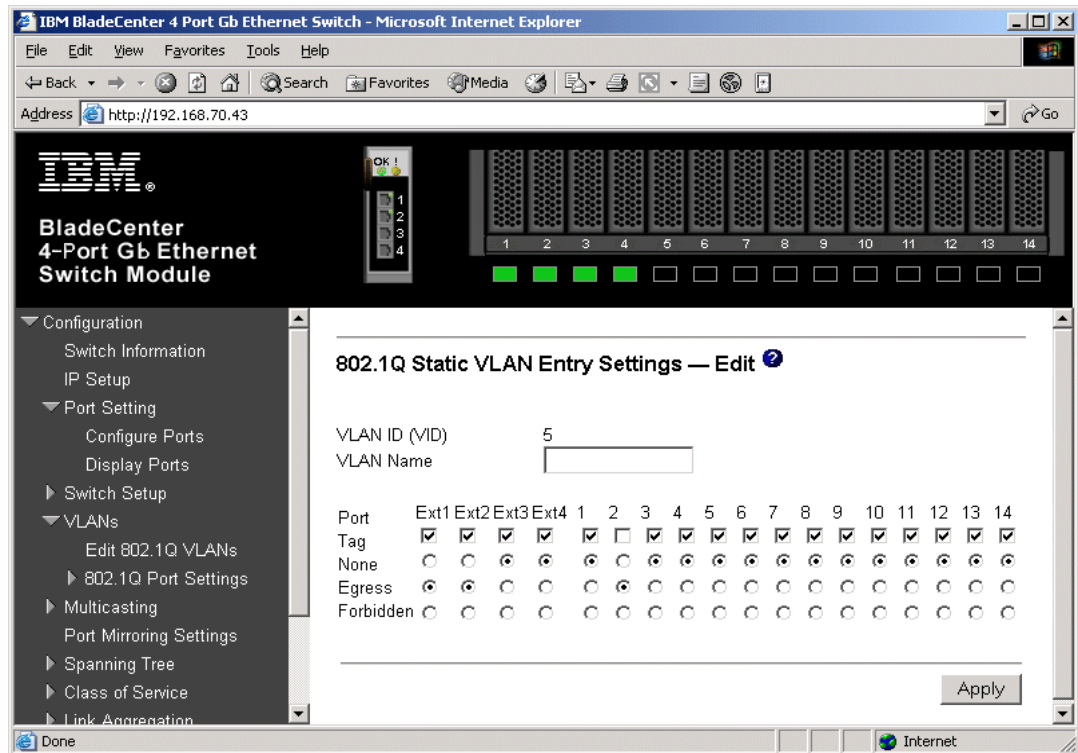


Figure 5-17 Verify VLAN 5

As in Figure 5-18 on page 105, verify VLAN 10 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting VLAN 10 and clicking **Edit**. Make sure ports 3 and 4 are set to Egress and that the Tag box is not checked for either port. Make sure the Ext1 and Ext2 interfaces are set for Egress (allows VLAN 10 traffic to pass through Ext interfaces) and that their associated **Tag** boxes are checked.

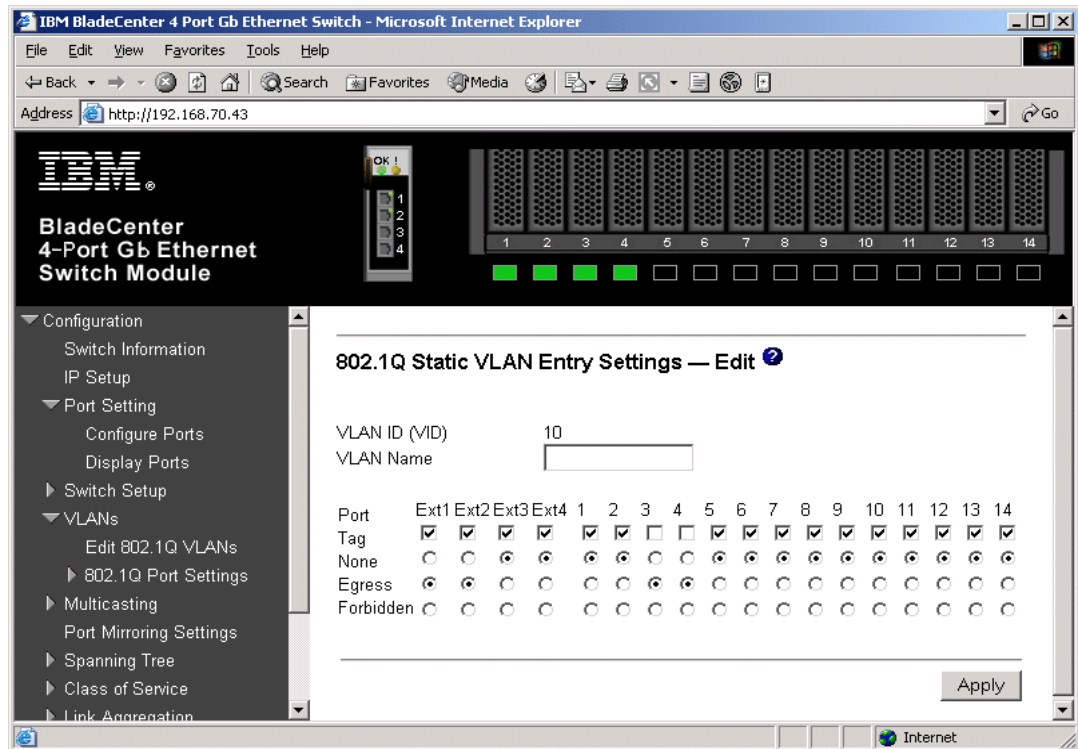


Figure 5-18 Verify VLAN 10

As in Figure 5-19 on page 106, verify Spanning Tree is operational by clicking **Configuration -> Spanning Tree -> STP Port Settings**. Based on the configurations in this example (Cisco-2 is root switch and setting of path costs) these ports should show Forwarding. The Cisco external switch of the Ext1 link should show blocking for that connection.

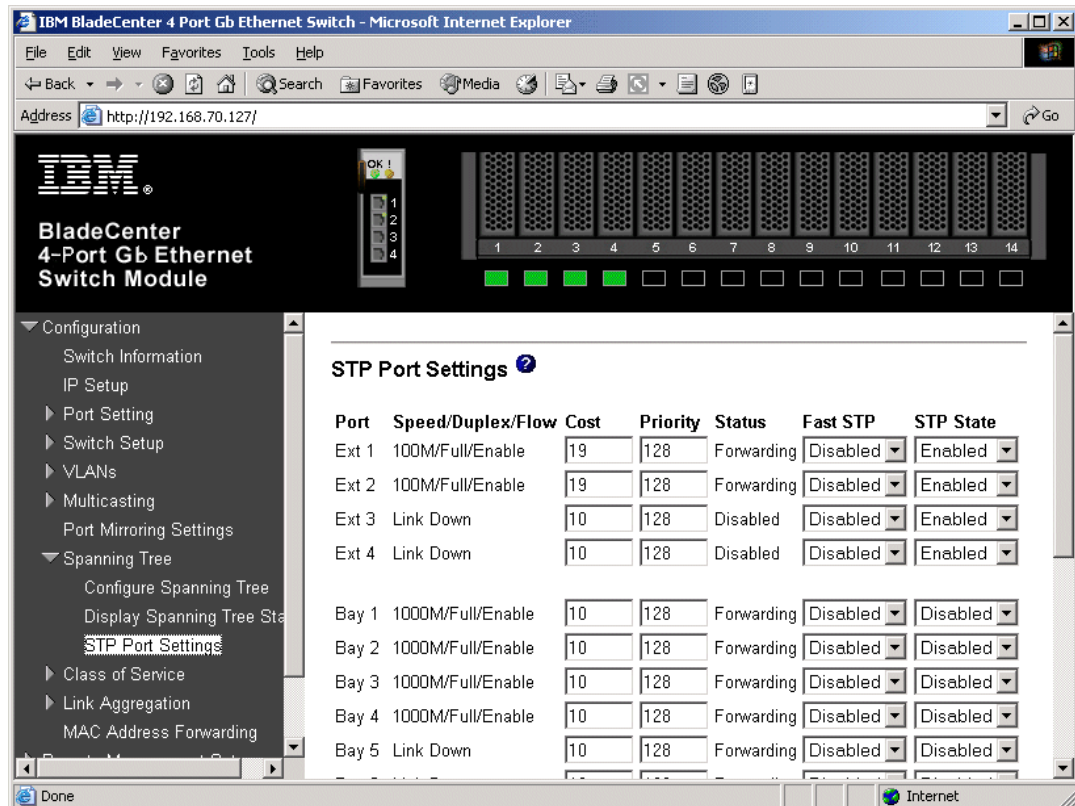


Figure 5-19 Verifying STP is forwarding to both Ext1 and Ext2

Verifying correct operation on the Cisco external switch

The following section includes some commands one can use to verify the desired configuration and operation of the Cisco equipment (reference Table 5-8).

Table 5-8 Verifying the configuration

Description and Comments	CatOS based switch	IOS based switch
Review running config for desired statements. Perform on both Cisco-1 and Cisco-2.	show config ► Review for the following: <ul style="list-style-type: none"> – set port speed 2/3 100 – set port duplex 2/3 full – set trunk 2/3 nonegotiate dot1q 1-1005,1025-4094 	show run ► Review for the following on int fa0/1: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport nonegotiate – duplex full – speed 100
Show speed and duplex. Perform on both Cisco-1 and Cisco-2.	show port status 2/3 ► Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100 	show int fa0/1 status ► Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100

Description and Comments	CatOS based switch	IOS based switch
Show trunking status. Link must be up before a trunk will come up. Perform on both Cisco-1 and Cisco-2.	show port trunk 2/3 ► Should show the following: – Mode = nonegotiate – Encapsulation = dot1q – Status = Trunking – Native VLAN = 1	show int fa0/1 trunk ► Should show the following: – Mode = on – Encapsulation = 802.1q – Status = Trunking – Native VLAN = 1
Show Spanning Tree status. In this configuration, both switches should show their ESM facing interfaces in a forwarding state.	show spantree 2/3 ► Perform on both Cisco-1 and Cisco-2. – State should show blocked for the link between Cisco-1 and the ESM for all VLANs. – Should show forwarding for all connections(2/3 and 2/34) on Cisco-2.	show spanning int fa0/1 ► Perform on both Cisco-1 and Cisco-2. – Sts should show blocked for the link between Cisco-1 and the ESM for all VLANs. – Should show forwarding for all connections (fa0/1 and fa0/24) on Cisco-2.
Ping the ESM. Where x.x.x.x is the IP address of the ESM (must be in same VLAN as subnet being pinged).	ping x.x.x.x ► Based on the Cat4K being at 192.168.70.201 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.	ping x.x.x.x ► Based on the Cat 3550 being at 192.168.70.201 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.
Attach a device to VLAN 5 and 10 on the Cisco switch and verify a ping to one of the blade servers in that same VLAN, works as desired	► For VLAN 5, attach a device to port 2/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work). ► For VLAN 10, attach a device to port 2/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).	► For VLAN 5, attach a device to port fa0/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work). ► For VLAN 10, attach a device to port fa0/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).

5.5.3 Single ESM, four port LACP aggregation to a single Cisco switch

This example (Figure 5-20 on page 108) shows a single ESM using all four external ports dynamically aggregated in to a single pipe to a single Cisco switch. It produces the maximum performance from a single ESM and also offers redundancy in the event of a link failure. It does not offer any redundancy in the event of a switch failure (either on the part of the ESM or the part of the Cisco switch). This configuration is suitable for environments that require maximum throughput with a single ESM, but are not heavily concerned with switch redundancy.




Figure 5-20 Single ESM with 4 port LACP aggregation in to a single Cisco switch

Step 1: Take down the link(s)

It is always advisable to disable the link(s) prior to making any configuration changes. Please see Table 5-1 on page 84 for procedures.

Step 2: Configure the IBM side of the link

This section walks you through the sequence of actions required to configure the ESM for this example (reference Table 5-9).

The following assumptions have been made for this example:

- ▶ The user is already logged in to the ESM with *root* level access.
- ▶ Commands are being performed in the sequence shown.
- ▶ The ESM is starting from a default config per the Example on page 85.
- ▶ Cisco switches being used are 10/100 based and we will be setting the ESM port to 100Mbps full duplex. This means a cross-over cable *must* be used for the link between the ESM and the Cisco switch.
 - If using switch ports that support 1000Mbps, the procedure to set speed will be slightly different from what is shown in the example.

Table 5-9 Configuring the ESM

Description and Comments	Actions to perform via Web interface to ESM
<p>Step 2.1- <i>Configure speed and duplex.</i></p> <p>As already noted, it will be necessary to use cross-over cables on the links between the ESM and the Cisco switch, as tuning off auto-negotiation also turns off the auto selection of MDI/MDI-X.</p> <p>Note that LACP <i>must</i> have full duplex connections to operate correctly.</p>	<ul style="list-style-type: none"> ▶ Click the top connector (Ext1) in the picture of the ESM at the top of the page. ▶ Change Speed/Duplex to 100/Full. ▶ Click Apply. ▶ Repeat for Ext2, Ext3 and Ext4.
<p>Step 2.2 - <i>Configure PVIDs</i></p> <p>This places the desired blade server ports in to the desired VLANs. If the VLAN does not exist, it will be created automatically.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> VLANs -> 802.1Q Port Settings -> Port VLAN ID (PVID). ▶ For Bay 2, change the PVID to 5. ▶ For Bays 3 and 4, change the PVID to 10. ▶ All other PVIDs should be set for 1. ▶ Click Apply.

Description and Comments	Actions to perform via Web interface to ESM
<p>Step 2.3 - <i>Configure 802.1Q trunking</i> This will allow Ext1 through Ext4 to carry traffic for VLAN 5 and VLAN 10 (already carrying VLAN 1 as an initial default).</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> VLANs -> Edit 802.1Q VLANs ▶ Select VLAN 5 and click Edit. ▶ Set Interface Ext1, Ext2, Ext3 and Ext4 to Egress and make sure Tag box is checked for each. ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked. ▶ Click Apply. ▶ Click Configuration -> VLANs -> Edit 802.1Q VLANs. ▶ Select VLAN 10 and click Edit. ▶ Set interface Ext1, Ext2, Ext3 and Ext4 to Egress and make sure the Tag box is checked for each. ▶ Make sure Ports 3 and 4 are set for Egress and the Tag box is unchecked for these two ports. ▶ Click Apply.
<p>Step 2.4- <i>Configure Link Aggregation</i>. This example makes use of LACP to dynamically negotiate link aggregation with the Cisco switch. Note that the Ext links used can not already be part of a different aggregation group, static or dynamic.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Setup IEEE 802.3ad Link Aggregation. ▶ Change LACP to Enabled and click Apply. ▶ Click Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Link Aggregation Port Settings. ▶ Set Configure External Port From to include all four ports: <ul style="list-style-type: none"> – Ext1 to Ext4. <ul style="list-style-type: none"> • This will set all 4 ports in to a single link. ▶ Change Mode to Enabled and click Apply.
<p>Step 2.5 - <i>Save ESM config to NVRAM</i> Failure to perform this step will result in all changes to the ESM being lost if the @server BladeCenter is powered off or the ESM is otherwise restarted.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> Maintenance -> Save Changes. ▶ Click Save Configuration. ▶ Click Ok when complete.

Step 3: Configuring the Cisco switch

The following assumptions have been made for this example (reference Table 5-10 on page 110):

- ▶ VLANs 1, 5 and 10 already exist in the Cisco configuration and will be part of the 802.1Q trunk to the ESM.
- ▶ The user is already logged in to the switch and the switch is in enable mode.
- ▶ The lowest available compatible ports are being used:
 - For the CatOS switch being used in this example, ports 2/3 through 2/6 are being used.
 - For the IOS switch being used in this example, ports fa0/1 through fa0/4 are being used.
- ▶ Commands are being performed in the sequence shown.
- ▶ The switch is starting from a default config per the Example on page 88 (IOS) or the Example on page 90 (CatOS).
- ▶ Cisco switch being used is 10/100 based and we will be setting the port to 100Mb full duplex.
 - If using switch ports that support 1000Mbps, the command to set speed will be different from what is shown in the example.

Table 5-10 Configuring the Cisco equipment

Description and Comments	CatOS based switch	IOS based switch
<p>Step 3.1</p> <p><i>Configure speed and duplex.</i></p>	<pre>set port speed 2/3-6 100 set port duplex 2/3-6 full</pre>	<pre>config t int range fa0/1 - 4 speed 100 duplex full</pre> <p>► Note that the <i>range</i> option on IOS is only available in newer revisions of code. If the range option is not available you will need to repeat steps 3.1 through 3.3 for each interface.</p>
<p>Step 3.2</p> <p><i>Configure 802.1Q trunking.</i></p> <p>Forces link to become an 802.1Q VLAN trunk.</p>	<pre>set trunk 2/3 nonegotiate dot1q set trunk 2/4 nonegotiate dot1q set trunk 2/5 nonegotiate dot1q set trunk 2/6 nonegotiate dot1q</pre>	<pre>switchport trunk encapsulation dot1q switchport mode trunk switchport nonegotiate</pre>
<p>Step 3.3</p> <p><i>Configure Link Aggregation.</i></p> <p>Note that if you do not have the option of setting the “channelprotocol” to “lacp” for CatOS, or the channel-group to “active” for IOS, then more then likely you will need to upgrade your code to support LACP.</p>	<pre>set channelprotocol lacp 2</pre> <p>► The trailing 2 in the above command tells CatOS to enable LACP for Module 2</p> <pre>set port lacp-channel 2/3-6 mode active</pre>	<pre>channel-group 1 mode active</pre> <p>► This will create a logical interface named <i>Port-Channel1</i> and place the interfaces fa0/1 through fa0/4 in to it.</p> <pre>end</pre>
<p>Step 3.4</p> <p><i>Save config to NVRAM.</i></p> <p>Only necessary on IOS based switches.</p>	<p><i>(does not apply)</i></p>	<pre>write mem</pre>

Step 4: Reconnecting the devices

This is the final step to bring the connection in to full operation. This will be the reverse of whatever procedure was used in Step 1. See Table 5-2 on page 84 for details on how to reestablish the links.

Step 5: Verifying the configuration

This section provides options for verifying the correct and desired operation.

Verifying correct operation on the BladeCenter ESM

Verify the configurations on the ESM look as follows:

As in Figure 5-21 on page 111, verify port state by clicking **Configuration -> Port Setting -> Display Ports**, and reviewing the status of Ext1 through Ext4 (State should be *Enabled*, Speed/Duplex should be *100/Full*, and Connection should show *100M/Full/802.3x* for all 4 Ext ports).

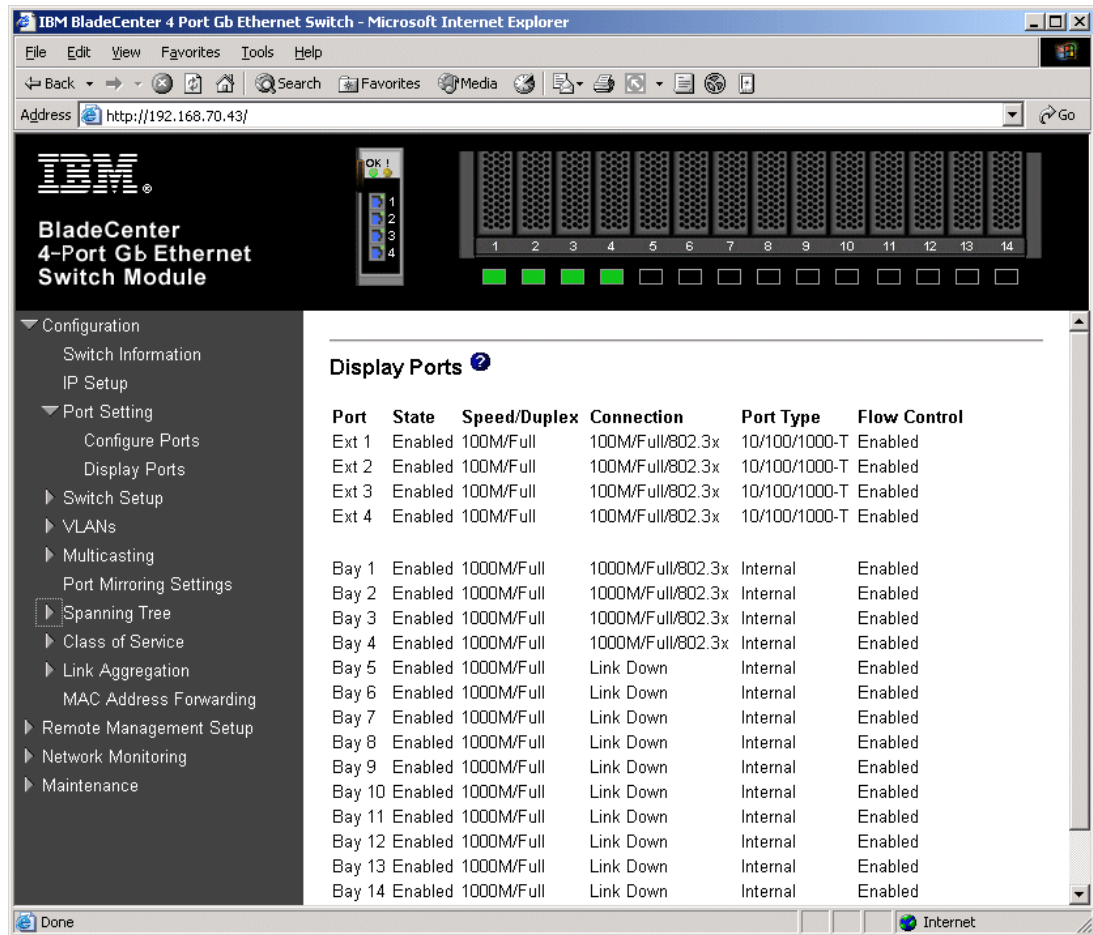


Figure 5-21 Verify ports are operational

As in Figure 5-22 on page 112, verify VLAN 1 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 1** and clicking **Edit**. Make sure the ports not on VLAN 1 (Ports 2, 3 and 4) are set to None and their associated **Tag** box is checked. Also note that the options for Ext ports 2, 3 and 4 are greyed out. Once an aggregation forms, you can only make changes to the lowest number port (in this case, Ext1), which will then be inherited by the other ports in the aggregation when the change is applied. Also note that if the aggregation had not yet formed (for example, if the cables were still not plugged in), then the other Ext ports would not be greyed out, and you could make changes to them. Of course, as soon as the aggregation came up, those changes would be overwritten by the lowest numbered port in the aggregation. The other settings could still be seen through the greyed out boxes, but they would not be getting used.

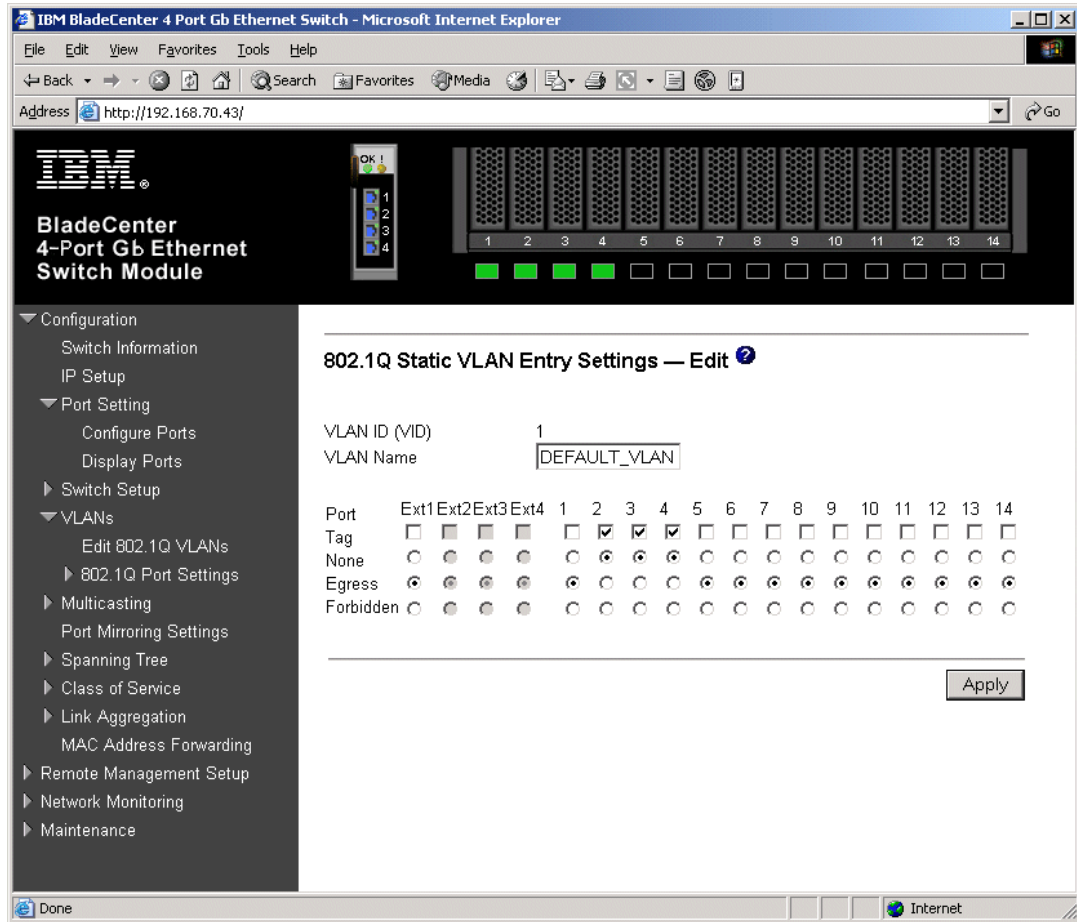


Figure 5-22 Verifying VLAN 1

As in Figure 5-23 on page 113, verify VLAN 5 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 5** and clicking **Edit**. Make sure **Port 2** is set to Egress and that the Tag box is not checked. Make sure the Ext1 interface is set for Egress (allows VLAN 5 traffic to pass through Ext1 and any aggregated link that is part of Ext1's aggregation) and that the box for **Tag** is checked. Also, note the same greying out of Ext2, 3 and 4 will also be seen here, if the aggregation has already become active.

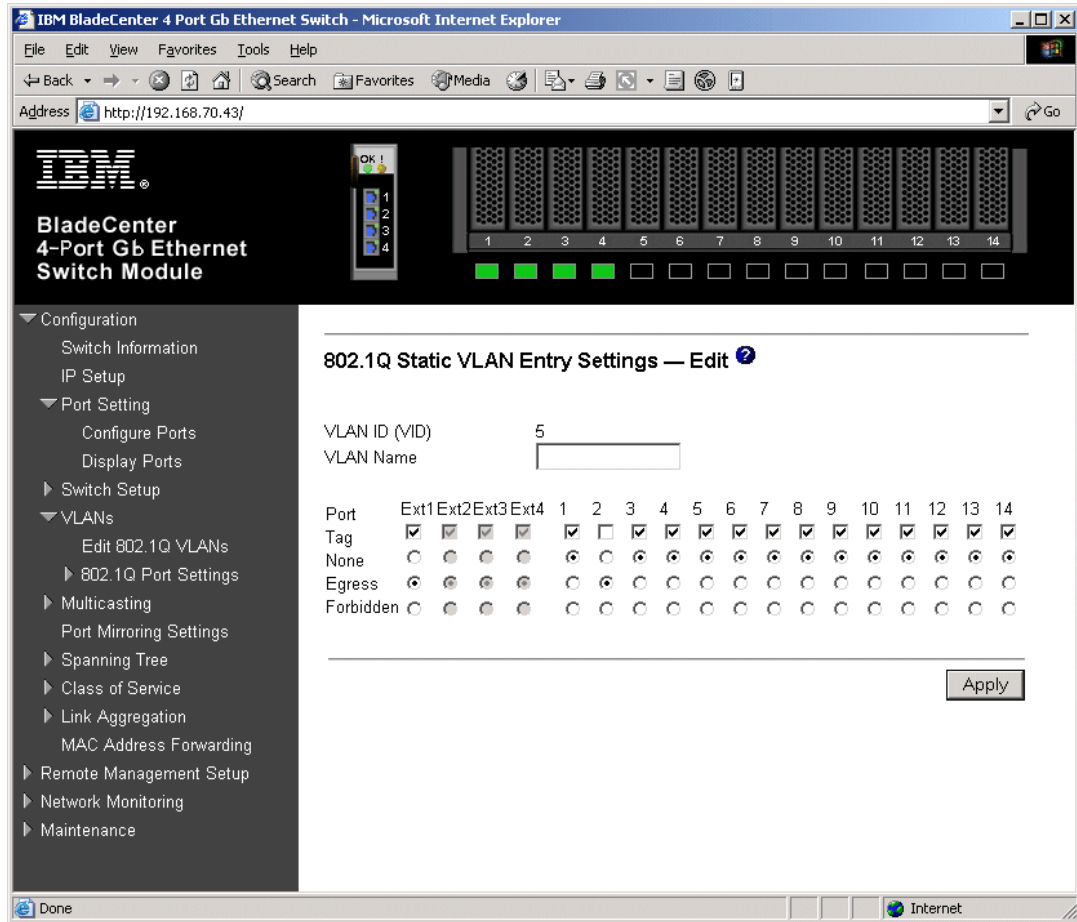


Figure 5-23 Verify VLAN 5

As in Figure 5-24 on page 114, verify VLAN 10 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 10** and clicking **Edit**. Make sure ports 3 and 4 are set to Egress and that the Tag box is not checked for either port. Make sure the Ext1 port is set for Egress (allows VLAN 10 traffic to pass through Ext1 and any aggregated link that is port of Ext1's aggregation) and that the box for **Tag** is checked. Here again we see the greying out of Ext2, 3 and 4 after an aggregation has formed.




Figure 5-24 Verify VLAN 10

As in Figure 5-25 on page 115, verify Spanning Tree is forwarding traffic for the aggregation by clicking **Configuration -> Spanning Tree -> STP Port Settings**. Check to see that the Status for Ext1 through Ext4 are forwarding. Notice here that you lose the ability to individually manage Spanning Tree options for all ports in the aggregation except the lowest numbered port (and that the aggregation handles Spanning Tree as a whole, rather than on a port by port basis). In this case, changes to Ext 1 will be inherited by the other ports in the aggregation. Again, this is only true if the aggregation has already formed.

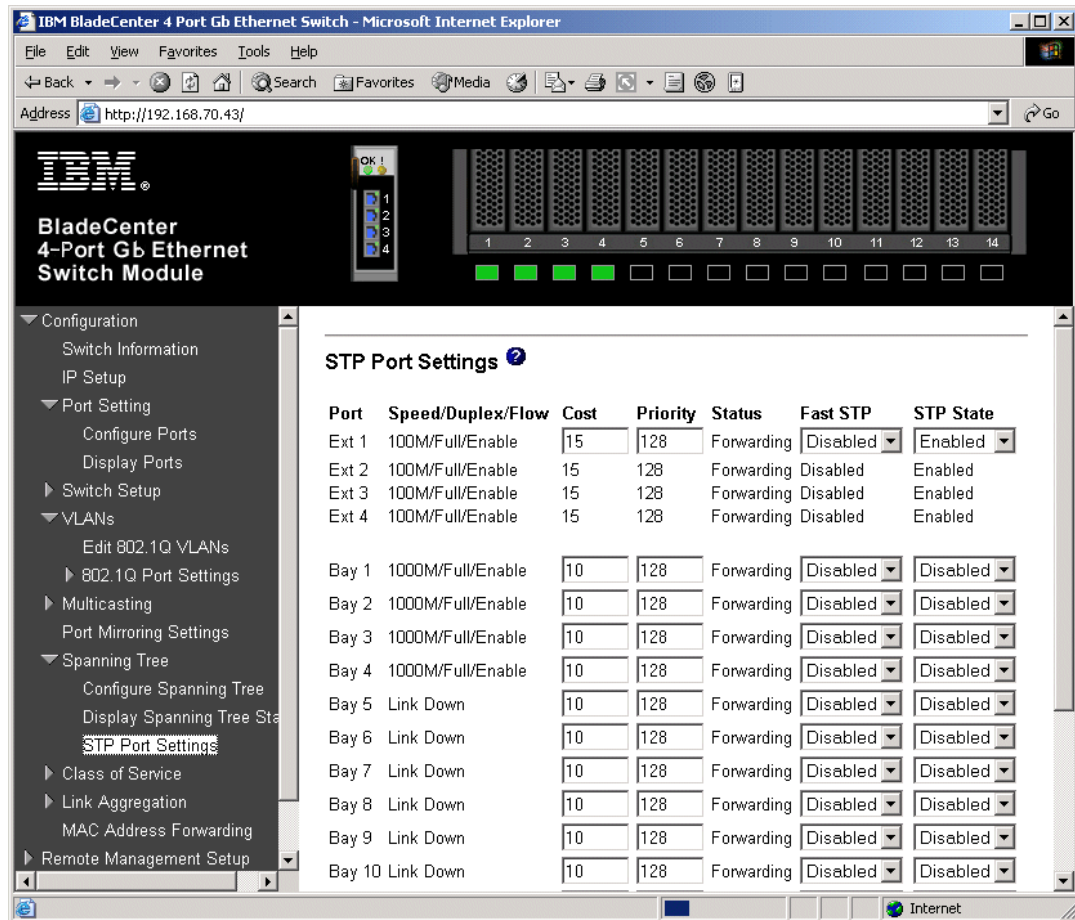


Figure 5-25 Verifying STP is forwarding on all 4 links in the aggregation

As in Figure 5-26 on page 116, one way to review the status of the aggregation is by clicking **Configuration -> Network Monitoring -> Applications Status -> Link Aggregation**. From here we can see that all 4 ports have joined the aggregation.

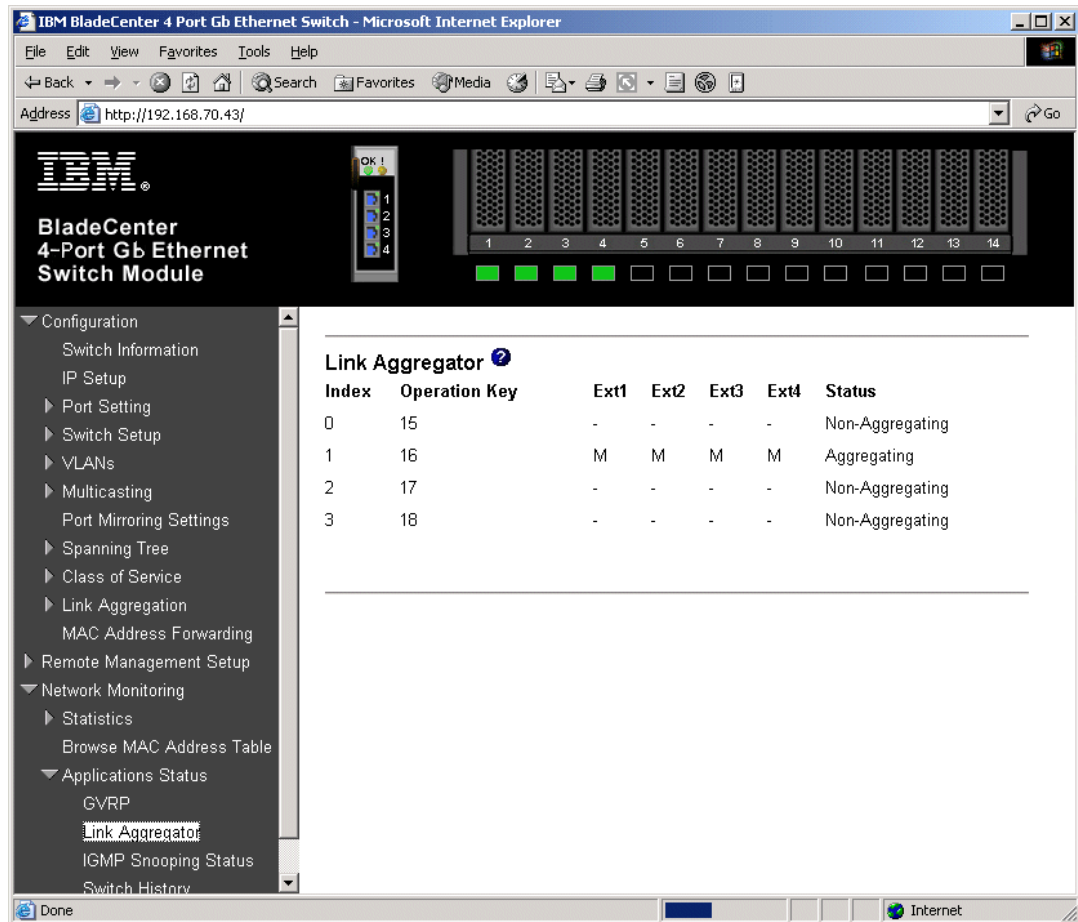


Figure 5-26 Link aggregation status

As in Figure 5-27 on page 117, another way to review the status of the aggregation is by clicking **Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Link Aggregation Port Setting** and noting that the Mode shows Enabled for all 4 ports (true whether the aggregation has formed or not), and the Status shows Active. (Shows Active after the aggregation has formed, shows Individual if the aggregation has not yet formed).

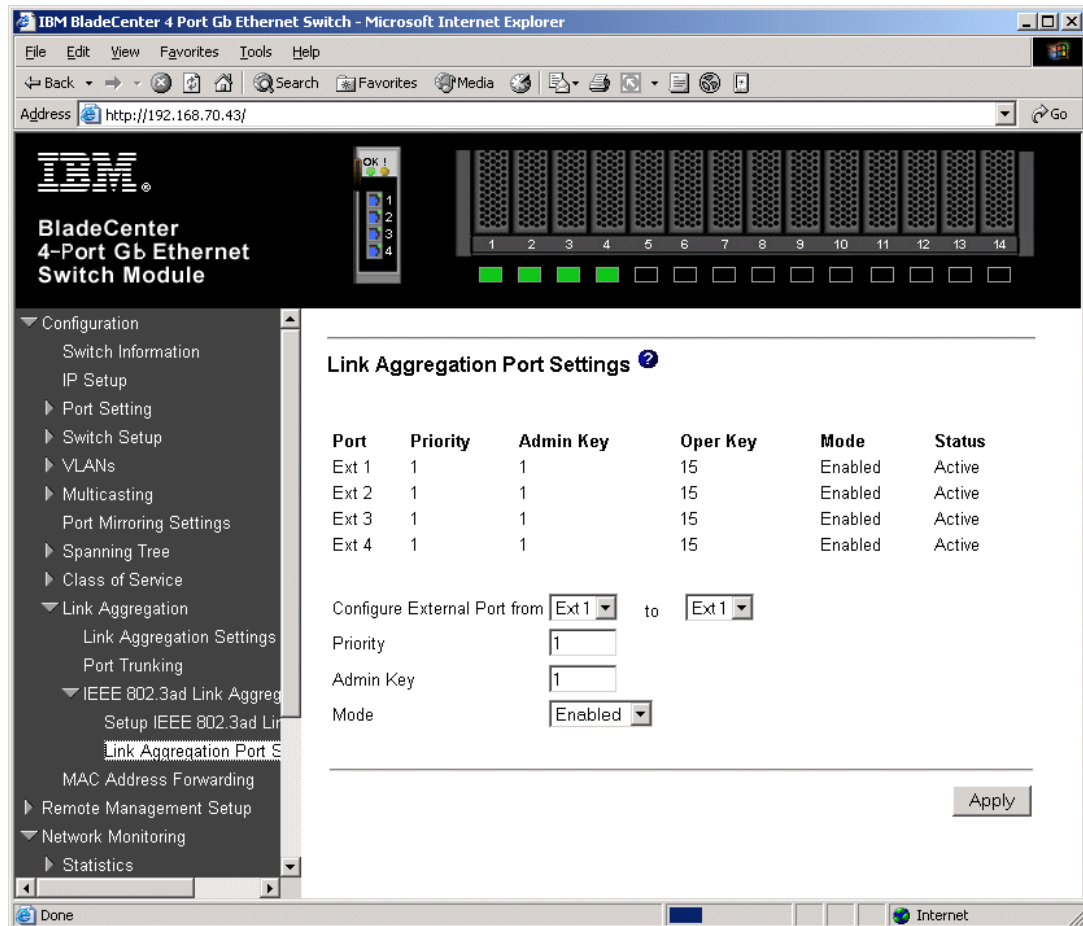


Figure 5-27 Link aggregation status, another view

As in Figure 5-28 on page 118, you can also verify the aggregation is operational by generating traffic across the link and clicking **Network Monitoring -> Statistics -> Port Utilization** and monitoring **Tx/sec** and **Rx/sec** for the four Ext ports. Depending on how much traffic is being generated, and from how many sources, the numbers will vary across the ports. For this example, 100,000, 1400 byte pings were being sent from the Cisco switch to the IP address of the ESM. In this case, Ext1 is being used to transmit and Ext4 is receiving (these could have been on different ports, or even transmitting and receiving on the same port). If one were to pull the cable for Ext1, the traffic would switch over to a different Ext port, usually with no loss of packets (usually under 1 second) (see Figure 5-29 on page 119).

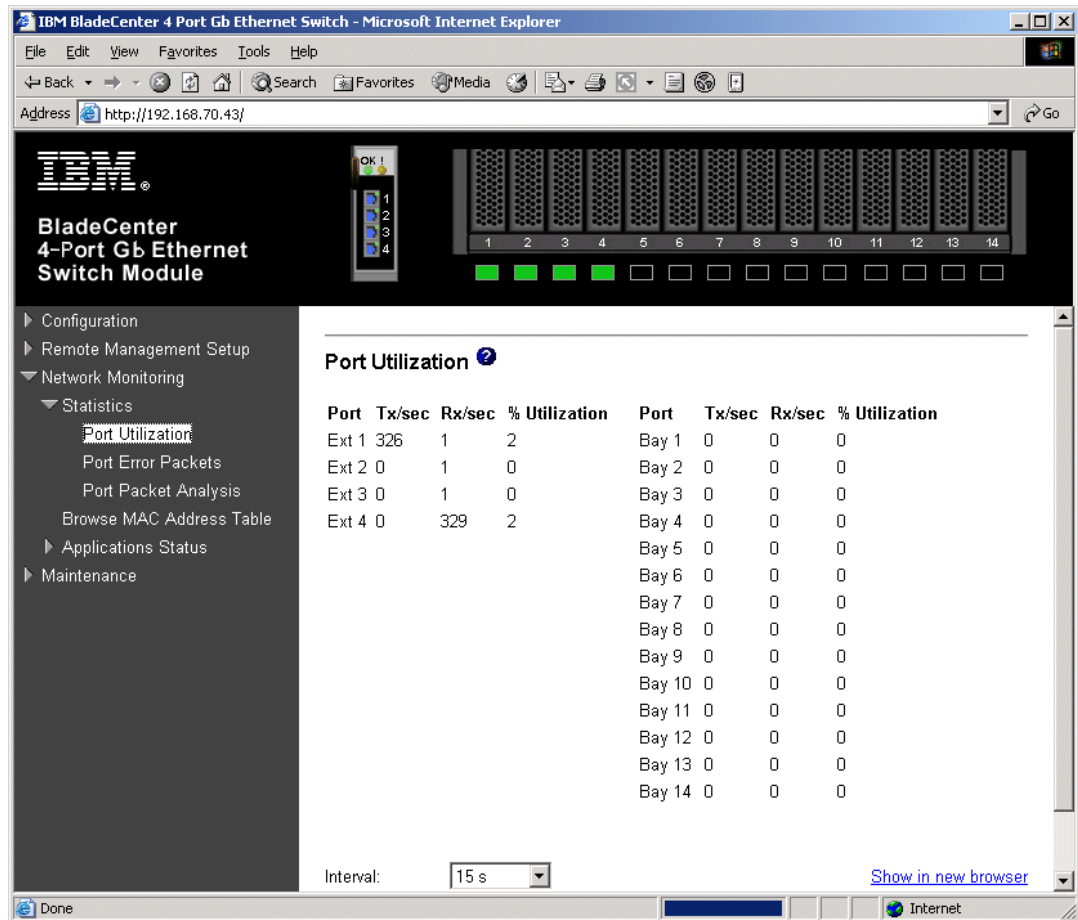


Figure 5-28 Example of aggregation load balancing

As in Figure 5-29 on page 119, with the same traffic flowing on the link, the cable to Ext1 is removed. The result is that the traffic previously carried on Ext1 is now on Ext2 (it could have gone to any of the available Ext ports in the aggregation).

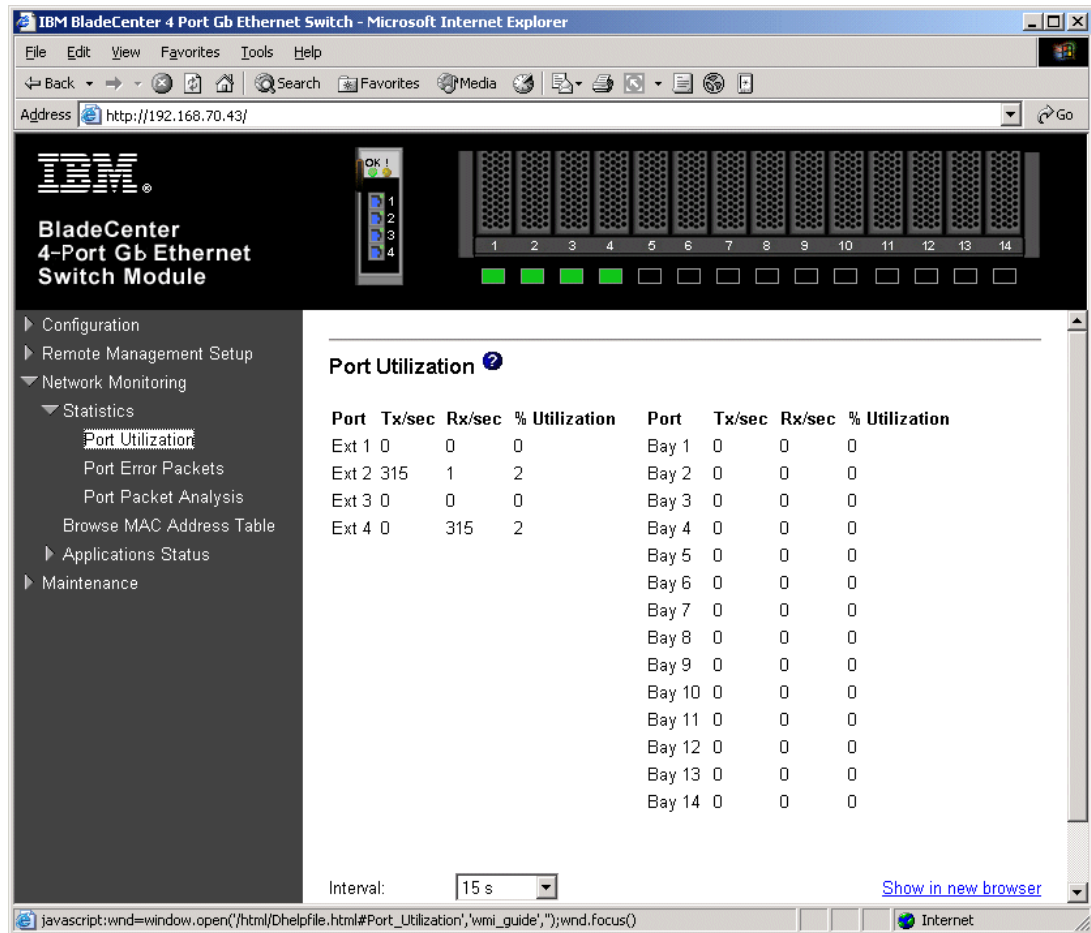


Figure 5-29 Traffic flow after cable to Ext1 removed

Verifying correct operation on the Cisco external switch

The following section includes some commands one can use to verify the desired configuration and operation of the Cisco equipment (reference Table 5-11 on page 120).

Table 5-11 Verifying the configuration and operation of the Cisco external switch of the connection

Description and Comments	CatOS based switch	IOS based switch
Review running config for desired statements.	show config <ul style="list-style-type: none"> ► Review for the following: <ul style="list-style-type: none"> – set channelprotocol lacp 2 – set port lacp-channel 2/3-6 52 <ul style="list-style-type: none"> • The number at the end may vary. – set port speed 2/3-6 100 – set port duplex 2/3-6 full – set trunk 2/3 nonegotiate dot1q 1-1005,1025-4094 <ul style="list-style-type: none"> • Should see a similar entry for each configured port (2/3 through 2/6). – set port lacp-channel 2/3-6 mode active 	show run <ul style="list-style-type: none"> ► Review for the following on interface Port-channel1: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport nonegotiate ► Note that the values in Port-channel1 may not show up if the aggregation has never come up since first being configured. ► Review for the following on int fa0/1 through fa0/4: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport nonegotiate – duplex full – speed 100 – channel-group 1 mode active
Show speed and duplex.	<ul style="list-style-type: none"> ► Do the following command on each interface, 2/3 through 2/6: show port status 2/3 ► Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100 	<ul style="list-style-type: none"> ► Do the following command on each interface, fa0/1 through fa0/4: show int fa0/1 status ► Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100 ► Note that if you do this on the Port-channel1 interface, the speed and duplex might show up as a-100 and a-full.
Show trunking status. Aggregation must be up before the trunk will come up.	show port trunk 2/3 <ul style="list-style-type: none"> ► Should show the following: <ul style="list-style-type: none"> – Mode = nonegotiate – Encapsulation = dot1q – Status = Trunking – Native VLAN = 1 	show int Port-channel1 trunk <ul style="list-style-type: none"> ► Should show the following: <ul style="list-style-type: none"> – Mode = on – Encapsulation = 802.1q – Status = Trunking – Native VLAN = 1
Review status of the aggregated link.	show port lacp-channel <ul style="list-style-type: none"> ► Should see all 4 ports in Channel Mode Active show port lacp-channel statistics <ul style="list-style-type: none"> ► Run several times in a row, should show LACP Pkts Transmitted and Received climbing slowly (Transmitted usually higher than Received). 	show etherchannel summary <ul style="list-style-type: none"> ► Should show the following: <ul style="list-style-type: none"> – Protocol = LACP – Ports fa0/1 -4 = (P) <ul style="list-style-type: none"> • (P) = part of an aggregation group show lacp counters <ul style="list-style-type: none"> ► Run several times in a row, should show LACPDUs Sent and Recv climbing slowly (Sent usually higher than Recv). show etherchannel port-channel <ul style="list-style-type: none"> ► Should show the following: <ul style="list-style-type: none"> – Port state = Port-channel Ag-Inuse – Protocol = LACP – All 4 ports in EC State Active

Description and Comments	CatOS based switch	IOS based switch
Ping the ESM. Where x.x.x.x is the IP address of the ESM (must be in same VLAN as subnet being pinged).	ping x.x.x.x ► Based on the Cat4K being at 192.168.70.202 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.	ping x.x.x.x ► Based on the Cat 3550 being at 192.168.70.200 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.
Attach a device to VLAN 5 and 10 on the Cisco switch and verify a ping to one of the blade servers in that same VLAN works as desired.	► For VLAN 5, attach a device to port 2/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work) ► For VLAN 10, attach a device to port 2/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).	► For VLAN 5, attach a device to port fa0/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work). ► For VLAN 10, attach a device to port fa0/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).

5.5.4 Single ESM, four port static aggregation to a single Cisco switch

This example (reference Figure 5-30) is identical to the example given in 5.5.3, “Single ESM, four port LACP aggregation to a single Cisco switch” on page 107 except that the aggregation is built in a static fashion. This example is unique in this chapter, in that it is the only one showing static aggregation. While LACP is the preferred method for aggregation, there are many older versions of IOS and CatOS that do not support LACP. Information from this example can be applied to any other example, if required, to accommodate older Cisco products.




Figure 5-30 Single ESM with 4 port static aggregation in to a single Cisco switch

Step 1: Take down the link(s)

It is always advisable to disable the link(s) prior to making any configuration changes. Please see Table 5-1 on page 84 for procedures.

Step 2: Configure the IBM side of the link

This section walks you through the sequence of actions required to configure the ESM for this example (reference Table 5-12 on page 122).

The following assumptions have been made for this example:

- The user is already logged in to the ESM with *root* level access.

- ▶ Commands are being performed in the sequence shown.
- ▶ The ESM is starting from a default config per the Example on page 85.
- ▶ Cisco switches being used are 10/100 based and we will be setting the ESM port to 100Mbps full duplex. This means a cross-over cable *must* be used for the link between the ESM and the Cisco switch.
 - If using switch ports that support 1000Mbps, the procedure to set speed will be slightly different from what is shown in the example.

Table 5-12 Configuring the ESM

Description and Comments	Actions to perform via Web interface to ESM
<p>Step 2.1- <i>Configure speed and duplex.</i> As already noted, it will be necessary to use cross-over cables on the links between the ESM and the Cisco switch, as tuning off auto-negotiation also turns off the auto selection of MDI/MDI-X. Note that LACP <i>must</i> have full duplex connections to operate correctly.</p>	<ul style="list-style-type: none"> ▶ Click the top connector (Ext1) in the picture of the ESM at the top of the page. ▶ Change Speed/Duplex to 100/Full. ▶ Click Apply. ▶ Repeat for Ext2, Ext3 and Ext4.
<p>Step 2.2 - <i>Configure PVIDs</i> This places the desired blade server ports in to the desired VLANs. If the VLAN does not exist, it will be created automatically.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> VLANs -> 802.1Q Port Settings -> Port VLAN ID (PVID). ▶ For Bay 2, change the PVID to 5. ▶ For Bays 3 and 4, change the PVID to 10. ▶ All other PVIDs should be set for 1. ▶ Click Apply.
<p>Step 2.3 - <i>Configure 802.1Q trunking</i> This will allow Ext1 through Ext4 to carry traffic for VLAN 5 and VLAN 10 (already carrying VLAN 1 as an initial default).</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> VLANs -> Edit 802.1Q VLANs ▶ Select VLAN 5 and click Edit. ▶ Set Interface Ext1, Ext2, Ext3 and Ext4 to Egress and make sure Tag box is checked for each. ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked. ▶ Click Apply. ▶ Click Configuration -> VLANs -> Edit 802.1Q VLANs. ▶ Select VLAN 10 and click Edit. ▶ Set interface Ext1, Ext2, Ext3 and Ext4 to Egress and make sure the Tag box is checked for each. ▶ Make sure Ports 3 and 4 are set for Egress and the Tag box is unchecked for these two ports. ▶ Click Apply.
<p>Step 2.4- <i>Configure Link Aggregation.</i> This example makes use of hard coded (static) link aggregation. Note that the Ext links used can not already be part of a different aggregation group, static or dynamic.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> Link Aggregation -> Port Trunking. ▶ Under Group ID 1, check each box for Ext1, Ext2, Ext3 and Ext4. ▶ Change Method for Group ID 1 to Enable, and click Apply.
<p>Step 2.5 - <i>Save ESM config to NVRAM</i> Failure to perform this step will result in all changes to the ESM being lost if the @server BladeCenter is powered off or the ESM is otherwise restarted.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> Maintenance -> Save Changes. ▶ Click Save Configuration. ▶ Click Ok when complete.

Step 3: Configuring the Cisco switch

The following assumptions have been made for this example (reference Table 5-13 on page 123):

- ▶ VLANs 1, 5 and 10 already exist in the Cisco configuration and will be part of the 802.1Q trunk to the ESM.
- ▶ The user is already logged in to the switch and the switch is in enable mode.
- ▶ The lowest available compatible ports are being used:
 - For the CatOS switch being used in this example, ports 2/3 through 2/6 are being used.
 - For the IOS switch being used in this example, ports fa0/1 through fa0/4 are being used.
- ▶ Commands are being performed in the sequence shown.
- ▶ The switch is starting from a default configuration as per “Cat 3550 (IOS based switch) base configuration” on page 88 (IOS) or “Cat 4003 (CatOS based switch) base configuration” on page 90 (CatOS).
- ▶ Cisco switch being used is 10/100 based and we will be setting the port to 100Mb full duplex.
 - If using switch ports that support 1000Mbps, the command to set speed will be different from what is shown in the example.

Table 5-13 Configuring the Cisco equipment

Description and Comments	CatOS based switch	IOS based switch
Step 3.1 <i>Configure speed and duplex.</i>	<pre>set port speed 2/3-6 100 set port duplex 2/3-6 full</pre>	<pre>config t int range fa0/1 - 4 speed 100 duplex full</pre> <p>▶ Note that the <i>range</i> option on IOS is only available in newer revisions of code. If the range option is not available you will need to repeat steps 3.1 through 3.3 for each interface.</p>
Step 3.2 <i>Configure 802.1Q trunking.</i> Forces link to become an 802.1Q VLAN trunk.	<pre>set trunk 2/3 nonegotiate dot1q set trunk 2/4 nonegotiate dot1q set trunk 2/5 nonegotiate dot1q set trunk 2/6 nonegotiate dot1q</pre>	<pre>switchport trunk encapsulation dot1q switchport mode trunk switchport nonegotiate</pre>
Step 3.3 <i>Configure static link aggregation.</i>	<pre>set port channel 2/3-6 mode on</pre>	<pre>channel-group 1 mode on end</pre>
Step 3.4 <i>Save config to NVRAM.</i> Only necessary on IOS based switches.	<i>(does not apply)</i>	<pre>write mem</pre>

Step 4: Reconnecting the devices

This is the final step to bring the connection in to full operation. This will be the reverse of whatever procedure was used in Step 1. See Table 5-2 on page 84 for details on how to reestablish the links.

Step 5: Verifying the configuration

This section provides options for verifying the correct and desired operation.

Verifying correct operation on the BladeCenter ESM

Verify the configurations on the ESM look as follows:

As in Figure 5-31 on page 124, verify the port state by clicking **Configuration -> Port Setting -> Display Ports**, and reviewing the status of Ext1 through Ext4 (State should be *Enabled*, Speed/Duplex should be *100/Full*, and Connection should show 100M/Full/802.3x for all 4 Ext ports).

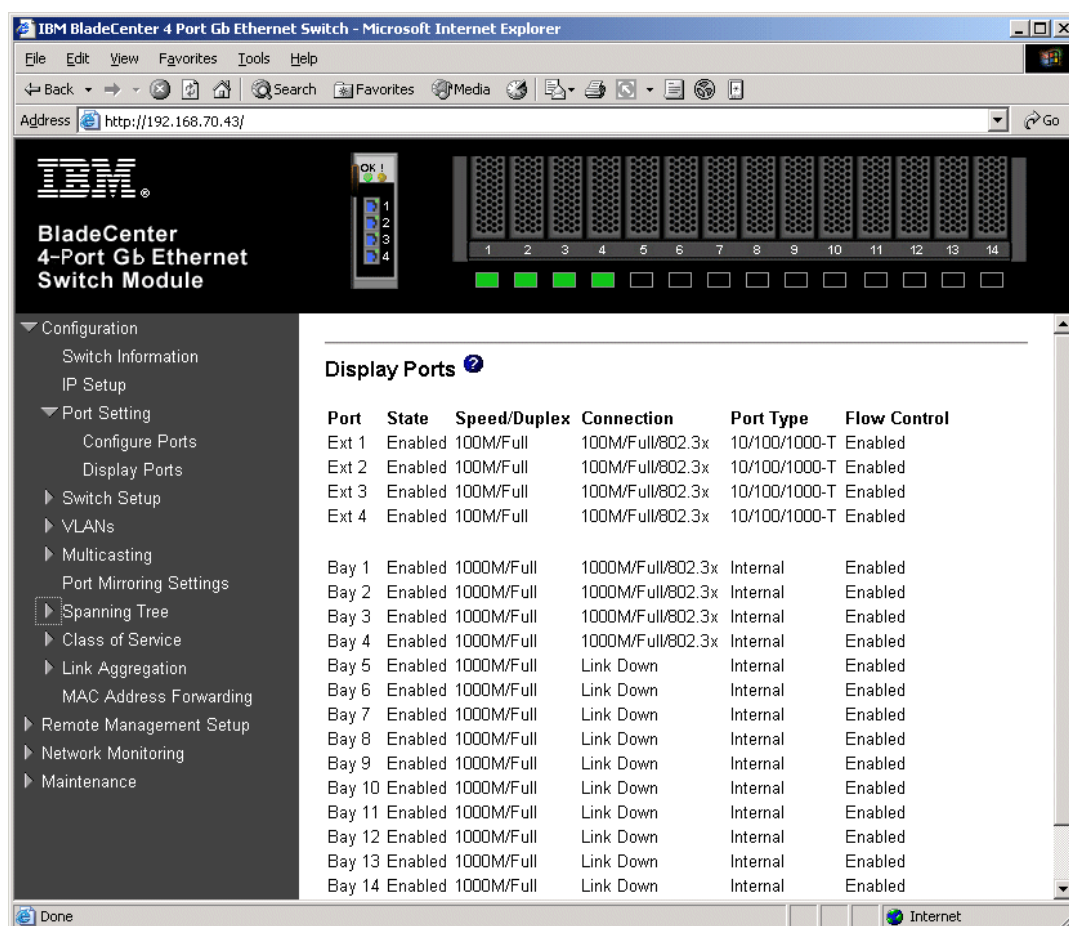


Figure 5-31 .Verify ports are operational

As in Figure 5-32 on page 125, verify VLAN 1 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 1** and clicking **Edit**. Make sure the ports not on VLAN 1 (Ports 2, 3 and 4) are set to None and their associated **Tag** box is checked. Also note that the options for Ext ports 2, 3 and 4 are greyed out. Unlike an LACP aggregation, when using static, this grey out occurs as soon as you configure the aggregation (links do not have to be up). Like the LACP configuration, you can only make changes to the lowest number port (in this case, Ext1), which will then be inherited by the other ports in the aggregation when the change is applied

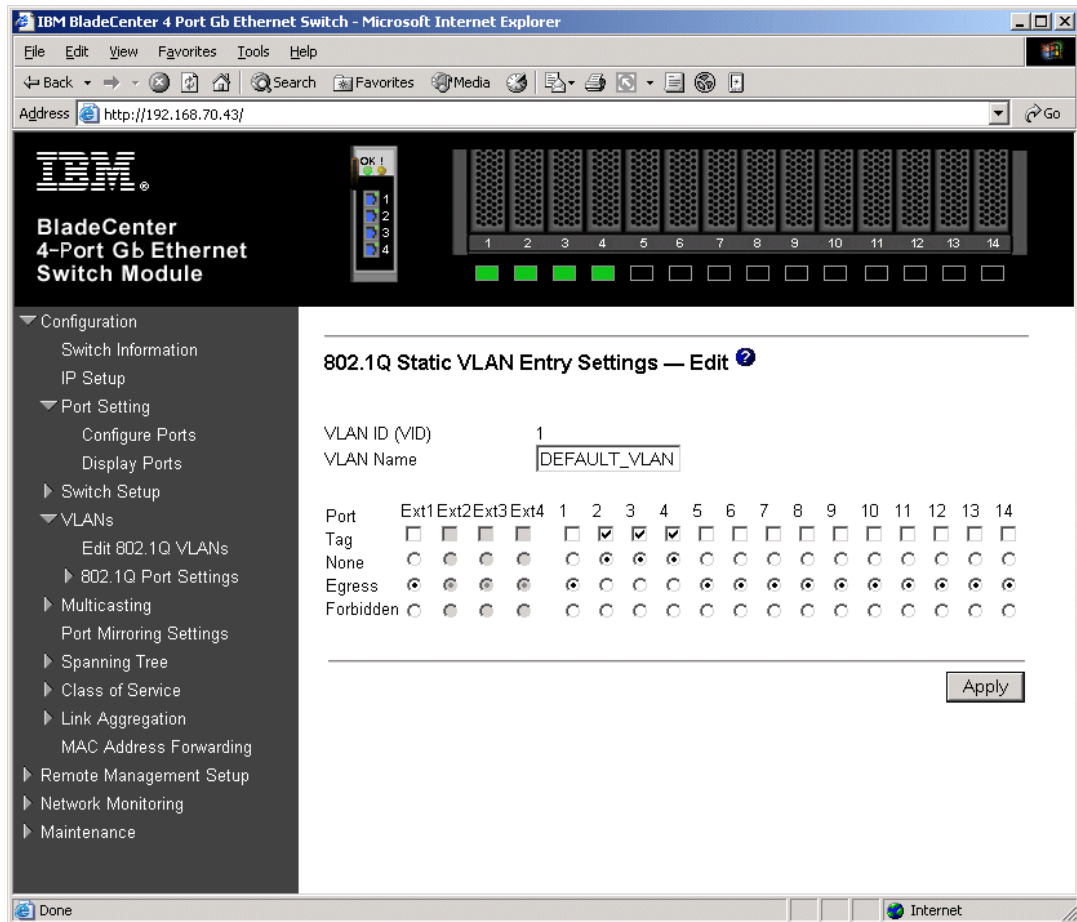


Figure 5-32 Verifying VLAN 1

As in Figure 5-33 on page 126, verify VLAN 5 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 5** and clicking **Edit**. Make sure **Port 2** is set to Egress and that the Tag box is not checked. Make sure the Ext1 interface is set for Egress (allows VLAN 5 traffic to pass through Ext1 and any aggregated link that is part of Ext1's aggregation) and that the box for **Tag** is checked. Also, note the same greying out of Ext2, 3 and 4 will also be seen here, regardless if the links in the aggregation are up or down.

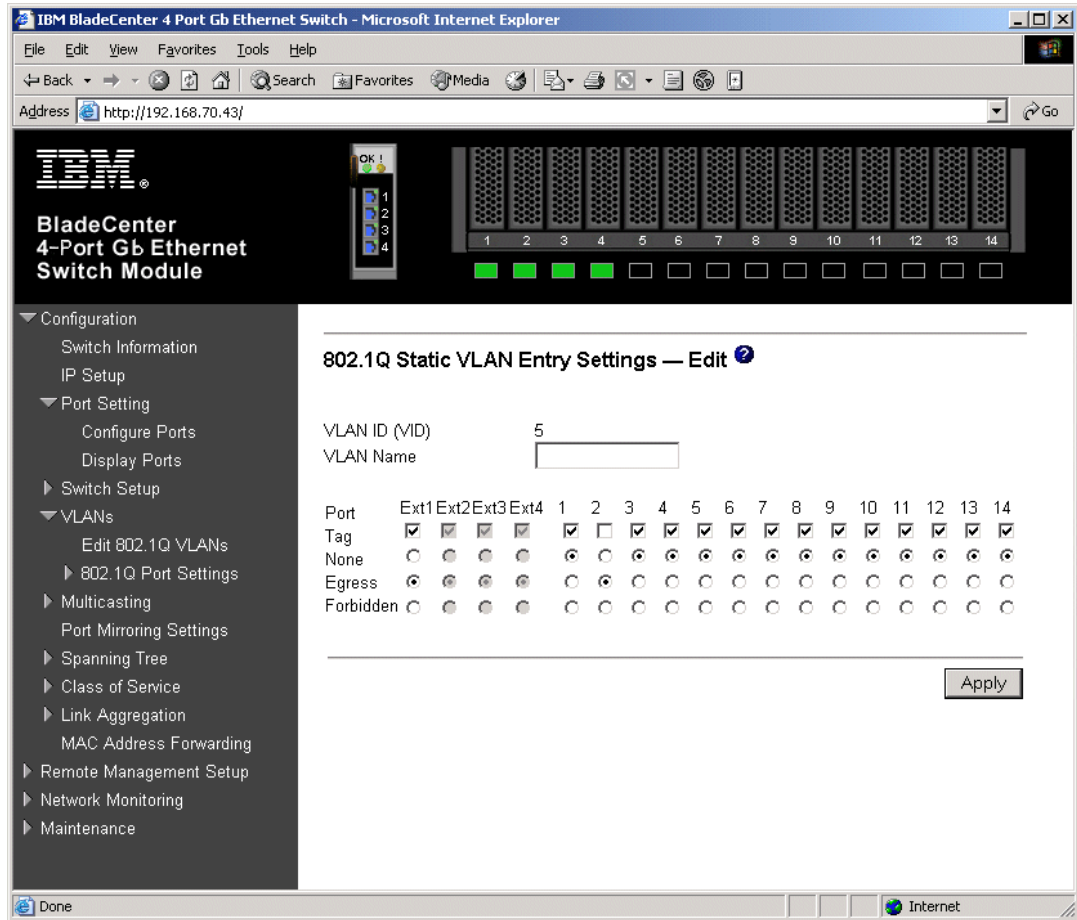


Figure 5-33 Verify VLAN 5

As in Figure 5-34 on page 127, verify VLAN 10 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 10** and clicking **Edit**. Make sure ports 3 and 4 are set to Egress and that the Tag box is not checked for either port. Make sure the Ext1 port is set for Egress (allows VLAN 10 traffic to pass through Ext1 and any aggregated link that is port of Ext1's aggregating) and that the box for **Tag** is checked. Here again we see the greying out of Ext2, 3 and 4, regardless of aggregation link status.

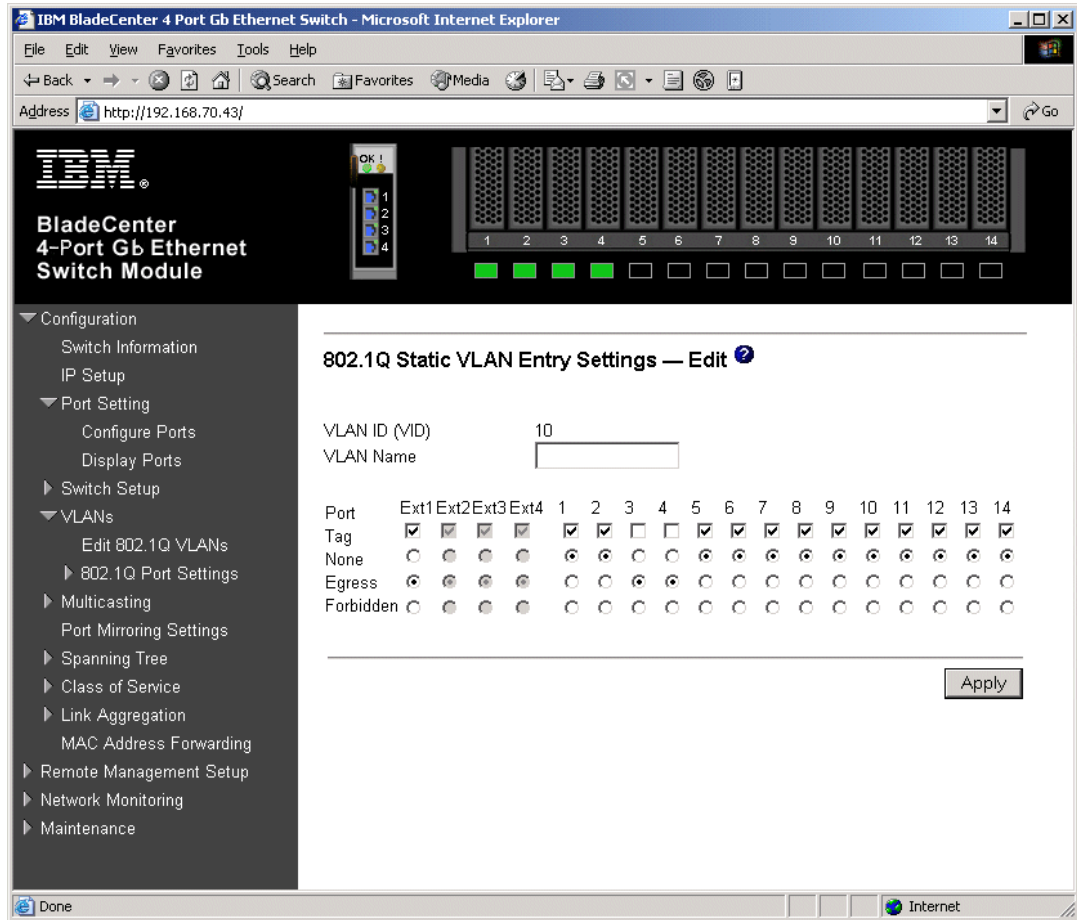


Figure 5-34 Verify VLAN 10

As in Figure 5-35 on page 128, verify Spanning Tree is forwarding traffic for the aggregation by clicking **Configuration -> Spanning Tree -> STP Port Settings**. Check to see that the Status for Ext1 through Ext4 are forwarding. Notice here that you lose the ability to individually manage Spanning Tree options for all ports in the aggregation except the lowest numbered port (and that the aggregation handles Spanning Tree as a whole, rather than on a port by port basis). In this case, changes to Ext 1 will be inherited by the other ports in the aggregation. Again, unlike LACP, static aggregation will show this locked down state (can't make changes to higher numbered links in the aggregation), whether the links in the aggregation are up or down.

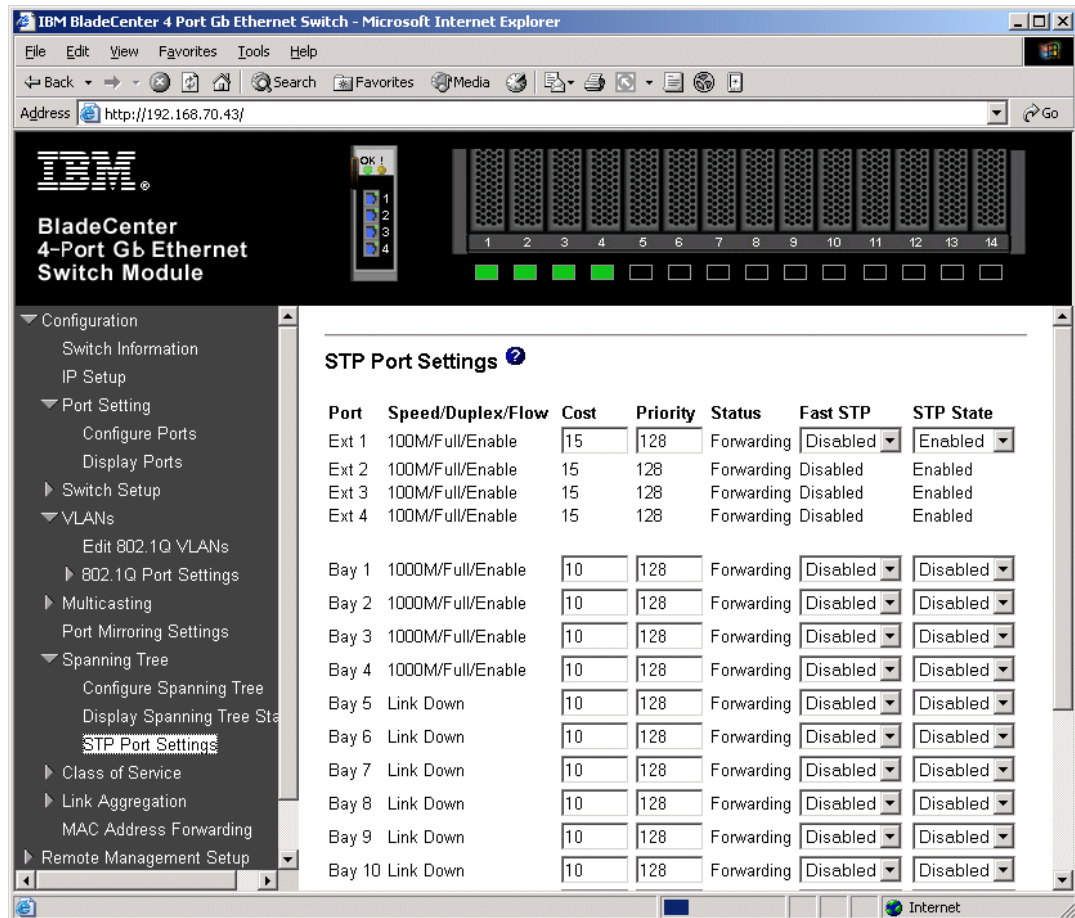


Figure 5-35 Verifying STP is forwarding on all 4 links in the aggregation

As in Figure 5-36 on page 129, one way you can *not* review the status of the aggregation is by clicking **Configuration -> Network Monitoring -> Applications Status -> Link Aggregation**. This window shows no aggregations exist, as it only shows aggregations that have been formed using LACP (and our current example is statically aggregated).

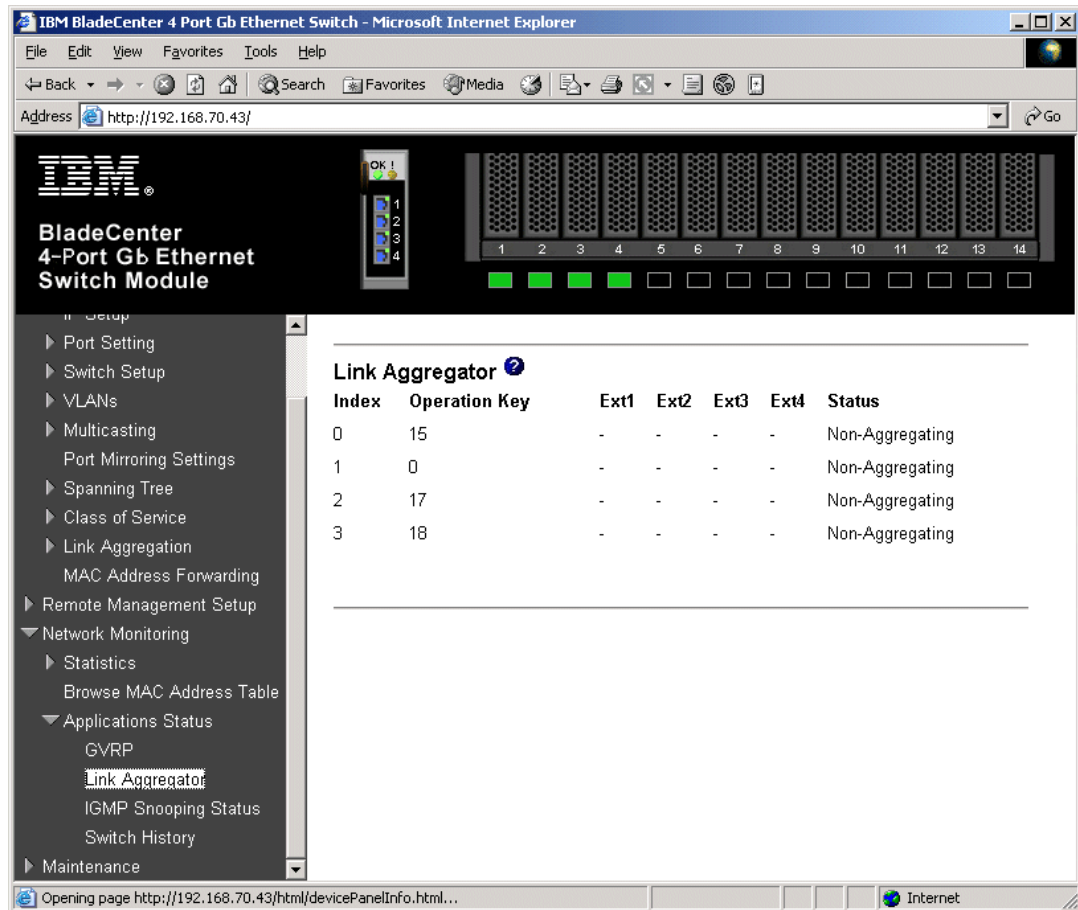


Figure 5-36 Link aggregation status

As in Figure 5-37 on page 130, you can review your static settings by going back to the original screen used to make the configuration (click **Configuration -> Link Aggregation -> Port Trunking**) and noting that the **Ext1** through **Ext4** are all checked, and that Method still shows Enabled).

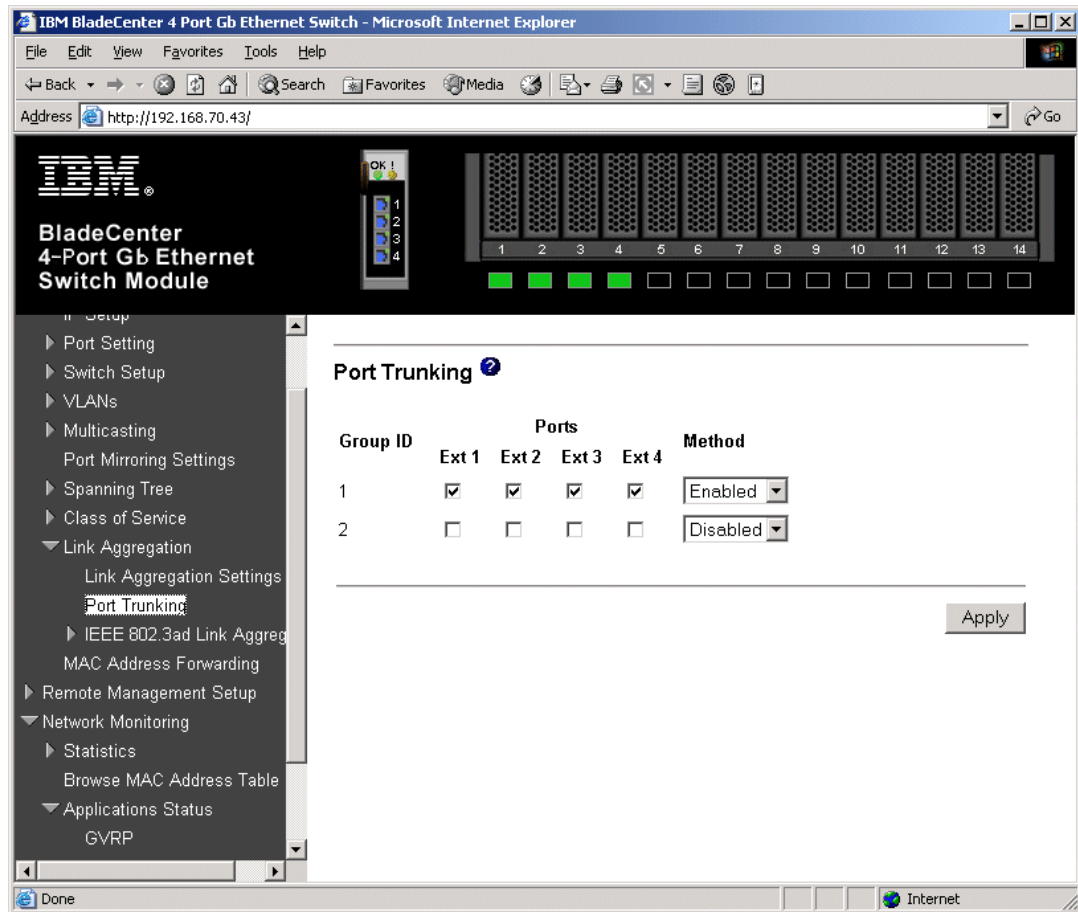


Figure 5-37 Link aggregation status, another view

As in Figure 5-38 on page 131, you can also verify the aggregation is operational by generating traffic across the link and clicking **Network Monitoring -> Statistics -> Port Utilization** and monitoring Tx/sec and Rx/sec for the four Ext ports. Depending on how much traffic is being generate, and from how many sources, the numbers will vary across the ports. For this example, 100,000, 1400 byte pings were being sent from the Cisco switch to the IP address of the ESM. In this case, Ext1 is being used to transmit and Ext4 is receiving (these could have been on different ports, or even transmitting and receiving on the same port). If one were to pull the cable for Ext1, the traffic would switch over to a different Ext port, usually with no loss of packets (usually under 1 second).




Figure 5-38 Example of aggregation load balancing

As in Figure 5-39 on page 132, with the same traffic flowing on the link, the cable to Ext1 is removed. The result is that the traffic previously carried on Ext1 is now on Ext2 (it could have gone to any of the available Ext ports in the aggregation).

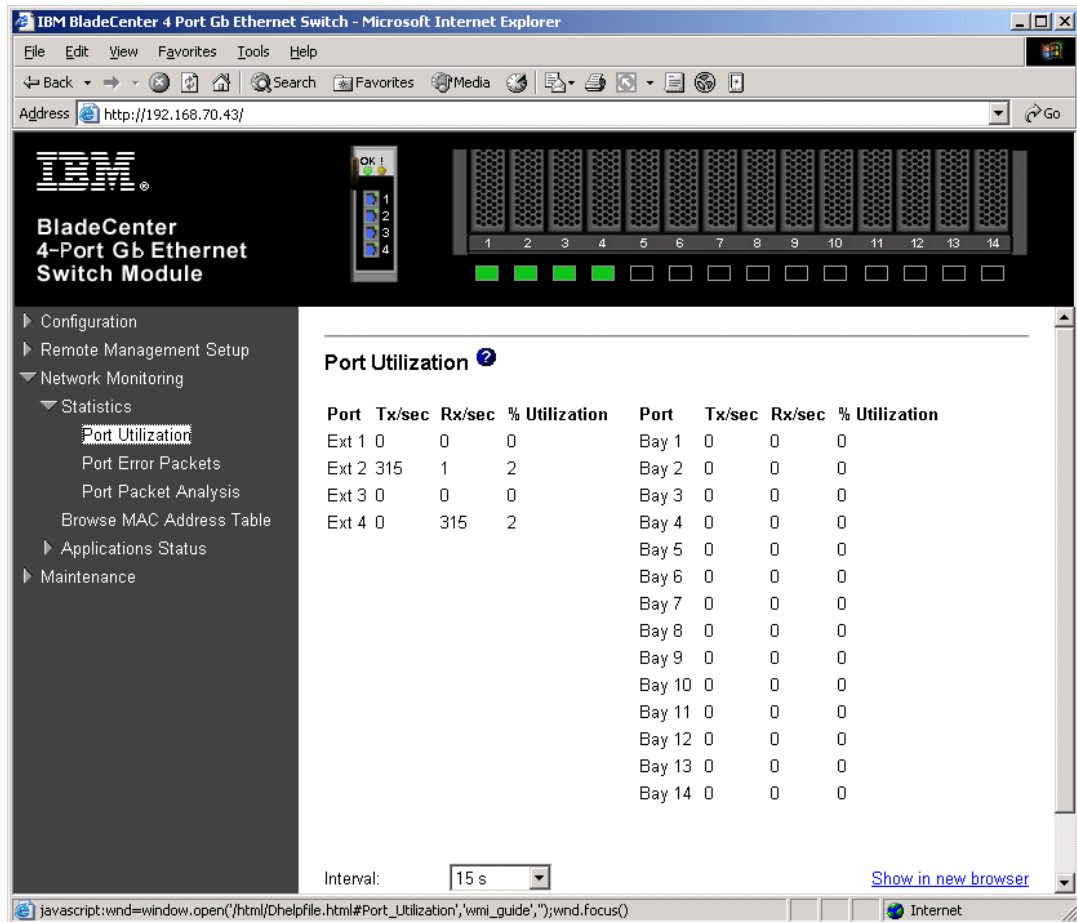


Figure 5-39 Traffic flow after cable to Ext1 removed

Verifying correct operation on the Cisco external switch

The following section includes some commands one can use to verify the desired configuration and operation of the Cisco equipment (reference Table 5-14 on page 133).

Table 5-14 Verifying the configuration and operation of the Cisco external switch of the connection

Description and Comments	CatOS based switch	IOS based switch
Review running config for desired statements.	show config ▶ Review for the following: <ul style="list-style-type: none"> – set port speed 2/3 100 – set port duplex 2/3 full – set trunk 2/3 negotiate dot1q 1-1005,1025-4094 – set port channel 2/3-6 mode on 	show run ▶ Review for the following on interface Port-channel1: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport negotiate ▶ Note that the values in Port-channel1 may not show up if the aggregation has never come up since first being configured. ▶ Review for the following on int fa0/1 through fa0/4: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport negotiate – duplex full – speed 100 – channel-group 1 mode on
Show speed and duplex.	▶ Do the following command on each interface, 2/3 through 2/6 show port status 2/3 ▶ Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100 	▶ Do the following command on each interface, fa0/1 through fa0/4 show int fa0/1 status ▶ Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100 ▶ Note that if you do this on the Port-channel1 interface, the speed and duplex might show up as a-100 and a-full.
Show trunking status. Aggregation must be up before the trunk will come up.	▶ Do the following command on each interface, 2/3 through 2/6 show port trunk 2/3 ▶ Should show the following: <ul style="list-style-type: none"> – Mode = negotiate – Encapsulation = dot1q – Status = Trunking – Native VLAN = 1 	show int Port-channel1 trunk ▶ Should show the following: <ul style="list-style-type: none"> – Mode = on – Encapsulation =802.1q – Status = Trunking – Native VLAN = 1
Review status of the aggregated link.	show port channel ▶ Should see the following: <ul style="list-style-type: none"> – Status = Connected – Channel Mode = on show channel traffic ▶ Should show how the links in the channel are being utilized.	show etherchannel summary ▶ Should show the following: <ul style="list-style-type: none"> – Protocol = - <ul style="list-style-type: none"> • (- = no protocol (static)). – Ports fa0/1 -4 = (P) <ul style="list-style-type: none"> • (P) = part of an aggregation group.
Ping the ESM. Where x.x.x.x is the IP address of the ESM (must be in same VLAN as subnet being pinged).	ping x.x.x.x ▶ Based on the Cat4K being at 192.168.70.202 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.	ping x.x.x.x ▶ Based on the Cat 3550 being at 192.168.70.200 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.

Description and Comments	CatOS based switch	IOS based switch
Attach a device to VLAN 5 and 10 on the Cisco switch and verify a ping to one of the blade servers in that same VLAN works as desired.	<ul style="list-style-type: none"> For VLAN 5, attach a device to port 2/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work) For VLAN 10, attach a device to port 2/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work). 	<ul style="list-style-type: none"> For VLAN 5, attach a device to port fa0/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work). For VLAN 10, attach a device to port fa0/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).

5.5.5 Single ESM, dual port LACP aggregation to two Cisco switches

This example (Figure 5-40) offers a compromise between performance and redundancy. It makes use of a single ESM with two ports aggregated to one Cisco switch, and the remaining two ports aggregated to a second Cisco switch. This configuration is suitable for those seeking a compromise between performance and redundancy. It still suffers from no protection in the event of an ESM failure.

Based on our specific configuration for this example (layer 2 network, Cisco-2 is root and port cost on both Cisco links toward the ESM set to 100) Spanning Tree will block the connection between Cisco-1 and the ESM, at the Cisco external switch until such time as the aggregation from the ESM to Cisco-2 (or Cisco-2 itself) goes down. The choice of root and port cost settings in this example was only made for this example, and may be a poor choice in a production network (certainly placing the root switch up against the ESM in a datacenter environment would not be very common). It is very important that any time an ESM is connected in a redundant fashion, the location of the root switch and various port costs in the network be accounted for to ensure correct data flow (need to prevent non-BladeCenter traffic from flowing through the ESM).




Figure 5-40 Single ESM with two 2 port LACP aggregated links to two different Cisco switches

Step 1: Take down the link(s)

It is always advisable to disable the link(s) prior to making any configuration changes. Please see Table 5-1 on page 84 for procedures.

Step 2: Configure the IBM side of the link

This section walks you through the sequence of actions required to configure the ESM for this example (reference Table 5-15 on page 135).

The following assumptions have been made for this example:

- ▶ The user is already logged in to the ESM with *root* level access.
- ▶ Commands are being performed in the sequence shown.
- ▶ The ESM is starting from a default config per the Example on page 85.
- ▶ Cisco switches being used are 10/100 based and we will be setting the ESM port to 100Mbps full duplex. This means a cross-over cable *must* be used for the link between the ESM and the Cisco switch.
 - If using switch ports that support 1000Mbps, the procedure to set speed will be slightly different from what is shown in the example.

Table 5-15 Configuring the ESM

Description and Comments	Actions to perform via Web interface to ESM
<p>Step 2.1- <i>Configure speed and duplex.</i> As already noted, it will be necessary to use cross-over cables on the links between the ESM and the Cisco switch, as tuning off auto-negotiation also turns off the auto selection of MDI/MDI-X. Note that LACP <i>must</i> have full duplex connections to operate correctly.</p>	<ul style="list-style-type: none"> ▶ Click the top connector (Ext1) in the picture of the ESM at the top of the page. ▶ Change Speed/Duplex to 100/Full. ▶ Click Apply. ▶ Repeat for Ext2, Ext3 and Ext4.
<p>Step 2.2 - <i>Configure PVIDs</i> This places the desired blade server ports in to the desired VLANs. If the VLAN does not exist, it will be created automatically.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> VLANs -> 802.1Q Port Settings -> Port VLAN ID (PVID). ▶ For Bay 2, change the PVID to 5. ▶ For Bays 3 and 4, change the PVID to 10. ▶ All other PVIDs should be set for 1. ▶ Click Apply.
<p>Step 2.3 - <i>Configure 802.1Q trunking</i> This will allow Ext1 through Ext4 to carry traffic for VLAN 5 and VLAN 10 (already carrying VLAN 1 as an initial default).</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> VLANs -> Edit 802.1Q VLANs ▶ Select VLAN 5 and click Edit. ▶ Set Interface Ext1, Ext2, Ext3 and Ext4 to Egress and make sure Tag box is checked for each. ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked. ▶ Click Apply. ▶ Click Configuration -> VLANs -> Edit 802.1Q VLANs. ▶ Select VLAN 10 and click Edit. ▶ Set interface Ext1, Ext2, Ext3 and Ext4 to Egress and make sure the Tag box is checked for each. ▶ Make sure Ports 3 and 4 are set for Egress and the Tag box is unchecked for these two ports. ▶ Click Apply.

Description and Comments	Actions to perform via Web interface to ESM
<p>Step 2.4- <i>Configure Link Aggregation</i>. This example makes use of LACP to dynamically negotiate link aggregation with the Cisco switch. Note that the Ext links used can not already be part of a different aggregation group, static or dynamic.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Setup IEEE 802.3ad Link Aggregation. ▶ Change LACP to Enabled and click Apply. ▶ Click Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Link Aggregation Port Settings. ▶ Set Configure External Port From to include the first two ports: <ul style="list-style-type: none"> – Ext1 to Ext2. <ul style="list-style-type: none"> • This will set these 2 ports in to a single link. ▶ Change Mode to Enabled and click Apply. ▶ Set Configure External Port From to include the second two ports: <ul style="list-style-type: none"> – Ext3 to Ext4. <ul style="list-style-type: none"> – Change the Admin Key to 2 <ul style="list-style-type: none"> • This will set these 2 ports in to their own link, separate from the Ext1/Ext2 link. ▶ Change Mode to Enabled and click Apply.
<p>Step 2.5 - <i>Save ESM config to NVRAM</i> Failure to perform this step will result in all changes to the ESM being lost if the @server BladeCenter is powered off or the ESM is otherwise restarted.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> Maintenance -> Save Changes. ▶ Click Save Configuration. ▶ Click Ok when complete.

Step 3: Configuring the Cisco switches

The following assumptions have been made for this example (reference Table 5-16):

- ▶ VLANs 1, 5 and 10 already exist in the Cisco configuration and will be part of the 802.1Q trunk to the ESM.
- ▶ The user is already logged in to the switch and the switch is in enable mode.
- ▶ The lowest available compatible ports are being used:
 - For the CatOS switch being used in this example, ports 2/3 and 2/4 are being used.
 - For the IOS switch being used in this example, ports fa0/1 and fa0/2 are being used.
- ▶ Commands are being performed in the sequence shown.
- ▶ The switch is starting from a default config per the Example on page 88 (IOS) or the Example on page 90 (CatOS).
- ▶ Cisco switch being used is 10/100 based and we will be setting the port to 100Mb full duplex.
 - If using switch ports that support 1000Mbps, the command to set speed will be different from what is shown in the example.

Table 5-16 Configuring the Cisco equipment

Description and Comments	CatOS based switch	IOS based switch
<p>Step 3.1 <i>Configure speed and duplex</i>. Need to perform on both Cisco-1 and Cisco-2 switch.</p>	<pre>set port speed 2/3-4 100 set port duplex 2/3-4 full</pre>	<pre>config t int range fa0/1 - 2 speed 100 duplex full</pre> <ul style="list-style-type: none"> ▶ Note that the <i>range</i> option on IOS is only available in newer revisions of code. If the range option is not available you will need to repeat steps 3.1 through 3.3 for each interface.

Description and Comments	CatOS based switch	IOS based switch
Step 3.2 <i>Configure 802.1Q trunking.</i> Need to perform on both Cisco-1 and Cisco-2 switch. Forces link to become an 802.1Q VLAN trunk.	set trunk 2/3 nonegotiate dot1q set trunk 2/4 nonegotiate dot1q	switchport trunk encapsulation dot1q switchport mode trunk switchport nonegotiate
Step 3.3 <i>Configure Spanning Tree port cost on individual ports.</i> Need to perform on both Cisco-1 and Cisco-2. Setting the port cost higher than default helps to prevent VLAN traffic between Cisco-1 and Cisco-2, other than VLAN 1, from being switched through the ESM. This results in a more optimal flow. For more information, review the section titled: "Guidelines and comments - Spanning Tree" on page 79	set spantree portcost 2/3-4 100 ► Note that for this particular design, leaving the links at their default cost will produce split traffic flows (see "Guidelines and comments - Spanning Tree" on page 79 for details on this behavior)	spanning-tree cost 100 ► Note that for this particular design, leaving the links at their default cost will produce split traffic flows (see "Guidelines and comments - Spanning Tree" on page 79 for details on this behavior)
Step 3.4 <i>Configure Link Aggregation.</i> Need to perform on both Cisco-1 and Cisco-2 switch. Note that if you do not have the option of setting the "channelprotocol" to "lACP" for CatOS, or the channel-group to "active" for IOS, then more then likely you will need to upgrade your code to support LACP.	set channelprotocol lACP 2 ► The trailing 2 in the above command tells CatOS to enable LACP for Module 2 set port lACP-channel 2/3-4 mode active	channel-group 1 mode active ► This will create a logical interface named <i>Port-Channel1</i> and place the interfaces fa0/1 and fa0/2 in to it. end
Step 3.5 <i>Setting port cost on the aggregated link.</i> Need to perform on both Cisco-1 and Cisco-2 switch.	(Does not apply) ► CatOS uses the values from the individual ports to meet this requirement.	int port-channel1 spanning-tree cost 100 ► See Step 3.3 for details. In this case, setting the cost for the whole aggregation as well as the individual ports.
Step 3.6 <i>Save config to NVRAM.</i> Only necessary on IOS based switches.	<i>(does not apply)</i>	write mem

Step 4: Reconnecting the devices

This is the final step to bring the connection in to full operation. This will be the reverse of whatever procedure was used in Step 1. See Table 5-2 on page 84 for details on how to reestablish the links.

Step 5: Verifying the configuration

This section provides options for verifying the correct and desired operation.

Certain elements of this verification depend on the existence of a layer 2 network behind the two Cisco switches. This was simulated for this example by tying the two switches together with an 802.1Q trunk link (see the configuration Example on page 88 for IOS based switches or the configuration Example on page 90 for CatOS based switches, for how this link was configured).

Verifying correct operation on the BladeCenter ESM

Verify the configurations on the ESM look as follows:

As in Figure 5-41, verify port state by clicking **Configuration -> Port Setting -> Display Ports**, and reviewing the status of Ext1 through Ext4 (State should be *Enabled*, Speed/Duplex should be *100/Full*, and Connection should show 100M/Full/802.3x for all 4 Ext ports).

IBM BladeCenter 4-Port Gb Ethernet Switch - Microsoft Internet Explorer

Address: <http://192.168.70.43/>

BladeCenter 4-Port Gb Ethernet Switch Module

Configuration

- Switch Information
- IP Setup
- Port Setting**
 - Configure Ports
 - Display Ports**
- Switch Setup
- VLANs
- Multicasting
- Port Mirroring Settings
- Spanning Tree
- Class of Service
- Link Aggregation
- MAC, Address, Forwarding
- Remote Management Setup
- Network Monitoring
- Maintenance

Display Ports

Port	State	Speed/Duplex	Connection	Port Type	Flow Control
Ext 1	Enabled	100M/Full	100M/Full/802.3x	10/100/1000-T	Enabled
Ext 2	Enabled	100M/Full	100M/Full/802.3x	10/100/1000-T	Enabled
Ext 3	Enabled	100M/Full	100M/Full/802.3x	10/100/1000-T	Enabled
Ext 4	Enabled	100M/Full	100M/Full/802.3x	10/100/1000-T	Enabled
Bay 1	Enabled	1000M/Full	1000M/Full/802.3x	Internal	Enabled
Bay 2	Enabled	1000M/Full	1000M/Full/802.3x	Internal	Enabled
Bay 3	Enabled	1000M/Full	1000M/Full/802.3x	Internal	Enabled
Bay 4	Enabled	1000M/Full	1000M/Full/802.3x	Internal	Enabled
Bay 5	Enabled	1000M/Full	Link Down	Internal	Enabled
Bay 6	Enabled	1000M/Full	Link Down	Internal	Enabled
Bay 7	Enabled	1000M/Full	Link Down	Internal	Enabled
Bay 8	Enabled	1000M/Full	Link Down	Internal	Enabled
Bay 9	Enabled	1000M/Full	Link Down	Internal	Enabled
Bay 10	Enabled	1000M/Full	Link Down	Internal	Enabled
Bay 11	Enabled	1000M/Full	Link Down	Internal	Enabled
Bay 12	Enabled	1000M/Full	Link Down	Internal	Enabled
Bay 13	Enabled	1000M/Full	Link Down	Internal	Enabled
Bay 14	Enabled	1000M/Full	Link Down	Internal	Enabled

Figure 5-41 Verify ports are operational

As in Figure 5-42 on page 139, verify VLAN 1 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 1** and clicking **Edit**. Make sure the ports not on VLAN 1 (Ports 2, 3 and 4) are set to None and their associated **Tag** box is checked. Also note that the options for Ext ports 2 and 4 are greyed out. Once an aggregation forms, you can only make changes to the lowest number port in the aggregation (in this case, Ext1 and Ext3), which will then be inherited by the other ports in the aggregation when the change is applied. Also of note, if the aggregation had not yet formed, (for example, if the cables were still not plugged in), then the other Ext ports would not be greyed out, and you could make changes to them. Of course as soon as the aggregation came up, those changes would be

overwritten by the lowest numbered port in the aggregation. The other settings could still be seen through the greyed out boxes, but they would not be getting used.

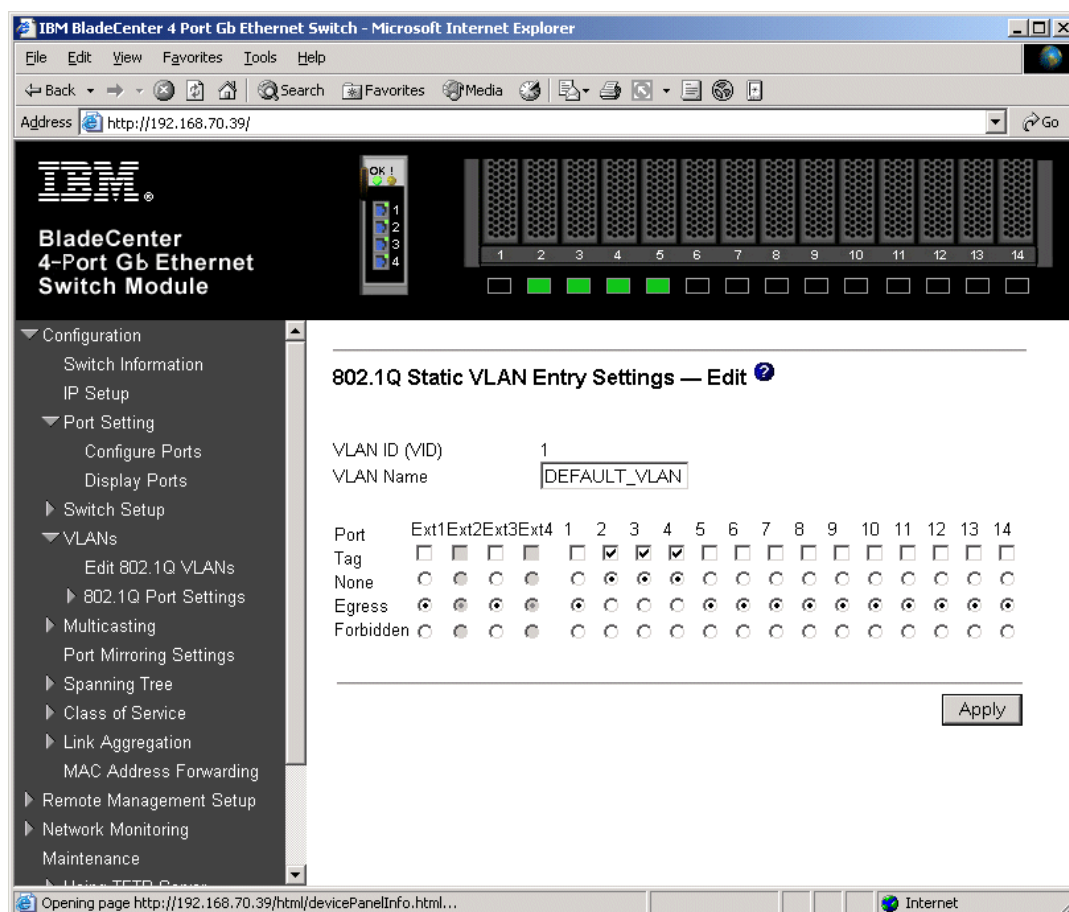


Figure 5-42 Verifying VLAN 1

As in Figure 5-43 on page 140, verify VLAN 5 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 5** and clicking **Edit**. Make sure **Port 2** is set to **Egress** and that the **Tag** box is not checked. Make sure the Ext1 and Ext3 interfaces are set for **Egress** (allows VLAN 5 traffic to pass through Ext1 and Ext3 and any aggregated link that is part of Ext1's and Ext3's aggregation) and that the box for **Tag** is checked. Also, note the same greying out of Ext2 and 4 will also be seen here, if the aggregation has already become active

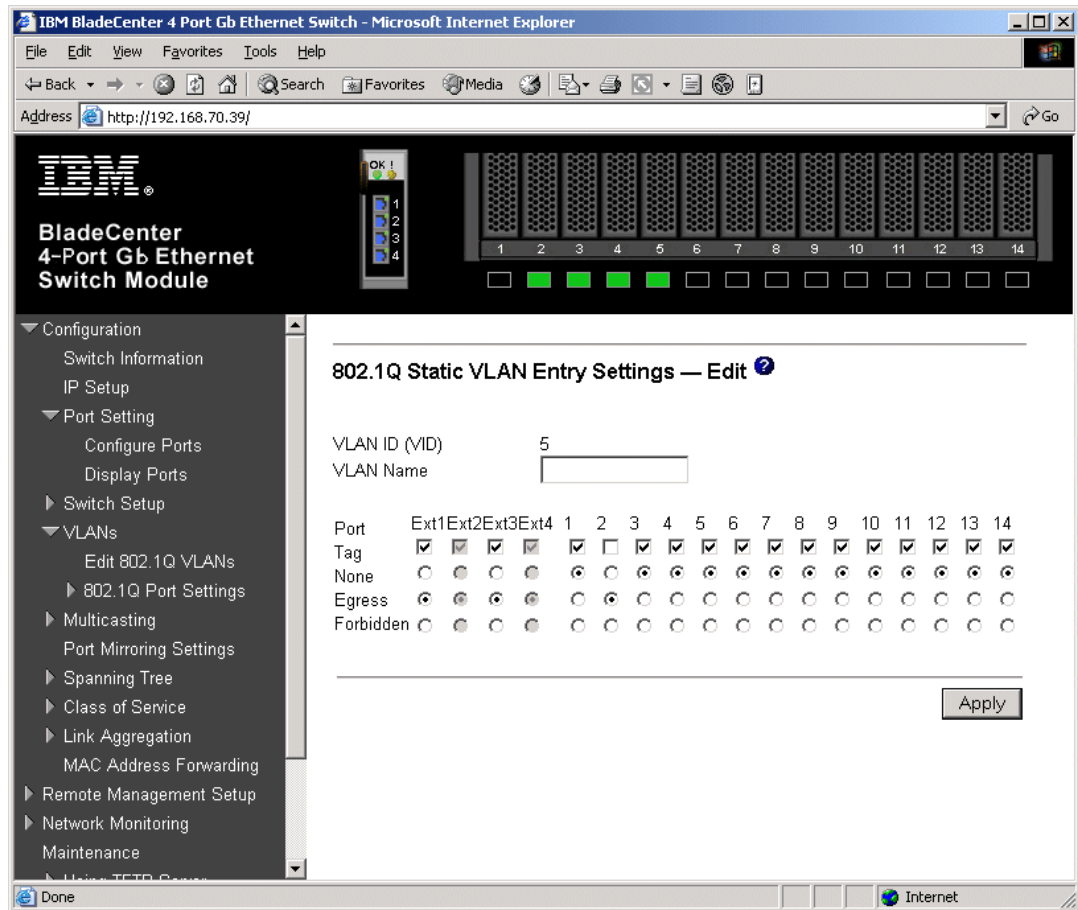


Figure 5-43 Verify VLAN 5

As in Figure 5-44 on page 141, verify VLAN 10 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 10** and clicking **Edit**. Make sure ports 3 and 4 are set to Egress and that the Tag box is not checked for either port. Make sure the Ext1 and Ext3 interfaces are set for Egress (allows VLAN 10 traffic to pass through Ext1 and Ext3 and any aggregated link that is part of Ext1's and Ext3's aggregation) and that the box for **Tag** is checked. Here again we see the greying out of Ext2 and Ext4 after an aggregation has formed.

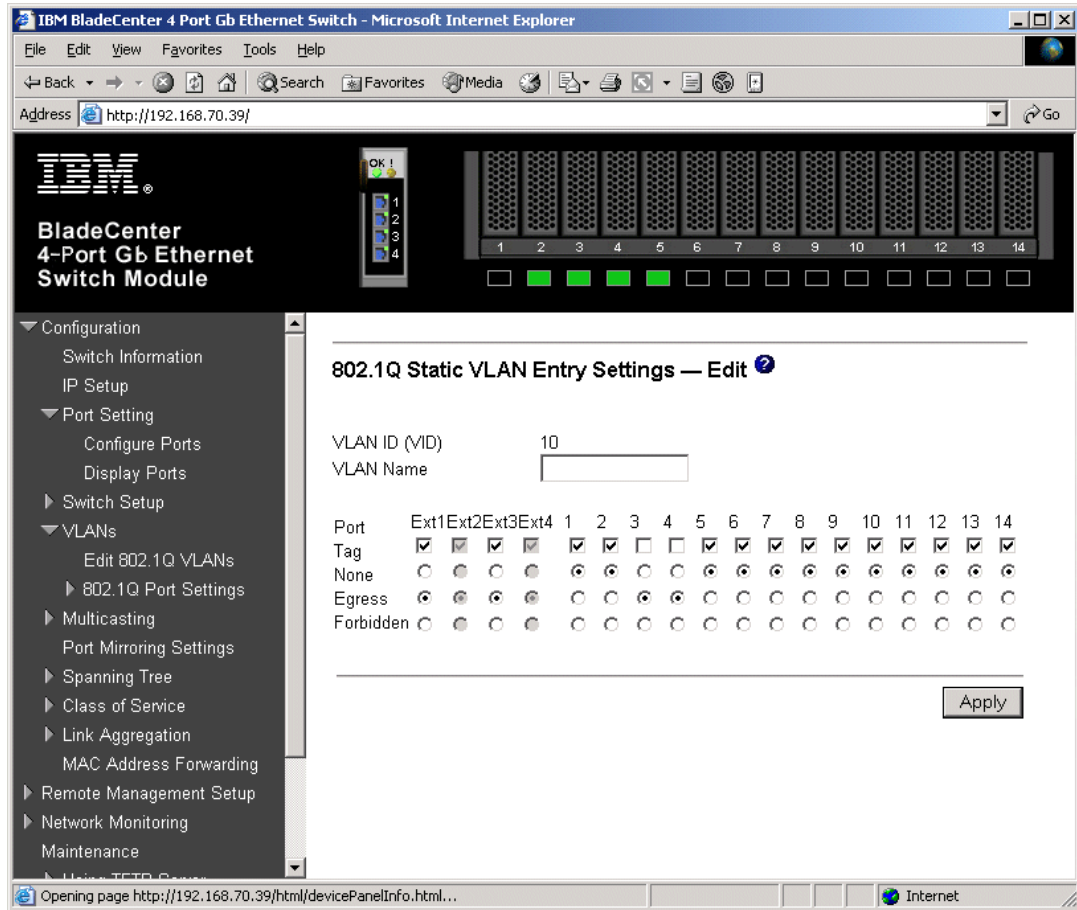


Figure 5-44 Verify VLAN 10

As in Figure 5-45 on page 142, verify Spanning Tree is forwarding traffic for the aggregation by clicking **Configuration -> Spanning Tree -> STP Port Settings**. Check the Status for Ext1 through Ext4. Based on the configuration of this example (root is Cisco-2 and certain ports on the Cisco external switch set for higher than default port cost), these ports should all show as Forwarding. May notice here that you lose the ability to individually manage Spanning Tree options for all ports in the aggregation except the lowest numbered port (and that the aggregation handles Spanning Tree as a whole, rather than on a port by port basis). In this case, changes to Ext 1 will be inherited by Ext 2 and changes to Ext 3 will be inherited by Ext 4. Again, this is only true if the aggregation has already formed.




Figure 5-45 Verifying STP is forwarding on all 4 links in the aggregation

As in Figure 5-46 on page 143, one way to review the status of the aggregation is by clicking **Configuration -> Network Monitoring -> Applications Status -> Link Aggregation**. From here we can see that Ext 1 and Ext 2 are on one aggregation and Ext3 and Ext4 are on a different aggregation.

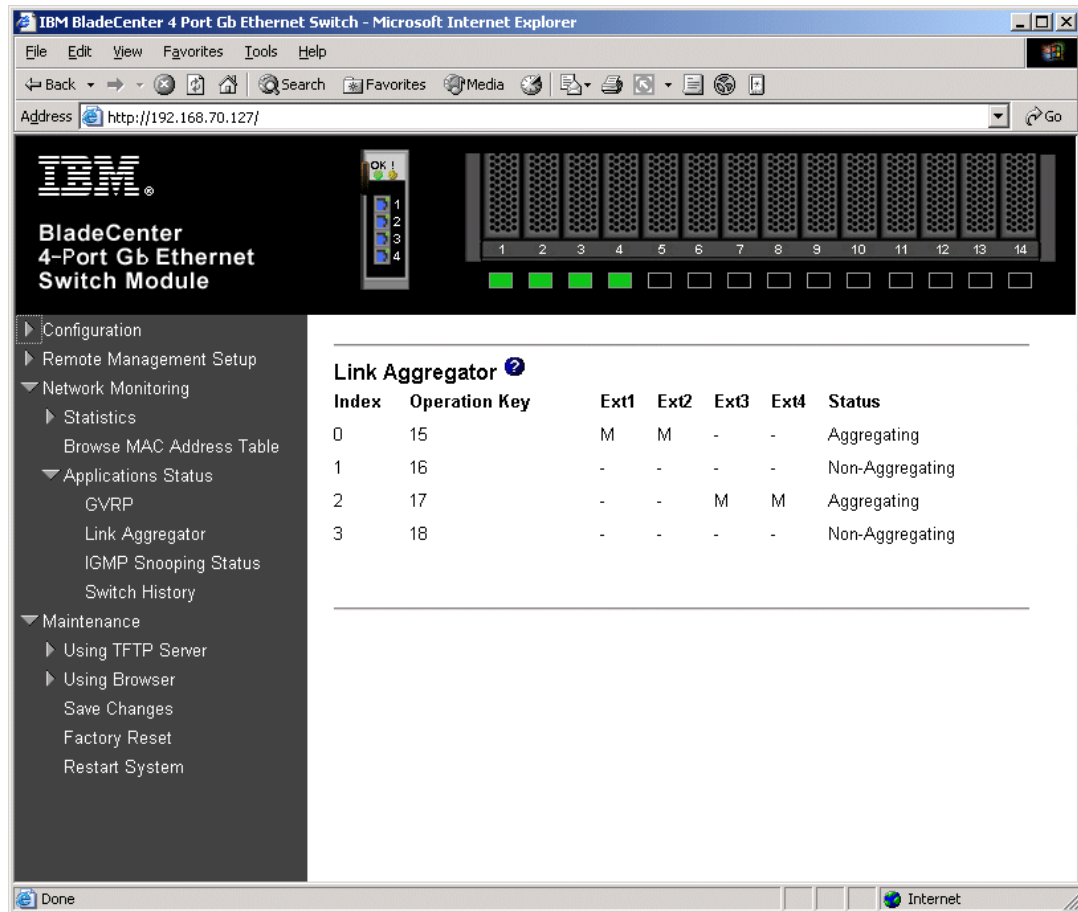


Figure 5-46 Link aggregation status

As in Figure 5-47 on page 144, another way to review the status of the aggregation can be done by clicking **Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Link Aggregation Port Setting** and noting that the Mode shows Enabled for all four ports (this is true whether the aggregation has formed or not), and the Status shows Active (shows Active after the aggregation has formed, shows Individual if the aggregation has not yet formed).

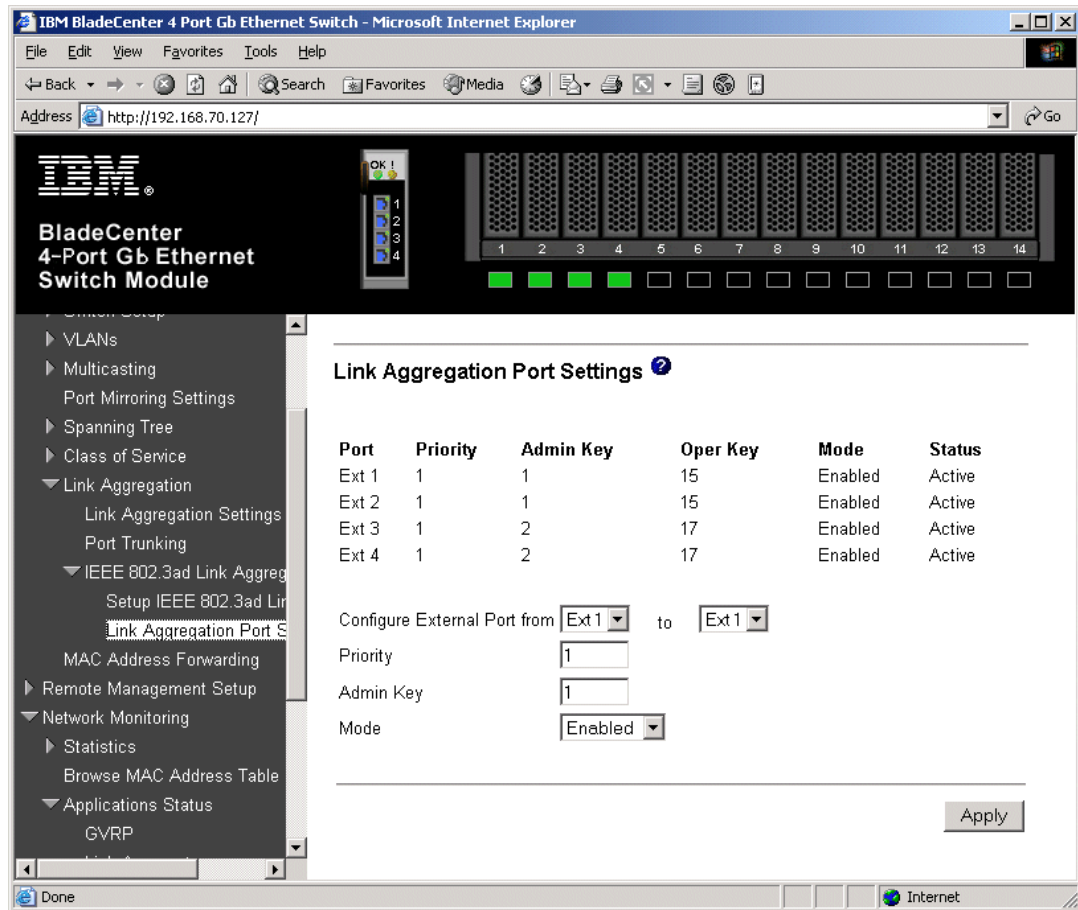


Figure 5-47 Link aggregation status, another view

As in Figure 5-48 on page 145, you can also verify the aggregation is operational by generating traffic across the link and clicking **Network Monitoring -> Statistics -> Port Utilization** and monitoring Tx/sec and Rx/sec for the four Ext ports. Depending on how much traffic is being generate, and from how many sources, and what aggregations may be blocking, the numbers will vary across the ports. For this example, 100,000, 1400 byte pings were being sent from the Cisco-2 switch to the IP address of the ESM. In this case, Ext3 is being used to transmit and Ext4 is receiving (these could have been on different ports, or even transmitting and receiving on the same port). If one were to pull the cable for Ext3, the traffic would switch over to a different Ext port on the same aggregation, usually with no loss of packets (usually under 1 second).

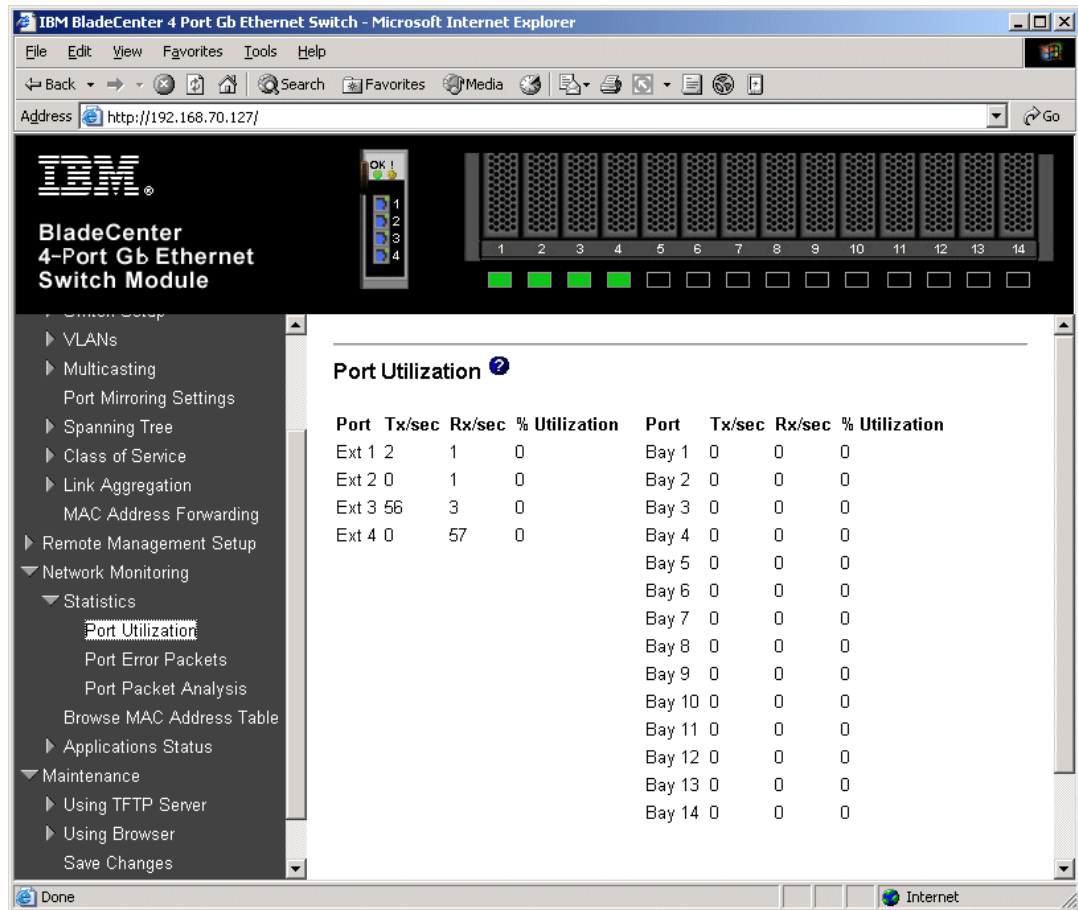


Figure 5-48 Example of aggregation load balancing

As in Figure 5-49 on page 146, with the same traffic flowing on the link, the cable to Ext3 is removed. The result is that the traffic previously carried on Ext3 is now on Ext4 (could have gone to any of the available Ext ports in the same aggregation).

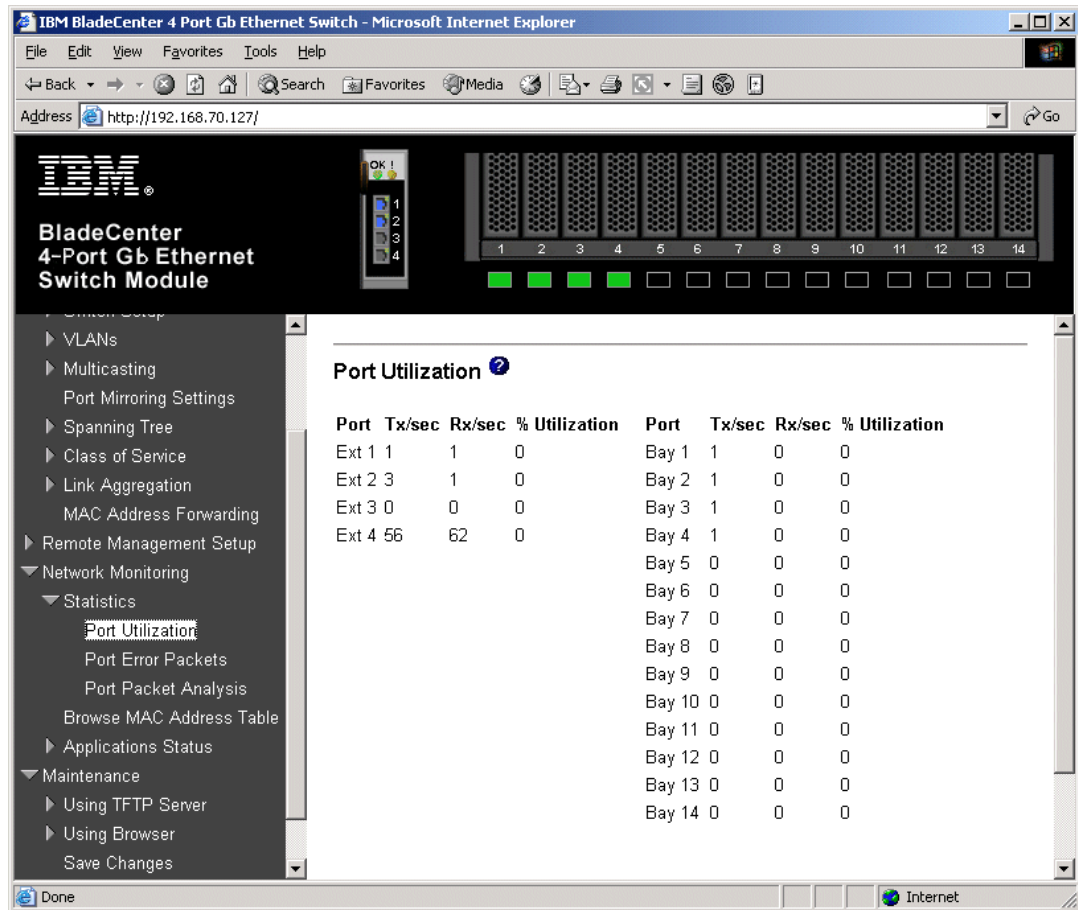


Figure 5-49 Traffic flow after cable to Ext3 removed

Verifying correct operation on the Cisco external switch

The following section includes some commands one can use to verify the desired configuration and operation of the Cisco equipment (reference Table 5-17 on page 147).

Table 5-17 Verifying the configuration and operation of the Cisco external switch of the connection

Description and Comments	CatOS based switch	IOS based switch
Review running config for desired statements. Perform on both Cisco-1 and Cisco-2.	show config <ul style="list-style-type: none"> ► Review for the following: <ul style="list-style-type: none"> – set channelprotocol lacp 2 – set port lacp-channel 2/3-6 120 <ul style="list-style-type: none"> • The number at the end may vary. • Even though we configured this to 2/3-4, it shows 2/3-6 here, this is normal. – set port speed 2/3-4 100 – set port duplex 2/3-4 full – set trunk 2/3 nonegotiate dot1q 1-1005,1025-4094 <ul style="list-style-type: none"> • Should see a similar entry for each configured port (2/3 and 2/4). – set port lacp-channel 2/3-6 mode active 	show run <ul style="list-style-type: none"> ► Review for the following on interface Port-channel1: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport nonegotiate ► Note that the values in Port-channel1 may not show up if the aggregation has never come up since first being configured. ► Review for the following on int fa0/1 and fa0/2: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport nonegotiate – duplex full – speed 100 – channel-group 1 mode active – spanning-tree cost 100
Show speed and duplex. Perform on both Cisco-1 and Cisco-2.	<ul style="list-style-type: none"> ► Do the following command on each interface, 2/3 and 2/4: show port status 2/3 <ul style="list-style-type: none"> ► Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN = Trunk – Duplex = Full – Speed = 100 	<ul style="list-style-type: none"> ► Do the following command on both interfaces, fa0/1 and fa0/2: show int fa0/1 status <ul style="list-style-type: none"> ► Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN = Trunk – Duplex = Full – Speed = 100 ► Note that if you do this on the Port-channel1 interface, the speed and duplex might show up as a-100 and a-full.
Show trunking status. Aggregation must be up before the trunk will come up. Perform on both Cisco-1 and Cisco-2.	show port trunk 2/3 <ul style="list-style-type: none"> ► Should show the following: <ul style="list-style-type: none"> – Mode = nonegotiate – Encapsulation = dot1q – Status = Trunking – Native VLAN = 1 	show int Port-channel1 trunk <ul style="list-style-type: none"> ► Should show the following: <ul style="list-style-type: none"> – Mode = on – Encapsulation = 802.1q – Status = Trunking – Native VLAN = 1
Review status of the aggregated links. Perform on both Cisco-1 and Cisco-2.	show port lacp-channel <ul style="list-style-type: none"> ► Should see both ports in Channel Mode Active show port lacp-channel statistics <ul style="list-style-type: none"> ► Run several times in a row, should show LACP Pkts Transmitted and Received climbing slowly (Transmitted usually higher than Received). 	show etherchannel summary <p>Should show the following:</p> <ul style="list-style-type: none"> ► Protocol = LACP ► Ports fa0/1 -2 = (P) <ul style="list-style-type: none"> – (P) = part of an aggregation group show lacp counters <ul style="list-style-type: none"> ► Run several times in a row, should show LACPDUs Sent and Recv climbing slowly (Sent usually higher than Recv). show etherchannel port-channel <ul style="list-style-type: none"> ► Should show the following: <ul style="list-style-type: none"> – Port state = Port-channel Ag-Inuse – Protocol = LACP – Both ports in EC State Active

Description and Comments	CatOS based switch	IOS based switch
<p>Show Spanning Tree status.</p> <p>Note that values shown here are very specific to the example configurations, and may not reflect production information.</p>	<p>show spantree 2/3 show spantree 2/4</p> <ul style="list-style-type: none"> ► Based on the configuration and cost setting in this example: <ul style="list-style-type: none"> – Cisco 1 switch should show all three VLANs Blocking for these two interfaces. – Cisco-2 switch (root) should show all three VLANs Forwarding for these two interfaces. – Interface 2/34 should also show forwarding for all three VLANs 	<p>show spanning int fa0/1 show spanning int fa0/1 show spanning int po1</p> <ul style="list-style-type: none"> ► Based on the configuration and cost settings in this example: <ul style="list-style-type: none"> – Cisco-1 switch should show BLK (Blocking) for all three VLANs, for each of these interfaces. – Cisco-2 switch (root) should show FWD (Forwarding) for all three VLANs, for each of these interfaces. – Interface fa0/24 should also show forwarding for all three VLANs
<p>Ping the ESM.</p> <p>Where x.x.x.x is the IP address of the ESM (must be in same VLAN as subnet being pinged).</p>	<p>ping x.x.x.x</p> <ul style="list-style-type: none"> ► Based on the Cat4K being at 192.168.70.202 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across. 	<p>ping x.x.x.x</p> <ul style="list-style-type: none"> ► Based on the Cat 3550 being at 192.168.70.200 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.
<p>Attach a device to VLAN 5 and 10 on the Cisco switch and verify a ping to one of the blade servers in that same VLAN works as desired.</p>	<ul style="list-style-type: none"> ► For VLAN 5, attach a device to port 2/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work) ► For VLAN 10, attach a device to port 2/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work). 	<ul style="list-style-type: none"> ► For VLAN 5, attach a device to port fa0/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work). ► For VLAN 10, attach a device to port fa0/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).
<p>Test redundancy.</p>	<ul style="list-style-type: none"> ► Experiment with shutting down the top two or bottom two ports for the ESM to force Spanning Tree to make the other path active. Verify with ping tests after Spanning Tree has stabilized. 	<ul style="list-style-type: none"> ► Experiment with shutting down the top two or bottom two ports for the ESM to force Spanning Tree to make the other path active. Verify with ping tests after Spanning Tree has stabilized.

5.5.6 Dual ESMs, each with a single link to the same Cisco switch

In this example (see Figure 5-50 on page 149), we discuss a basic configuration that includes dual ESMs, each with a single link to a single Cisco switch. This configuration offers minimal performance and limited redundancy and might be used for initial installation and testing of a dual BladeCenter ESM, or in an environment that does not require maximum performance. It does offer limited redundancy in the form of two ESMs but depends on the Operating Systems installed on the blade servers to perform redundancy in the event of an ESM failure.

One thing to note here is that the spanning tree is not an issue in this configuration, as each ESM only has one connection in to the layer 2 network.




Figure 5-50 Dual ESMs each with a single link to a single Cisco switch

For this example, reference the configuration steps as presented in “Single ESM, single link to a single Cisco switch” on page 91. The only difference is that you will have to configure both ESM’s instead of just one (follow the same procedure for both), and on the Cisco external switch, you will configure a second port (port 2/4 or fa0/2, depending if the switch is running CatOS or IOS) (also follow the same procedure on both ports in the Cisco switch).

The verification procedure will be the same on the ESM side, except that you will need to log in to both ESMs and verify the desired configuration on each.

For the Cisco external switch of the connection, you can follow the same verification procedure as well, modified slightly to review each port instead of just one port.

The results should be the same for both the ESM and the Cisco switch as they were in 5.5.1, “Single ESM, single link to a single Cisco switch” on page 91.

One item to repeat here is that there will be no Spanning Tree blocking in this configuration, as each of the ESMs is a completely standalone access switch, and do not have any internal connections to one another (except through the blade servers). It is theoretically possible to obtain and install software onto a blade server that would make it (the blade server) operate as a two port switch, in which case Spanning Tree may come in to play (as now the ESM’s are being tied together both inside the @server BladeCenter and external to the @server BladeCenter). The discussion of this interaction is beyond the scope of this document.

5.5.7 Dual ESMs with a single link to two different Cisco switches

The example (reference Figure 5-51 on page 150) listed in this section is almost identical to the one in section 5.5.1 “Single ESM, single link to a single Cisco switch” on page 91, with the exception of having two ESM’s and two external switches instead of one. It offers minimal performance and some redundancy. Possible uses for this configuration might be the initial installation and testing of a @server BladeCenter deployment or environments that do not require maximum performance, but do require redundancy. Note that no ports will be in a

Spanning Tree blocking state in this configuration as each ESM only has a single connection in to the layer 2 network.




Figure 5-51 Dual ESMs each with a single link to two separate Cisco switches

Reference the configuration example in 5.5.1, “Single ESM, single link to a single Cisco switch” on page 91. For both the ESM and the Cisco switches, the configuration will be the same, you will just have to perform the set of operations twice, once for ESM-1 and Cisco-1 and once for the ESM-2 and Cisco-2.

The verification procedure will be the same on the ESM side, except you’ll need to log in to each ESM and verify the desired configuration.

For the Cisco external switch of the connection, you can follow the same verification procedure as well, modified slightly to review each port instead of just one port.

The results should be the same for both the ESM and the Cisco switch as they were in 5.5.1, “Single ESM, single link to a single Cisco switch” on page 91.

As already noted, like the last example, there will be no Spanning Tree blocking in this configuration, as each of the ESMs is a completely standalone access switch, and do not have any internal connections to one another (except through the blade servers). It is theoretically possible to obtain and install software onto a blade server that would make it (the blade server) operate as a two port switch, in which case Spanning Tree may come in to play (as now the ESM’s are being tied together both inside the @server BladeCenter and external to the @server BladeCenter). The discussion of this interaction is beyond the scope of this document.

5.5.8 Dual ESMs each with one link to separate Cisco switches

This example (reference Figure 5-52 on page 151) begins to show some high redundancy, but still with limited performance. It contains dual ESMs each cross connected to a pair of Cisco switches. Possible uses for this configuration are in environments that are not concerned with performance but require high availability for their network connections.

Based on our specific configuration for this example (layer 2 network, Cisco-2 is root and port cost on both Cisco links toward each ESM is set to 100) Spanning Tree will normally block the connection between Cisco-1 and the ESMs, at the Cisco external switch. The choice of root

and port cost settings in this example were only made for this example, and may be a poor choice in a production network (certainly placing the root switch up against the ESM in a datacenter environment would not be very common). It is very important that any time an ESM is connected in a redundant fashion, that the location of the root switch and various port costs in the network be accounted for to ensure correct data flow (need to prevent non-BladeCenter traffic from flowing through the ESM).




Figure 5-52 Dual ESMs with cross connected single links to two different Cisco switches

Configuration and verification for this example is similar to 5.5.2, “Single ESM, single link to two Cisco switches” on page 98, but contains enough differences that a whole new procedure is presented here.

Step 1: Take down the link(s)

It is always advisable to disable the link(s) prior to making any configuration changes. Please see Table 5-1 on page 84 for procedures.

Step 2: Configure the IBM side of the link

This section walks you through the sequence of actions required to configure the ESM for this example (reference Table 5-18 on page 152).

The following assumptions have been made for this example:

- ▶ The user is already logged in to the ESMs with *root* level access.
- ▶ The following cabling was used (see Figure 5-53 on page 152) for a diagram:
 - ESM-1, Ext1 goes to Cisco-1 fa0/1 (or 2/3 if CatOS).
 - ESM-1, Ext2 goes to Cisco-2 fa0/1 (or 2/3 if CatOS).
 - ESM-2, Ext1 goes to Cisco-2 fa0/2 (or 2/4 if CatOS).
 - ESM-2, Ext2 goes to Cisco-1 fa0/2 (or 2/4 if CatOS)

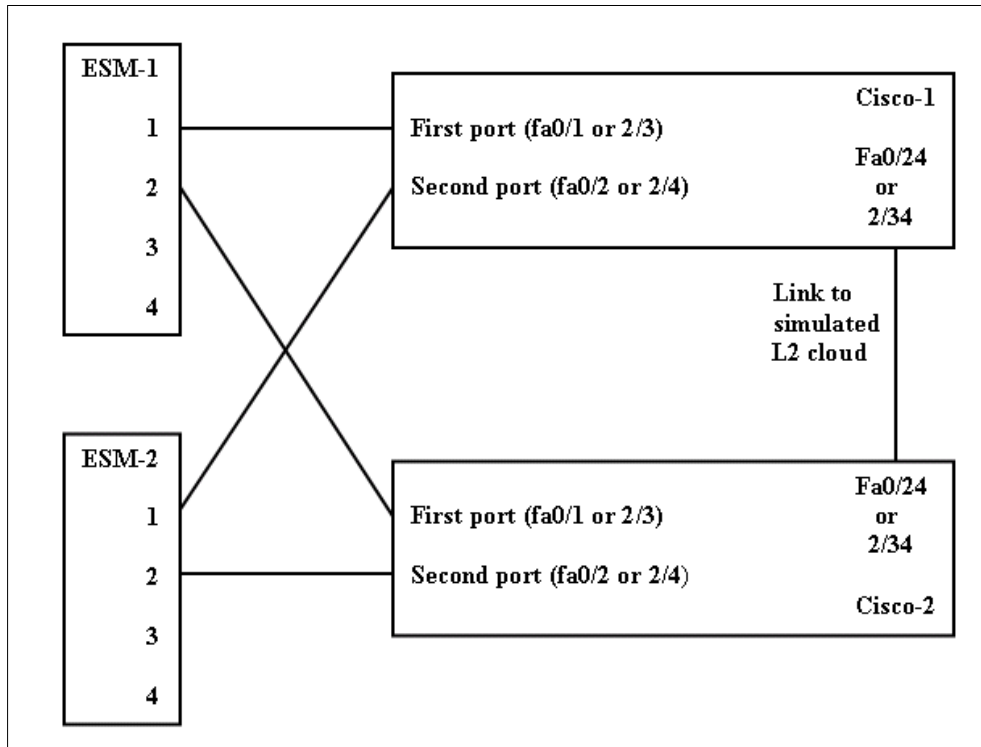


Figure 5-53 Cabling layout for Example 5-52 on page 151

- ▶ Commands are being performed in the sequence shown.
- ▶ The ESM is starting from a default config per the Example on page 85.
- ▶ Cisco switches being used are 10/100 based and we will be setting the ESM port to 100Mbps full duplex. This means a cross-over cable *must* be used for the link between the ESM and the Cisco switch.
 - If using switch ports that support 1000Mbps, the procedure to set speed will be slightly different from what is shown in the example.

Table 5-18 Configuring the ESMs

Description and Comments	Actions to perform via Web interface to ESM
Step 2.1- Configure speed and duplex. As already noted, it will be necessary to use a cross-over cable on the link between the ESM and the Cisco switch, as tuning off auto-negotiation also turns off the auto selection of MDI/MDI-X.	On ESM-1: <ul style="list-style-type: none"> ▶ Click the top connector (Ext1) in the picture of the ESM at the top of the page. ▶ Change Speed/Duplex to 100/Full. ▶ Click Apply. ▶ Repeat for Ext2 interface. Repeat process for ESM-2
Step 2.2 - Configure PVIDs This places the desired blade server ports in to the desired VLANs. If the VLAN does not exist, it will be created automatically.	On ESM-1: <ul style="list-style-type: none"> ▶ Click Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID). ▶ For Bay 2, change the PVID to 5. ▶ For Bays 3 and 4, change the PVID to 10. ▶ All other PVIDs should be set for 1. ▶ Click Apply. Repeat process for ESM-2

Description and Comments	Actions to perform via Web interface to ESM
<p>Step 2.3 - <i>Configure 802.1Q trunking</i> This will allow Ext1 and Ext2 to carry traffic for VLAN 5 and VLAN 10 (already carrying VLAN 1 as an initial default).</p>	<p>On ESM-1:</p> <ul style="list-style-type: none"> ▶ Click Configuration->VLANs->Edit 802.1Q VLANs. ▶ Select VLAN 5 and click Edit. ▶ Set Interface Ext1 to Egress and make sure Tag box is checked. ▶ Set Interface Ext2 to Egress and make sure Tag box is checked. ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked. ▶ Click Apply. ▶ Click Configuration->VLANs->Edit 802.1Q VLANs. ▶ Select VLAN 10 and click Edit. ▶ Set interface Ext1 to Egress and make sure Tag box is checked. ▶ Set Interface Ext2 to Egress and make sure Tag box is checked. ▶ Make sure Ports 3 and 4 are set for Egress and the Tag box is unchecked for these two ports. ▶ Click Apply. <p>Repeat process for ESM-2</p>
<p>Step 2.4 - <i>Save ESM config to NVRAM</i> Failure to perform this step will result in all changes to the ESM being lost if the @server BladeCenter is powered off or the ESM is otherwise restarted.</p>	<p>On ESM-1:</p> <ul style="list-style-type: none"> ▶ Click Configuration -> Maintenance -> Save Changes. ▶ Click Save Configuration. ▶ Click Ok when complete. <p>Repeat process for ESM-2</p>

Step 3: Configuring the Cisco switches

The following assumptions have been made for this example (reference Table 5-19):

- ▶ VLANs 1, 5 and 10 already exist in the Cisco configurations and will be part of the 802.1Q trunk to the ESM.
- ▶ The user is already logged in to the switches and the switches are in enable mode.
- ▶ The port connections are as described in Step 2 of this example.
- ▶ Commands are being performed in the sequence shown.
- ▶ The switches are starting from a default config per the Example on page 88 (IOS) or the Example on page 90 (CatOS).
- ▶ Cisco switches being used are 10/100 based and we will be setting the ports to 100Mb full duplex.
 - If using switch ports that support 1000Mbps, the command to set speed will be different from what is shown in the example.

Table 5-19 Configuring the Cisco equipment

Description and Comments	CatOS based switch	IOS based switch
<p>Step 3.1 - <i>Configure speed and duplex.</i></p>	<ul style="list-style-type: none"> ▶ Perform the following two commands on both Cisco-1 and Cisco-2: set port speed 2/3-4 100 set port duplex 2/3-4 full 	<ul style="list-style-type: none"> ▶ Perform the following four commands on both Cisco-1 and Cisco-2: config t int range fa0/1 - 2 speed 100 duplex full
<p>Step 3.2 - <i>Configure 802.1Q trunking.</i> Forces link to become an 802.1Q VLAN trunk.</p>	<ul style="list-style-type: none"> ▶ Perform the following two commands on both Cisco-1 and Cisco-2: set trunk 2/3 negotiate dot1q set trunk 2/4 negotiate dot1q 	<ul style="list-style-type: none"> ▶ Perform the following three commands on both Cisco-1 and Cisco-2: switchport trunk encap dot1q switchport mode trunk switchport nonegotiate

Description and Comments	CatOS based switch	IOS based switch
<p>Step 3.3 <i>Configure Spanning Tree port cost.</i> Setting the port cost higher than default performs two operations for this example:</p> <ul style="list-style-type: none"> ► It helps to ensure a more balanced default flow (default flow when all links are operating normally). While not necessarily important for this example, in a production network, it makes sense to force conditions so that you know which ports will be blocking. ► It helps to prevent VLAN traffic between Cisco-1 and Cisco-2, other than VLAN 1, from being switched through the ESM. This forcing of the port cost helps to promote a more optimal flow. <p>For more information, review the section titled: "Guidelines and comments - Spanning Tree" on page 79</p>	<ul style="list-style-type: none"> ► Set the port cost on Cisco-1 and Cisco-2 to control default flow: set spantree portcost 2/3-4 100 ► Note that for this particular design, the port costs on all links is 19. In this situation, normal traffic flow from Cisco-1 to Cisco-2 would all be through 2/34. The port cost being set here is to get in to the habit of forcing port cost when using redundant connections from an ESM in to the production network. 	<ul style="list-style-type: none"> ► Set the port cost on Cisco-1 and Cisco-2 to control default flow: spanning-tree cost 100 end ► Note that for this particular design, the port costs on all links is 19. In this situation, normal traffic flow from Cisco-1 to Cisco-2 would all be through fa0/24. The port cost being set here is to get in to the habit of forcing port cost when using redundant connections from an ESM in to the production network.
<p>Step 3.4 <i>Save config to NVRAM.</i> Only necessary on IOS based switches.</p>	<i>(does not apply)</i>	write mem

Step 4: Reconnecting the devices

This is the final step to bring the connection in to full operation. This will be the reverse of whatever procedure was used in Step 1. See Table 5-2 on page 84 for details on how to reestablish the links.

Step 5: Verifying the configuration

This section provides options for verifying the correct and desired operation.

Certain elements of this verification depend on the existence of a layer 2 network behind the two Cisco switches. This was simulated for this example by tying the two switches together with an 802.1Q trunk link (see the configuration Example on page 88 for IOS based switches or the configuration Example on page 90 for CatOS based switches, for how this link was configured).

Verifying correct operation on the BladeCenter ESM

Verify the configurations on the ESM look as follows:

As in Figure 5-54 on page 155, verify port state by clicking **Configuration -> Port Setting -> Display Ports**, and reviewing the status of Ext1 and Ext2 (State should be *Enabled*, Speed/Duplex should be *100/Full*, and Connection should show *100M/Full/802.3x*). Perform on both ESM-1 and ESM-2.

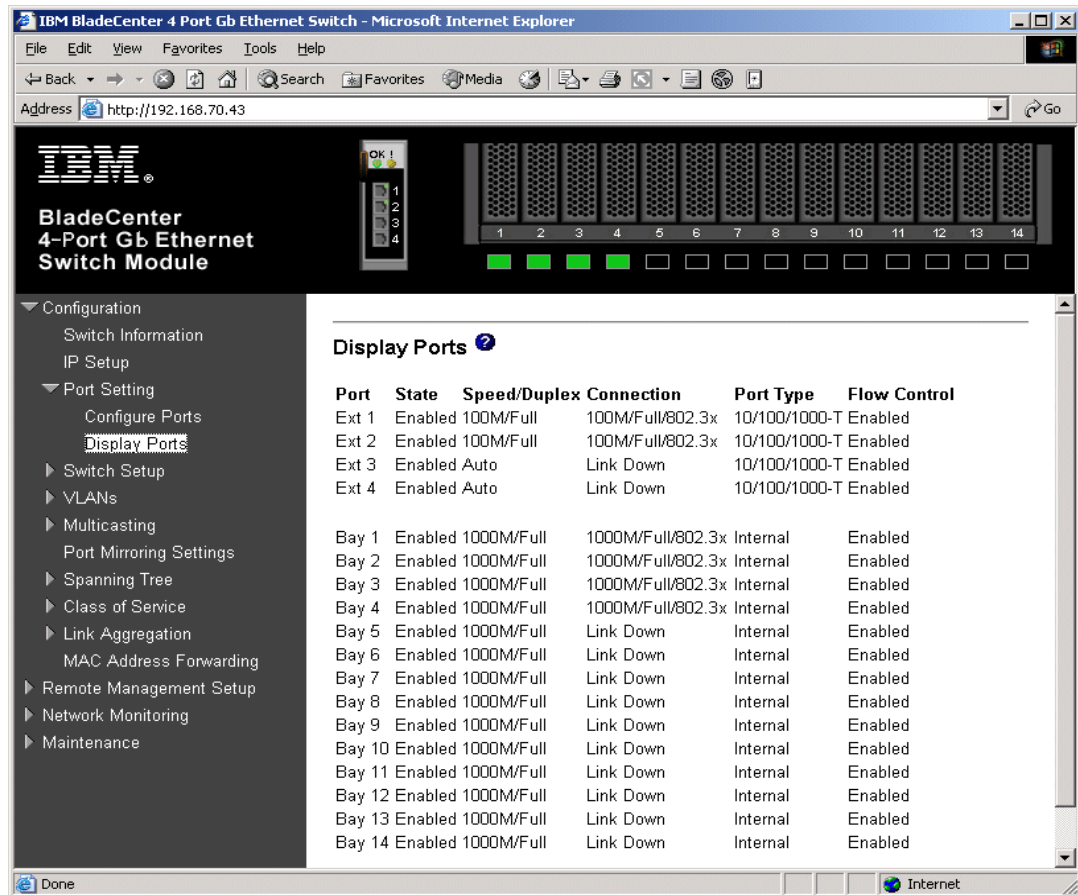


Figure 5-54 Verify ports are operational

As in Figure 5-55 on page 156, verify VLAN 1 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting VLAN 1 and clicking **Edit**. Make sure the ports not on VLAN 1 (Ports 2, 3 and 4) are set to None and their associated **Tag** box is checked. Perform on both ESM-1 and ESM-2.

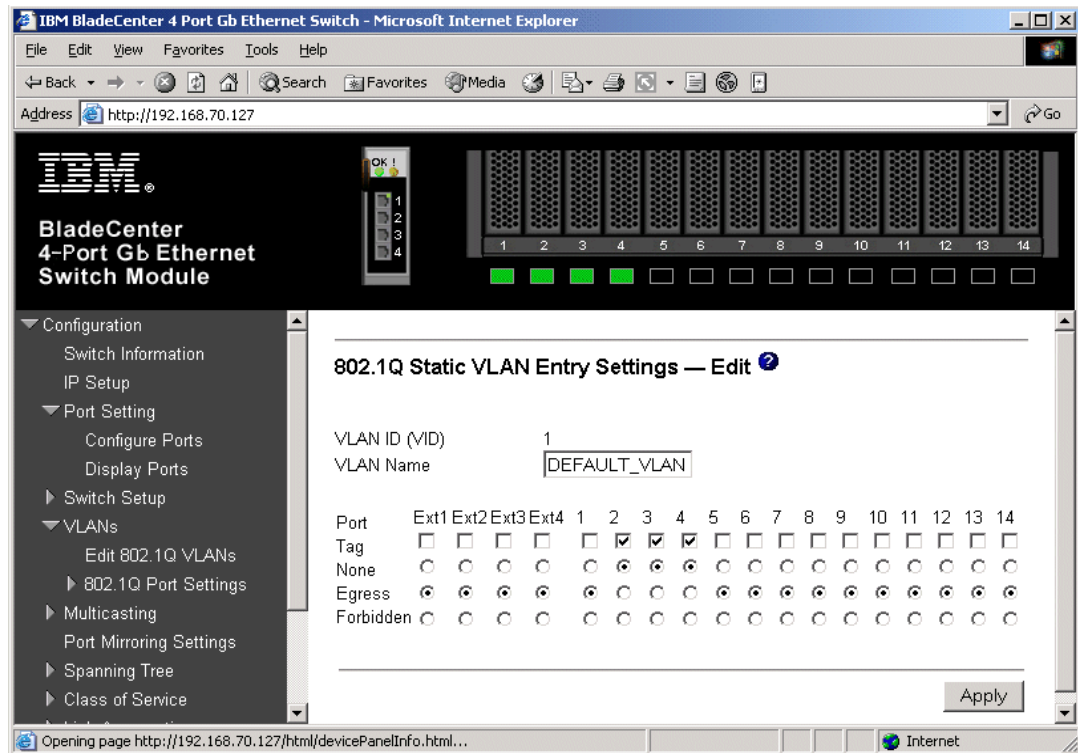


Figure 5-55 Verifying VLAN 1

As in Figure 5-56 on page 157, verify VLAN 5 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting VLAN 5 and clicking **Edit**. Make sure port 2 is set to Egress and that the Tag box is not checked. Make sure the Ext1 and Ext2 interfaces are set for Egress (allows VLAN 5 traffic to pass through these Ext interfaces) and that their associated **Tag** boxes are checked. Perform on both ESM-1 and ESM-2.

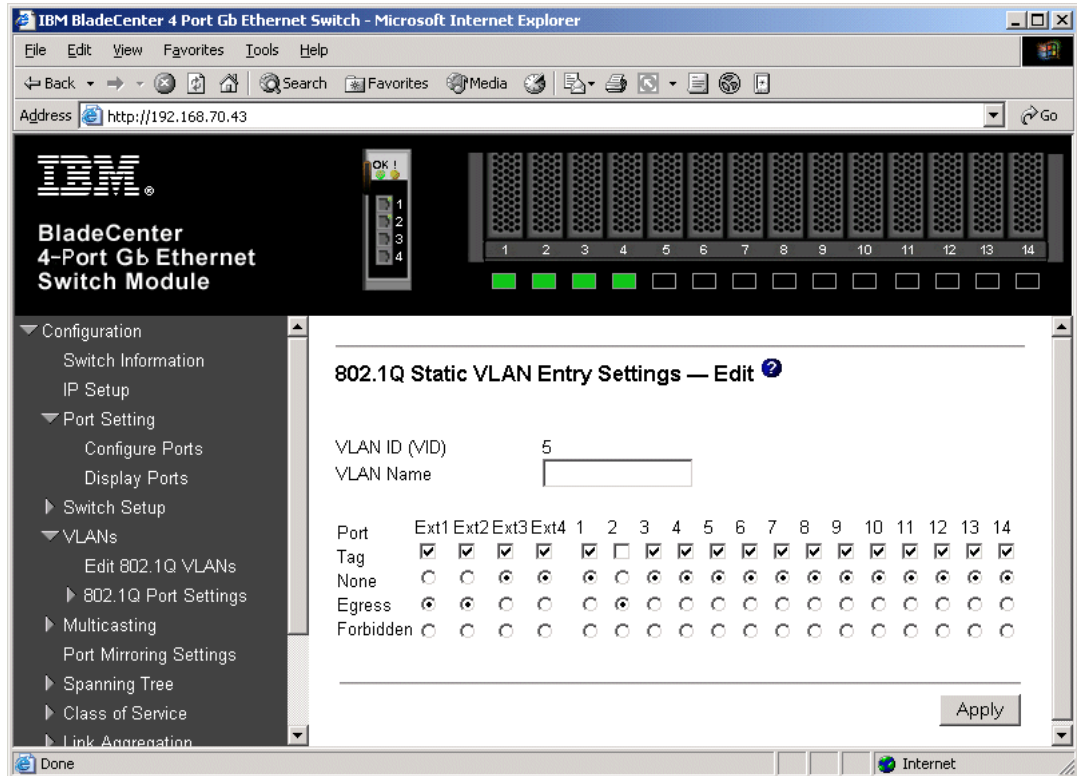


Figure 5-56 Verify VLAN 5

As in Figure 5-57 on page 158, verify VLAN 10 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting VLAN 10 and clicking **Edit**. Make sure ports 3 and 4 are set to Egress and that the Tag box is not checked for either port. Make sure the Ext1 and Ext2 interfaces are set for Egress (allows VLAN 10 traffic to pass through Ext interfaces) and that their associated **Tag** boxes are checked. Perform on both ESM-1 and ESM-2.

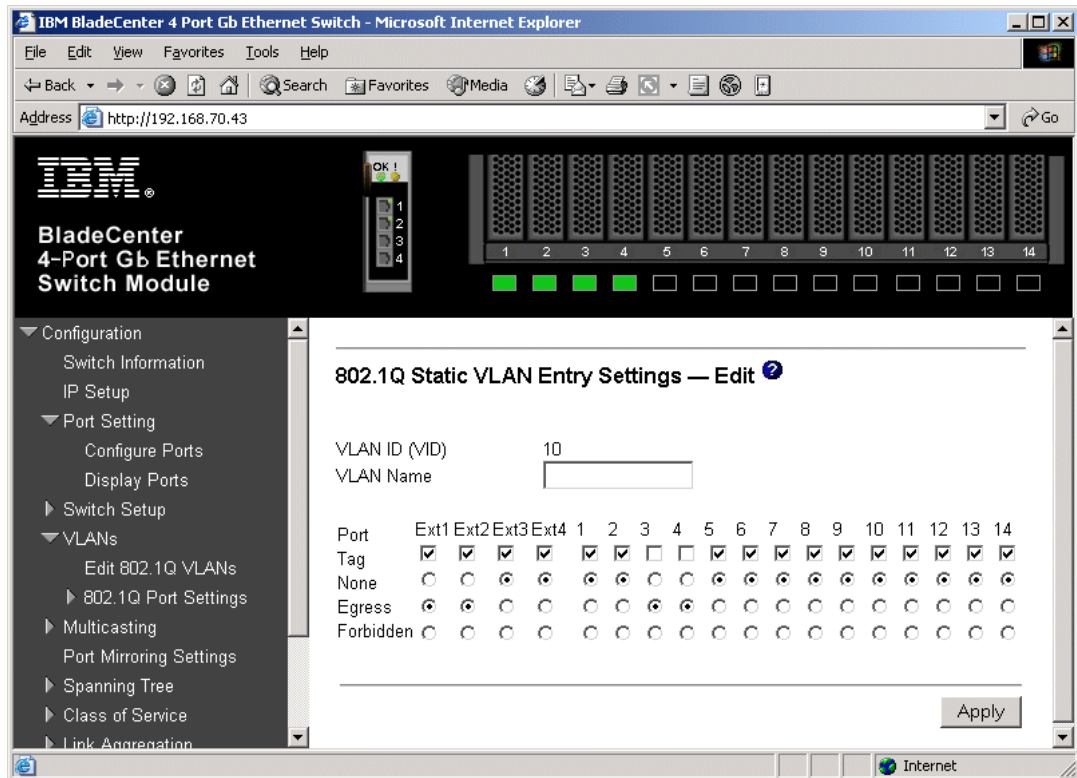


Figure 5-57 Verify VLAN 10

As in Figure 5-58 on page 159, verify Spanning Tree is operational by clicking **Configuration -> Spanning Tree -> STP Port Settings**. Based on the configurations used in this example (Cisco-2 is root and Cisco links to ESM set for higher than default port cost) both Ext1 and Ext2 should show as forwarding (the links on the Cisco external switch of Cisco-1 to each ESM should show blocking). Should see similar results on both ESM-1 and ESM-2.

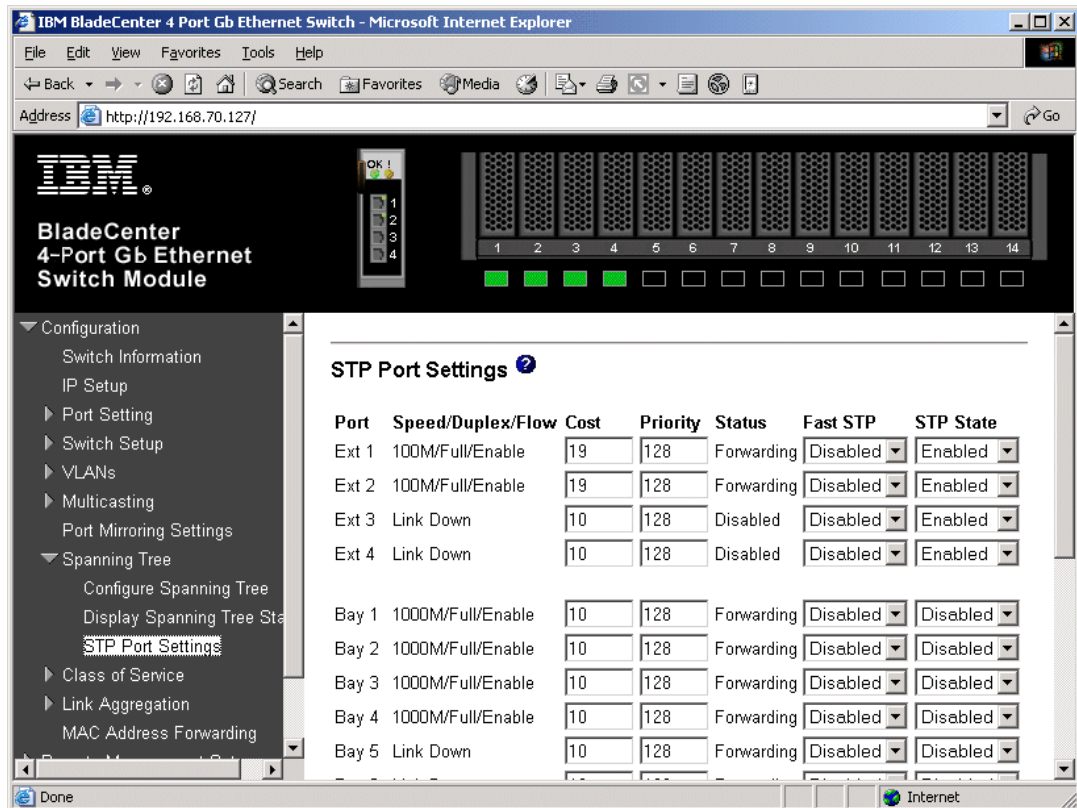


Figure 5-58 Verifying STP is forwarding to Ext1

Verifying correct operation on the Cisco external switch

The following section includes some commands one can use to verify the desired configuration and operation of the Cisco equipment (reference Table 5-20).

Table 5-20 Verifying the configuration

Description and Comments	CatOS based switch	IOS based switch
Review running config for desired statements. Perform on both Cisco-1 and Cisco-2.	show config ► Review both Cisco-1 and Cisco -2 for the following: <ul style="list-style-type: none"> – set port speed 2/3-4 100 – set port duplex 2/3-4 full – set trunk 2/3 nonegotiate dot1q 1-1005,1025-4094 – set trunk 2/4 nonegotiate dot1q 1-1005,1025-4094 	show run ► Review both Cisco-1 and Cisco-2 for the following on int fa0/1 and fa0/2: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport nonegotiate – duplex full – speed 100
Show speed and duplex. Perform on both Cisco-1 and Cisco-2.	show port status 2/3 show port status 2/4 ► Should show the following for both ports: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100 	show int fa0/1 status show int fa0/2 status ► Should show the following for both ports: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100

Description and Comments	CatOS based switch	IOS based switch
Show trunking status. Link must be up before a trunk will come up. Perform on both Cisco-1 and Cisco-2.	show port trunk 2/3 show port trunk 2/4 ► Should show the following for both ports: – Mode = nonegotiate – Encapsulation = dot1q – Status = Trunking – Native VLAN = 1	show int fa0/1 trunk show int fa0/2 trunk ► Should show the following both ports: – Mode = on – Encapsulation =802.1q – Status = Trunking – Native VLAN = 1
Show Spanning Tree status. Perform on both Cisco-1 and Cisco-2.	show spantree 2/3 show spantree 2/4 ► Based on our example config, cross connected links should be in a blocking state, in either the ESM or the Cisco switch.	show spanning int fa0/1 show spanning int fa0/2 ► Based on our example config, cross connected links should be in a blocking state, in either the ESM or the Cisco switch.
Show Spanning Tree status. In this configuration, both switches should show their ESM facing interfaces in a forwarding state.	show spantree 2/3 show spantree 2/4 ► Perform on both Cisco-1 and Cisco-2. – State should show blocked for the links between Cisco-1 and the ESMs for all VLANs. – Should show forwarding for all connections(2/3, 2/4 and 2/34) on Cisco-2.	show spanning int fa0/1 show spanning int fa0/2 ► Perform on both Cisco-1 and Cisco-2. – Sts should show blocked for the links between Cisco-1 and the ESMs for all VLANs. – Should show forwarding for all connections (fa0/1, fa0/2 and fa0/24) on Cisco-2.
Ping the ESM. Where x.x.x.x is the IP address of the ESM (must be in same VLAN as subnet being pinged).	ping x.x.x.x ► Based on the Cat4K being at 192.168.70.201 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.	ping x.x.x.x ► Based on the Cat 3550 being at 192.168.70.201 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.
Attach a device to VLAN 5 and 10 on the Cisco switch and verify a ping to one of the blade servers in that same VLAN, works as desired	► For VLAN 5, attach a device to port 2/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work). ► For VLAN 10, attach a device to port 2/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).	► For VLAN 5, attach a device to port fa0/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work). ► For VLAN 10, attach a device to port fa0/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).
Test redundancy.	► Experiment with redundancy by taking down Cisco-2 (which would normally be carrying all traffic) and make sure Cisco-1 takes over. Verify with ping tests after Spanning Tree has stabilized. Should be able to still ping both ESMs from Cisco-1.	► Experiment with redundancy by taking down Cisco-2 (which would normally be carrying all traffic) and make sure Cisco-1 takes over. Verify with ping tests after Spanning Tree has stabilized. Should be able to still ping both ESMs from Cisco-1.

5.5.9 Dual ESMs with four port LACP aggregation to different Cisco switches

This example (reference Figure 5-59 on page 161) offers the maximum performance available from the @server BladeCenter, as well as fairly high redundancy, depending on the configuration of the Operating Systems running on the blade servers within the @server BladeCenter. It makes use of two ESMs, each with all 4 ports aggregated in to a single link,

and each going to a separate Cisco switch. No ports will be in a Spanning Tree blocking state with this configuration as each ESM only has a single (albeit aggregated) connection in to the layer 2 network).




Figure 5-59 Dual ESM's each with 4 port LACP aggregation to two separate Cisco switches

This example will use the same configuration and verification procedures as Example 5.5.3 on page 107. As with the other dual ESM examples, you will need to run through the configurations and verifications twice, once of the ESM-1/Cisco-1 pair, and once for the ESM-2/Cisco-2 pair.

5.5.10 Dual ESMs with two port LACP aggregation to two Cisco switches

This example (reference Figure 5-60 on page 162) offers a good compromise between performance and high availability. It is made up of dual ESMs, each with two, two port aggregated links, going to two separate Cisco switches. The 2 port aggregation provides for higher performance than a single link, and the second 2 port aggregation provides for full redundancy on any link or switch failures. Possible uses might be high availability and performance environments.

Based on our specific configuration for this example (layer 2 network, Cisco-2 is root and port cost on both Cisco links/aggregations toward the ESM set to 100) Spanning Tree will block the connection between Cisco-1 and the ESM, at the Cisco external switch. The choices of root and port cost settings in this example were only made for this example, and may be a poor choice in a production network (certainly placing the root switch up against the ESM in a datacenter environment would not be very common). It is very important that any time an ESM is connected in a redundant fashion, that the location of the root switch and various port costs in the network be accounted for to ensure correct data flow (need to prevent non-BladeCenter traffic from flowing through the ESMs).




Figure 5-60 Dual ESM's with cross connected LACP aggregated links to two separate Cisco switches

Configuration and verification for this example is similar to 5.5.5, “Single ESM, dual port LACP aggregation to two Cisco switches” on page 134, but has enough differences that a whole new procedure is presented here.

Step 1: Take down the link(s)

It is always advisable to disable the link(s) prior to making any configuration changes. Please see Table 5-1 on page 84 for procedures.

Step 2: Configure the IBM side of the link

This section walks you through the sequence of actions required to configure the ESM for this example (reference Table 5-21 on page 163).

The following assumptions have been made for this example:

- ▶ The user is already logged in to the ESMs with *root* level access.
- ▶ The following cabling was used (see Figure 5-61 on page 163) for a diagram:
 - ESM-1, Ext1 goes to Cisco-1 fa0/1 (or 2/3 if CatOS).
 - ESM-1, Ext2 goes to Cisco-1 fa0/2 (or 2/4 if CatOS).
 - ESM-1, Ext3 goes to Cisco-2 fa0/1 (or 2/3 if CatOS).
 - ESM-1, Ext4 goes to Cisco-2 fa0/2 (or 2/4 if CatOS).
 - ESM-2, Ext1 goes to Cisco-1 fa0/3 (or 2/5 if CatOS).
 - ESM-2, Ext2 goes to Cisco-1 fa0/4 (or 2/6 if CatOS).
 - ESM-2, Ext3 goes to Cisco-2 fa0/3 (or 2/5 if CatOS).
 - ESM-2, Ext4 goes to Cisco-2 fa0/4 (or 2/6 if CatOS)

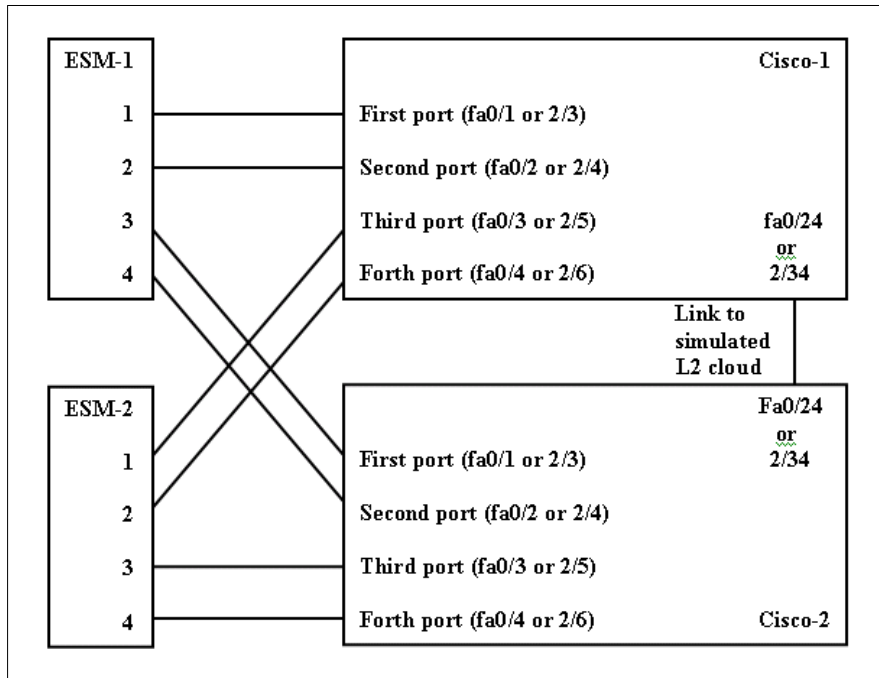


Figure 5-61 Cabling layout for Example 5-52 on page 151

- Commands are being performed in the sequence shown.
- The ESM is starting from a default config per the Example on page 85.
- Cisco switches being used are 10/100 based and we will be setting the ESM port to 100Mbps full duplex. This means a cross-over cable *must* be used for the link between the ESM and the Cisco switch.
 - If using switch ports that support 1000Mbps, the procedure to set speed will be slightly different from what is shown in the example.

Table 5-21 Configuring the ESM

Description and Comments	Actions to perform via Web interface to ESM
Step 2.1- <i>Configure speed and duplex.</i> As already noted, it will be necessary to use a cross-over cable on the links between the ESMs and the Cisco switches, as tuning off auto-negotiation also turns off the auto selection of MDI/MDI-X.	On ESM-1: <ul style="list-style-type: none"> ► Click the top connector (Ext1) in the picture of the ESM at the top of the page. ► Change Speed/Duplex to 100/Full. ► Click Apply. ► Repeat for Ext2 through Ext4 interfaces. Repeat process on ESM-2.
Step 2.2 - <i>Configure PVIDs</i> This places the desired blade server ports in to the desired VLANs. If the VLAN does not exist, it will be created automatically.	On ESM-1: <ul style="list-style-type: none"> ► Click Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID). ► For Bay 2, change the PVID to 5. ► For Bays 3 and 4, change the PVID to 10. ► All other PVIDs should be set for 1. ► Click Apply. Repeat process on ESM-2.

Description and Comments	Actions to perform via Web interface to ESM
<p>Step 2.3 - <i>Configure 802.1Q trunking</i> This will allow all Ext ports to carry traffic for VLAN 5 and VLAN 10 (already carrying VLAN 1 as an initial default).</p>	<p>On ESM-1:</p> <ul style="list-style-type: none"> ▶ Click Configuration->VLANs->Edit 802.1Q VLANs. ▶ Select VLAN 5 and click Edit. ▶ Set Interface Ext1 through Ext4 to Egress and make sure Tag box is checked. ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked. ▶ Click Apply. ▶ Click Configuration->VLANs->Edit 802.1Q VLANs. ▶ Select VLAN 10 and click Edit. ▶ Set interface Ext1 through Ext4 to Egress and make sure Tag box is checked. ▶ Make sure Ports 3 and 4 are set for Egress and the Tag box is unchecked for these two ports. ▶ Click Apply. <p>Repeat process on ESM-2.</p>
<p>Step 2.4- <i>Configure Link Aggregation.</i> This example makes use of LACP to dynamically negotiate link aggregation with the Cisco switch. Note that the Ext links used can not already be part of a different aggregation group, static or dynamic.</p>	<p>On ESM-1</p> <ul style="list-style-type: none"> ▶ Click Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Setup IEEE 802.3ad Link Aggregation. ▶ Change LACP to Enabled and click Apply. ▶ Click Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Link Aggregation Port Settings. ▶ Set Configure External Port From to include the first two ports: <ul style="list-style-type: none"> – Ext1 to Ext2. <ul style="list-style-type: none"> • This will set these 2 ports in to a single link. ▶ Change Mode to Enabled and click Apply. ▶ Set Configure External Port From to include the second two ports: <ul style="list-style-type: none"> – Ext3 to Ext4. – Change the Admin Key to 2 <ul style="list-style-type: none"> • This will set these 2 ports in to their own link, separate from the Ext1/Ext2 link. ▶ Change Mode to Enabled and click Apply. <p>Repeat process for ESM-2</p>
<p>Step 2.5 - <i>Save ESM config to NVRAM</i> Failure to perform this step will result in all changes to the ESM being lost if the @server BladeCenter is powered off or the ESM is otherwise restarted.</p>	<ul style="list-style-type: none"> ▶ Click Configuration -> Maintenance -> Save Changes. ▶ Click Save Configuration. ▶ Click Ok when complete. <p>Repeat process on ESM-2.</p>

Step 3: Configuring the Cisco switches

The following assumptions have been made for this example (reference Table 5-22 on page 165):

- ▶ VLANs 1, 5 and 10 already exist in the Cisco configuration and will be part of the 802.1Q trunk to the ESMs.
- ▶ The user is already logged in to the switches and the switches are in enable mode.
- ▶ The port connections are as described in Step 2 of this example.
- ▶ Commands are being performed in the sequence shown.
- ▶ The switches are starting from a default config per the Example on page 88 (IOS) or the Example on page 90 (CatOS).
- ▶ Cisco switches being used are 10/100 based and we will be setting the port to 100Mb full duplex.

- If using switch ports that support 1000Mbps, the command to set speed will be different from what is shown in the example.

Table 5-22 Configuring the Cisco equipment

Description and Comments	CatOS based switch	IOS based switch
Step 3.1 - <i>Configure speed and duplex.</i>	<ul style="list-style-type: none"> ▶ Perform the following two commands on both Cisco-1 and Cisco-2: set port speed 2/3-6 100 set port duplex 2/3-6 full 	<ul style="list-style-type: none"> ▶ Perform the following four commands on both Cisco-1 and Cisco-2: config t int range fa0/1 - 4 speed 100 duplex full
Step 3.2 - <i>Configure 802.1Q trunking.</i> Forces link to become an 802.1Q VLAN trunk.	<ul style="list-style-type: none"> ▶ Perform the following four commands on both Cisco-1 and Cisco-2: set trunk 2/3 nonegotiate dot1q set trunk 2/4 nonegotiate dot1q set trunk 2/5 nonegotiate dot1q set trunk 2/6 nonegotiate dot1q 	<ul style="list-style-type: none"> ▶ Perform the following three commands on both Cisco-1 and Cisco-2: switchport trunk encap dot1q switchport mode trunk switchport nonegotiate
Step 3.3 <i>Configure Spanning Tree port cost.</i> Setting the port cost higher than default performs two operations for this example: <ul style="list-style-type: none"> ▶ Forces all connections to a known Spanning Tree state (Blocking or Forwarding). ▶ Helps to prevent VLAN traffic between Cisco-1 and Cisco-2, other than VLAN 1, from being switched through the ESM. This forcing of the port cost helps to promote a more predictable flow. For more information, review the section titled: "Guidelines and comments - Spanning Tree" on page 79.	<ul style="list-style-type: none"> ▶ Set the port cost on Cisco-1 and Cisco-2 to control default flow: set spantree portcost 2/3-6 100 ▶ Note that for this particular design, leaving the links at their default cost will produce split traffic flows (see "Guidelines and comments - Spanning Tree" on page 79 for details on this behavior) 	<ul style="list-style-type: none"> ▶ Set the port cost on Cisco-1 and Cisco-2 to control default flow: spanning-tree cost 100 ▶ Note that for this particular design, leaving the links at their default cost will produce split traffic flows (see "Guidelines and comments - Spanning Tree" on page 79 for details on this behavior)
Step 3.4 <i>Configure Link Aggregation.</i> Need to perform on both Cisco-1 and Cisco-2 switch. Note that if you do not have the option of setting the "channelprotocol" to "lACP" for CatOS, or the channel-group to "active" for IOS, then more then likely you will need to upgrade your code to support LACP.	set channelprotocol lACP 2 <ul style="list-style-type: none"> ▶ The trailing 2 in the above command tells CatOS to enable LACP for Module 2 set port lACP-channel 2/3-4 mode active set port lACP-channel 2/5-6 mode active	int range fa0/1 - 2 channel-group 1 mode active <ul style="list-style-type: none"> ▶ This will create a logical interface named <i>Port-Channel1</i> and place the interfaces fa0/1 and fa0/2 in to it. int range fa0/3 - 4 channel-group 2 mode active <ul style="list-style-type: none"> ▶ This will create a logical interface named <i>Port-Channel2</i> and place the interfaces fa0/3 and fa0/4 in to it.

Description and Comments	CatOS based switch	IOS based switch
<p>Step 3.5</p> <p><i>Setting port cost on the aggregated link.</i></p> <p>Need to perform on both Cisco-1 and Cisco-2 switch.</p>	<p><i>(Does not apply)</i></p> <ul style="list-style-type: none"> ► CatOS uses the values from the individual ports to meet this requirement. 	<ul style="list-style-type: none"> ► Set the channel cost on Cisco-1 and Cisco-2 to control default flow: <pre>int port-channel1 spanning-tree cost 100 int port-channel2 spanning-tree cost 100</pre> ► See Step 3.3 for details. In this case, setting the cost for the whole aggregation as well as the individual ports.
<p>Step 3.6</p> <p><i>Save config to NVRAM.</i></p> <p>Only necessary on IOS based switches.</p>	<p><i>(does not apply)</i></p>	<p><code>write mem</code></p>

Step 4: Reconnecting the devices

This is the final step to bring the connection in to full operation. This will be the reverse of whatever procedure was used in Step 1. See Table 5-2 on page 84 for details on how to reestablish the links.

Step 5: Verifying the configuration

This section provides options for verifying the correct and desired operation.

Certain elements of this verification depend on the existence of a layer 2 network behind the two Cisco switches. This was simulated for this example by tying the two switches together with an 802.1Q trunk link (see the configuration shown in “Cat 3550 (IOS based switch) base configuration” on page 88 for IOS based switches or the configuration shown in “Cat 4003 (CatOS based switch) base configuration” on page 90 for CatOS based switches, for details on how this link was configured).

Verifying correct operation on the BladeCenter ESM

Verify the configurations on the ESM look as follows:

As in Figure 5-62 on page 167, verify port state by clicking **Configuration -> Port Setting -> Display Ports**, and reviewing the status of Ext1 through Ext4 (State should be *Enabled*, Speed/Duplex should be *100/Full*, and Connection should show 100M/Full/802.3x for all 4 Ext ports). Repeat for both ESMs.

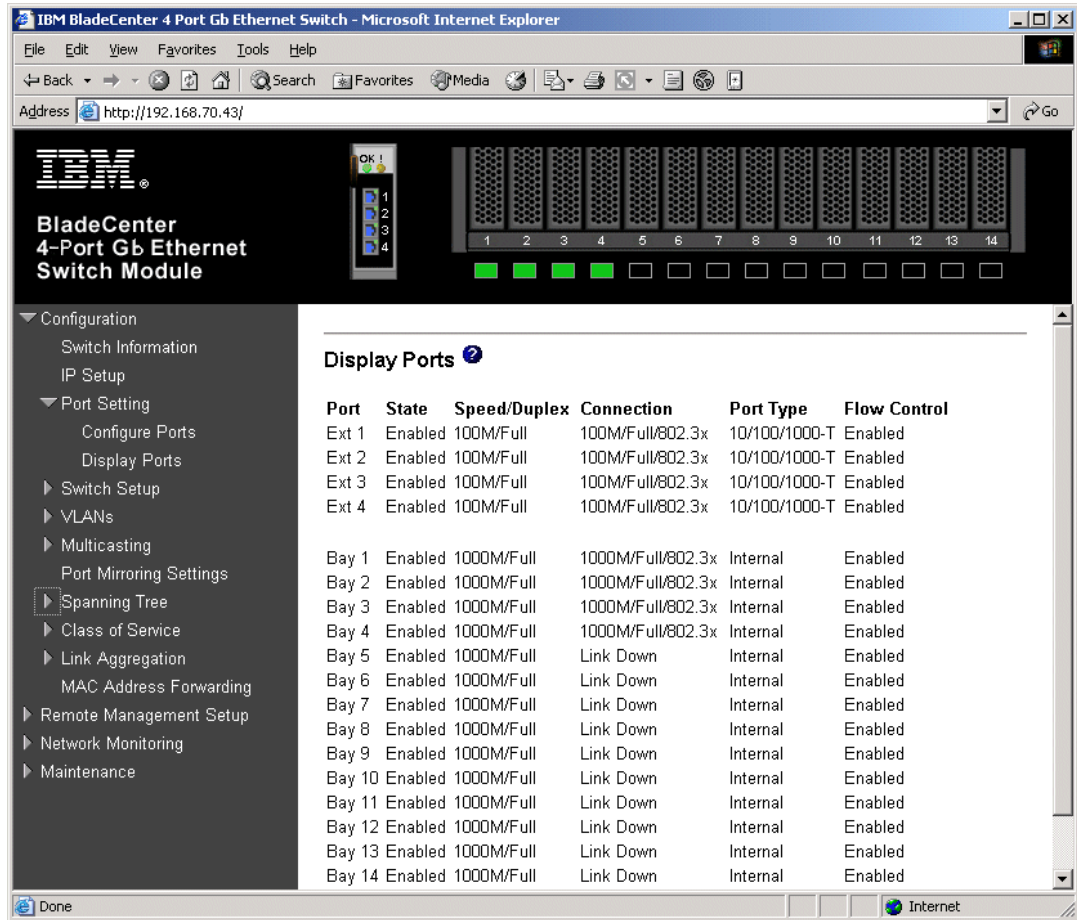


Figure 5-62 Verify ports are operational

As in Figure 5-63 on page 168, verify VLAN 1 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 1** and clicking **Edit**. Make sure the ports not on VLAN 1 (Ports 2, 3 and 4) are set to None and their associated **Tag** box is checked. Also note that the options for Ext ports 2 and 4 are greyed out. Once an aggregation forms, you can only make changes to the lowest number port in the aggregation (in this case, Ext1 and Ext3), which will then be inherited by the other ports in the aggregation when the change is applied. Also of note, if the aggregation had not yet formed, (for example, if the cables were still not plugged in), then the other Ext ports would not be greyed out, and you could make changes to them. Of course as soon as the aggregation came up, those changes would be overwritten by the lowest numbered port in the aggregation. The other settings could still be seen through the greyed out boxes, but they would not be getting used. Repeat for both ESMs.

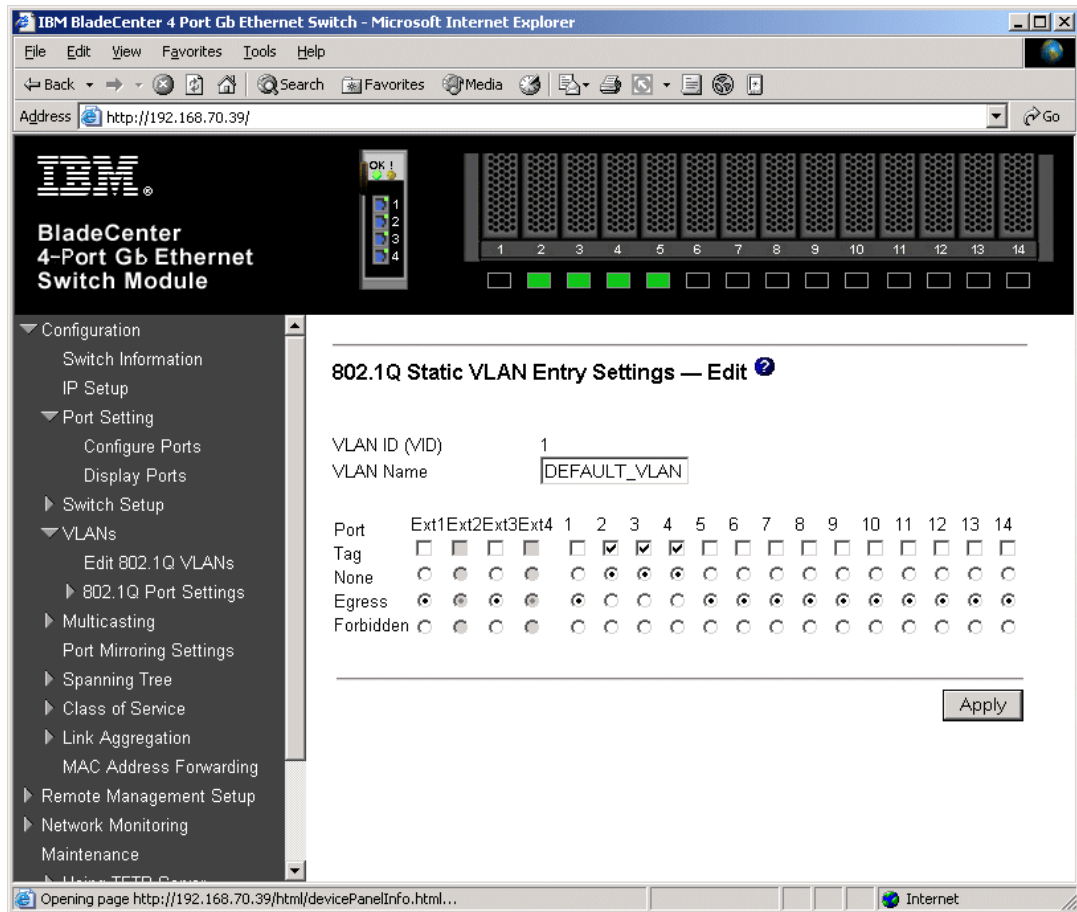


Figure 5-63 Verifying VLAN 1

As in Figure 5-64 on page 169, verify VLAN 5 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 5** and clicking **Edit**. Make sure Port 2 is set to Egress and that the Tag box is not checked. Make sure the Ext1 and Ext3 interfaces are set for Egress (allows VLAN 5 traffic to pass through Ext1 and Ext3 and any aggregated link that is part of Ext1's and Ext3's aggregation) and that the box for **Tag** is checked. Also, note the same greying out of Ext2 and 4 will also be seen here, if the aggregation has already become active. Repeat for both ESMs

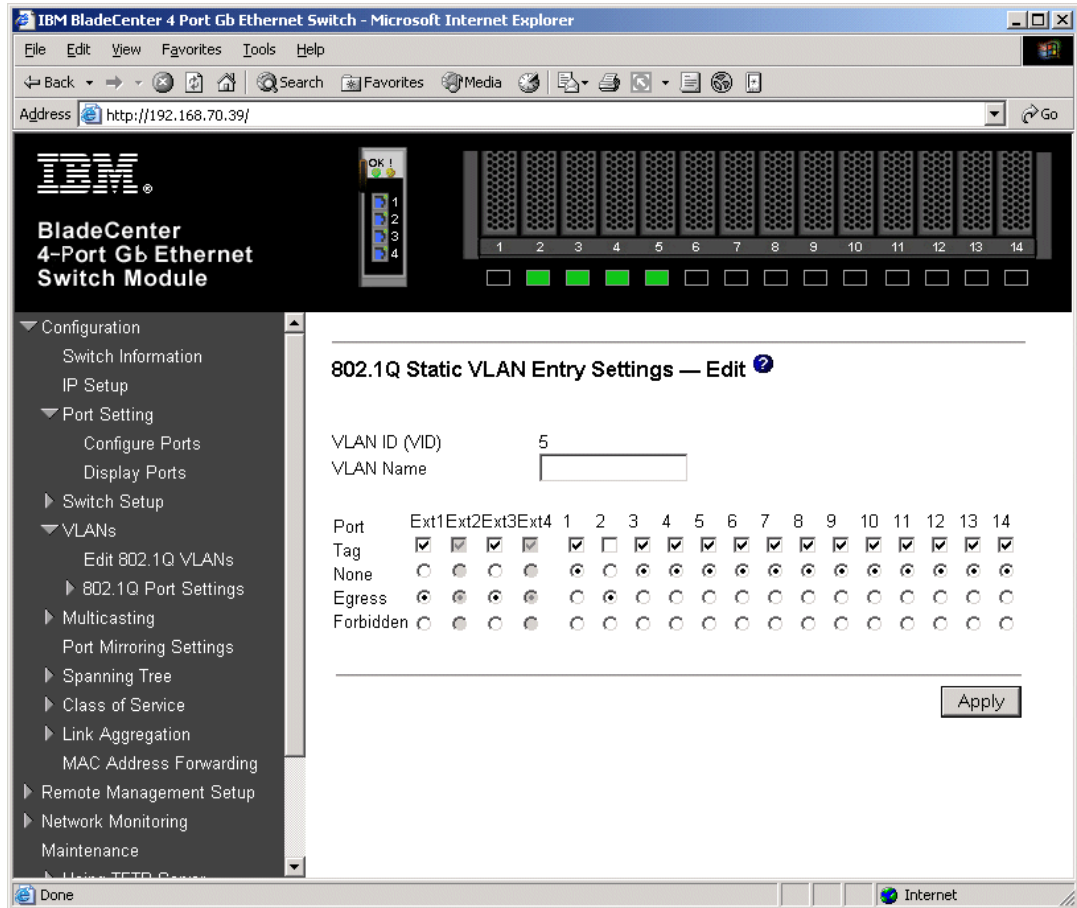


Figure 5-64 Verify VLAN 5

As in Figure 5-65 on page 170, verify VLAN 10 configurations by clicking **Configuration -> VLANs -> Edit 802.1Q VLANs**, selecting **VLAN 10** and clicking **Edit**. Make sure ports 3 and 4 are set to Egress and that the Tag box is not checked for either port. Make sure the Ext1 and Ext3 interfaces are set for Egress (allows VLAN 10 traffic to pass through Ext1 and Ext3 and any aggregated link that is part of Ext1's and Ext3's aggregation) and that the box for **Tag** is checked. Here again we see the greying out of Ext2 and Ext4 after an aggregation has formed. Repeat for both ESMs.




Figure 5-65 Verify VLAN 10

As in Figure 5-66 on page 171, verify Spanning Tree is forwarding traffic for the aggregation by clicking **Configuration -> Spanning Tree -> STP Port Settings**. Check the Status for Ext1 through Ext4. Based on the configuration of this example (root is Cisco-2 and all ports on the Cisco external switch set for higher than default port cost), these ports should all show as Forwarding. Notice here that you lose the ability to individually manage Spanning Tree options for all ports in the aggregation except the lowest numbered port (and that the aggregation handles Spanning Tree as a whole, rather than on a port by port basis). In this case, changes to Ext 1 will be inherited by Ext 2 and changes to Ext 3 will be inherited by Ext 4. Again, this is only true if the aggregation has already formed. Repeat for both ESMs.

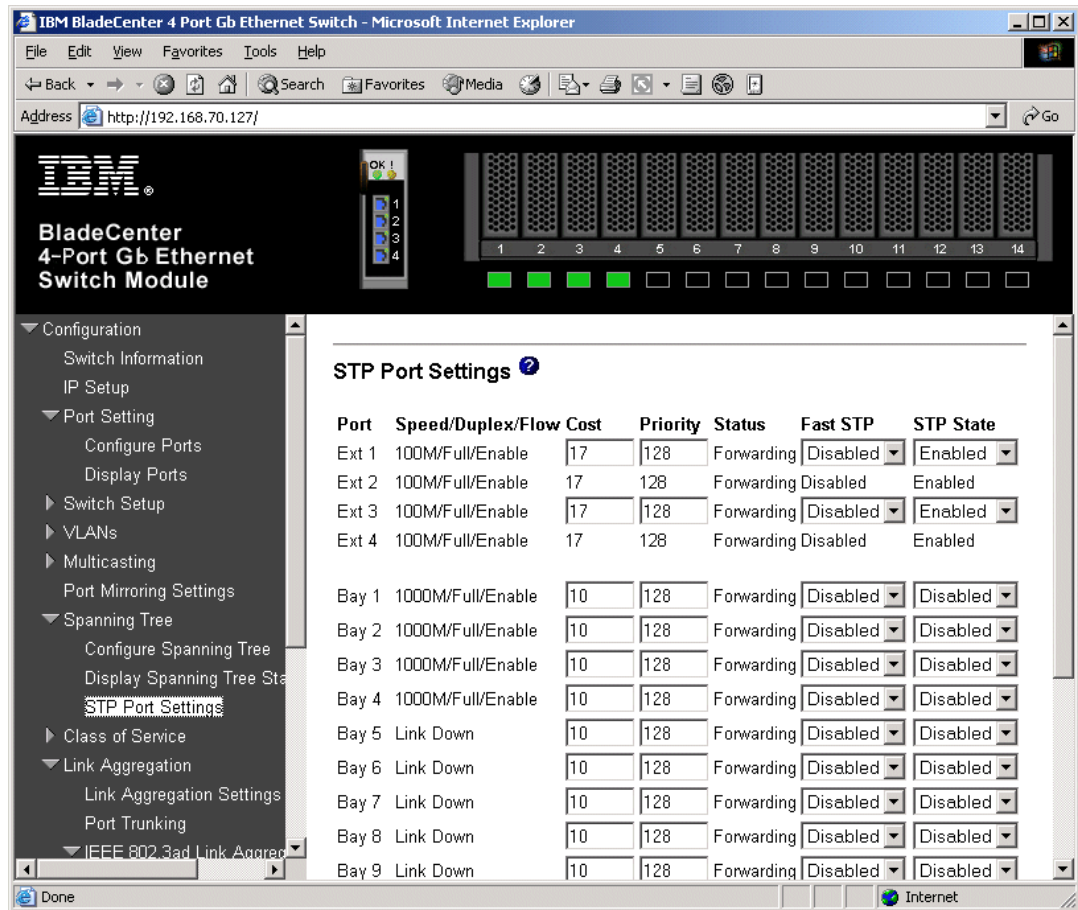


Figure 5-66 Verifying STP is forwarding on all 4 links in the aggregation

As in Figure 5-67 on page 172, one way to review the status of the aggregation is by clicking **Configuration -> Network Monitoring -> Applications Status -> Link Aggregation**. From here we can see that Ext 1 and Ext 2 are on one aggregation and Ext3 and Ext4 are on a different aggregation. Repeat for both ESMs.

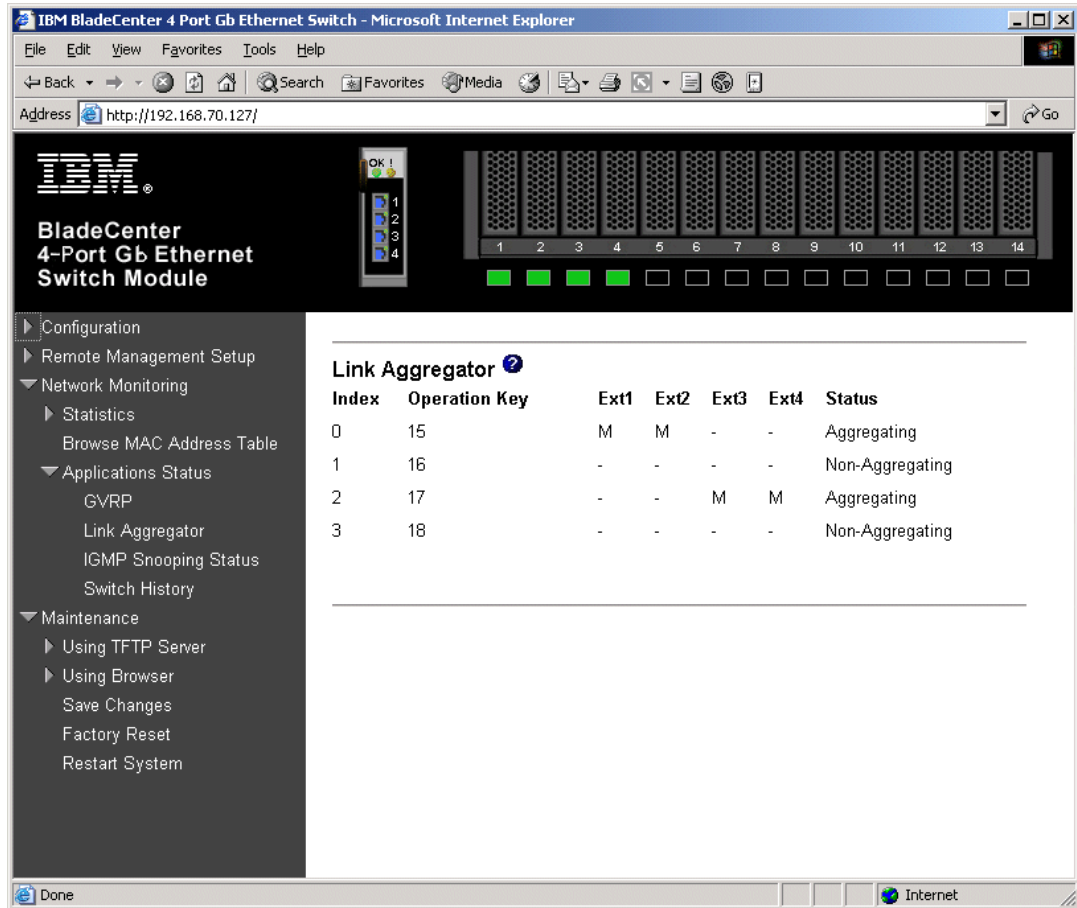


Figure 5-67 Link aggregation status

As in Figure 5-68 on page 173, another way to review the status of the aggregation can be done by clicking **Configuration -> Link Aggregation -> IEEE 802.3ad Link Aggregation -> Link Aggregation Port Setting** and noting that the Mode shows Enabled for all four ports (this is true whether the aggregation has formed or not), and the Status shows Active (shows as Active after the aggregation has formed, shows as Individual if the aggregation has not yet formed). Repeat for both ESMs.




Figure 5-68 Link aggregation status, another view

As in Figure 5-69 on page 174, you can also verify the aggregation is operational by generating traffic across the link and clicking **Network Monitoring -> Statistics -> Port Utilization** and monitoring Tx/sec and Rx/sec for the four Ext ports. Depending on how much traffic is being generate, and from how many sources, and what aggregations may be blocking, the numbers will vary across the ports. For this example, 100,000, 1400 byte pings were being sent from the Cisco-2 switch to the IP address of the ESM1. In this case, Ext3 is being used to transmit and Ext4 is receiving (these could have been on different ports, or even transmitting and receiving on the same port). If one were to pull the cable for Ext3, the traffic would switch over to a different Ext port on the same aggregation, usually with no loss of packets (usually under 1 second).

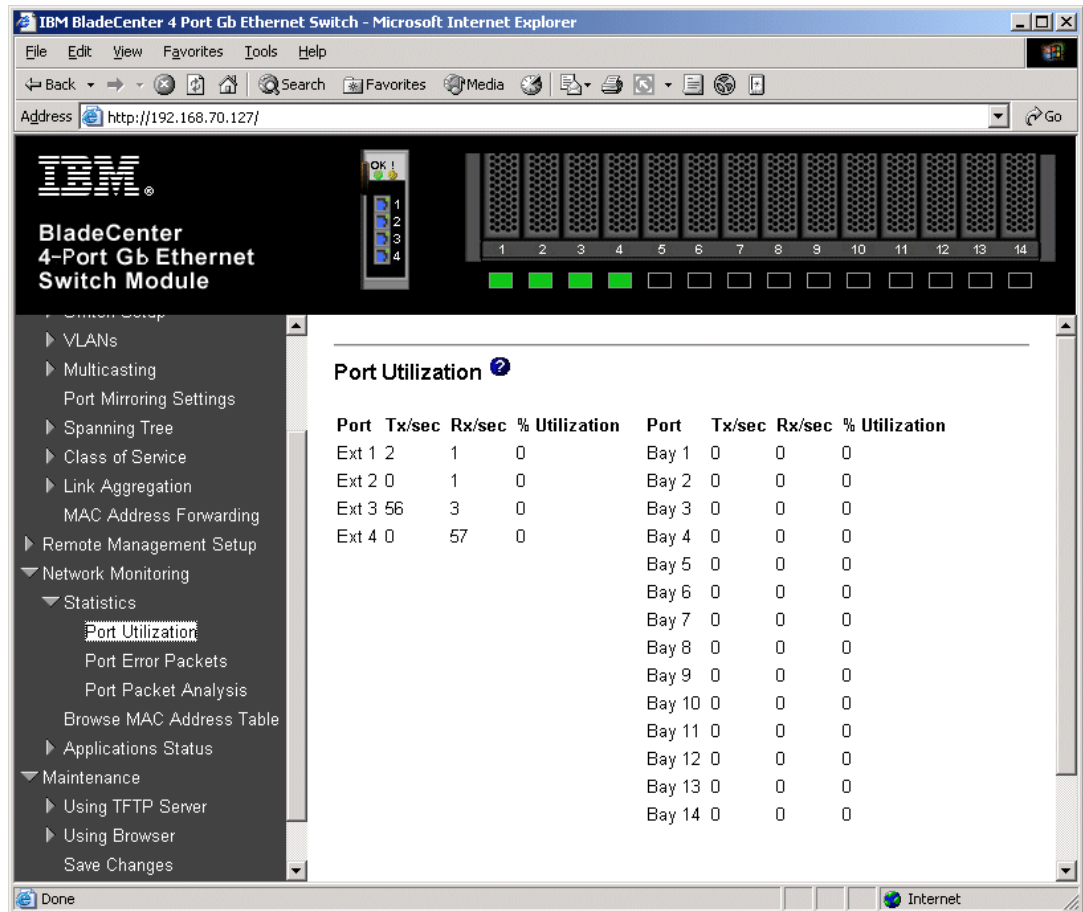


Figure 5-69 Example of aggregation load balancing

As in Figure 5-70 on page 175, with the same traffic flowing on the link, the cable to Ext3 is removed. The result is that the traffic previously carried on Ext3 is now on Ext4 (it could have gone to any of the available Ext ports in the same aggregation).




Figure 5-70 Traffic flow after cable to Ext3 removed

Verifying correct operation on the Cisco external switch

The following section includes some commands one can use to verify the desired configuration and operation of the Cisco equipment (reference Table 5-23 on page 176).

Table 5-23 Verifying the configuration and operation of the Cisco external switch of the connection

Description and Comments	CatOS based switch	IOS based switch
<p>Review running config for desired statements. Perform on both Cisco-1 and Cisco-2.</p>	<p>show config</p> <ul style="list-style-type: none"> ► Review for the following: <ul style="list-style-type: none"> – set channelprotocol lacp 2 – set port lacp-channel 2/3-6 120 <ul style="list-style-type: none"> • The number at the end may vary. – set port speed 2/3-6 100 – set port duplex 2/3-6 full – set trunk 2/3 negotiate dot1q 1-1005,1025-4094 <ul style="list-style-type: none"> • Should see a similar entry for each configured port (2/3 through 2/6). – set spantree portcost 2/3-6 100 – set spantree portvlancost 2/3 cost 99 – set spantree portvlancost 2/4 cost 99 – set spantree portvlancost 2/5 cost 99 – set spantree portvlancost 2/6 cost 99 <ul style="list-style-type: none"> • The portvlancost settings are a result of the <i>set spantree portcost</i> commands – set port lacp-channel 2/3-6 mode active <ul style="list-style-type: none"> • Notice that even though we are using 2/3-4 as an aggregation pair and 2/5-6 as an aggregation pairs, that the CatOS shows them together on the command line. They will still be used as desired but this can be confusing. This is also true for the <i>port lacp-channel</i> command. 	<p>show run</p> <ul style="list-style-type: none"> ► Review for the following on interface Port-channel1 and Port-channel2: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport nonegotiate – spanning-tree cost 100 ► Note that the values in Port-channel1 and 2 may not show up if the aggregation has never come up since first being configured. ► Review for the following on int fa0/1 through fa0/4: <ul style="list-style-type: none"> – switchport trunk encapsulation dot1q – switchport mode trunk – switchport nonegotiate – duplex full – speed 100 – channel-group 1 mode active <ul style="list-style-type: none"> • Channel-group 1 will be on fa0/1 and fa0/2 • Channel-group 2 will be on fa0/3 and fa0/4 – spanning-tree cost 100
<p>Show speed and duplex.</p>	<ul style="list-style-type: none"> ► Do the following command on each interface, 2/3 through 2/6 on both Cisco-1 and Cisco-2: show port status 2/3 ► Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100 	<ul style="list-style-type: none"> ► Do the following command on fa0/1 through fa0/4 on both Cisco-1 and Cisco-2: show int fa0/1 status ► Should show the following: <ul style="list-style-type: none"> – Status = Connected – VLAN - Trunk – Duplex = Full – Speed = 100 <ul style="list-style-type: none"> • Note that if you do this on the Port-channel1 or Port-channel2 interface, the speed and duplex might show up as a-100 and a-full.

Description and Comments	CatOS based switch	IOS based switch
<p>Show trunking status. Aggregation must be up before the trunk will come up.</p>	<ul style="list-style-type: none"> Do the following command on each interface, 2/3 through 2/6 on both Cisco-1 and Cisco-2: show port trunk 2/3 Should show the following: <ul style="list-style-type: none"> Mode = nonegotiate Encapsulation = dot1q Status = Trunking Native VLAN = 1 	<ul style="list-style-type: none"> Do the following command on fa0/1 through fa0/4 on both Cisco-1 and Cisco-2: show int fa0/1 trunk Should show the following: <ul style="list-style-type: none"> Mode = on Encapsulation =802.1q Status = trunk-inbndl (Po1) or (Po2) Native VLAN = 1
<p>Review status of the aggregated links. Perform on both Cisco-1 and Cisco-2.</p>	<p>show port lacp-channel</p> <ul style="list-style-type: none"> Should see all ports in Channel Mode Active (2/3 and 2/4 part of one aggregation and 2/5 and 2/6 part of a second aggregation). Note that you may also see the other Cisco switch listed in this output as the directly attached neighbor for one of the aggregations. <p>show port lacp-channel statistics</p> <ul style="list-style-type: none"> Run several times in a row, should show LACP Pkts Transmitted and Received climbing slowly (Transmitted usually higher than Received) for both channels. 	<p>show etherchannel summary</p> <ul style="list-style-type: none"> Should show the following: <ul style="list-style-type: none"> Protocol = LACP (for both channels) Ports fa0/1 -2 = (P) <ul style="list-style-type: none"> (P) = part of an aggregation group Ports fa0/3 -4 = (P) <ul style="list-style-type: none"> (P) = part of an aggregation group <p>show lacp counters</p> <ul style="list-style-type: none"> Run several times in a row, should show LACPDUs Sent and Recv climbing slowly (Sent usually higher than Recv) for both channels. <p>show etherchannel port-channel</p> <ul style="list-style-type: none"> Should show the following for both channels: <ul style="list-style-type: none"> Port state = Port-channel Ag-Inuse Protocol = LACP Both ports in EC State Active
<p>Show Spanning Tree status. Note that values shown here are very specific to the example configurations, and may not reflect production information.</p>	<p>show spantree 2/3-6</p> <ul style="list-style-type: none"> Based on the configuration and cost setting in this example: <ul style="list-style-type: none"> Cisco 1 switch should show all three VLANs for each aggregation pair Blocking for these interfaces. Cisco-2 switch (root) should show all three VLANs for each aggregation pair Forwarding for these interfaces. Both Cisco-1 and Cisco-2 should show forwarding for the simulated layer 2 link (2/34). 	<p>show spanning int fa0/1 show spanning int fa0/1 show spanning int po1</p> <ul style="list-style-type: none"> Based on the configuration and cost settings in this example: <ul style="list-style-type: none"> Cisco-1 switch should show BLK (Blocking) for all three VLANs, for each of these interfaces. Cisco-2 switch (root) should show FWD (Forwarding) for all three VLANs, for each of these interfaces. Both Cisco-1 and Cisco-2 should show forwarding for the simulated layer 2 link (fa0/24).
<p>Ping the ESM. Where x.x.x.x is the IP address of the ESM (must be in same VLAN as subnet being pinged).</p>	<p>ping x.x.x.x</p> <ul style="list-style-type: none"> Based on the Cat4K being at 192.168.70.202 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across. 	<p>ping x.x.x.x</p> <ul style="list-style-type: none"> Based on the Cat 3550 being at 192.168.70.200 on VLAN 1 and the ESM being at 192.168.70.127 on VLAN 1, should be able to ping across.

Description and Comments	CatOS based switch	IOS based switch
Attach a device to VLAN 5 and 10 on the Cisco switch and verify a ping to one of the blade servers in that same VLAN works as desired.	<ul style="list-style-type: none"> ► For VLAN 5, attach a device to port 2/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work) ► For VLAN 10, attach a device to port 2/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work). 	<ul style="list-style-type: none"> ► For VLAN 5, attach a device to port fa0/15 and attempt to ping across to the IP address on the blade server in slot 2 (note that IP subnets must match for ping to work). ► For VLAN 10, attach a device to port fa0/10 and attempt to ping across to the IP address on the blade server in slots 3 or 4 (note that IP subnets must match for ping to work).
Test redundancy.	<ul style="list-style-type: none"> ► Experiment with redundancy by taking down Cisco-2 (which would normally be carrying all traffic) and make sure Cisco-1 takes over. Verify with ping tests after Spanning Tree has stabilized. Should be able to still ping both ESMs from Cisco-1. Cisco-1 should also show all of its trunk ports as forwarding now. 	<ul style="list-style-type: none"> ► Experiment with redundancy by taking down Cisco-2 (which would normally be carrying all traffic) and make sure Cisco-1 takes over. Verify with ping tests after Spanning Tree has stabilized. Should be able to still ping both ESMs from Cisco-1. Cisco-1 should also show all of its trunk ports as forwarding now.

5.6 Troubleshooting ESM connections to Cisco devices

A section on troubleshooting could realistically be an entire document in its own right. This section covers some basic items to look for when problems are being experienced, but should not be considered a comprehensive resource for troubleshooting this environment.

Both IBM and Cisco have many excellent tools and documents to assist with troubleshooting, and it is recommended that these resources be utilized to their fullest.

There is no substitute for training and experience. This section, while offering some hints and tips, should not be considered a substitute for this training and experience. It is highly recommend that problems be directed to individuals with experience and training in the areas that are experiencing issues.

5.6.1 Troubleshooting specifics

This section includes issues that are unique to deploying the ESM in a Cisco environment.

Troubleshooting: Port will not come up

Link light not lighting on the ESM or the Cisco switch for the port in question.

On the ESM Web interface, clicking **Configuration->Port Setting->Display Port** shows the status of the connection as *Link Down*.

Cause:

The most likely cause is incorrect use of cable type (cross-over or straight-through).

Resolution:

As discussed elsewhere in this chapter, the ESM can support either a cross-over or a straight-through cable, as long as the port is configured for auto-negotiate for speed/duplex. When auto-negotiate is disabled (by hard coding the speed/duplex to something other than Auto), the auto-negotiation for cable type is also disabled, and the only cable that will work on

a switch-to-switch link is a cross-over cable. Based on this, it is always recommended to use a cross-over cable to attach the ESM to any switch, to ensure the link will always stay up, regardless of the speed/duplex setting.

Troubleshooting: LACP Aggregation link will not form

Aggregation link fails to pass traffic. Indications from both the ESM and Cisco external switch are that the link has failed to come ready. ESM may show all links as “individual”

Cause 1:

It has been noted that on occasion, depending on the order that the links come up, that the aggregation link fails to complete. This has not been noticed with static aggregation, only with dynamic (LACP) aggregation.

Resolution:

There are several possible approaches available to resolve this issue. The preferred method is to go into Enable mode on the Cisco switch and shutdown/disable all of the ports involved in the aggregation, and then bring them all back up at once. Table 5-1 on page 84 and Table 5-2 on page 84 provide the commands necessary to stop and start a group of interfaces with a single command. If access to the Cisco switches is not available, but you have physical access to the cables, unplug all of the cables in the aggregation link, and plug them back in, one at a time, starting with the lowest numbered port (from the perspective of the ESM) and proceeding to the highest numbered port (for example, if ports 3 and 4 were part of a two port link aggregation group, then plug in port 3 first, followed by port 4).

Cause 2:

A common cause is that one of the links in the aggregation is configured differently (for example, EXT1 is configured for 100/Full while Ext2 is configured for Auto).

Resolution:

When a link refuses to form, review the characteristics of all of the connections in the link (on both the Cisco external switch and the ESM side). Such things as trunk type and VLAN's carried can impact the formation of a dynamic (LACP) aggregation. All ports in an aggregation must be configured identically.

Cause 3:

Down level code on the Cisco switch. Prior to LACP becoming available, Cisco created a proprietary aggregation protocol named PAgP. This protocol is not compatible with LACP.

Resolution:

Upgrade code or use static aggregation. Check release notes to make sure your version of code supports LACP. Another way to spot non-support for LACP is lack of certain command options (for CatOS, not having the option of setting the “channelprotocol” to “lACP” is an indicator of non-LACP support. For IOS, not being able to set the “channel-group” to “active” is an indicator of non-LACP support).

Troubleshooting: Unable to access bottom ESM after Factory Reset

After performing a factory reset to the bottom ESM, and waiting for sufficient time to let it reboot, you still can not ping or otherwise contact it in any way (not through the Ext ports or through the Management Module).

Cause:

The cause is unknown, observed in lab during testing.

Resolution:

Re-seat the bottom ESM (pull it partially out of the bay, wait ten seconds and plug it back in to the bay). After waiting approximately one minute, you should now be able to ping through normal pathways.



Deploying the IBM @server BladeCenter in a Nortel environment

The features and functionality of the ESM component of the IBM @server BladeCenter allow the device to be connected to a network based on Nortel Networks equipment. Highlighted in this document are the BayStack 380-24T, a Smgigabit aggregation layer switch, and the Passport 8600, a core L2/L3 switch. Depending on the network design philosophy and/or business requirements, designs may call for one or the other.

In general, the BayStack 380-24T is a low cost option for providing a large amount of throughput to the network from the ESM(s). Depending on the particular solution implemented, it can provide a good deal of resiliency as well as throughput. In general, the amount of throughput is tempered by the level of resiliency required, especially when Spanning Tree is implemented as the failover mechanism.

The Passport 8600 is a higher cost option, but also higher performing option. From one perspective, the ESM already fulfills a typical aggregation layer switch function, thus moving the @server BladeCenter one step closer to the core network from an architecture perspective. Thus it makes sense to consider how an @server BladeCenter can connect to a core switch that provides L2/L3 capabilities and added levels of resiliency. The Passport 8600 can provide the greatest level of network reliability, including sub-second failover, while actively utilizing every single ESM external port on a dual ESM configuration for maximum throughput as well. A Passport 8600 based solution would most likely be used to connect to many @server BladeCenter chassis units in the same configuration as in a server farm or data center application.

6.1 Products test specifications

All configurations were tested with the following equipment and code revisions:

- ▶ IBM @server BladeCenter 8677
 - Blade servers
 - One with Windows 2000 server with SP3
 - Three with Red Hat Linux 7.2
 - One Management module
 - Two Ethernet Switch Modules (ESM)
 - ESM (FRU 59P6620):
 - Boot PROM version: 00.00.04
 - Firmware version: 00.00.62
 - Hardware version: Rev 2
- ▶ BayStack 380-24T
 - 24 10/100/1000Base-T ports
 - Software version: 2.0.1
 - Firmware version: 2.0.0.12
- ▶ Passport 8600 Switch 1
 - 8010 ten slot chassis
 - 8690SF (slot 5)
 - 8608GT (slot 1)
 - 8608 (slot 8)
 - Software version: 3.3.2.0
- ▶ Passport 8600 Switch 2
 - 8003 three slot chassis
 - 8691SF (slot 3)
 - 8608GT (slot 1)
 - 8608GT (slot 2)
 - Software version: 3.3.2.0

Note: Before you start using your Nortel Networks switch, you should ensure that the switch is set back to its factory defaults.

6.2 Nortel Networks feature descriptions

Below is a partial list of features supported by Nortel Networks equipment, some of which are proprietary to Nortel Networks. These features are listed here because they relate to the specific network designs being discussed in this document. Included in this section are caveats for the interworking of these features with IBM @server BladeCenter server.

Multi-Link Trunking

Multi-Link Trunking (MLT) is Nortel Networks' method of link aggregation that allows multiple Ethernet trunks to be aggregated together in order to provide a single logical trunk. An MLT provides the combined bandwidth of the multiple links, load sharing between these links as well as the physical layer protection against the failure of any single link. Although MLT is a proprietary algorithm, it is completely interoperable with most link aggregation protocols including 802.3ad static mode (no LACP) and Etherchannel.

Split MLT

Split-MLT (SMLT) is defined as an MLT, which terminates on two different aggregation switches. This provides nodal redundancy with sub-second failure recovery. Additional terms in relation to the SMLT feature are defined as follows:

SMLT aggregation switch – One of two switches performing the role of SMLT aggregator that connects to multiple wiring closet switches, edge switches or Customer Premise Equipment (CPE) devices.

IST (Inter Switch Trunk) – One or more parallel point-to-point links that connect two Aggregation switches together. The two aggregation switches utilize this channel to share information so that they may operate as a single logical switch. There can be only one IST per SMLT aggregation switch.

SMLT Client – A switch or server located at the edge of the network, such as in a wiring closet or CPE. An SMLT Client switch must be able to perform link aggregation (such as with MLT, 802.3ad static mode, or some other compatible method) but does not require any SMLT intelligence.

The IBM @server BladeCenter ESM module does not need to support SMLT for it to serve as an SMLT Client. As a client, links will be aggregated using 802.3ad static mode, and the ESM will be unaware that the individual links in the group terminate on different SMLT aggregation switches. This will be described further in the configuration that includes SMLT. For now it is only important to note that this feature interoperates with any standards compliant 802.3ad device acting in static mode.

Auto-MDI/MDI-X

Nortel Networks refers to this feature as Autopolarity. The terms can be considered interchangeable in the context of this document. The BayStack 380-24T switch supports Autopolarity as does the ESM.

Multiple Spanning Tree Groups

The Passport 8600 supports multiple Spanning Tree Groups, but for purposes of interoperability with the ESM, only one should be used. Note that SMLT requires Spanning Tree to be disabled entirely on the ports between the ESM and the Passport 8600. The other scenario described shows Spanning Tree being disabled altogether. This note concerning Spanning Tree Groups should be considered if for some reason it is deemed necessary to incorporate Spanning Tree into a network design that includes ESM's directly connected to Passport 8600's. The BayStack 380-24T does not support multiple Spanning Tree Groups.

Quality of Service

Both BayStack 380-24T and Passport 8600 have QoS features. The configuration of these features is beyond the scope of this document.

Pause Frames

All 1000Base-T interfaces on the Passport 8600 support receiving and sending pause frames (in hardware). However, certain blades disable the sending of pause frames (in software) due to the fact that the blades are not oversubscribed. This is because these blades will always be able to receive the offered load no matter how great. Thus it is normal, and expected, that the ESM, when connected to such blades, will display that 802.1x as “Asymmetric.” Passport 8600 blades that are oversubscribed will correctly negotiate “Symmetric” pause frame support on the links.

6.3 Limitations of configuration examples

The purpose of this manual is to show basic network designs and accompanying configurations. There are many advanced features and hardware options that could be incorporated into such network scenarios. For example, a Web Switching Module (WSM) could be added to a Passport 8600 to load balance Web traffic to and from an IBM *@server* BladeCenter server. However, such configurations are beyond the scope of this Redpaper.

6.4 Preliminary configuration for examples

6.4.1 General recommendations

This section includes high level comments on network design and recommendations for physical connection of Nortel Networks devices to an ESM.

Usage of Spanning Tree Protocol

Spanning Tree has been the traditional L2 method of providing resiliency. In many cases, Spanning Tree becomes the limiting factor in network resiliency. This is why many Nortel Networks resiliency features, such as SMLT, do not depend on Spanning Tree, and in fact these features surpass the failover performance of Spanning Tree by orders of magnitude. It should be noted that Spanning Tree is not employed in many of the following network designs.

Spanning Tree is also often used to safeguard against improper cabling or end-users that create loops by connecting two drops to a hub, etc. Since the BayStack 380-24T and Passport 8600 are not performing a wiring closet switch role in any of these scenarios, there is no threat of an unsuspecting end user inadvertently creating a loop on these particular switches. However it is up to the network administrator to determine the exact nature of risks posed to the network. Knowing when to use Spanning Tree and when not to is of critical importance. In a number of the scenarios described here, Spanning Tree is shown as disabled (or Fast Learning). It does not mean that it cannot be enabled (except in the SMLT design, Spanning Tree must be disabled), just that it does not serve a specific purpose in providing resiliency for that scenario.

Autonegotiation

Autonegotiation allows for a number of features other than speed and duplex settings to be enabled. In many cases, features specifically require autonegotiation. Because of this, it is recommended to use autonegotiation when connecting Nortel Networks equipment to another device that supports standards based autonegotiation. While it seems intuitive that

manual configuration is the safer way to configure two devices that are connected to each other, this rule of thumb turns out to have the opposite effect in most cases. Human intervention is far more likely to introduce a misconfiguration, and this is compounded by the fact that many network administrators assume that only one of the devices need to be manually configured, while leaving the other to autonegotiate (resulting in a duplex mismatch). Furthermore, resiliency features such as RFI and FEFI (not applicable to 1000Base-T) require autonegotiation in order to be enabled, just like the autopolarity feature. Manual setting of ports thus sacrifices such features, resulting in less network resiliency rather than more. Manual setting should be reserved on a case-by-case basis for known interoperability problems.

Cabling

Since autonegotiation is being used to detect link speed and duplex settings, auto-MDI/MDI-X will allow use of either a crossover or straight cable. It is preferable, however, to use a crossover cable since this optionally allows manual setting of ports, if desired.

6.4.2 Base configuration options common to all examples

The following are some configuration options established that are common to all of the examples. These are only for demonstration purposes in the examples, and more than likely will not be duplicated in your particular environment.

All configurations will have three VLANs configured: VLAN 1, VLAN 5 and VLAN 10. All VLANs should be tagged.

All configurations assume that all VLANs will be carried on all trunks. When using Spanning Tree in a scenario, the BayStack 380-24T or Passport 8600 will be preferred over the ESM to become the root bridge. The existing network may already have a root bridge that is more preferred over the BayStack 380-24T or Passport 8600. In any event, when using Spanning Tree it is important to plan the location of the root bridge to ensure optimal traffic flows. The ESM should not be allowed to become the root bridge

The following blade servers internal to the @server BladeCenter will be placed in the following VLANs during the configuration stage of each example:

- ▶ BladeServer 1: VLAN 1
- ▶ BladeServer 2: VLAN 5
- ▶ BladeServer 3: VLAN 10
- ▶ BladeServer 4: VLAN 10

6.4.3 Basic configuration procedures

There are three ways to configure the BayStack 380-24T, each of which could be used to setup the scenarios described later in this chapter. In order to initially configure a BayStack 380-24T an IP address must be assigned, through the console port. The management interface on the console port is called CI Menu and is shown below. Once an IP address is assigned, the console port could be used to complete the configuration. Alternatively, telnet could be used, which presents the same CI Menu interface, Web-based Management could be used, or Device Manager could be used. Both Device Manager and the Web-based Interface present GUIs to the user. In this guide, the Web-based Interface configuration screens are shown for configuration tasks beyond the initial switch setup of the BayStack 380-24T.

The Passport 8600 has two main user interfaces, the CLI and Device Manager. The CLI is accessible through the console interface and telnet. As with the BayStack 380-24T, Device Manager is a GUI based utility that can be used to manage and configure the Passport 8600. In the configuration examples of this guide that include the Passport 8600, the CLI will be used to configure the switch. The initial configuration of IP address must be performed via the console interface. After this, telnet could be used for the remaining steps, though this is not required or explicitly noted in the configuration steps.

When configuring each of the scenarios below, it is recommended that the configuration be performed on both devices prior to connecting the cabling between ESM and Nortel Networks switch (BayStack 380-24T or Passport 8600). This will ensure that Spanning Tree loops are not formed inadvertently as the steps are followed. Alternatively, the ports can be manually disabled through any of the management interfaces mentioned above. To do this on the BayStack 380-24T via the Web-based interface, go to **Configuration -> Port Management**, select **Disabled** and click **Submit** for each port that connects to the ESM as shown below:

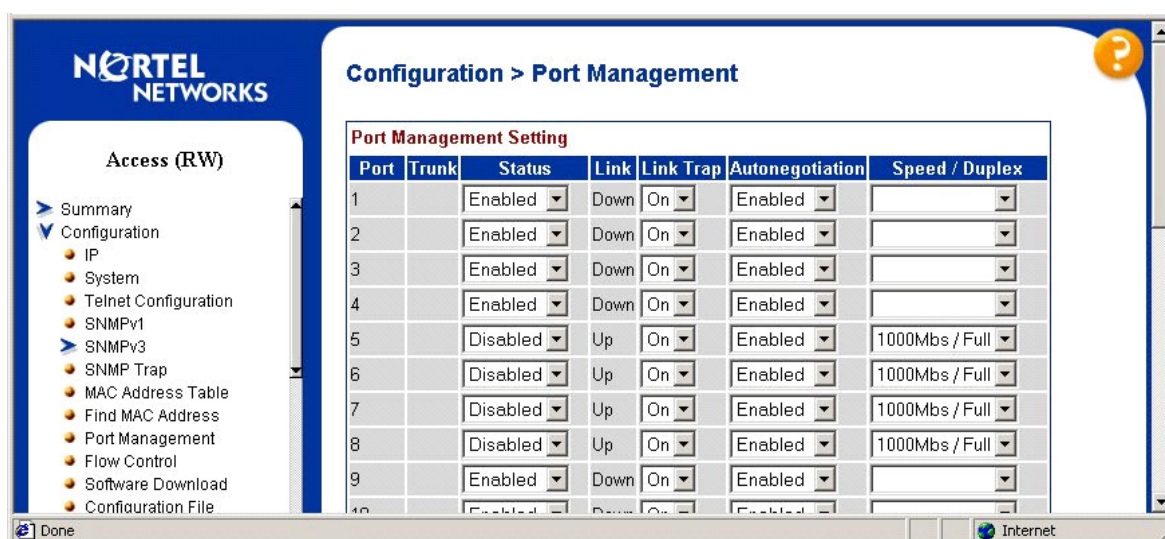


Figure 6-1 Port Management configuration window

The link status LEDs for each port should begin flashing. If any ports are currently members of a MLT group, the MLT group status will have to first be set to Disabled on the Application -> MultiLink Trunk -> Group screen before the ports can be disabled. Remember to re-enable the trunk after enabling the ports.

On the Passport 8600, port(s) can be disabled using the following command:

```
config ethernet <portlist> state disable
```

Ports are re-enabled using the same command substituting the word “enable” as the state.

Also when making configuration changes this is the safest practice. For example, if the network grows to the point where it is necessary to change the configuration, then it is conceivable that a loop may be formed while performing configuration tasks. It is important that you understand the significance of each configuration change on a live network. While it is possible to migrate from one scenario to another without introducing such problems depending on the order of steps taken, this approach should only be used with extreme caution.

6.4.4 Base configuration tasks for the ESM

Minor base configuration changes need to be applied to the ESM.

The only change that may be necessary after the Factory Reset, is the need to enable the external ports (Ext 1 through Ext 4). Under certain circumstances, the external ports are set to Disable by default after a factory reset, and will need to be set to Enable for each external interface. To configure these ports via the Web interface, log in to the IP address of the ESM with root privileges and click the top connector (**Ext 1**) in the picture of the ESM at the top of the page. Set the State to Enable and click **Apply**. Repeat for each of the four external interfaces (Ext1 through Ext4).

Configure Ports — Ext 1 ?

Connection	Link Down
State	Enabled
Speed/Duplex	Auto
Flow Control	Enabled

Figure 6-2 ESM showing Ext 1 set to enable (should be set the same for other EXT ports)

The following set of screen shots show some of the default configurations on the ESM after a factory default reset is performed.

From **Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID)**:

Port VLAN ID (PVID) ?

Port	PVID	Port	PVID
Ext 1	1	Bay 1	1
Ext 2	1	Bay 2	1
Ext 3	1	Bay 3	1
Ext 4	1	Bay 4	1
		Bay 5	1
		Bay 6	1
		Bay 7	1
		Bay 8	1
		Bay 9	1
		Bay 10	1
		Bay 11	1
		Bay 12	1
		Bay 13	1
		Bay 14	1

Figure 6-3 ESM showing default PVID settings

From **Configuration->VLANs->Edit 802.1Q VLANs**, select **VLAN 1**, click **Edit**.

Port VLAN ID (PVID) ?

Port	PVID	Port	PVID
Ext 1	1	Bay 1	1
Ext 2	1	Bay 2	1
Ext 3	1	Bay 3	1
Ext 4	1	Bay 4	1
		Bay 5	1
		Bay 6	1
		Bay 7	1
		Bay 8	1
		Bay 9	1
		Bay 10	1
		Bay 11	1
		Bay 12	1
		Bay 13	1
		Bay 14	1

Figure 6-4 ESM showing default settings for VLAN 1

From **Configuration->Link Aggregation->IEEE 802.3ad Link Aggregation->Link Aggregation Port Settings**:

Link Aggregation Port Settings ?

Port	Priority	Admin Key	Oper Key	Mode	Status
Ext 1	128	15	15	Disabled	Individual
Ext 2	128	16	16	Disabled	Individual
Ext 3	128	17	17	Disabled	Individual
Ext 4	128	18	18	Disabled	Individual

Configure External Port from Ext 1 to Ext 1

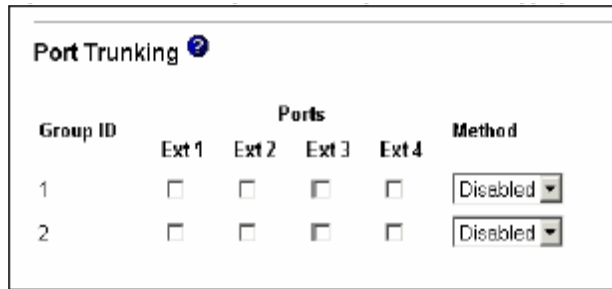
Priority1

Admin Key1

ModeDisabled

Figure 6-5 ESM showing default settings for Link Aggregation LACP

From **Configuration->Link Aggregation->Port Trunking**:



The image shows a 'Port Trunking' configuration window. It has a title bar with a question mark icon. Below the title bar, there is a table with the following structure:

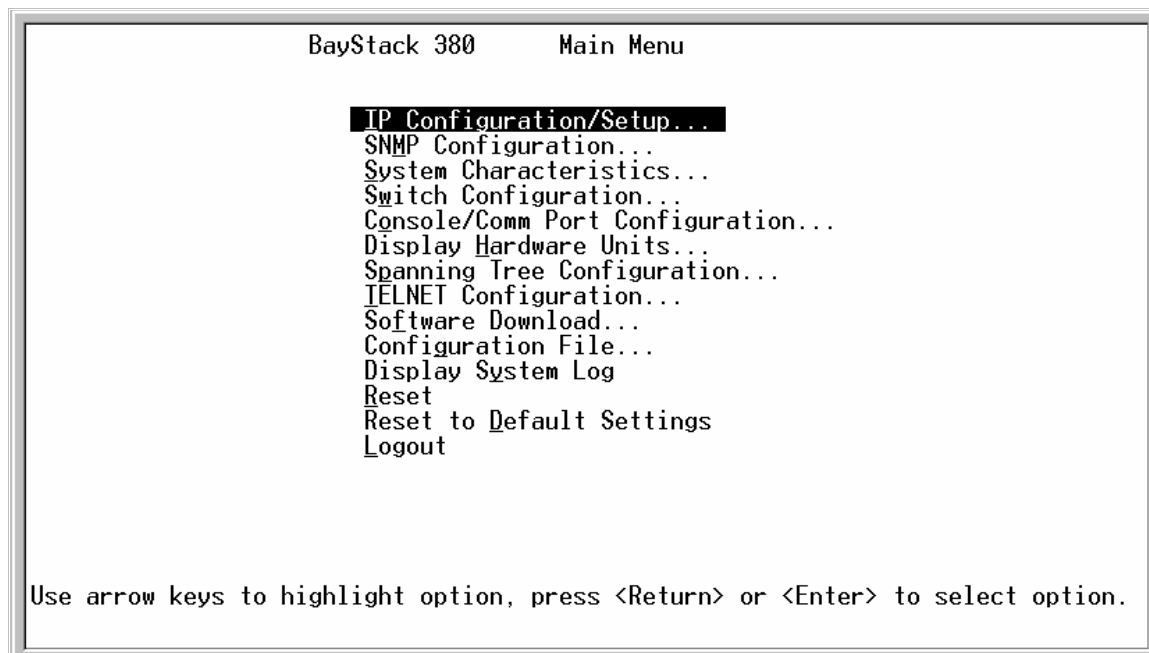
Group ID	Ports				Method
	Ext 1	Ext 2	Ext 3	Ext 4	
1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled
2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled

Figure 6-6 ESM showing default settings for static link aggregation

6.4.5 Base configuration tasks for BayStack 380-24T:

This configuration procedure assumes factory default settings. The console interface is the only interface that can be used to initially administer since the switch has no IP address.

The Main Menu is shown in Figure 6-7.



The image shows a terminal window titled 'BayStack 380 Main Menu'. The menu options are listed as follows:

```
BayStack 380      Main Menu

IP Configuration/Setup...
SNMP Configuration...
System Characteristics...
Switch Configuration...
Console/Comm Port Configuration...
Display Hardware Units...
Spanning Tree Configuration...
TELNET Configuration...
Software Download...
Configuration File...
Display System Log
Reset
Reset to Default Settings
Logout
```

Use arrow keys to highlight option, press <Return> or <Enter> to select option.

Figure 6-7 BayStack 380 Main Menu

Choose **IP Configuration/Setup** from the Main Menu. Assign an IP address to the switch in Figure 6-8:

IP Configuration/Setup			
BootP Request Mode: [BootP Disabled]			
	Configurable	In Use	Last BootP
In-Band Switch IP Address:	[192.168.47.250]	192.168.47.250	0.0.0.0
In-Band Subnet Mask:	[255.255.255.0]	255.255.255.0	0.0.0.0
Default Gateway:	[192.168.47.1]	192.168.47.1	0.0.0.0
IP Address to Ping:	[0.0.0.0]		
Start Ping:	[No]		
Use space bar to display choices, press <Return> or <Enter> to select choice. Press Ctrl-R to return to previous menu. Press Ctrl-C to return to Main Menu.			

Figure 6-8 IP Configuration/Setup window

Set Switch Passwords. Go to “Console/Comm Port Configuration” as shown in Figure 6-9.

Console/Comm Port Configuration	
Comm Port Data Bits:	8 Data Bits
Comm Port Parity:	No Parity
Comm Port Stop Bits:	1 Stop Bit
Console Port Speed:	[9600 Baud]
Console Switch Password Type:	[Local Password]
Telnet/Web Switch Password Type:	[Local Password]
Read-Only Switch Password:	[guest]
Read-Write Switch Password:	[Nortel]
Primary RADIUS Server:	[0.0.0.0]
Secondary RADIUS Server:	[0.0.0.0]
UDP RADIUS Port:	[1645]
RADIUS Shared Secret:	[]
Enter text, press <Return> or <Enter> when complete. Press Ctrl-R to return to previous menu. Press Ctrl-C to return to Main Menu.	

Figure 6-9 Console/Comm Port Configuration

Now that the switch has an IP interface configured the Web interface can be used to configure the switch through a Web browser. To log in to Web interface, use the Username “RW” and the correct password.

Go to **Configuration -> System** as shown in Figure 6-10. Define the system information and click **Submit**.

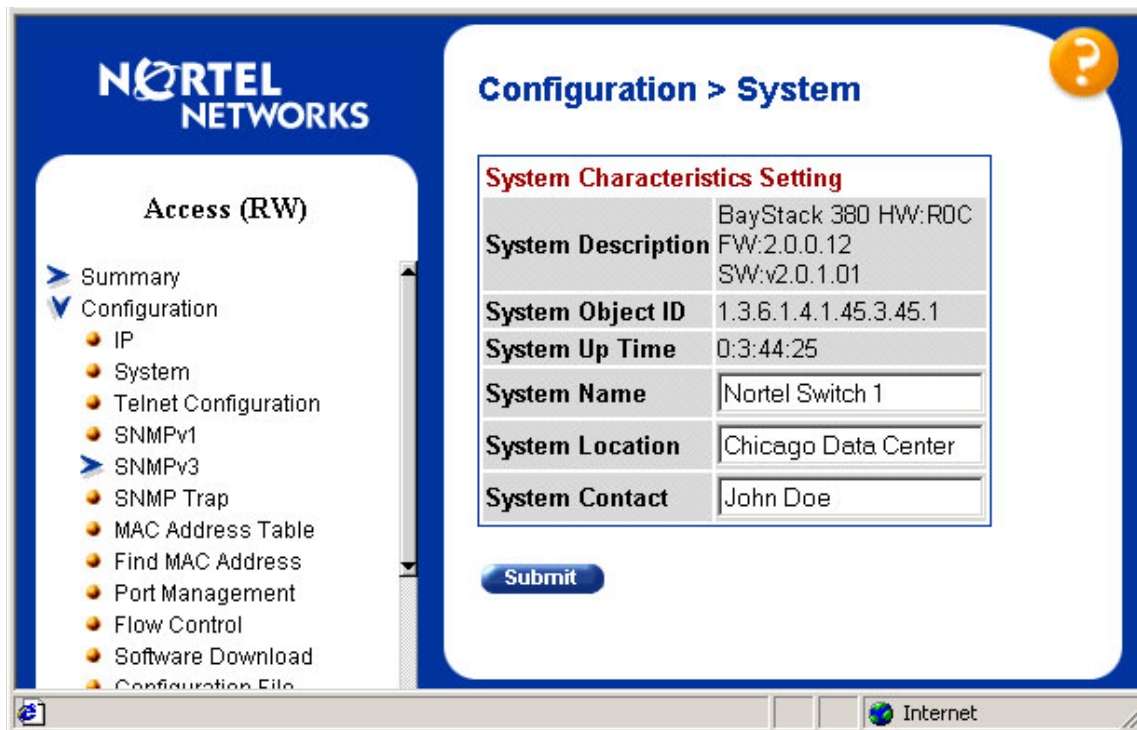


Figure 6-10 Configuration > System window

Go to **Application -> VLAN -> VLAN Configuration**. Create VLAN 5 as shown in Figure 6-11 and click **Create VLAN**.

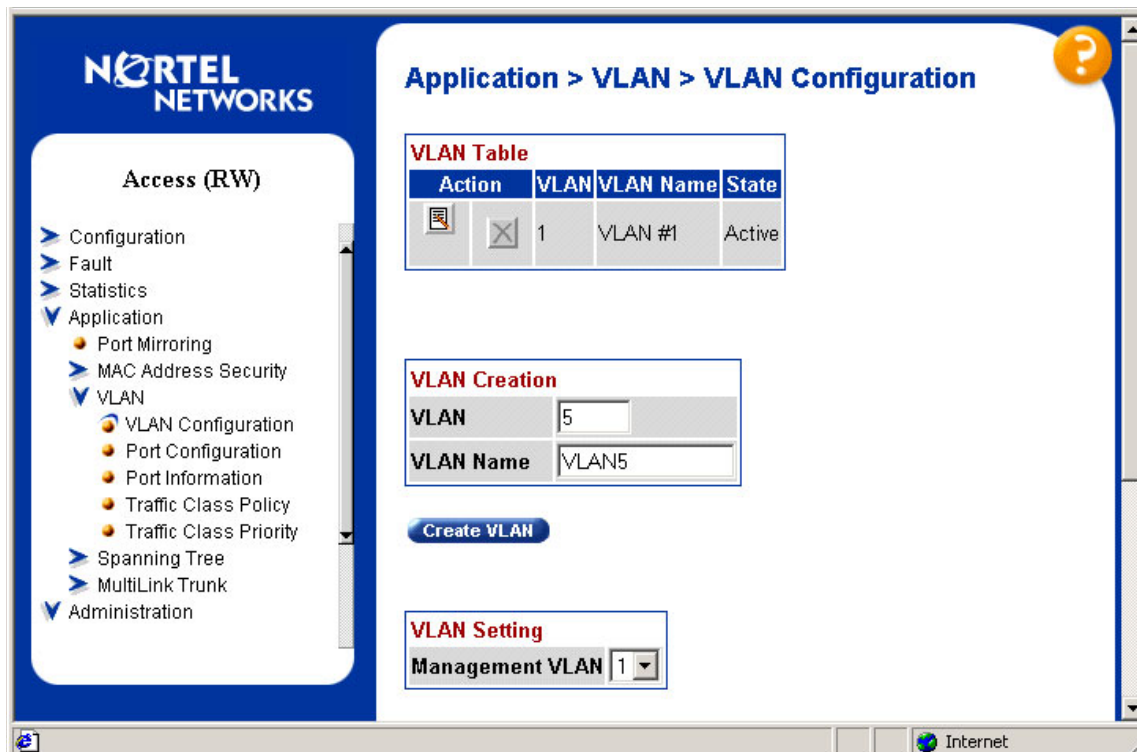


Figure 6-11 Application > VLAN > VLAN Configuration

The spanning tree bridge priority default of the BS380 (0x8000h) is the same as the default of the ESM. In scenarios where Spanning Tree is to be used for redundancy, it is desirable that the ESM not be elected as the root bridge. To ensure that this does not happen, the Bridge Priority of the BayStack 380 should be lowered. Even though some configurations recommend not running Spanning Tree between ESM and BayStack 380 at all, it is good practice to go ahead and modify the bridge priority.

Go to **Application -> Spanning Tree -> Bridge Information** (Figure 6-12). Set the STP bridge priority to a lower value (0x7000).

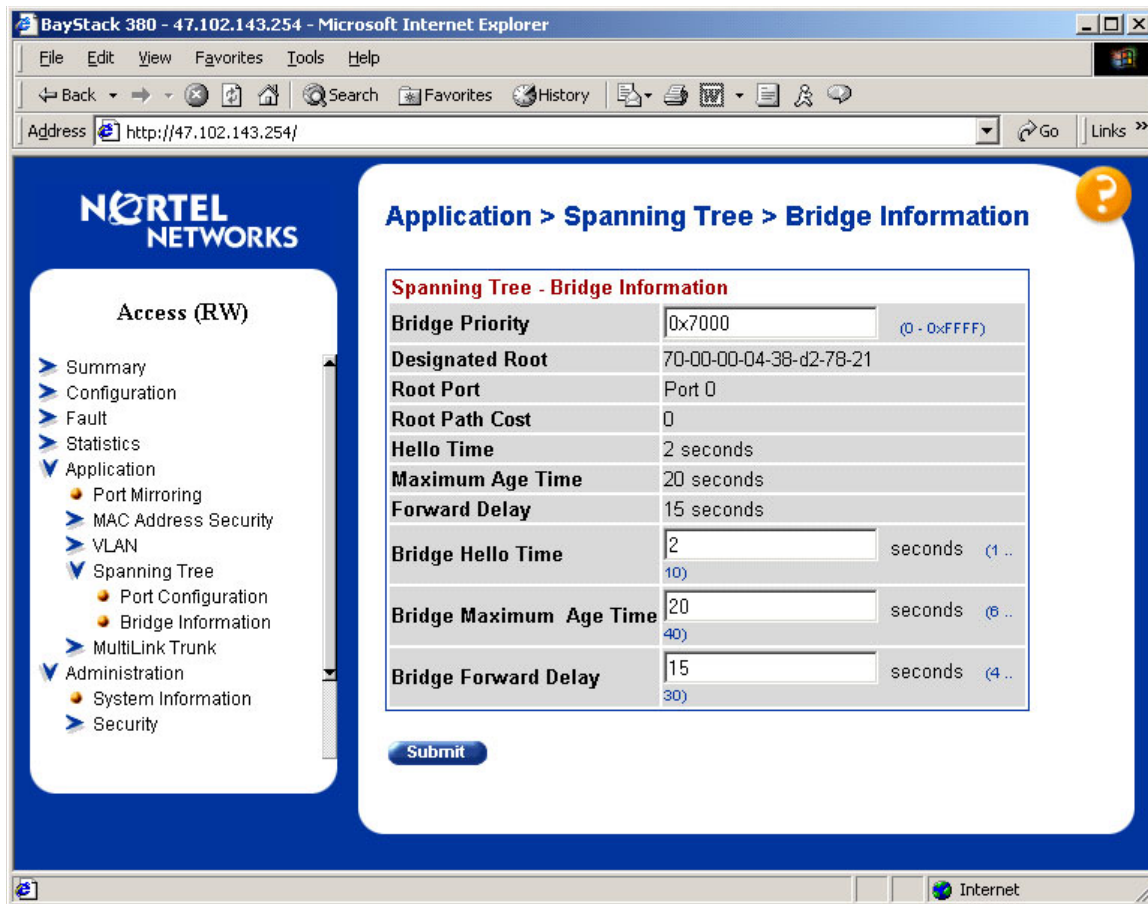


Figure 6-12 Application > Spanning Tree > Bridge Information

6.4.6 Base configuration tasks for Passport 8600

This configuration procedure assumes factory default settings. The commands below highlights the steps required to default an 8600. The out-of-band management IP should not be affected by defaulting the Passport 8600.

```
config bootconfig flags factorydefault true
save boot
boot -y
```

Once the 8600 has been rebooted the following command should be issued to save the default configuration:

```
save config
```

The console interface is the only interface that can be used to initially administer since the switch has no IP address. The following commands can be used to assign an IP address to the management interface on the switch fabric. This is an out of band port.

```
Passport-8610:5# config bootconfig net mgmt
Passport-8610:5/config/bootconfig/net/mgmt# ip 192.168.47.252/24 cpu-slot 5
Passport-8610:5/config/bootconfig/net/mgmt#
```

Figure 6-13 Console interface

The next set of commands will set the system level attributes, such as switch name, location, and contact. Also shown below is a password change for the administration user account, “rwa”. Other passwords should be changed as well for security reasons, but this is not shown here. Those usernames are “rw”, “ro”, “l2”, and “l3”.

```
Passport-8610:5# config sys set
Passport-8610:5/config/sys/set# name PP8600-1
PP8600-1:5/config/sys/set# location SanFran Data Center
PP8600-1:5/config/sys/set# contact Jane Doe
PP8600-1:5/config/sys/set# back
PP8600-1:5/config/sys# top
PP8600-1:5#
PP8600-1:5# config cli
PP8600-1:5/config/cli# password rwa 123456

Enter the old password : ****
Enter the New password : ****
Re-enter the New password : ****

Password changed successfully
PP8600-1:5/config/cli# top
PP8600-1:5#
```

Figure 6-14 Console interface

VLANs 5 and 10 will be created next. VLAN 1 is the default VLAN so it is not necessary to explicitly create it.

```

PP8600-1:5# config vlan 5 create
PP8600-1:5/config/vlan/5/create# byport 1
PP8600-1:5/config/vlan/5/create# name VLAN5
PP8600-1:5/config/vlan/5/create# back
PP8600-1:5/config/vlan/5# back
PP8600-1:5/config# vlan 10 create
PP8600-1:5/config/vlan/10/create# byport 1
PP8600-1:5/config/vlan/10/create# name VLAN10
PP8600-1:5/config/vlan/10/create# top
PP8600-1:5#
PP8600-1:5# config stg 1
PP8600-1:5/config/stg/1# priority 24567
PP8600-1:5/config/stg/1# top
PP8600-1:5#

```

Figure 6-15 Console interface

6.5 Configuration examples

This section contains the various design scenarios for connecting an IBM eServer BladeCenter ESM to a network based on Nortel Networks equipment.

6.5.1 Single ESM with Link Aggregation to Single BayStack 380-24T

This connectivity scenario covers the most basic design. A single ESM is connected to a single BayStack 380. Since the ESM is aggregating as many as 14 servers each with a Gigabit connection to the ESM, uplink capacity to the rest of the network is important. That is why this basic scenario incorporates trunking four Gigabit links between ESM and BayStack 380 (see Figure 6-16). This offers the maximum throughput between ESM and BayStack 380. Fewer links can be used if desired, and the following configuration is easily modified accordingly.

This topology is resilient against single link or port failures, but not ESM or switch failure. This design is useful for networks that need the maximum throughput of one ESM but where redundancy of the switches is not a concern.




Figure 6-16 IBM eServer BladeCenter and BayStack 380-24T

Configuring the ESM

This section walks through the sequence of actions (Figure 6-1) required to configure the ESM for this example. The following assumptions have been made for this example:

- ▶ The user is already logged into the ESM with root level access.
- ▶ Ports EXT1 through EXT4 on the ESM in Switch Module Bay 1 are being used between IBM and Nortel.
- ▶ Commands are being performed in the sequence shown.
- ▶ Nortel Networks switches are BayStack 380-24T gigabit switches. Autonegotiate will result in ports operating at 1000Base-T.

Table 6-1 Configuring the ESM

Description and Comments	Description and Comments
<p>Step 1 - Configure PVIDs</p> <p>This places the desired blade server ports into the desired VLANs. If the VLAN does not exist, it will be created automatically</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID) ▶ For Bay 2, change the PVID to 5 ▶ For Bays 3 and 4, change the PVID to 10 ▶ All other PVIDs should be set for 1 ▶ Click Apply.
<p>Step 2 - Configure 802.1Q trunking</p> <p>This will allow EXT 1 through 4 to carry traffic for VLANs 1, 5, and 10.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 1 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 1 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Select VLAN 5 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 10 and click Edit ▶ Set interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure the Tag boxes are checked ▶ Make sure Ports 3 and 4 are set for Egress and the Tag boxes are unchecked for these two ports ▶ Click Apply
<p>Step 3 - Configure Link Aggregation</p> <p>This will allow EXT 1 through 4 to be bundled together into an aggregate link.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->Port Trunking ▶ Set Ext1, Ext2, Ext3, and Ext4 as members of Group 1. ▶ Set Method to Enabled ▶ Click Apply

Step 4 - Ensure LACP is Disabled.	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->IEEE 802.3ad Link Aggregation->Setup IEEE 802.3ad Link Aggregation. ▶ Set LACP from Enabled to Disabled. ▶ Click Apply
-----------------------------------	--

Configuring the BayStack 380-24T

Step 1 – Configure Ports

Click **Application -> VLAN -> Port Configuration**. Configure the links to be tagged. Port names can also be assigned. Click **Submit**.

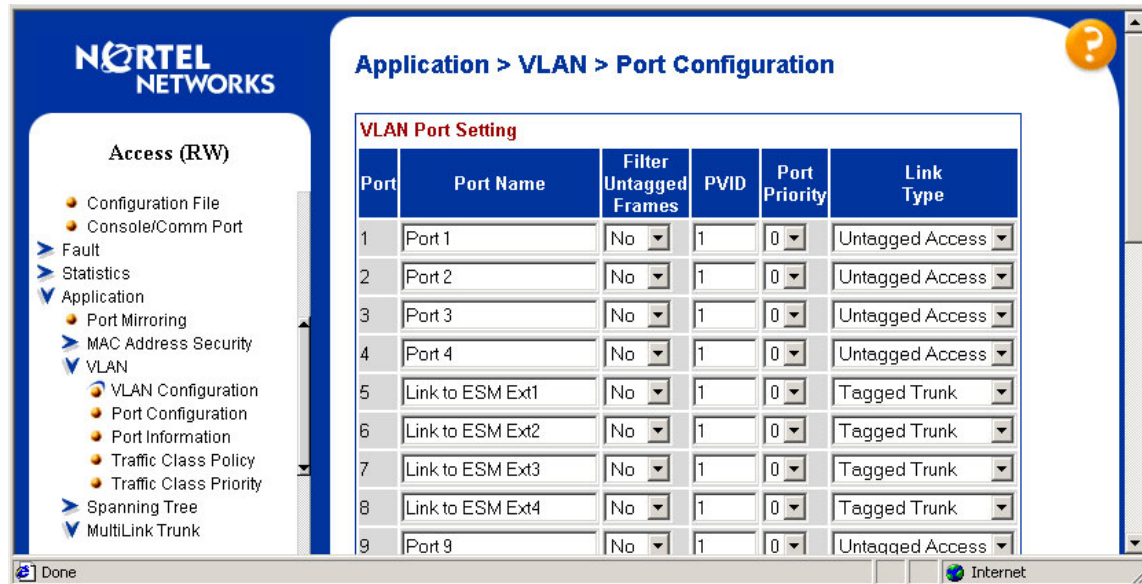


Figure 6-17 Application > VLAN > Port Configuration window

Step 2 – Add VLANs to each link member

Click **Application -> VLAN -> VLAN Configuration**. Click the action button next to VLAN 5. Add the ports to VLAN 5 as shown in Figure 6-18 and click **Submit**. Repeat this for VLAN 10 and click **Submit**.

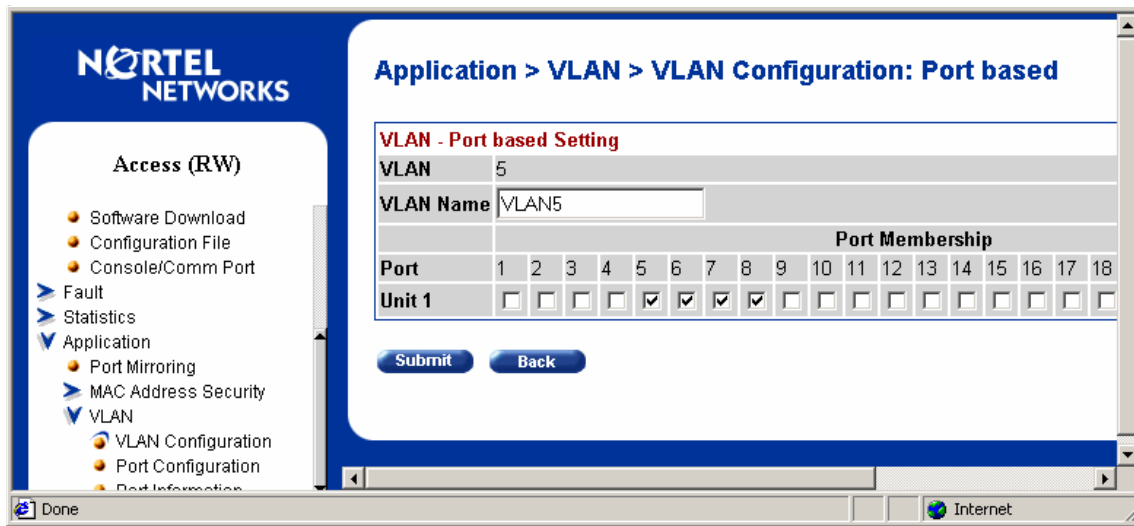


Figure 6-18 Application > VLAN > VLAN Configuration: Port based window

Step 3 – Configure MLT

Click **Application** -> **MultiLink Trunk** -> **Group**. Specify ports 5, 6, 7, and 8, and STP Learning to Fast (Figure 6-19). Click **Submit**. The trunk is not enabled at this point. It must explicitly be enabled by changing the setting in the table to Enabled and clicking **Submit**.

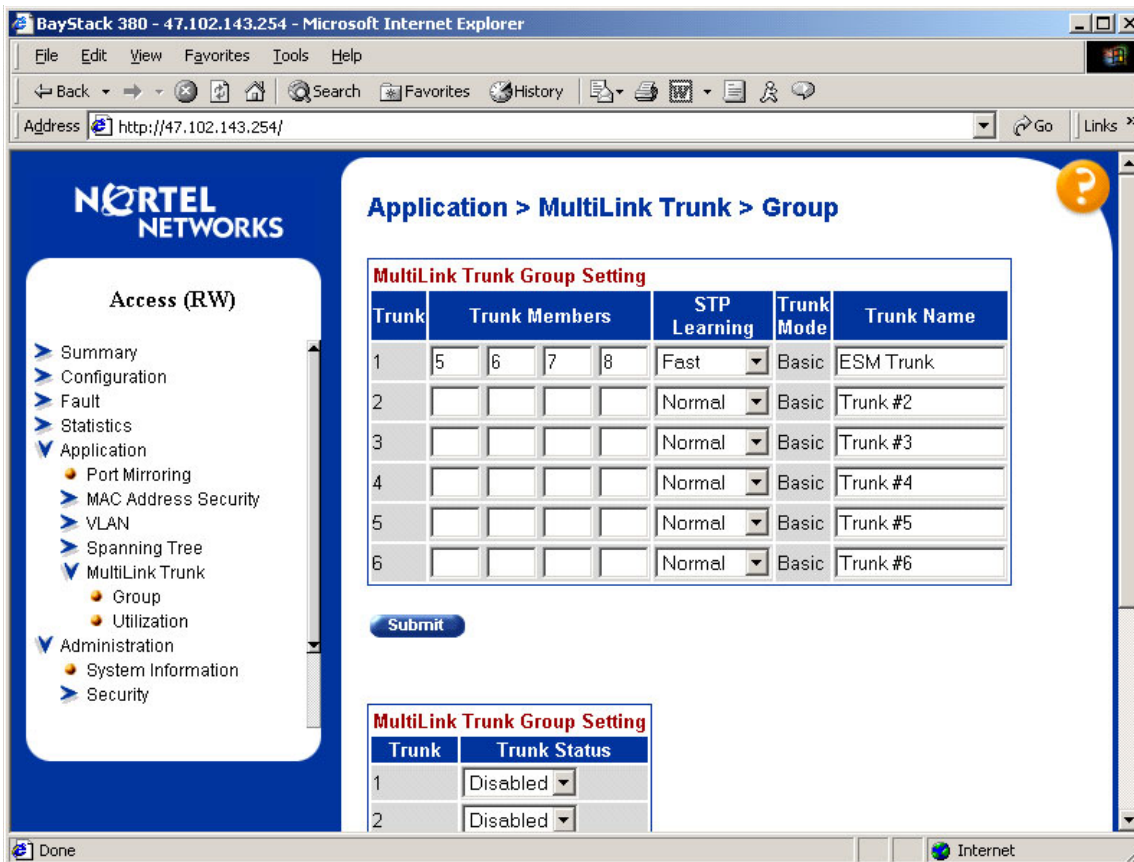


Figure 6-19 Application > MultiLink Trunk > Group window

6.5.2 Validation of ESM configuration

To validate the setup of the ESM, please refer to “Step 5: Verifying the configuration” on page 94.

6.5.3 Validation of BayStack 380-24T Configuration

This section contains some steps that can be taken to quickly verify the configuration of the BayStack 380-24T.

Check port status, speed, and duplex settings from the Configuration -> Port Management screen. Note that if autonegotiate is enabled, the Speed/Duplex box displays the negotiated speed/duplex setting. If autonegotiate is disabled, the Speed/Duplex box reflects the manual setting (Figure 6-20).

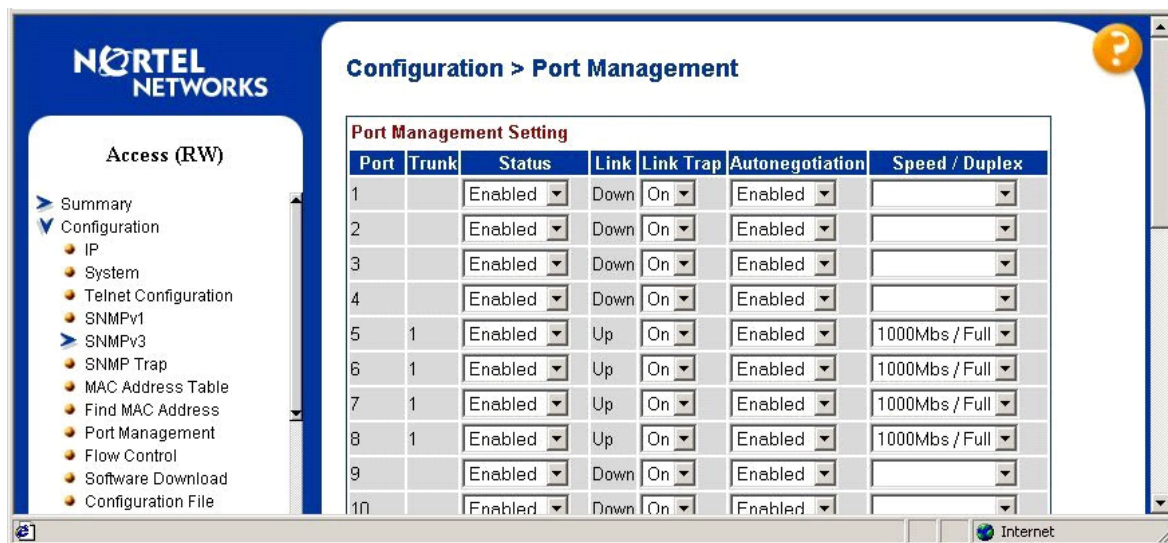


Figure 6-20 Port Management configuration window

Verify Spanning Tree forwarding state for the ports that connect to the ESM. Go to **Application -> Spanning Tree -> Port Configuration** and confirm that the trunk ports are in a forwarding state (Figure 6-21).

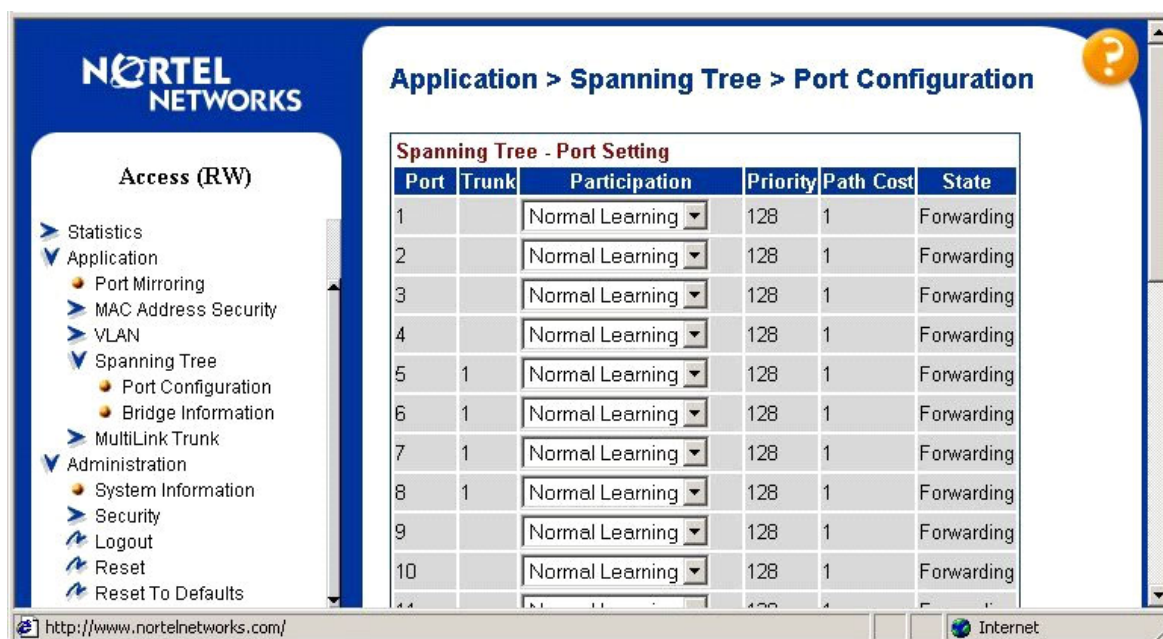


Figure 6-21 Port Configuration window

Other steps can be taken to verify the proper operation. If ESM's management VLAN is the same VLAN as the BayStack 380-24Ts management VLAN, you can ping from BayStack 380-24T to ESM. The BayStack 380-24T ping utility is only available from the CI Menu interface (console or Telnet). From the main menu, choose IP Configuration/Setup. Set the address to ping and then start the ping by choosing the corresponding option. This is shown below (Figure 6-22).

IP Configuration/Setup

BootP Request Mode: [BootP Disabled]

	Configurable	In Use	Last BootP
In-Band Switch IP Address:	[47.102.143.254]	47.102.143.254	0.0.0.0
In-Band Subnet Mask:	[255.255.255.0]	255.255.255.0	0.0.0.0
Default Gateway:	[0.0.0.0]	0.0.0.0	0.0.0.0
IP Address to Ping:	[192.168.47.250]		
Start Ping:	[Yes]		

Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu. Press Ctrl-C to return to Main Menu.

Figure 6-22 IP Configuration/Setup

Also a BayStack 380-24T port can be assigned to a particular VLAN so that a workstation can be connected to the port in order to ping the corresponding server.

Not all steps must be performed to verify proper configuration; these are just possible tests that can be used.

6.5.4 Single ESM with Link Aggregation to Dual BayStack 380-24T's

This network design offers greater protection against equipment failure, but at the cost of available throughput. Specifically, it incorporates a redundant BayStack 380, but only two Gigabit links are carrying traffic as the other two are blocked by Spanning Tree. This configuration (Figure 6-23) does not provide protection against ESM failure. It can be used when the extra throughput is not as big a concern as the possibility of a BayStack 380 failure.

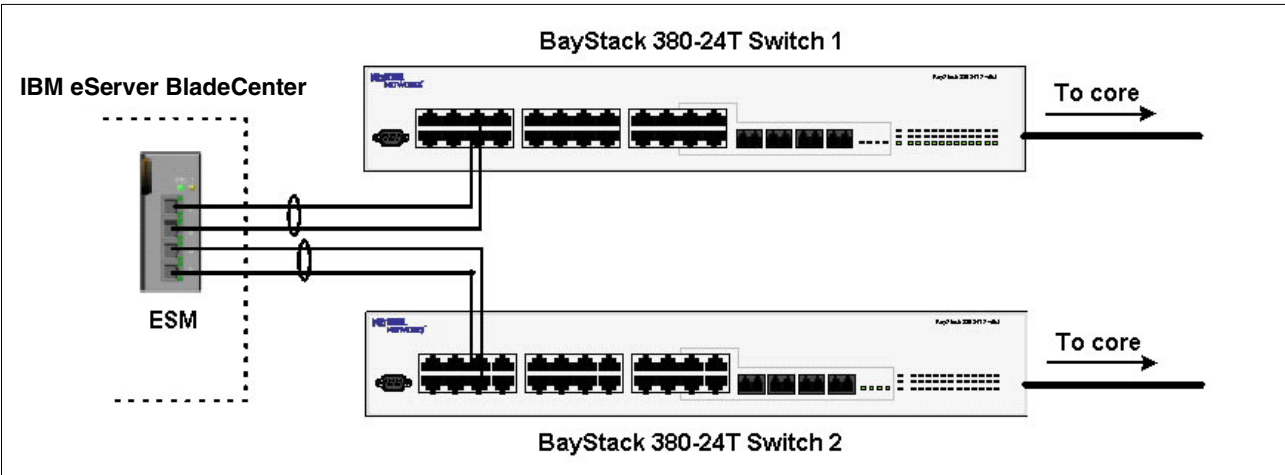


Figure 6-23 IBM eServer BladeCenter and dual BayStack 380-24T switches

Configuring the ESM

- This section walks through the sequence of actions (Figure 6-2) required to configure the ESM for this example. The following assumptions have been made for this example:
- ▶ The user is already logged into the ESM with root level access.
 - ▶ Ports EXT1 and EXT2 on the ESM in Switch Module Bay 1 are being used as the link between the @server BladeCenter and Nortel Switch 1.
 - ▶ Ports EXT3 and EXT4 on the ESM in Switch Module Bay 1 are being used as the link between the @serverBladeCenter and Nortel Switch 2.
 - ▶ Commands are being performed in the sequence shown.
 - ▶ Nortel Networks switches are BayStack 380-24T gigabit switches. Autonegotiate will result in ports operating at 1000Base-T.

Table 6-2 Configuring the ESM

Description and Comments	Description and Comments
<p>Step 1 - <i>Configure PVIDs</i></p> <p>This places the desired blade server ports into the desired VLANs. If the VLAN does not exist, it will be created automatically</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID) ▶ For Bay 2, change the PVID to 5 ▶ For Bays 3 and 4, change the PVID to 10 ▶ All other PVIDs should be set for 1 ▶ Click Apply

<p>Step 2 - <i>Configure 802.1Q trunking</i></p> <p>This will allow EXT 1 through 4 to carry traffic for VLANs 1, 5, and 10.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 1 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 1 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Select VLAN 5 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 10 and click Edit ▶ Set interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure the Tag boxes are checked ▶ Make sure Ports 3 and 4 are set for Egress and the Tag boxes are unchecked for these two ports ▶ Click Apply
<p>Step 3 - <i>Configure Link Aggregation</i></p> <p>This will allow EXT1 and 2 to be bundled together into an aggregate link, and EXT3 and 4 to be bundled together.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->Port Trunking ▶ Set Ext1 and Ext2 as members of Group 1. ▶ Set Ext3 and Ext4 as members of Group 2 ▶ Set Method to Enabled ▶ Click Apply
<p>Step 4 - Ensure LACP is Disabled.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->IEEE 802.3ad Link Aggregation->Setup IEEE 802.3ad Link Aggregation. ▶ Set LACP from Enabled to Disabled. Click Apply

Configuring the BayStack 380-24T

The configuration of Switch 2 will be similar to Switch 1. Only the configuration of Switch 1 is displayed here.

Step 1 – Configure Ports

Click **Application -> VLAN -> Port Configuration** (Figure 6-24). Configure the ports to be tagged. Port names can also be assigned. Click **Submit**.

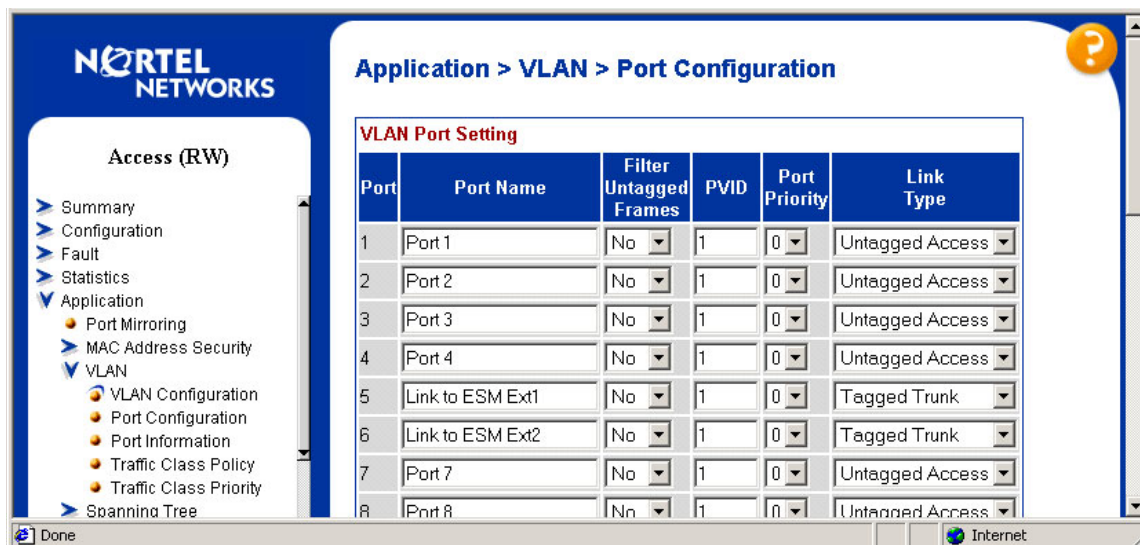


Figure 6-24 Application > VLAN > Port Configuration window

Step 2 – Add VLANs to each link member

Click **Application -> VLAN -> VLAN Configuration**. Click the action button next to VLAN 5. Add the ports to VLAN 5 as shown in Figure 6-25 and click **Submit**. Repeat this for VLAN 10 and click **Submit**.

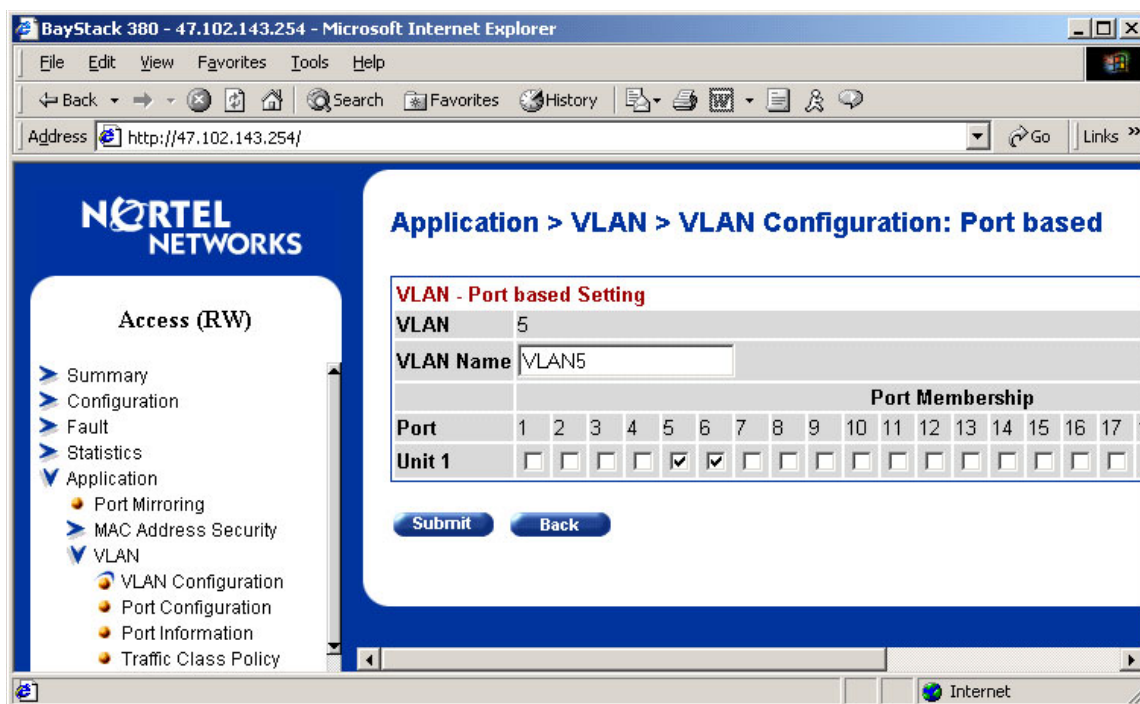


Figure 6-25 Application > VLAN > VLAN Configuration: Port based

Step 3 – Configure MLT

Click **Application -> MultiLink Trunk -> Group**. Specify ports 5 and 6, and leave STP Learning as "Normal" (Figure 6-26). Click **Submit**. The trunk is not enabled at this point. It must

explicitly be enabled by changing the setting in the table to Enabled and clicking **Submit**.

Application > MultiLink Trunk > Group

MultiLink Trunk Group Setting

Trunk	Trunk Members	STP Learning	Trunk Mode	Trunk Name
1	5 6	Normal	Basic	ESM Trunk
2		Normal	Basic	Trunk #2
3		Normal	Basic	Trunk #3
4		Normal	Basic	Trunk #4
5		Normal	Basic	Trunk #5
6		Normal	Basic	Trunk #6

Submit

Figure 6-26 Application > MultiLink Trunk > Group window

6.5.5 Validation of ESM configuration

To validate the setup of the ESM, please refer to “Step 5: Verifying the configuration” on page 101.

6.5.6 Validation of BayStack 380-24T configuration

This section contains some steps that can be taken to quickly verify the configuration of the BayStack 380-24Ts. Only Switch 1 is shown below though both BayStack switches should show the same essential information.

Check port status, speed, and duplex settings from the **Configuration -> Port Management** screen. Note that if autonegotiate is enabled, the Speed/Duplex box displays the negotiated speed/duplex setting. If autonegotiate is disabled, the Speed/Duplex box reflects the manual setting (Figure 6-27).

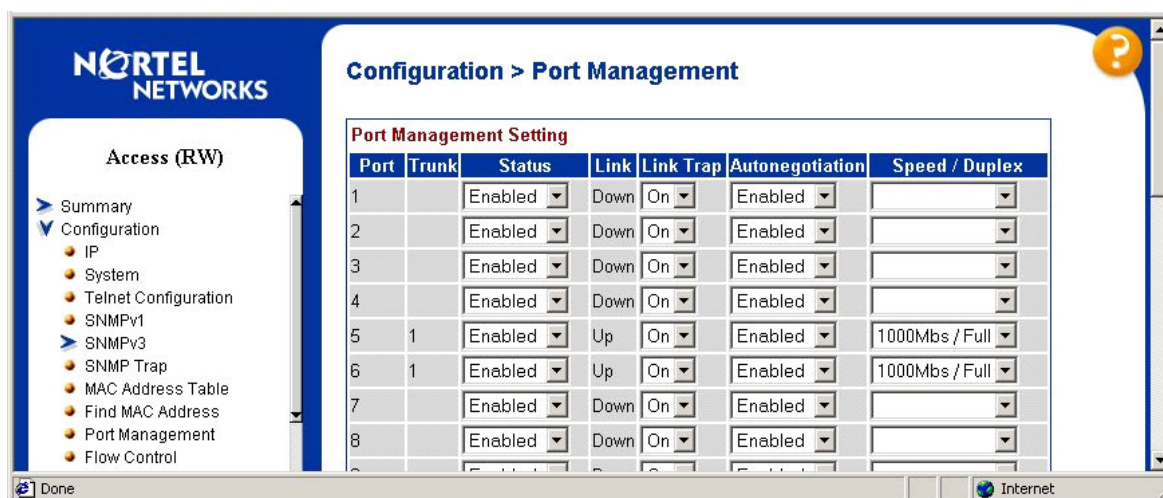


Figure 6-27 Port Management configuration window

Verify Spanning Tree forwarding state for the ports that connect to the ESM. Go to **Application -> Spanning Tree -> Port Configuration** and confirm that the trunk ports are in a forwarding state. Even though one of the ESM's link aggregation groups should show a blocking state, both BayStack switches will be in the forwarding state (Figure 6-28). This is the proper result for Spanning Tree Protocol.

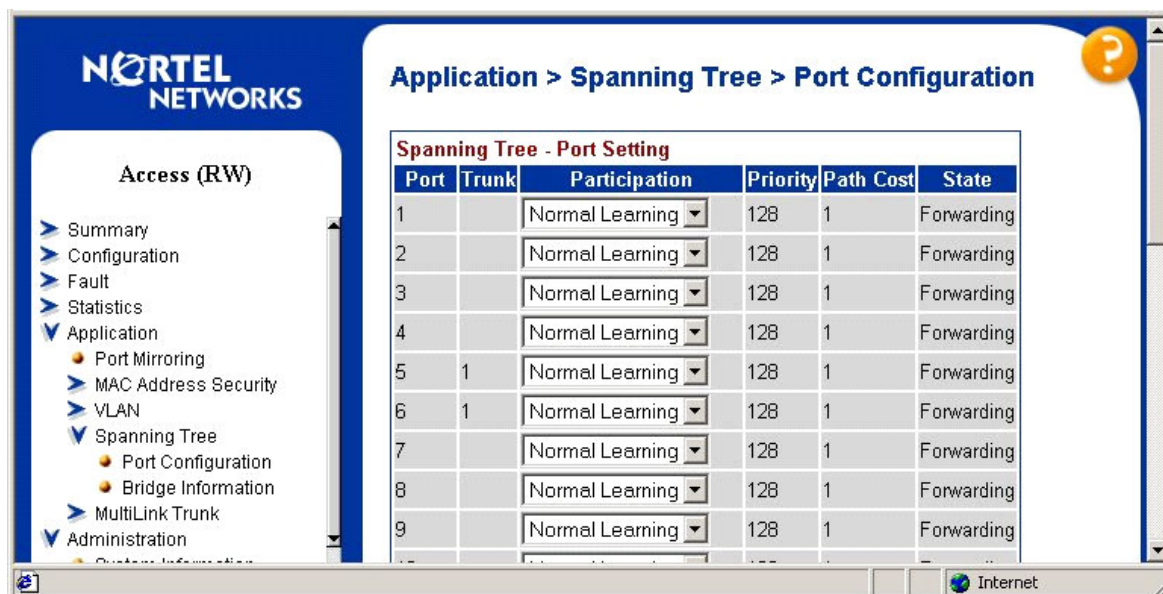


Figure 6-28 Port Configuration window

Other steps can be taken to verify the proper operation. If ESM's management VLAN is the same VLAN as the BayStack 380-24Ts management VLAN, you can ping from BayStack 380-24T to ESM. The BayStack 380-24T ping utility is only available from the CI Menu interface (console or Telnet). From the main menu, choose IP Configuration/Setup. Set the address to ping and then start the ping by choosing the corresponding option. This is shown below (Figure 6-29):

IP Configuration/Setup

BootP Request Mode: [BootP Disabled]

	Configurable	In Use	Last BootP
In-Band Switch IP Address:	[47.102.143.254]	47.102.143.254	0.0.0.0
In-Band Subnet Mask:	[255.255.255.0]	255.255.255.0	0.0.0.0
Default Gateway:	[0.0.0.0]	0.0.0.0	0.0.0.0
IP Address to Ping:	[192.168.47.250]		
Start Ping:	[Yes]		

Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu. Press Ctrl-C to return to Main Menu.

Figure 6-29 IP Configuration/ Setup window

Also, a BayStack 380-24T port can be assigned to a particular VLAN so that a workstation can be connected to the port in order to ping the corresponding server.

Not all steps must be performed to verify proper configuration; these are just possible tests that can be used.

6.5.7 Dual ESMs with Two Port Aggregation to Dual BayStack 380-24Ts

In this scenario (Figure 6-30), two ESMs and two BayStack 380s are connected by a series of two link trunks. Two of the four trunks will be blocked by Spanning Tree. Any one switch or ESM can fail as well as several ports or links, and it will still survive. This reliability is at the cost of throughput. This configuration can be used when a high degree of resiliency is required, but the maximum possible throughput is not required.

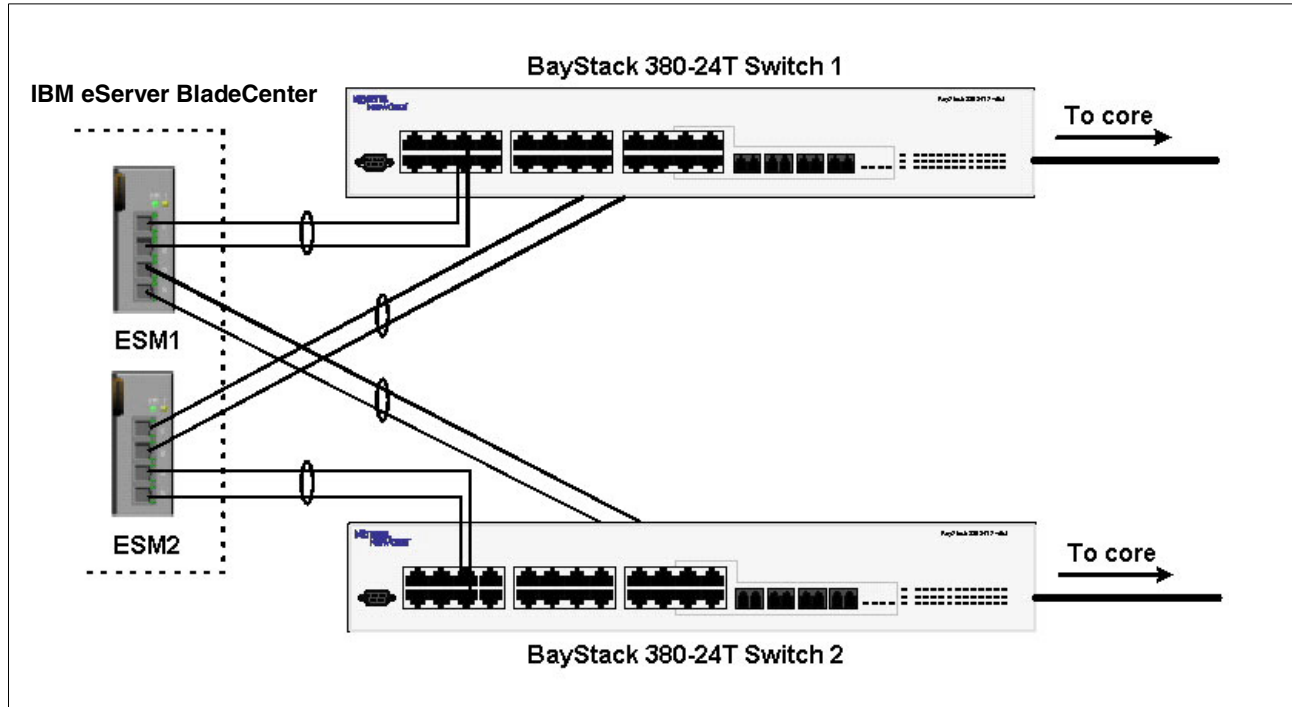


Figure 6-30 Dual ESMs and dual BayStack 380 switches

Configuring the ESM

This section walks through the sequence of actions required to configure the ESM for this example. The following assumptions have been made for this example:

- ▶ The user is already logged into the ESM with root level access.
- ▶ Ports EXT1 and EXT2 on the ESM in Switch Module Bay 1 are being used as the link between the @server BladeCenter and Nortel Switch 1.
- ▶ Ports EXT3 and EXT4 on the ESM in Switch Module Bay 1 are being used as the link between the @server BladeCenter and Nortel Switch 2.
- ▶ Ports EXT1 and EXT2 on the ESM in Switch Module Bay 2 are being used as the link between the @server BladeCenter and Nortel Switch 1.
- ▶ Ports EXT3 and EXT4 on the ESM in Switch Module Bay 1 are being used as the link between the @server BladeCenter and Nortel Switch 2.
- ▶ Commands are being performed in the sequence shown.
- ▶ Nortel Networks switches are BayStack 380-24T gigabit switches. Autonegotiate will result in ports operating at 1000Base-T.

The following steps apply to configuring the first ESM. Essentially the same steps are used to configure the second ESM as well.

Table 6-3 Configuring the ESMs

Description and Comments	Description and Comments
<p>Step 1 - <i>Configure PVIDs</i></p> <p>This places the desired blade server ports into the desired VLANs. If the VLAN does not exist, it will be created automatically</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID) ▶ For Bay 2, change the PVID to 5 ▶ For Bays 3 and 4, change the PVID to 10 ▶ All other PVIDs should be set for 1 ▶ Click Apply
<p>Step 2 - <i>Configure 802.1Q trunking</i></p> <p>This will allow EXT 1 through 4 to carry traffic for VLANs 1, 5, and 10.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 1 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 1 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Select VLAN 5 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 10 and click Edit ▶ Set interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure the Tag boxes are checked ▶ Make sure Ports 3 and 4 are set for Egress and the Tag boxes are unchecked for these two ports ▶ Click Apply
<p>Step 3 - <i>Configure Link Aggregation</i></p> <p>This will allow EXT1 and 2 to be bundled together into an aggregate link, and EXT3 and 4 to be bundled together.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->Port Trunking ▶ Set Ext1 and Ext2 as members of Group 1 ▶ Set Ext3 and Ext4 as members of Group 2 ▶ Set Method to Enabled ▶ Click Apply
<p>Step 4 - Ensure LACP is Disabled.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->IEEE 802.3ad Link Aggregation->Setup IEEE 802.3ad Link Aggregation. ▶ Set LACP from Enabled to Disabled. Click Apply

Configuring the BayStack 380-24T

The configuration of Switch 2 will be similar to Switch 1. Only the configuration of Switch 1 is displayed here.

Step 1 – Configure Ports

Click **Application** -> **VLAN** -> **Port Configuration** (Figure 6-31). Configure the ports to be tagged. Port names can also be assigned. Click **Submit**.

Port	Port Name	Filter Untagged Frames	PVID	Port Priority	Link Type
1	Port 1	No	1	0	Untagged Access
2	Port 2	No	1	0	Untagged Access
3	Port 3	No	1	0	Untagged Access
4	Port 4	No	1	0	Untagged Access
5	Link to ESM1 Ex1	No	1	0	Tagged Trunk
6	Link to ESM1 Ex2	No	1	0	Tagged Trunk
7	Link to ESM2 Ex1	No	1	0	Tagged Trunk
8	Link to ESM2 Ex2	No	1	0	Tagged Trunk
9	Port 9	No	1	0	Untagged Access

Figure 6-31 Application > VLAN > Port Configuration window

Step 2 – Add VLANs to each link member

Click **Application** -> **VLAN** -> **VLAN Configuration**. Click the action button next to VLAN 5. Add the ports to VLAN 5 as shown in Figure 6-32 and click **Submit**. Repeat this for VLAN 10 and click **Submit**.

VLAN - Port based Setting	
VLAN	5
VLAN Name	VLAN5
Port Membership	
Port	1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18
Unit 1	<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>

Figure 6-32 Application > VLAN > VLAN Configuration: Port based window

Step 3 – Configure MLT

Click **Application** -> **MultiLink Trunk** -> **Group** (Figure 6-33). Specify ports 5 and 6, and leave STP Learning as “Normal”. Click **Submit**. Next specify ports 7 and 8 as a second trunk with STP Learning as “Normal”. Click **Submit**. The trunks are not enabled at this point. Both trunks must explicitly be enabled by changing the setting in the table to Enabled and clicking **Submit**.

Trunk	Trunk Members	STP Learning	Trunk Mode	Trunk Name
1	5 6	Normal	Basic	ESM1 Trunk
2	7 8	Normal	Basic	ESM2 Trunk
3		Normal	Basic	Trunk #3
4		Normal	Basic	Trunk #4
5		Normal	Basic	Trunk #5
6		Normal	Basic	Trunk #6

Trunk	Trunk Status
1	Disabled
2	Disabled

Figure 6-33 Application > MultiLink Trunk > Group window

6.5.8 Validation of ESM configuration

To validate the setup of the ESM, please refer to “Step 5: Verifying the configuration” on page 110.

6.5.9 Validation of BayStack 380-24T configuration

This section contains some steps that can be taken to quickly verify the configuration of the BayStack 380-24Ts. Only Switch 1 is shown in Figure 6-34 although both BayStack switches should show the same essential information.

Check port status, speed, and duplex settings from the Configuration -> Port Management screen. Note that if autonegotiate is enabled, the Speed/Duplex box displays the negotiated speed/duplex setting. If autonegotiate is disabled, the Speed/Duplex box reflects the manual setting.

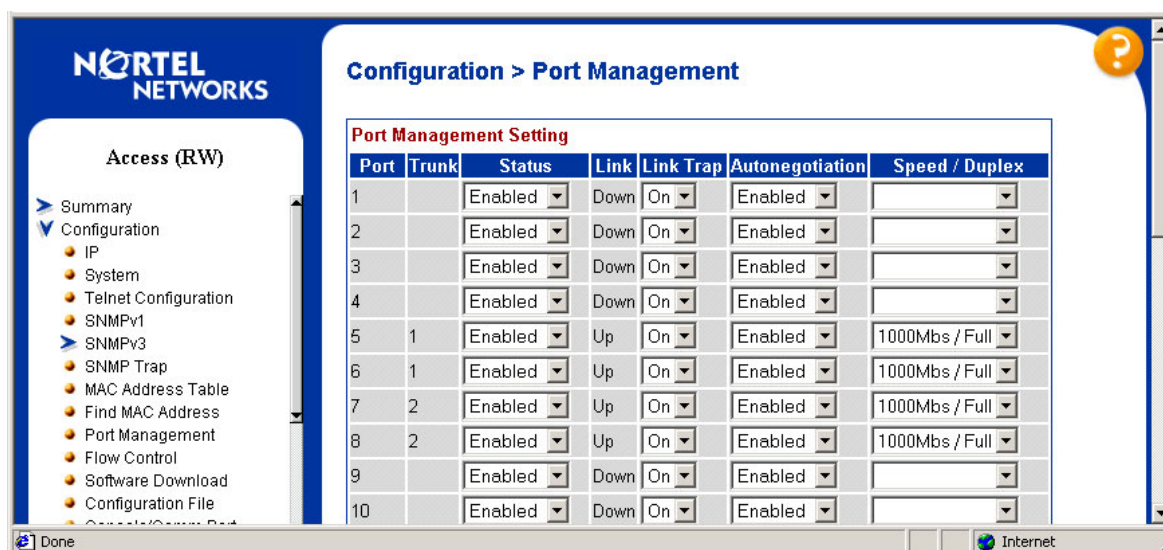


Figure 6-34 Port Management window

Verify Spanning Tree forwarding state for the ports that connect to the ESM. Go to **Application -> Spanning Tree -> Port Configuration** and confirm that the trunk ports are in a forwarding state. Even though some of the ESMs' link aggregation groups should show a blocking state, both BayStack switches will be in the forwarding state (Figure 6-35). This is the proper result for Spanning Tree Protocol.

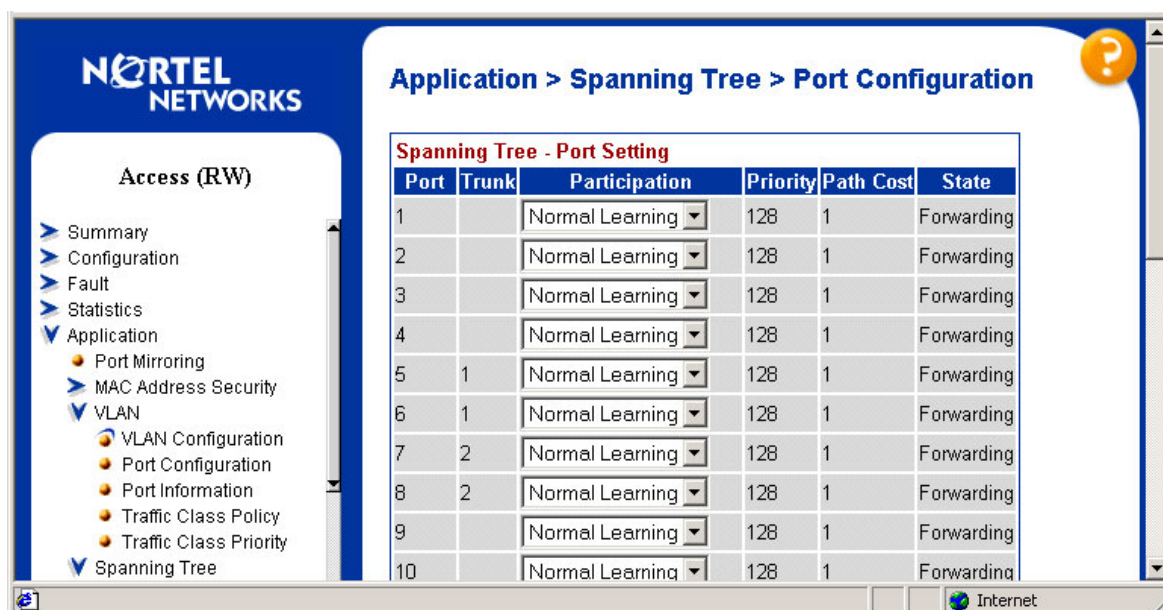


Figure 6-35 Port Configuration window

Other steps can be taken to verify the proper operation. If ESM's management VLAN is the same VLAN as the BayStack 380-24Ts management VLAN, you can ping from BayStack 380-24T to ESM. The BayStack 380-24T ping utility is only available from the CI Menu interface (console or Telnet). From the main menu, choose **IP Configuration/Setup**. Set the address to ping and then start the ping by choosing the corresponding option. This is shown in Figure 6-36.

IP Configuration/Setup			
BootP Request Mode: [BootP Disabled]			
	Configurable	In Use	Last BootP
	-----	-----	-----
In-Band Switch IP Address:	[47.102.143.254]	47.102.143.254	0.0.0.0
In-Band Subnet Mask:	[255.255.255.0]	255.255.255.0	0.0.0.0
Default Gateway:	[0.0.0.0]	0.0.0.0	0.0.0.0
IP Address to Ping:	[192.168.47.250]		
Start Ping:	[Yes]		

Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu. Press Ctrl-C to return to Main Menu.

Figure 6-36 IP Configuration/Setup window

Also, a BayStack 380-24T port can be assigned to a particular VLAN so that a workstation can be connected to the port in order to ping the corresponding server.

Not all steps must be performed to verify proper configuration; these are just possible tests that can be used.

6.5.10 Dual ESMs with Four Port Aggregation to Dual BayStack 380-24Ts

This topology (Figure 6-37) shows the maximum possible throughput for an IBM @server BladeCenter server equipped with two ESMs. It does not sacrifice very much resiliency in gaining the throughput, but rather it moves the resiliency mechanisms to the servers themselves. This configuration can be used in networks which need the highest level of performance.

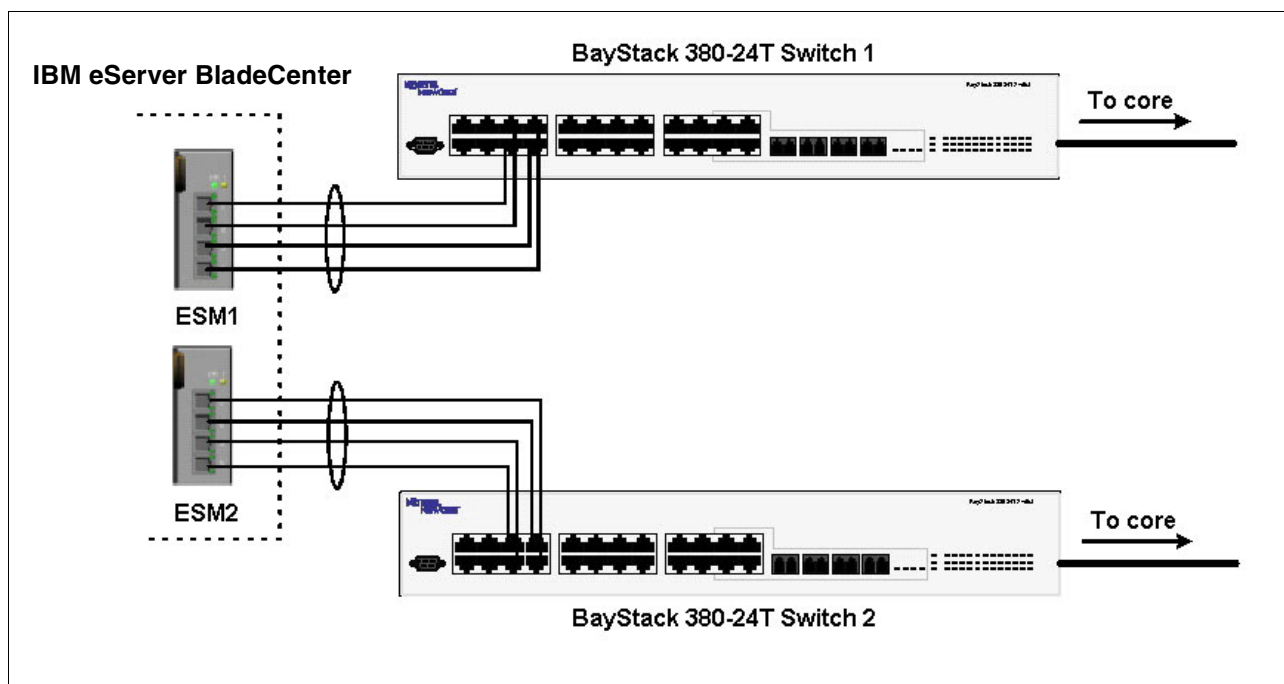


Figure 6-37 Dual ESMs and Dual BayStack 380 switches

Configuring the ESM

This section walks through the sequence of actions required to configure the ESM for this example. The following assumptions have been made for this example:

- ▶ The user is already logged into the ESM with root level access.
- ▶ Ports EXT1 through EXT4 on the ESM in Switch Module Bay 1 are being used between IBM and Nortel Switch 1.
- ▶ Ports EXT1 through EXT4 on the ESM in Switch Module Bay 2 are being used between IBM and Nortel Switch 2.
- ▶ The same VLAN setup is being used on both ESMs (this is not a requirement, the second ESM could be configured to use different VLANs)
- ▶ Commands are being performed in the sequence shown.
- ▶ Nortel Networks switches are BayStack 380-24T gigabit switches. Autonegotiate will result in ports operating at 1000Base-T.

The following steps apply to configuring the first ESM. Essentially, the same steps are used to configure the second ESM as well.

Table 6-4 Configuring the ESM

Description and Comments	Description and Comments
<p>Step 1 - <i>Configure PVIDs</i></p> <p>This places the desired blade server ports into the desired VLANs. If the VLAN does not exist, it will be created automatically</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID) ▶ For Bay 2, change the PVID to 5 ▶ For Bays 3 and 4, change the PVID to 10 ▶ All other PVIDs should be set for 1 ▶ Click Apply

<p>Step 2 - <i>Configure 802.1Q trunking</i> This will allow EXT 1 through 4 to carry traffic for VLANs 1, 5, and 10.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 1 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 1 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Select VLAN 5 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 10 and click Edit ▶ Set interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure the Tag boxes are checked ▶ Make sure Ports 3 and 4 are set for Egress and the Tag boxes are unchecked for these two ports ▶ Click Apply
<p>Step 3 - <i>Configure Link Aggregation</i> This will allow EXT 1 through 4 to be bundled together into an aggregate link.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->Port Trunking ▶ Set Ext1, Ext2, Ext3, and Ext4 as members of Group 1. ▶ Set Method to Enabled ▶ Click Apply
<p>Step 4 - Ensure LACP is Disabled.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->IEEE 802.3ad Link Aggregation->Setup IEEE 802.3ad Link Aggregation. ▶ Set LACP from Enabled to Disabled. Click Apply

Configuring the BayStack 380-24T

The configuration of Switch 2 will be similar to Switch 1. Only the configuration of Switch 1 is displayed here.

Step 1 – Configure Ports

Click **Application -> VLAN -> Port Configuration** (Figure 6-38). Configure the links to be tagged. Port names can also be assigned. Click **Submit**.

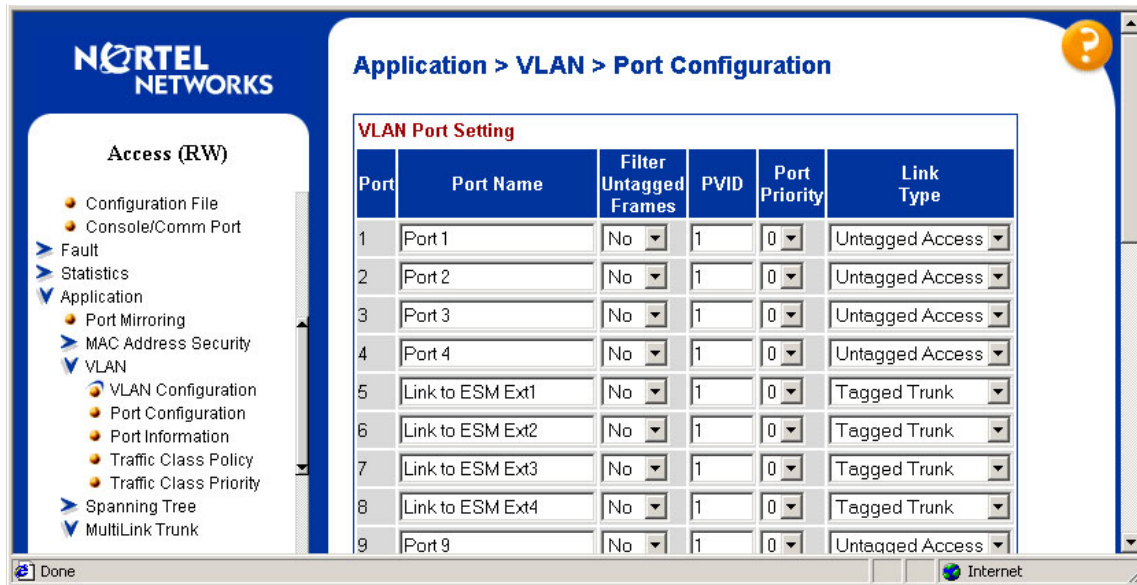


Figure 6-38 Application > VLAN > Port Configuration window

Step 2 – Add VLANs to each link member

Click **Application -> VLAN -> VLAN Configuration** (Figure 6-39). Click the action button next to VLAN 5. Add the ports to VLAN 5 as shown below and click **Submit**. Repeat this for VLAN 10 and click **Submit**.

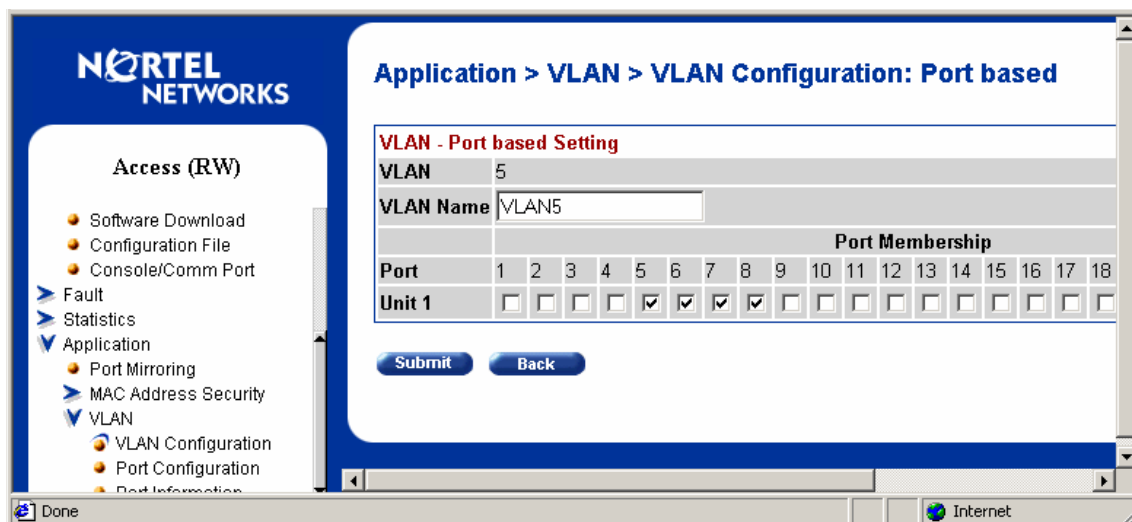


Figure 6-39 Application > VLAN > VLAN Configuration: Port based window

Step 3 – Configure MLT

Click **Application -> MultiLink Trunk -> Group** (Figure 6-40). Specify ports 5, 6, 7, and 8, and STP Learning to Fast. Click **Submit**. The trunk is not enabled at this point. It must explicitly be enabled by changing the setting in the table to Enabled and clicking **Submit**.

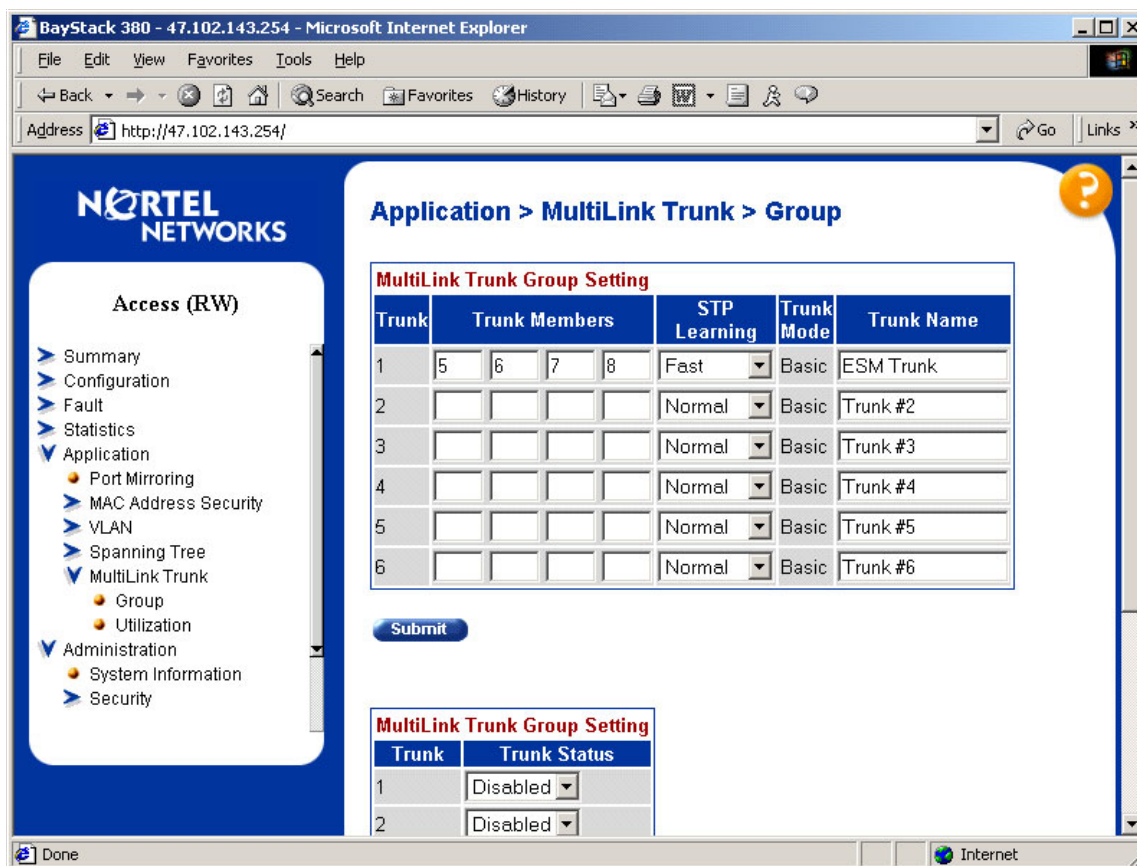


Figure 6-40 Application > MultiLink Trunk > Group window

Note: In the following sections these configuration examples were provided by Nortel Networks, however, the ITSO was unable to test and verify them during the scope of this residency.

6.5.11 Validation of ESM configuration

To validate the setup of the ESM, please refer to “Step 5: Verifying the configuration” on page 123.

6.5.12 Validation of BayStack 380-24T configuration

This section contains some steps that can be taken to quickly verify the configuration of the BayStack 380-24T. Validation should be performed on both switches, and the methods listed here are applicable to both switches.

Check port status, speed, and duplex settings from the **Configuration -> Port Management** screen. Note that if autonegotiate is enabled, the Speed/Duplex box displays the negotiated

speed/duplex setting. If autonegotiate is disabled, the Speed/Duplex box reflects the manual setting (Figure 6-41).

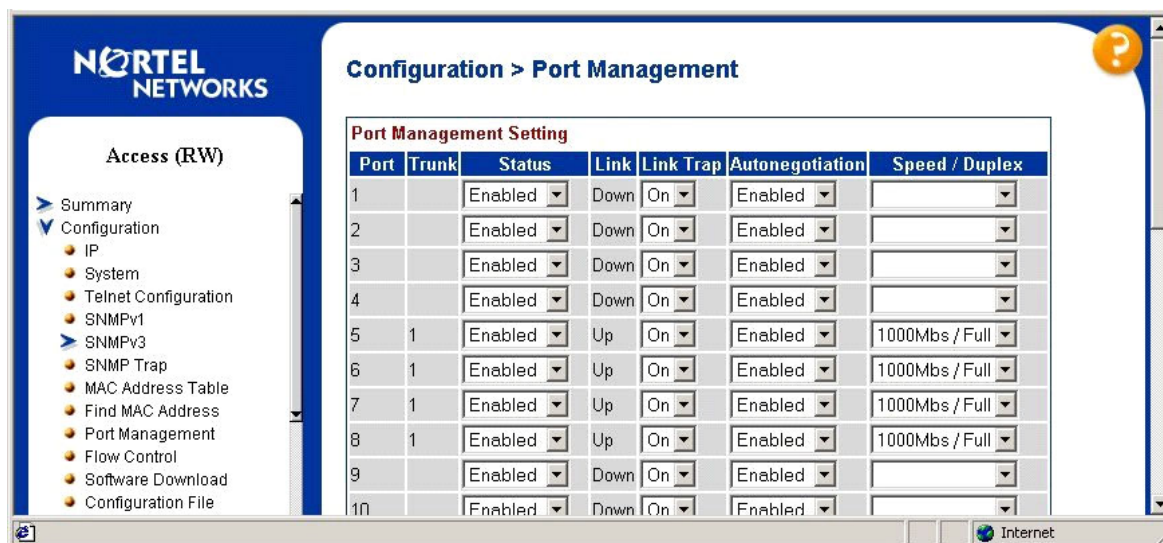


Figure 6-41 Port Management configuration window

Verify Spanning Tree forwarding state for the ports that connect to the ESM. Go to **Application -> Spanning Tree -> Port Configuration** and confirm that the trunk ports are in a forwarding state (Figure 6-42).

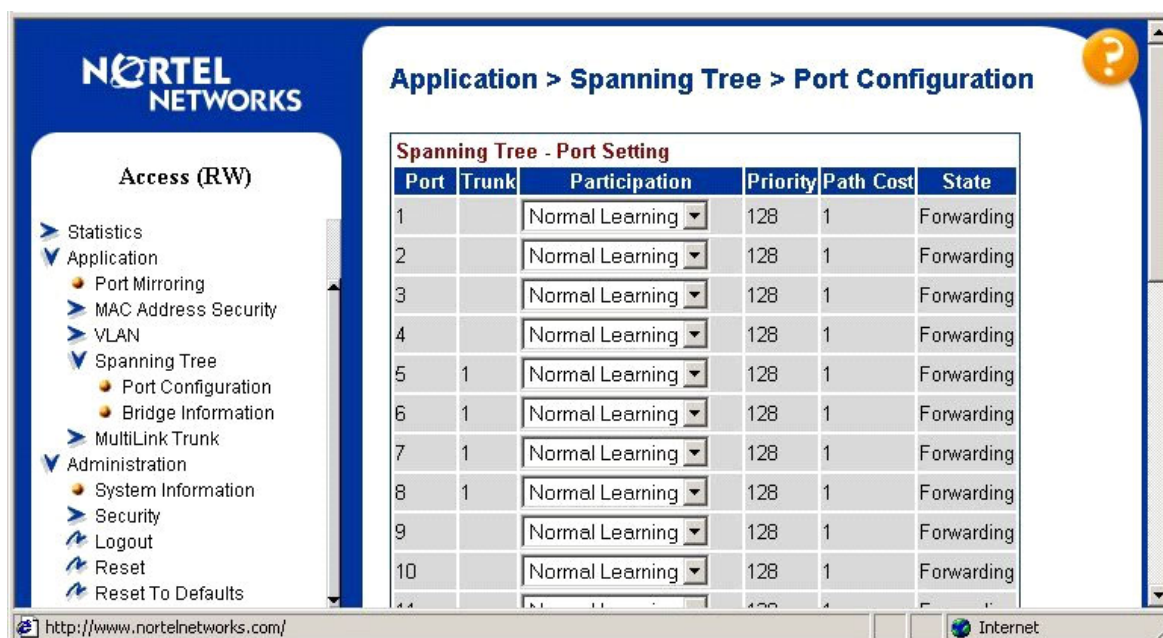


Figure 6-42 Port Configuration window

Other steps can be taken to verify the proper operation. If ESM's management VLAN is the same VLAN as the BayStack 380-24Ts management VLAN, you can ping from BayStack 380-24T to ESM. The BayStack 380-24T ping utility is only available from the CI Menu interface (console or Telnet). From the main menu, choose IP Configuration/Setup. Set the address to ping and then start the ping by choosing the corresponding option. This is shown below (Figure 6-43):

IP Configuration/Setup

BootP Request Mode: [BootP Disabled]

	Configurable	In Use	Last BootP
In-Band Switch IP Address:	[47.102.143.254]	47.102.143.254	0.0.0.0
In-Band Subnet Mask:	[255.255.255.0]	255.255.255.0	0.0.0.0
Default Gateway:	[0.0.0.0]	0.0.0.0	0.0.0.0
IP Address to Ping:	[192.168.47.250]		
Start Ping:	[Yes]		

Use space bar to display choices, press <Return> or <Enter> to select choice.
Press Ctrl-R to return to previous menu. Press Ctrl-C to return to Main Menu.

Figure 6-43 IP Configuration/Setup window

Also a BayStack 380-24T port can be assigned to a particular VLAN so that a workstation can be connected to the port in order to ping the corresponding server.

Not all steps must be performed to verify proper configuration; these are just possible tests that can be used.

6.5.13 Dual ESMs with Four Port Aggregation each to Single Passport 8600

Many network architectures include an aggregation layer. The BayStack 380-24T serves as a gigabit aggregation layer switch in many networks. By the same measure, the ESM is performing an aggregation function since it has 14 internal gigabit links and only 4 external gigabit links. Consider a BladeCenter server scenario where the 14 servers are capable of utilizing most of the four port (gigabit speed) aggregate trunk to an uplink device on a sustained basis. Now consider the expanded view of this sample network diagram if BayStack 380-24Ts are used to connect the ESMs to the network. See Figure 6-44.




Figure 6-44 Dual ESMs with four port aggregation each to single Passport 8600

Since the aggregate links between ESM and BayStack 380-24T is mostly utilized, therefore the aggregate links between the BayStack 380-24T and the Passport 8600 will be mostly utilized. In this case, it does not make much sense to use additional ports on the BayStack 380-24T to connect to other servers or end stations, otherwise oversubscription on the aggregate links between BayStack 380-24T and Passport 8600 will be likely. In this scenario, it makes much more sense to directly connect the ESM to the core switch and bypass the BayStack 380-24T altogether. Hence, Figure 6-45 illustrates the next configuration example.

In short, this design is best implemented when all links are utilized highly enough to necessitate being connected to a core switch located in the data center. No ports will be in a Spanning Tree blocking state. This scenario does not provide resiliency in the case of a Passport 8600 failure, though ESM and multiple link/port failures can be survived.




Figure 6-45 Dual ESMs with four port aggregation each to single Passport 8600

Configuring the ESM

This section walks through the sequence of actions required to configure the ESM for this example. The following assumptions have been made for this example:

- ▶ The user is already logged into the ESM with root level access.
- ▶ Ports EXT1 through EXT4 on the ESM in Switch Module Bay 1 are being used between IBM and Nortel Passport 8600 Switch 1.
- ▶ Ports EXT1 through EXT4 on the ESM in Switch Module Bay 2 are being used between IBM and Nortel Passport 8600 Switch 1.
- ▶ The same VLAN setup is being used on both ESMs (this is not a requirement; the second ESM could be configured to use different VLANs)
- ▶ Commands are being performed in the sequence shown.
- ▶ The Nortel switch is a Passport 8600. Autonegotiate will result in ports operating at 1000Base-T.

The following steps apply to configuring the first ESM. Essentially the same steps are used to configure the second ESM as well.

Table 6-5 Configuring the ESM

Description and Comments	Description and Comments
Step 1 - <i>Configure PVIDs</i> This places the desired blade server ports into the desired VLANs. If the VLAN does not exist, it will be created automatically	<ul style="list-style-type: none">▶ Click Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID)▶ For Bay 2, change the PVID to 5▶ For Bays 3 and 4, change the PVID to 10▶ All other PVIDs should be set for 1▶ Click Apply

<p>Step 2 - <i>Configure 802.1Q trunking</i></p> <p>This will allow EXT 1 through 4 to carry traffic for VLANs 1, 5, and 10.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 1 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 1 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Select VLAN 5 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 10 and click Edit ▶ Set interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure the Tag boxes are checked ▶ Make sure Ports 3 and 4 are set for Egress and the Tag boxes are unchecked for these two ports ▶ Click Apply
<p>Step 3 - <i>Configure Link Aggregation</i></p> <p>This will allow EXT 1 through 4 to be bundled together into an aggregate link.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->Port Trunking ▶ Set Ext1, Ext2, Ext3, and Ext4 as members of Group 1. ▶ Choose Disabled Method ▶ Click Apply
<p>Step 4 - <i>Ensure LACP is Disabled</i></p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->IEEE 802.3ad Link Aggregation->Setup IEEE 802.3ad Link Aggregation ▶ Set LACP from Enabled to Disabled. ▶ Click Apply

Configuring the Passport 8600

This section walks through the sequence of actions required to configure the Passport 8600 for this example.

Step 1 – Configure MLT

To do this you will need to create both MLT groups, enable tagging, and add specific ports. In this example, the 8608GT blade, which has 8 1000Base-T ports, is installed in slot 8. The first four ports are part of the first MLT group, and the last four ports are part of the second MLT group. After ports are added, VLANs are added to the MLT groups. The configuration is performed as follows:

```
PP8600-1:5#
PP8600-1:5# config mlt 1
PP8600-1:5/config/mlt/1# create
PP8600-1:5/config/mlt/1# name "ESM1 Trunk"
PP8600-1:5/config/mlt/1# perform-tagging enable
PP8600-1:5/config/mlt/1# add ports 8/1-8/4
PP8600-1:5/config/mlt/1# add vlan 1
PP8600-1:5/config/mlt/1# add vlan 5
PP8600-1:5/config/mlt/1# add vlan 10
PP8600-1:5/config/mlt/1# top
PP8600-1:5#
PP8600-1:5# config mlt 2
PP8600-1:5/config/mlt/2# create
PP8600-1:5/config/mlt/2# name "ESM2 Trunk"
PP8600-1:5/config/mlt/2# perform-tagging enable
PP8600-1:5/config/mlt/2# add ports 8/5-8/8
PP8600-1:5/config/mlt/2# add vlan 1
PP8600-1:5/config/mlt/2# add vlan 5
PP8600-1:5/config/mlt/2# add vlan 10
PP8600-1:5/config/mlt/2# top
PP8600-1:5#
```

Figure 6-46 Console interface

In this configuration Spanning Tree is running on the trunks to the two ESMs. Neither will end up in the blocking state since there is no loop. So you may additionally enable Faststart for the interfaces without consequence. To do this, use:

```
config ethernet <portlist> stg 1 faststart enable
```

6.5.14 Validation of ESM configuration

To validate the setup of the ESM, please refer to “Step 5: Verifying the configuration” on page 137.

6.5.15 Validation of Passport 8600 Configuration

This section contains some steps that can be taken to quickly verify the configuration of the Passport 8600.

Check port status, speed, and duplex settings using the following command:

```
show ports info name <portlist>
```

If Autonegotiate is enabled, the information shown will be the negotiated speed/duplex settings. In Figure 6-47 is an example output of the command given the configuration provided in this scenario:

```
PP8600-1:5#
PP8600-1:5# show ports info name 8/1-8/8
```

Port Name						
PORT NUM	NAME	DESCRIPTION	OPERATE STATUS	OPERATE DUPLX	OPERATE SPEED	VLAN
8/1		1000BaseT	up	full	1000	Tagged
8/2		1000BaseT	up	full	1000	Tagged
8/3		1000BaseT	up	full	1000	Tagged
8/4		1000BaseT	up	full	1000	Tagged
8/5		1000BaseT	up	full	1000	Tagged
8/6		1000BaseT	up	full	1000	Tagged
8/7		1000BaseT	up	full	1000	Tagged
8/8		1000BaseT	up	full	1000	Tagged

```
PP8600-1:5#
```

Figure 6-47 Command output

Verify Spanning Tree forwarding state for the ports that connect to the ESMs. To do this, the `show port info stg main <portlist>`

command will be used. Confirm that the trunk ports are all in the forwarding state. The results of this command should look like the following (Figure 6-48):

```
PP8600-1:5#
PP8600-1:5# show port info stg main 8/1-8/8
```

Port Stg								
SID	PORT_NUM	PRIO	STATE	ENABLE STP	FASTSTART	PATHCOST	FORWARD TRANSITION	CHANGE DETECTION
1	8/1	128	forwarding	true	true	1	2	true
1	8/2	128	forwarding	true	true	1	2	true
1	8/3	128	forwarding	true	true	1	2	true
1	8/4	128	forwarding	true	true	1	2	true
1	8/5	128	forwarding	true	true	1	2	true
1	8/6	128	forwarding	true	true	1	2	true
1	8/7	128	forwarding	true	true	1	2	true
1	8/8	128	forwarding	true	true	1	2	true

```
PP8600-1:5#
```

Figure 6-48 Command output

Other steps can be taken to verify the proper operation. For example the Passport 8600's out-of-band management interface can be connected to a port that is a member of VLAN 1 (assuming the ESM's management VLAN is 1). If this is done, it is possible to ping from Passport 8600 to ESM. To ping from the Passport 8600 management interface to ESM use the following command:

```
ping <ip_addr>
```

An example is shown below (Figure 6-49):

```
PP8600-1:5#  
PP8600-1:5# ping 192.168.47.250  
192.168.47.250 is alive  
PP8600-1:5#
```

Figure 6-49 Command output

Also a Passport 8600 port can be assigned to a particular VLAN so that a workstation can be connected to the port in order to ping the corresponding server.

Not all steps must be performed to verify proper configuration; these are just possible tests that can be used.

6.5.16 Dual ESMs with Four Port SMLT to Dual Passport 8600s

This last configuration best meets the needs of a mission critical data center where 24/7 uptime is a requirement. Several advantages exist in this scenario, the first advantage is that by using SMLT there is no need to use Spanning Tree to block loops since SMLT inherently removes any loops, thus all ports are forwarding all the time for maximum throughput. The second advantage is subsecond failover, since the two Passports share a common forwarding database either Passport 8600 could fail or any link pairs could fail, and failover would occur in less than a second. If an ESM fails, resiliency would be provided by the redundancy mechanism of the BladeCenter servers.

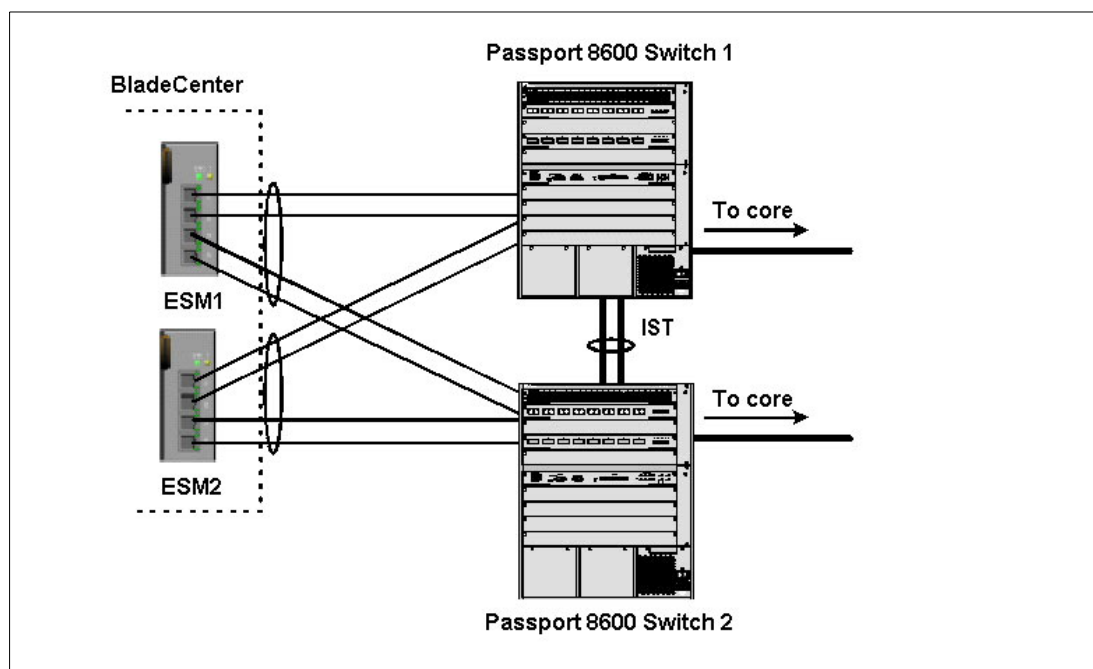


Figure 6-50 Dual ESMs with four port SMLT to dual Passport 8600s

Configuring the ESM

This section walks through the sequence of actions required to configure the ESM for this example. The following assumptions have been made for this example:

- ▶ The user is already logged into the ESM with root level access.
- ▶ Ports EXT1 and EXT2 on the ESM in Switch Module Bay 1 are connected to ports 8/1 and 8/2 on Passport 8600 Switch 1.
- ▶ Ports EXT3 and EXT4 on the ESM in Switch Module Bay 1 are connected to ports 2/1 and 2/2 on Passport 8600 Switch 2.
- ▶ Ports EXT1 and EXT2 on the ESM in Switch Module Bay 2 are connected to ports 8/3 and 8/4 on Passport 8600 Switch 1.
- ▶ Ports EXT3 and EXT4 on the ESM in Switch Module Bay 1 are connected to ports 2/3 and 2/4 on Passport 8600 Switch 2.
- ▶ The same VLAN setup is being used on both ESMs (this is not a requirement, the second ESM could be configured to use different VLANs)
- ▶ Commands are being performed in the sequence shown.
- ▶ The Nortel switches are Passport 8600s with 1000Base-T ports. Autonegotiate will result in ports operating at 1000Base-T.

The following steps apply to configuring the first ESM. Essentially the same steps are used to configure the second ESM as well.

Table 6-6 Configuring the ESM

Description and Comments	Description and Comments
<p>Step 1 - Configure PVIDs</p> <p>This places the desired blade server ports into the desired VLANs. If the VLAN does not exist, it will be created automatically</p>	<ul style="list-style-type: none"> ▶ Configuration->VLANs->802.1Q Port Settings->Port VLAN ID (PVID) ▶ For Bay 2, change the PVID to 5 ▶ For Bays 3 and 4, change the PVID to 10 ▶ All other PVIDs should be set for 1 ▶ Click Apply
<p>Step 2 - Configure 802.1Q trunking</p> <p>This will allow EXT 1 through 4 to carry traffic for VLANs 1, 5, and 10.</p>	<ul style="list-style-type: none"> ▶ Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 1 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 1 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Select VLAN 5 and click Edit ▶ Set Interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure Tag boxes are checked ▶ Make sure Port 2 is set for Egress and the Tag box is unchecked ▶ Click Apply ▶ Click Configuration->VLANs->Edit 802.1Q VLANs ▶ Select VLAN 10 and click Edit ▶ Set interface Ext1, Ext2, Ext3, and Ext4 to Egress and make sure the Tag boxes are checked ▶ Make sure Ports 3 and 4 are set for Egress and the Tag boxes are unchecked for these two ports ▶ Click Apply
<p>Step 3 – Disable Spanning Tree Protocol</p> <p>It is a requirement for SMLT Clients that Spanning Tree Protocol is disabled.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Spanning Tree->STP Port Settings. ▶ Set STP Status to Disabled for Ext1, Ext2, Ext3, and Ext4 ▶ Click Apply
<p>Step 4 - Configure Link Aggregation</p> <p>This will allow EXT 1 through 4 to be bundled together into an aggregate link.</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->Port Trunking ▶ Set Ext1, Ext2, Ext3, and Ext4 as members of Group 1. ▶ Choose Disabled Method ▶ Click Apply
<p>Step 5 - Ensure LACP is Disabled</p>	<ul style="list-style-type: none"> ▶ Click Configuration->Link Aggregation->IEEE 802.3ad Link Aggregation->Setup IEEE 802.3ad Link Aggregation ▶ Set LACP from Enabled to Disabled. ▶ Click Apply

Configuring the Passport 8600

This section walks through the sequence of actions required to configure the two Passport 8600 switches for this example. The following assumptions have been made for this example:

- ▶ Ports 1/1, 1/2, 1/3, 1/4 on Passport 8600 Switch 1 are connected to ports 1/1, 1/2, 1/3, 1/4 on Passport 8600 Switch 2. These links will be the IST trunk.
- ▶ Ports EXT1 and EXT2 on the ESM in Switch Module Bay 1 are connected to ports 8/1 and 8/2 on Passport 8600 Switch 1.
- ▶ Ports EXT3 and EXT4 on the ESM in Switch Module Bay 1 are connected to ports 2/1 and 2/2 on Passport 8600 Switch 2.
- ▶ Ports EXT1 and EXT2 on the ESM in Switch Module Bay 2 are connected to ports 8/3 and 8/4 on Passport 8600 Switch 1.
- ▶ Ports EXT3 and EXT4 on the ESM in Switch Module Bay 1 are connected to ports 2/3 and 2/4 on Passport 8600 Switch 2.
- ▶ Commands are being performed in the sequence shown. Be sure to note which switch the commands are being applied to.

Step 1 – Configure the IST on Passport 8600 Switch 1

To do this you will need to first configure a separate VLAN for IST communications. This VLAN and its subnet should remain local to the IST link.

An MLT group needs to be created, tagging enabled, and specific ports added. In this example, an 8608GT blade, which has 8 1000Base-T ports, is installed in slot 1 on Passport 8600 Switch 1 (the same blade is also installed in slot one of Passport 8600 Switch 2). Note: The first four ports of slot 1 on both Passport 8600s are connected to form the IST. This is strictly for purposes of keeping this example easy to follow. In a real network design, it is more robust to place IST links on different blades in the chassis to be able to survive an entire blade failure.

After this, VLANs need to be added and the MLT designated as an IST. Lastly the CP-limit feature is disabled on the IST links. These commands are shown in Figure 6-51.


```

PP8600-1:5#
PP8600-1:5# config vlan 1000 create
PP8600-1:5/config/vlan/1000/create# byport 1
PP8600-1:5/config/vlan/1000/create# name IST
PP8600-1:5/config/vlan/1000# ip create 10.1.1.1/24
PP8600-1:5/config/vlan/1000/create# top
PP8600-1:5#
PP8600-1:5# config mlt 1
PP8600-1:5/config/mlt/1# create
PP8600-1:5/config/mlt/1# name "IST"
PP8600-1:5/config/mlt/1# perform-tagging enable
PP8600-1:5/config/mlt/1# add ports 1/1-1/4
PP8600-1:5/config/mlt/1# add vlan 1000
PP8600-1:5/config/mlt/1# add vlan 1
PP8600-1:5/config/mlt/1# add vlan 5
PP8600-1:5/config/mlt/1# add vlan 10
PP8600-1:5/config/mlt/1# ist
PP8600-1:5/config/mlt/1/ist# create ip 10.1.1.2 vlan-id 1000

INFO : IST is created and enabled.
       The spanning tree protocol is disabled on the port(s) with IST enabled!

PP8600-1:5/config/mlt/1/ist# top
PP8600-1:5# config ethernet 1/1-1/4 cp-limit disable
PP8600-1:5#

```

Figure 6-51 Console interface

Step 2 – Configure the IST on Passport 8600 Switch 2

This step is much like the first step except you will now be configuring the IST on the second switch. The same notes and caveats from step 1 apply to this step too. The commands to configure Passport 8600 Switch 2 are shown in Figure 6-52:

```

PP8600-2:3#
PP8600-2:3# conf vlan 1000 create
PP8600-2:3/config/vlan/1000/create# byport 1
PP8600-2:3/config/vlan/1000/create# name IST
PP8600-2:3/config/vlan/1000/create# ip create 10.1.1.2/24
PP8600-2:3/config/vlan/1000/create# top
PP8600-2:3#
PP8600-2:3# conf mlt 1
PP8600-2:3/config/mlt/1# create
PP8600-2:3/config/mlt/1# name "IST"
PP8600-2:3/config/mlt/1# perform-tagging enable
PP8600-2:3/config/mlt/1# add ports 1/1-1/4
PP8600-2:3/config/mlt/1# add vlan 1000
PP8600-2:3/config/mlt/1# add vlan 1
PP8600-2:3/config/mlt/1# add vlan 5
PP8600-2:3/config/mlt/1# add vlan 10
PP8600-2:3/config/mlt/1# ist
PP8600-2:3/config/mlt/1/ist# create ip 10.1.1.1 vlan-id 1000

INFO : IST is created and enabled.
       The spanning tree protocol is disabled on the port(s) with IST enabled!

PP8600-2:3/config/mlt/1/ist# top
PP8600-2:3# config ethernet 1/1-1/4 cp-limit disable
PP8600-2:3#

```

Figure 6-52 Console interface

Step 3 – Configure the SMLT trunks on Passport 8600 Switch 1

To do this you will need to create two MLT groups, enable tagging, and add specific ports. In this example, the 8608GT blade, which is has 8 1000Base-T ports, is installed in slot 8 of

Passport 8600 Switch 1. The first two ports, 8/1 and 8/2, are part of the first SMLT group, and the next two ports, 8/3 and 8/4 are part of the second SMLT group. After ports are added, VLANs are added to the SMLT groups. Lastl, the trunk is designated as an SMLT group. The configuration is performed as follows.

```
PP8600-1:5#
PP8600-1:5# config mlt 2
PP8600-1:5/config/mlt/2# create
PP8600-1:5/config/mlt/2# name "ESM1 Trunk"
PP8600-1:5/config/mlt/2# perform-tagging enable
PP8600-1:5/config/mlt/2# add ports 8/1-8/2
PP8600-1:5/config/mlt/2# add vlan 1
PP8600-1:5/config/mlt/2# add vlan 5
PP8600-1:5/config/mlt/2# add vlan 10
PP8600-1:5/config/mlt/2# smlt create smlt-id 1

INFO : The spanning tree protocol is disabled on the port(s)
       while configure the port(s) with SMLT!

PP8600-1:5/config/mlt/2# top
PP8600-1:5#
PP8600-1:5# config mlt 3
PP8600-1:5/config/mlt/3# create
PP8600-1:5/config/mlt/3# name "ESM2 Trunk"
PP8600-1:5/config/mlt/3# perform-tagging enable
PP8600-1:5/config/mlt/3# add ports 8/3-8/4
PP8600-1:5/config/mlt/3# add vlan 1
PP8600-1:5/config/mlt/3# add vlan 5
PP8600-1:5/config/mlt/3# add vlan 10
PP8600-1:5/config/mlt/3# smlt create smlt-id 2

INFO : The spanning tree protocol is disabled on the port(s)
       while configure the port(s) with SMLT!

PP8600-1:5/config/mlt/3# top
PP8600-1:5#
```

Figure 6-53 Console interface

Step 4 – Configure the SMLT trunks on Passport 8600 Switch 2

To do this you will need to create two MLT groups, enable tagging, and add specific ports. In this example, the 8608GT blade, which is has 8 1000Base-T ports, is installed in slot 2 of Passport 8600 Switch 2. The first two ports, 2/1 and 2/2, are part of the first SMLT group, and the next two ports, 2/3 and 2/4 are part of the second SMLT group. After ports are added, VLANs are added to the SMLT groups. Lastly the trunk is designated as an SMLT group. The configuration is performed as follows.

```

PP8600-2:3#
PP8600-2:3# config mlt 2
PP8600-2:3/config/mlt/2# create
PP8600-2:3/config/mlt/2# name "ESM1 Trunk"
PP8600-2:3/config/mlt/2# perform-tagging enable
PP8600-2:3/config/mlt/2# add ports 2/1-2/2
PP8600-2:3/config/mlt/2# add vlan 1
PP8600-2:3/config/mlt/2# add vlan 5
PP8600-2:3/config/mlt/2# add vlan 10
PP8600-2:3/config/mlt/2# smlt create smlt-id 1

INFO : The spanning tree protocol is disabled on the port(s)
       while configure the port(s) with SMLT!

PP8600-2:3/config/mlt/2# top
PP8600-2:3#
PP8600-2:3# config mlt 3
PP8600-2:3/config/mlt/3# create
PP8600-2:3/config/mlt/3# name "ESM2 Trunk"
PP8600-2:3/config/mlt/3# perform-tagging enable
PP8600-2:3/config/mlt/3# add ports 2/3-2/4
PP8600-2:3/config/mlt/3# add vlan 1
PP8600-2:3/config/mlt/3# add vlan 5
PP8600-2:3/config/mlt/3# add vlan 10
PP8600-2:3/config/mlt/3# smlt create smlt-id 2

INFO : The spanning tree protocol is disabled on the port(s)
       while configure the port(s) with SMLT!

PP8600-2:3/config/mlt/3# top
PP8600-2:3#

```

Figure 6-54 Console interface

6.5.17 Validation of ESM configuration

To validate the setup of the ESM, please refer to “Step 5: Verifying the configuration” on page 154.

6.5.18 Validation of Passport 8600 Configuration

This section contains some steps that can be taken to quickly verify the configuration of the SMLT setup of the two Passport 8600s. The same basic verification steps discussed in the previous scenario apply as follow as before. Note that Spanning Tree was disabled on all SMLT and IST trunk groups, so it may not be necessary to check for forwarding state on the Passport 8600.

Check port status, speed, and duplex settings:

```
show ports info name <portlist>
```

If autonegotiate is enabled, the information shown will be the negotiated speed/duplex settings. For sample output refer back to the previous scenario.

Check the status of the IST trunk using the following command:

```
show mlt ist info
```

This will confirm that the IST trunk is up and receiving heartbeat messages. Sample output of what you should see is shown next (Figure 6-55).

```
PP8600-1:5#
PP8600-1:5# show mlt ist info
```

```
=====
                        Mlt IST Info
=====
```

MLT ID	IP ADDRESS	VLAN ID	ENABLE IST	IST STATUS
1	10.1.1.2	1000	true	up

```
=====
PP8600-1:5#
```

Figure 6-55 Command output

More detailed information on the IST will be shown with the following command:

```
show mlt ist stat
```

The status of the SMLT can be verified using the

```
show mlt smlt info
```

command. You should see the “current type” field match the “admin type” field. In the case of SMLT, both should have a value of “smlt”. A working configuration is shown below (Figure 6-56).

```
PP8600-1:5#
PP8600-1:5# show mlt smlt info
```

```
=====
                        Mlt SMLT Info
=====
```

MLT ID	SMLT ID	ADMIN TYPE	CURRENT TYPE
2	1	smlt	smlt
3	2	smlt	smlt

```
=====
PP8600-1:5#
```

Figure 6-56 Command output

Other steps can be taken to verify the proper operation. For example the Passport 8600's out-of-band management interface can be connected to a port that is a member of VLAN 1 (assuming the ESM's management VLAN is 1). If this is done, it is possible to ping from Passport 8600 to ESM. To ping from the Passport 8600 management interface to ESM use the following command:

```
ping <ip_addr>
```

Also a Passport 8600 port can be assigned to a particular VLAN so that a workstation can be connected to the port in order to ping the corresponding server. To fully ensure that the SMLT is working under failure conditions, start a continuous ping from workstation to BladeCenter server, and begin disconnecting cables between Passport 8600 switches and an ESM. As long as at least one IST link is up and as long as at least one link between ESM and either of the two Passports is up, pings should continue to get through. At most one ping should drop, but in many cases not even a single ping will drop.

Not all steps must be performed to verify proper configuration; these are just possible tests that can be used.

6.6 Troubleshooting ESM Connections to Nortel Networks Devices

Troubleshooting a network can range from a simple problem resolution to a lengthy process. Covering all possible steps to isolate and fix problems is beyond the scope of this document. The purpose of this section is to give some helpful tips to solving some of the more obvious problems. Other tools and documents exist that will aid in troubleshooting both IBM and Nortel Networks equipment. This document should not be treated as a substitute for training and experience. It is recommended that problems that are not easily resolved by the suggestions contained in this document be directed to individuals with experience and training in the areas that are experiencing issues.

Issue: Link is down

Link lights are off, or administration interface to either device is showing the link as down.

Resolution:

Verify that autonegotiate is enabled on both ends of the link. If it is not, or if manual setting of speed/duplex is required for some reason, verify that both sides are configured the same way and that a crossover cable is being used. Also check to make sure that ports are not administratively disabled on either device.

Issue: Port Aggregation does not work

Links are up, but Spanning Tree is blocking the ports, or other similar symptoms

Resolution:

Make sure that the ESM 802.3ad link aggregation groups are configured to operate in static mode, and that LACP is not operational.

Issue: Traffic for a specific VLAN or all VLANs will not pass between ESM and Nortel device

Links are up, but pings between a workstation and BladeCenter server fail.

Resolution:

Ensure proper VLAN configurations. Make sure that VLANs are added to the correct trunk groups. Make sure that tagging is enabled for all VLANs on the ESM and on the Nortel device.

Glossary

802.1ad. IEEE definition of Link Aggregation. Now incorporated into IEEE 802.3-2002

802.1D. IEEE standard describing the function of MAC bridges and switches, including the Spanning Tree Protocol

802.1Q. IEEE standard describing VLAN tagging and priority tagging

Aging Time. the length of time that switch waits before deleting a MAC address from its forwarding table if no frames originate from that MAC address

Backpressure. special sequence of bits generated by a receiver to prevent the transmitter from sending frames in half-duplex connection

BPDU. Bridge Protocol Data Unit - message sent by switches for establishing and maintaining the spanning tree

CatOS. Catalyst Operating System common on some Cisco switches

Cut-through. mode of forwarding frames with no buffering by the switch. Can only be used when all links are at the same speed

Egress port. a port where traffic exits the ESM

ESM. Ethernet Switch Module for the @serverBladeCenter

Full-Duplex. type of switch connection whereby two devices can transmit simultaneously

GVRP. GARP VLAN Registration Protocol - provides 802.1Q compliant VLAN pruning and dynamic VLAN creation

Half-Duplex. type of switch connection whereby only one device can transmit at any time

Ingress port. a port where traffic enters the ESM

IOS. Cisco Internetworking Operating System common on most Cisco routers and some Cisco switches

LACP. Link Aggregation Control Protocol - creates dynamic link aggregates using protocol information that is exchanged between switches

Link Aggregation. Combining multiple links into a single high-bandwidth virtual link. Accomplished through either Port Trunking or LACP

MM. Management module for the @serverBladeCenter

Native VLAN. Non-tagged VLAN on an 802.1Q trunk

NVRAM nonvolatile random access memory is a form of static random access memory whose contents are saved when a computer is turned off or loses its external power source.

Port Mirroring. Configuring two ports such that frames transmitted and received from one port are copied to another port. On the mirroring (destination) port, data is only sent, not received. Intended to be used for network troubleshooting and analysis

Port Trunking. Manually configuring a static link aggregate where no protocol information is exchanged between switches

PVID. Port VLAN Identifier- VID assigned to a switch port for tagging untagged frames that arrive at that port

Spanning Tree. Prevents loops in a network by blocking redundant links between switches, unless said links are configured as a link aggregate

Store and Forward. mode of forwarding where a switch buffers received frames before transmitting. Required when transmitting on links of different speeds, or when frames must be modified by the switch

Tagging Inserting VLAN and priority information into an Ethernet frame

Untagging. Removing tag header from an 802.1Q Ethernet frame. Necessary when downstream device does not support tagging

VID. VLAN Identifier - information in a tag header to identify what VLAN a frame belongs in

VLAN Trunking. configuring multiple VLANs on one port or Link Aggregate.

VLAN. Virtual Local Area Network - a logical segmentation of a network

VPD vital product data is information about a device that is stored on a computer's hard disk (or the device itself) that allows the device to be administered at a system or network level.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this Redpaper.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 236. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *The Cutting Edge: IBM eServer BladeCenter*, REDP3581
- ▶ *IBM eServer BladeCenter Systems Management*, REDP3582
- ▶ *Deploying Citrix MetaFrame on IBM eServer BladeCenter*, REDP3583
- ▶ *Deploying Lotus Domino on IBM eServer BladeCenter*, REDP3584
- ▶ *Deploying Microsoft Exchange on IBM eServer BladeCenter*, REDP3585
- ▶ *Deploying Samba on IBM eServer BladeCenter*, REDP3595
- ▶ *Deploying Apache on IBM eServer BladeCenter*, REDP3588

Other publications

These publications are also relevant as further information sources:

- ▶ *IBM @server BladeCenter Type 8677 Installation and User's Guide*
- ▶ *IBM 4-Port GB Ethernet Switch Module for BladeCenter Installation Guide*
- ▶ *IBM @server BladeCenter Management Module User's Guide*

Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ IBM eServer xSeries products
<http://www.ibm.com/eserver/xseries/>
- ▶ Documentation for Cisco switches
<http://www.cisco.com/en/US/products/hw/switches/>
- ▶ Documentation for Cisco routers
<http://www.cisco.com/en/US/products/hw/routers/>
- ▶ IEEE LAN specifications online
<http://standards.ieee.org/reading/ieee/std/lanman/>
- ▶ Documentation for troubleshooting Spanning Tree issues
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a00800951ac.shtml
- ▶ Layer 2 technology documents
<http://www.cisco.com/pcgi-bin/Support/browse/index.pl?i=Technologies&f=1324>

- ▶ BayStack 380-24T Product Literature
http://www.nortelnetworks.com/products/02/bstk/switches/baystack_380/doclib.html
- ▶ Technical documentation for BayStack 380-24T
<http://www130.nortelnetworks.com/cgi-bin/eserv/cs/main.jsp?cscat=documentation&tranProduct=11501>
- ▶ Passport 8600 Product Literature
<http://www.nortelnetworks.com/products/01/passport/lan/doclib.html#8600>
- ▶ Technical documentation for Passport 8600
<http://www130nortelnetworks.com/cgi-bin/eserv/cs/main.jsp?cscat=documentation&tranProduct=9015>
- ▶ IEEE LAN specifications online
<http://standards.ieee.org/reading/ieee/std/lanman/>
- ▶ xSeries products
<http://www.ibm.com/eserver/xseries/>
- ▶ Cisco Connection Online
http://www.cisco.com/en/US/partner/netsol/ns110/ns53/ns224/networking_solutions_packages_list.html
- ▶ xSeries products
<http://www.ibm.com/eserver/xseries/>

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Index

Numerics

10/100/1000 Mbps connections 8
1000BASE-T 9
1000Base-T 200, 212
100BASE-TX 9
100-ohm STP 9
10BASE-T 9
802.1D Spanning Tree support 8
802.1P/Q MIB 9
802.1Q Port Settings 187
802.1Q Tagged 8
802.3ad 183
8677 2

A

aging time 31–32
aging timer 39
ANSI/IEEE 802.3 NWay auto-negotiation 9
application-specific integrated circuits (ASICs) 30
auto-MDI/MDI-X 183, 185
autonegotiate 200, 206, 212
autonegotiation 184
autopolarity 183
autosensing 8

B

backbone 8
backpressure 31
bandwidth 5
base configuration 187
BayStack 380-24T 181–182, 185, 194, 200
blade server 8
blade servers 31
BladeCenter 181
blades 25
boot PROM 182
BOOTP server 23
bootstrap protocol (BOOTP) 9
bridge 39
bridge forward delay 40
bridge hello time 40
bridge max age 40
bridge priority 40
bridging 30
broadcast domains 33
broadcast frames 33

C

Carrier Sense Multiple Access with Collision Detection (CSMA/CD) 31
chassis 2, 7, 12, 25
collision domains 30

collisions 30
Configure 802.1Q trunking 195
configure 802.1Q trunking 201
Configure Link Aggregation 195
configure Link Aggregation 201
configure MLT 197
configure PVID 200, 212
Configure PVIDs 195
configuring the BayStack 380-24T 196, 208, 214
configuring the ESM 195, 200, 206, 212
copper ports 8
CRC error 32
cross-over cable 13
crossover cable 185
CSMA/CD 31
Customer Premise Equipment (CPE) 183
cut-through 32

D

data traffic 6
default addresses 3
default gateway 6
design scenarios 194
destination machine 30
DHCP server 12, 23
drivers 6
dynamic host configuration protocol (DHCP) 9

E

egress 7, 31
EIA/TIA-568 100-ohm STP 9
EIA/TIA-568B 100-ohm STP 9
enable external Ethernet ports 16
ESM 194
ESM IP address 23
ESM management session 23
ESM telnet interface 17
ESM Web interface 17
Ethernet activity 5
Ethernet connectivity 2
Ethernet daughter card 4
Ethernet link 5
Ethernet port 12
Ethernet switch error 5
Ethernet Switch Module (ESM) 182
Ethernet switch module (ESM) 2, 29
ethernet switch module (ESM) 11
Ethernet switches 30
Ethernet switching 2, 30
excessive flooding 32
external copper ports 8
external Ethernet interface 5
external management module port 23

F

- factory reset 187
- Fast Learning 184
- filtering data frames 30
- filtering database 38
- firmware 6, 21
- firmware version 182
- flash memory 8
- floods 32
- flow control 31
- forwarding table 30
- forwarding table age time 8
- full-duplex 31

G

- general switch information 6
- Generic Attribute Registration Protocol (GARP) 8
- gigabit aggregation layer switch 181

H

- half-duplex 31
- http
 - [//www.cisco.com/en/US/partner/net-sol/ns110/ns53/ns224/networking_solutions_packages_list.html](http://www.cisco.com/en/US/partner/net-sol/ns110/ns53/ns224/networking_solutions_packages_list.html) 75, 236
- HTTP Web interface 14
- hypertext transfer protocol (HTTP) 8

I

- IBM BladeCenter 4-Port Ethernet Switch Module 1
- IBM BladeCenter 8677 182
- IBM Director 14, 17, 25
- identification label 3
- IEEE 802.1d 36
- IEEE 802.1d Spanning Tree Algorithm 36, 40
- IEEE 802.1P Tagged Packets 9
- IEEE 802.1Q 7
- IEEE 802.1Q Tagged VLAN 9
- IEEE 802.1Q VLANs 40
- IEEE 802.3 10BASE-T Ethernet 9
- IEEE 802.3ab 1000BASE-T 9
- IEEE 802.3ad Link Aggregation 188
- IEEE 802.3u 100BASE-TX Fast Ethernet 9
- IEEE 802.3x 31
- IEEE 802.3x Full-duplex Flow Control 9
- IEEE 802.3z Gigabit Ethernet 9
- in-band 8
- ingress 7, 31
- initial recommended switch settings 14
- interface MIB 9
- internal full-duplex 10/100 Mbps ports 8
- internal full-duplex gigabit ports 8
- Internet group management protocol (IGMP) 8
- IP address 23
- IST (Inter Switch Trunk) 183

J

- Java 14
- JavaScript 1.2 14

L

- L2 method 184
- LED 4–5
- Link Aggregation 188, 194
- link aggregation 5, 8
- link aggregation protocols 183
- link trunks 206
- logical network segment 7
- logical network topology 7
- longer timer value 39
- loop-free network 36, 38

M

- MAC address 3, 30
- MAC addresses 7
- management information base (MIB) 9
- management module 2, 8, 12, 16
- management module Web interface 17
- management network 13
- management processor assistant 26
- management session 25
- management station 12
- MDI-X port 13
- media access control (MAC) 8
- media access control (MAC) address 3
- media dependent interface (MDI) 13
- mini-RMON MIB 9
- misconfiguration 185
- multi-link trunking (MLT) 183

N

- network design 200
- network management 8
- network monitoring 6
- network packets 7
- network topology 38
- nonvolatile random access memory (NVRAM) 17
- Nortel Networks 181
- NVRAM 21

O

- out of band port 193
- out-of-band 8

P

- packet forwarding 7
- packet header 7
- Passport 8600 181–183, 185, 192
- path cost 40
- pause frames 184
- Port Aggregation 211
- port priority 40
- port statistics 6

- Port VLAN ID (PVID) 187
- power-on self-test (POST) 5
- primary link 7
- priority queues 8
- processor blades 2
- production network 13
- protocols 8

Q

- Quality of Service (QoS) 183

R

- random-access memory (RAM) 8
- Red Hat Linux 182
- Redbooks Web site 236
 - Contact us xi
- redundant switch 200
- remote management 6
- remote monitoring (RMON) 9
- root 20
- root bridge 38

S

- secondary link 7
- security 19
- servers 8
- set intervals 23
- shared media 30
- shorter timer value 39
- simple network management protocol (SNMP) 8
- SMLT aggregation switch 183
- SMLT Client 183
- SNMP 6
- SNMP traps 16
- snooping support 8
- source learning 30
- source machine 30
- Spanning Tree 184, 206
- spanning tree algorithm 36, 38
- spanning tree bridge priority 192
- Spanning Tree Groups, 183
- spanning tree parameters 38, 40
- Spanning Tree Protocol 184
- spanning tree protocol 7, 37
- Spanning Tree Protocol (STP) 9
- spanning tree sequence 37
- spanning tree topology 39
- Split-MLT (SMLT) 183
- standards 8
- static IP addresses 24
- stations 8
- store-and-forward 8, 32
- straight cable 185
- subnet 13
- switch forwarding table 31
- switch information panel 18
- switch maintenance 6
- switch management 6, 8

- switch module 12
- switch parameters 6
- switch TCP/IP address 6

T

- tagging 7
- telnet remote console 8
- Topology Change Acknowledgment (TCA) 39
- Topology Change Notification (TCN) 39
- topology change notification protocol 38
- transmission method 8
- traps 6
- trivial file transfer protocol (TFTP) 8

U

- unshielded twisted pair (UTP) 13
- untagging 7
- User guides 6
- UTP Category 3 9
- UTP Category 4 9
- UTP Category 5 9
- UTP Category 5e 9

V

- virtual local area network (VLAN) 7–8, 32
- Virtual local area networks (VLANs) 40
- vital product data (VPD) 25
- VLAN 33, 185, 187, 193
- VLAN Registration Protocol (GVRP) 8
- VLAN tags 7

W

- Web browser 6, 14, 190
- Web Switching Module (WSM) 184
- Web-based management 8



IBM @server BladeCenter Networking Options



**Step-by-step
configuration
instructions for the
BladeCenter Ethernet
Switch Module**

**Working network
configurations to get
your IBM @server
BladeCenter solution
up and running quickly**

**Helpful information for
rapid deployment in a
Cisco or Nortel network
infrastructure**

This IBM Redpaper will help you install, tailor and configure the new IBM @server BladeCenter and its Ethernet switch module in various network environments.

In this Redpaper, we discuss the IBM BladeCenter 4-Port Gb Ethernet Switch Module. We discuss the features and functions of the Ethernet module and how the switch is managed.

This Redpaper is designed to help you with the initial network configurations to ensure that your Ethernet switch module is configured correctly and operable so that you can begin immediate communications across the network. In this Redpaper, you will find working configurations that have been specified and tested in our labs.

Also featured in this Redpaper are several configuration examples of the Ethernet switch module deployed in Cisco and Nortel environments.

The intent of this Redpaper is to introduce networking terms and concepts to the users and administrators of the IBM @server BladeCenter product. Our explanations of these terms and concepts are meant to give the non-networking professional an overall view of the switching and bridging environment and not substitute in-depth training for networking fundamentals. The examples illustrated in this paper are meant to provide real “working” examples of networking configurations where an @server BladeCenter is deployed; of course, many other configurations are possible.

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