



BladeCenter T 4-Port Gb
Ethernet Switch Module

User's Guide





@server

BladeCenter T 4-Port Gb
Ethernet Switch Module
User's Guide

Note: Before using this information and the product it supports, read the general information in Appendix A, "Getting help and technical assistance," on page 233, and Appendix B, "Notices," on page 235.

First Edition (July 2004)

© Copyright International Business Machines Corporation 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Safety	ix
Chapter 1. Introducing the IBM @server BladeCenter T 4-Port Gb Ethernet	
Switch Module	1
Specifications and features	2
Related documentation	5
Notices and statements used in this book	6
Major components of the Ethernet switch module.	6
Chapter 2. Information panel LEDs and internal and external ports	9
Information panel LEDs	9
Internal and external ports.	10
Chapter 3. Switch management and operating concepts	11
Ethernet switch module overview	11
BladeCenter T unit configuration and operation	12
Ethernet switch module management and control	12
IP addresses and SNMP community names	13
Traps	14
Management information bases (MIB)	15
Port mirroring	16
Simple Network Management Protocol (SNMP)	16
Authentication	16
Switching concepts	17
Packet forwarding	17
Spanning Tree Protocol (STP)	18
Virtual Local Area Networks (VLAN)	19
Notes about VLANs on the Ethernet switch module	19
IEEE 802.1Q VLANs.	19
IEEE 802.1Q VLAN packet forwarding	20
IEEE 802.1Q VLAN tags	21
Port VLAN ID	22
Tagging and untagging	23
Ingress filtering and egress rules	23
IEEE 802.1Q VLAN configuration	24
Protocol-based VLANs (PBVLANS)	24
Static MAC filtering	25
Generic Attribute Registration Protocol (GARP)	25
GARP VLAN Registration Protocol (GVRP)	25
GARP Multicast Registration Protocol (GMRP)	26
Internet Group Management Protocol (IGMP) snooping	27
Link Aggregation (LAG).	28
Static LAGs	28
Distribution method	28
Dynamic Host Configuration Protocol (DHCP)	29
Routing concepts	30
IP mapping	30
Routing Information Protocol (RIP).	30
Open Shortest Path First (OSPF)	31
VLAN routing	32
Packet processing at the bridge layer	32
Packet processing at the routing layer	32
Virtual Router Redundancy Protocol (VRRP)	33

BOOTP/ DHCP relay agent	33
Router discovery	34
Security	35
IEEE 802.1X.	35
Local authentication	36
RADIUS authentication	36
Secure Shell (SSH)	37
Secure Socket Layer (SSL)	38
Quality of Service (QoS)	38
Bandwidth provisioning	39
Access Control Lists (ACL)	39
IP multicast concepts	40
Internet Group Management Protocol (IGMP).	40
Distance Vector Multicast Routing Protocol (DVMRP).	40
Protocol Independent Multicast - Dense Mode (PIM-DM)	42
Protocol Independent Multicast - Sparse Mode (PIM-SM)	43
Understanding and Troubleshooting the Spanning Tree Protocol.	44
Spanning Tree Protocol (STP) operation	44
Creating a stable topology.	46
IEEE 802.1D STP port states	47
IEEE 802.1w STP port states	48
Setting user-changeable STP parameters	49
Illustration of STP	50
Discarding state	52
Learning state	53
Forwarding state	55
Disabled state	56
Troubleshooting STP.	57
Spanning Tree Protocol Failure	57
Full/half duplex mismatch	57
Unidirectional link	58
Packet corruption	60
Resource errors	60
Identifying a data loop	60
Avoiding network problems	61
Chapter 4. Web-based network management	63
Introduction	63
Remotely managing the switch module	63
Getting started	64
System.	66
ARP cache	67
Inventory information.	67
Configuration	68
System description	69
Network connectivity	69
Telnet	71
User accounts	71
Login configuration	73
Login session	75
Login summary.	75
User login.	76
Forwarding database.	77
Configuration	77
Search	77
Logs.	79

Message log	79
Event log	80
Port	80
Configuration	81
Summary	83
Mirroring	85
SNMP	85
Community configuration	86
Trap receiver configuration	87
Trap receiver summary	88
Supported Management Information Bases (MIB)	89
Statistics	89
Switch detailed	89
Switch summary	91
Port detailed	93
Port summary	98
System utilities	99
Save all applied changes	99
System reset	100
Reset configuration to defaults.	100
Reset passwords to defaults	101
Download file to switch	101
Upload file from switch	103
Ping	104
Trap manager	104
Trap flags	104
Trap log	106
Switching	106
VLAN	106
Configuration	107
Status.	109
Port configuration	110
Port summary	111
Reset configuration	111
Protocol-based VLAN	112
Configuration	112
Summary	113
Filters	114
MAC filter configuration	114
MAC filter summary.	115
GARP	115
Status.	116
Switch configuration	117
Port configuration	118
IGMP snooping	119
Configuration and status	119
Interface configuration.	120
LAG	120
Configuration	120
Status.	122
MFDB.	122
MFDB table	123
GMRP table	124
IGMP snooping table	124
Stats	125
Spanning tree	125

Switch configuration/status	126
Common Spanning Tree (CST) configuration/status	127
CST port configuration/status	128
Statistics	130
Routing	130
ARP	131
ARP create	131
ARP table configuration	132
IP	133
Configuration	133
Statistics	134
Interface configuration	137
OSPF	138
OSPF configuration	139
Area configuration	141
Stub area summary	143
Area range configuration	144
Interface statistics	145
Interface configuration	146
Neighbor table	149
Neighbor status	150
LSDB	152
Virtual link configuration	153
Virtual link summary	156
BOOTP/DHCP relay	157
Configuration	157
Status	158
RIP	159
Configuration	159
Interface summary	160
Interface configuration	161
Router discovery	162
Configuration	163
Status	164
Router	164
Route table	165
Best routes table	166
Route entry configuration	167
Route preferences configuration	168
VLAN routing	169
Configuration	169
Summary	170
VRRP	170
VRRP configuration	171
Virtual router configuration	171
Virtual router status	173
Virtual router statistics	174
Class of service	175
802.1p priority mapping	176
Security	176
Port access control	176
Configuration	177
Port configuration	177
Port status	179
Port summary	182
Statistics	183

Login	184
Port access privileges	185
Port access summary	185
RADIUS	186
Configuration	186
Server configuration	187
RADIUS statistics	188
Server statistics	188
Accounting server configuration	190
Accounting server statistics	191
Clear statistics	192
Secure HTTP	192
Configuration	192
Secure Shell	193
Configuration	193
QoS	194
Access Control Lists	194
Configuration	195
Summary	196
Rule configuration	196
Bandwidth provisioning	198
Bandwidth profile configuration	198
Bandwidth profile summary	199
Traffic class configuration	200
Traffic class summary	201
Interface allocation summary	201
IP multicast	202
DVMRP	202
Global configuration	203
Interface configuration	203
Configuration summary	204
Next hop summary	206
Prune summary	206
Route summary	207
IGMP	208
Global configuration	208
Interface configuration	208
Configuration summary	210
Cache information	212
Multicast	212
Global configuration	213
Interface configuration	214
MRoute summary	214
Static routes configuration	215
Static routes summary	216
Admin boundary configuration	217
Admin boundary summary	218
Mdebug	218
Mrinfo run	219
Mrinfo show	219
Mstat run	220
Mstat show	220
Mtrace config	221
Mtrace run	222
Mtrace show	222
PIM-DM	223

Global configuration	223
Interface configuration	224
Interface summary	224
PIM-SM	225
Global configuration	226
Global status	226
Interface configuration	227
Interface summary	228
Component summary	229
RP set summary	230
Candidate RP summary	231
Logout	231
Appendix A. Getting help and technical assistance	233
Before you call	233
Using the documentation	233
Getting help and information from the World Wide Web	233
Software service and support	234
Hardware service and support	234
Appendix B. Notices	235
Edition notice	235
Trademarks	236
Important notes	236
Product recycling and disposal	237
Battery return program	237
Electronic emission notices	238
Federal Communications Commission (FCC) statement	238
Industry Canada Class A emission compliance statement	238
Australia and New Zealand Class A statement	238
United Kingdom telecommunications safety requirement	238
European Union EMC Directive conformance statement	239
Taiwanese Class A warning statement	239
Chinese Class A warning statement	239
Japanese Voluntary Control Council for Interference (VCCI) statement	239
Index	241

Safety

Before installing this product, read the Safety Information.

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الأمنية

Antes de instalar este produto, leia as Informações de Segurança.

在安装本产品之前，请仔细阅读 **Safety Information** (安全信息)。

安裝本產品之前，請先閱讀「安全資訊」。

Prije instalacije ovog produkta obavezno pročitajte Sigurnosne Upute.

Před instalací tohoto produktu si přečtěte příručku bezpečnostních instrukcí.

Læs sikkerhedsforskrifterne, før du installerer dette produkt.

Lees voordat u dit product installeert eerst de veiligheidsvoorschriften.

Ennen kuin asennat tämän tuotteen, lue turvaohjeet kohdasta Safety Information.

Avant d'installer ce produit, lisez les consignes de sécurité.

Vor der Installation dieses Produkts die Sicherheitshinweise lesen.

Πριν εγκαταστήσετε το προϊόν αυτό, διαβάστε τις πληροφορίες ασφαλείας (safety information).

לפני שתתקינו מוצר זה, קראו את הוראות הבטיחות.

A termék telepítése előtt olvassa el a Biztonsági előírásokat!

Prima di installare questo prodotto, leggere le Informazioni sulla Sicurezza.

製品の設置の前に、安全情報をお読みください。

본 제품을 설치하기 전에 안전 정보를 읽으십시오.

Пред да се инсталира овој продукт, прочитајте информацијата за безбедност.

Les sikkerhetsinformasjonen (Safety Information) før du installerer dette produktet.

Przed zainstalowaniem tego produktu, należy zapoznać się z książką "Informacje dotyczące bezpieczeństwa" (Safety Information).

Antes de instalar este produto, leia as Informações sobre Segurança.

Перед установкой продукта прочтите инструкции по технике безопасности.

Pred inštaláciou tohto zariadenia si pečítajte Bezpečnostné predpisy.

Pred namestitvijo tega proizvoda preberite Varnostne informacije.

Antes de instalar este producto, lea la información de seguridad.

Läs säkerhetsinformationen innan du installerar den här produkten.

Important:

All caution and danger statements in this documentation begin with a number. This number is used to cross reference an English caution or danger statement with translated versions of the caution or danger statement in the *IBM Safety Information* book.

For example, if a caution statement begins with a number 1, translations for that caution statement appear in the *IBM Safety Information* book under statement 1.

Be sure to read all caution and danger statements in this documentation before performing the instructions. Read any additional safety information that comes with your server or optional device before you install the device.

Statement 1:



DANGER

Electrical current from power, telephone, and communication cables is hazardous.

To avoid a shock hazard:

- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **Connect all power cords to a properly wired and grounded electrical outlet.**
- **Connect to properly wired outlets any equipment that will be attached to this product.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following table when installing, moving, or opening covers on this product or attached devices.**

To Connect:

1. Turn everything OFF.
2. First, attach all cables to devices.
3. Attach signal cables to connectors.
4. Attach power cords to outlet.
5. Turn device ON.

To Disconnect:

1. Turn everything OFF.
2. First, remove power cords from outlet.
3. Remove signal cables from connectors.
4. Remove all cables from devices.

Statement 3:



CAUTION:

When laser products (such as CD-ROMs, DVD drives, fiber optic devices, or transmitters) are installed, note the following:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.



DANGER

Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following.


Laser radiation when open. Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam.

Statement 5:



CAUTION:

The power control button on the device and the power switch on the power supply do not turn off the electrical current supplied to the device. The device also might have more than one power cord. To remove all electrical current from the device, ensure that all power cords are disconnected from the power source.



Statement 6:



CAUTION:

If you install a strain-relief bracket option over the end of the power cord that is connected to the device, you must connect the other end of the power cord to an easily accessible power source.

Statement 8:



CAUTION:

Never remove the cover on a power supply or any part that has the following label attached.



Hazardous voltage, current, and energy levels are present inside any component that has this label attached. There are no serviceable parts inside these components. If you suspect a problem with one of these parts, contact a service technician.

Statement 13:



DANGER

Overloading a branch circuit is potentially a fire hazard and a shock hazard under certain conditions. To avoid these hazards, ensure that your system electrical requirements do not exceed branch circuit protection requirements. Refer to the information that is provided with your device for electrical specifications.

Statement 14:



CAUTION:

Hazardous voltage, current, and energy levels might be present. Only a qualified service technician is authorized to remove the covers where the following label is attached.



Chapter 1. Introducing the IBM @server BladeCenter T 4-Port Gb Ethernet Switch Module

The IBM® @server BladeCenter™ T 4-Port Gb Ethernet Switch Module high-performance Ethernet switch module is ideally suited for networking environments that require superior microprocessor performance, efficient memory management, flexibility and reliable data storage. The BladeCenter T 4-Port Gb Ethernet Switch Module is one of up to two switch modules that can be installed in the IBM BladeCenter T unit.

This *User's Guide* contains information about:

- Specifications and features of the Ethernet switch module
- Configuring the Ethernet switch module
- Using the Web interface

The product name, machine type, serial number, and the Media Access Control (MAC) address are located on the identification label on the side of the Ethernet switch module. See "Major components of the Ethernet switch module" on page 6 for an illustration showing 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.

Specifications and features

The IBM BladeCenter T 4-Port Gb Ethernet Switch Module has the following features:

- Ports
 - Four external 1000BASE-T ports for making 10/100/1000 Mbps connections to a backbone, end stations, and servers
 - Eight internal full-duplex gigabit ports, one connected to each of the blade servers in the BladeCenter T unit
 - Two internal full-duplex 100 Mbps ports connected to the BladeCenter T management module
- Performance features
 - Support for packet sizes from 64 bytes to 1518 bytes
 - Transmission method: Store-and-forward
 - Packet filtering/forwarding rate
 - Full-wire speed for all connections. 148800 packets per second (pps) per port (for 100 Mbps)
 - 1488100 pps per port (for 1000 Mbps)
 - Media Access Control (MAC) address learning: automatic update. Supports 12 K MAC address
 - Forwarding table age time: Maximum age: 10 to 1,000,000 seconds. Default is 300 seconds
 - Support for 3584 concurrent VLANs
 - Switch topology: Star
- Standards

The following standards apply to the Ethernet switch module.

- Switching support
 - IEEE 802.3 10BASE-T Ethernet
 - IEEE 802.3 Auto-negotiation
 - IEEE 802.3u 100BASE-TX Fast Ethernet
 - IEEE 802.3z Gigabit Ethernet
 - IEEE 802.3ab 1000BASE-T
 - IEEE 802.1Q Tagged VLAN
 - IEEE 802.1p Priority
 - Protocol-based VLANs
 - Port-based VLANs
 - GARP
 - GMRP
 - GVRP
 - IEEE 802.3ac - VLAN Tagging
 - IEEE 802.3ad - Link Aggregation
 - IEEE 802.1s - Spanning Tree
 - IEEE 802.1w - Rapid Spanning Tree
 - IEEE 802.1X - Port Based Authentication
 - IEEE 802.3X - Flow Control
 - RFC 768 - UDP
 - RFC 783 - TFTP
 - RFC 791 - IP
 - RFC 792 - ICMP
 - RFC 793 - TCP
 - RFC 826 - ARP
 - RFC 1321 - Message Digest Algorithm
 - RFC 2131 - DHCP Client
 - RFC 2865 - RADIUS Client
 - RFC 2866 - RADIUS Accounting
 - RFC 2868 - RADIUS Attributes for Tunnel Protocol Support

- RFC 2869 - RADIUS Extensions
- RFC 2869bis - RADIUS Support for Extensible Authentication Protocol (EAP)
- Advanced Layer 2 functionality
 - Broadcast Storm Recovery
 - Multicast Storm Recovery
 - Independent VLAN Learning (IVL) support
 - Port Mirroring
 - IGMP Snooping
 - Static MAC Filtering
- Switch facilities
 - Event and Error Logging Facility
 - Run-time and Configuration Download Capability
 - PING Utility
- Routing support
 - RFC 826 - Ethernet ARP
 - RFC 894 - Transmission of IP Datagrams over Ethernet Networks
 - RFC 896 - Congestion Control in IP/TCP Networks
 - RFC 1058 - RIP v1
 - RFC 1256 - ICMP Router Discovery Messages
 - RFC 1519 - CIDR
 - RFC 1812 - Requirements for IP Version 4 Routers
 - RFC 2082 - RIP-2 MD5 Authentication
 - RFC 2328 - OSPF v2
 - RFC 2338 - VRRP - Virtual Router Redundancy Protocol
 - RFC 2453 - RIP v2
 - RFC 3046 - DHCP/BootP Relay
 - VLAN Routing
- Quality of Service (QoS) support
 - Bandwidth Provisioning
 - Maximum Burst Rate (MBR)
 - Per Port (Interface)
 - Per VLAN
 - Access Control Lists
 - Inbound Filtering
 - Source IP
 - Destination IP
 - Source L4 Port
 - Destination L4 Port
- Multicast
 - RFC 2236 - IGMP v2
 - Draft-ietf-idmr-dvmrp-v3-10 - DVMRP
 - Draft-ietf-v2-dm-03 - PIM-DM
 - RFC 2362 - PIM-SM
 - RFC 1112 - Host Extensions for IP Multicasting
 - RFC 2365 - Administratively Scoped Boundaries
 - IP Multicast Traceroute
- Management
 - RMON - Groups 1, 2, 3 and 9 supported
 - Simple Network Management Protocol (SNMP) versions 1, 2 and 3
 - Flash memory for software upgrades, done using trivial file transfer protocol (TFTP)
 - Supports Web-based management
 - HTML 4.0 Specification - December, 1997
 - Java™ Script 1.3
 - Java 1.3

- RFC 2068 - HTTP/1.1 protocol as updated by draft-ietf-http-v11-spec-rev-03
- HTML/2.0 Forms with file upload extensions
- Command Line Interface (CLI) with the following features
 - Scripting capability
 - Command completion
 - Context sensitive help
 - Multi-session Telnet Server
- RFC 854 - Telnet
- RFC 855 - Telnet Option
- RFC 1155 - SMI v1
- RFC 1157 - SNMP
- RFC 1212 - Concise MIB Definitions
- RFC 1901 - Community-based SNMP v2
- RFC 1905 - Protocol Operations for SNMP v2
- RFC 1906 - Transport Mappings for SNMP v2
- RFC 1907 - Management Information Base for SNMP v2
- RFC 1908 - Coexistence between SNMP v1 and SNMP v2
- RFC 2295 - Remote Variant Selection; RSVP/1.0 State Management “cookies”
 - draft-ietf-http-state-mgmt-05
- RFC 2571 - Architecture for Describing SNMP Management Frameworks
- RFC 2572 - Message Processing and Dispatching for SNMP
- RFC 2573 - SNMP v3 Applications
- RFC 2574 - User Based Security Model for SNMP v3
- RFC 2575 - View-based Access Control Model for SNMP
- RFC 2576 - Coexistence between SNMP v1, v2, and v3
- RFC 2580 - Conformation statements for SMI v2
- Configurable management VLAN
 - Secure Socket Layer (SSL) 3.0 and Transport Layer Security (TLS) 1.0
 - RFC 2246 - The TLS Protocol, Version 1.0
 - RFC 2818 - HTTP over TLS
 - RFC 2346 - AES Ciphersuites for TLS
 - Secure Shell (SSH) 1.5 and 2.0
 - Draft-ietf-secsh-transport-16 - SSH Transport Layer Protocol
 - Draft-ietf-secsh-userauth-17 - SSH Authentication Protocol
 - Draft-ietf-secsh-connect-17 - SSH Connection Protocol
 - Draft-ietf-secsh-architecture-14 - SSH Protocol Architecture
 - Draft-ietf-secsh-publickeyfile-03 - SECSH Public Key File Format
 - Draft-ietf-secsh-dh-group-exchange-04 - Diffie-Hellman Group Exchange for the SSH Transport Layer Protocol
- MIBs supported
 - Switching MIBs
 - RFC 1213 - MIB-II
 - RFC 1493 - Bridge MIB
 - RFC 1643 - Ethernet-like MIB
 - RFC 2674 - VLAN MIB
 - RFC 2618 - RADIUS Authentication Client MIB
 - RFC 2620 - RADIUS Accounting MIB
 - RFC 2819 - RMON Groups 1, 2, 3 and 9
 - IEEE 802.1X MIB (IEEE 802.1-PAE-MIB)
 - Enterprise MIB
 - Routing MIBs
 - RFC 1724 - RIP v2 MIB Extension
 - RFC 1850 - OSPF MIB
 - RFC 2233 - The Interfaces Group MIB using SMI v2
 - RFC 2787 - VRRP MIB
 - QoS / SNMP support in Enterprise MIBs

- Available through management module
- Private MIBs for full configuration of ACL and Bandwidth Provisioning functionality
- Multicast MIBs
 - RFC 2933 - IGMP MIB
 - RFC 2934 - PIM MIB for IP v4
 - Draft-ietf-idmr-dvmrp-mib-11.txt - DVMRP
- Network cables
 - 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)
 - UTP Category 5 (100 meters maximum)
 - EIA/TIA-568B 100-ohm STP (100 meters maximum)

Related documentation

This *User's Guide* provides information about the IBM @server BladeCenter T 4-Port Gb Ethernet Switch Module, including information about features, how to configure the Ethernet switch module, and how to get help.

In addition to this *User's Guide*, the following documentation is provided in PDF on the IBM *BladeCenter T Documentation* CD that comes with your IBM BladeCenter T unit:

- *Safety Information*
This document contains translated caution and danger statements. Each caution and danger statement that appears in the documentation has a number that you can use to locate the corresponding statement in your language in the *Safety Information* document.
- *BladeCenter T 4-Port Gb Ethernet Switch Module Installation Guide*
This document contains information about setting up the Ethernet switch module.
- *BladeCenter T 4-Port Gb Ethernet Switch Module CLI Reference Guide*
This document contains information about setting up and configuring the Ethernet switch module using the command-line interface.
- *BladeCenter T Types 8720 and 8730 Installation and User's Guide*
This document contains instructions for setting up and configuring the BladeCenter T unit and basic instructions for installing some options. It also contains general information about the BladeCenter T unit.
- *BladeCenter T Management Module Installation Guide*
This document contains information about getting started, first-time connection to the management module, and how to configure the management module in a BladeCenter T unit.
- *BladeCenter T Management Module User's Guide*
This document contains instructions for using the Web interface to configure the management modules in a BladeCenter T unit.
- *BladeCenter HS20 Type 8832 Installation and User's Guide*
This document contains instructions for setting up a BladeCenter HS20 blade server and basic instructions for installing some options. It also contains general information about the blade server.

- *Hardware Maintenance Manual and Troubleshooting Guides*

The BladeCenter T unit and the BladeCenter HS20 product each have a customized *Hardware Maintenance Manual and Troubleshooting Guide*. These documents contain information to help you solve problems yourself, and they contain information for service technicians.

Additional documents might be included on the IBM *BladeCenter T Documentation* CD.

Your Ethernet switch module might have features that are not described in the documentation that you received with the BladeCenter T unit. The documentation might be updated occasionally to include information about those features, or technical updates might be available to provide additional information that is not included in your documentation. These updates are available from the IBM Web site. Complete the following steps to check for updated documentation and technical updates:

1. Go to <http://www.ibm.com/pc/support/>.
2. In the **Learn** section, click **Online publications**.
3. On the “Online publications” page, in the **Brand** field, select **Servers**.
4. In the **Family** field, select **BladeCenter T**.
5. Click **Continue**.

Notices and statements used in this book

The caution and danger statements that appear in this book are also in the multilingual *Safety Information Book* on the IBM *BladeCenter T Documentation* CD. Each statement is numbered to refer to the corresponding statement in the *Safety Information Book*.

The following notices and statements are used in this book:


- **Note:** These notices provide important tips, guidance or advice.
- **Important:** These notices provide information or advice that might help you avoid inconvenient or problematic situations.
- **Attention:** These notices indicate possible damage to programs, devices or data. An attention notice is placed just before the instruction or situation in which damage could occur.
- **Caution:** These statements indicate situations that can be potentially hazardous to you. A caution statement is placed just before the description of a potentially hazardous procedure, step or situation.
- **Danger:** These statements indicate situations that can be potentially lethal or extremely hazardous to you. A danger statement is placed just before the description of a potentially lethal or extremely hazardous procedure, step or situation.

Major components of the Ethernet switch module

The orange color on components and labels on the Ethernet switch module and on the platform identifies hot-swap or hot-plug components. You can install or remove these components while the system is running, provided that your system is configured to support this function.

The following illustration shows the major components of your Ethernet switch module.

Note: The illustrations in this document may differ slightly from your hardware.




For more information about the components of the information panel, see Chapter 2, “Information panel LEDs and internal and external ports,” on page 9. For more information about the MAC address, see “IP addresses and SNMP community names” on page 13.

Chapter 2. Information panel LEDs and internal and external ports

This chapter describes the information panel LEDs (also known as indicators) and the internal and external ports on the Ethernet switch module.

Information panel LEDs

The information panel of the Ethernet switch module consists of LEDs and four external RJ-45 connectors. The following illustration shows the LEDs on the switch module. A description of each LED follows the illustration.



Note: The illustrations in this document may differ slightly from your hardware.

OK LED: This green LED is at the top left of the information panel as shown in the illustration. When this LED is lit, it indicates that the switch module has passed the power-on self-test (POST) and is operational.


Ethernet switch error / Fault indicator (!) LED: This amber LED is at the top right of the information panel as shown in the illustration. When this LED is lit, it indicates that the switch module has a fault.

Ethernet link: This green link status LED is at the top right of each external 10/100/1000 Mbps connector as shown in the illustration. When one of these LEDs is lit, it indicates that there is a connection (or link) to a device on the corresponding connector.

Ethernet activity: This green activity LED is at the bottom right of each external 10/100/1000 Mbps connector as shown in the illustration. When one of these LEDs is flashing, it indicates that data is being received or transmitted (that is, activity is occurring) on the corresponding connector. The flashing frequency is proportional to the amount of traffic on that port.

Internal and external ports

The information panel of the Ethernet switch module consists of LEDs and four external RJ-45 connectors, as shown in the following illustration.



The switch module contains the following internal and external connectors:

- Two internal full-duplex 100 Mbps ports that are connected to the BladeCenter T unit management module.
- Eight internal full-duplex 1000 Mbps ports that are connected to the eight blade servers in the BladeCenter T unit.
- Four external RJ-45 Ethernet connectors for 10/100/1000 Mbps connections to external Ethernet devices such as backbones, end stations, and servers. These connectors are identified as Ext1, Ext2, Ext3, and Ext4 in the switch configuration menus and are labeled 1 through 4 on the switch module, as shown in the previous illustration.

Chapter 3. Switch management and operating concepts

This chapter presents many of the concepts and features that are used to manage the Ethernet switch module.

Ethernet switch module overview

This section provides information that you should be familiar with when managing and configuring the internal switch modules. If you are familiar with Ethernet switches, you will recognize the industry-standard parameters and terminology that are used in this document. It is important to understand the operating environment of the Ethernet switch module with regard to the internal switches.

Ethernet switch modules are hot-swappable subsystems that provide Ethernet switching capabilities within the BladeCenter T unit. The primary purpose of the Ethernet switch module is to provide Ethernet interconnectivity among the blade servers, management modules, and the external network infrastructure.

The BladeCenter T unit can be configured with up to two Ethernet switch modules, supporting up to eight blade servers. Internal ports 1 through 8 on the Ethernet switch module correspond to blade servers 1 through 8 (numbered top to bottom when viewed from the front of the BladeCenter T unit). Each Ethernet 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 Ethernet switch-module configuration menus, and the connectors for these ports are labeled 1 through 4 on the Ethernet switch module (see Chapter 2, "Information panel LEDs and internal and external ports," on page 9 for an illustration).

Depending on the application, the external Ethernet interfaces can be configured to meet a variety of requirements for bandwidth or function. The Ethernet switch module has been preconfigured with default parameter settings that can be used with some typical installations. Most installations will need some configuration of parameters. Information about the initial software configuration is in the *CLI Reference Guide* on the IBM *BladeCenter T Documentation CD*.

BladeCenter T unit configuration and operation

Each Ethernet switch module is an integral subsystem in a BladeCenter T unit. Each BladeCenter T unit has one or two management modules for management and control. The Ethernet switch module has two 100 Mbps internal Ethernet ports that can be accessed only by the management modules. To prevent inadvertent changes, these management ports are “hidden” and do not appear in the port configuration and status screens. The factory default settings will permit management and control access to the Ethernet switch module only through a single 100 Mbps Ethernet port on the management module. You can use the four external 10/100/1000 Mbps Ethernet ports on the Ethernet switch module for management and control of the module by selecting this mode as an option through the management-module configuration utility program (see the applicable *Installation Guide* and *User's Guide* on the IBM *BladeCenter T Documentation CD* for more information).

Ethernet switch module management and control

This document describes the user interfaces, screens, parameters and other information that you need for remote management and control of your Ethernet switch module. Complete the following initial configuration steps:

1. Connect the Ethernet port of the management module to a 100 Mbps network (with access to a management station) or directly to a management station.
2. Initially configure the management module with the applicable IP addresses for network access (see the *IBM BladeCenter T Types 8720 and 8730 Installation and User's Guide* publication on the IBM *BladeCenter T Documentation CD* for more information).
3. An IP address is assigned to the Ethernet switch module automatically through a DHCP server.

When a transmission control protocol/Internet protocol (TCP/IP) communication path has been established with the Ethernet switch module through the Ethernet port of the management module, you can perform a series of management and control tasks. These tasks are in the following categories:

- Configuration
- Modification of the Ethernet switch module parameter settings
- Remote management setup
- Network monitoring
 - Automatically receive error alerts (traps)
 - View and reset port traffic statistics
 - Monitor data traffic on selected output ports
- Maintenance
 - Update the software on the Ethernet switch module
 - View and configure the message and event logs
 - Restore factory default settings

The Ethernet switch module supports three primary management and control user interfaces. A built-in Web browser interface is the primary interface (see Chapter 4, “Web-based network management,” on page 63 for detailed information). You can invoke the Web browser interface from the management and configuration utility program, along with the Telnet interface, which provides a command-line interface (CLI) (see the *CLI Reference Guide* on the IBM *BladeCenter T Documentation* CD for detailed information). Both interfaces provide access to the same switch information and control parameters.

In addition, you can access an extensive set of both standard and enterprise MIB objects through SNMP protocols.


IP addresses and SNMP community names

Each Ethernet switch module must be assigned its own Internet protocol (IP) address, which is used for communication with a Simple Network Management Protocol (SNMP) network manager or other TCP/IP application. The Ethernet switch module factory-default IP address is 10.90.90.9x, where x is determined by the number of the I/O-module bay into which you have installed the Ethernet switch module, as shown in Table 1.

Table 1. Factory-default IP addresses based on I/O-module bay numbers

I/O-module bay number	Factory-default IP address
1	10.90.90.91
2	10.90.90.92
3	10.90.90.94
4	10.90.90.97


The following illustration shows the I/O-module bay locations at the rear of the BladeCenter T unit.



You can change the factory-default switch-module IP address to meet the requirements of your networking address scheme.

The switch module also has a unique, factory-assigned media access control (MAC) address. The switch-module MAC address is located on one side of the switch module, on the same label as the serial number, as shown in the following illustration.

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



The Ethernet switch-module MAC address can be displayed using the Inventory Information panel on page 67, or by using the CLI command **show inventory**.

If the network-management station and Ethernet switch modules are on different IP networks, the management packets must go through a router to reach the network manager. You can set the IP address for the gateway router using the CLI and management-module Web interfaces.

For security, you can specify the IP addresses of the network managers that are permitted to manage the Ethernet switch module by using the SNMP Community Configuration panel on page 86, or by using the **config snmpcommunity ipaddr** CLI command. You can also change the default SNMP community strings in the switch module and set the access rights of these community strings.

Traps

Traps are messages that alert you of certain events that occur on the switch module. The events can be as serious as a restart (for example, the switch module was turned off accidentally) or less serious, such as a port-status change. The switch module generates traps and sends them to the network manager (a trap recipient).

Trap recipients are special users of the network who are given certain rights and access to oversee the maintenance of the network. Trap recipients will receive traps that are sent from the switch module; they might take certain actions to avoid future failure or breakdown of the network.

You can also specify which network managers can receive traps from the switch module by entering a list of the IP addresses of authorized network managers. You can enter up to four trap-recipient IP addresses and four corresponding SNMP community strings.

SNMP community strings function as passwords. The community string that is entered for an IP address must be set prior to use in the management station software, or a trap will be sent.

The following trap types can be sent from the switch module to a trap recipient:

Cold start This trap indicates that the Ethernet switch module is turned on and initialized, the software settings are reconfigured, and hardware systems are restarted. A cold start saves configuration settings to nonvolatile random-access memory (NVRAM) and uses these settings to reconfigure the switch module.

Note: If you select **Reset Configuration to Defaults** in the Web interface, all user-defined configuration settings are erased. See page 100 for more information.

Warm start This trap indicates that the switch module has been restarted.

Authentication failure

This trap indicates that someone has tried to log on to the switch module using an invalid SNMP community string. The switch module automatically stores the source IP address of the unauthorized user.

Topology change (Spanning Tree Protocol (STP))

This trap indicates that one or more of the configured ports have changed from the learning state to the forwarding state, or from the forwarding state to the blocking state.

Link up This trap indicates that the link state of a port has changed from link down to link up.

Link down This trap indicates that the link state of a port has changed from link up to link down.

Management information bases (MIB)

Management and counter information is stored in the Ethernet switch module in the management information base (MIB). The switch module uses the standard MIB-II module. You can retrieve values for MIB objects using SNMP-based network management software. In addition to the standard MIB-II module, the switch module also supports its own proprietary enterprise MIB. You can retrieve this MIB by specifying the object identifier (OID) of the MIB as the network manager. MIB parameters can be either read-only or read/write.

Read-only MIB parameters can be either constants that are programmed into the switch module or variables that change while the switch module is in operation. Examples of read-only constants are the number of ports and type of ports. Examples of read-only variables are the statistics counters, such as the number of errors that have occurred, or how much data (in KB) has been received and forwarded through a port.

Read/write MIB parameters are usually related to user-customized configurations. Examples of these are the switch-module IP address, Spanning Tree Protocol (STP) parameters, and port status.

If you use non-IBM SNMP software to manage the switch module, you can request access to the enterprise MIB from the vendor. If your software provides functions to browse or modify MIBs, you can change values of MIB variables that enable the write operation. To change the values of the variables, you must know the MIB OIDs and set the values individually.

Port mirroring

Port mirroring can be useful for monitoring and troubleshooting the network. The Ethernet switch module enables you to copy packets that were transmitted and received on one of the eight internal ports and redirect a copy of these packets to one of the four external ports. You can connect a monitoring-and-troubleshooting device, such as a sniffer or an RMON probe, to the external port and view details about these packets.

Simple Network Management Protocol (SNMP)

The Simple Network Management Protocol (SNMP) is an open system interconnection (OSI) layer 7 (application layer) protocol for remotely monitoring and configuring network devices. SNMP enables network management stations to read and modify the settings of gateways, routers, switches, and other network devices. You can use SNMP to perform many of the same functions on a directly connected console, or you can use it within an integrated network-management software package such as IBM NetView®. SNMP provides the following capabilities:

- Sending and receiving SNMP packets using IP
- Collecting information about the status and current configuration of network devices
- Modifying the configuration of network devices

The switch module has a software program, called an *agent*, that processes SNMP requests, but the user program that makes the requests and collects the responses runs on a separate management station (a designated computer on the network). The SNMP agent and the user program both use the user datagram protocol/Internet protocol (UDP/IP) to exchange packets.

Authentication

The authentication protocol ensures that both the SNMP agent in the switch module and the remote user SNMP application program discard packets from unauthorized users. Authentication is accomplished by using community strings that function as passwords. The remote user SNMP application and the SNMP agent of the switch module must use the same community string. SNMP community strings of up to 20 characters can be entered using the **snmp community** CLI commands, that are described in the *CLI Reference Guide* on the IBM *BladeCenter T Documentation* CD.

Switching concepts

This section introduces the concepts and protocols that are related to the switching functionality of the Ethernet switch module.

Packet forwarding

The switch module uses a forwarding table to store the information that it collects about the location of devices on the network. The table holds destination MAC addresses and the destination port numbers of these devices. Packets that are sent to known addresses are transmitted only through relevant destination ports, thus reducing network traffic. For example, if port 1 receives a packet that is destined for a station on port 2, the switch module transmits that packet through port 2 only and transmits nothing through the other ports. Creating the table is referred to as learning the network topology.

An aging timer is used to make sure that the table is updated if devices are moved. Dynamic entries, which are learned by the switch by observing network traffic, are deleted from the table if they are not accessed within the aging time. Static entries that are entered by a network administrator are not subject to the aging process.

The aging time can be from 10 through 1000000 seconds, with a default value of 300 seconds. Setting the value too high might cause some entries in the table to become out of date. This might cause the switch module to make incorrect packet-forwarding decisions. If the aging time is too short, entries might be aged out too soon and have to be relearned. While the entries are being relearned, received packets whose source addresses cannot be found in the forwarding table will be transmitted through all of the ports on the switch, increasing the network traffic.

Spanning Tree Protocol (STP)

The Institute of Electrical and Electronics Engineers (IEEE) 802.1D Spanning Tree Protocol (STP) enables the blocking of links between switches that form loops within the network. When multiple links between switches are detected, a primary link is established. Duplicated links are blocked from use and become standby links. The protocol allows the duplicate links to be used if a primary link fails. When the STP is configured and enabled, primary links are established, and duplicated links are blocked automatically. The reactivation of the blocked links, at the time of a primary link failure, is automatic.

This automatic network reconfiguration provides maximum uptime to network users. It is possible to cause serious degradation of the performance of the network if the spanning tree is incorrectly configured. Read the following information before making any changes from the default values.

The switch-module STP performs the following functions:

- Creates a single spanning tree from any combination of switching or bridging elements
- Automatically reconfigures the spanning tree to compensate for the failure, addition, or removal of any element in the tree
- Reconfigures the spanning tree without operator intervention

Incorrect configuration of the switch-module external ports or incorrect cabling of the external ports to another switch device can create duplicate links that might cause network loops. Consult your network administrator for details about the configuration requirements for your system.

The single spanning tree that is created by the Spanning Tree Algorithm is referred to as the Common Spanning Tree (CST) in some of the commands that are described in this document.

The original Spanning Tree Algorithm that is defined in IEEE 802.1D has been updated to allow for faster reconfiguration in the event of a change to network topology or configuration parameters. This new protocol is defined in IEEE 802.1w as Rapid Reconfiguration and is based on the ability of the bridging device to recognize ports that are full-duplex and ports that are connected directly to end stations. The IEEE 802.1 standards committee recommends the use of IEEE 802.1w in preference to IEEE 802.1D, except when you are running certain protocols (such as LLC2 and NETBEUI) that are sensitive to the slightly increased probability of frame misordering. The Ethernet switch module defaults to IEEE 802.1D operation, but it can be configured to use the algorithm and protocols that are defined in IEEE 802.1w instead.

IEEE 802.1D has been further revised in IEEE 802.1s, which incorporates IEEE 802.1w and defines a multiple Spanning Tree Protocol along with an IEEE 802.1D compatibility mode. The Ethernet switch module defaults to IEEE 802.1D compatibility mode operation, but it can be configured to use the algorithm and protocols that are defined in IEEE 802.1w instead. Where this document refers to IEEE 802.1D, note that the reference is to IEEE 802.1D compatibility mode.

For additional information about both forms of the Spanning Tree Protocol, see the *CLI Reference Guide* on the IBM *BladeCenter T Documentation CD*.

Virtual Local Area Networks (VLAN)

A virtual local area network (VLAN) is a network topology configured according to a logical scheme rather than the physical layout. VLANs can be used to combine any collection of blade servers into an autonomous user group that appears as a group within one or more BladeCenter T units. VLANs also logically segment the blade servers into different broadcast domains so that packets are forwarded only between blade servers and the four external ports within the VLAN. Typically, although not necessarily, a VLAN corresponds to a particular subnet.

VLANs can enhance performance by conserving bandwidth and improve security by limiting traffic to specific domains.

Notes about VLANs on the Ethernet switch module

No matter what basis is used to uniquely identify blade servers and assign these nodes VLAN membership, packets *cannot* cross VLANs without a network device performing a routing function between the VLANs.

The switch module supports only IEEE 802.1Q VLANs. The port untagging function can be used to remove the 802.1Q tag from packet headers to maintain compatibility with devices that are tag-unaware.

The switch module default is to assign all blade servers and the four external ports to a single 802.1Q VLAN named DEFAULT with a VLAN ID (VID) of 1.

The switch module can be configured to enable a wide variety of VLAN configurations among the various external ports.

IEEE 802.1Q VLANs

The following terms are relevant to VLANs and important with respect to understanding how VLANs function:

Tagging	The act of adding 802.1Q VLAN information to the header of a packet.
Untagging	The act of stripping 802.1Q VLAN information out of the packet header.
Ingress port	A port on a switch where packets are flowing into the switch and where VLAN decisions must be made.
Egress port	A port on a switch where packets are flowing out of the switch, either to another switch or to an end station, and where tagging decisions must be made.

The Ethernet switch module implements IEEE 802.1Q VLANs, which requires tagging. This enables them to span the entire network (provided that all switches on the network are IEEE 802.1Q-compliant).

VLANs enable a network to be segmented to reduce the size of broadcast domains. All packets entering a VLAN will be forwarded (over IEEE 802.1Q enabled switches) only to the stations that are members of that VLAN. This includes broadcast packets, multicast packets and unicast packets from unknown sources.

VLANs can also provide a level of security to your network. IEEE 802.1Q VLANs will deliver packets only between stations that are members of the VLAN.

Any port can be configured as either tagging or untagging. The untagging feature of IEEE 802.1Q VLANs enables VLANs to work with legacy switches that do not recognize VLAN tags in packet headers (tag-unaware devices). The tagging feature enables VLANs to span multiple 802.1Q-compliant switches through a single physical connection and enables the Spanning Tree Protocol to be enabled on all ports and work normally.

The IEEE 802.1Q standard restricts the forwarding of untagged packets to the VLAN of which the receiving port is a member.

The main characteristics of IEEE 802.1Q are as follows:

- Assigns packets to VLANs by filtering
- Assumes the presence of a single global spanning tree
- Uses an explicit tagging scheme with one-level tagging

IEEE 802.1Q VLAN packet forwarding

The switch module makes packet-forwarding decisions based on the following types of rules:

Ingress rules Rules relevant to the classification of received packets belonging to a VLAN.

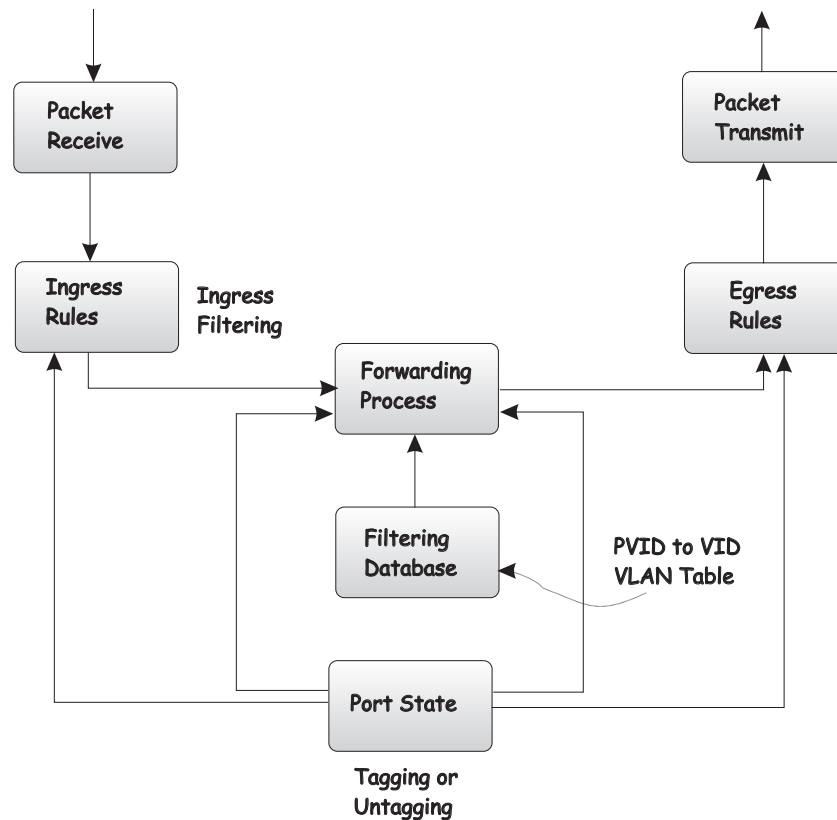
Forwarding rules between ports

The switch module decides whether to filter or forward the packet.

Egress rules The switch module determines whether the packet must be sent tagged or untagged.

The following illustration shows the 802.1Q VLAN packet-forwarding decision-making process of the switch module. For more information about packet forwarding, see “Packet forwarding” on page 17. For more information about port VLAN IDs (PVIDs), see “Port VLAN ID” on page 22. For more information about tagging and untagging, see “Tagging and untagging” on page 23. For more information about port states, see the *CLI Reference Guide* on the IBM *BladeCenter T Documentation CD*.

802.1Q Packet Forwarding



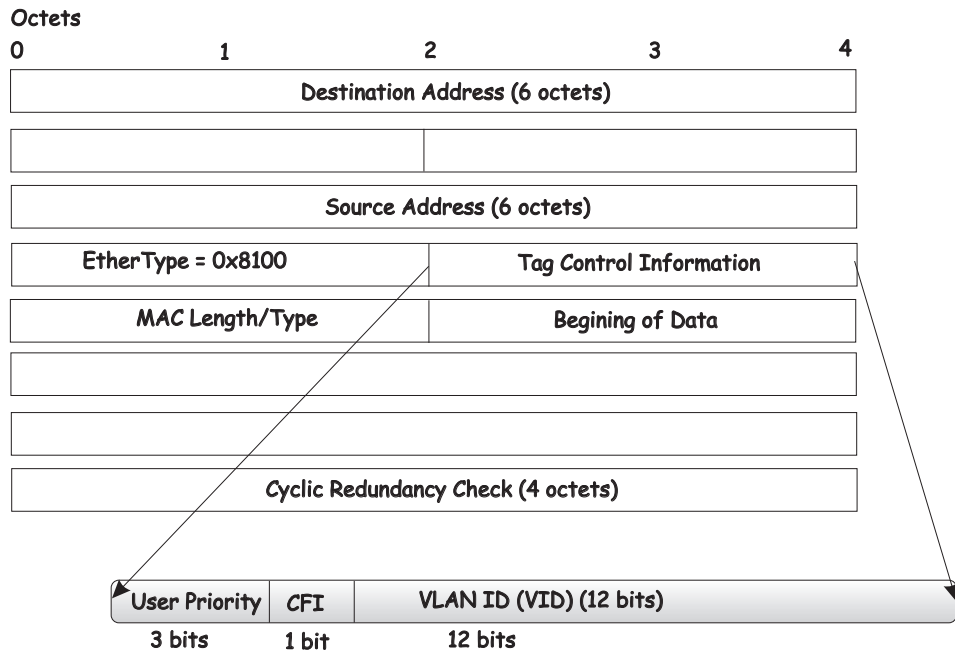
IEEE 802.1Q VLAN tags

The following illustration shows the 802.1Q VLAN tag. Four additional octets are inserted between the source MAC address and the packet's EtherType field. Their presence is indicated by a value of 0x8100 in the two bytes following the MAC address, in the VLAN tag's EtherType field, indicating that the packet carries an IEEE 802.1Q/802.1p tag. The tag is contained in the following 2 octets and consists of 3 bits of user priority, 1 bit of Canonical Format Identifier (CFI)¹ and 12 bits of VLAN ID (VID). The 3 bits of user priority are used according to the protocols that are defined in IEEE 802.1p (now part of IEEE 802.1D). The VID is the VLAN identifier and its use is defined by the 802.1Q standard. Because the VID is 12 bits long, 4094 unique VLANs can be identified.

The tag is inserted into the packet header, increasing the length of the entire packet by 4 octets. All of the information that was originally contained in the packet is retained.

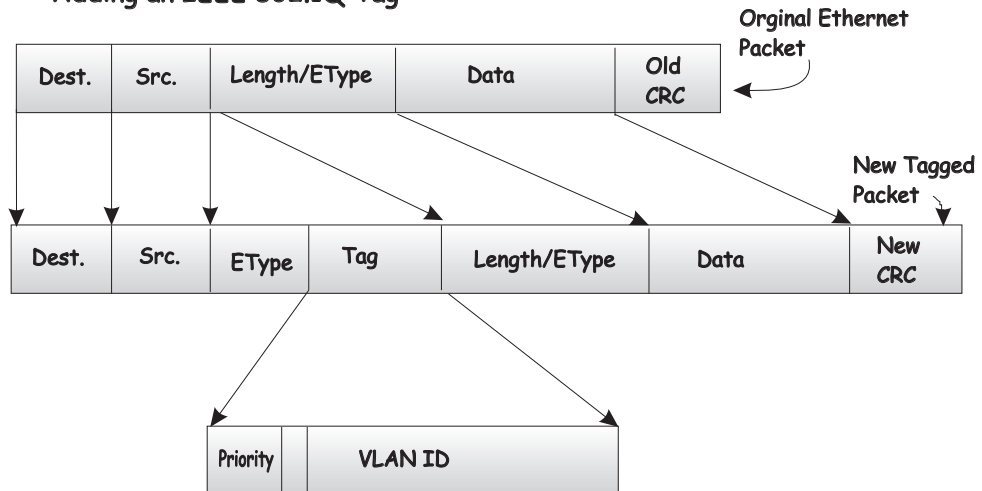
1. CFI is used for encapsulating Token Ring packets so that they can be carried across Ethernet backbones.

IEEE 802.1Q Tag



The **EtherType** and **VLAN ID** are inserted after the MAC source address, but before the original **EtherType/Length** or **Logical Link Control**. Because the packet is now longer than it was originally, the cyclic redundancy check (CRC) must be recalculated.

Adding an IEEE 802.1Q Tag



Port VLAN ID

Packets that are tagged (are carrying the 802.1Q VID information) can be transmitted from one 802.1Q compliant network device to another with the VLAN information intact. This enables 802.1Q VLANs to span network devices (and indeed, the entire network, if all network devices are 802.1Q compliant).

Not all network devices are 802.1Q compliant. These devices are referred to as tag-unaware. 802.1Q devices are referred to as tag-aware.

Before the adoption of 802.1Q VLANs, port-based and MAC-based VLANs were in common use. These VLANs relied upon a port VLAN ID (PVID) to forward packets. A packet received on a port would be assigned that port PVID and then be forwarded to the port that corresponded to the packet destination address (found in the switch forwarding table). If the PVID of the port that receives the packet is different from the PVID of the port that is to transmit the packet, the switch module will drop the packet.

A switch port can have only one PVID but it can have as many VLANs as the switch module has memory in its VLAN table to store them.

Tagging and untagging

Every port on an 802.1Q compliant switch can be configured to admit or discard packets that are received without a tag. Untagged packets that are admitted will be tagged with the port's PVID.

Every port on an 802.1Q compliant switch can also be configured to transmit packets with or without tags. Ports with tagging enabled will leave the 802.1Q tag received with the packet or inserted by the ingress port unchanged. Ports with untagging enabled will strip the 802.1Q tag from all packets that it transmits. Untagging is used to send packets from an 802.1Q-compliant network device to a noncompliant one.

Ingress filtering and egress rules

A port on a switch where packets are flowing into the switch and VLAN decisions must be made is referred to as an ingress port. If ingress filtering is enabled for a port, the switch will examine the VLAN information in the packet header (if present) to decide whether to forward the packet. If ingress filtering is disabled, packets will not be dropped based on their VLAN classification.

If ingress filtering is enabled and the packet is tagged with VLAN information, the ingress port will determine whether the ingress port itself is a member of the tagged VLAN. If it is not, the packet will be dropped. If the ingress port is a member of the 802.1Q VLAN, the packet is passed to the forwarding function.

If the packet is not tagged with VLAN information, the ingress port will tag the packet with its own PVID as a VID (if the port is configured to accept untagged packets) and pass it to the forwarding function. The forwarding function determines the destination port. If the destination, or egress, port is a member of the same VLAN as the packet the destination port transmits the packet on its attached network segment. If the egress port is not a member of the VLAN, the packet is dropped.

IEEE 802.1Q VLAN configuration

The switch module initially configures one VLAN (VID = 1) named DEFAULT. The factory default setting assigns all ports on the switch module to VLAN 1. As new VLANs are configured, their respective member ports are removed from VLAN 1. In addition, the VLAN ID value of 4095 is reserved for internal use. Following is additional configuration information:

- Packets cannot cross VLANs. If a member of one VLAN is to connect to a member of another VLAN, the link must be through an external router.
- All VLANs must be configured before the IP interface is set up. An IP addressing scheme must then be established and implemented when the IP interface is set up on the switch module.
- If no VLANs are configured on the switch module, all packets will be forwarded to any destination port. Packets with unknown source addresses will be flooded to all ports. Broadcast and multicast packets will also be flooded to all ports.
- An IP interface consists of two parts: a subnet mask and a network address. You will need to use this information when you configure the IP interface of the switch module.

Protocol-based VLANs (PBVLANS)

The main purpose of Protocol-based VLANs (PBVLANS) is to selectively process packets based on their upper-layer protocol by setting up protocol-based filters. Packets are bridged through user-specified ports based on their protocol.

In PBVLANS, the VLAN classification of a packet is based on its protocol (IP, IPX, etc.). PBVLANS help optimize network traffic because protocol-specific broadcast messages are sent only to end stations using that protocol. End stations do not receive unnecessary traffic and bandwidth is used more efficiently. It is a flexible method that provides a logical grouping of users. An IP subnet or an IPX network, for example, can each be assigned its own VLAN.

In port-based VLAN classification, the Port VLAN Identifier (PVID) is associated with the physical ports. The VLAN ID (VID) for an untagged packet is equal to the PVID of the port. In port- and protocol-based VLAN classifications, multiple VIDs are associated with each of the physical ports. Each VID is also associated with a protocol. The ingress rules used to classify incoming packets include the use of the packet's protocol in addition to the PVID to determine the VLAN to which the packet belongs. This approach requires one VID on each port for each protocol for which the filter is desired.

To configure PBVLAN support perform the following steps:

- Create and name a group.
- Assign one or more of the protocols – IP, IPX or ARP – to the group.
- Assign a VID to it.
- Specify the port(s) to which it applies.

If a tagged packet is received on a port in a PBVLAN group it will be processed using normal IEEE 802.1Q rules. If an untagged or priority-tagged packet is received, and the port is a member of a group with the matching protocol, the packet will be assigned the group's VID, otherwise it will be dropped. If no VID has been configured for the group, the PVID will be used.

Static MAC filtering

Static MAC Filtering allows you to add a small number (in the order of hundreds) of unicast or multicast MAC addresses directly to the forwarding database. Associated with each Static MAC address is a set of destination ports and VLAN information.

Any packet with a particular Static MAC Address in a particular VLAN is admitted only if the ingress port is in the set of source ports, otherwise the packet is dropped. On the egress side the packet, if admitted, is sent out of all the ports that are in the set of destination ports.

Upon ingress, each packet's destination MAC address is compared against the forwarding database. If the address is not in the table, the packet is flooded within the VLAN. If the address is in the table, then it is checked to see if it has been defined as a filter. If the MAC address is not defined as a filter, forwarding is performed as a normal parced address.

If the specific destination MAC address is defined as a filter, the packet is forwarded to the set of destination ports that are defined in the filter.

Static entries are never aged and can be removed only by user command.

Note: Even though the above discussion pertains to the forwarding database, MAC filters are not configured and displayed as part of the forwarding database; they are configured and displayed separately.

Generic Attribute Registration Protocol (GARP)

This protocol is used to exchange information between GARP participants to register and de-register attribute values within a bridged LAN. When a GARP participant declares or withdraws an attribute, the attribute value is recorded with the applicant state machine for that attribute for the port from which the declaration or withdrawal was made. Registration occurs only on ports that receive the GARP PDU containing a declaration or withdrawal. De-registration occurs only if all GARP participants connected to the same LAN segment as the port withdraw the declaration.

GARP VLAN Registration Protocol (GVRP)

GVRP (GARP VLAN Registration Protocol) is used to propagate VLAN membership information throughout the network. GVRP is based on the Generic Attribute Registration Protocol (GARP), which defines a method of propagating a defined attribute (i.e. VLAN membership) throughout the network. GVRP allows both end stations and the switch module to issue and revoke declarations relating to membership in VLANs. The Ethernet switch module complies with the specifications in IEEE 802.1D and IEEE 802.1Q.

End stations that participate in GVRP register VLAN membership via GARP Protocol Data Unit (GPDU) messages. Networking devices that implement the GVRP protocol and enable GVRP then process the GPDUs. The VLAN registration is made in the context of the port that receives the GPDU. The switch module propagates this VLAN membership on all of its other ports in the active topology. Thus, the end station's VLAN ID is propagated throughout the network.

GARP Multicast Registration Protocol (GMRP)

Networking devices use the GARP Multicast Registration Protocol to dynamically register (and de-register) Group membership information with other networking devices attached to the same segment and across all the bridged LAN devices that support Extended Filtering Services.

The operation of GMRP relies upon the services provided by the GARP. The information registered, de-registered and disseminated via GMRP is in the following forms:

Group Membership Information

This indicates that there exists one or more GMRP participants which are members of a particular Group, and carries the group MAC address(es) associated with this Group. Registration of group membership information allows networking devices to be made aware that frames destined for these group MAC address(es) should be forwarded in the direction of registered members of the group. Forwarding of frames destined for the group MAC address(es) occurs on ports on which such membership registration has been received.

Group Service Requirements Information

This indicates that one or more GMRP participants require Forward all Groups or Forward Unregistered to be the default filtering behavior. Registration of group services requirement information allows networking devices to be made aware that any of their ports that can forward frames in the direction from which the group service requirement information has been received should modify their default group behavior in accordance with the group service requirement.

When the switch module receives GMRP PDUs it will update the multicast table with a new entry or modify an existing entry with the new information. The switch module will forward multicast packets through only those ports for which GMRP has created a group registration entry (for that multicast address).

GMRP registrations are specific to a VLAN, which allows the Group filtering behavior for one VLAN to be independent of the Group filtering behavior for other VLANs. The same ingress rules are applied to GMRP PDUs as to other packets. Therefore:

- GMRP frames with no VLAN classification (i.e., untagged or priority-tagged GMRP frames) are discarded if the Acceptable Frame Types parameter for the Port is set to Admit Only VLAN-tagged frames. Otherwise, they are classified according to the PVID (Port VLAN ID) for the Port.
- VLAN-tagged GMRP frames are classified according to the VID carried in the tag header.
- If Ingress Filtering is enabled, and if the Port is not in the Member set for the GMRP frame's VLAN classification, then the frame is discarded.

The VLAN classification thus associated with received GMRP PDUs establishes the VLAN context for the received PDU, and identifies the GARP participant instance to which the PDU is directed. GMRP PDUs transmitted by GMRP participants are VLAN-classified according to the VLAN context associated with that participant. GMRP Participants in VLAN networking devices apply the same egress rules that are defined for the transmission port. Therefore:

- GMRP PDUs are transmitted through a port only if the port is a member of the VLAN concerned.
- GMRP PDUs are transmitted as VLAN-tagged frames or untagged frames, in accordance with the state of the Untagged Set for that port for the VLAN concerned. Where VLAN-tagged frames are transmitted, the VID field of the tag header carries the VLAN Context Identifier value.

Internet Group Management Protocol (IGMP) snooping

Internet Group Management Protocol (IGMP) snooping is a feature that allows a switch to forward multicast traffic intelligently on the switch. Multicast IP traffic is traffic destined to a host group. Host groups are identified by class D IP addresses, which range from 224.0.0.0 to 239.255.255.255. Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic. This prevents the switch from broadcasting the traffic to all ports and possibly affecting network performance.

Note that the IP address range 224.0.0.1 through 224.0.0.255 is reserved for routing protocols and other low-level topology discovery or maintenance protocols. For example, the address 224.0.0.1 is the “all hosts” address, and 224.0.0.2 indicates all routers on this subnet. Also, only the least significant 23 bits of the IP address are mapped to MAC addresses, so, for example, 225.0.0.123 and 239.128.0.123 and similar IP multicast addresses all map to MAC address 01-00-5E-00-00-7B (for Ethernet). Therefore, a switch using IGMP Snooping can collapse the IP multicast group memberships into a single Ethernet multicast group.

A traditional Ethernet network can be physically separated into different network segments to prevent overload of the shared media. Bridges and switches connect these segments. When a packet with a broadcast or multicast destination address is received, the switch will forward a copy into each of the remaining network segments in accordance with IEEE 802.1D. Eventually, the packet is made accessible to all nodes connected to the network.

This approach works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach can lead to less efficient use of network bandwidth, particularly when the packet is intended for only a small number of nodes. Packets will be flooded onto network segments where no node has any interest in receiving the packet. The problem of wasting bandwidth is even worse when the LAN segment is not shared, for example in full duplex links.

Allowing switches to snoop IGMP packets is one way to solve this problem. The switch uses the information in the IGMP packets as they are being forwarded throughout the network to determine which segments should receive packets directed to particular group addresses.

Group addresses are stored in the Multicast Forwarding Database (MFDB). An IGMP address will be removed from the database if a report for it is not received within the query interval. An interface can be removed from an IGMP group in response to an IGMP Leave Group message.

Link Aggregation (LAG)

The Ethernet switch module supports Link Aggregation (LAG), or port trunking. Port trunks (aggregated ports) can be used to increase the bandwidth of a network connection or to ensure fault recovery. You can configure up to two trunk connections (combining two to four ports into one fat pipe) between any two BladeCenter T units or other Layer 2 switches. However, before making any physical connections between devices, use the Link Aggregation commands to specify the ports that will belong to the trunking group on both switches.

When using a port trunk, note that:

- The ports used in a trunk must all be of the same speed (100 Mbps or 1000 Mbps) and operate in full-duplex mode only.
- The ports that can be assigned to the same trunk have certain other restrictions, as described in this section.
- Each port can be assigned to only one trunk group, either a static or dynamic group.
- The ports at both ends of a connection must be configured as trunk ports.
- All of the ports in a trunk have to be treated as a whole when moved from/to, added, or deleted from a VLAN.
- The Spanning Tree Protocol (STP) will treat all the ports in a trunk as a whole.
- Enable the trunk before connecting any cable between the switches to avoid creating a data loop.
- Disconnect all trunk port cables or disable the trunk ports before removing a port trunk to avoid creating a data loop.

Trunking can be set as a static or a dynamic port/group using the IEEE 802.3ad Link Aggregation commands. When trunking is enabled, a blue border will be placed around the ports on the link-status panel. See page 65 for more information about the link-status panel.

Static LAGs

When you create a LAG, the member links will attempt to exchange LACPDU with their partners. If a link does not receive a LACPDU within 3 seconds, it will come up with default values. If a LACPDU is later received with different values, the link will drop out of the LAG. When all member links have dropped out, the LAG will reconfigure itself with the new values from the received LACPDUs.

It is important that when you configure LAGs, you should configure the LAGs and enable STP on both partner devices before connecting the cables.

Distribution method

Link aggregation, or port trunking, enables several ports to be grouped together and to act as a single link. This gives a bandwidth that is a multiple of a single-link bandwidth.

Port trunking is most commonly used to link a bandwidth intensive network device or devices, such as a server, to the backbone of a network.

The switch module offers link aggregation on four external ports for up to two static trunk groups or two LACP 802.3ad link aggregation groups. The trunked ports can be non-continuous (that is, have non-sequential port numbers). All of the ports in the group must be members of the same VLAN. In addition, the trunked ports must connect at the same speed in full-duplex mode.

Load balancing is automatically applied to the ports in the trunked group, and a link failure within the group causes the network traffic to be directed to the remaining links in the group.

The STP will treat a port trunking group as a single link on the switch level. On the port level, the STP will use the port parameters of the Master Port in the calculation of port cost and in determining the state of the port trunking group. If two redundant port trunking groups are configured on the switch module, STP will block one entire group in the same way STP will block a single port that has a redundant link.

Dynamic Host Configuration Protocol (DHCP)

The Dynamic Host Configuration Protocol (DHCP) can reduce the administrative burden of assigning and maintaining IP address information. DHCP provides reliable and simple TCP/IP network configuration, ensures that address conflicts do not occur, and helps to conserve the use of IP addresses through centralized management of address allocation.

Dynamic address allocation enables a client to be assigned an IP address from a pool of free addresses. Each address is assigned with a lease and a lease expiration period. The client must renew the lease to continue using the assigned address. Dynamically assigned addresses can be returned to the free address pool if the computer is not being used, if it is moved to another subnet, or if its lease expires. Usually, network policy ensures that the same IP address is assigned to a client each time and that addresses returned to the free address pool are reassigned.

When the address lease expires, the DHCP client enters the renewing state. The client sends a request message to the DHCP server that provided the address. The DHCP server sends an acknowledgement that contains the new lease and configuration parameters. The client then updates its configuration values and returns to the bound state.

When the DHCP client is in the renewing state, it must release its address immediately in the rare event that the DHCP server sends a negative acknowledgment. The DHCP server sends this message to inform a client that it has incorrect configuration information, forcing it to release its current address and acquire new information.

If the DHCP client cannot successfully renew its lease, the client enters a rebinding state. The client then sends a request message to all DHCP servers in its range, attempting to renew its lease. Any DHCP server that can extend the lease sends an acknowledgment containing the extended lease and updated configuration information. If the lease expires or if a DHCP server responds with a negative acknowledgment, the client must release its current configuration and then return to the initializing state.

If your DHCP client uses more than one network adapter to connect to multiple networks, this protocol is followed for each adapter that you want to configure for TCP/IP. Multi-homed systems are selectively configured for any combination of system interfaces.

When a DHCP-enabled computer is restarted, it sends a message to the DHCP server with its current configuration information. The DHCP server either confirms this configuration or sends a negative reply so that the client must begin the initializing state again. System startup might therefore result in a new IP address for a client computer, but neither the user nor the network administrator has to take any action in the configuration process.

Before loading TCP/IP with an address acquired from the DHCP server, DHCP clients check for an IP address conflict by sending an Address Resolution Protocol (ARP) request containing the address. If a conflict is found, TCP/IP does not start, and the user receives an error message. The conflicting address should be removed from the list of active leases, or it should be excluded until the conflict is identified and resolved.

Routing concepts

This section introduces the concepts and protocols relevant to the routing functionality of the Ethernet switch module.

IP mapping

The IP mapping component provides the following functions:

- The ARP Mapping Layer
- The routing table used by all routing protocols
- The IP Forwarding Layer that forwards all received IP packets

Routing Information Protocol (RIP)

The Routing Information Protocol, or RIP, is a long-standing protocol used by routers for exchanging route information. RIP is a distance vector protocol where each route is characterized by the number of gateways, or hops, a packet must traverse to reach its intended destination. Categorized as an interior gateway protocol, RIP operates within the scope of an autonomous system. RIP is a simple protocol, and its usefulness is limited to moderate sized networks whose physical interconnections are of similar type and speed. The Berkeley Unix 'routed' program is considered the standard RIP implementation.

Two versions of RIP are currently defined:

RIPv1:

RIPv1 is defined in RFC 1058 and establishes the base operating characteristics for the protocol. It does not include the concept of subnets, and routing messages are sent to all stations attached to the physical medium.

RIPv2:

RFC 1723 describes enhancements to the first RIP version and is known as RIPv2. These enhancements place additional information in the RIP routing messages (such as a subnet mask field) and add an authentication scheme to improve security for updating route tables. RIPv2 routing messages are sent to a specific multicast address to reduce the processing burden on systems not participating in this protocol. The following enhancements are described in RFC 2453:

- An implementation of RIP must use split horizon and should use split horizon with poisoned reverse (under management control) to avoid routing loops.

- An implementation of RIP must implement triggered update for deleted routes and can implement triggered updates for new routes or changes of routes. RIP implementations must also limit the rate at which triggered updates can be transmitted.
- An implementation of RIP should support host routes.

RIP is designed such that its routers send the contents of their routing table every 30 seconds to each adjacent router. These periodic updates allow routers to determine which routes remain active in the route table; absence of a route from the updates causes the route to be declared unusable after 180 seconds have elapsed, and to be removed from the table after an additional 120 seconds passes without the route appearing in an update message.

The Ethernet switch module supports both RIPv1 and RIPv2, as well as the relevant MIBs.

Open Shortest Path First (OSPF)

The major alternative to RIP is the Open Shortest Path First (OSPF) protocol, which is used within larger autonomous networks in preference to RIP. OSPF is a link-state protocol that multicasts table updates only when a change has taken place and transmits only the changed portion of the table. To give preferences to certain routes, OSPF uses both administratively assigned costs for a router and link-states as metrics. In addition, OSPF supports variable-length subnet masks.

OSPF can operate within a hierarchy. The largest entity within the hierarchy is the autonomous system (AS), a collection of networks under a common administration sharing a common routing strategy. This is sometimes called a routing domain. An AS can be divided into a number of areas or groups of contiguous networks and attached hosts. Routers within the same area share the same information, so they have identical topological databases. Information is sent in the form of link-state advertisements (LSAs) to all other routers within the same hierarchical area. An area's topology is not visible to routers outside the area.

Because of area partitioning there are two types of OSPF routing:

Intra-area routing

When a source and destination are in the same area.

Inter-area routing

When a source and destination are in different areas, in which case an OSPF backbone is required.

The Ethernet switch module supports OSPF version 2 in accordance with RFC 2328, together with a compatibility mode for the RFC 1583 OSPF specification, which allows interoperability with OSPF version 2 routers using the older implementation.

Note: OSPF RFC 1583 compatibility mode is enabled by default. All routers in a domain should have the same 1583 compatibility mode setting to prevent routing loops.

VLAN routing

The Ethernet switch module incorporates both 802.1Q VLAN bridging and routing functions. The internal bridging function can be an interface to the routing function and the routing function can be an interface to the bridging function.

The VLAN Routing function has the following characteristics:

- From the perspective of the network, the VLAN function and the router function act as two independent entities.
- VLANs can be ports, or nodes, on the internal router function. To the internal router function, a VLAN for which routing is enabled looks just like any other routing interface.
- The internal router function can be a port on the internal bridge function. To the internal bridge function, the internal router function is treated largely as any other bridge interface.
- A port cannot operate as both a router port and an 802.1Q bridge port.

Packet processing is split into two: processing at the bridge layer and processing at the router layer.

Packet processing at the bridge layer

Even with the addition of VLAN routing, all normal 802.1Q processing (ingress rules, VLAN association, etc.) is performed if the interface on which the packet arrived is enabled for bridging. On ingress, each packet is associated with a VLAN and admitted or dropped accordingly. If an admitted packet is a unicast packet, it is forwarded to the interface with which the destination address is associated in the Layer 2 unicast MAC address table. With the advent of VLAN routing, that interface can now be the internal bridge-router interface, or IBRI. When VLAN routing is configured, the MAC address associated with the internal bridge-router interface will always be placed in the Layer 2 unicast MAC tables if any VLAN is enabled for routing.

If the packet is a multicast packet, it is forwarded to all ports in the VLAN, and to the IBRI, if and only if, the packet is received on a VLAN which is routed. The IBRI has to be a recipient of multicast packets as the router might process these packets. The network behavior will be the same as if the router were a host attached to one of the switch's ports.

When the bridge forwards a packet to the IBRI, layer 2 egress rules are essentially ignored. The end effect is as if the IBRI were a member of all VLANs from an egress perspective.

Packet processing at the routing layer

At the routing layer, a VLAN enabled for routing is viewed largely as any other interface. A packet is associated with the VLAN on which it arrived if the frame arrived on a VLAN which was routed and is a routable frame. Otherwise, it is associated by the router with the physical interface on which it arrived. The configuration parameters of the routing interface, either the physical interface or the VLAN routing interface, apply.

Virtual Router Redundancy Protocol (VRRP)

An end station running IP needs to know the address of its first hop router. While some network administrators choose to install dynamic router discovery protocols such as DHCP, others prefer to statically allocate router addresses. If the router identified by such a statically allocated address goes down, the end station loses connectivity. The Virtual Router Redundancy Protocol (VRRP) is designed to provide backup for the failing router without requiring any action on the part of the end station. It is based on the concept of having more than one router recognize the same IP address. One of the routers is elected the “master” router and handles all traffic sent to the specified virtual router IP address. If the master router fails, one of the backup routers will be elected in its place and will start handling traffic sent to the address. This change will be transparent to end stations.

VRRP increases the availability of the default path without requiring configuration of dynamic routing or router discovery protocols on every end station. The greater default path availability is accomplished by using any of the virtual router IP addresses on the LAN as the default first hop router for the end stations.

Multiple virtual routers can be defined on a single router interface on the Ethernet switch module, but only one IP address can be assigned to a virtual router.

BOOTP/ DHCP relay agent

In the majority of network configurations, BOOTP/DHCP clients and their associated servers do not reside on the same IP network or subnet. Therefore, some kind of third-party agent is required to transfer BOOTP/DHCP messages between clients and servers. Such an agent is known as a BOOTP/DHCP relay agent.

The BOOTP/DHCP relay agent relays BOOTP/DHCP requests between subnets. The agent relays requests from a subnet without a BOOTP/DHCP server to a server or next hop agent on another subnet.

Unlike a router which switches IP packets transparently, a BOOTP/DHCP relay agent processes BOOTP/DHCP messages and generates new BOOTP/DHCP messages as a result.

Router discovery

The router discovery messages do not constitute a routing protocol. Instead, the router discovery messages enable hosts to discover the existence of neighboring routers through the use of router advertisement. Router advertisement is unsuitable for determining the best route to a particular destination. If a host chooses a poor first-hop router for a particular destination, it should receive an ICMP Redirect from that router, identifying a better one.

In router discovery, a router periodically multicasts a Router Advertisement from each of its multicast interfaces, announcing the IP address(es) of that interface. Hosts discover the addresses of their neighboring routers simply by listening for advertisements. Since a host knows the address of the neighbors, the host can send IP datagrams beyond its directly attached subnet.

Typically, this is accomplished by one of two methods:

- Reading a list of one or more router addresses from a (possibly remote) configuration file at startup time.
- On multicast links, some hosts also discover router addresses by listening to routing protocol traffic.

Both of these methods have serious drawbacks. Configuration files must be maintained manually, which presents a significant administrative burden. Also, configuration files are unable to track dynamic changes in router availability. Eavesdropping on routing traffic requires that hosts recognize the particular routing protocols in use. The protocols in use will vary from subnet to subnet and are subject to change at any time. So eavesdropping requires that each host is able to recognize a range of protocols

When a host attached to a multicast link starts up, it can multicast a router solicitation to ask for immediate advertisements, rather than waiting for the next periodic ones to arrive. In the event no advertisements are forthcoming, the host can retransmit the solicitation a small number of times. Afterwards, the host must desist from sending any more solicitations. Any additional routers are eventually discovered by reception of their periodic (unsolicited) advertisements.

If you enable routing on your Ethernet switch module, support for the router discovery protocol will also be enabled.

Security

Local Area Networks (LANs) are often deployed in environments that permit the attachment of unauthorized devices. The networks also permit unauthorized users to attempt to access the LAN through existing equipment. In such environments, you might want to restrict access to the services offered by the LAN. This section introduces the concepts associated with the two forms of security available on the Ethernet switch module: Local Authentication and Remote Authentication Dial-In User Service (RADIUS). These mechanisms are used to authenticate user access to the switch module and conform to the specifications in IEEE 802.1X.

IEEE 802.1X

Port-based network access control makes use of the physical characteristics of LAN infrastructures to provide a means of authenticating and authorizing devices attached to a LAN port. Port-based network access control prevents access to the port in cases in which the authentication and authorization process fails.

Access control is achieved by enforcing authentication of entities seeking access to a port on the switch module. These entities are referred to as supplicants. The result of the authentication process determines whether the supplicant is authorized to access services on that controlled port.

A Port Access Entity (PAE) can adopt two different roles in an access control interaction:

Authenticator A port that enforces authentication before allowing access.

Supplicant A port that attempts to access services offered by an authenticator.

Additionally, there is a third role:

Authentication server

Performs the authentication function necessary to check the credentials of the Supplicant on behalf of the Authenticator.

All three roles are required to complete the authentication process.

The Ethernet switch module operates in the authenticator role only. The authenticator PAE is responsible for submitting information received from the supplicant to the authentication server in order for the credentials to be checked, which will determine the authorization state of the port. The authenticator PAE controls the authorized/unauthorized state of the controlled port depending on the outcome of the authentication process. Authentication messages use the Extensible Authentication Protocol (EAP).

A port can take one of two states:

Controlled Traffic will only be exchanged if the port is in the Authorized state.

Uncontrolled Allows the uncontrolled exchange of EAP over IEEE 802 LANs (EAPoL) PDUs between the Authenticator and Supplicant.

A controlled port is configured by management to be in one of three states:

ForceUnauthorized

The port is set to the unauthorized state.

ForceAuthorized

The port is set to the authorized state.

Auto

The port's state will be set based on the outcome of authentication exchanges between the Supplicant, Authenticator and the Authentication server. This is the default port state when port-based access control is enabled.

Local authentication

Local authentication matches a user ID/password combination received from the supplicant to the switch module's local database. The switch module will transmit an EAP-Request/Identity packet to the supplicant to obtain the combination, and if a match is found will then send an EAP-Request/MD5 packet to the supplicant. The supplicant's MD5 response is sent to the authenticator for validation. A match results in a successful authentication of the port.

Note: The switch module's Authenticator supports only the EAP-MD5 authentication type for local authentication.

RADIUS authentication

When Remote Authentication Dial-In User Service (RADIUS) authentication is used, the authenticator basically becomes a pass through to facilitate communication between the supplicant and the RADIUS server. The authenticator encapsulates the EAP messages exchanged between the supplicant and the server in either EAPoL or RADIUS frames (depending on the direction of the frame). The authenticator determines the authorization status of the port based on RADIUS Access-Accept or Access-Reject frames. The authenticator switch also needs to send and process all applicable RADIUS attributes.

Secure Shell (SSH)

Interactive login is widely used as a means to control and/or configure an entity across a network. For decades the Telnet protocol, and its cousin rlogin, have provided this capability. However, these protocols permit the transmission of sensitive information over unprotected networks. The current standard for providing interactive login in a secure fashion is the Secure SHell (SSH).

Table 2. Secure Shell Feature Details

SSH Feature	Component Type
Connection Type	Interactive Login
Authentication Method	Password
Ciphers	<ul style="list-style-type: none">• 3DES-CBC• Blowfish-CBC• Twofish128-CBC• AES128-CBC
Hash Algorithms	<ul style="list-style-type: none">• MD5• SHA-1• SHA-1-96
Key Exchange Methods	Diffie-Hellman
Compression Algorithms	<ul style="list-style-type: none">• zlib• none (i.e. no compression)
Public Key Algorithms	<ul style="list-style-type: none">• SSH-DSA• SSH-RSA
SSH Protocol Versions	<ul style="list-style-type: none">• SSH 2.0• SSH 1.5

Secure Socket Layer (SSL)

Managing devices with a web browser has been standard practice for several years. Unfortunately standard HTTP transactions are no more secure than Telnet. The solution is the use of the Secure Sockets Layer (SSL) protocol which provides a means of abstracting an encrypted connection between two stations. Once established, such a connection is virtually no different to use than an unsecured connection. This allows an established protocol (e.g. HTTP) to operate in a secure manner on an open network.

Table 3. Secure Sockets Layer Details

SSL Feature	Component Type
Protocols Secured	HTTP
Ciphers	<ul style="list-style-type: none">• RC4• DES• 3DES
Hash Algorithms	<ul style="list-style-type: none">• MD5• SHA-1
Key Exchange Methods	<ul style="list-style-type: none">• Diffie-Hellman• RSA
SSL Protocol Versions	<ul style="list-style-type: none">• TLS 1.0• SSL 3.0

Quality of Service (QoS)

The Quality of Service (QoS) features of the Ethernet switch module allow you to allocate network bandwidth according to the needs of the network users. This section will give you an overview of the methods available.

Quality of Service technologies are intended to provide guaranteed, timely, delivery of specific application data to a particular destination. In contrast, standard IP-based networks are designed to provide “best effort” data delivery service. Best effort service implies that the network will attempt to deliver the data in a timely fashion, although there is no guarantee. During times of congestion, packets can be delayed, sent sporadically, or dropped. For typical Internet applications, such as electronic mail and file transfer, a slight degradation in service is acceptable and in many cases is unnoticeable. Conversely, any degradation of service has undesirable effects on applications with strict timing requirements, such as voice or multimedia.

QoS is a means of providing consistent, predictable data delivery by distinguishing packets that have strict timing requirements from those that are more tolerant of delay. Packets with strict timing requirements are given special treatment in a QoS-capable network. To accomplish this, all elements of the network must be QoS-capable. If one node is unable to meet the necessary timing requirements, this creates a deficiency in the network path and the performance of the entire packet flow is compromised.

Bandwidth provisioning

Bandwidth provisioning allows you to deliver varying levels of allocated bandwidth to users sharing the same physical interface. By mapping a subscriber's traffic profile to a predefined policy and then actively provisioning the maximum bandwidth consumed by that subscriber, you can provide enhanced service offerings to your users. Bandwidth provisioning reduces the risk of network congestion and prevents a small number of applications or users from consuming all the available bandwidth.

Bandwidth provisioning provides Maximum Burst Rate (MBR) management for an interface and a flexible framework for defining and extending traffic classes. It allows you to allocate bandwidth by mapping a subscriber's traffic profile (e.g. source/destination IP address, traffic type) to a prescribed policy. Bandwidth provisioning actively provisions maximum bandwidth. For example, bandwidth provisioning can enable monitoring and management of bandwidth for VLAN traffic based on VLAN class IDs over an interface.

To run bandwidth provisioning you need to define Bandwidth Allocation Profiles (BAPs) and Traffic Classes (TCs), and then associate the two:

Bandwidth Allocation Profile

A transmission link definition which specifies a Bandwidth Bucket Identifier, as well as maximum bandwidth allowances.

Traffic Class The definition of the traffic to which a set of rules will apply. A class is defined by specifying a VLAN Identifier and an interface number, along with the class priority.

A default BAP, which you cannot modify, is assigned to all new TCs. Any given BAP can be assigned to multiple TCs. When you have defined the BAPs and TCs, and attached BAPs to the TCs, VLAN traffic on the specified interfaces will not exceed the maximum configured bandwidth.

Access Control Lists (ACL)

You use Access Control Lists (ACLs) to control the traffic entering or exiting a network, for example where two networks are connected, or an internal network is connected through a firewall router to the Internet. This allows you to ensure that only authorized users have access to specific resources while blocking off any unwarranted attempts to reach them.

You can use ACLs to:

- Provide traffic flow control
- Restrict the contents of routing updates
- Determine which types of traffic will be forwarded or blocked
- Provide network security

An ACL consists of one or more rules or filtering criteria. A packet is accepted or rejected based on whether or not it matches the criteria. After you create the set of rules for an ACL, you attach the ACL to a physical port. Filtering is done on inbound traffic.

An ACL rule can apply to any one or more of the following fields:

- Source IP address
- Source Port (Layer 4)
- Destination IP
- Destination Port (Layer 4)
- IP Protocol Number

An 'implicit deny' rule is added to the end of every ACL. This means that if a packet does not match any of the rules you have defined it will be dropped.

To apply an ACL to a LAG interface, the ACL must be applied to all member ports in a LAG.

IP multicast concepts

This section offers a high-level description of the multicast support provided by the Ethernet switch module.

Note: The functions that are described in this section are available on the BladeCenter T unit only. All of the functions rely on IGMP to provide basic routing information.

Internet Group Management Protocol (IGMP)

The Internet Group Management Protocol (IGMP) is used to communicate multicast routing information between multicast-capable hosts and routers within a subnet. When a router boots up, it transmits IGMP Query messages. If a router is already present in the network and acting as Querier, the new router will transition to Non-Querier Mode unless it has the lower IP address.

Hosts with applications that wish to subscribe to a multicast group send IGMP Group Membership Reports on receipt of query messages, and Leave Group Membership Reports to leave a group. The routers use these messages to construct and maintain the Multicast Forwarding Database. The router will not delete a group until all hosts have unsubscribed from the group or no host responds to a Group Specific Query.

If more than one interface on the router is connected to the same subnet, only one of those interfaces can be configured to:

- Listen for multicast reports to all IP hosts
- Listen for multicast reports from all IP hosts
- Transmit multicast packets to all IP hosts
- Receive multicast packets from all IP hosts

In addition, you cannot configure both IGMP and IGMP snooping on the same interface.

There are two versions of IGMP: IGMPv1 and IGMPv2. You must configure all routers in a subnet to run the same version of the protocol.

Distance Vector Multicast Routing Protocol (DVMRP)

Distance Vector Multicast Routing Protocol (DVMRP) is a dense mode multicast protocol and is applicable for use in networks where bandwidth is relatively plentiful and there is at least one multicast group member in each subnet. DVMRP assumes that all hosts are part of a multicast group until it is informed of multicast group changes. When the dense-mode multicast router is informed of a group membership change, the multicast delivery tree is pruned. DVMRP uses a distributed routing algorithm to build per-source-group multicast trees. It is also called Broadcast and Prune Multicasting protocol. It dynamically generates per-source-group multicast trees using Reverse Path Multicasting. Trees are calculated and updated dynamically to track membership of individual groups.

The DVMRP protocol operates as follows:

- The first message for any (source, group) pair is forwarded to the entire multicast network, with respect to the packet's time-to-live (TTL).
- TTL restricts the area to be flooded by the message.
- Any leaf router which does not have members on directly attached subnetworks sends back prune messages to the upstream router.
- The branch which transmitted a prune message is deleted from the delivery tree.

DVMRP transmits two types of messages:

Protocol Messages

Probe Messages, Prune Messages, Graft and Graft Acknowledgements, etc., which are sent as multicast packets. Most protocol messages are addressed to ALL-DVMRP-ROUTERS. If a router in the path does not support multicasting, it will encapsulate these packets in unicast IP packets. Encapsulated packets are tunneled through non-multicast routers and decapsulated by the multicast router at the end of the tunnel.

Routing Messages

These are unicast messages used to exchange routing information. They are typically sent by a neighboring router when it is restarted, has lost its routing information, or has come up from a down state. They are also transmitted when a graft is sent in response to a new member's subscription to an existing group, or when a new group is subscribed.

When messages arrive, the reverse path to the source of the message is discovered by examining the routing table. If the message arrived on the interface that would be used to transmit the message back to the source, the message is transmitted to downstream routers. Otherwise the message is not on the optimal delivery tree and so the packet is dropped. By doing this, loops and duplicates are filtered. DVMRP discovers its neighbors by periodically sending probe messages with a TTL of one (1). These probe messages contain the list of neighboring DVMRP routers from which a probe has been received. With this, a two-way neighbor relationship is established.

When two different routers are connected to the same multi-access network, there is a possibility of getting duplicate packets. This is overcome by nominating a Designated Forwarder (DF) for each network. The router with the least metric will be the designated forwarder. In case of a tie, the router with the lower IP address will be the designated forwarder. After a designated forwarder is elected, the multicast trees are built. If the destination group of a packet exists in the local database and the router on which the packet arrived is the designated forwarder, then that interface is added into the downstream interface list.

The edge routers remove interfaces with no group members from their multicast trees. If all the downstream interfaces are removed, the router sends a prune message to its upstream neighbors. Every prune message has a lifetime, after which the interface is joined back onto the delivery tree. If the unwanted datagrams still appear, the prune is initiated again. If a host joins a previously pruned branch of a tree, the DVMRP routers will use graft messages to cancel the previous prunes.

Protocol Independent Multicast - Dense Mode (PIM-DM)

Protocol Independent Multicast (PIM) protocols are not dependent on any particular unicast routing protocols to construct forwarding information for multicast packets, although unicast information is needed for forwarding packets. The Dense Mode (DM) version of PIM is applicable for networks with relatively plentiful bandwidth and with at least one multicast member in each subnet.

PIM-DM assumes that all hosts are part of a multicast group and forwards packets to hosts until informed that group membership has changed. A group membership change results in the multicast delivery tree being pruned.

The PIM-DM protocol operates as follows:

- The first message for any (source, group) pair is forwarded to the entire multicast network, with respect to the time-to-live (TTL) value in the packet.
- TTL restricts the area flooded by the packet.
- All leaf routers with no members in a directly attached subnet send prune messages to the upstream router.
- Any branch for which a prune message is received is deleted from the delivery tree.

PIM-DM uses Reverse Path Forwarding (RPF), which is the fundamental concept in multicast routing that enables routers to correctly forward multicast messages down the distribution tree. RPF makes use of the existing unicast routing table to determine the upstream and downstream neighbors and build a source-based shortest-path distribution tree. A router forwards a multicast message only if the multicast message is received on the upstream interface. This RPF check helps to guarantee that the distribution tree will be loop-free.

The multicast messages contain the source and group information so that downstream routers can build up their multicast forwarding tables. If the source goes inactive, the tree is torn down. Multicast messages arriving at a router over the proper receiving interface (i.e., the interface that provides the shortest path back to the source) are forwarded on all downstream interfaces until unnecessary branches of the tree are explicitly pruned. In addition to the prune messages, PIM-DM uses graft messages and assert messages. Graft messages are used when a new host wants to join a group, and assert messages are used to shut off duplicate flows.

Before you can enable PIM-DM on your Ethernet switch module, you must first enable routing and IGMP both for the switch module and for the relevant interface(s). The switch module implements PIM-DM Version 2.

Protocol Independent Multicast - Sparse Mode (PIM-SM)

As with PIM-DM, PIM-SM is not dependent on any particular unicast routing protocols to construct the forwarding information it uses for multicast packets, although unicast information is needed for actual forwarding. The Sparse Mode (SM) version of PIM is applicable for networks with relatively limited bandwidth and where group membership is widely distributed across regions.

Note: PIM-SM packet forwarding rates are severely limited since all packets are forwarded by the switch module's control processor.

Sparse mode protocols begin with the assumption that few routers in the network will be involved in any given multicast path. Sparse mode routers minimize network traffic by adding branches to the tree only when explicitly requested to do so. Therefore, sparse mode protocols such as PIM-SM are better suited to WANs than are dense mode protocols.

PIM-SM uses the following concepts:

Rendezvous Point (RP)

The root of a shared distribution tree down which all multicast traffic flows.

Designated Router (DR)

Responsible for sending join messages to the RP for group members and for sending register messages to the RP for sources.

PIM-SM uses shared trees by default and implements source-based trees for efficiency. The data threshold rate is used to toggle between trees. PIM-SM assumes that none of the hosts want multicast traffic unless they specifically ask for it. It creates a shared distribution tree centered on a defined RP from which source traffic is relayed to the receivers. Senders first send the multicast data to the RP, which, in turn, sends the data down the shared tree to the receivers. Shared trees centered on an RP do not necessarily provide the shortest or optimal path. In such cases, PIM-SM provides a means to switch to more efficient source-specific trees.

As soon as an active source sends a packet to its DR, the DR is responsible for registering the source with the RP and requesting that the RP build a tree back to that DR. The DR encapsulates the multicast data from the source in a Register message and unicasts that data to the RP.

PIM-SM uses the explicit join model whereby hosts wishing to receive multicast packets send PIM join messages to the RP via the DR. When a DR gets a membership indication for a new group it looks up the RP associated with the group and sends a join message to the RP.

The router can switch to a source's shortest path tree if the data rate of packets received from a specific source over the shared tree exceeds the threshold value during a specified time interval. The routers (RP and the last hop DR of the receiver) dynamically create a source-specific shortest path tree using join and prune messages and stop traffic from flowing down the shared RP tree (using Register Stop Messages when the RP has no downstream receivers for the group or that particular source) when the data rate reaches a threshold value.

Understanding and Troubleshooting the Spanning Tree Protocol

This appendix provides details about how the Spanning Tree Protocol and Algorithm work and describes how to troubleshoot them.

Spanning Tree Protocol (STP) operation

Spanning Tree Protocol (STP) is used in a bridged LAN environment to reduce the physical network to a stable logical topology with no data loops that still allows for the existence of redundant connections. The topology is calculated by the bridges that interconnect the individual LAN segments, and is recalculated when physical or parameter changes occur. Each bridge in the network has a unique bridge identifier, which is used to determine the root bridge of the spanning tree. Where more than one bridge on the same LAN segment offers connectivity to the root bridge, one bridge is selected as the designated bridge and one port on that bridge becomes the root port, providing access to the root bridge.

Two versions of STP are supported by the Ethernet switch module, both of which are defined in IEEE 802.1s. The first version is IEEE 802.1D compatibility mode, set as the factory default. The second version is Rapid Reconfiguration mode, originally defined in IEEE 802.1w. Rapid Reconfiguration uses the ability of the bridging device to recognize full-duplex links (point-to-point) and ports connected to end stations (edge ports) to offer faster transitions to the forwarding state. The **config spanningtree forceversion** command is used to switch from IEEE8021D operation to IEEE 802.1w operation. The two versions of the protocol can interoperate within the same LAN: it is not necessary for all bridges to run the same version. Where IEEE 802.1D is mentioned in this document, you should understand that the switch is actually operating in IEEE 802.1D compatibility mode according to the protocol specified in IEEE 802.1s.

Both versions of the Spanning Tree Algorithm (STA) create a single spanning tree for an entire network within which there is at most one route between any two end stations, and will automatically reconfigure the tree when necessary. The topology created by the algorithm is influenced by user-configurable parameters, but care should be taken when changing these parameters from the factory defaults.

The following table shows the user-configurable STP parameters for the bridge.

Table 4. STP parameters – bridge

Parameter	Description	Default value
Bridge identifier (Not user-configurable except by setting the priority as described in this table)	A combination of the Bridge Priority and the switch MAC address. The 16-bit priority parameter is concatenated with the 48-bit Ethernet MAC address.	32768 + MAC
Bridge Priority	A relative priority for each bridge. The lower the number the higher the priority and the greater the likelihood of the bridge being elected as the root bridge.	32768
Bridge hello time	The length of time between broadcasts of the hello message.	2 seconds
Bridge maxage time	The length of time before topology information or information from BPDUs is discarded because it has aged out.	20 seconds
Bridge forward delay time	The amount of time spent by a port in the discarding states waiting for a BPDU that might return the port to the discarding state if the bridge is in IEEE 802.1D compatibility mode or if operPointToPointMAC and operEdgePort are both False.	15 seconds

The following table shows the user-configurable STP parameters for the ports on the bridge.

Table 5. STP port parameters

Variable	Description	Default value
Port priority	The relative priority for each port. The lower the number the higher the priority and the greater the likelihood of the port being elected as the root port.	128
Port path cost	A value used by STP to evaluate paths.	auto (calculated based on the link speed)

Creating a stable topology

For STP to arrive at a stable network topology, the following information is used:

- A unique identifier for each bridge
- An identifier for each bridge port
- The path cost to the root bridge associated with each bridge port

STP communicates between bridges on the network using bridge protocol data units (BPDUs). There are two types of BPDUs:

- Configuration messages containing a spanning tree priority vector describing the transmitter view of the spanning tree topology
- Topology Change Notification (TCN) messages

Each BPDU has the following information:

- The unique identifier of the bridge that the transmitting bridge currently recognizes as the root bridge
- The path cost to the root from the transmitting port
- The port identifier of the transmitting port

The bridge sends BPDUs to communicate and construct the spanning-tree topology. All bridges connected to the LAN on which a packet is transmitted will receive the BPDU. BPDUs are not directly forwarded by the bridge, but the receiving bridge uses the information in the frame to calculate the topology and, if it changes, to initiate a BPDU transmission.

The communication between bridges through BPDUs causes the following results:

- The bridge with the lowest numerical identifier is elected as the root bridge.
- Each bridge calculates its root path cost by adding the path costs for each port receiving frames on the lowest cost path to the root bridge.
- The port on each bridge with the lowest root path cost for that bridge becomes the root port of that bridge (in the event of a tie the port with the lowest numerical port identifier is chosen).
- For each LAN the bridge with the lowest root path cost is selected as the designated bridge (in the event of a tie, the bridge with the lowest numerical bridge identifier is chosen) and the port connecting that bridge to the LAN becomes the designated port (in the event of a tie, the port with the lowest numerical port identifier is chosen).
- In the IEEE 802.1D standard, ports that are not selected as root or designated ports do not forward frames and are known as alternate ports.
- In the IEEE 802.1w standard, a port that offers an alternate path to the root bridge but is not selected as the root does not forward frames and is known as an alternate port. Ports that offer an alternate connection to the same LAN as a designated port do not forward frames and are known as backup ports.

If all bridges have STP enabled with default settings, the bridge with the lowest MAC address in the network will become the root bridge. By increasing the priority (lowering the priority number) of a bridge, STP can be forced to select that bridge as the root bridge.

When STP is enabled using the default parameters, the path between source and destination stations in a switched network might not be ideal. For example,

connecting higher-speed links to a port that has a higher number than the current root port can cause a root-port change. The goal is to make the fastest link the root port.

IEEE 802.1D STP port states

The BPDUs take some time to pass through a network. This propagation delay can result in topology changes in which a port that changed directly from a discarding state to a forwarding state could create temporary data loops. Ports must wait for new network topology information to propagate throughout the network before starting to forward packets. They must also wait for the packet lifetime to expire for BPDU packets that were forwarded based on the old topology. The forward delay timer is used to ensure that the network topology stabilizes after a topology change. In addition, STP specifies a series of states a port must go through to further ensure that a stable network topology is created after a topology change.

Each port on a bridge using STP exists in one of the following four states:

- | | |
|-------------------|---|
| Discarding | The port is blocked from forwarding or receiving packets. For additional information, see “Discarding state” on page 52. |
| Learning | The port is adding addresses to its forwarding database but not yet forwarding packets. For additional information, see “Learning state” on page 53. |
| Forwarding | The port is forwarding packets. For additional information, see “Forwarding state” on page 55. |
| Disabled | The port responds only to network management messages and must return to the discarding state first. For additional information, see “Disabled state” on page 56. Note that the STP port state of disabled applies only to the role of the port within the spanning tree, and should not be confused with the port administrative state of enabled or disabled. |

A port changes from one state to another as follows:

- From initialization (switch startup) to discarding
- From discarding to learning or to disabled
- From learning to forwarding or to disabled
- From forwarding to disabled or to discarding
- From disabled to discarding

When you enable STP, every port on every bridge in the network goes through the discarding state and then goes through the learning state at startup. If properly configured, each port stabilizes to the forwarding or discarding state.

No packets (except BPDUs and LACPDUs) are forwarded from, or received by, STP enabled ports until the forwarding state is enabled for that port.

IEEE 802.1w STP port states

The IEEE 802.1w protocol definition speeds up the reconfiguration of the spanning tree using two new mechanisms:

- Bridges exchange explicit acknowledgement frames
- Ports can be configured to transition directly to the forwarding state when the bridge is reinitialized – this is applicable for edge ports

The number of port states were reduced from five to three, specified in the original IEEE 802.1D standard:

Discarding The port is blocked from forwarding or receiving packets and does not add information to the forwarding database.

Learning The port is adding addresses to its forwarding database but not yet forwarding packets.

Forwarding The port is adding addresses to its forwarding database and is forwarding packets.

Table 6. Relationship between IEEE 802.1D and IEEE 802.1w port states

IEEE 802.1D port state	Admin. bridge port state	MAC operational	IEEE 802.1w port state	Active topology port role
Disabled	Disabled	False	Discarding	Excluded, disabled
Disabled	Enabled	False	Discarding	Excluded, disabled
Blocking	Enabled	True	Discarding	Excluded, alternate or backup
Listening	Enabled	True	Discarding	Included, root or designated
Learning	Enabled	True	Learning	Included, root or designated
Forwarding	Enabled	True	Forwarding	Included, root or designated

Setting user-changeable STP parameters

The next table shows the default spanning-tree configuration.

Table 7. Default STP parameters

Feature	Default value
Enable state	STP enabled for all ports
Port priority	128
Port cost	auto
Bridge priority	32768

The factory default settings are compatible with the majority of installations, and it is advisable to keep the default settings as set at the factory unless it is absolutely necessary to change them. The user-changeable parameters in the bridge are as follows:

Priority You can set a priority for the bridge from 0 to 65535. A value of 0 indicates the highest priority.

Hello Time The hello time can be from 1 to 10 seconds. This is the interval between two transmissions of BPDU packets sent by the root bridge to tell all other bridges that it is indeed the root bridge. If you set a hello time for your bridge, and it is not the root bridge, the set hello time will be used if and when your bridge becomes the root bridge.

Note: The hello time cannot be longer than the Max Age. Otherwise, a configuration error will occur.

Max. Age The Max. Age can be from 6 to 40 seconds. At the end of the Max. Age, if a BPDU has still not been received from the root bridge, your bridge will start sending its own BPDU to all other bridges for permission to become the root bridge. If your bridge has the lowest bridge identifier, it will become the root bridge.

Forward Delay The Forward Delay can be from 4 to 30 seconds. For IEEE 802.1D operation this is the time that any port on the bridge spends in the learning state while moving from the discarding state to the forwarding state. For IEEE 802.1w operation this is the time that a designated port on the bridge spends in the learning state while moving from the disabled state to the forwarding state when both `operPointToPointMAC` and `operEdgePort` are false.

Note: Observe the following formulas when setting the previously described parameters:

- $\text{Max. Age} \leq 2 \times (\text{Forward Delay} - 1 \text{ second})$
- $\text{Max. Age} \geq 2 \times (\text{Hello Time} + 1 \text{ second})$

Port Priority You can set a port priority from 0 to 240. The lower the number, the greater the probability that the port will be chosen as the root port.

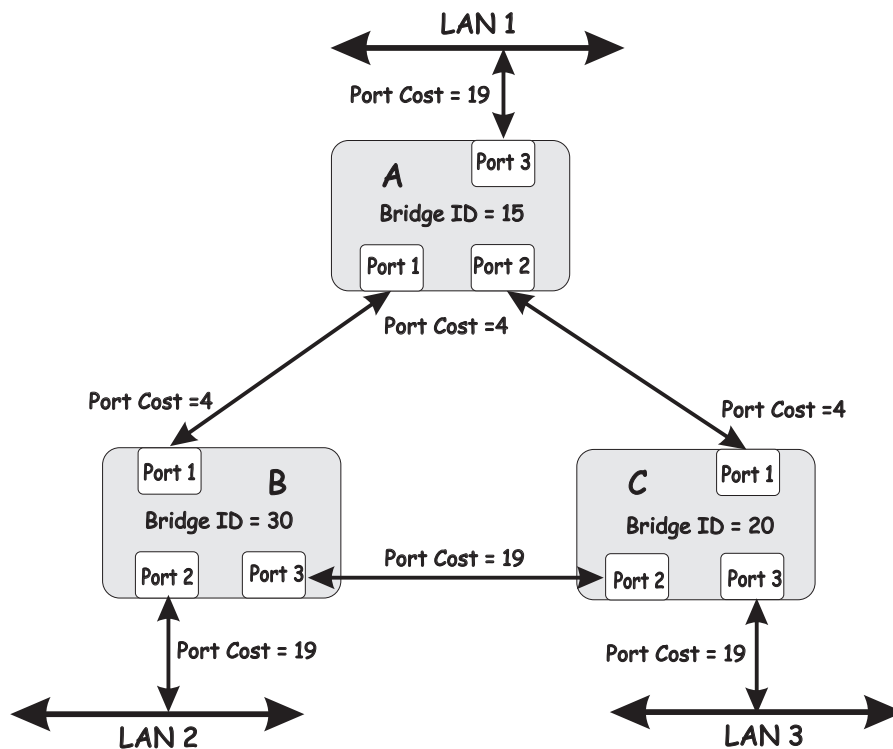
Port Path Cost You can set a port cost from 1 to 200000000, or specify auto. The lower the number, the greater the probability that the port will be chosen to forward packets. If you specify auto the switch will assign the port cost based on the link speed.

Illustration of STP

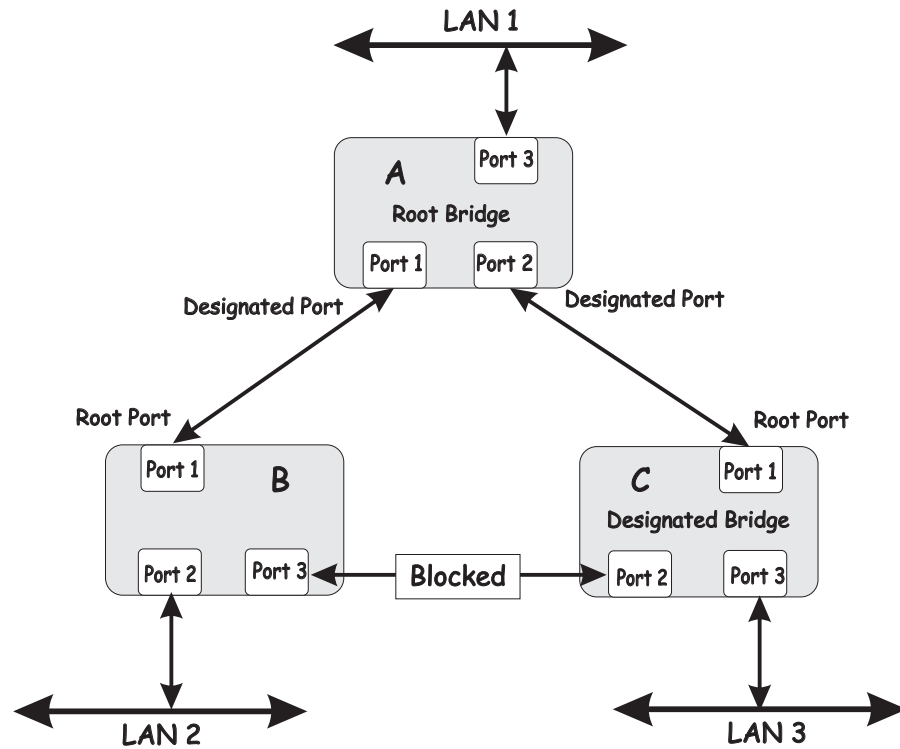
A simple illustration of three bridges (or three switches) connected in a loop is depicted in this section. In this example, you can anticipate some major network problems if the STP assistance is not applied. If bridge A broadcasts a packet to bridge B, bridge B will broadcast it to bridge C, and bridge C will broadcast it back to bridge A, and so on. The broadcast packet will be passed indefinitely in a loop, potentially causing a network failure.

STP can be applied as shown in the following illustration. In this example, STP breaks the loop by blocking the connection between bridges B and C. The decision to block a particular connection is based on the STP calculation of the most current bridge and port settings. If bridge A broadcasts a packet to bridge C, bridge C will drop the packet at port 2, and the broadcast will end there.

Setting up an STP using values other than the defaults can be complex. Therefore, keep the default factory settings and the STP will automatically assign root bridges, ports and block loop connections. However, influencing STP to choose a particular bridge as the root bridge using the Priority setting, or influencing STP to choose a particular port to block using the Port Priority and Port Cost settings, is relatively straightforward.



Note: In this example, only the default STP values are used.



The bridge with the lowest bridge ID (bridge A) was elected the root bridge, and the ports were selected to give a high port cost between bridges B and C.

Note also that the example network topology is intended to provide redundancy to protect the network against a link or port failure, not a switch failure or removal. For example, a failure of bridge A would isolate LAN 1 from connecting to LAN 2 or LAN 3.

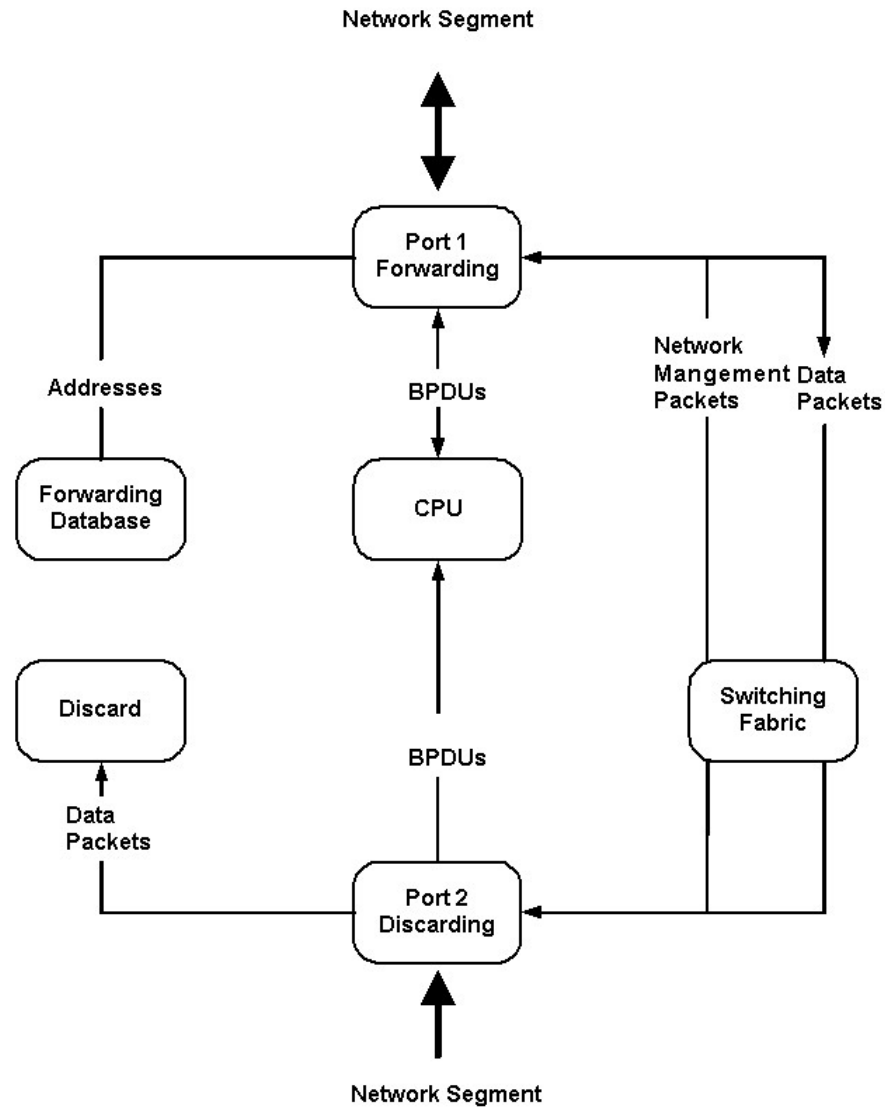
Discarding state

A port in the discarding state does not forward packets. When the switch is started, a BPDU is sent to each port in the bridge, putting these ports in the discarding state. A bridge initially assumes it is the root; it then begins the exchange of BPDUs with other bridges. This will determine which bridge in the network is the best choice for the root bridge. If there is only one bridge on the network, no BPDU exchange occurs, the forward delay timer expires, and the ports move to the learning state. All STP enabled ports enter the discarding state following the bridge startup.

A port in the discarding state does the following:

- Discards packets received from the network segment to which it is attached.
- Discards packets sent from another port on the bridge for forwarding.
- Does not add addresses to its forwarding database.
- Receives BPDUs and directs them to the central processing unit (CPU).
- Does not transmit BPDUs from the CPU.

The following illustration shows the actions that occur when a port is in the discarding state.



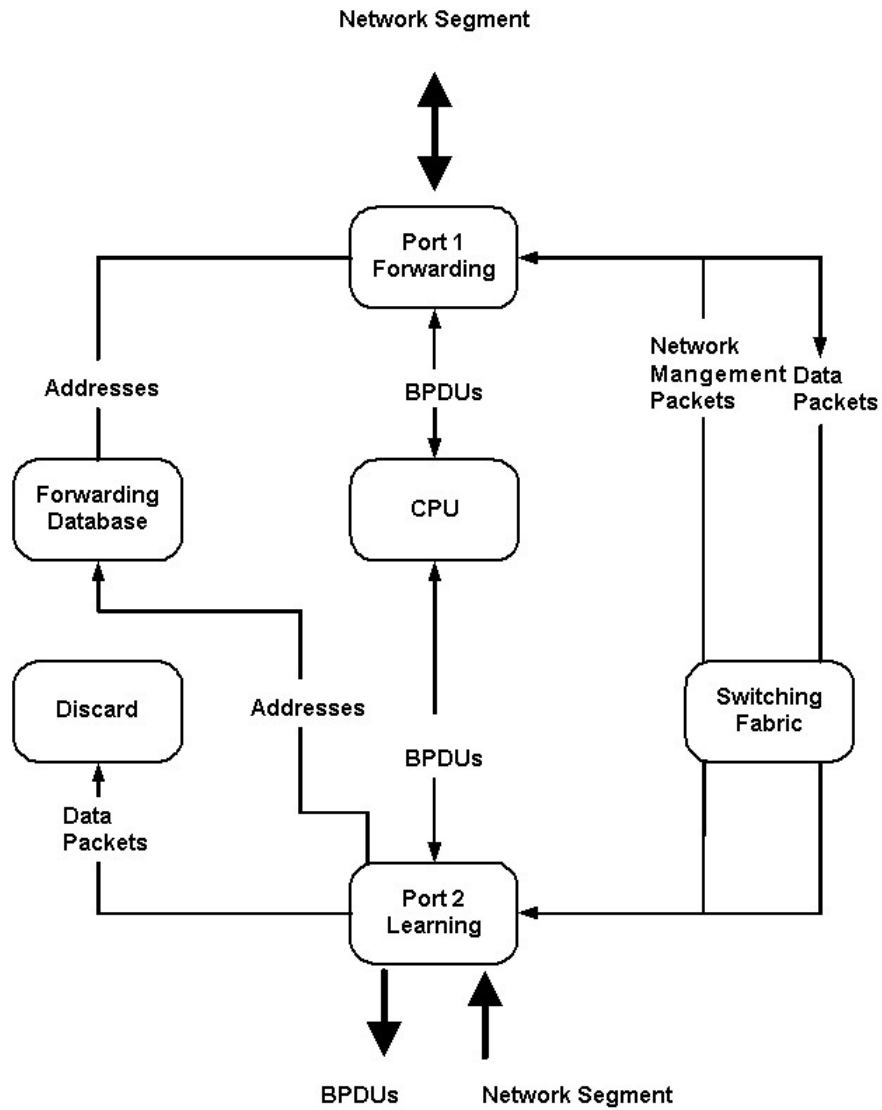
Learning state

A port in the learning state prepares to participate in frame forwarding. The port enters the learning state from the discarding state. A port will move from learning to forwarding when its forward delay timer expires.

A port in the learning state does the following:

- Discards frames received from the network segment to which it is attached.
- Discards packets sent from another port on the bridge for forwarding.
- Learns station location information from the source address of packets and adds this information to its forwarding database.
- Receives BPDUs for the CPU and transmits BPDUs from the CPU.

The following illustration shows the actions that occur when a port is in the learning state.



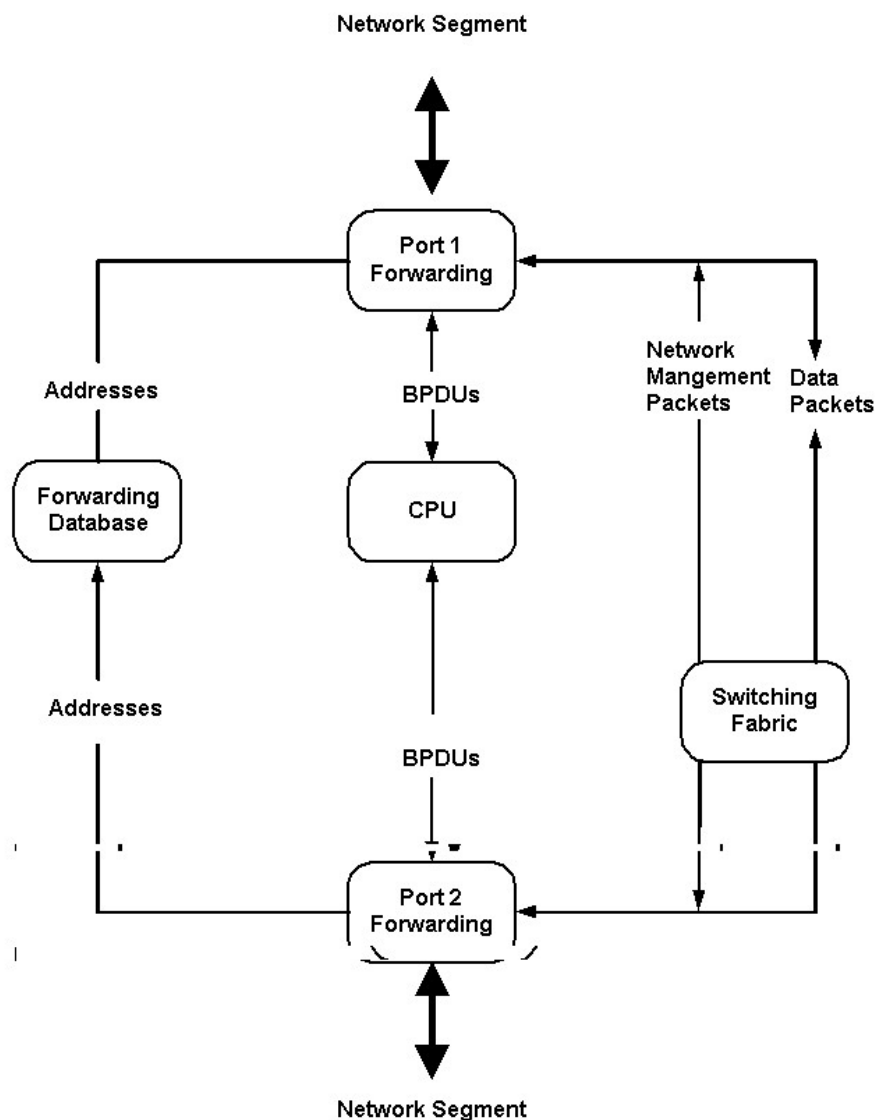
Forwarding state

A port in the forwarding state forwards packets. The port enters the forwarding state from the learning state when the forward delay timer expires.

A port in the forwarding state does the following:

- Forwards packets received from the network segment to which it is attached.
- Forwards packets sent from another port on the bridge for forwarding.
- Incorporates station location information into its address database.
- Receives BPDUs and directs them to the system CPU.
- Transmits BPDUs from the system CPU.
- Receives and responds to network management messages.

The following illustration shows the actions that occur when a port is in the forwarding state.



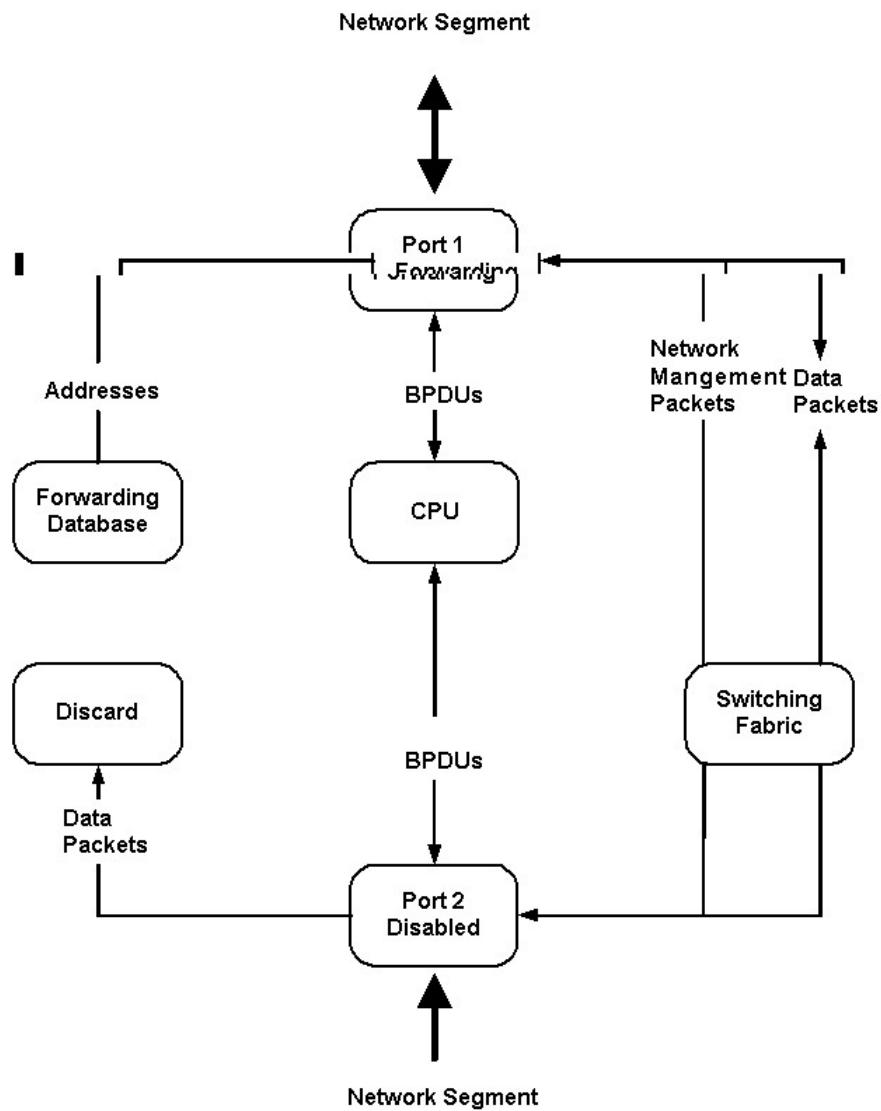
Disabled state

A port in the disabled state does not participate in frame forwarding or STP. A port in the disabled state is virtually non-operational. Note that this STP port state should not be confused with the port administrative state.

A disabled port does the following:

- Discards packets received from the network segment to which it is attached.
- Discards packets sent from another port on the bridge for forwarding.
- Does not add addresses to its forwarding database.
- Neither receives nor transmits BPDUs.

The following illustration shows the actions that occur when a port is in the disabled state.

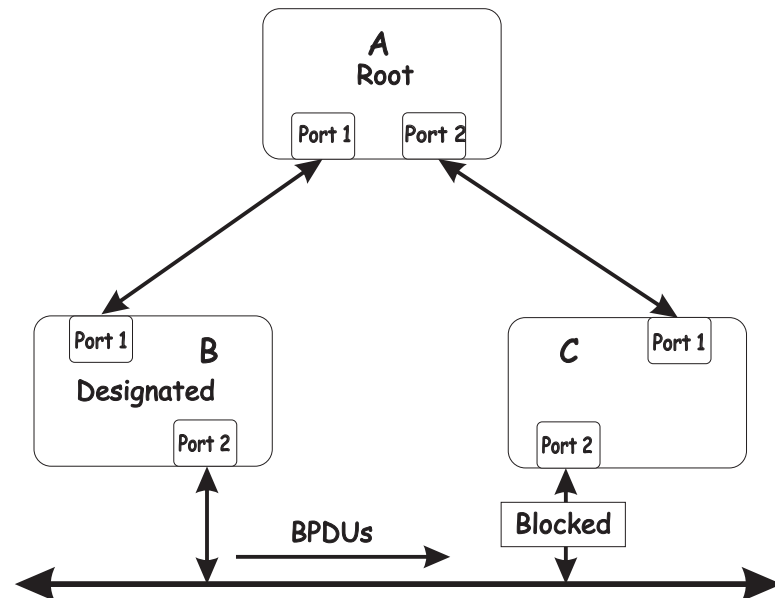


Troubleshooting STP

This section describes how to troubleshoot the STP.

Spanning Tree Protocol Failure

A failure in the Spanning Tree Algorithm generally results in a bridging loop. This is caused by a port that should be in the discarding state but is instead forwarding packets.



In this example, B has been elected as the designated bridge and port 2 on bridge C is in the discarding state. The election of B as the designated bridge is determined by the exchange of BPDUs between bridges B and C. Bridge B had a better spanning tree priority vector than bridge C. Bridge B continues sending BPDUs that advertise its superiority over the other bridges on this LAN. If bridge C fails to receive these BPDUs for longer than the Max. Age time (default of 20 seconds), it could start to change its port 2 from the discarding state to the forwarding state.

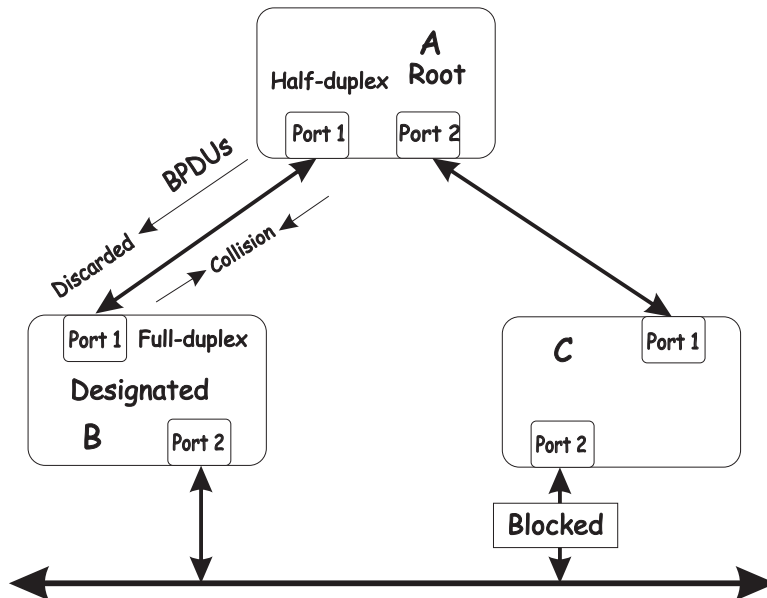
Note: To remain in the discarding state, a port must continue to receive BPDUs that advertise superior paths.

There are several circumstances in which the algorithm can fail, mostly related to the loss of a large number of BPDUs. These situations will cause a port in the discarding state to change to the forwarding state.

Full/half duplex mismatch

A mismatch in the duplex state of two ports is a very common configuration error for a point-to-point link. If one port is configured as full duplex and the other port is left in auto-negotiation mode, the second port will end up in half-duplex because ports

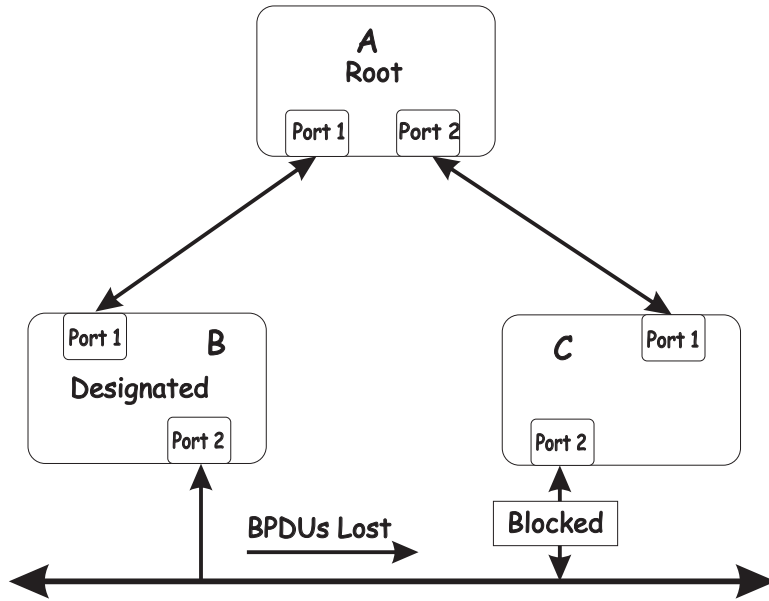
explicitly configured as half- or full-duplex do not negotiate.



In the preceding example, port 1 on bridge B is configured as a full-duplex port and port 1 on bridge A is either configured as a half-duplex port or is left in auto-negotiation mode. Because port 1 on bridge B is configured as a full-duplex port, it does not test for carrier sense when accessing the link. Bridge B will then start sending packets even if bridge A is using the link. A will then detect collisions and begin to run the flow control algorithm. If there is enough traffic between bridges B and A, all packets (including BPDUs) will be dropped. If the BPDUs sent from bridge A to bridge B are dropped for longer than the Max. Age, bridge B will lose its connection to the root (bridge A) and will unblock its connection to bridge C. This will create a data loop.

Unidirectional link

Unidirectional links can be caused by an undetected failure in one side of a fiber cable or by a problem with a port transceiver. Any failure that enables a link to remain up while providing one-way communication is very likely to cause a Spanning Tree Protocol failure.



In this example, port 2 on bridge B can receive but not transmit packets. Port 2 on bridge C should be in the discarding state, but since it can no longer receive BPDUs from port 2 on bridge B, it will change to the forwarding state. If the failure exists at boot time, STP will not converge on a stable topology and restarting the bridges will have no effect.

Note: In the previous example, restarting the bridges will provide a temporary resolution.

This type of failure is difficult to detect because the Link-state LEDs for Ethernet links rely on the transmit side of the cable to detect a link. If a unidirectional failure on a link is suspected, it is usually necessary to go to the console or other management software and look at the packets received and transmitted for the port. For example, a unidirectional port will have many packets transmitted but none received, or vice versa.

Packet corruption

Packet corruption can lead to the same type of failure. If a link is experiencing a high rate of physical errors, a large number of consecutive BPDUs can be dropped and a port in the discarding state would change to the forwarding state. The discarding port would have to have the BPDUs dropped for 50 seconds (at the default settings) and a single BPDU would reset the timer. If the Max. Age is set too low, this time is reduced.

Resource errors

The switch performs its switching and routing functions primarily in hardware, using specialized application-specific integrated circuits (ASICs). STP is implemented in software and is thus reliant upon the speed of the CPU and other factors to converge. If the CPU is over utilized, it is possible that BPDUs might not be sent in a timely fashion. STP is generally not very CPU intensive and is given priority over other processes, so this type of error is rare.

It can be seen that very low values for the Max. Age and the Forward Delay can result in an unstable spanning tree. The loss of BPDUs can lead to data loops. The diameter of the network can also cause problems. The default values for STP give a maximum network diameter of about seven. This means that two bridges in the network cannot be more than seven hops apart. Part of this diameter restriction is the BPDU age field. As BPDUs are propagated from the root bridge to the leaves of the spanning tree, each bridge increments the age field. When this field is beyond the maximum age, the packet is discarded. For large diameter networks, STP convergence can be very slow.

Identifying a data loop

Broadcast storms have a very similar effect on the network-to-data loops, but broadcast storm controls in modern bridges have been (along with subnetting and other network practices) very effective in controlling broadcast storms. The best way to determine if a data loop exists is to capture traffic on a saturated link and check whether similar packets are seen multiple times.

Generally, if all the users of a domain are unable to connect to the network at the same time, a data loop is the cause. In this case, the port utilization data will have unusually high values.

The priority for most cases is to restore connectivity as soon as possible. The simplest remedy is to manually disable all of the ports that provide redundant links. Disabling the ports one at a time and then checking for the restoration of user

connectivity will identify the link that is causing the problem, if sufficient time is available. Connectivity will be restored immediately after disabling a data loop.

Avoiding network problems

To help your network operate more efficiently, you can avoid or minimize network problems, as described in this section.

- Know where the root is located.

Although the STP can elect a root bridge, a well-designed network has an identifiable root for each VLAN. Careful setup of the STP parameters results in the selection of this best bridge as the root for each VLAN. Redundant links can then be built into the network. STP is well-suited to maintaining connectivity in the event of a device failure or removal, but is poorly suited to designing networks.

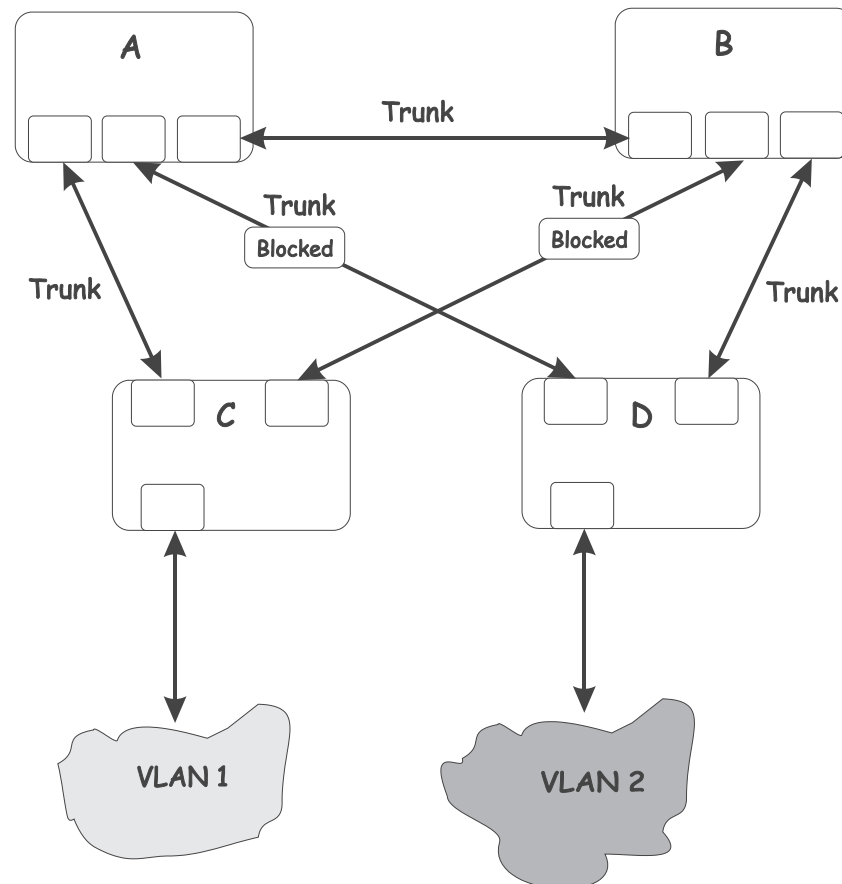
- Know which links are redundant.

Organize the redundant links and tune the port cost parameters of STP to force those ports into the discarding state.


For each VLAN, know which ports should be discarding in a stable network. A network illustration that shows each physical loop in the network and which ports break which loops is extremely helpful.

- Minimize the number of ports in the discarding state.

A single discarding port changing to the forwarding state at an inappropriate time can cause a large part of a network to fail. Limiting the number of blocked ports helps to limit the risk of an inappropriate change.



This is a common network design. Through trunks, bridges C and D have redundant links to backbone bridges A and B. Trunks, by default, carry all the VLAN traffic from VLAN 1 and VLAN 2. Therefore, bridge C is not only receiving traffic for VLAN 1, but also unnecessary broadcast and multicast traffic for VLAN 2. Bridge C is also discarding one port for VLAN 2. Thus, there are three redundant paths between bridges A and B, and two blocked ports per VLAN. This increases the chance of a data loop.



In this example, the VLAN definitions are extended to bridges A and B. This gives only a single blocked port per VLAN and enables the removal of all redundant links by removing bridge A or B from the network.

Chapter 4. Web-based network management

This chapter describes how to use the Web-based network management module to access and configure the internal switching software.

Important: Before you configure the Ethernet switch module, be sure that the management modules in your BladeCenter T unit are correctly configured. In addition, to access and manage your switch module from an external environment, you might need to enable certain features, such as the external ports and external management over all ports. See the applicable *Installation* and *User's Guide* publications on the IBM *BladeCenter T Documentation* CD for more information.

Introduction

The Ethernet switch module offers an embedded Hypertext Markup Language (HTML), Web-based interface that enables you to manage the switch through a standard browser, such as Opera, Netscape Navigator/Communicator or Microsoft® Internet Explorer. The Web browser acts as an access tool and can communicate directly with the switch using the HTTP protocol.

Note: This Web-based management module does not accept Chinese language input (or other double-byte character-set languages).

The Web-based management module and the Telnet program are different ways to access and configure the same internal switching software. Thus, all the settings that you encounter in Web-based management are the same as those found in the Telnet program. If your system application requires that you use the Telnet program, see the *CLI Reference Guide* on the IBM *BladeCenter T Documentation* CD for additional information.

This chapter explains the menus and parameters used by the web management interface. Note that your browser window might not exactly match the window illustrations in this guide.

Remotely managing the switch module

The Ethernet switch module supports two remote-access modes for management through Ethernet connections. You can select the mode that is best suited for your platform's environment. The switch module has an internal Ethernet path to the management module and the four external Ethernet ports on the switch module.

- The default mode uses the internal path to the management module only. In this mode, the remote-access link to the management console must be attached to the 100 Mbps Ethernet port on the management module. With this mode, the IP addresses and Simple Network Management Protocol (SNMP) parameters of the switch modules can be assigned manually through the management and configuration program. This mode enables the system administrator to provide a secure LAN for management of the platform's subsystems separately from the data network.

Important: With this mode, the Ethernet switch module does not respond to remote-management commands through the four external Ethernet ports on the switch module.

See the applicable *Installation* and *User's Guide* on the IBM *BladeCenter T Documentation* CD for additional instructions for configuring the switch module for this mode of operation.

- The system administrator can select to enable remote management of the Ethernet switch module through the four external Ethernet ports on the switch module, instead of, or in addition to, access through the management module. This mode can be enabled only through the management module configuration interface. When this mode is enabled, the external Ethernet ports will support both management traffic and BladeCenter T unit application data traffic. Also, the Ethernet switch module can transmit DHCP request frames through the external Ethernet ports.

This mode enables the switch module's IP addresses to reside on a different subnet than the management modules. This is useful when the switch modules are to be managed and controlled as part of the overall network infrastructure, while maintaining secure management of other BladeCenter T unit subsystems through the management module. However, management access to the Ethernet switch module link will be lost if its IP address is not on the same subnet as the management module. This chapter contains additional instructions for configuring the Ethernet switch module for this mode of operation.

The two previously described modes are only applicable to the Ethernet switch module. The management module can be remotely accessed only through the 100 Mbps Ethernet port on the management module.

Getting started

The first step in getting started using Web-based management for your switch is to install a web browser on the endstation you will be using. The web browser will allow you to connect to the switch and read the management screens. Some popular browsers are Opera, Netscape Navigator/Communicator and Microsoft Internet Explorer. Follow the installation instructions for the browser.

The switch module will acquire its IP address from a DHCP server.


You are now ready to begin managing your switch by simply running the browser installed on your computer and pointing it to the IP address defined for the device. The URL in the address bar should have the following format and contain information similar to: `http://123.123.123.123`, where the numbers `123.123.123.123` represent the IP address of the switch.

Depending on which browser you are using, a Login hyperlink displays:

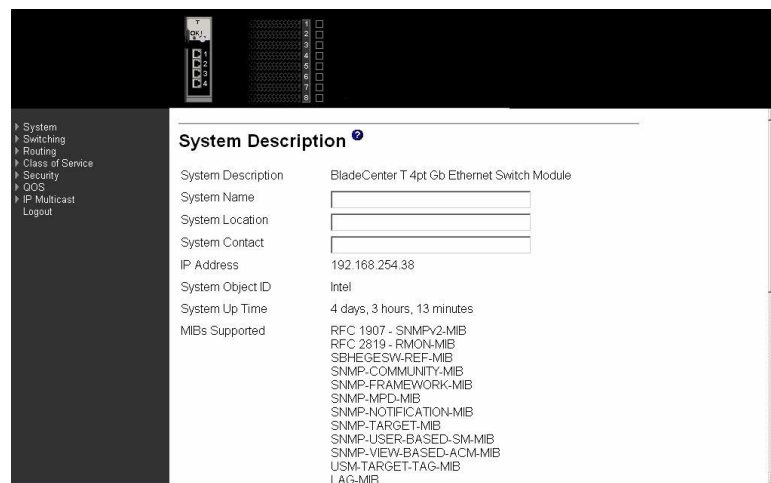
[Login](#)

System Name
System Location
System Contact

Click on Login, and a dialog box similar to the following will open:



Enter “**USERID**” in the User name field and enter “**PASSW0RD**” (with a zero in place of the O) in the **Password** field. Click the **OK** button. This opens the main page in the management module.



Note: The User name and Password fields ARE case sensitive. To increase system security, set a password after you log onto the switch module for the first time and be sure to store the new password in a safe location.

If java mode is enabled for the switch (the default is enabled) the link-status panel at the top of the screen shows a real-time display of the switch module, as shown below. The link-status panel reflects the current link status of the four external ports and the eight internal ports. You can change the java mode on the Network Connectivity Configuration menu (see “Network connectivity” on page 69).




The panel on the left side of the screen displays the main menu. The main menu contains:

- System
- Switching

- Routing
- Class of service
- Security
- QOS
- IP multicast
- Logout

All of these main menu options (except Logout) have sub-menus, some of which have further sub-menus, as shown below. All of the Web-based switch module management features are accessed from these sub-menus and are described in the remainder of this chapter.



When you first log on to the switch, you will see the System Description details in the center of the screen. For more details on the information displayed see 69.

System Description [?]	
System Description	BladeCenter T 4pt Gb Ethernet Switch Module
System Name	<input type="text"/>
System Location	<input type="text"/>
System Contact	<input type="text"/>
IP Address	192.168.254.38
System Object ID	Intel
System Up Time	4 days, 1 hours, 33 minutes
MIBs Supported	RFC 1907 - SNMPV2-MIB RFC 2819 - RMON-MIB SBHEGESW-REF-MIB SNMP-COMMUNITY-MIB SNMP-FRAMEWORK-MIB SNMP-MPD-MIB SNMP-NOTIFICATION-MIB SNMP-TARGET-MIB SNMP-USER-BASED-SM-MIB SNMP-VIEW-BASED-ACM-MIB USM-TARGET-TAG-MIB LAG-MIB

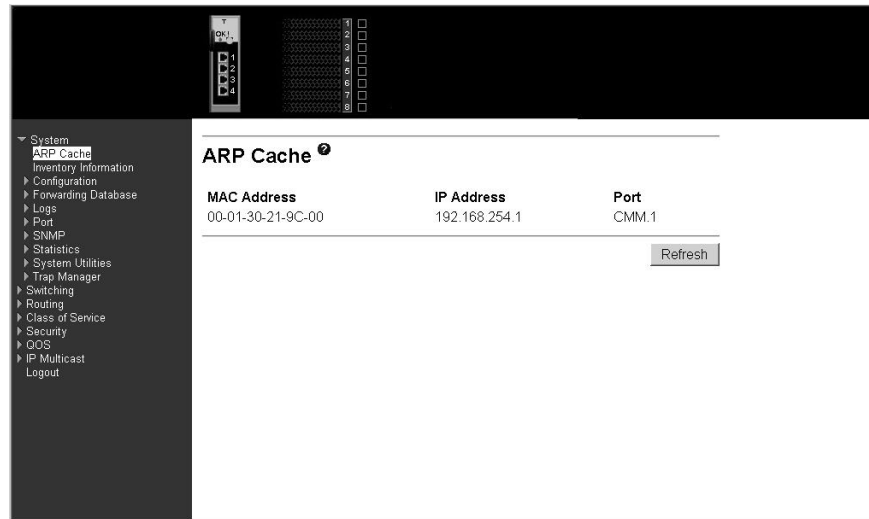
System

The switch's system menu provides access to the following panels and menus:

- Address Resolution Protocol (ARP) cache
- Inventory information
- Configuration
- Forwarding database
- Logs
- Port
- SNMP
- Statistics
- System utilities
- Trap manager

ARP cache

This panel displays the connectivity between the switch and other devices. The ARP cache identifies the Media Access Control (MAC) addresses of the IP stations communicating with the switch.



MAC Address A unicast MAC address of a device on a subnet attached to one of the switch's interfaces for which the switch has forwarding and/or filtering information. The format is six two-digit hexadecimal numbers separated by hyphens; for example, 01-23-45-67-89-AB.

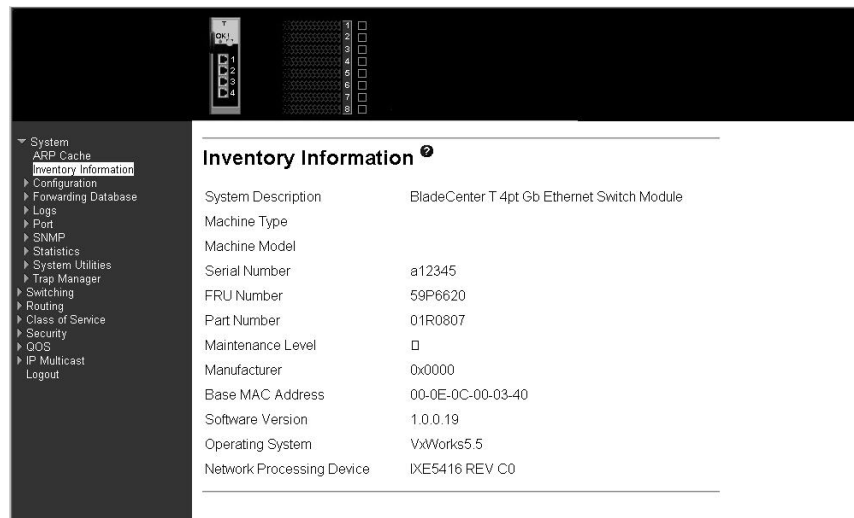
IP Address The IP address associated with the MAC address.

Port The identification of the port being used for the connection.

Click the Refresh button to retrieve and display the database again, starting with the first entry in the table.

Inventory information

This panel displays inventory information for the switch.



System Description

The product name of this switch.

Machine Type The machine type of this switch.

Machine Model

The model within the machine type.

Serial Number

The unique box serial number for this switch.

FRU Number The field-replaceable unit number.

Part Number The manufacturing part number.

Maintenance Level

The identification of the hardware change level.

Manufacturer The code that identifies the manufacturer, displayed as two two-digit hexadecimal numbers.

Base MAC Address

The burned-in, universally administered, MAC address of this switch, displayed as six two-digit hexadecimal numbers separated by hyphens.

Software Version

The release.version.maintenance number of the code currently running on the switch.

Operating System

The operating system currently running on the switch.

Network Processing Device

The network processor hardware.

Additional Packages

The list of optional software packages installed on the switch, if any. For example, Quality of Service or Multicast.

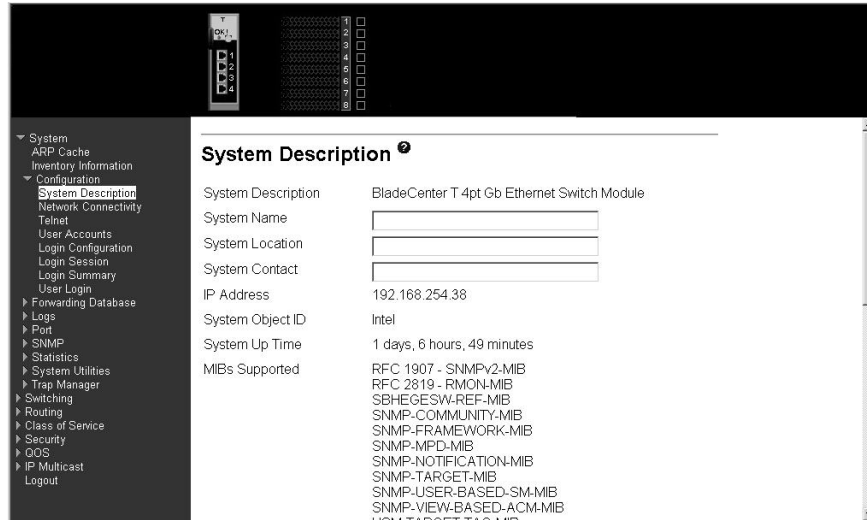
Configuration

The Configuration menu gives you access to panels used for switch module management. The options are:

- System description
- Network connectivity
- Telnet
- User accounts
- Login configuration
- Login session
- Login summary
- User login

System description

This panel displays and allows configuration of switch system information.



System Description

The product name of this switch.

System Name

The name used to identify this switch. The range for name is from 1 to 31 alphanumeric characters.

System Location

The physical location of this switch. Can be up to 31 alphanumeric characters. The factory default is blank.

System Contact

The person or organization responsible for this switch. Can be up to 31 alphanumeric characters. The factory default is blank.

IP Address The IP address of the interface. The factory default value is 0.0.0.0.

System Object ID

The base object ID for the switch's enterprise MIB.

System Up Time

The time in days, hours and minutes since the last reboot.

MIBs Supported

The list of MIBs supported by the management agent running on this switch.

Click the Apply button to update the switch with the values on the screen. If you want the switch to retain the new values across a power cycle you must perform a save.

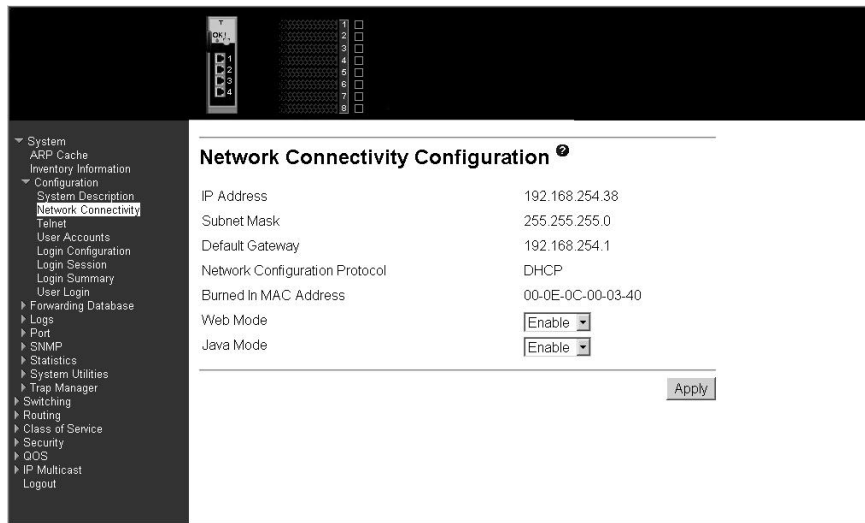
Network connectivity

This panel displays network configuration settings necessary for in-band connectivity. The network interface is the logical interface used for in-band connectivity with the switch via any of the switch's front panel ports. The configuration parameters associated with the switch's network interface do not affect the configuration of the front panel ports through which traffic is switched or routed.

To access the switch over a network it must first configure its IP information (IP address, subnet mask and default gateway) via DHCP.

When you have established in-band connectivity, you can change the IP information using any of the following:

- Terminal interface via telnet or SSH connections
- SNMP-based management
- Web-based management



IP Address The IP address of the interface. The factory default value is 10.90.90.9x.

Subnet Mask The IP subnet mask for this interface. The factory default value is 255.255.255.0.

Default Gateway The default IP gateway address for this interface. The factory default value is 0.0.0.0.

Network Configuration Protocol Indicates what network protocol was used on the last, or current power-up cycle, if any. The factory default method is none.

Burned In MAC Address The burned-in MAC address used for in-band connectivity if you choose not to configure a locally administered address.

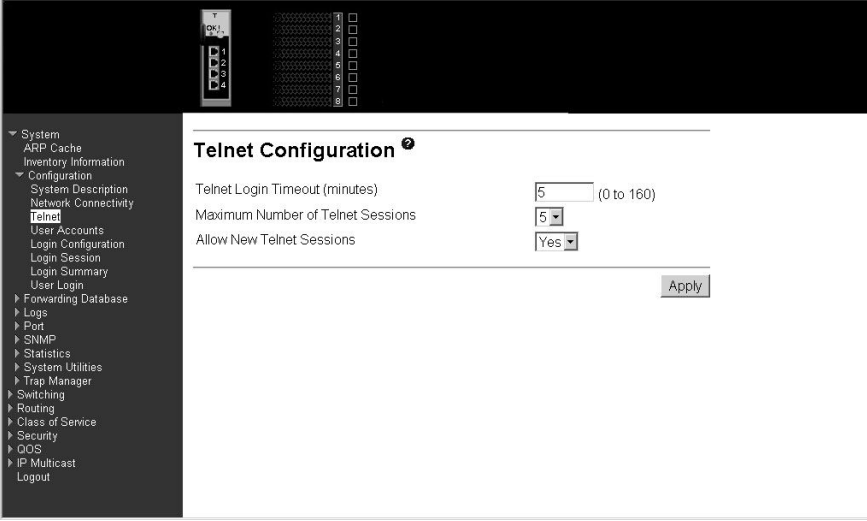
Web Mode Specify whether the switch can be accessed from a web browser through TCP port 80. If you choose to Enable web mode you will be able to manage the switch from a web browser. The factory default is Enabled.

Java Mode Enable or Disable the java applet that displays a picture of the switch at the top right of the screen. If you run the applet you will be able to click on the picture of the switch to select configuration screens instead of using the navigation tree at the left side of the screen. The factory default is Enabled.

Click the Apply button to update the switch with new values. If you want the switch to retain the new values across a power cycle you must perform a save.

Telnet

Use this panel to configure Telnet settings.



The screenshot shows a web-based configuration interface for a network device. On the left is a dark sidebar with a tree view of configuration categories. The main area is titled "Telnet Configuration" and contains three settings:

- Telnet Login Timeout (minutes):** A text input field containing the value "5", with "(0 to 160)" to its right.
- Maximum Number of Telnet Sessions:** A pull-down menu showing "5".
- Allow New Telnet Sessions:** A pull-down menu showing "Yes".

An "Apply" button is located at the bottom right of the configuration area.

Telnet Login Timeout (minutes)

Specify how many minutes of inactivity should occur on a Telnet or SSH session before the switch logs off. A zero means there will be no timeout. You can enter any number from 0 to 160. The factory default is 5.

Maximum Number of Telnet Sessions

Use the pull-down menu to select how many simultaneous Telnet and SSH sessions will be allowed. The maximum is 5, with 5 being the factory default.

Allow New Telnet Sessions

Indicates whether new Telnet sessions are allowed. If you set this to no, new Telnet sessions will not be allowed. The factory default is yes.

Click the Apply button to update the switch with new values. If you want the switch to retain the new values across a power cycle you must perform a save.

User accounts

Use this panel to reconfigure an existing user account or to create a new one. This panel is only available for the user with Read/Write privileges, herein referred to as

admin.

User Use this pull-down menu to select one of the existing accounts, or select Create to add a new one, provided the maximum of five Read-only accounts has not been reached.

User Name The name the user will use to login using the serial port, Telnet or Web. It can be up to eight alphanumeric characters and is not case-sensitive. Six user names can be defined, including the Read-only user “GUEST” which cannot be changed. The admin user will enter USERID (all caps, case sensitive) in this field.

Password Enter the optional new or changed password for the account. It will not display as it is typed, only asterisks (*) will show. The password is up to eight alphanumeric characters and is case-sensitive. Default for GUEST is blank and for the admin is “PASSWORD” (please note the use of zero instead of “O”).

Confirm Password Enter the password again to confirm that you entered it correctly. The information entered in this field will not display, but will show as asterisks (*).

Access Mode Displays whether the user is able to change parameters on the switch (Read/Write) or is only able to view them (Read-only). As a factory default, admin has Read/Write access and GUEST has Read-only access. There can be one Read/Write user and up to five Read-only users.

SNMP v3 Access Mode Indicates the SNMPv3 access privileges for the user account. If the value is set to Read/Write, the SNMPv3 user will be able to set and retrieve parameters on the system. If the value is set to Read-only, the SNMPv3 user will only be able to retrieve parameter information. The SNMPv3 access mode can be different from the CLI and Web access mode.

Authentication Protocol The protocol (if any) used to authenticate the user. This field specifies the protocol to be used to authenticate a user account. The valid authentication protocols are None, MD5 or SHA. If MD5

or SHA are specified, the user login password will be used as the SNMPv3 authentication password.

Encryption Protocol

Specify the SNMPv3 Encryption Protocol settings for the selected user account. The valid encryption protocols are None or DES. If you select the DES protocol you must enter a key in the Encryption Key field. The key can be up to 16 characters long. If None is specified for the protocol, the Encryption Key is ignored.

Encryption Key

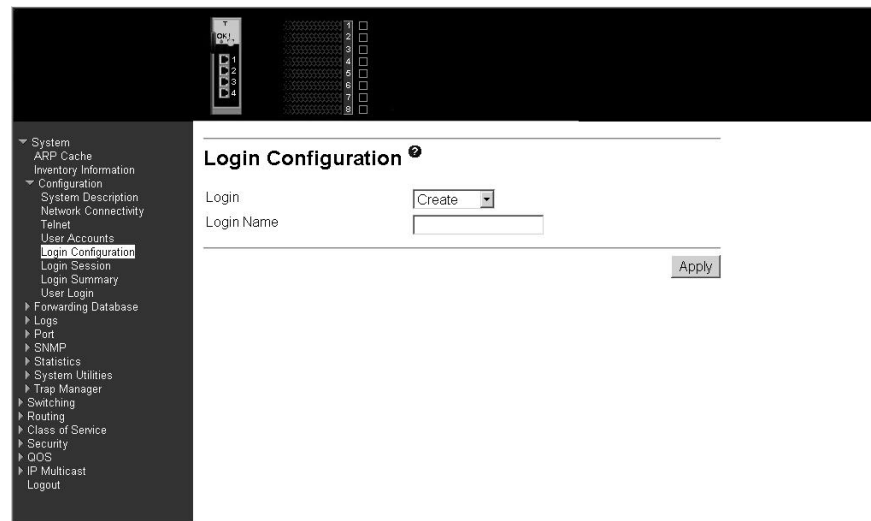
If you selected DES in the Encryption Protocol field, enter the SNMPv3 Encryption Key here. Otherwise this field is ignored. Valid keys are 0 to 15 characters long. The Apply checkbox must be checked in order to change the Encryption Protocol and Encryption Key.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Click the Delete button to delete the displayed user; this button is only visible when you have selected a user account with Read-only access. You cannot delete the Read/Write user.

Login configuration

Use this panel to configure login lists. A login list specifies the authentication method(s) you want used to validate switch or port access for the users associated with the list. The pre-configured users, admin and GUEST, are assigned to a pre-configured list named defaultList, which you cannot delete. All newly created users are also assigned to the defaultList until you specifically assign them to a different list.



Login

Select the authentication login list you want to configure. Select Create to define a new login list. When you create a new login list, Local is set as the initial authentication method.

Login Name

If you are creating a new login list, enter the name you want to assign. It can be up to 15 alphanumeric characters long and is not

case sensitive. The pull-down menus you use to specify authentication methods only appear after you create a list by entering a name.

Method 1 Use the pull-down menu to select the method that should appear first in the selected authentication login list. If you select a method that does not time out as the first method, such as local, no other method will be tried, even if you have specified more than one method. Note that this parameter will not appear when you first create a new login list. The options are:

Local The user's locally stored ID and password will be used for authentication.

Radius

The user's ID and password will be authenticated using the RADIUS server instead of locally.

Reject The user is never authenticated.

Undefined

The authentication method is unspecified (this might not be assigned as the first method).

Method 2 Use the pull-down menu to select the method, if any, that should appear second in the selected authentication login list. This is the method that will be used if the first method times out. If you select a method that does not time out as the second method, the third method will not be tried. Note that this parameter will not appear when you first create a new login list.

Method 3 Use the pull-down menu to select the method, if any, that should appear third in the selected authentication login list. Note that this parameter will not appear when you first create a new login list.

Click the Apply button to cause the changes made on this screen to take effect on the switch. If you want the switch to retain the new values across a power cycle, you must perform a save.

Click the Delete button to remove the selected authentication login list from the configuration. The delete will fail if the selected login list is assigned to any user (including the default user) for system login or IEEE 802.1X port access control. You can use this button only if you have Read/Write access.

Login session

This panel displays the details for all user login sessions.

ID	User Name	Connection From	Idle Time	Session Time
00	USERID	EIA-232	30:52:45	30:53:30

ID The ID of this row.

User Name The user name of user made the session.

Connection From
The user is connected from which machine.

Idle Time The idle session time.

Session Time The total session time.

Click the Refresh button to update the information on the page.

Login summary

This panel displays a list of all users set up for each authentication login list.

Login	Method List	Login Users	802.1x Port Security Users
defaultList	local	USERID GUEST default	USERID GUEST default

Login Identifies the authentication login list summarized in this row.

Method List The ordered list of methods configured for this login list.

Login Users The users you assigned to this login list on the User Login Configuration screen. This list is used to authenticate the users for system login access.

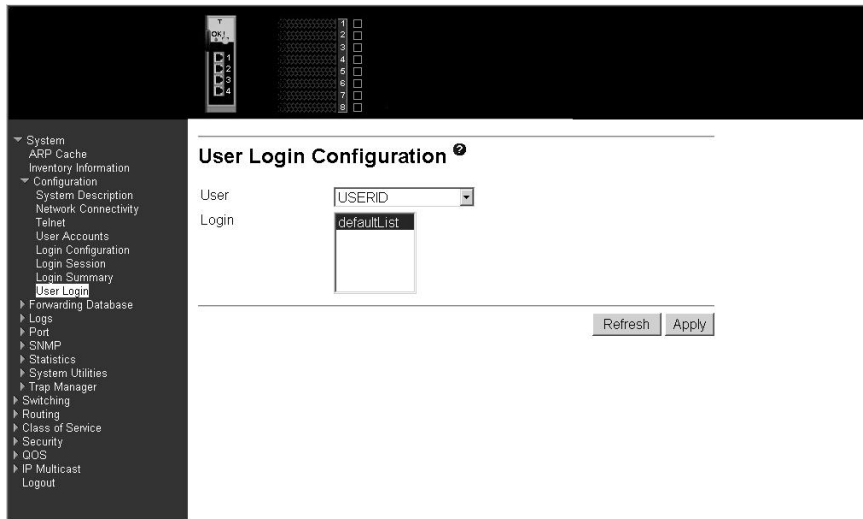
802.1X Port Security Users

The users you assigned to this login list on the Port Access Control User Login Configuration screen. This list is used to authenticate the users for port access, using the IEEE 802.1X protocol.

Click the Refresh button to update the information on the page.

User login

Use this panel to assign a user to an authentication login list.



Each configured user is assigned to a login list that specifies how the user should be authenticated when attempting to access the switch or a port on the switch. After creating a new user account on the User Account screen, you should assign that user to a login list for the switch using this screen and, if necessary, to a login list for the ports using the Port Access Control User Login Configuration screen. If you need to create a new login list for the user, you would do so on the Login Configuration screen.

The pre-configured users are assigned to a pre-configured list named defaultList, which you cannot delete. All newly created users are also assigned to the defaultList until you specifically assign them to a different list.

A user that does not have an account configured on the switch is termed the default or non-configured user. If you assign the non-configured user to a login list that specifies authentication via the RADIUS server, you will not need to create an account for all users on each switch. However, by default the non-configured user is assigned to defaultList, which by default uses local authentication.

User Select the user you want to assign to a login list. Note that you must always associate the admin user with the default list. This forces the admin user to always be authenticated locally to prevent full lockout from switch configuration. If you assign a user to a login list that requires remote authentication, the user's access to the switch from all CLI, web, and Telnet sessions will be blocked until the authentication is complete.

Login Select the authentication login list you want to assign to the user for system login.

Click the Refresh button to update the information on the page.

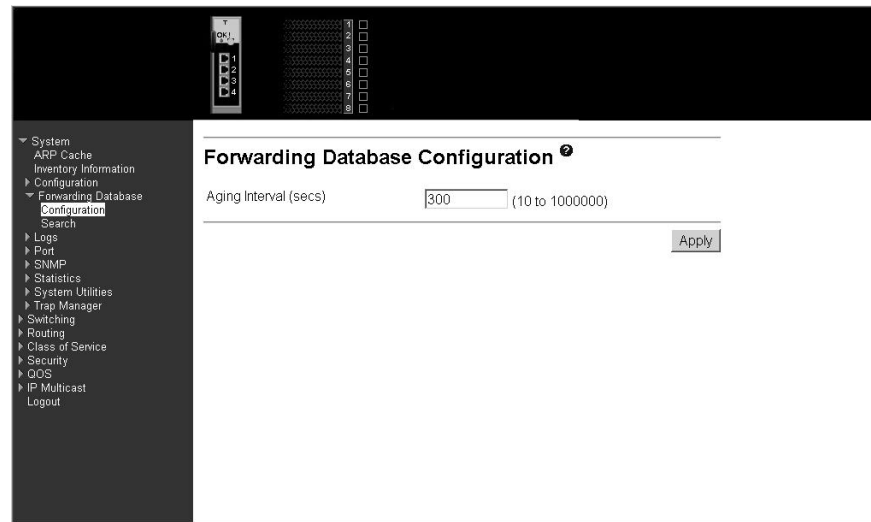
Click the Apply button to cause the changes made on this screen to take effect on the switch click. If you want the switch to retain the new values across a power cycle, you must perform a save.

Forwarding database

The first option on this menu is the Configuration panel, which allows you to configure the forwarding database aging interval. The second option is the Search panel, which displays the forwarding database entries specified by a MAC address or filter you enter.

Configuration

Use this panel to configure the forwarding database aging interval.



Aging Interval (secs)

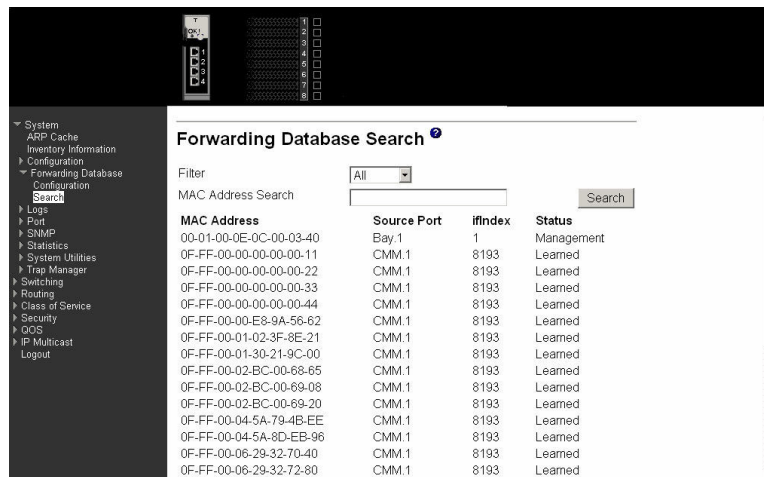
The forwarding database contains static entries, which are never aged out, and dynamically learned entries, which are removed if they are not updated within a specified time. You specify that time by entering a value for the Aging Interval. Enter any number of seconds between 10 and 1000000. IEEE 802.1D recommends a default of 300 seconds, which is the factory default.

Click the Apply button to cause the changes made on this screen to take effect on the switch. If you want the switch to retain the new values across a power cycle, you must perform a save.

Search

This panel displays the forwarding database entries. You can specify a filter to determine which addresses are displayed or a MAC address to display the table entry for the requested MAC address (and all entries following the requested MAC

address).



Filter Specify the entries you want displayed from the pull-down menu. When a choice is made the list is automatically refreshed with the selected filter. Filter choices are:

Learned

Only MAC addresses that have been learned will be displayed.

All The entire table will be displayed.

MAC Address Search

You can also search for an individual MAC address. Enter the two byte hexadecimal Virtual Local Area Network (VLAN) ID followed by the six byte hexadecimal MAC address in two-digit groups separated by hyphens; for example, 01-23-00-67-89-AB-CD-EF where 01-23 is the VLAN ID and 45-67-89-AB-CD-EF is the MAC address. Then click the Search button. If the address exists, that entry will be displayed as the first entry followed by the remaining (greater) MAC addresses. An exact match is required.

MAC Address A unicast MAC address for which the switch has forwarding and/or filtering information. The format is a two byte hexadecimal VLAN ID number followed by a six byte MAC address with each byte separated by hyphens, for example 00-01-00-23-45-67-89-AB.

Source Port The port where this address was learned – i.e. the port through which the MAC address can be reached. In the above example, MM refers to the Management Module ports.

ifIndex The ifIndex of the MIB interface table entry associated with the port.

Status The status of this entry. The possible values are:

Learned

The entry was learned by observing the source MAC addresses of incoming traffic, and is currently in use.

Management

The value of the corresponding instance is also the value of an existing instance of dot1d StaticAddress. Currently this is used when enabling VLANs for routing.

Self The MAC address of one of the switch's physical interfaces.

GMRP Learned

The value of the corresponding instance was learned via GARP Multicast Registration Protocol (GMRP).

Other The value of the corresponding instance does not fall into one of the other categories.

Click the Search button to search for the specified MAC address.

Click the Refresh button to retrieve and display the database again starting with the first entry in the table.

Logs

This menu provides access to the following two logs:

- Message log
- Event log

The message log tracks non-critical error information, while the event log tracks critical event information.

Message log

This panel displays the message log maintained by the switch. The message log contains switch trace information that records non-critical problems. Message log information is not retained across a switch reset and wraps after 512 entries.



Time The time the event occurred, calculated from the time the switch was last reset, in days, hours, minutes and seconds.

File The source code filename identifying the code that detected the event.

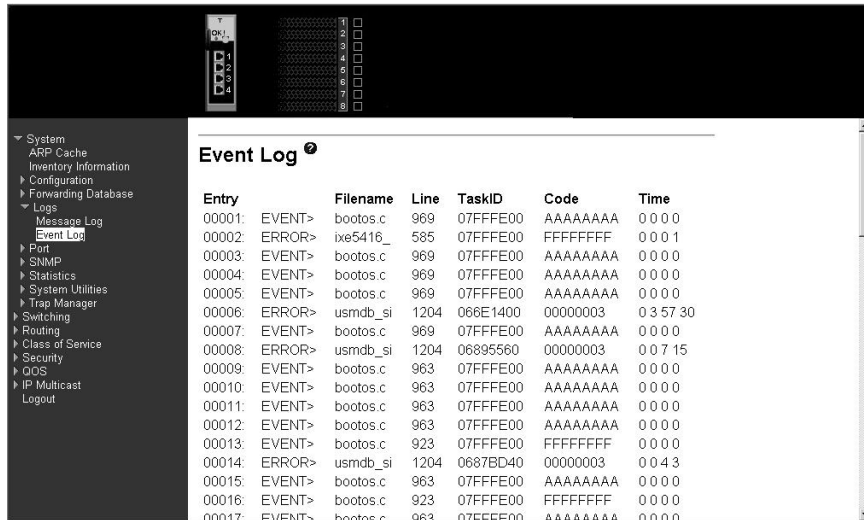
Line The line number within the source file of the code that detected the event.

Description An explanation of the problem being reported.

Click the Refresh button to retrieve and display the database again starting with the first entry in the table.

Event log

This panel displays the event log, which is used to hold error messages for critical events. After the event has been logged and the updated log has been saved in FLASH memory, the switch will be reset. The log can hold at least 2,000 entries and is erased when an attempt is made to add an entry after it is full. The event log is preserved across switch resets.



Entry	Filename	Line	TaskID	Code	Time	
00001:	EVENT>	bootos.c	969	07FFFE00	AAAAAAAA	00:00
00002:	ERROR>	ixe5416_	585	07FFFE00	FFFFFFFF	00:01
00003:	EVENT>	bootos.c	969	07FFFE00	AAAAAAAA	00:00
00004:	EVENT>	bootos.c	969	07FFFE00	AAAAAAAA	00:00
00005:	EVENT>	bootos.c	969	07FFFE00	AAAAAAAA	00:00
00006:	ERROR>	usmdb_si	1204	068E1400	00000003	03:57:30
00007:	EVENT>	bootos.c	969	07FFFE00	AAAAAAAA	00:00
00008:	ERROR>	usmdb_si	1204	06895560	00000003	00:07:15
00009:	EVENT>	bootos.c	963	07FFFE00	AAAAAAAA	00:00
00010:	EVENT>	bootos.c	963	07FFFE00	AAAAAAAA	00:00
00011:	EVENT>	bootos.c	963	07FFFE00	AAAAAAAA	00:00
00012:	EVENT>	bootos.c	963	07FFFE00	AAAAAAAA	00:00
00013:	EVENT>	bootos.c	923	07FFFE00	FFFFFFFF	00:00
00014:	ERROR>	usmdb_si	1204	0687BD40	00000003	00:04:3
00015:	EVENT>	bootos.c	963	07FFFE00	AAAAAAAA	00:00
00016:	EVENT>	bootos.c	923	07FFFE00	FFFFFFFF	00:00
00017:	EVENT>	bootos.c	963	07FFFE00	AAAAAAAA	00:00

Entry The number of the entry within the event log. The most recent entry is first.

Filename The source code filename identifying the code that detected the event.

Line The line number within the source file of the code that detected the event.

TaskID The OS-assigned ID of the task reporting the event.

Code The event code passed to the event log handler by the code reporting the event.

Time The time the event occurred, measured from the previous reset, in days, hours, minutes and seconds.

Click the Refresh button to retrieve and display the database again starting with the first entry in the table.

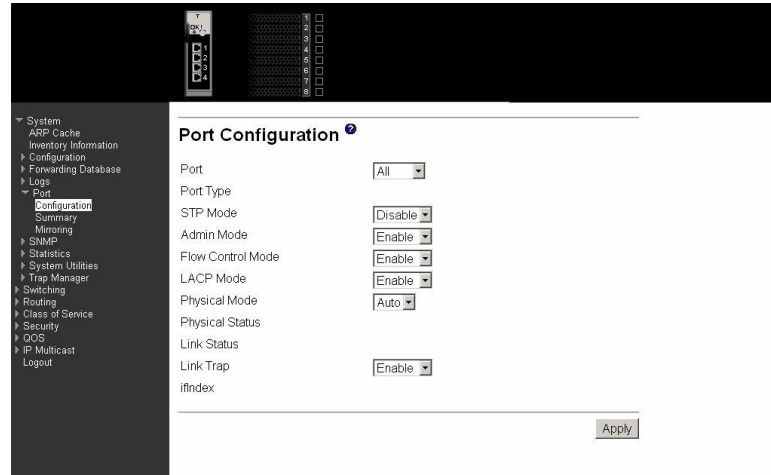
Port

This menu provides access to port configuration and display options, including:

- Configuration
- Summary
- Mirroring

Configuration

Use this panel to enable or disable one or more ports. The port will only participate in the network when it is enabled.



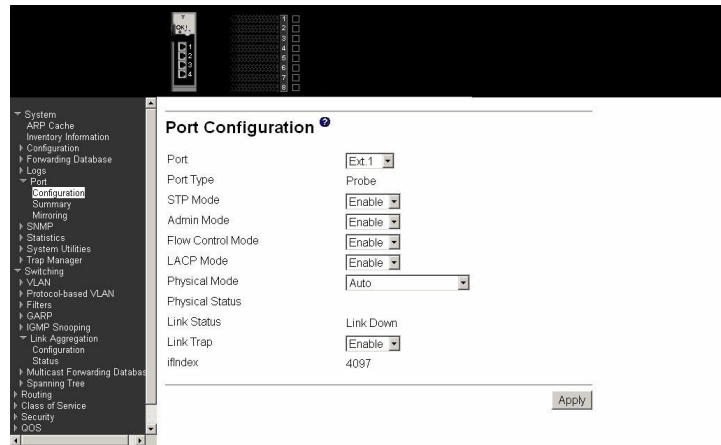
Port

Selects the interface for which data is to be displayed or configured.

Port Type

For normal and LAG ports this field will be blank. Otherwise the possible values are:

Probe Monitoring port, participating in Port Mirroring. Following is how this panel displays when the port type is Probe.



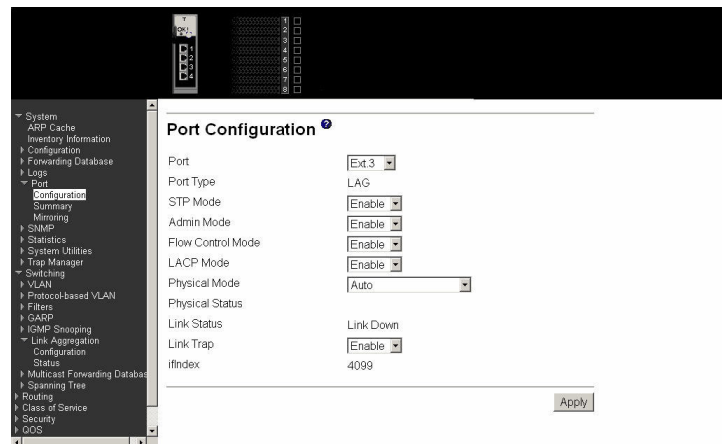
Mirrored

Port being mirrored.

LAG

Member of a Link Aggregation (LAG) trunk. Following is

how this panel displays when the port type is LAG.



STP Mode Select the Spanning Tree Protocol (STP) Administrative Mode for the port or LAG. The possible values are Enabled and Disabled.

Admin Mode Use the pull-down menu to select the port control administration state. You must select Enabled if you want the port to participate in the network. The factory default is Enabled.

Flow Control Mode Use the pull-down menu to Enable or Disable flow control for the port. The factory default is Disabled.

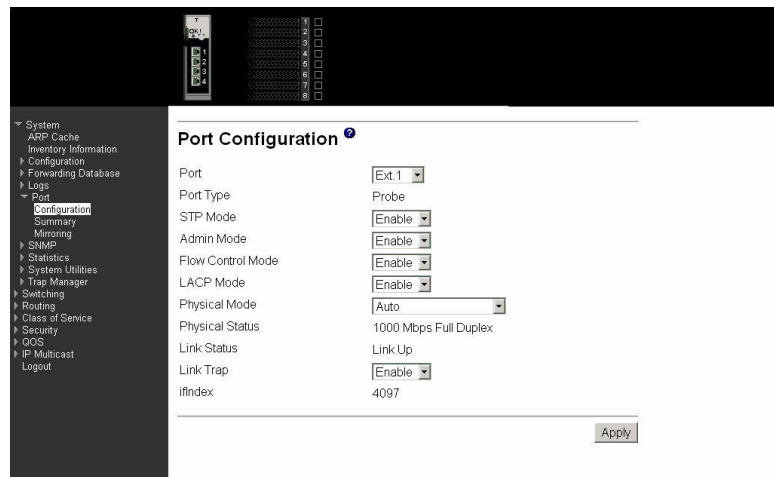
Mode Selects the Link Aggregation Control Protocol administration state. The mode must be Enabled in order for the port to participate in Link Aggregation. It can be Enabled or Disabled by selecting the corresponding line on the pull-down entry field. The factory default is Enabled.

Physical Mode Use the pull-down menu to select the port's speed and duplex mode. If you select auto the duplex mode and speed will be set by the auto-negotiation process. Note that the port's maximum capability (full duplex and 100 Mbps) will be advertised. Otherwise, your selection will determine the port's duplex mode and transmission rate. The factory default is auto. You can use this menu only for the external ports.

Physical Status Indicates the port speed and duplex mode. This field only displays if the Link Status is Up.

Link Status Indicates whether the Link is Up or Down. Following is how this

panel displays when the link status is link up.



Link Trap This object determines whether or not to send a trap when link status changes. The factory default is Enabled.

ifIndex The ifIndex of the interface table entry associated with this port.

Click the Apply button to update the switch with the values you entered. If you want the switch to retain the new values across a power cycle you must perform a save.

Summary

This panel displays the status of all ports in the box.

Port	Port Type	STP Mode	Forwarding State	Port Role	Admin Mode	Flow Control Mode
Bay.1		Disabled	Disabled	Disabled Port	Enable	Enable
Bay.2		Disabled	Disabled	Disabled Port	Enable	Enable
Bay.3		Disabled	Disabled	Disabled Port	Enable	Enable
Bay.4		Disabled	Disabled	Disabled Port	Enable	Enable
Bay.5		Disabled	Disabled	Disabled Port	Enable	Enable
Bay.6		Disabled	Disabled	Disabled Port	Enable	Enable
Bay.7		Disabled	Disabled	Disabled Port	Enable	Enable
Bay.8		Disabled	Disabled	Disabled Port	Enable	Enable
Ext.1		Enabled	Disabled	Disabled Port	Enable	Enable
Ext.2		Enabled	Disabled	Disabled Port	Enable	Enable
Ext.3	LAG	Enabled	Disabled	Disabled Port	Enable	Enable
Ext.4	LAG	Enabled	Disabled	Disabled Port	Enable	Enable
LAG.1		Enabled	Disabled	Disabled Port	Enable	Enable

Port Identifies the physical port.

Port Type If not blank, this field indicates that this port is a special type of port. The possible values are:

Mirrored

Port being mirrored.

Probe

Probe port, participating in Port Mirroring.

LAG Member of a link aggregation trunk.

STP Mode The Administrative Mode for the port or LAG. The possible values are Enabled and Disabled.

Forwarding State

The port's current spanning tree state. This state controls what action a port takes on receipt of a frame. If the bridge detects a malfunctioning port it will place that port into the Broken state. The other four states are defined in IEEE 802.1s as:

- Disabled
- Manual Forwarding
- Learning
- Forwarding

Port Role

Each Enabled bridge port is assigned a port role. The port role will be one of the following values: Root Port, Designated Port, Alternate Port, Backup Port, Master Port or Disabled Port.

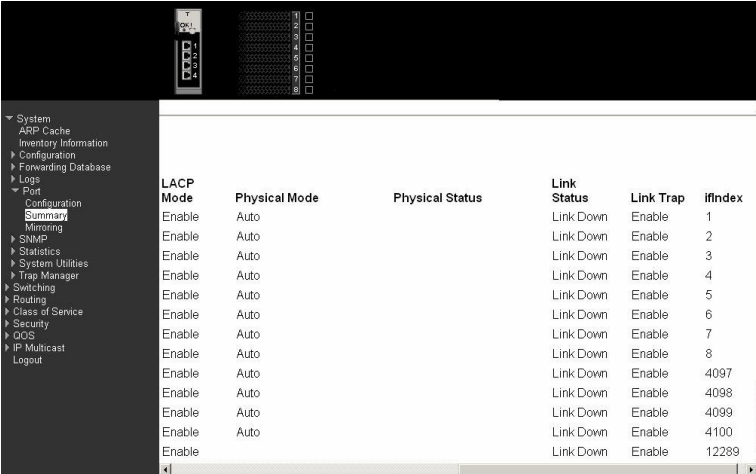
Admin Mode

Displays the port administration mode. The port must be Enabled in order for it to be allowed into the network. The factory default is Enabled.

Control Mode

Displays whether flow control is Enabled or Disabled on this port.

The following displays the right side of the panel. Descriptions of these fields follow.



LACP Mode	Physical Mode	Physical Status	Link Status	Link Trap	ifindex
Enable	Auto		Link Down	Enable	1
Enable	Auto		Link Down	Enable	2
Enable	Auto		Link Down	Enable	3
Enable	Auto		Link Down	Enable	4
Enable	Auto		Link Down	Enable	5
Enable	Auto		Link Down	Enable	6
Enable	Auto		Link Down	Enable	7
Enable	Auto		Link Down	Enable	8
Enable	Auto		Link Down	Enable	4097
Enable	Auto		Link Down	Enable	4098
Enable	Auto		Link Down	Enable	4099
Enable	Auto		Link Down	Enable	4100
Enable	Auto		Link Down	Enable	12289

LACP Mode

Displays whether Link Aggregation Control Protocol (LACP) is Enabled or Disabled on this port.

Physical Mode

Displays the selected port speed and duplex mode. If auto-negotiation support is selected, then the duplex mode and speed will be set by the auto-negotiation process. Note that the port's maximum capability will be advertised. Otherwise, you must enter the port's speed and duplex mode manually. The factory default is auto.

Physical Status

Indicates the current port speed and duplex mode.

Link Status

Indicates whether the link is Up or Down.

Link Trap

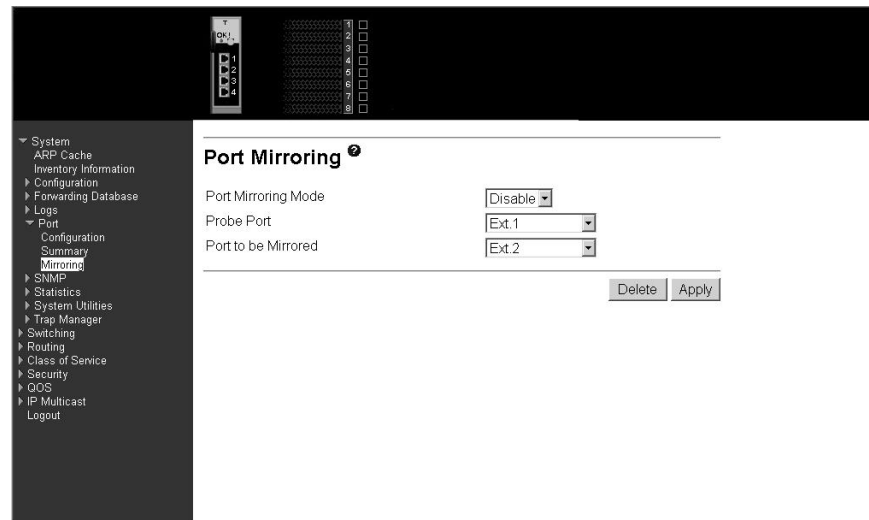
Indicates whether or not a trap will be sent when link status changes. The factory default is Enabled.

ifIndex

Indicates the ifIndex of the interface table entry associated with this port.

Mirroring

This panel displays the port mirroring information for the switch module.



Port Mirroring Mode

Select the Port Mirroring Mode by selecting the corresponding line on the pull-down entry field. The factory default is Disabled.

Probe Port

The interface you want to act as the Probe. Once configured there is no network connectivity on the probe port. The probe port will not forward or receive any traffic. The probe tool attached to the probe port will not be able to ping the switch or through the switch, and nobody will be able to ping the probe tool.

Port to be Mirrored

The interface selected as the Mirror. Every packet seen at the mirrored port is copied to the probe port. That includes all packets received and admitted, received and dropped, and transmitted out of the mirrored port.

Click the Delete button to remove the Port Mirroring configuration. The mode must be Disabled before the configuration can be deleted.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

SNMP

This menu provides access to the following Simple Network Management Protocol (SNMP) options:

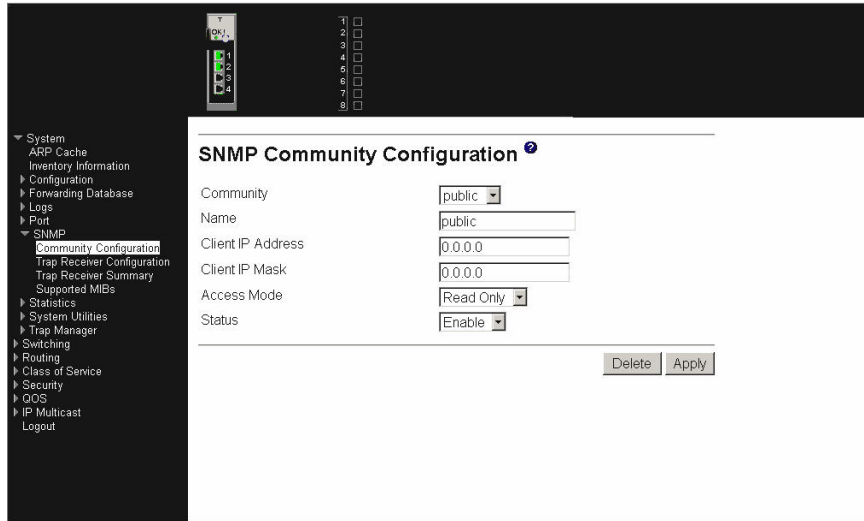
- Community configuration
- Trap receiver configuration
- Trap receiver summary
- Supported MIBs

Community configuration

By default, two SNMP Communities exist:

- private, with Read/Write privileges and status set to Enable
- public, with Read-only privileges and status set to Enable

These are well-known communities; you can use this menu to change the defaults or to add other communities. Only the communities that you define using this menu will have access to the switch using the SNMPv1 and SNMPv2c protocols. Only those communities with Read-Write privileges will have access to this menu via SNMP. Use this panel when you are using the SNMPv1 or SNMPv2c protocol; if you want to use SNMP v3 you should use the User Accounts menu.



Six communities are supported. You can add, change or delete communities. The switch does not have to be reset for changes to take effect.

The SNMP agent of the switch complies with SNMPv1 (for more about the SNMP specification, see the SNMP RFCs). The SNMP agent sends traps through TCP/IP to an external SNMP manager based on the SNMP configuration (the trap receiver and other SNMP community parameters).

Community Use this pull-down menu to select one of the existing community names, or select Create to add a new one.

Name A community name is associated with the switch and with a set of SNMP managers that manage it with a specified privileged level. The length of the name can be up to 16 case-sensitive characters. There are two default community names: public (with Read-only access) and private (with Read/Write access). You can replace these default community names with unique identifiers for each community. The default values for the remaining four community names are blank. Community names in the SNMP community table must be unique. If you make multiple entries using the same community name, the first entry is kept and processed and all duplicate entries are ignored.

Client IP Address Enter the IP address (or portion thereof) from which this device will accept SNMP packets with the associated community name. The requesting entity's IP address is ANDed with the Client IP mask before being compared to the Client IP address. Note that if the

Client IP mask is set to 0.0.0.0, an IP address of 0.0.0.0 matches all IP addresses. The default value is 0.0.0.0.

Client IP Mask

Enter the mask to be ANDed with the requesting entity's IP address before comparison with the Client IP address. If the result matches the Client IP address then the address is an authenticated IP address. For example, if the IP address = 9.47.128.0 and the corresponding Subnet Mask = 255.255.255.0, a range of incoming IP addresses would match, i.e. the incoming IP address could equal 9.47.128.0 - 9.47.128.255. The default value is 0.0.0.0.

Access Mode

Specify the access level for this community by selecting Read/Write or Read-only from the pull-down menu. This field restricts access to switch information.

Status

Specify the status of this community by selecting Enable or Disable from the pull-down menu. This field activates or deactivates an SNMP community. If a community is Enabled, an SNMP manager associated with this community is allowed to access the switch. If the community is Disabled, no SNMP requests using this community name are accepted. In this case the SNMP manager associated with this community cannot manage the switch until the Status is changed back to Enable.

Click the Delete button to delete the currently selected Community Name. If you want the switch to retain the new values across a power cycle, you must perform a save.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Trap receiver configuration

Use this panel to assign a new IP address to a specified trap receiver community. The maximum length of name is 16 case-sensitive alphanumeric characters.

IP addresses in the SNMP trap receiver table must be unique. If you make multiple entries using the same IP address, the first entry is retained and processed. All duplicate entries are ignored.



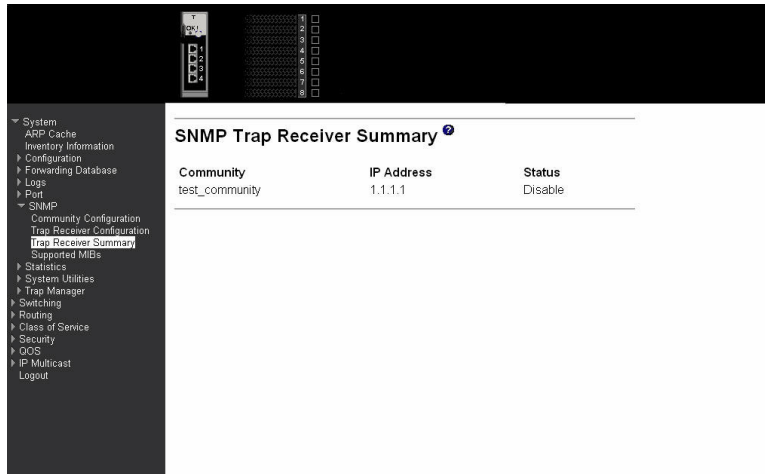
- Community** This field adds an SNMP trap receiver community name and associated IP address.
- Name** Enter the community string for the SNMP trap packet to be sent to the trap manager. This can be up to 16 characters and is case sensitive.
- IP Address** Enter the IP address to receive SNMP traps from this device.
- Status** This field Enables or Disables the SNMP trap receiver identified by trap receiver community name and IP address. Enabled trap receivers are active (able to receive traps). Disabled trap receivers are inactive (not able to receive traps).

Click the Delete button to delete the currently selected Community Name. If you want the switch to retain the new values across a power cycle, you must perform a save.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Trap receiver summary

This panel displays information about SNMP trap receivers. Trap messages are sent across a network to an SNMP Network Manager. These messages alert the manager to events occurring within the switch or on the network. Up to six trap receivers are supported at the same time.



- Community** Displays the community string for the SNMP trap packet to be sent to the trap manager. Note that trap receiver communities and SNMP communities are separate and distinct.
- IP Address** Displays the IP address to receive SNMP traps from this device.
- Status** Indicates whether traps are currently Enabled for this community:
 - Enable** Traps will be sent.
 - Disable** Traps will not be sent.

Supported Management Information Bases (MIB)

This panel displays a list of all the MIBs supported by the switch.



Name	Description
RFC 1907 - SNMPV2-MIB	The MIB module for SNMPv2 entities
RFC 2819 - RMON-MIB	Remote Network Monitoring Management Information Base
SBHEGESW-REF-MIB	SBHEGESW Reference
SNMP-COMMUNITY-MIB	This MIB module defines objects to help support coexistence between SNMPv1, SNMPv2, and SNMPv3.
SNMP-FRAMEWORK-MIB	The SNMP Management Architecture MIB
SNMP-MPD-MIB	The MIB for Message Processing and Dispatching
SNMP-NOTIFICATION-MIB	The Notification MIB Module
SNMP-TARGET-MIB	The Target MIB Module
SNMP-USER-BASED-SM-MIB	The management information definitions for the SNMP User-based Security Model
SNMP-VIEW-BASED-ACM-MIB	The management information definitions for the View-based Access Control Model for SNMP.

Name The RFC number if applicable and the name of the MIB.

Description The RFC title or MIB description.

Click the Refresh button to retrieve and display the database again starting with the first entry in the table.

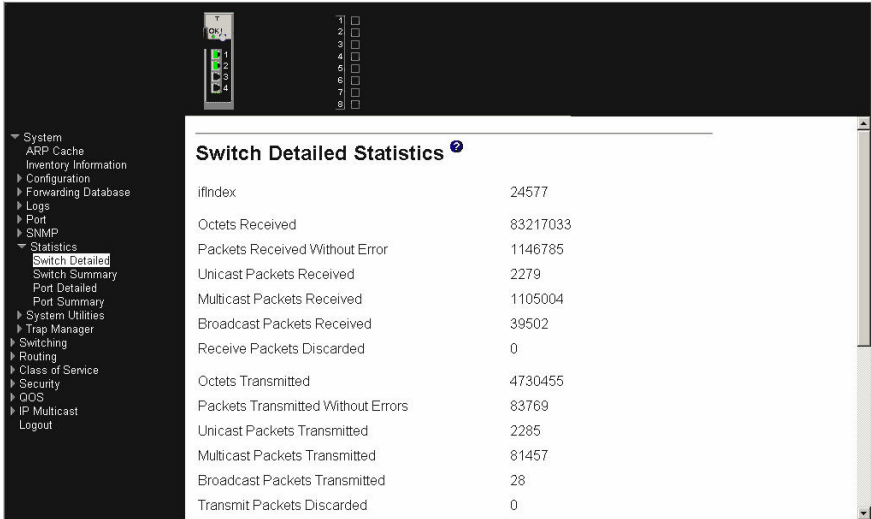
Statistics

This menu provides access to menu options that display various switch statistics, including:

- Switch detailed
- Switch summary
- Port detailed
- Port summary

Switch detailed

This panel displays detailed statistics for all CPU traffic.



Name	Value
IfIndex	24577
Octets Received	83217033
Packets Received Without Error	1146785
Unicast Packets Received	2279
Multicast Packets Received	1105004
Broadcast Packets Received	39502
Receive Packets Discarded	0
Octets Transmitted	4730455
Packets Transmitted Without Errors	83769
Unicast Packets Transmitted	2285
Multicast Packets Transmitted	81457
Broadcast Packets Transmitted	28
Transmit Packets Discarded	0

ifIndex This object indicates the ifIndex of the interface table entry associated with the processor of this switch.

Received

Octets Received

The total number of octets of data received by the processor (excluding framing bits but including FCS octets).

Packets Received Without Error

The total number of packets (including broadcast packets and multicast packets) received by the processor.

Unicast Packets Received

The number of subnetwork-unicast packets delivered to a higher-layer protocol.

Multicast Packets Received

The number of packets received that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.

Broadcast Packets Received

The number of packets received that were directed to a broadcast address. Note that this number does not include packets directed to the multicast address.

Receive Packets Discarded

The number of inbound packets that were chosen to be discarded even though no errors had been detected that would prevent their being deliverable to a higher-layer protocol. One possible reason for discarding a packet could be to free up buffer space.

Transmitted

Octets Transmitted

The total number of octets of data transmitted on the network including framing bits.

Packets Transmitted Without Errors

The total number of packets that have been transmitted on the network without an error occurring.

Unicast Packets Transmitted

The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent.

Multicast Packets Transmitted

The total number of packets that higher-level protocols requested be transmitted to a multicast address, including those that were discarded or not sent.

Broadcast Packets Transmitted

The total number of packets that higher-level protocols requested be transmitted to a broadcast address, including those that were discarded or not sent.

Transmit Packets Discarded

The number of outbound packets that were chosen to be discarded even though no errors had been detected. One possible reason for discarding a packet could be to free up buffer space.

Table Entries:

Most Address Entries Ever Used

The highest number of Forwarding Database Address Table entries used by this switch module since the last reboot.

Address Entries In Use

The number of learned and static Forwarding Database Address Table entries currently in use by this switch module.

Maximum VLAN Entries

The maximum number of VLANs allowed on the switch module.

Most VLAN Entries Ever Used

The highest number of VLANs that have been active on this switch module since the last reboot.

Static VLAN Entries

The number of VLANs currently active on this switch module that were created statically.

Dynamic VLAN Entries

The number of VLANs currently active on this switch module that were created by GARP VLAN Registration Protocol (GVRP) registration.

VLAN Deletes

The number of VLANs that have been created and then deleted on this switch module since the last reboot.

Time Since Counters Last Cleared:

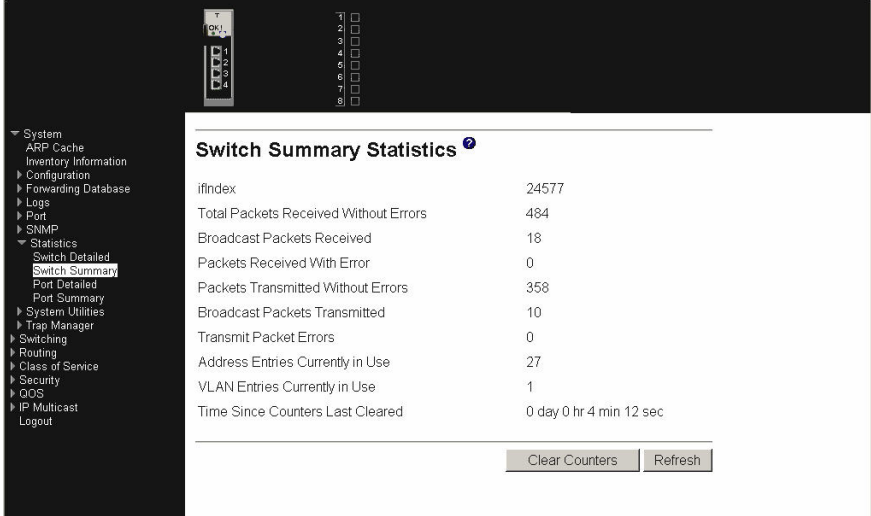
The elapsed time in days, hours, minutes and seconds since the statistics for this port were last cleared.

Click the Clear Counters button to clear all the counters, resetting all summary and switch detailed statistics to defaults, except for the counts of discarded packets, which cannot be cleared.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

Switch summary

This panel displays a summary of the statistics for CPU traffic.



The screenshot shows a web-based network management interface. On the left is a navigation menu with categories like System, Configuration, Forwarding Database, Logs, Port, SNMP, Statistics, and Switching. The 'Statistics' section is expanded, showing 'Switch Detailed' and 'Switch Summary' (which is selected). The main content area is titled 'Switch Summary Statistics' and contains a table of statistics. At the bottom of the table are two buttons: 'Clear Counters' and 'Refresh'.

Switch Summary Statistics	
ifIndex	24577
Total Packets Received Without Errors	484
Broadcast Packets Received	18
Packets Received With Error	0
Packets Transmitted Without Errors	358
Broadcast Packets Transmitted	10
Transmit Packet Errors	0
Address Entries Currently in Use	27
VLAN Entries Currently in Use	1
Time Since Counters Last Cleared	0 day 0 hr 4 min 12 sec

ifIndex This object indicates the ifIndex of the interface table entry associated with the processor of this switch.

Total Packets Received Without Errors

The total number of packets (including multicast and broadcast packets) received by the processor without an error occurring.

Broadcast Packets Received

The total number of packets received that were directed to the broadcast address. Note that this does not include multicast packets.

Packets Received With Error

The number of inbound packets that contained errors that prevented them being delivered to a higher-layer protocol.

Packets Transmitted Without Errors

The total number of packets transmitted from the switch module without an error occurring.

Broadcast Packets Transmitted

The total number of packets that higher-layer protocols requested to be transmitted to the broadcast address, including those that were discarded or not sent.

Transmit Packet Errors

The number of outbound packets that could not be transmitted because of errors.

Address Entries Currently In Use

The number of learned and static Forwarding Database Address Table entries currently in use by this switch module.

VLAN Entries Currently In Use

The number of VLANs currently in the VLAN table on this switch module.

Time Since Counters Last Cleared

The elapsed time in days, hours, minutes and seconds since the statistics for this port were last cleared.

Click the Clear Counters button to clear all the counters, resetting all summary and switch detailed statistics to defaults, except for the counts of discarded packets, which cannot be cleared.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

Port detailed

This panel displays detailed statistics for a specified port.

Port	Bay 1
ifIndex	1
Octets Received	0
Packets Received 64 Octets	0
Packets Received 65-127 Octets	0
Packets Received 128-255 Octets	0
Packets Received 256-511 Octets	0
Packets Received 512-1023 Octets	0
Packets Received 1024-1518 Octets	0
Packets Received 1519-1522 Octets	0
Packets Received > 1522 Octets	0
Total Packets Received Without Errors	0
Unicast Packets Received	0
Multicast Packets Received	0
Broadcast Packets Received	0
Total Packets Received with MAC Errors	0
Jabbers Received	0

Port Use this field to select the port for which to display statistics. Click the down arrow to display the list of ports from which to choose.

ifIndex This object indicates the ifIndex of the interface table entry associated with this port.

Packets Received:

Octets Received

The total number of octets of data (including those in bad packets) received on the network (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of Ethernet utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects should be sampled before and after a common interval.

Packets Received 64 Octets

The total number of packets (including bad packets) received that were 64 octets in length (excluding framing bits but including FCS octets).

Packets Received 65-127 Octets

The total number of packets (including bad packets) received that were between 65 and 127 octets in length (excluding framing bits but including FCS octets).

Packets Received 128-255 Octets

The total number of packets (including bad packets) received that were between 128 and 255 octets in length (excluding framing bits but including FCS octets).

Packets Received 256-511 Octets

The total number of packets (including bad packets) received that were between 256 and 511 octets in length (excluding framing bits but including FCS octets).

Packets Received 512-1023 Octets

The total number of packets (including bad packets) received that were between 512 and 1023 octets in length (excluding framing bits but including FCS octets).

Packets Received 1024-1518 Octets

The total number of packets (including bad packets) received that were between 1024 and 1518 octets in length (excluding framing bits but including FCS octets).

Packets Received 1519-1522 Octets

The total number of packets (including bad packets) received that were between 1519 and 1522 octets in length (excluding framing bits but including FCS octets).

Packets Received >1522 Octets

The total number of packets (including bad packets) received that were >1522 octets in length (excluding framing bits but including FCS octets).

Total Packets Received Without Error

Total Packets Received Without Errors

The total number of packets received that were without error.

Unicast Packets Received

The number of subnetwork-unicast packets delivered to a higher-layer protocol.

Multicast Packets Received

The number of packets received that were directed to a multicast address. Note that this number does not include packets directed to the broadcast address.

Broadcast Packets Received

The number of packets received that were directed to a broadcast address. Note that this number does not include packets directed to the multicast address.

Total Packets Received with MAC Errors

Total Packets Received with MAC Errors

The total number of inbound packets that contained errors that prevented them from being delivered to a higher-layer protocol.

Jabbers Received

The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). Note that this definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.

Fragments/Undersize Received

The total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets).

Alignment Errors

The total number of packets received that had a length (excluding framing

bits but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with a non-integral number of octets.

Rx FCS Errors

The total number of packets received that had a length (excluding framing bits but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with an integral number of octets.

Total Received Packets Not Forwarded

802.3x Pause Frames Received

A count of MAC Control frames received on this interface with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.

Total Packets Transmitted (Octets)

Total Packets Transmitted (Octets)

The total number of octets of data (including those in bad packets) transmitted on the network (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of Ethernet utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects should be sampled before and after a common interval.

Packets Transmitted 64 Octets

The total number of packets (including bad packets) transmitted that were 64 octets in length (excluding framing bits but including FCS octets).

Packets Transmitted 65-127 octets

The total number of packets (including bad packets) transmitted that were between 65 and 127 octets in length (excluding framing bits but including FCS octets).

Packets Transmitted 128-255 Octets

The total number of packets (including bad packets) transmitted that were between 128 and 255 octets in length (excluding framing bits but including FCS octets).

Packets Transmitted 256-511 Octets

The total number of packets (including bad packets) transmitted that were between 256 and 511 octets in length (excluding framing bits but including FCS octets).

Packets Transmitted 512-1023 Octets

The total number of packets (including bad packets) transmitted that were between 512 and 1023 octets in length (excluding framing bits but including FCS octets).

Packets Transmitted 1024-1518 Octets

The total number of packets (including bad packets) transmitted that were between 1024 and 1518 octets in length (excluding framing bits but including FCS octets).

Packets Transmitted 1519-1522 Octets

The total number of packets (including bad packets) transmitted that were between 1519 and 1530 octets in length (excluding framing bits but including FCS octets).

Max Info

The maximum size of the information (non-MAC) field that this port will receive or transmit.

Total Packets Transmitted Successfully

Total Packets Transmitted Successfully

The total number of packets that have been transmitted by this port to its segment without an error occurring.

Unicast Packets Transmitted

The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent.

Multicast Packets Transmitted

The total number of packets that higher-level protocols requested be transmitted to a multicast address, including those that were discarded or not sent.

Broadcast Packets Transmitted

The total number of packets that higher-level protocols requested be transmitted to a broadcast address, including those that were discarded or not sent.

Total Transmit Errors

Total Transmit Errors

The sum of Single, Multiple and Excessive Collisions.

Tx FCS Errors

The total number of packets transmitted that had a length (excluding framing bits but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with an integral number of octets.

Tx Oversized

The total number of packets that exceeded the maximum permitted frame size. This counter has a maximum increment rate of 815 counts per second at 10 Mbps.

Underrun Errors

The total number of packets discarded because the transmit FIFO buffer became empty during frame transmission.

Total Transmit Packets Discarded

Total Transmit Packets Discarded

The sum of single collision frames discarded, multiple collision frames discarded, and excessive collision frames discarded.

Single Collision Frames

The number of successfully transmitted packets which encountered exactly one collision.

Multiple Collision Frames

The number of successfully transmitted packets which encountered more than one collision.

Excessive Collision Frames

The number of packets which were not successfully transmitted because of excessive collisions.

STP BPDUs Received

The number of STP BPDUs (Bridge Protocol Data Units) received by the spanning tree layer.

STP BPDUs Transmitted

The number of STP BPDUs transmitted from the spanning tree layer.

RSTP BPDUs Received

The number of RSTP BPDUs received at the selected port.

RSTP BPDUs Transmitted

The number of RSTP BPDUs transmitted from the selected port.

802.3x Pause Frames Transmitted

A count of MAC Control frames transmitted on this interface with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.

GVRP PDUs Received

The number of GVRP PDUs received by the Generic Attribute Registration Protocol (GARP) layer.

GVRP PDUs Transmitted

The number of GVRP PDUs transmitted by the GARP layer.

GVRP Failed Registrations

The number of times attempted GVRP registrations could not be completed.

GMRP PDUs Received

The number of GMRP PDUs received by the GARP layer.

GMRP PDUs Transmitted

The number of GMRP PDUs transmitted by the GARP layer.

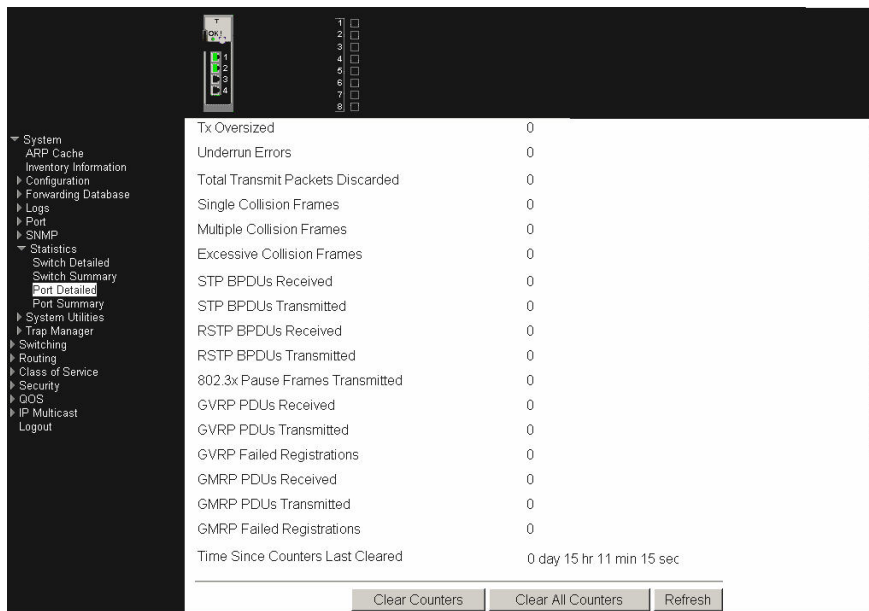
GMRP Failed Registrations

The number of times attempted GMRP registrations could not be completed.

Time Since Counters Last Cleared

The elapsed time in days, hours, minutes and seconds since the statistics for this port were last cleared.

The following displays the bottom of the panel, showing the buttons available.



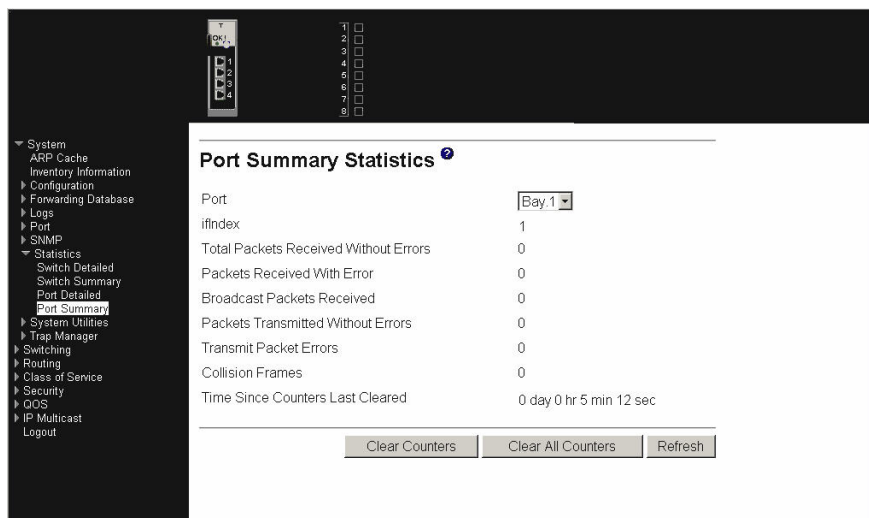
Click the Clear Counters button to clear all the counters, resetting all statistics for this port to default values.

Click the Clear All Counters button to clear all the counters for all ports, resetting all statistics for all ports to default values.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

Port summary

This panel displays a summary of the statistics for a specified port.



Port Use this field to select the port for which to display statistics. Click the down arrow to display the list of ports from which to choose.

ifIndex This object indicates the ifIndex of the interface table entry associated with this port on an adapter.

Total Packets Received Without Errors

The total number of packets (including multicast and broadcast packets) received on this port without an error occurring.

Packets Received With Error

The number of inbound packets that contained errors that prevented them being delivered to a higher-layer protocol.

Broadcast Packets Received

The total number of packets received that were directed to the broadcast address. Note that this does not include multicast packets.

Packets Transmitted Without Errors

The total number of packets transmitted from the interface without an error occurring.

Transmit Packet Errors

The number of outbound packets that could not be transmitted because of errors.

Collision Frames

The best estimate of the total number of collisions on this Ethernet segment.

Time Since Counters Last Cleared

The elapsed time in days, hours, minutes and seconds since the statistics for this port were last cleared.

Click the Clear Counters button to clear all the counters, resetting all statistics for this port to default values.

Click the Clear All Counters button to clear all the counters for all ports, resetting all statistics for all ports to default values.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

System utilities

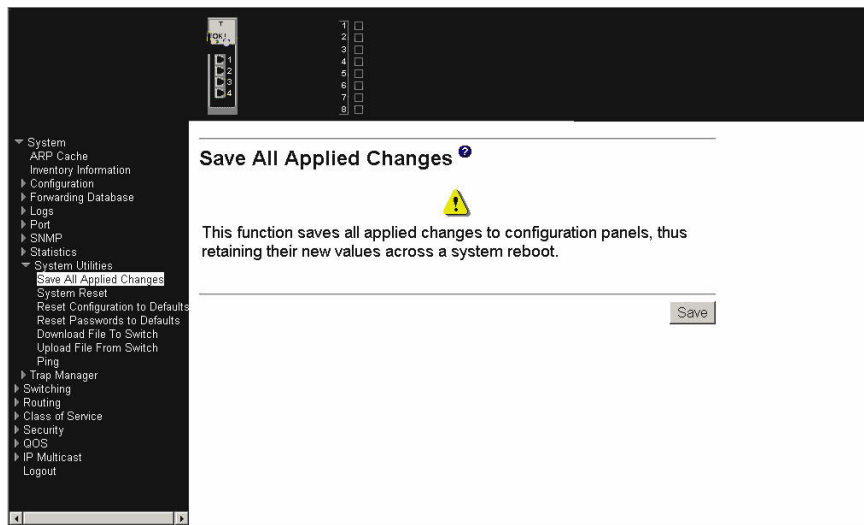
This menu provides access to several systems-related panels. These include:

- Save all applied changes
- System reset
- Reset configuration to default
- Reset passwords to default
- Download file to switch
- Upload file from switch
- Ping

Save all applied changes

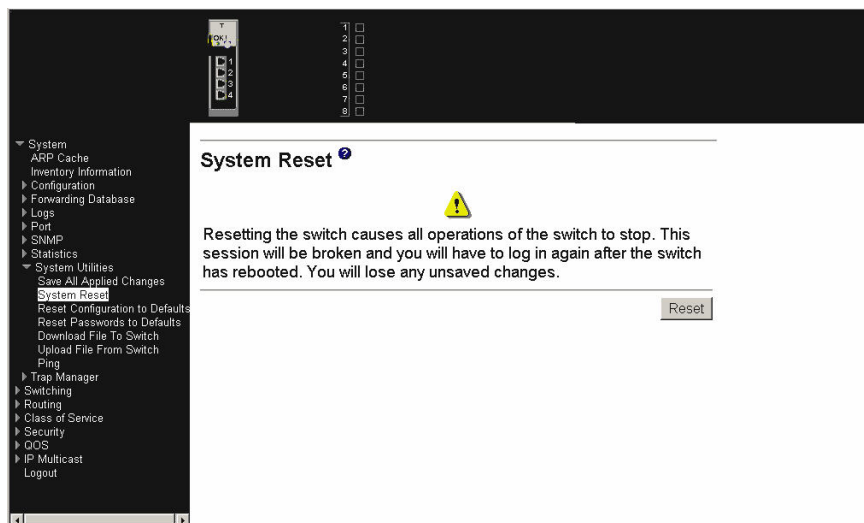
Click the Save button to have configuration changes you have made saved across a system reboot. All changes submitted since the previous save or system reboot

will be retained by the switch.



System reset

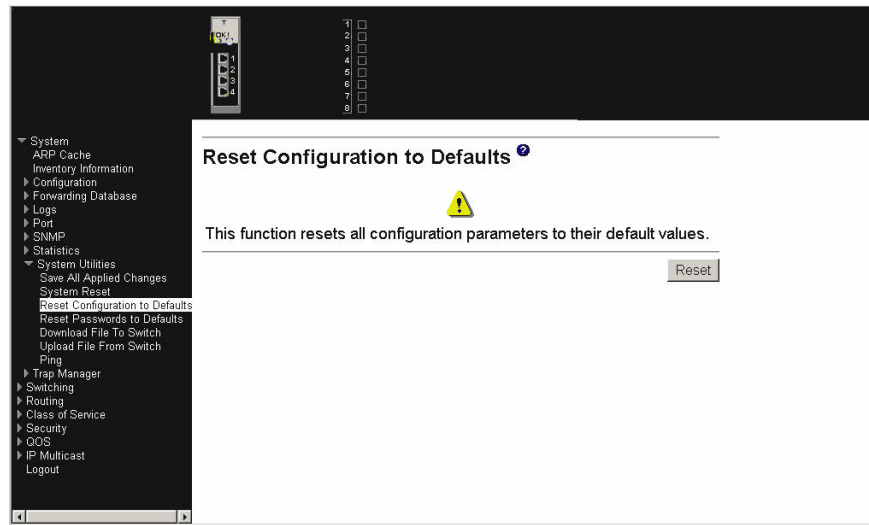
Click the Reset button to reset the switch without powering off. Reset means that all network connections are terminated and the boot code executes. The switch uses the stored configuration to initialize the switch. You are prompted to confirm that the reset should proceed. A successful reset is indicated by the LEDs on the switch.



Reset configuration to defaults

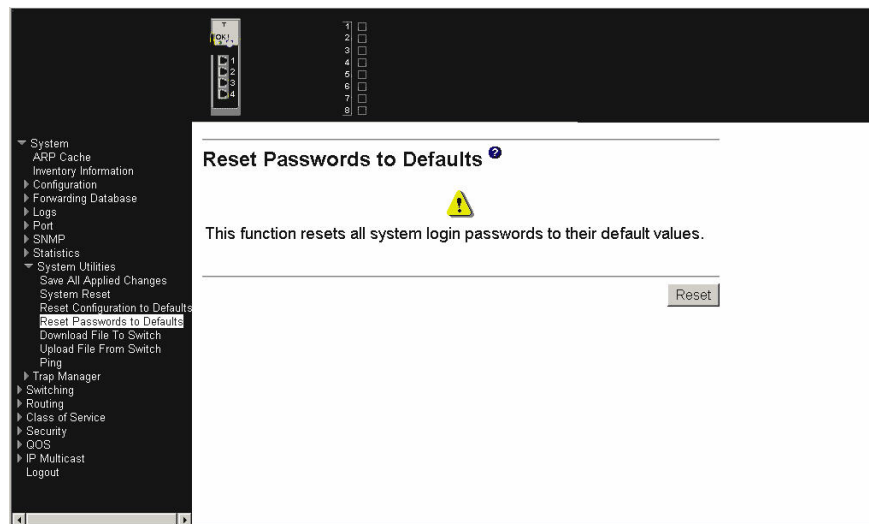
Click the Reset button to reset the configuration of the switch module to the factory defaults. The switch is automatically reset when this command is processed. All configuration changes that you have made, including those saved to NVRAM, will

be lost. You are prompted to confirm that the reset should proceed.



Reset passwords to defaults

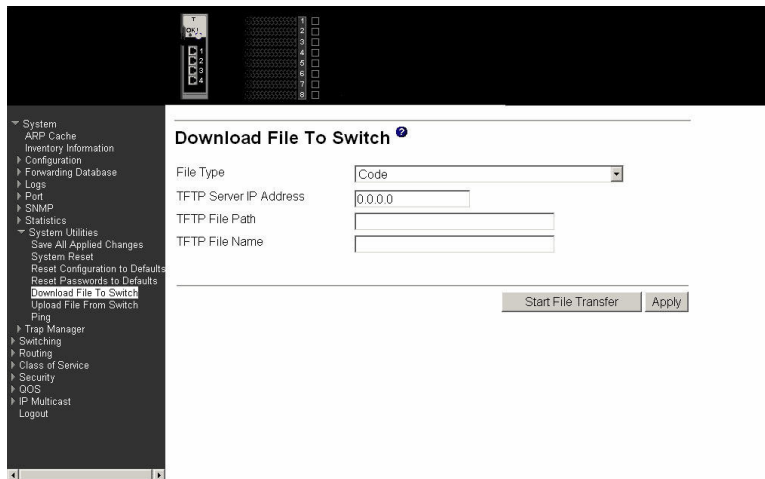
Click the Reset button to reset all user passwords to the factory defaults (since only the ADMIN can set passwords, this is blank). You are prompted to confirm that the password reset should proceed.



Download file to switch

Use this panel to configure the information needed to download a file to the switch.

Note: TFTP server software must be running on the management station for the TFTP upload or download to work.



File Type

Specify the type of file to be downloaded to the switch:

- Config
- Code
- SSH RSA1 Key File
- SSH RSA2 Key PEM File
- SSH DSA Key PEM File
- SSL Trusted Root Certificate PEM File
- SSL Server Certificate PEM File
- SSL DH Weak Encryption Parameter PEM File
- SSL DH Strong Encryption Parameter PEM File

TFTP Server IP Address

Enter the IP address of the TFTP server. The factory default is 0.0.0.0.

TFTP File Path

This field specifies the directory path on the TFTP server where the file to be downloaded to the switch is located. The switch will retain the last file path used.

TFTP File Name

This field specifies the name of the file that is to be downloaded to the switch. The switch will remember the last file name used.

The last row of the table is used to display information about the progress of the file transfer. The screen will refresh automatically until the file transfer completes.

The Ethernet switch module software supports the use of a TFTP client. The TFTP client path statement requirement is server dependent. A path statement is generally required to setup the TFTP client; however, the client path can remain blank. See the example of the path setup.

TFTP Upload Example:

The TFTP upload example details three scenarios for TFTP client-to-server file transfer. Each scenario involves uploading the config.bin file from the switch to the location c:\tftp\ on the server. The different scenarios are detailed below:

Table 8. TFTP Upload Scenarios

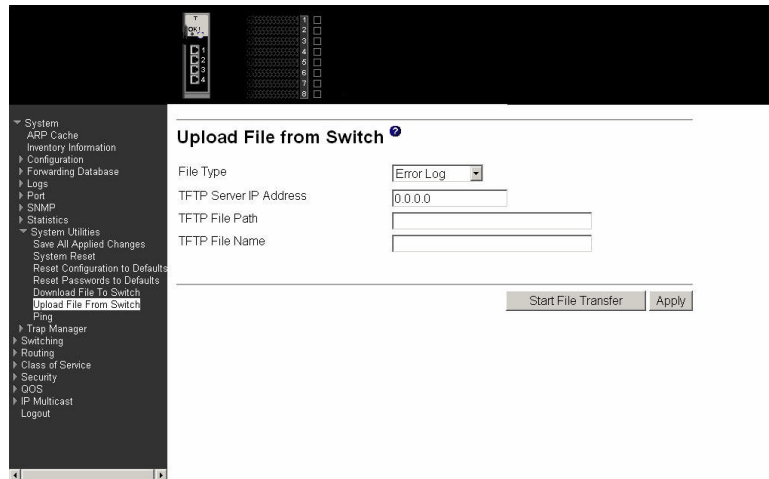
TFTP Server path	TFTP Client path
c:\tftp\	blank
c:\	tftp\
c:	\tftp\

Click the Start File Transfer button to apply any changes made to the fields and initiate the download.

Click the Apply button to send the updated screen to the switch; this does not perform the file download.

Upload file from switch

Use this panel to configure the information needed to upload a file from the switch. See the previous menu option “Download file to switch” on page 101 for more information about specifying TFTP File Paths and Names.



File Type

This field sets the type of file to be uploaded from the switch. The datatype is one of the following:

- Configuration
- Error Log
- Message Log
- Trap Log (default)

TFTP Server IP Address

Enter the IP address of the TFTP server. The factory default is 0.0.0.0.

TFTP File Path

This field specifies the directory path on the TFTP server where the file to be uploaded from the switch is to be located. The switch will remember the last file path used.

TFTP File Name

This field specifies the name of the file that is to be uploaded from the switch. The switch will remember the last file name used.

The last row of the table is used to display information about the progress of the file transfer. The screen will refresh automatically until the file transfer completes.

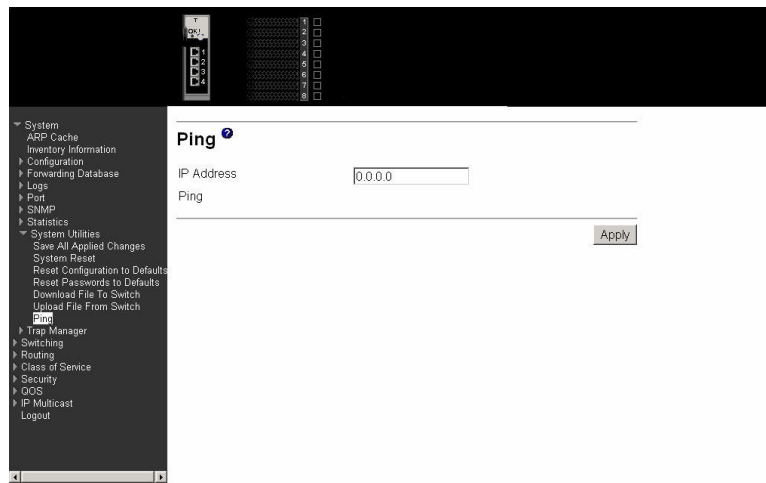
Click the Start File Transfer button to apply any changes made to the fields and initiate the upload.

Click the Apply button to send the updated screen to the switch; this does not perform the file upload. This command is valid only when the transfer mode is TFTP.

Ping

Use this panel to have the switch transmit a Ping request to a specified IP address. This checks whether the switch can communicate with a particular IP device. When you click the Apply button, the switch will send three pings and the results will be displayed in the Ping field, below the IP address.

The switch can be pinged from any IP workstation with which the switch is connected through the default VLAN (VLAN 1), as long as there is a physical path between the switch and the workstation.



IP Address Enter the IP address of the station you want the switch to ping. The initial value is blank. The IP address you enter is not retained across a power cycle.

Ping Displays the results of the ping. If a reply to the ping is not received, you will see No Reply Received from IP xxx.xxx.xxx.xxx, otherwise you will see Reply received from IP xxx.xxx.xxx.xxx: (send count = 3, receive count = n).

Click the Apply button to initiate the ping.

Trap manager

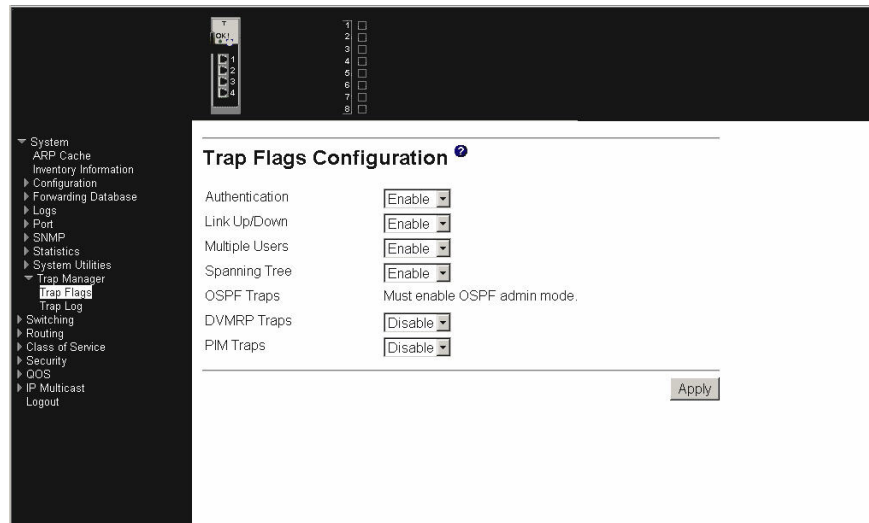
The following trap-related panels are available from this menu:

- Trap flags
- Trap log

Trap flags

This panel displays trap conditions. When the condition identified by an active trap is encountered by the switch, a trap message will be sent to any enabled SNMP Trap Receivers and a message will be written to the trap log. Cold and warm start

traps are always enabled.



Authentication

Indicates whether authentication failure traps will be sent (Enable) or not (Disable). This field Enables or Disables the Authentication Flag, which determines whether a trap message is sent when the switch detects an authentication failure. The factory default is Enabled.

Link Up/Down

Indicates whether a trap will be sent when the link status changes from Up to Down or vice versa. This field Enables or Disables Link Up/Down traps for the entire switch. When Enabled, link trap messages are sent only if the Link Trap flag associated with the affected port is also set to Enabled.

Multiple Users

Indicates whether a trap will be sent when the same user ID is logged into the switch more than once at the same time (either via Telnet or the serial port). This field Enables or Disables Multiple User traps. When Enabled, a multiple user trap message is sent when a user logs in to the terminal interface (EIA 232 or Telnet) and there is an existing terminal interface session for the same user account.

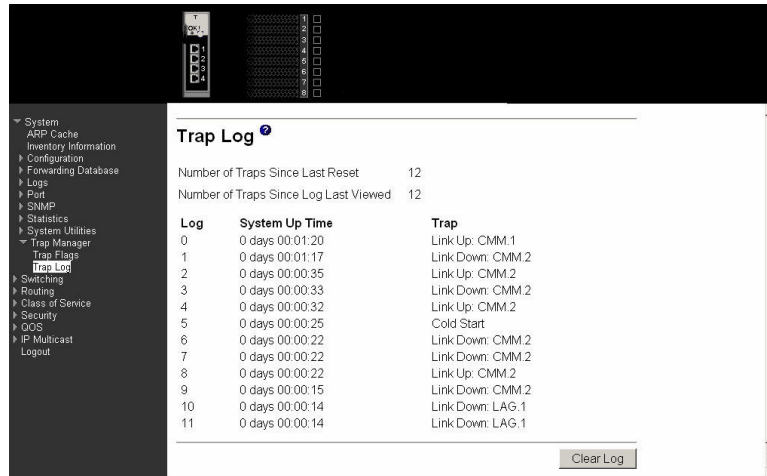
Spanning Tree

Indicates whether spanning tree traps will be sent. This field Enables or Disables STP traps. When Enabled, topology change notification trap messages will be sent.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. These changes will not be retained across a power cycle unless a save is performed.

Trap log

This panel displays the entries in the trap log.



Number of Traps Since Last Reset

The number of traps that have occurred since the last time the switch was reset.

Number of Traps Since Log Last Viewed

The number of traps that have occurred since the traps were last displayed. Displaying the traps by any method (terminal interface display, Web display, upload file from switch, etc.) will cause this counter to be cleared to 0.

Log

The sequence number of this trap.

System Up Time

The time at which this trap occurred, expressed in days, hours, minutes and seconds since the last reboot of the switch

Trap

Information identifying the trap.

Click the Clear Log button to clear all entries in the log. Subsequent displays of the log will only show new log entries.

Switching

This menu provides access to all the switch-related processing screens. Options on this menu are:

- VLAN
- Protocol-based VLAN
- Filters
- GARP
- IGMP snooping
- Link aggregation
- Multicast forwarding database
- Spanning tree

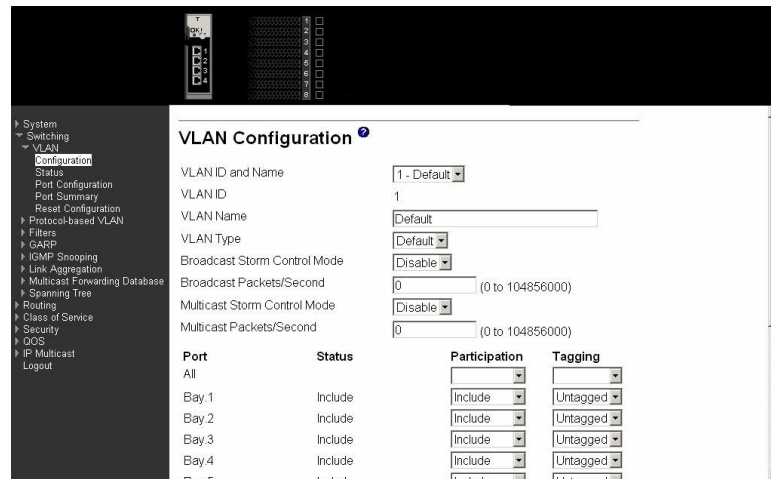
VLAN

This menu provides access to Virtual Local Area Network (VLAN) configuration, displays status and displays summary information. Menu options are:

- Configuration
- Status
- Port configuration
- Port summary
- Reset configuration

Configuration

This panel displays detailed information, including interface information, for a specific VLAN. You also use it to create new VLANs.



VLAN ID and Name

Select the VLAN to display from the pop-down menu, or select Create to set up a new VLAN. When Create is selected the VLAN ID field changes from non-configurable to configurable.

VLAN ID

There is a VLAN Identifier (VLAN ID) associated with each VLAN. Use this field to create a new VLAN and assign it an ID. The ID is a number in the range of 2 to 4094 (ID 1 is reserved for the default VLAN).

VLAN Name

A string associated with this VLAN as a convenience. It can be up to 16 alphanumeric characters long, including blanks. The default is blank. VLAN ID 1 always has a name of Default. Use this field to change an existing Name. This field is optional.

VLAN Type

What type of VLAN this is. A VLAN can be:

- the Default VLAN (VLAN ID = 1).
- a Static VLAN, one that you create using this panel or the **config vlan create** command.
- a Dynamic VLAN, one that is created by GVRP registration.

To change a VLAN from Dynamic to Static, use this panel or the **config vlan makestatic** command.

Broadcast Storm Control Mode

Configures broadcast storm control mode on the VLAN. To Enable broadcast storm control on this VLAN, select Enable from the pull-down list. If storm control is Enabled, storms are controlled by counting the number of broadcast packets within a certain time

period. If a count limit is exceeded, the packets are discarded. Only 64 combined broadcast and multicast storm rules are allowed to be configured at one time.

Broadcast Packets/Second

The rate at which the broadcast packets will begin being discarded. The valid range is 0 to 104856000 packets per second.

Multicast Storm Control Mode

Configures multicast storm control on the VLAN. To Enable multicast storm control on this VLAN, select Enable from the pull-down list. This command Enables or Disables multicast storm control for a particular VLAN. If storm control is Enabled, storms are controlled by counting the number of multicast packets within a certain time period. If a count limit is exceeded, the packets are discarded. Only 64 combined broadcast and multicast storm rules are allowed to be configured at one time.

Multicast Packets/Second

The rate level at which the multicast packets will begin being discarded. The valid range is 0 to 104856000 packets per second.

Port

Indicates which port is associated with the fields on this line.

Status

Displays the current degree of participation of this port in this VLAN. The permissible values are:

Include This port is always a member of this VLAN. This is equivalent to registration fixed in the IEEE 802.1Q standard.

Exclude This port is never a member of this VLAN. This is equivalent to registration forbidden in the IEEE 802.1Q standard.

Autodetect This port will not participate in this VLAN unless a GVRP join request is received on this port. This is equivalent to registration normal in the IEEE 802.1Q standard.

Participation

Use the pull-down menu to configure the degree of participation of this port in this VLAN. The permissible values are:

Include This port is always a member of this VLAN. This is equivalent to registration fixed in the IEEE 802.1Q standard.

Exclude This port is never a member of this VLAN. This is equivalent to registration forbidden in the IEEE 802.1Q standard.

Autodetect This port will not participate in this VLAN unless a GVRP join request is received on this port. This is equivalent to registration normal in the IEEE 802.1Q standard.

Tagging

Use the pull-down menu to configure the tagging behavior of this port in this VLAN. The default is untagged.

Tagged All frames transmitted for this VLAN will be tagged.

Untagged All frames transmitted for this VLAN will be untagged.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle you must perform a save.

Status

This panel displays information about all configured VLANs.

VLAN ID	VLAN Name	VLAN Type	Broadcast Storm Control Mode	Broadcast Packets/Second	Multicast Storm Control Mode	Multicast Packets/Second
1	Default	Default	Disable		Disable	
10		Static	Disable		Disable	
20		Static	Disable		Disable	
30		Static	Disable		Disable	
40		Static	Disable		Disable	
50		Static	Disable		Disable	
60		Static	Disable		Disable	

VLAN ID There is a VLAN Identifier (VLAN ID) associated with each VLAN. The range of the VLAN ID is 1 to 4094.

VLAN Name A string associated with this VLAN as a convenience. It can be up to 16 alphanumeric characters long, including blanks. The default is blank. VLAN ID 1 always has a name of Default. This field is optional.

VLAN Type What type of VLAN this is. A VLAN can be:

- the Default VLAN (VLAN ID = 1).
- a static VLAN, one that you have created.
- a Dynamic VLAN, one that is created by GVRP registration.

To change a VLAN from Dynamic to Static, use the VLAN Configuration panel or the **config vlan makestatic** command.

Broadcast Storm Control Mode

This field shows the mode of broadcast storm control on the VLAN. If storm control is Enabled, storms are controlled by counting the number of broadcast packets within a certain time period. If a count limit is exceeded, the packets are discarded.

Broadcast Packets/Second

The rate level at which the broadcast packets will begin being discarded.

Multicast Storm Control Mode

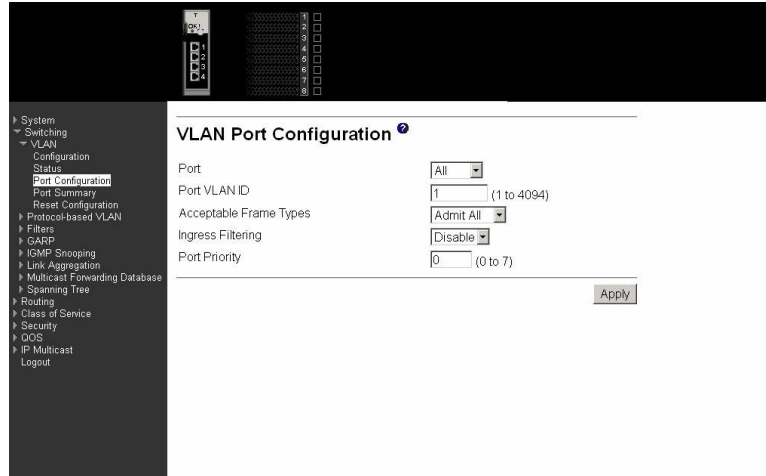
This field shows the mode of multicast storm control on the VLAN. If storm control is Enabled, storms are controlled by counting the number of multicast packets within a certain time period. If a count limit is exceeded, the packets are discarded.

Multicast Packets/Second

The rate level at which the multicast packets will begin being discarded.

Port configuration

Use this panel to configure the VLAN behavior for a specific interface in a VLAN.



Port Select the port you want to configure from the pull-down menu.

Port VLAN ID Specify the VLAN ID you want assigned to untagged or priority tagged frames received on this port. The VLAN ID must be that of a VLAN you have already created. The factory default is 1.

Acceptable Frame Types

Specify how you want the port to handle untagged and priority tagged frames. If you select VLAN only, the port will discard any untagged or priority tagged frames it receives. If you select Admit All, untagged and priority tagged frames received on the port will be accepted and assigned the value of the Port VLAN ID for this port. Whichever you select, VLAN tagged frames will be forwarded in accordance with the IEEE 802.1Q VLAN standard. The factory default is Admit All.

Ingress Filtering

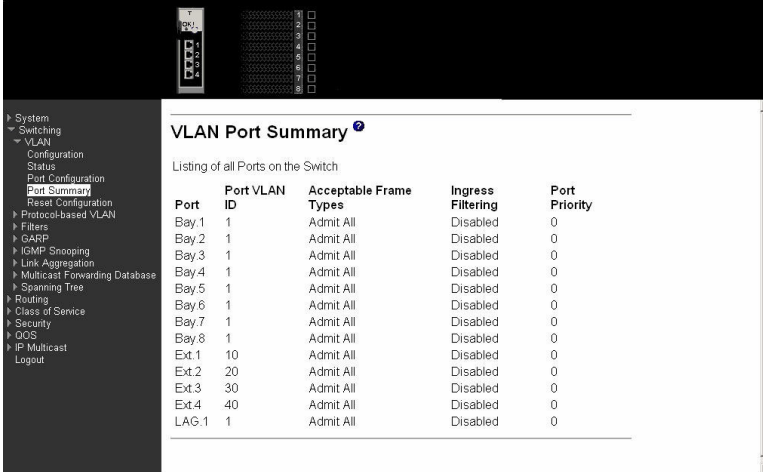
Specify how you want the port to handle tagged frames. If you Enable Ingress Filtering on the pull-down menu, a tagged frame will be discarded if this port is not a member of the VLAN identified by the VLAN ID in the tag. If you select Disable from the pull-down menu, all tagged frames will be accepted. The factory default is Disable.

Port Priority Specify the default 802.1p priority for the port.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle you must perform a save.

Port summary

This panel displays VLAN information for all ports on the switch.



Port	Port VLAN ID	Acceptable Frame Types	Ingress Filtering	Port Priority
Bay.1	1	Admit All	Disabled	0
Bay.2	1	Admit All	Disabled	0
Bay.3	1	Admit All	Disabled	0
Bay.4	1	Admit All	Disabled	0
Bay.5	1	Admit All	Disabled	0
Bay.6	1	Admit All	Disabled	0
Bay.7	1	Admit All	Disabled	0
Bay.8	1	Admit All	Disabled	0
Ext.1	10	Admit All	Disabled	0
Ext.2	20	Admit All	Disabled	0
Ext.3	30	Admit All	Disabled	0
Ext.4	40	Admit All	Disabled	0
LAG.1	1	Admit All	Disabled	0

Port Indicates which port is associated with the fields on this line.

Port VLAN ID The VLAN ID that this port will assign to untagged frames or priority tagged frames received on this port if the acceptable frame types parameter is set to Admit All. The factory default is 1.

Acceptable Frame Types

The types of frames that can be received on this port. The options are VLAN Only and Admit All. When set to VLAN Only, untagged frames or priority tagged frames received on this port are discarded. When set to Admit All, untagged frames or priority tagged frames received on this port are accepted and assigned the value of the Port VLAN ID for this port. With either option, VLAN tagged frames are forwarded in accordance with the IEEE 802.1Q VLAN specification.

Ingress Filtering

Specifies whether Ingress Filtering is Enabled or Disabled on this port. When Enabled, a frame is discarded if this port is not a member of the VLAN with which the frame is associated. In a tagged frame, the VLAN is identified by the VLAN ID in the tag. In an untagged frame, the VLAN is the Port VLAN ID specified for the port that received this frame. When Disabled, all frames are accepted and forwarded in accordance with the 802.1Q VLAN bridge specification. The factory default is Disabled.

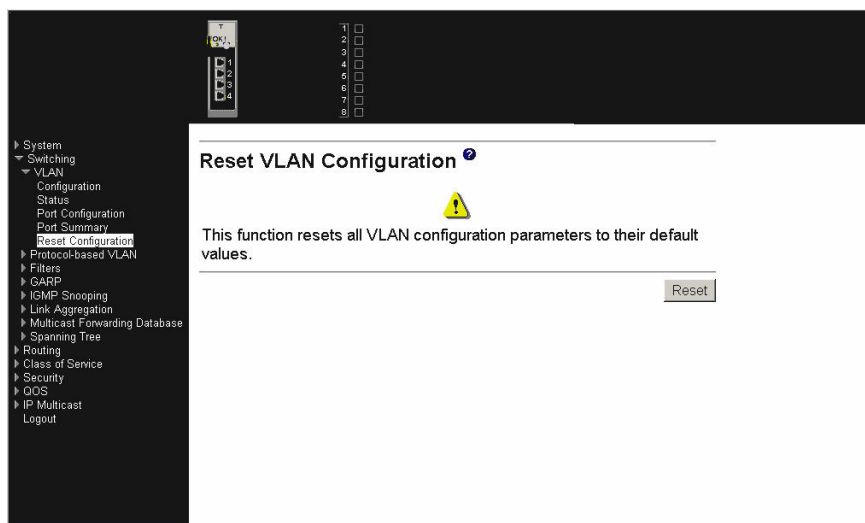
Port Priority The VLAN Port Priority that this port will assign to untagged frames received on this port.

Reset configuration

All VLAN configuration parameters are reset to their factory default values if you click the Reset button and confirm your selection on the next screen. Also, all VLANs, except for the default VLAN, will be deleted. The factory default values are:

- All ports are assigned to the default VLAN of 1.
- All ports are configured with a PVID of 1.
- All ports are configured to an Acceptable Frame Types value of Admit All Frames.
- All ports are configured with Ingress Filtering disabled.
- All ports are configured to transmit only untagged frames.

- GVRP is disabled on all ports and all dynamic entries are cleared.
- GVRP is disabled for the switch and all dynamic entries are cleared.



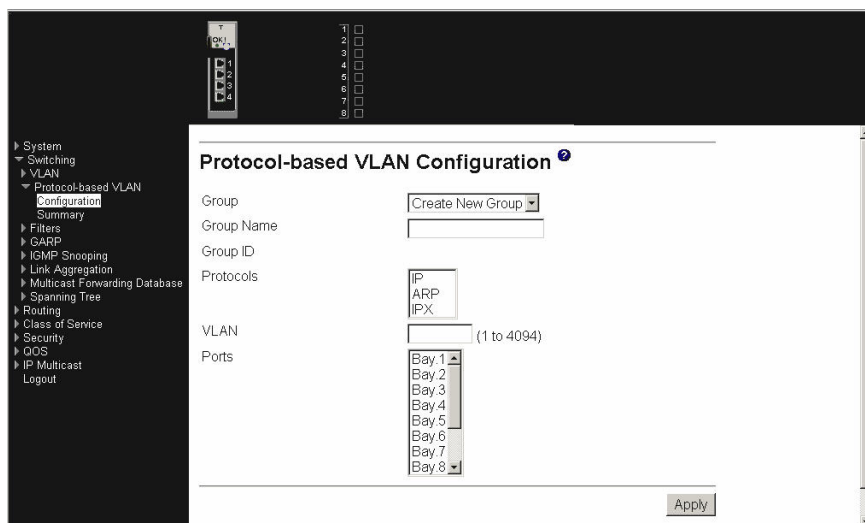
Protocol-based VLAN

This menu provides access to two protocol-based VLAN screens:

- Configuration
- Summary

Configuration

Use this panel to add a protocol-based VLAN group to the switch module, or reconfigure or delete an existing group. When a new group is created, it will be assigned a Group ID that will be used to identify it in subsequent processing.



Group Use this pull-down menu to select one of the existing PBVLANS, or select Create to add a new one. A Group ID number will be assigned automatically when you create a new group. You can create up to 128 groups.

Group Name Use this field to assign a name to a new group. You can enter up to 16 characters.

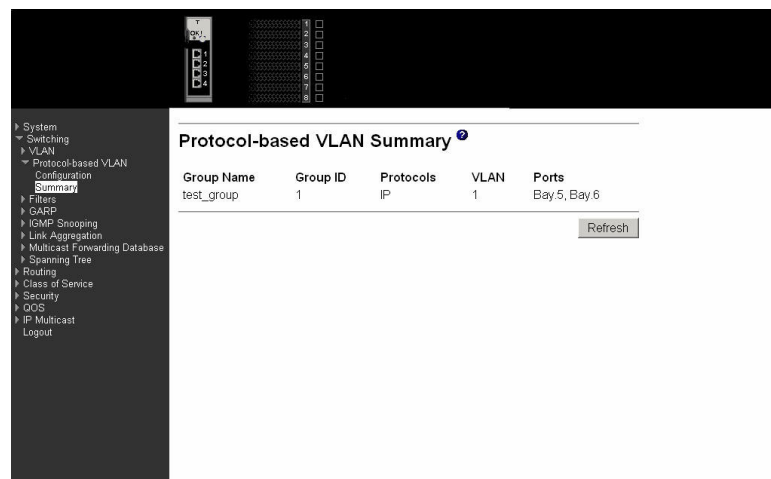
- Group ID** A number used to identify the group. A Group ID is automatically assigned when you create a group.
- Protocols** Select the protocols you want to be associated with the group. There are three configurable protocols: IP, ARP, IPX. Hold down the control key to select more than one protocol.
- IP** IP is a network layer protocol that provides a connectionless service for the delivery of data.
 - ARP** ARP is a low-level protocol that dynamically maps network layer addresses to physical medium access control (MAC) addresses.
 - IPX** The Internetwork Packet Exchange (IPX) is a connectionless datagram network-layer protocol that forwards data over a network.
- VLAN** VLAN can be any number in the range of 2 to 4094. All the ports in the group will assign this VLAN ID to untagged packets received for the protocols you included in this group.
- Ports** Select the interface(s) you want to be included in the group. Note that an interface can belong only to one group for a protocol. If you have already added interface 0.1 to a group for IP, you cannot add it to another group that also includes IP, although you could add it to a new group for IPX.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Delete Group button to remove the protocol-based VLAN group identified by the value in the Group ID field. Again, if you want the switch to retain the deletion across a power cycle, you must perform a save.

Summary

This panel displays the protocol-based VLAN information for all groups.



- Group Name** The name associated with the group. Group names can be up to 16 characters long. The maximum number of groups allowed is 128.
- Group ID** The number used to identify the group. It was automatically assigned when you created the group.

Protocols The protocols that belong to the group. There are three configurable protocols: IP, IPX, ARP.

IP IP is a network layer protocol that provides a connectionless service for the delivery of data.

ARP ARP is a low-level protocol that dynamically maps network layer addresses to physical medium access control (MAC) addresses.

IPX The Internetwork Packet Exchange (IPX) is a connectionless datagram Network-layer protocol that forwards data over a network.

VLAN This field indicates the VLAN associated with this protocol group. All ports in the group will assign this VLAN ID to untagged packets received for the protocols identified for the group.

Ports This field lists the port interface(s) that are associated with this protocol group. Note that an interface can belong only to one group for a specific protocol.

Click the Refresh button to update the screen with the latest information.

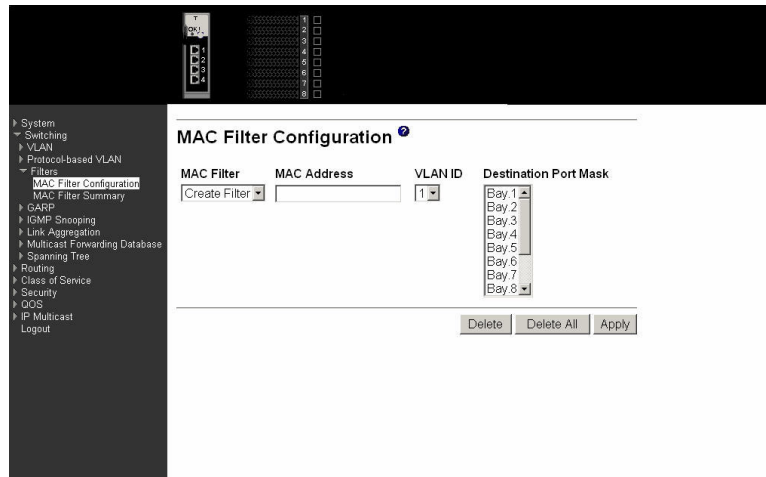
Filters

This menu provides access to two MAC filter screens:

- MAC filter configuration
- MAC filter summary

MAC filter configuration

Use this panel to add a static MAC filter entry for a MAC address and VLAN pair, update existing filter information, or delete one or more configured filters.



MAC Filter This is the list of MAC address and VLAN ID pairings for all configured filters. To change the port mask(s) for an existing filter, select the entry you want to change. To add a new filter, select Create Filter from the top of the list. Up to 48 static MAC filters can be created.

MAC Address The MAC address of the filter in the format 00-01-1A-B2-53-4D. You can change this field only when you have selected the Create Filter option. You cannot define filters for these MAC addresses:

- 00-00-00-00-00-00
- 01-80-C2-00-00-00 to 01-80-C2-00-00-0F
- 01-80-C2-00-00-20 to 01-80-C2-00-00-21
- FF-FF-FF-FF-FF-FF

VLAN ID The VLAN ID used with the MAC address to fully identify packets you want filtered. You can change this field only when you have selected the Create Filter option and you can select only a configured VLAN.

Destination Port Mask

Select the ports you want included in the filter from the pull-down menu. Packets with the MAC address and VLAN ID you selected will only be transmitted out of ports that are in the list.

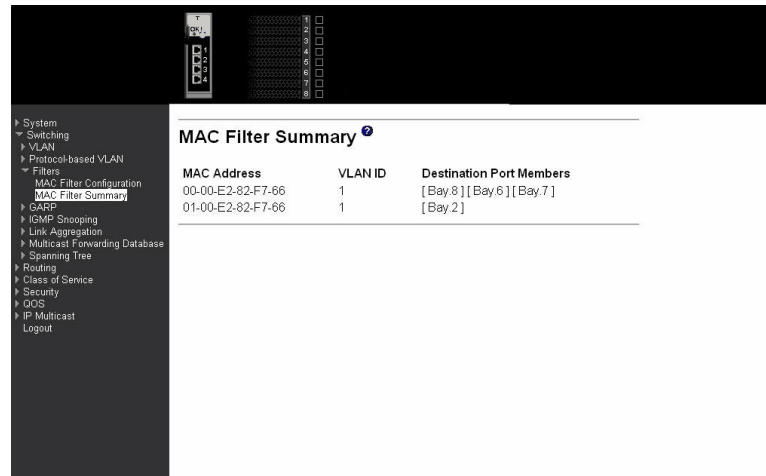
Click the Delete button to remove the currently selected filter.

Click the Delete All button to remove all configured filters.

Click the Apply button to update the switch with the values on the screen. If you want the switch to retain the new values across a power cycle you must perform a save.

MAC filter summary

This panel displays the Static MAC filtering information.



MAC Address The MAC address of the filter in the format 00-01-1A-B2-53-4D.

VLAN ID The VLAN ID associated with the filter.

Destination Port Members

A list of the ports to which packets with the MAC address and VLAN ID can be forwarded.

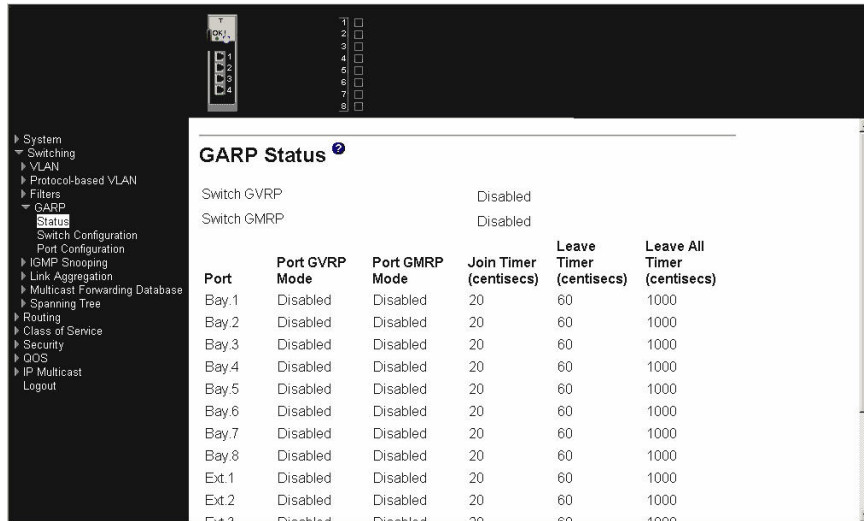
GARP

This menu provides access to the Generic Attribute Registration Protocol (GARP) summary and configuration panels. Menu options are:

- Status
- Switch configuration
- Port configuration

Status

This screen shows the GARP Status for the switch and for the individual ports. Note that the timers are only relevant when the status for a port shows as Enabled.



The screenshot shows the 'GARP Status' configuration page. The navigation tree on the left includes: System, Switching, VLAN, Protocol-based VLAN, Filters, GARP, Status, Switch Configuration, Port Configuration, IGMP Snooping, Link Aggregation, Multicast Forwarding Database, Spanning Tree, Routing, Class of Service, Security, QoS, IP Multicast, and Logout. The main content area displays the following information:

Switch GVRP: Disabled
Switch GMRP: Disabled

Port	Port GVRP Mode	Port GMRP Mode	Join Timer (centiseconds)	Leave Timer (centiseconds)	Leave All Timer (centiseconds)
Bay.1	Disabled	Disabled	20	60	1000
Bay.2	Disabled	Disabled	20	60	1000
Bay.3	Disabled	Disabled	20	60	1000
Bay.4	Disabled	Disabled	20	60	1000
Bay.5	Disabled	Disabled	20	60	1000
Bay.6	Disabled	Disabled	20	60	1000
Bay.7	Disabled	Disabled	20	60	1000
Bay.8	Disabled	Disabled	20	60	1000
Ext.1	Disabled	Disabled	20	60	1000
Ext.2	Disabled	Disabled	20	60	1000
Ext.3	Disabled	Disabled	20	60	1000

Switch GVRP

Indicates whether the GVRP administrative mode for this switch is Enabled or Disabled. The factory default is Disabled.

Switch GMRP

Indicates whether the GMRP administrative mode for this switch is Enabled or Disabled. The factory default is Disabled.

Port Indicates which port is associated with the fields on this line.

Port GVRP Mode

Indicates whether the GVRP administrative mode for the port is Enabled or Disabled. The factory default is Disabled.

Port GMRP Mode

Indicates whether the GMRP administrative mode for the port is Enabled or Disabled. The factory default is Disabled.

Join Timer (centiseconds)

Specifies the time between the transmission of GARP PDUs registering (or re-registering) membership for a VLAN or multicast group in centiseconds. An instance of this timer exists for each GARP participant for each port. Permissible values are 10 to 100 centiseconds (0.1 to 1.0 seconds). The factory default is 20 centiseconds (0.2 seconds).

Leave Timer (centiseconds)

Specifies the time to wait after receiving an unregister request for a VLAN or multicast group before deleting the associated entry, in centiseconds. This allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service. An instance of this timer exists for each GARP participant for each port. Permissible values are 20 to 600 centiseconds (0.2 to 6.0 seconds). The factory default is 60 centiseconds (0.6 seconds).

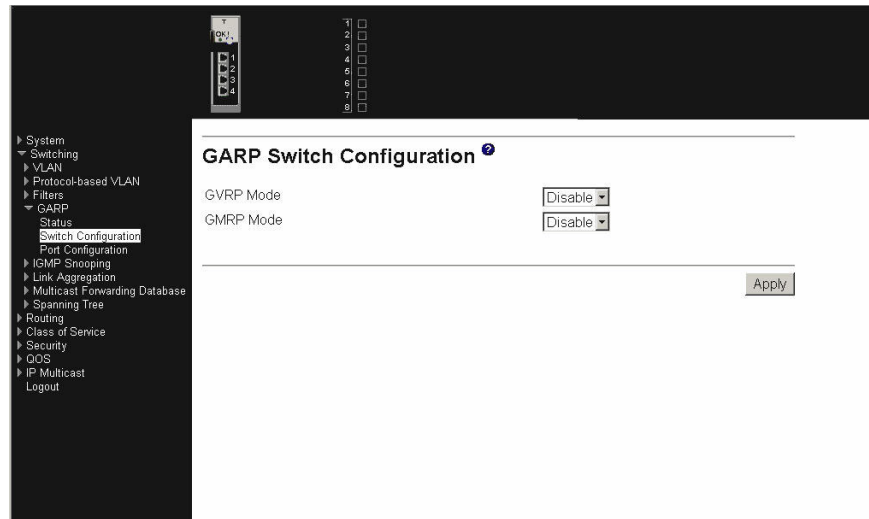
Leave All Timer (centiseconds)

This Leave All Time controls how frequently LeaveAll PDUs are generated. A LeaveAll PDU indicates that all registrations will shortly be deregistered. Participants will need to rejoin in order to maintain registration. An instance

of this timer exists for each GARP participant for each port. The Leave All Period Timer is set to a random value in the range of LeaveAllTime to $1.5 * \text{LeaveAllTime}$. Permissible values are 200 to 6000 centiseconds (2 to 60 seconds). The factory default is 1000 centiseconds (10 seconds).

Switch configuration

Use this panel to Enable or Disable GVRP and GMRP for this switch. Note: It can take up to 10 seconds for GARP configuration changes to take effect.



GVRP Mode

Choose the GVRP administrative mode for the switch by selecting Enable or Disable from the pull-down menu. The factory default is Disable.

GMRP Mode

Choose the GMRP administrative mode for the switch by selecting Enable or Disable from the pull-down menu. The factory default is Disable.

Click the Apply button to update the switch with the values you enter. If you want the switch to retain the new values across a power cycle you must perform a save.

Port configuration

Use this panel to specify GARP detail for one or all ports. Note: It can take up to 10 seconds for GARP configuration changes to take effect.

The screenshot shows the 'GARP Port Configuration' web interface. On the left is a navigation menu with the following items: System, Switching, VLAN, Protocol-based VLAN, Filters, GARP, Status, Switch Configuration, Port Configuration, IGMP Snooping, Link Aggregation, Multicast Forwarding Database, Spanning Tree, Routing, Class of Service, Security, QoS, IP Multicast, and Logout. The main content area is titled 'GARP Port Configuration' and contains the following fields:

- Port: All (dropdown menu)
- Port GVRP Mode: Disable (dropdown menu)
- Port GMRP Mode: Disable (dropdown menu)
- GARP Timers**
 - Join Timer (centiseocs): 20 (range: 10 to 100)
 - Leave Timer (centiseocs): 60 (range: 20 to 600)
 - Leave All Timer (centiseocs): 1000 (range: 200 to 6000)

An 'Apply' button is located at the bottom right of the configuration area.

Port Select the port you want to configure from the pull-down list, or select all ports.

Port GVRP Mode

Specify the GVRP administrative mode for the port by selecting Enable or Disable from the pull-down menu. If you select Disable, the protocol will not be active and the Join Time, Leave Time and Leave All Time will have no effect. The factory default is Disable.

Port GMRP Mode

Specify the GMRP administrative mode for the port by selecting Enable or Disable from the pull-down menu. If you select Disable, the protocol will not be active, and Join Time, Leave Time and Leave All Time have no effect. The factory default is Disable.

Join Timer (centiseocs)

Specify the time between the transmission of GARP PDUs registering (or re-registering) membership for a VLAN or multicast group in centiseocs. Enter a number between 10 and 100 (0.1 to 1.0 seconds). The factory default is 20 centiseocs (0.2 seconds). An instance of this timer exists for each GARP participant for each port.

Leave Timer (centiseocs)

Specify the time to wait after receiving an unregister request for a VLAN or multicast group before deleting the associated entry, in centiseocs. This allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service. Enter a number between 20 and 600 (0.2 to 6.0 seconds). The factory default is 60 centiseocs (0.6 seconds). An instance of this timer exists for each GARP participant for each port.

Leave All Timer (centiseocs)

The Leave All Timer controls how frequently LeaveAll PDUs are generated. A LeaveAll PDU indicates that all registrations will shortly be deregistered. Participants will need to rejoin in order to maintain registration. The Leave All Period Timer is set to a random value in the range of LeaveAllTime to 1.5*LeaveAllTime. The timer is specified in centiseocs. Enter a number

between 200 and 6000 (2 to 60 seconds). The factory default is 1000 centiseconds (10 seconds). An instance of this timer exists for each GARP participant for each port.

Click the Apply button to update the switch with the values you enter. If you want the switch to retain the new values across a power cycle you must perform a save.

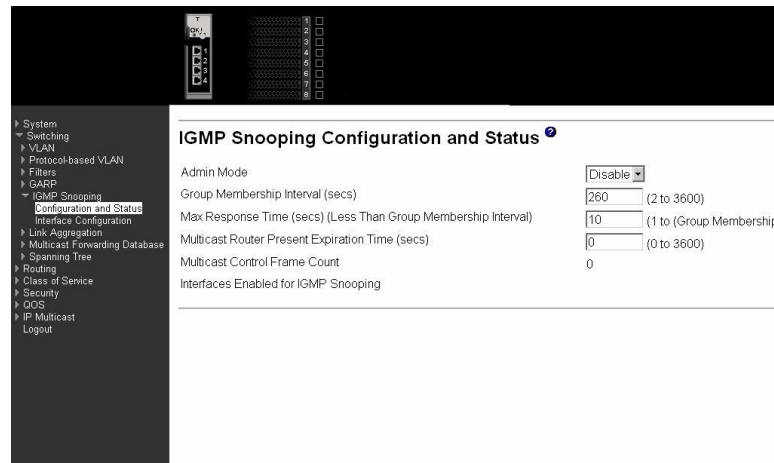
IGMP snooping

This menu provides access to the Internet Group Management Protocol (IGMP) snooping configuration and status screens. Menu options are:

- Configuration and status
- Interface configuration

Configuration and status

Use this menu to configure the parameters for IGMP snooping, which is used to build forwarding lists for multicast traffic.



Admin Mode

Select the administrative mode for IGMP snooping for the switch from the pull-down menu. The default is Disable.

Group Membership Interval (secs)

Specify the amount of time you want the switch to wait for a report for a particular group on a particular interface before it deletes that interface from the group. Enter a value between 1 and 3600 seconds. The default is 260 seconds.

Max Response Time (secs) (Less Than Group Membership Interval)

Specify the amount of time you want the switch to wait after sending a query on an interface because it did not receive a report for a particular group on that interface. Enter a value between 1 and 3600 seconds. The default is 10 seconds. The configured value must be less than the Group Membership Interval.

Multicast Router Present Expiration Time (secs)

Specify the amount of time you want the switch to wait to receive a query on an interface before removing it from the list of interfaces with multicast routers attached. Enter a value between 0 and 3600 seconds. The default is 0 seconds. A value of zero indicates an infinite timeout, i.e. no expiration.

Multicast Control Frame Count

The number of multicast control frames that are processed by the CPU.

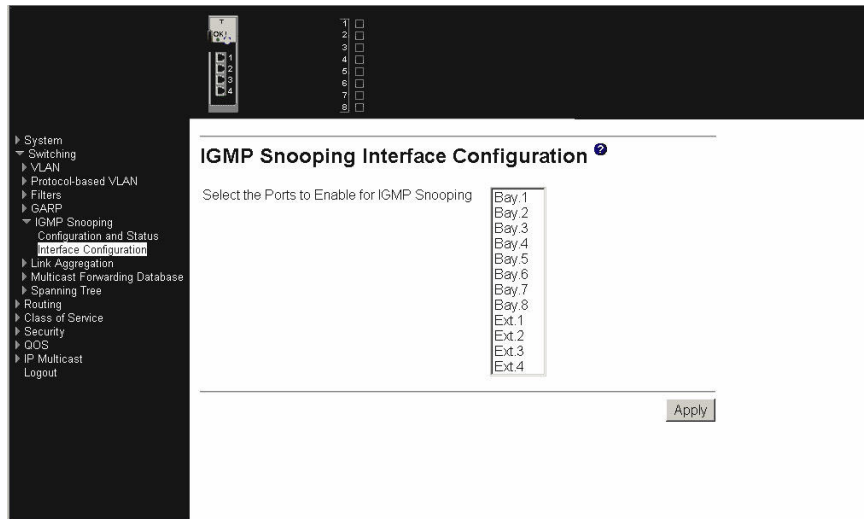
Interfaces Enabled for IGMP Snooping

A list of all the interfaces currently enabled for IGMP snooping.

Click the Apply button to update the switch with the values you enter. If you want the switch to retain the new values across a power cycle you must perform a save.

Interface configuration

Use this panel to specify on which ports to enable IGMP snooping.



Select the Ports to Enable for IGMP Snooping

The multiple select box lists all physical and LAG interfaces. Those interfaces currently enabled for IGMP snooping are shown as selected. Select all the interfaces you want enabled and deselect all those you want Disabled.

Click the Apply button to update the switch with the values you enter. If you want the switch to retain the new values across a power cycle you must perform a save.

LAG

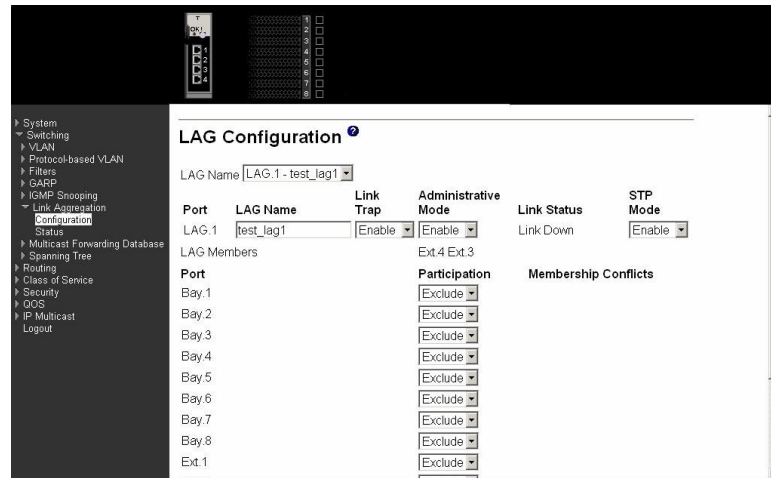
This menu provides access to the Link Aggregation (LAG) configuration and status screens. Menu options are:

- Configuration
- Status

Configuration

Use this panel to configure a new LAG, assign a name to it and generate a logical port number for it. The logical port number will be displayed after the LAG has been

created.



LAG Name (Create)

Use this pull-down menu to select one of the existing LAGs, or select Create to add a new one. There can be a maximum of 6 LAGs. This is an alphanumeric string up to 15 characters in length.

Port Displays the logical port number associated with this LAG Name.

LAG Name

Enter a name for the LAG you are creating. Name is an alphanumeric string of up to 15 characters. You can also use this field to modify the name that was associated with a LAG when it was created.

Link Trap

Enables or Disables link trap notifications for the specified LAG.

Administrative Mode

This field Enables or Disables the specified LAG(s).

Link Status

Indicates whether the Link is Up or Down.

STP Mode

Sets the STP mode for the specified LAG(s).

Port Identifies a physical port. To add the port to the LAG select Include from the Participation column. There can be a maximum of 8 member ports in a LAG.

Participation

For each port specify whether it is to be included as a member of this LAG or not. The default is exclude. There can be a maximum of 8 ports assigned to a LAG.

Membership Conflicts

Shows ports that are already members of other LAGs. A port can be a member of only one LAG at a time. If the entry is blank, it is not currently a member of any LAG.

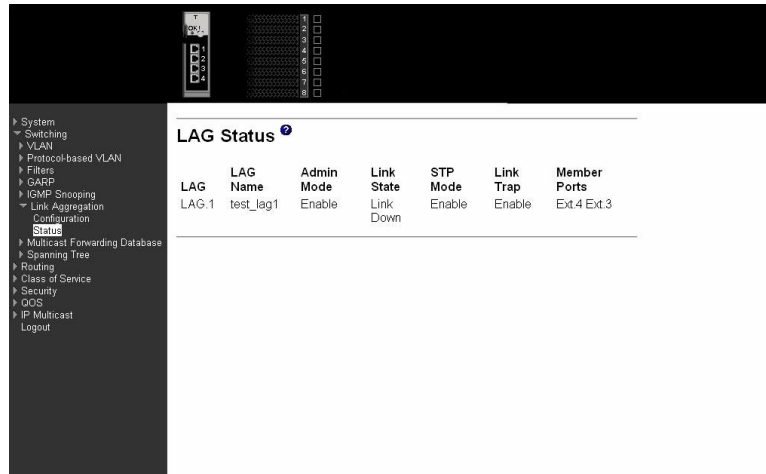
Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

Click the Apply button to update the switch with the values you enter. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Delete button to remove the currently selected LAG. All ports that were members of this LAG are removed from the LAG and included in the default VLAN. This field will not appear when a new LAG is being created.

Status

This panel displays an overview of all LAGs on the switch.



LAG The logical port identifier of the LAG, in the format lag.port.

LAG Name The name of this LAG.

Admin Mode The administrative mode. The factory default is Enabled.

Link State Indicates whether the link is Up or Down.

STP Mode The Spanning Tree Protocol Administrative Mode associated with the LAG. The possible values are:

Disable

Spanning tree is Disabled for this LAG.

Enable

Spanning tree is Enabled for this LAG.

Link Trap Indicates whether or not a trap will be sent when link status changes. The factory default is Enabled.

Member Ports

A listing of the ports that are members of this LAG, in port notation. There can be a maximum of 8 ports assigned to a LAG.

MFDB

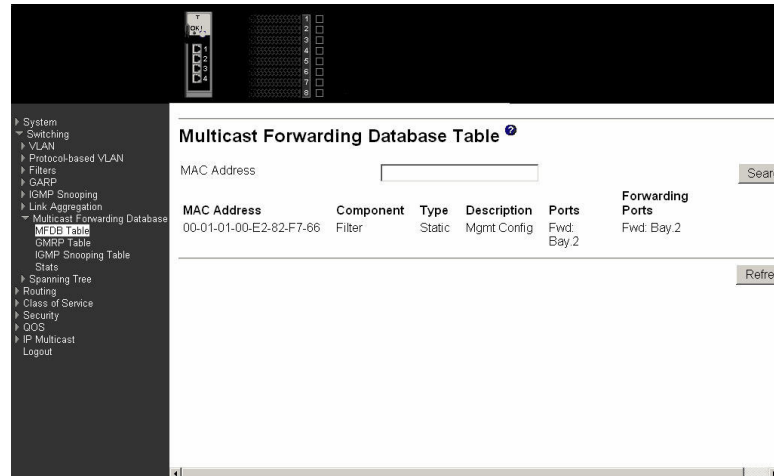
The Multicast Forwarding Database (MFDB) holds the port membership information for all active multicast address entries. The key for an entry consists of a VLAN ID and MAC address pair. Entries can contain data for more than one protocol.

Options on this menu are:

- MFDB table
- GMRP table
- IGMP snooping table
- Stats

MFDB table

Use this panel to display entries from the MFDB.



MAC Address Enter a MAC address and VLAN pair for which the switch has forwarding and/or filtering information. The format is two two-digit hexadecimal numbers representing the VLAN and six two-digit hexadecimal numbers representing the MAC address, separated by hyphens; for example, 00-01-00-23-45-67-89-AB.

After you have entered a MAC address click the Search button and the data associated with the address will be displayed. Otherwise, all entries will be displayed.

Component The component that is responsible for this entry in the Multicast Forwarding Database. Possible values are IGMP Snooping, GMRP, and Static Filtering.

Type This displays the type of the entry. Static entries are those that are configured by the end user. Dynamic entries are added to the table as a result of a learning process or protocol.

Description The text description of this multicast table entry.

Ports The list of interfaces that are designated for forwarding (Fwd:) and filtering (Flt:).

Forwarding Ports

The forwarding list is derived from combining all the component's forwarding interfaces and removing the interfaces that are listed as the static filtering interfaces.

Click the Refresh button to update the screen with the latest information.

GMRP table

This panel displays the GMRP entries in the MFDB table.



MAC Address A MAC address and VLAN pair for which the switch has forwarding and/or filtering information. The format is two two-digit hexadecimal numbers representing the VLAN and six two-digit hexadecimal numbers representing the MAC address, separated by hyphens; for example, 00-01-00-23-45-67-89-AB.

Type Displays the type of the entry. Static entries are those that are configured by the end user. Dynamic entries are added to the table as a result of a learning process or protocol.

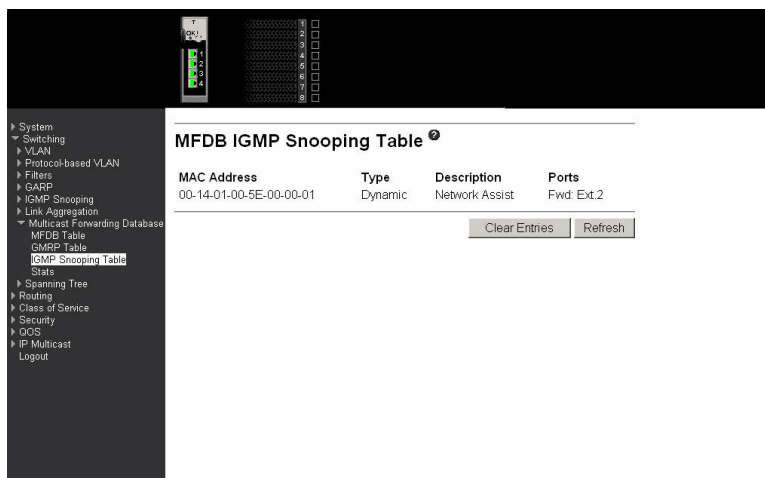
Description The text description of this multicast table entry. Possible values are Management Configured, Network Configured and Network Assisted.

Ports The list of interfaces that are designated for forwarding (Fwd:) and filtering (Flt:).

Click the Refresh button to update the screen with the latest information.

IGMP snooping table

This panel displays the IGMP snooping entries in the MFDB.



MAC Address A MAC address and VLAN pair for which the switch has forwarding

and/or filtering information. The format is two two-digit hexadecimal numbers representing the VLAN and six two-digit hexadecimal numbers representing the MAC address, separated by hyphens; for example, 00-01-00-23-45-67-89-AB.

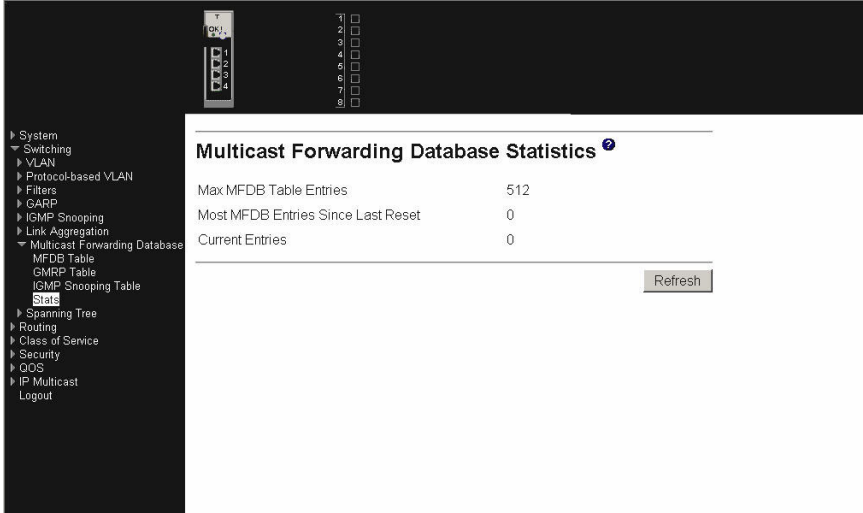
- Type** Displays the type of the entry. Static entries are those that are configured by the end user. Dynamic entries are added to the table as a result of a learning process or protocol.
- Description** The text description of this multicast table entry. Possible values are Management Configured, Network Configured and Network Assisted.
- Ports** The list of interfaces that are designated for forwarding (Fwd:) and filtering (Flt:).

Click the Clear Entries button to tell the IGMP Snooping component to delete all of its entries from the multicast forwarding database.

Click the Refresh button to update the screen with the latest information.

Stats

This panel displays the MFDB statistics.



The screenshot shows a web-based network management interface. On the left is a dark sidebar with a navigation menu. The main content area is titled "Multicast Forwarding Database Statistics" and contains a table with three rows of statistics. A "Refresh" button is located at the bottom right of the statistics table.

Multicast Forwarding Database Statistics	
Max MFDB Table Entries	512
Most MFDB Entries Since Last Reset	0
Current Entries	0

Max MFDB Table Entries

Displays the total number of entries possible in the MFDB table.

Most MFDB Entries Since Last Reset

Displays the largest number of entries that have been present in the MFDB table since last reset. This value is also known as the MFDB high-water mark.

Current Entries

Displays the current number of entries in the MFDB table.

Click the Refresh button to update the screen with the latest information.

Spanning tree

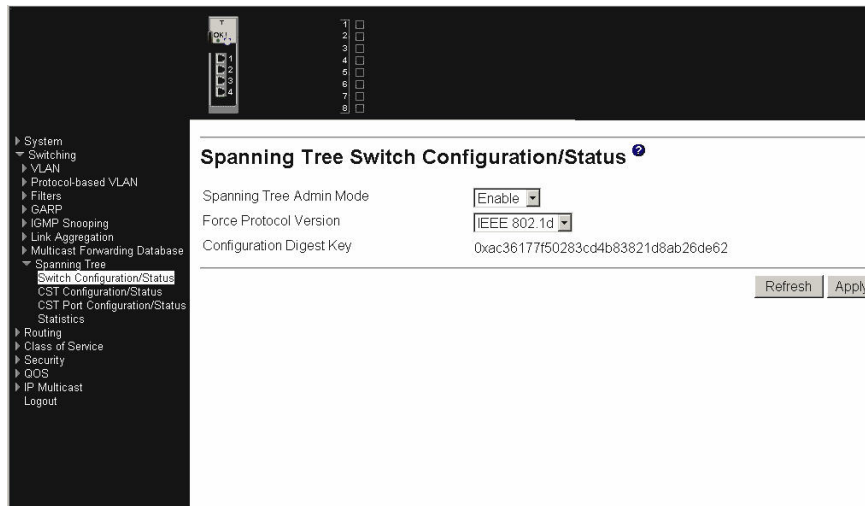
This menu provides access to spanning tree-related configuration and status screens. Menu options are:

- Switch configuration/status

- CST configuration/status
- CST port configuration/status
- Statistics

Switch configuration/status

Use this panel to configure the spanning tree parameters for the switch.



Spanning Tree Admin Mode

Select Enable or Disable from the pull-down menu to specify whether spanning tree operation is Enabled on the switch.

Force Protocol Version

Specify the version of the Spanning Tree Protocol (STP) you want the switch to use. The options are IEEE 802.1D (standard) and IEEE 802.1w (Rapid Reconfiguration).

Configuration Digest Key

A derived value identifying the configuration.

Click the Refresh button to update the screen with the most recent data.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Common Spanning Tree (CST) configuration/status

Use this panel to configure or display the bridge parameters for the Spanning Tree Algorithm.

Parameter	Value	Range
Bridge Priority	32768	(0 to 61440)
Bridge Max Age (secs)	20	(6 to 40)
Bridge Hello Time (secs)	2	(1 to 10)
Bridge Forward Delay (secs)	15	(4 to 30)
Bridge Identifier	80-00-00-0E-0C-00-03-40	
Time Since Topology Change	0 day 0 hr 0 min 0 sec	
Topology Change Count	0	
Topology Change	False	
Designated Root	80-00-00-0E-0C-00-03-40	
Root Path Cost	0	
Root Port	00:00	
Max Age (secs)	20	
Forward Delay (secs)	15	
Hold Time (secs)	?	

Bridge Priority

Specifies the bridge priority. The value can be between 0 and 61440. It is set in multiples of 4096. For example, if you enter any value between 0 and 4095, it will be set to 0. If you enter any value between 4096 and $(2 \times 4096 - 1)$ it will be set to 4096. The default priority is 32768.

Bridge Max Age (secs)

Specifies the bridge maximum age timeout value. The value can be between 1 and 40, and should be less than or equal to $((2 \times \text{Bridge Forward Delay}) - 1)$ and greater than or equal to $(2 \times (\text{Bridge Hello Time} + 1))$. The default value is 15.

Bridge Hello Time (secs)

Specifies the bridge hello timeout value, with the value being less than or equal to $((\text{Bridge Max Age} / 2) - 1)$. The default hello time value is 2.

Bridge Forward Delay (secs)

Specifies the time the bridge will spend in Listening and Learning mode before starting to forward packets. Bridge Forward Delay must be greater than or equal to $((\text{Bridge Max Age} / 2) + 1)$. The time range is from 4 seconds to 30 seconds and the default value is 15.

Bridge Identifier

The bridge identifier. The bridge priority is concatenated with the base MAC address of the bridge to create the identifier.

Time Since Topology Change

The time in seconds since the spanning tree topology last changed.

Topology Change Count

Number of times the spanning tree topology has changed.

Topology Change

The value of the topology change parameter for the switch

indicating if a topology change is in progress on any port on the bridge. It takes a value if True or False.

Designated Root

The bridge identifier of the root bridge.

Root Path Cost

Path Cost to the Designated Root for this bridge instance.

Root Port

Port to access the Designated Root.

Max Age (secs)

Path Cost to the Designated Root for this bridge instance.

Forward Delay (secs)

Derived value of the Root Port Bridge Forward Delay parameter.

Hold Time (secs)

Minimum time between transmission of Configuration BPDUs.

CST Regional Root

Priority and base MAC address of the Common Spanning Tree Regional Root.

CST Path Cost

Path Cost to the CST tree Regional Root.

Click the Refresh button to update the screen with the most recent data.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle you must perform a save.

CST port configuration/status

Use this panel to configure a particular port within the CST.

Spanning Tree CST Port Configuration/Status	
Port	Bay 1
Port Priority	128 (0 to 240)
Admin Edge Port	Disable
Port Path Cost	0 (0 to 200000000) 0 = Auto
Auto-calculate Port Path Cost	Enabled
Port ID	80.01
Port Up Time Since Counters Last Cleared	0 day 2 hr 10 min 4 sec
Port Mode	Disabled
Port Forwarding State	Disabled
Port Role	Disabled Port
Designated Root	80-00-00-0E-0C-00-03-40
Designated Cost	0
Designated Bridge	80-00-00-0E-0C-00-03-40

Port Select one of the physical or LAG interfaces from the pull-down menu.

Port Priority Specify the priority for the selected port. The port priority is set in multiples of 16, and the range is 0 to 240.

Admin Edge Port Select Enable to specify the port as an Edge Port within the CST. Disable is the default.

Port Path Cost

Set the Path Cost to a new value for the specified port. The range is 1 to 200000000.

Auto-calculate Port Path Cost

Displays whether the path cost is automatically calculated (Enabled) or not (Disabled). Path cost will be calculated based on the link speed of the port if the configured value for Port Path Cost is zero.

Port ID

The port identifier for the specified port. It is created by concatenating the port priority with the interface number of the port.

Port Up Time Since Counters Last Cleared

Time since the counters were last cleared, displayed in Days, Hours, Minutes, and Seconds.

Port Mode

STP Administrative Mode associated with the port or LAG. The possible values are Enable or Disable.

Port Forwarding State

The Forwarding State of this port.

Port Role

Each Enabled bridge port is assigned a Port Role within the spanning tree. The port role will be one of the following values: Root Port, Designated Port, Alternate Port, Backup Port, Master Port or Disabled Port.

Designated Root

Root Bridge for the spanning tree.

Designated Cost

Path Cost offered to the LAN by the Designated Port.

Designated Bridge

Bridge Identifier of the bridge with the Designated Port.

Designated Port

Port Identifier on the Designated Bridge that offers the lowest cost to the LAN.

Topology Change Acknowledge

Identifies whether the next BPDU to be transmitted for this port would have the topology change acknowledgement flag set. It is either True or False.

Hello Time (secs)

Configured value of the hello timer.

Edge Port

Indicates whether the port is Enabled as an edge port. It takes the value Enabled or Disabled.

Point-to-point MAC

Derived value of the point-to-point status.

CST Regional Root

Bridge Identifier of the CST Regional Root. It is made up using the bridge priority and the base MAC address of the bridge.

CST Path Cost

Path Cost to the CST Regional Root.

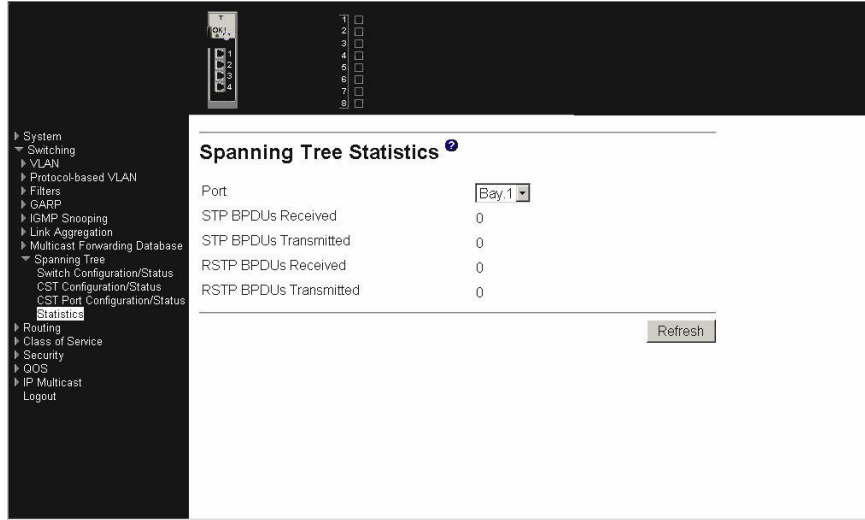
Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Force button to force the port to send out 802.1w BPDUs.

Click the Refresh button to update the screen with the most recent data.

Statistics

This panel displays BPDU statistics for the selected port.



Port Select the port for which information is to be displayed.

STP BPDUs Received

Number of STP BPDUs received at the selected port.

STP BPDUs Transmitted

Number of STP BPDUs transmitted from the selected port.

RSTP BPDUs Received

Number of Rapid Reconfiguration BPDUs received at the selected port.

RSTP BPDUs. Transmitted

Number of Rapid Reconfiguration BPDUs transmitted from the selected port.

Click the Refresh button to update the screen with the most recent data.

Routing

This menu provides access to the following routing-related menus:

- ARP
- Internet Protocol (IP)
- Open Shortest Path First (OSPF)
- BOOTP/DHCP relay
- Routing Information Protocol (RIP)
- Router discovery
- Router
- VLAN
- Virtual Router Redundancy Protocol (VRRP)

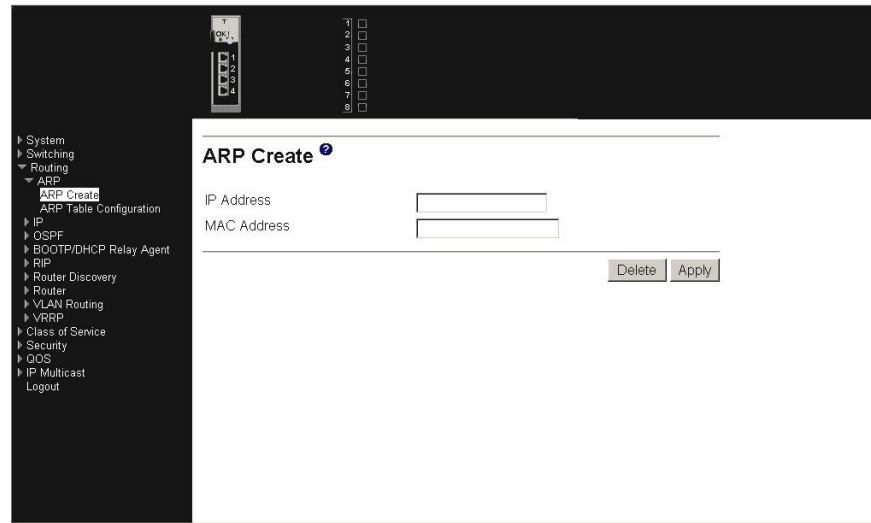
ARP

This menu provides access to the following ARP data entry panels:

- ARP create
- ARP table configuration

ARP create

Use this panel to add an entry to the ARP table.



The screenshot shows a web-based network management interface. On the left is a dark sidebar with a navigation menu. The main content area is titled "ARP Create" and contains two input fields: "IP Address" and "MAC Address". Below these fields are two buttons: "Delete" and "Apply". The sidebar menu includes items like System, Switching, Routing, ARP (expanded), ARP Table Configuration, IP, OSPF, BOOTP/DHCP Relay Agent, RIP, Router Discovery, Router, VLAN Routing, VRRP, Class of Service, Security, QoS, IP Multicast, and Logout.

IP Address Enter the IP address you want to add. It must be the IP address of a device on a subnet attached to one of the switch's existing routing interfaces.

MAC Address The unicast MAC address of the device. Enter the address as six two-digit hexadecimal numbers separated by dashes, for example 00-06-29-32-81-40.

Click the Delete button to delete the specified ARP entry. You only need to enter the IP address of the ARP Entry to delete the entry. If you want the switch to retain the deletion across a power cycle, you must perform a save.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

ARP table configuration

Use this panel to change the configuration parameters for the ARP Table. You can also use this panel to display the contents of the table.

The screenshot shows the 'ARP Table Configuration' web interface. On the left is a navigation tree with the following items: System, Switching, Routing, ARP (expanded), ARP Create, ARP Table Configuration (selected), IP, OSPF, BOOTP/DHCP Relay Agent, RIP, Router Discovery, Router, VLAN Routing, VRRP, Class of Service, Security, QoS, IP Multicast, and Logout. The main configuration area is titled 'ARP Table Configuration' and contains the following fields:

- Age Time (secs): 1200 (15 to 3600)
- Response Time (secs): 1 (1 to 10)
- Retries: 4 (1 to 10)
- Cache Size: 3072 (10 to 3072)

Below the fields is a table with the following columns: IP Address, MAC Address, Port, and Type. An 'Apply' button is located at the bottom right of the table area.

Age Time (secs)

Enter the value for the ARP entry ageout time. You must enter a valid integer, which represents the number of seconds it will take for an ARP entry to age out. The range for this field is 15 to 3600 seconds. The default value for Age Time is 1200 seconds.

Response Time (secs)

Enter the value for the ARP response timeout. You must enter a valid integer, which represents the number of seconds the switch will wait for a response to an ARP request. The range for this field is 1 to 10 seconds. The default value for Response Time is 1 second.

Retries

Enter an integer which specifies the maximum number of times an ARP request will be retried. The range for this field is 1 to 10. The default value for Retries is 4.

Cache Size

Enter an integer which specifies the maximum number of entries for the ARP cache. The range for this field is 10 to 3072. The default value for Cache Size is 3072.

IP Address

The IP address of a device on a subnet attached to one of the switch's routing interfaces.

MAC Address

The unicast MAC address for the device. The format is six two-digit hexadecimal numbers separated by dashes, for example 00-06-29-32-81-40.

Port

The routing interface associated with the ARP entry.

Type

The type of the ARP entry. Options are:

Local

An ARP entry associated with one of the switch's routing interface's MAC addresses.

Gateway

An ARP entry whose IP address is that of a router

Static

An ARP entry configured by the user.

Dynamic An ARP entry which has been learned by the router.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

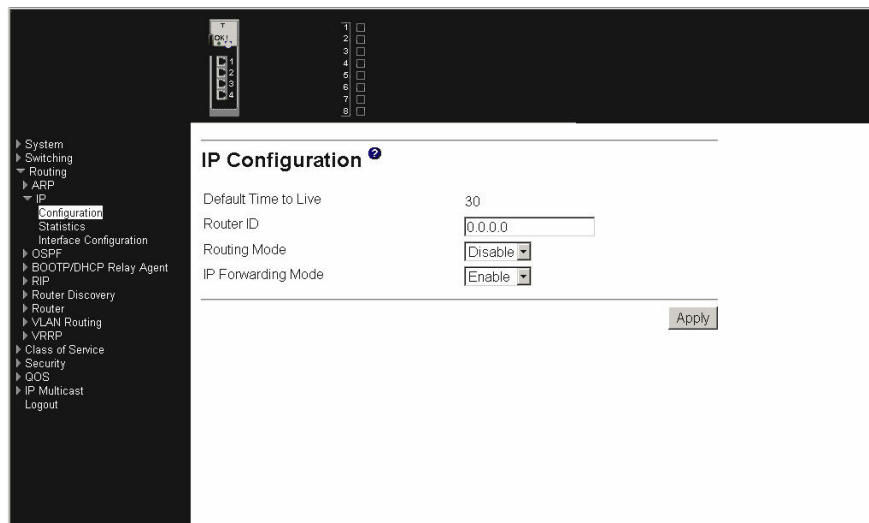
IP

This menu provides access to the following Internet Protocol (IP) entry panels:

- Configuration
- Statistics
- Interface configuration

Configuration

Use this panel to configure routing parameters for the switch as opposed to an interface.



Default Time to Live

The default value inserted into the Time-To-Live (TTL) field of the IP header of datagrams originated by the switch, if a TTL value is not supplied by the transport layer protocol.

Router ID

Enter a 32-bit integer in dotted decimal format that will uniquely identify this router within an Autonomous System (AS). Before you change the Router ID you must disable OSPF. After you enter the new Router ID and click on the Apply button, you must re-enable OSPF to have the change take effect. The default value is 0.0.0.0, although this is not a valid Router ID. When you change the Router ID to a valid address, you cannot reset it to 0.0.0.0.

Routing Mode

Select Enable or Disable from the pull-down menu. You must Enable routing for the switch before you can route through any of the interfaces. The default value is Disable.

IP Forwarding Mode

Select Enable or Disable from the pull-down menu. This Enables or Disables the forwarding of IP frames. The default value is Enable.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Statistics

This panel displays IP statistics as specified in RFC 1213.

IP Statistics	
IpInReceives	12335
IpInHdrErrors	0
IpInAddrErrors	0
IpForwDatagrams	0
IpInUnknownProtos	0
IpInDiscards	0
IpInDelivers	12335
IpOutRequests	12282
IpOutDiscards	0
IpOutNoRoutes	0
IpReasmTimeout	0
IpReasmReqds	0
IpReasmOKs	0
IpReasmFails	0

IpInReceives The total number of input datagrams received from interfaces, including those received in error.

IpInHdrErrors The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.

IpInAddrErrors The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (e.g., 0.0.0.0) and addresses of unsupported Classes (e.g., Class E). For entities which are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.

IpForwDatagrams The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP Gateways, this counter will include only those packets which were Source-Routed via this entity, and the Source-Route option processing was successful.

IpInUnknownProtos The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.

IpInDiscards The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.

- IpInDelivers** The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
- IpOutRequests** The total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in IpForwDatagrams.
- IpOutDiscards** The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
- IpOutNoRoutes** The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this 'no-route' criterion. Note that this includes any datagrams which a host cannot route because all of its default gateways are down.
- IpReasmTimeout** The maximum number of seconds received fragments are held while they are awaiting reassembly at this entity.
- IpReasmReqds** The number of IP fragments received needing reassembly at this entity.
- IpReasmOKs** The number of IP datagrams successfully re-assembled.
- IpReasmFails** The number of failures detected by the IP re-assembly algorithm (for any reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IP fragments since some algorithms can lose track of the number of fragments by combining them as they are received.
- IpFragOKs** The number of IP datagrams that have been successfully fragmented at this entity.
- IpFragFails** The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not because, for example, their Don't Fragment flag was set.
- IpFragCreates** The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
- IpRoutingDiscards** The number of routing entries which were chosen to be discarded even though they are valid. One possible reason for discarding such an entry could be to free-up buffer space for other routing entries.
- IcmpInMsgs** The total number of ICMP messages the entity received. Note that this counter includes all those counted by IcmpInErrors.
- IcmpInErrors** The number of ICMP messages the entity received but determined as having ICMP-specific errors (e.g., bad ICMP checksums, bad length, etc.).

IcmpInDestUnreachs

The number of ICMP Destination Unreachable messages received.

IcmpInTimeExcds

The number of ICMP Time Exceeded messages received.

IcmpInParmProbs

The number of ICMP Parameter Problem messages received.

IcmpInSrcQuenchs

The number of ICMP Source Quench messages received.

IcmpInRedirects

The number of ICMP Redirect messages received.

IcmpInEchos The number of ICMP Echo (request) messages received.

IcmpInEchoReps

The number of ICMP Echo Reply messages received.

IcmpInTimestamps

The number of ICMP Timestamp (request) messages received.

IcmpInTimestampReps

The number of ICMP Timestamp Reply messages received.

IcmpInAddrMasks

The number of ICMP Address Mask Request messages received.

IcmpInAddrMaskReps

The number of ICMP Address Mask Reply messages received.

IcmpOutMsgs The total number of ICMP messages this entity attempted to send. Note that this counter includes all those counted by IcmpOutErrors.

IcmpOutErrors

The number of ICMP messages this entity did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there might not be any error types that contribute to this counter's value.

IcmpOutDestUnreachs

The number of ICMP Destination Unreachable messages sent.

IcmpOutTimeExcds

The number of ICMP Time Exceeded messages sent.

IcmpOutParmProbs

The number of ICMP Parameter Problem messages sent.

IcmpOutSrcQuenchs

The number of ICMP Source Quench messages sent.

IcmpOutRedirects

The number of ICMP Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects.

IcmpOutEchos

The number of ICMP Echo (request) messages sent.

IcmpOutEchoReps

The number of ICMP Echo Reply messages sent.

IcmpOutTimestamps

The number of ICMP Timestamp (request) messages.

IcmpOutTimestampReps

The number of ICMP Timestamp Reply messages sent.

IcmpOutAddrMasks

The number of ICMP Address Mask Request messages sent.

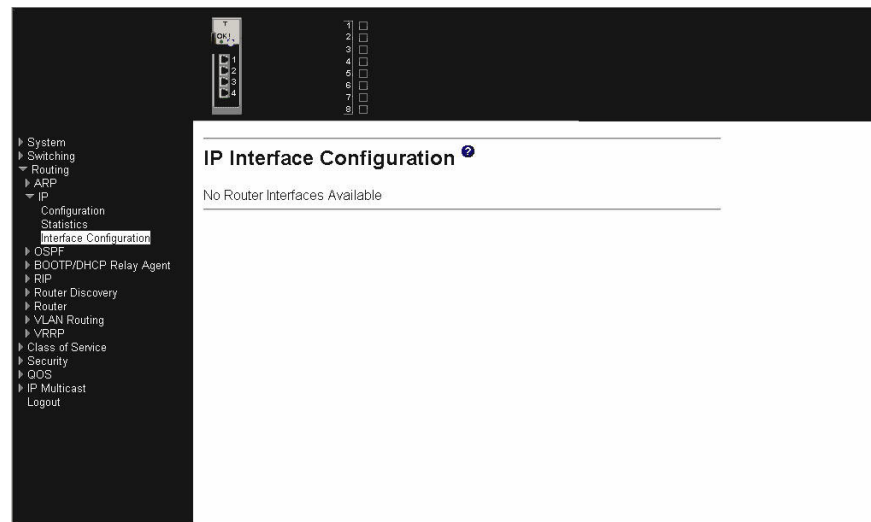
IcmpOutAddrMaskReps

The number of ICMP Address Mask Reply messages sent.

Click the Refresh button to update the data on the screen with the present state of the data in the switch.

Interface configuration

Use this panel to configure or display routing parameters for the selected interface.



Port Select the interface for which data is to be displayed or configured.

IP Address Enter the IP address for the interface.

Subnet Mask Enter the subnet mask for the interface. This is also referred to as the subnet/network mask, and defines the portion of the interface's IP address that is used to identify the attached network.

Routing Mode Setting this Enables or Disables routing for an interface. The default value is Enable.

Administrative Mode The Administrative Mode of the interface. The default value is Enable.

Forward Net Directed Broadcasts Select how network directed broadcast packets should be handled. If you select Enable from the pull-down menu network directed broadcasts will be forwarded. If you select Disable they will be dropped. The default value is Disable.

Active State The state of the specified interface is either Active or Inactive. An interface is considered active if it the link is up and it is in forwarding state.

MAC Address The burned-in physical address of the specified interface. The format is six two-digit hexadecimal numbers separated by dashes, for example 00-06-29-32-81-40.

Maximum Transmission Unit

Enter the Maximum Transmission Unit (MTU) size (in bytes) for the interface. The default value is 1500. For the standard implementation the maximum value is 1500 and the minimum value is 576 bytes.

Encapsulation Type

Select the link layer encapsulation type for packets transmitted from the specified interface from the pull-down menu. The possible values are Ethernet and SNAP. The default is Ethernet.

Click the Delete IP address button to delete the IP address from the interface. If you want the switch to retain the deletion across a power cycle, you must perform a save.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

OSPF

This menu provides access to the following Open Shortest Path First (OSPF)-related configuration and display panels:

- OSPF configuration
- Area configuration
- Stub area summary
- Area range configuration
- Interface statistics
- Interface configuration
- Neighbor table
- Neighbor status
- Link state database
- Virtual link configuration
- Virtual link summary

OSPF configuration

Use this panel to configure and display basic OSPF parameters for the switch operating as an OSPF router.

The screenshot shows the OSPF Configuration page in a web-based network management interface. On the left is a navigation tree with categories like System, Switching, Routing, ARP, IP, and OSPF. The OSPF Configuration page is active, displaying several configuration parameters:

- Router ID: 0.0.0.0
- Admin Mode: Disable (dropdown menu)
- ASBR Mode: Disable (dropdown menu)
- RFC 1583 Compatibility: Enable (dropdown menu)
- ABR Status: (checkbox)
- Exit Overflow Interval (secs): 0 (range: 0 to 2147483647)
- External LSA Count: (input field)
- External LSA Checksum: (checkbox)
- New LSAs Originated: (checkbox)
- LSAs Received: (checkbox)
- External LSDB Limit: -1 (range: -1 to 2147483647, -1 = No Limit)

An 'Apply' button is located at the bottom right of the configuration area.

Router ID The 32-bit integer in dotted decimal format that uniquely identifies the router within the Autonomous System (AS). The Router ID is set on the IP Configuration page. If you want to change the Router ID you must first disable OSPF. After you set the new Router ID, you must re-enable OSPF to have the change take effect. The default value is 0.0.0.0, although this is not a valid Router ID.

Admin Mode Select Enable or Disable from the pull-down menu. If you select Enable, OSPF will be activated for the switch. The default value is Disable. You must configure a Router ID before OSPF can become operational. You do this on the IP Configuration page or by issuing the CLI command: **config router id**. NOTE: When OSPF is initialized on the router, it will remain initialized until the router is reset.

ASBR Mode Select Enable or Disable from the pull-down menu. If you select Enable you are specifying that the router is an Autonomous System Border Router (ASBR). The default value is Disable. You will only be offered the pull-down menu if OSPF Admin Mode is Disable.

RFC 1583 Compatibility Select Enable or Disable from the pull-down menu to specify the preference rules that will be used when choosing among multiple AS-external Link State Advertisements (LSAs) advertising the same destination. If you select Enable, the preference rules will be those defined by RFC 1583. If you select Disable, the preference rules will be those defined in Section 16.4.1 of the OSPF-2 standard (RFC 2328), which will prevent routing loops when AS-external-LSAs for the same destination have been originated from different areas. The default value is Enable. To prevent routing loops, you should select Disable, but only if all OSPF routers in the routing domain are capable of operating according to RFC 2328.

ABR Status The values of this are Enabled or Disabled. Enabled implies that the router is an area border router. Disabled implies that it is not an

Area Border Router (ABR). External LSA Count - The number of external (LS type 5) LSAs in the Link State Database (LSD).

Exit Overflow Interval (secs)

Enter the number of seconds that, after entering overflow state, the router should wait before attempting to leave overflow state. This allows the router to again originate non-default AS-external-LSAs. If you enter 0, the router will not leave overflow state until restarted. The range is 0 to 2147483647 seconds.

External LSA Count

The number of external (LS type 5) LSAs in the Link-State Database (LSDB).

External LSA Checksum

The sum of the LS checksums of the external LSAs contained in the LSDB. This sum can be used to determine if there has been a change in a router's LSDB, and to compare the LSDBs of two routers.

New LSAs Originated

In any given OSPF area, a router will originate several LSAs. Each router originates a router-LSA. If the router is also the Designated Router for any of the area's networks, it will originate network-LSAs for those networks. This value represents the number of LSAs originated by this router.

LSAs Received

The number of LSAs received that were determined to be new instantiations. This number does not include newer instantiations of self-originated LSAs.

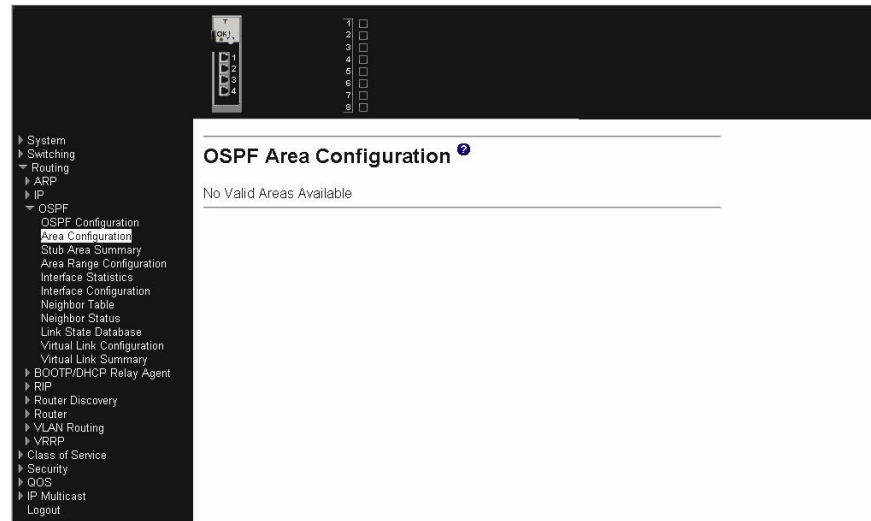
External LSDB Limit

Enter the maximum number of non-default AS-external LSA entries that should be stored in the LSDB. This parameter establishes the external LSDB limit for OSPF. If you enter -1, then there will be no limit. When the number of non-default AS-external-LSAs in the router's LSDB reaches the external LSDB limit, the router will enter the overflow state. The router will never hold more than the external LSDB limit non-default AS-external-LSAs in its database. You MUST set the external LSDB limit to the same value in all routers attached to the OSPF backbone and/or any regular OSPF area. The range is -1 to 2147483647.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Area configuration

Use this panel to configure and display the parameters for an OSPF area.



Area Select the area to be configured or displayed.

Area ID

The area whose data is displayed.

Aging Interval (secs)

The LSA aging timer interval.

External Routing

A definition of the router's capabilities for the area, including whether or not AS-external-LSAs are flooded into/throughout the area. If the area is a stub area, then these are the possible options for which you can configure the external routing capability, otherwise the only option is Import External LSAs.

- Import External LSAs - Import and propagate external LSAs
- Import No LSAs - Do not import and propagate external LSAs

SPF Runs

The number of times that the intra-area route table has been calculated using this area's LSDB. This is typically done using Dijkstra's algorithm.

Authentication Type

Specifies if summary LSAs are imported into the stub area.

Area Border Router Count

The total number of ABRs reachable within this area. This is initially zero, and is calculated in each SPF Pass.

Area LSA Count

The total number of LSAs in this area's LSDB, excluding AS External LSAs.

Area LSA Checksum

The 32-bit unsigned sum of the LSA's LS checksums contained in this area's LSDB. This sum excludes external (LS type 5) LSAs. The sum can be used to determine if there has been a change in a router's LSDB, and to compare the LSDB of two routers.

Mode This field tells you whether the area is or is not a stub area. If the area can be a stub area, a Create Stub Area button will be displayed. If you have

configured the area as a stub area a Delete Stub Area button will be displayed. Otherwise neither button will be displayed.

The following fields are only available when Stub Area is selected.

Import Summary LSAs

Select Enable or Disable from the pull-down menu. If you select Enable, summary LSAs will be imported into stub areas. This field is only present and configurable when the selected area is a stub area.

Metric Value Enter the metric value you want applied for the default route advertised into the stub area. Valid values range from 1 to 16,777,215. This field is only present and configurable when the selected area is a stub area.

Metric Type This field is only present and configurable when the selected area is a stub area. Select the type of metric specified in the Metric Value field.

- OSPF Metric - Regular OSPF metric.
- Comparable Cost - External Type 1 metrics that are comparable to the OSPF metric.
- Non-comparable Cost - External Type 2 metrics that are assumed to be larger than the cost of the OSPF metric.

Type of Service

The type of service associated with the stub metric. The switch supports Normal only. This field is present only when the selected area is a stub area.

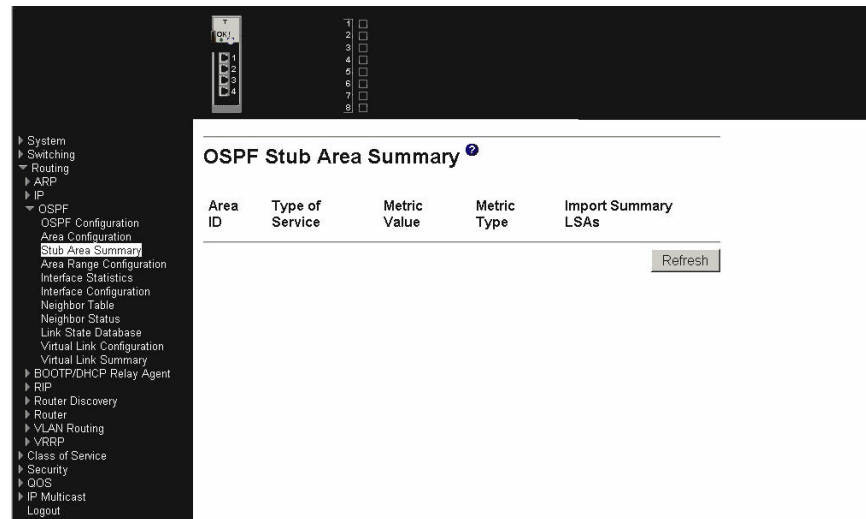
Click the Create Stub Area button to configure the area as a stub area. This button is present only when the selected area is able to become a stub area. Area '0.0.0.0' cannot be a stub area.

Click the Delete Stub Area button to delete the stub area designation. The area will be returned to normal state. This button will be present only when the selected area is a stub area.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Stub area summary

Use this panel to configure and display the parameters for an OSPF stub area.



Area ID

The Area ID of the Stub area.

Type of Service

The type of service associated with the stub metric. The switch supports Normal only.

Metric Value

Set the metric value you want applied for the default route advertised into the area. Valid values range from 1 to 16,777,215.

Metric Type

The type of metric for the stub area where valid types are:

- OSPF Metric - Regular OSPF metric.
- Comparable Cost - External Type 1 metrics that are comparable to the OSPF metric.
- Non-comparable Cost - External Type 2 metrics that are assumed to be larger than the cost of the OSPF metric.

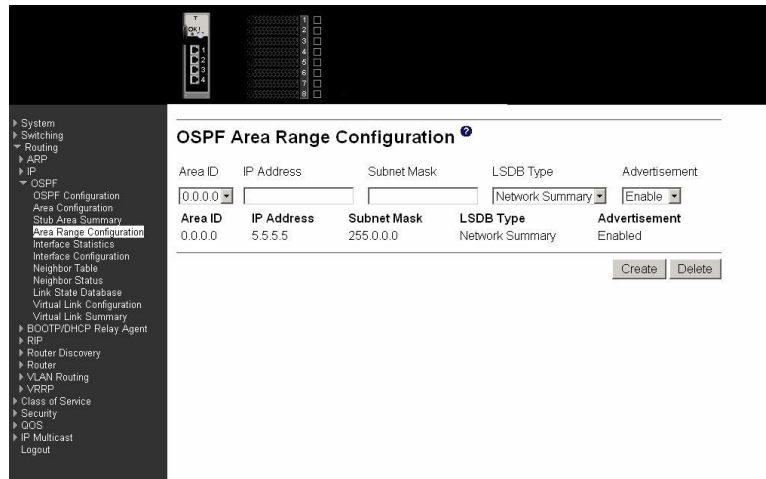
Import Summary LSAs

Whether the import of Summary LSAs is Enabled or Disabled.

Click the Refresh button to refresh the data on the screen to the current values from the switch.

Area range configuration

Use this panel to configure an IP address range for a selected OSPF area, and to specify the type of LSA. You can also use this panel to display an already configured area range.



Area ID Selects the area for which data is to be configured.

IP Address Enter the IP address for the address range for the selected area.

Subnet Mask Enter the Subnet Mask for the address range for the selected area.

LSDB Type Select the type of Link Advertisement associated with the specified area and address range. The default type is Network Summary.

Advertisement Select Enable or Disable from the pull-down menu. If you selected Enable the address range will be advertised outside the area via a Network Summary LSA. The default is Enable.

Area ID The OSPF area.

IP Address The IP address of an address range for the area.

Subnet Mask The Subnet Mask of an address range for the area.

LSDB Type The Link Advertisement type for the address range and area.

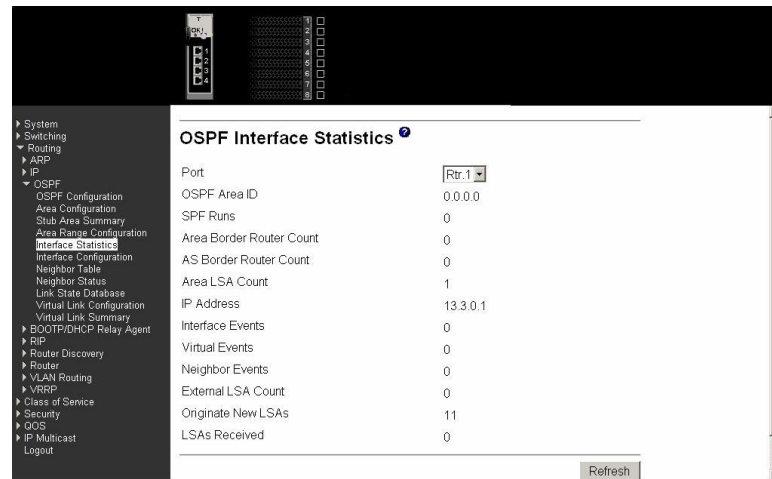
Advertisement The Advertisement mode for the address range and area.

Click the Create button to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed. The new address range will be added to the display in the non-configurable data area.

Click the Delete button to remove the specified address range from the area configuration.

Interface statistics

Use this panel to display the OSPF statistics associated with an interface.



Port Select the interface for which data is to be displayed.

OSPF Area ID The OSPF area to which the selected router interface belongs. An OSPF Area ID is a 32-bit integer in dotted decimal format that uniquely identifies the area to which the interface connects.

SPF Runs The number of times that the intra-area route table has been calculated using this area's LSDB.

Area Border Router Count The total number of area border routers reachable within this area. This is initially zero, and is calculated in each SPF Pass.

AS Border Router Count The total number of ASBRs reachable within this area. This is initially zero, and is calculated in each SPF Pass.

Area LSA Count The total number of LSAs in this area's LSDB, excluding AS External LSAs.

IP Address The IP address of the interface.

Interface Events The number of times the specified OSPF interface has changed its state, or an error has occurred.

Virtual Events The number of state changes or errors that have occurred on this virtual link.

Neighbor Events The number of times this neighbor relationship has changed state, or an error has occurred.

External LSA Count The number of external (LS type 5) LSAs in the LSDB.

Originate New LSAs The number of new LSAs that have been originated. In any given OSPF area, a router will originate several LSAs. Each router

originates a router-LSA. If the router is also the Designated Router for any of the area's networks, it will originate network-LSAs for those networks.

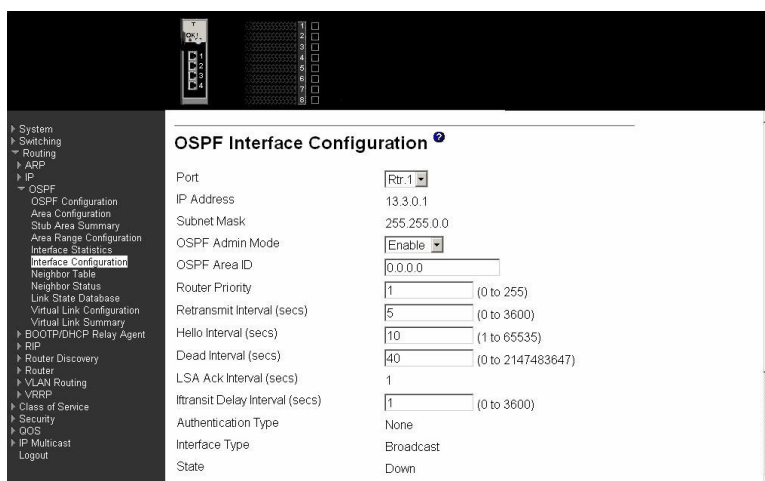
LSAs Received

The number of LSAs that have been received that have been determined to be new instantiations. This number does not include newer instantiations of self-originated LSAs.

Click the Refresh button to Refresh the data on the screen with the present state of the data in the switch.

Interface configuration

Use this panel to configure and display the OSPF parameters for a router interface.



Port Select the interface for which data is displayed or configured.

IP Address The IP address of the interface.

Subnet Mask The subnet/network mask indicating the portion of the IP interface address that identifies the attached network.

OSPF Admin Mode

You can select Enable or Disable from the pull-down menu. The default value is Disable. You can configure OSPF parameters without enabling OSPF Admin Mode, but they will have no effect until you enable Admin Mode. The following information will be displayed only if the Admin Mode is enabled: State, Designated Router, Backup Designated Router, Number of Link Events, LSA Ack Interval, and Metric Cost. For OSPF to be fully functional, you must enter a valid IP address and Subnet Mask via the Interface IP Configuration page or through the CLI command: **config ip interface network**.

It is important to remember that when OSPF is initialized on the router, it will remain initialized until the router is reset.

OSPF Area ID Enter the 32-bit integer in dotted decimal format that uniquely identifies the OSPF area to which the selected router interface connects. If you assign an Area ID which does not exist, the area will be created with default values.

Router Priority

Enter the OSPF priority for the selected interface. The priority of the Neighbor is a priority integer from 0 to 255. The default is 1, which is the highest router priority. A value of "0" indicates that the router is not eligible to become the designated router on this network.

Retransmit Interval (secs)

Enter the OSPF retransmit interval for the specified interface. This is the number of seconds between LSAs for adjacencies belonging to this router interface. This value is also used when retransmitting database descriptions and link-state request packets. Valid values range from 0 to 3600 seconds (1 hour). The default is 5 seconds.

Hello Interval (secs)

Enter the OSPF Hello Interval for the specified interface in seconds. This parameter must be the same for all routers attached to a network. Valid values range from 1 to 65,535. The default is 10 seconds.

Dead Interval (secs)

Enter the OSPF dead interval for the specified interface in seconds. This specifies how long a router will wait to see a neighbor router's Hello packets before declaring that the router is down. This parameter must be the same for all routers attached to a network. This value should be a multiple of the Hello Interval. Valid values range from 0 to 2147483647. The default is 40.

LSA Ack Interval (secs)

The number of seconds between LSA Acknowledgment packet transmissions, which must be less than the Retransmit Interval.

Iftransit Delay Interval (secs)

Enter the OSPF Transit Delay for the specified interface. This specifies the estimated number of seconds it takes to transmit a link state update packet over the selected interface. Valid values range from 0 to 3600 seconds (1 hour). The default value is 1 second.

Authentication Type

You can select an authentication type other than none by clicking on the Configure button. You will then see a new screen, where you can select the authentication type from the pull-down menu. The choices are:

None This is the initial and default authentication state. If you select this option from the pull-down menu on the second screen you will be returned to the first screen and no authentication protocols will be run.

Simple

If you select Simple you will be prompted to enter an authentication key. This key will be included, in the clear, in the OSPF header of all packets sent on the network. All routers on the network must be configured with the same key.

Encrypt

If you select Encrypt you will be prompted to enter both an authentication key and an authentication ID. Encryption uses the MD5 Message-Digest algorithm. All routers on the network must be configured with the same key and ID.

Interface Type

The OSPF interface type, which will always be broadcast.

State

The state of the OSPF interface. The possible values are:

Down

This is the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable. In this state, interface parameters will be set to their initial values. All interface timers will be disabled, and there will be no adjacent routers associated with the interface.

Loopback

In this state, the router's interface to the network is looped back either in hardware or software. The interface is unavailable for regular data traffic. However, it might still be desirable to gain information on the quality of this interface, either through sending ICMP pings to the interface or through something like a bit error test. For this reason, IP packets can still be addressed to an interface in Loopback state. To facilitate this, such interfaces are advertised in router-LSAs as single host routes, whose destination is the IP interface address.

Waiting

The router is trying to determine the identity of the (Backup) Designated Router for the network by monitoring received Hello Packets. The router is not allowed to elect a Backup Designated Router or a Designated Router until it transitions out of Waiting state. This prevents unnecessary changes of (Backup) Designated Router.

Designated router

This router is itself the Designated Router on the attached network. Adjacencies are established to all other routers attached to the network. The router must also originate a network-LSA for the network node. The network-LSA will contain links to all routers (including the Designated Router itself) attached to the network

Backup designated router

This router is itself the Backup Designated Router on the attached network. It will be promoted to Designated Router if the present Designated Router fails. The router establishes adjacencies to all other routers attached to the network. The Backup Designated Router performs slightly different functions during the Flooding Procedure, as compared to the Designated Router

Other designated router

The interface is connected to a broadcast or NBMA network on which other routers have been selected to be the Designated Router and Backup Designated Router either. The router attempts to form adjacencies to both the Designated Router and the Backup Designated Router.

Designated Router

The identity of the Designated Router for this network, in the view of the Advertising Router. The Designated Router is identified by its

router ID. The value 0.0.0.0 means that there is no Designated Router. This field is only displayed if the OSPF admin mode is Enabled.

Backup Designated Router

The identity of the Backup Designated Router for this network, in the view of the advertising router. The Backup Designated Router is identified here by its router ID. Set to 0.0.0.0 if there is no Backup Designated Router. This field is only displayed if the OSPF admin mode is Enabled.

Number of Link Events

This is the number of times the specified OSPF interface has changed its state. This field is only displayed if the OSPF admin mode is Enabled.

Metric Cost

Enter the value on this interface for the cost. The range for the metric cost is between 0 and 65,535. Metric Cost is only configurable/displayed if OSPF is initialized on the interface.

Click the Configure Authentication button to display a new screen where you can select the authentication method for the virtual link.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Neighbor table

Use this panel to display information about the selected interface's neighbor router.

Router ID	IP Address	Neighbor Interface Index
192.168.112.50	9.1.1.2	Rtr.1

Port

Selects the interface for which data is to be displayed or configured.

Router ID

A 32-bit integer in dotted decimal format representing the neighbor interface.

IP Address

The IP address of the neighboring router's interface to the attached network. It is used as the destination IP address when protocol packets are sent as unicasts along this adjacency. Also used in router-LSAs as the Link ID for the attached network if the neighboring router is selected to be designated router. The Neighbor IP address is learned when Hello packets are received

from the neighbor. For virtual links, the Neighbor IP address is learned during the routing table build process.

Neighbor Interface Index

A port identifying the neighbor interface index.

Refresh the data on the screen with the present state of the data in the switch.

Neighbor status

This panel displays the OSPF neighbor configuration for a selected neighbor ID. When a particular neighbor ID is specified, detailed information about a neighbor is given. The information below will only be displayed if OSPF is Enabled and the interface has a neighbor. The IP address is the IP address of the neighbor.



Port Selects the interface for which data is to be displayed or configured.

Neighbor IP Address

Selects the IP address of the neighbor for which data is to be displayed.

Router ID A 32-bit integer in dotted decimal format that identifies the neighbor router.

Options

The optional OSPF capabilities supported by the neighbor. The neighbor's optional OSPF capabilities are also listed in its Hello packets. This enables received Hello packets to be rejected (i.e., neighbor relationships will not even start to form) if there is a mismatch in certain crucial OSPF capabilities.

Router Priority

Displays the OSPF priority for the specified neighbor. The priority of a neighbor is a priority integer from 0 to 255. A value of "0" indicates that the router is not eligible to become the designated router on this network.

State

The state of a neighbor can be the following:

Down This is the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor. On Non-Broadcast Multi-Access (NBMA)

networks, Hello packets can still be sent to Down neighbors, although at a reduced frequency.

Attempt

This state is only valid for neighbors attached to NBMA networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor. This is done by sending the neighbor Hello packets at intervals of Hello Interval.

Init

In this state, a Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor (i.e., the router itself did not appear in the neighbor's Hello packet). All neighbors in this state (or greater) are listed in the Hello packets sent from the associated interface.

2-Way

In this state, communication between the two routers is bidirectional. This has been assured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (Backup) Designated Router is selected from the set of neighbors in state 2-Way or greater.

Exchange Start

This is the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial DD sequence number. Neighbor conversations in this state or greater are called adjacencies.

Exchange

In this state the router is describing its entire LSDB by sending Database Description packets to the neighbor. In this state, Link State Request Packets can also be sent asking for the neighbor's more recent LSAs. All adjacencies in Exchange state or greater are used by the flooding procedure. These adjacencies are fully capable of transmitting and receiving all types of OSPF routing protocol packets.

Loading

In this state, Link State Request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the Exchange state.

Full

In this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router-LSAs and network-LSAs.

Events

The number of times this neighbor relationship has changed state, or an error has occurred.

Permanence

This variable displays the status of the entry. Dynamic and Permanent refer to how the neighbor became known.

Hello Suppressed

This indicates whether Hellos are being suppressed to the neighbor.

Retransmission Queue Length

The current length of the retransmission queue.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

LSDB

Use this panel to display the contents of the OSPF LSDB.

Router ID	Area ID	LSA Type	LS ID	Age	Sequence	Checksum
3.3.3.3	0.0.0.0	Network Summary	13.0.0.0	748	2147483651	51056
3.3.3.3	0.0.0.1	Router Links	3.3.3.3	754	2147483651	4880
3.3.3.3	0.0.0.1	Network Summary	0.0.0.0	794	2147483650	14106
3.3.3.3	0.0.0.2	Network Summary	13.0.0.0	748	2147483651	51056

Router ID

The 32-bit integer in dotted decimal format that uniquely identifies the router within the AS. The Router ID is set on the IP Configuration page. If you want to change the Router ID you must first disable OSPF. After you set the new Router ID, you must re-enable OSPF for the change take effect. The default value is 0.0.0.0, although this is not a valid Router ID.

Area ID

The ID of an OSPF area to which one of the router interfaces is connected. An Area ID is a 32-bit integer in dotted decimal format that uniquely identifies the area to which an interface is connected.

LSA Type

The format and function of the LSA. One of the following:

- Router Links
- Network Links
- Network Summary
- ASBR Summary
- AS-external

LS ID

The Link State ID identifies the piece of the routing domain that is described by the advertisement. The value of the LS ID is determined by the advertisement's LS type.

Age

The time since the LSA was first originated, in seconds.

Sequence

The sequence number field is a signed 32-bit integer. It is used to detect old and duplicate LSAs. The larger the sequence number, the more recent the advertisement.

Checksum

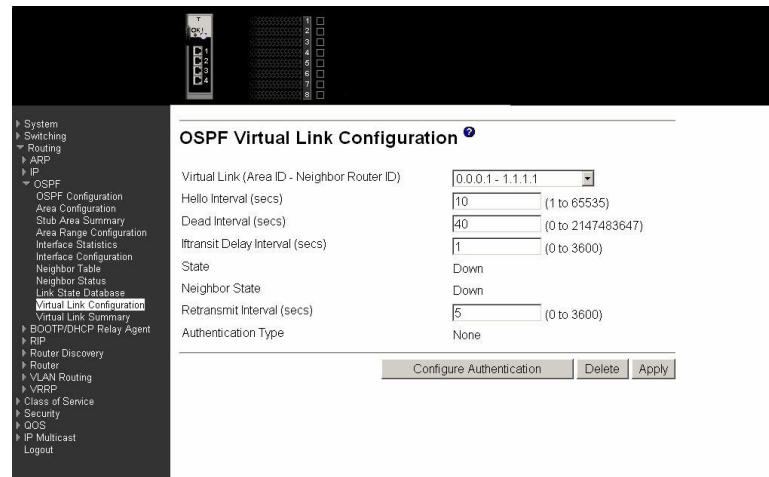
The checksum is used to detect data corruption of an advertisement. This corruption can occur while an advertisement is being flooded, or while it is being held in a router's memory. This

field is the checksum of the complete contents of the advertisement, except the LS age field.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch

Virtual link configuration

Use this panel to configure a virtual link to a neighboring router.



Virtual Link (Area ID and Neighbor Router ID)

Select the virtual link for which you want to display or configure data. It consists of the Area ID and Neighbor IP address.

Neighbor Router ID

Enter the neighbor portion of a Virtual Link specification. Virtual links can be configured between any pair of area border routers having interfaces to a common (non-backbone) area. You only enter this ID when you are creating a new virtual link.

Hello Interval (secs)

Enter the OSPF Hello Interval for the specified interface in seconds. This parameter must be the same for all routers attached to a network. Valid values range from 1 to 65,535. The default is 10 seconds.

Dead Interval (secs)

Enter the OSPF Dead Interval for the specified interface in seconds. This specifies how long a router will wait to see a neighbor router's Hello packets before declaring that the router is down. This parameter must be the same for all routers attached to a network. This value should be a multiple of the Hello Interval (e.g. 4). Valid values range from 1 to 2147483647. The default is 40.

Iftransit Delay Interval (secs)

Enter the OSPF Transit Delay for the specified interface. This specifies the estimated number of seconds it takes to transmit a link state update packet over the selected interface. Valid values range from 1 to 3600 seconds (1 hour). The default value is 1 second.

State

Following are the possible states:

Down This is the initial interface state. In this state, the lower level

protocols have indicated that the interface is unusable. Interface parameters will be set to their initial values. All interface timers will be disabled and there will be no adjacencies associated with the interface.

Waiting

The router is trying to determine the identity of the Designated Router and its Backup by monitoring received Hello packets. To prevent unnecessary changes, the router is not allowed to elect a Designated Router or a Backup Designated Router until it transitions out of Waiting State.

Point-to-Point

The interface is operational, and is connected to the virtual link. Upon entering this state the router attempts to form an adjacency with the neighboring router. Hello packets are sent to the neighbor every Hello Interval seconds.

Designated Router

Indicates the Designated Router on the attached network. Adjacencies are established to all other routers attached to the network. The router must also originate a network LSA for the network node. The network LSA will contain links to all routers (including the Designated Router itself) attached to the network.

Backup Designated Router

Indicates the Backup Designated Router on the attached network. The Backup will be promoted to Designated Router if the current Designated Router fails. The router establishes adjacencies to all other routers attached to the network. The Backup Designated Router performs functions differently during the Flooding Procedure than the Designated Router.

Other Designated Router

The router(s) on the interface not selected as either the Backup or Designated routers. The interface is connected to a broadcast or NBMA network. The Other Designated Router attempts to form adjacencies to both the Designated Router and the Backup Designated Router.

Neighbor State

The state of the Virtual Neighbor Relationship.

Retransmit Interval (secs)

Enter the OSPF Retransmit Interval for the specified interface. This is the number of seconds between LSAs for adjacencies belonging to this router interface. This value is also used when retransmitting database descriptions and link-state request packets. Valid values range from 1 to 3600 seconds (1 hour). The default is 5 seconds.

Authentication Type

You can select an authentication type other than none by clicking on the Configure Authentication button. You will then see a new screen, where you can select the authentication type from the pull-down menu. The choices are:

None This is the initial and default authentication state. If you

select this option from the pull-down menu on the second screen you will be returned to the first screen.

Simple

If you select Simple you will be prompted to enter an authentication key. This key will be included, in the clear, in the OSPF header of all packets sent on the network. All routers on the network must be configured with the same key.

Encrypt

If you select Encrypt you will be prompted to enter both an authentication key and an authentication ID. Encryption uses the MD5 Message-Digest algorithm. All routers on the network must be configured with the same key and ID.

Authentication Key

Enter the OSPF Authentication Key for the specified interface. If you do not choose to use authentication you will not be prompted to enter a key. If you choose simple authentication you cannot use a key of more than 8 octets. If you choose encrypt the key can be up to 16 octets long. The key value will only be displayed if you are logged on with Read/Write privileges, otherwise it will be displayed as asterisks.

Authentication ID

Enter the ID to be used for authentication. You will only be prompted to enter an ID when you select Encrypt as the authentication type. The ID is a number between 0 and 255, inclusive.

Click the Configure Authentication button to display a new screen where you can select the authentication method for the virtual link.

Click the Delete button to remove the specified virtual link from the router configuration.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Virtual link summary

Use this panel to display the parameters configured for the specified virtual link.

Area ID	Neighbor Router ID	Hello Interval (secs)	Dead Interval (secs)	Retransmit Interval (secs)	iftransit Delay Interval (secs)
0.0.0.1	1.1.1.1	10	40	5	1

Area ID The Area ID portion of the virtual link identification for which data is displayed. The Area ID and Neighbor Router ID together define a virtual link.

Neighbor Router ID The neighbor portion of the virtual link identification. Virtual links can be configured between any pair of area border routers having interfaces to a common (non-backbone) area.

Hello Interval (secs) The OSPF Hello Interval for the virtual link in units of seconds. The value for Hello Interval must be the same for all routers attached to a network.

Dead Interval (secs) The OSPF Dead Interval for the virtual link in units of seconds. This specifies how long a router will wait to see a neighbor router's Hello packets before declaring that the router is down. This parameter must be the same for all routers attached to a common network, and should be a multiple of the Hello Interval.

Retransmit Interval (secs) The OSPF retransmit interval for the virtual link in units of seconds. This specifies the time between LSAs for adjacencies belonging to this router interface. This value is also used when retransmitting database descriptions and link-state request packets.

iftransit Delay Interval (secs) The OSPF Transit Delay for the virtual link in units of seconds. It specifies the estimated number of seconds it takes to transmit a link state update packet over this interface.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch

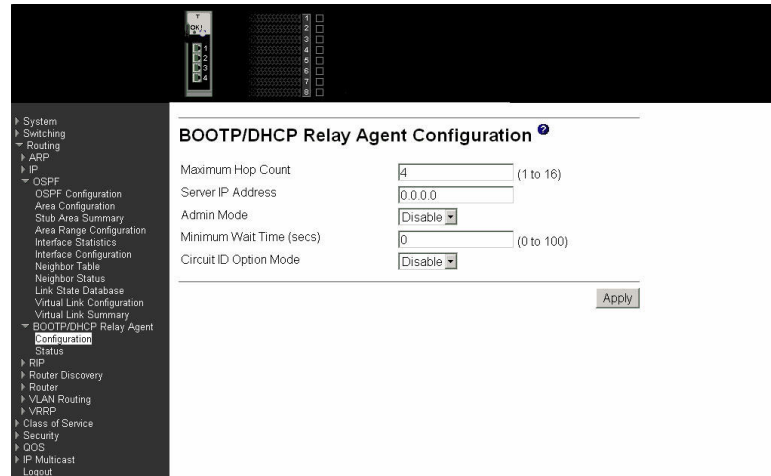
BOOTP/DHCP relay

This menu provides access to the following BOOTP/DHCP Relay Agent entry panels:

- Configuration
- Status

Configuration

Use this panel to configure the parameters for a BOOTP/DHCP server or relay agent.



Maximum Hop Count

Maximum number of hops a client request can go without being discarded.

Server IP Address

IP address of BOOTP/DHCP server or the IP address of the next BOOTP/DHCP Relay Agent.

Admin Mode Administrative mode of the relay can be either Enabled or Disabled. When Admin mode is Enabled the relay forwards the BOOTP/DHCP requests to the IP address given by Server IP address.

Minimum Wait Time

The Minimum time in seconds the client should wait before its request can be forwarded to the BOOTP/DHCP server.

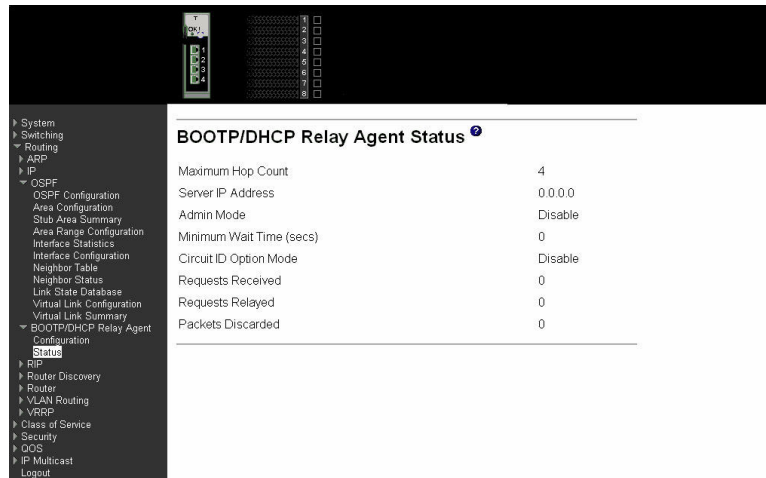
Circuit ID Option Mode

This is the Relay agent option which can be either Enabled or Disabled. When Enabled adds Relay agent options to the request before forwarding it to Server and removes the options before forwarding the reply to clients.

Click the Apply button to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

Status

This panel displays the parameters specified on the Relay Agent Configuration screen.



Maximum Hop Count

Maximum number of Hops a client request can go without being discarded.

Server IP Address

IP address of BOOTP/DHCP server or the IP address of the next BOOTP/DHCP Relay Agent.

Admin Mode Administrative mode of the relay can be either Enabled or Disabled. When Admin mode is Enabled the relay forwards the BOOTP/DHCP requests to the IP address given by Server IP address.

Minimum Wait Time (secs)

The Minimum time in seconds the client should wait before its request can be forwarded to the BOOTP/DHCP server.

Circuit ID Option Mode

This is the Relay agent option which can be either Enabled or Disabled. When Enabled adds Relay agent options to the request before forwarding it to Server and removes the options before forwarding the reply to clients.

Requests Received

Total number of BOOTP/DHCP requests received from all clients from the system bootup time.

Requests Delayed

Total number of BOOTP/DHCP requests forwarded to Server from the system bootup time.

Packets Discarded

Total number of BOOTP/DHCP packets discarded by this relay agent from the system bootup time.

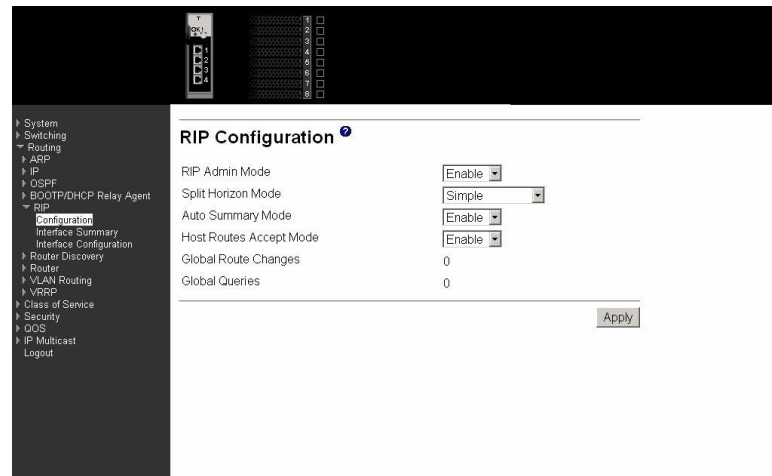
RIP

This menu provides access to the following Routing Information Protocol (RIP) entry panels:

- Configuration
- Interface summary
- Interface configuration

Configuration

Use this panel to configure and display the RIP parameters for the switch operating as a RIP router.



RIP Admin Mode

Select Enable or Disable from the pull-down menu. If you select Enable RIP will be Enabled for the switch. The default is Disable.

Split Horizon Mode

Select none, simple or poison reverse from the pull-down menu. Split horizon is a technique for avoiding problems caused by including routes in updates sent to the router from which the route was originally learned. The options are:

None No special processing for this case.

Simple

A route will not be included in updates sent to the router from which it was learned. This is the default.

Poisoned Reverse

A route will be included in updates sent to the router from which it was learned, but the metric will be set to infinity.

Auto Summary Mode

Select Enable or Disable from the pull-down menu. If you select Enable groups of adjacent routes will be summarized into single entries, in order to reduce the total number of entries. The default is Enable.

Host Routes Accept Mode

Select Enable or Disable from the pull-down menu. If you select Enable the router will be accept host routes. The default is Enable.

Global Route Changes

The number of route changes made to the IP Route Database by RIP. This does not include the refresh of a route's age.

Global Queries

The number of responses sent to RIP queries from other systems.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Interface summary

Use this panel to display the RIP parameters for the selected router interface.

Port	IP Address	Send Version	Receive Version	RIP Admin Mode	Link State
Rtr.1	13.3.0.1	RIP-2	Both	Enable	Link Down
Rtr.2	0.0.0.0			Disable	Link Down
Rtr.3	0.0.0.0			Disable	Link Down
Rtr.4	13.0.0.1	RIP-2	Both	Enable	Link Down

Port The physical interface for which the information is being displayed.

IP Address The IP address of the router interface.

Send Version The RIP version to which RIP control packets sent from the interface conform. The value is one of the following:

RIP-1 RIP version 1 packets will be sent using broadcast.

RIP-1c

RIP version 1 compatibility mode. RIP version 2 formatted packets will be transmitted using broadcast.

RIP-2 RIP version 2 packets will be sent using multicast. This is the default.

None RIP control packets will not be transmitted.

Receive Version

Which RIP version controls packets will be accepted by the interface. The value is one of the following:

RIP-1 Only RIP version 1 formatted packets will be received.

RIP-2 Only RIP version 2 formatted packets will be received.

Both Packets will be received in either format. This is the default.

None No RIP control packets will be received.

RIP Admin Mode

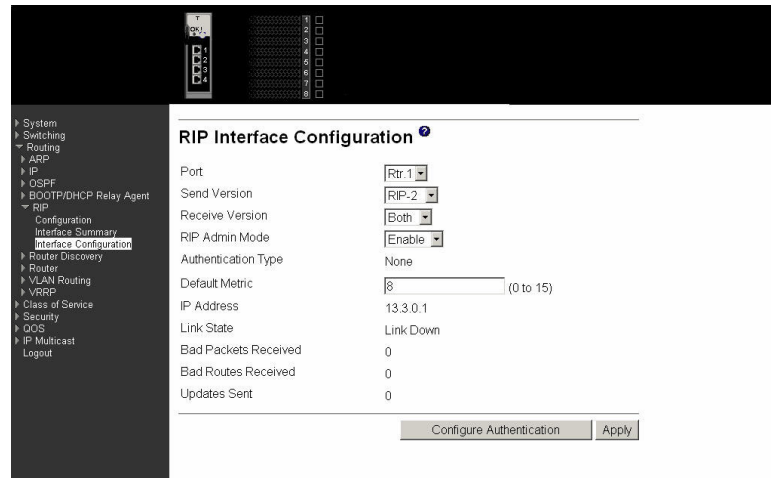
Whether RIP is Enabled or Disabled on the interface.

Link State Whether the RIP interface is up or down.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

Interface configuration

Use this panel to configure the selected interface as a RIP router interface.



Port Select the interface for which data is to be configured.

Send Version Select the version of RIP control packets the interface should send from the pull-down menu. The value is one of the following:

RIP-1 RIP version 1 packets will be sent using broadcast.

RIP-1c

RIP version 1 compatibility mode. RIP version 2 formatted packets will be transmitted using broadcast.

RIP-2 RIP version 2 packets will be sent using multicast. This is the default.

None RIP control packets will not be transmitted.

Receive Version

Which RIP version controls packets will be accepted by the interface. The value is one of the following:

RIP-1 Accept only RIP version 1 formatted packets.

RIP-2 Accept only RIP version 2 formatted packets.

Both Accept packets in either format. This is the default.

None No RIP control packets will be accepted.

RIP Admin Mode

Select Enable or Disable from the pull-down menu. Before you Enable RIP version 1 or version 1c on an interface, you must first Enable network directed broadcast mode on the corresponding interface. The default value is Disable.

Authentication Type

You can select an authentication type other than none by clicking

on the Configure Authentication button. You will then see a new screen, where you can select the authentication type from the pull-down menu. The choices are:

None This is the initial and default authentication state. If you select this option from the pull-down menu on the second screen you will be returned to the first screen and no authentication protocols will be run.

Simple

If you select Simple you will be prompted to enter an authentication key. This key will be included, in the clear, in the OSPF header of all packets sent on the network. All routers on the network must be configured with the same key.

Encrypt

If you select Encrypt you will be prompted to enter both an authentication key and an authentication ID. Encryption uses the MD5 Message-Digest algorithm. All routers on the network must be configured with the same key and ID.

Default Metric

The metric that is to be used for the default route entry in RIP updates originated on this interface. The range for the default metric is between 0 and 15. Note that a metric value of 0 suppresses default route originations (although a default route can be propagated on this interface from another router). A metric value of 1 instructs the router to always advertise a default route entry with a metric of 1 in its route update messages, which could adversely affect network operation.

IP Address The IP address of the router interface.

Link State Indicates whether the RIP interface is up or down.

Bad Packets Received

The number of RIP response packets received by the RIP process which were subsequently discarded for any reason.

Bad Routes Received

The number of routes, in valid RIP packets, which were ignored for any reason (e.g. unknown address family, or invalid metric).

Updates Sent The number of triggered RIP updates actually sent on this interface. This explicitly does NOT include full updates sent containing new information.

Click the Configure Authentication button to display a new screen where you can select the authentication method for the virtual link.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

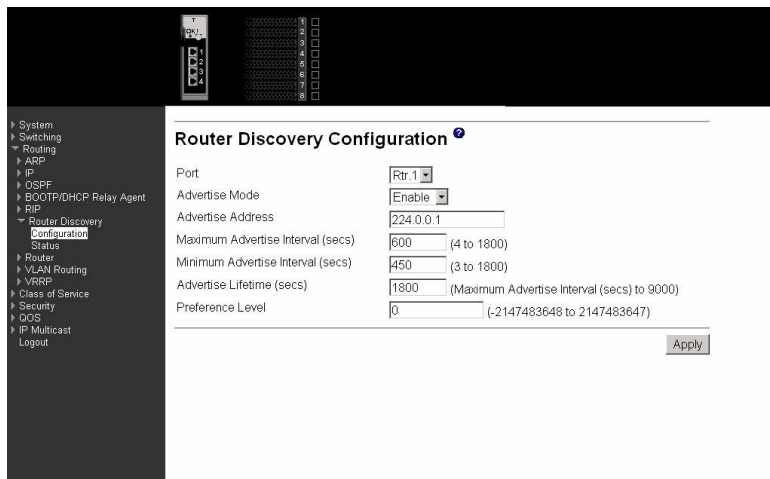
Router discovery

This menu provides access to the following Router (rtr) discovery entry panels:

- Configuration
- Status

Configuration

Use this panel to configure the selected interface to transmit router advertisements.



The screenshot shows a web-based configuration interface for a network device. On the left is a navigation menu with categories like System, Switching, Routing, ARP, IP, OSPF, BOOTP/DHCP Relay Agent, RIP, Router Discovery, Status, Router, VLAN Routing, VRRP, Class of Service, Security, QoS, IP Multicast, and Logout. The 'Router Discovery' section is expanded. The main area is titled 'Router Discovery Configuration' and contains the following fields:

Port	<input type="text" value="Rtr 1"/>
Advertise Mode	<input type="text" value="Enable"/>
Advertise Address	<input type="text" value="224.0.0.1"/>
Maximum Advertise Interval (secs)	<input type="text" value="600"/> (4 to 1800)
Minimum Advertise Interval (secs)	<input type="text" value="450"/> (3 to 1800)
Advertise Lifetime (secs)	<input type="text" value="1800"/> (Maximum Advertise Interval (secs) to 9000)
Preference Level	<input type="text" value="0"/> (-2147483648 to 2147483647)

An 'Apply' button is located at the bottom right of the configuration area.

Port Select the router interface for which data is to be configured.

Advertise Mode

Select Enable or Disable from the pull-down menu. If you select Enable, Router Advertisements will be transmitted from the selected interface.

Advertise Address

Enter the IP address to be used to advertise the router.

Maximum Advertise Interval (secs)

Enter the maximum time (in seconds) allowed between router advertisements sent from the interface.

Minimum Advertise Interval (secs)

Enter the minimum time (in seconds) allowed between router advertisements sent from the interface.

Advertise Lifetime (secs)

Enter the value (in seconds) to be used as the lifetime field in router advertisements sent from the interface. This is the maximum length of time that the advertised addresses are to be considered as valid router addresses by hosts.

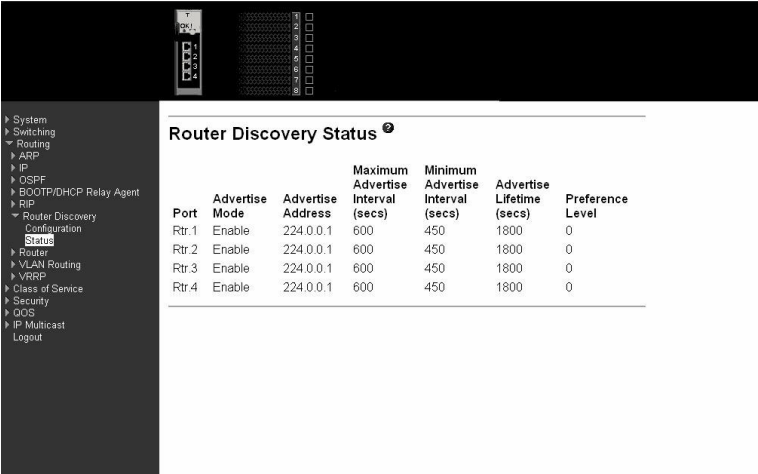
Preference Level

Specify the preference level of the router as a default router relative to other routers on the same subnet. Higher numbered addresses are preferred. You must enter an integer.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Status

This panel displays the router discovery details specified on the Router Discovery Configuration screen.



Port	Advertise Mode	Advertise Address	Maximum Advertise Interval (secs)	Minimum Advertise Interval (secs)	Advertise Lifetime (secs)	Preference Level
Rtr.1	Enable	224.0.0.1	600	450	1800	0
Rtr.2	Enable	224.0.0.1	600	450	1800	0
Rtr.3	Enable	224.0.0.1	600	450	1800	0
Rtr.4	Enable	224.0.0.1	600	450	1800	0

Port The router interface for which data is displayed.

Advertise Mode

The values are Enable or Disable. Enable denotes that Router Discovery is Enabled on that interface.

Advertise Address

The IP address used to advertise the router.

Maximum Advertise Interval (secs)

The maximum time (in seconds) allowed between router advertisements sent from the interface.

Minimum Advertise Interval (secs)

The minimum time (in seconds) allowed between router advertisements sent from the interface.

Advertise Lifetime (secs)

The value (in seconds) used as the lifetime field in router advertisements sent from the interface. This is the maximum length of time that the advertised addresses are to be considered as valid router addresses by hosts.

Preference Level

The preference level of the router as a default router relative to other routers on the same subnet. Higher numbered addresses are preferred.

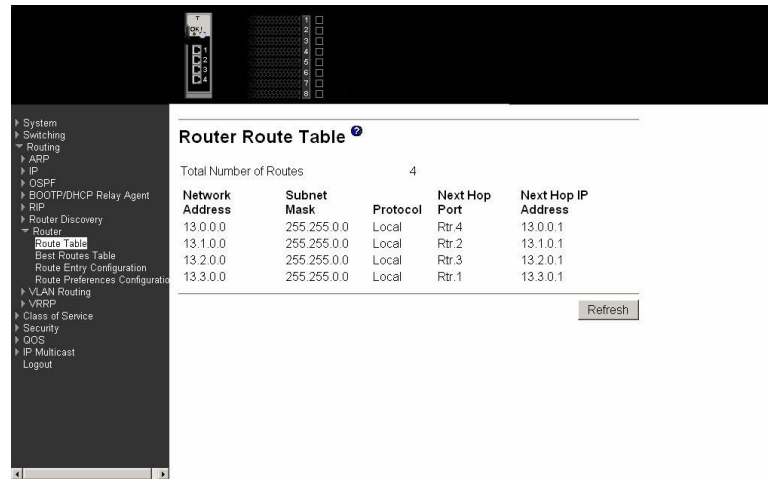
Router

This menu provides access to the following Route table entry panels:

- Route table
- Best routes table
- Route entry configuration
- Route preferences configuration

Route table

Use this panel to display the entire contents of the route table.



Total Number of Routes

The total number of routes in the route table.

Network Address

The IP route prefix for the destination.

Subnet Mask The subnet/network mask that indicates the portion of the IP interface address that identifies the attached network.

Protocol Identifies the protocol that added the route to the table. One of the following:

- Local
- Static
- Default
- MPLS
- OSPF Intra
- OSPF Inter
- OSPF Type-1
- OSPF Type-2
- RIP
- BGP4

Next Hop Port

The router interface used to forward traffic to the destination.

Next Hop IP Address

The IP address of the next router in the path towards the destination. The next router will always be one of the adjacent neighbors or the IP address of the local interface for a directly attached network.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

Best routes table

Use this panel to display the best routes to destinations with routes in the route table.

Network Address	Subnet Mask	Protocol	Next Hop Port	Next Hop IP Address
13.0.0.0	255.255.0.0	Local	Rtr:4	13.0.0.1
13.1.0.0	255.255.0.0	Local	Rtr:2	13.1.0.1
13.2.0.0	255.255.0.0	Local	Rtr:3	13.2.0.1
13.3.0.0	255.255.0.0	Local	Rtr:1	13.3.0.1

Total Number of Routes

The total number of routes in the route table.

Network Address

The IP route prefix for the destination.

Subnet Mask The subnet/network mask that indicates the portion of the IP interface address that identifies the attached network.

Protocol Identifies the protocol that added the route to the table. One of the following:

- Local
- Static
- Default
- MPLS
- OSPF Intra
- OSPF Inter
- OSPF Type-1
- OSPF Type-2
- RIP
- BGP4

Next Hop Port

The router interface used to forward traffic to the destination.

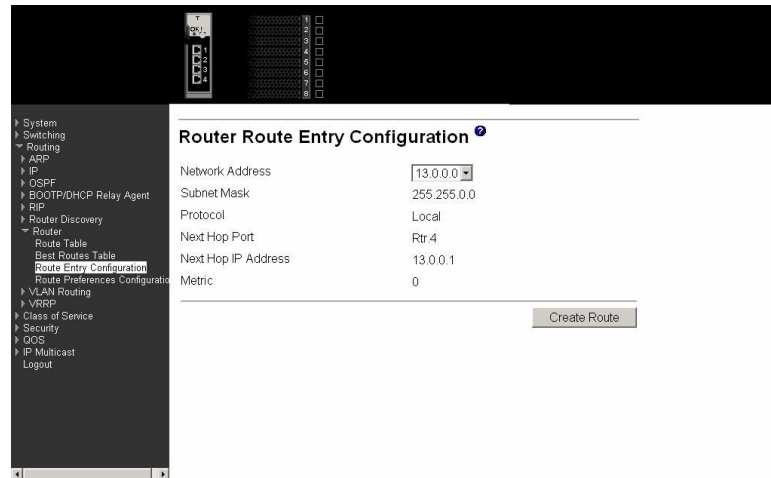
Next Hop IP Address

The IP address of the next router in the path towards the destination. The next router will always be one of the adjacent neighbors or the IP address of the local interface for a directly attached network.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

Route entry configuration

Use this panel to display the route to the specified network address. If create is selected on this panel you can manually enter the route information.



Network Address

Specifies the IP route prefix for the destination. To create a route a valid routing interface must exist and the next hop IP address must be on the same network as the routing interface. Routing interfaces are created on the IP Interface Configuration page. Valid next hop IP addresses can be viewed on the Route Table screen.

Subnet Mask Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.

Protocol This field tells which protocol created the specified route. The possibilities are one of the following:

- Local
- Static
- Default
- MPLS
- OSPF Intra
- OSPF Inter
- OSPF Type-1
- OSPF Type-2
- RIP
- BGP4

Next Hop Port

The outgoing router interface to use when forwarding traffic to the destination.

Next Hop IP Address

The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router will always be one of the adjacent neighbors or the IP address of the local interface for a directly attached network. When

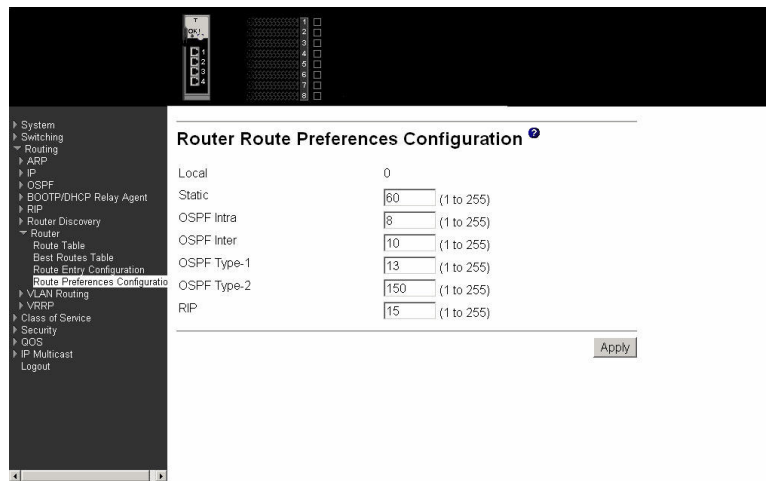
creating a route, the next hop IP must be on the same network as the routing interface. Valid next hop IP addresses can be seen on the Route Table page.

Metric Administrative cost of the path to the destination. If no value is entered, default is 1. The range is 0 - 255.

Click the Create Route button to go to a separate page where a route can be created.

Route preferences configuration

Use this panel to configure the default preference for each protocol (e.g. 60 for static routes, 170 for BGP). These values are arbitrary in the range of route metrics, but are independent of route metrics. Most routing protocols use a route metric to determine the shortest path known to the protocol, independent of any other protocol. The best route to a destination is chosen by selecting the route with the lowest preference value. When there are multiple routes to a destination, the preference values are used to determine the preferred route. If there is still a tie, the route with the best route metric will be chosen. To avoid problems with mismatched metrics (i.e. RIP and OSPF metrics are not directly comparable) you should configure different preference values for each of the protocols.



Local The next hop is the final destination

Static Manually configured route

OSPF Intra OSPF intra-area route

OSPF Inter OSPF inter-area route

OSPF Type-1 OSPF external route

OSPF Type-2 OSPF route learned from another protocol

RIP RIP route

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

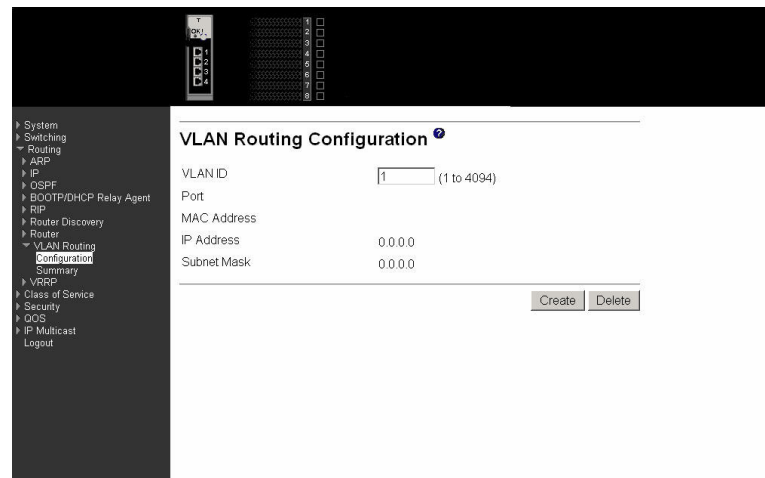
VLAN routing

This menu provides access to the following VLAN entry panels:

- Configuration
- Summary

Configuration

Use this panel to configure a new VLAN. Enter a new VLAN ID in the field labeled VLAN ID. Click on the Create button. The page will be updated to display the interface and MAC address assigned to this new VLAN. The IP address and Subnet Mask fields will be 0.0.0.0. Note the interface assigned to the VLAN. Use the index pane to change to the IP Interface Configuration page. Select the interface assigned to the VLAN. The IP address and Subnet Mask fields will be 0.0.0.0. Enter the IP address and subnet mask for the VLAN. Select the Apply button. Change back to the VLAN Routing Summary page. The new VLAN should appear in the table with the correct IP address and subnet mask assigned.



VLAN ID Enter the ID of a VLAN you want to configure for VLAN Routing. Initially, the field will display the ID of the first VLAN. After you enter a new VLAN ID and click on the Create button the non-configurable data will be displayed. See below for detailed instructions on how to use that data to complete the configuration of the VLAN.

Port The interface identification assigned to the routing VLAN.

MAC Address The MAC Address assigned to the VLAN Routing Interface.

IP Address The configured IP address of the VLAN Routing Interface. This will be 0.0.0.0 when the VLAN Routing Interface is first configured and must be entered on the IP Interface Configuration page.

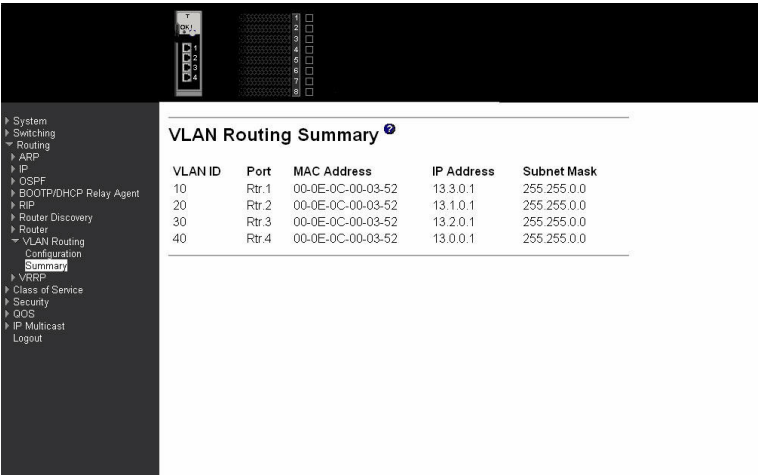
Subnet Mask The configured Subnet Mask of the VLAN Routing Interface. This will be 0.0.0.0 when the VLAN Routing Interface is first configured and must be entered on the IP Interface Configuration page.

Click the Create button to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

Click the Delete button to remove the VLAN Routing Interface SPECIFIED in the VLAN ID input field from the router configuration.

Summary

Use this panel to display the configuration information for a VLAN.



VLAN Routing Summary

VLAN ID	Port	MAC Address	IP Address	Subnet Mask
10	Rtr.1	00-0E-0C-00-03-52	13.3.0.1	255.255.0.0
20	Rtr.2	00-0E-0C-00-03-52	13.1.0.1	255.255.0.0
30	Rtr.3	00-0E-0C-00-03-52	13.2.0.1	255.255.0.0
40	Rtr.4	00-0E-0C-00-03-52	13.0.0.1	255.255.0.0

VLAN ID The ID of the VLAN whose data is displayed in the current table row.

Port The port assigned to the VLAN Routing Interface.

MAC Address The MAC Address assigned to the VLAN Routing Interface.

IP Address The configured IP address of the VLAN Routing Interface. This will be 0.0.0.0 when the VLAN Routing Interface is first configured and must be entered on the IP Interface Configuration page.

Subnet Mask The configured Subnet Mask of the VLAN Routing Interface. This will be 0.0.0.0 when the VLAN Routing Interface is first configured and must be entered on the IP Interface Configuration page.

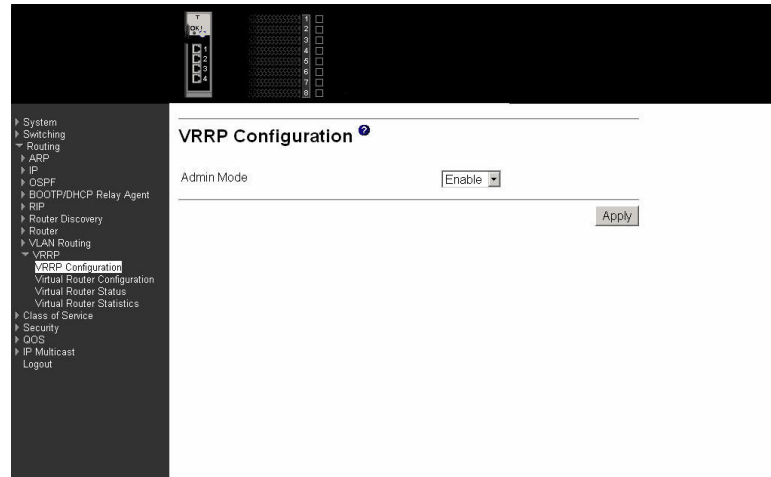
VRRP

This menu provides access to the following Virtual Router Redundancy Protocol (VRRP) entry and display panels:

- VRRP configuration
- Virtual router configuration
- Virtual router status
- Virtual router statistics

VRRP configuration

Use this panel to enable or disable Virtual Router support for the switch.

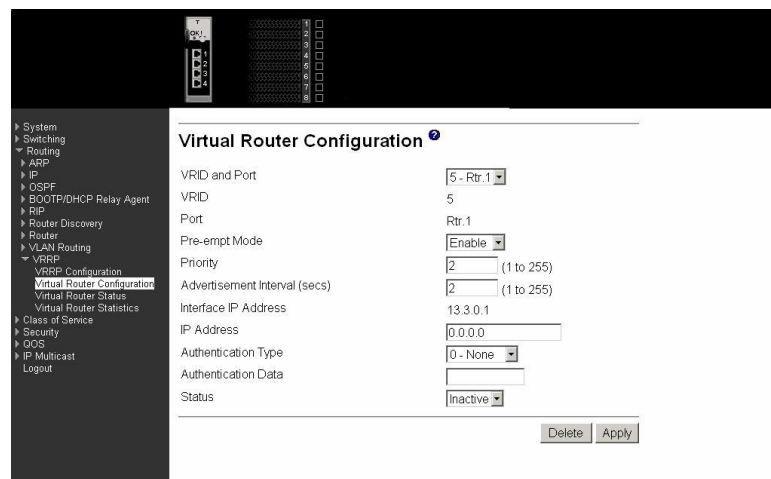


Admin Mode This sets the administrative status of VRRP in the router to active or inactive. Select Enable or Disable from the pull-down menu. The default is Disable.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Virtual router configuration

Use this panel to configure a new Virtual Router, or to display or configure the parameters for an existing Virtual Router.



VRID and Port

Select Create from the pull-down menu to configure a new Virtual Router, or select one of the existing Virtual Routers, listed by interface number and Virtual Router ID (VRID).

VRID

This field is only configurable if you are creating new Virtual Router, in which case enter a VRID number between 1 and 255.

Port This field is only configurable if you are creating a new Virtual Router, in which case select the port for the new Virtual Router from the pull-down menu.

Preempt Mode

Select Enable or Disable from the pull-down menu. If you select Enable a backup router will preempt the master router if it has a priority greater than the master Virtual Router's priority, provided the master is not the owner of the Virtual Router IP address. The default is Enable.

Priority

Enter the priority value to be used by the VRRP router in the election for the master Virtual Router. Enter a number between 1 and 255. The default value is 100.

Advertisement Interval (secs)

Enter the time, in seconds, between the transmission of advertisement packets by this Virtual Router. Enter a number between 1 and 255. The default value is 1 second.

Interface IP Address

Indicates the IP address associated with the selected interface.

IP Address

Enter the IP address associated with the Virtual Router. The default is 0.0.0.0.

Authentication Type

Select the type of authentication for the Virtual Router from the pull-down menu. The default is None. The choices are:

- 0-None - No authentication will be performed.
- 1-Simple - Authentication will be performed using a text password.

Authentication Data

If you selected simple authentication, enter the password.

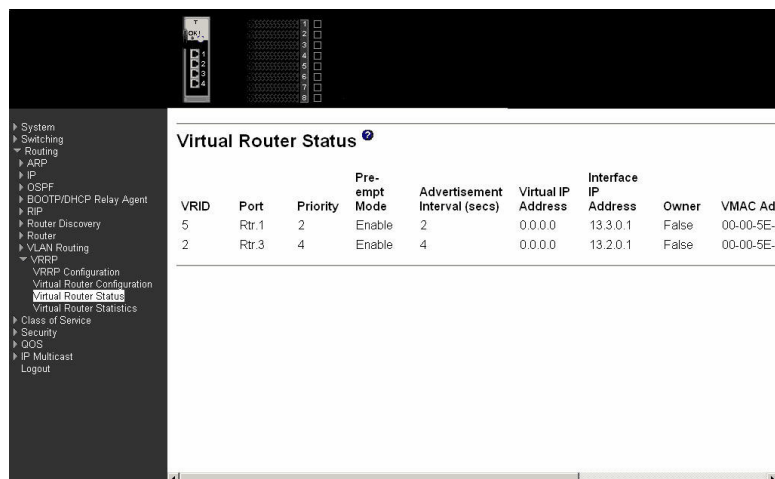
Status

Select active or inactive from the pull-down menu to start or stop the operation of the Virtual Router. The default is inactive.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Virtual router status

Use this panel to display information associated with Virtual Routers that have been configured on the switch.



VRID	Port	Priority	Pre-empt Mode	Advertisement Interval (secs)	Virtual IP Address	Interface IP Address	Owner	VMAC Address
5	Rtr.1	2	Enable	2	0.0.0.0	13.3.0.1	False	00-00-5E-00-00-00
2	Rtr.3	4	Enable	4	0.0.0.0	13.2.0.1	False	00-00-5E-00-00-00

VRID The Virtual Router Identifier.

Port Indicates the interface associated with the VRID.

Priority The priority value to be used by the VRRP router in the election for the master Virtual Router.

Preempt Mode Enable indicates that if the Virtual Router is a backup router it will preempt the master router if it has a priority greater than the master Virtual Router's priority, provided the master is not the owner of the Virtual Router IP address. Disable indicates that if the Virtual Router is a backup router it will not preempt the master router even if its priority is greater.

Advertisement Interval (secs) The time, in seconds, between the transmission of advertisement packets by this Virtual Router.

Virtual IP Address The IP address associated with the Virtual Router.

Interface IP Address The actual IP address associated with the interface used by the Virtual Router.

Owner Set to True if the Virtual IP address and the Interface IP address are the same, otherwise set to False. If this parameter is set to True, the Virtual Router is the owner of the Virtual IP address, and will always win an election for master router when it is active.

VMAC Address The virtual MAC Address associated with the Virtual Router, composed of a 24 bit organizationally unique identifier, the 16 bit constant identifying the VRRP address block and the 8 bit VRID.

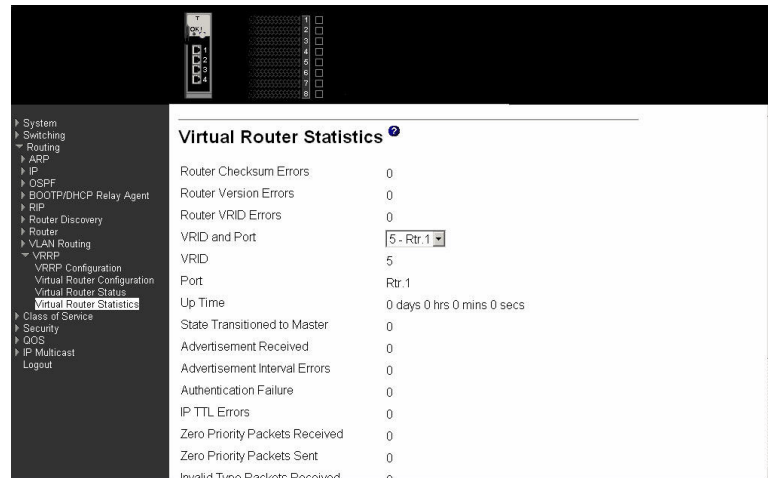
Auth Type The type of authentication in use for the Virtual Router. Options are None and Simple.

State	The current state of the Virtual Router. Options are Initialize, Master and Backup.
Status	The current status of the Virtual Router. Options are Inactive and Active.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

Virtual router statistics

Use this panel to display statistics associated with the Virtual Routers.



Router Checksum Errors

The total number of VRRP packets received with an invalid VRRP checksum value.

Router Version Errors

The total number of VRRP packets received with an unknown or unsupported version number.

Router VRID Errors

The total number of VRRP packets received with an invalid VRID for this Virtual Router.

VRID and Port

Select the existing Virtual Router, listed by interface number and VRID, for which you want to display statistical information.

VRID

The VRID for the selected Virtual Router.

Port

The port for the selected Virtual Router.

Up Time

The time, in days, hours, minutes and seconds, that has elapsed since the Virtual Router transitioned to the initialized state.

State Transitioned to Master

The total number of times that this Virtual Router's state has transitioned to Master.

Advertisement Received

The total number of VRRP advertisements received by this Virtual Router.

Advertisement Interval Errors

The total number of VRRP advertisement packets received for which the advertisement interval was different than the one configured for the local Virtual Router.

Authentication Failure

The total number of VRRP packets received that did not pass the authentication check.

IP TTL Errors The total number of VRRP packets received by the Virtual Router with IP TTL (Time-To-Live) not equal to 255.

Zero Priority Packets Received

The total number of VRRP packets received by the Virtual Router with a priority of '0'.

Zero Priority Packets Sent

The total number of VRRP packets sent by the Virtual Router with a priority of '0'.

Invalid Type Packets Received

The number of VRRP packets received by the Virtual Router with an invalid value in the type field.

Address List Errors

The total number of packets received for which the address list does not match the locally configured list for the Virtual Router.

Invalid Authentication Type

The total number of packets received with an unknown authentication type.

Authentication Type Mismatch

The total number of packets received with an authentication type different to the locally configured authentication method.

Packet Length Errors

The total number of packets received with a packet length less than the length of the VRRP header.

Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

Class of service

This menu contains one option – 802.1p priority mapping.

802.1p priority mapping

Use this panel to specify how IEEE 802.1p priority classes are to be mapped to the switch's internal traffic classes.

User Priority	Traffic Class
0	2
1	0
2	1
3	3
4	4
5	5
6	6
7	7

Apply

User Priority

The 802.1p user priority to be mapped.

Traffic Class

Use the pull-down menus to select the internal traffic class for each user priority.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Security

This menu describes the web menus used to configure and manage the security features of the Ethernet switch module. These features include:

- Port access control
- RADIUS
- Secure HTTP
- Secure shell

Port access control

The Port Access Control menu provides access to configuration, status and summary screens:

- Configuration
- Port configuration
- Port status
- Port summary
- Statistics
- Login
- Port access privileges
- Port access summary

Configuration

Use this panel to enable or disable authentication support on the switch. In disabled mode, the IEEE 802.1X configuration is retained and can be changed, but it is not activated.



Administrative Mode

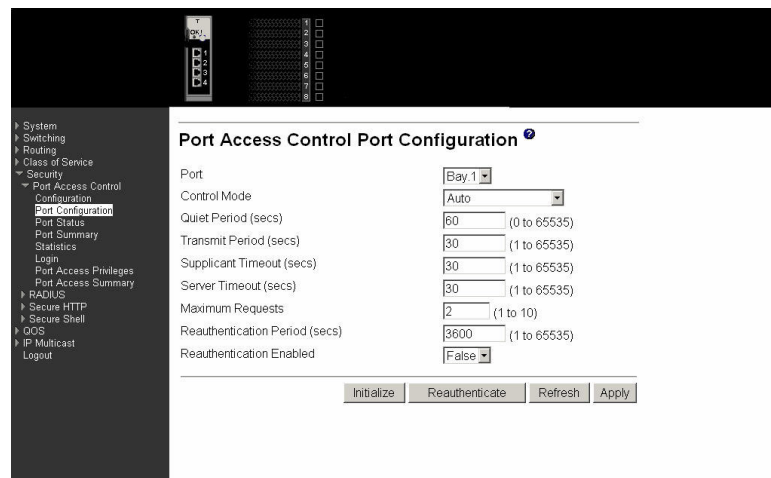
Lists the two options for administrative mode: Enable and Disable. The default value is Disable.

Click the Cancel button to reset the page to display the administrative mode that is currently configured by the selected unit.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Port configuration

Use this panel to begin the initialization or the reauthentication sequence on the selected port.



Port

Select the port to be configured. When the selection is changed, a screen refresh will occur causing all fields to be updated for the newly selected port. All physical interfaces are valid.

Control Mode Lists the options for control mode. The control mode is only set if the port is in Link Up status. The options are:

Force Unauthorized

The authenticator Port Access Entity (PAE) unconditionally sets the controlled port to unauthorized.

Force Authorized

The authenticator PAE unconditionally sets the controlled port to authorized mode.

Auto The authenticator PAE sets the controlled port mode to reflect the result of the authentication exchanges between the supplicant, authenticator and authentication server.

Quiet Period (secs)

Configures the quiet period for the selected port. This command sets the value, in seconds, of the timer used by the authenticator state machine on this port to define periods of time during which it will not attempt to acquire a supplicant. The quiet period is the period for which the authenticator does not attempt to acquire a supplicant after a failed authentication exchange with the supplicant. The quiet period range is 0 to 65535. A quiet period value of 0 means that the authenticator state machine will never acquire a supplicant. The default value is 60.

Transmit Period (secs)

Configures the transmit period for the selected port. The transmit period is the value, in seconds, of the timer used by the authenticator state machine on the specified port to determine when to send an Extensible Authentication Protocol Over LAN (EAPOL) EAP Request/Identity frame to the supplicant. The transmit period range is 1 to 65535. The default value is 30.

Supplicant Timeout (secs)

Specify the supplicant timeout for the selected port. The supplicant timeout is the value, in seconds, of the timer used by the authenticator state machine on this port to timeout the supplicant. The supplicant timeout range is 1 to 65535. The default value is 30. Changing the value will not change the configuration until the Apply button is clicked.

Server Timeout (secs)

Specify the server timeout for the selected port. The server timeout is the value, in seconds, of the timer used by the authenticator on this port to timeout the authentication server. The server timeout range is 1 to 65535. The default value is 30.

Maximum Requests

Specify the maximum requests for the selected port. The maximum requests value is the maximum number of times the authenticator state machine on this port will retransmit an EAPOL EAP Request/Identity before timing out the supplicant. The maximum requests range is 1 to 10. The default value is 2.

Reauthentication Period (secs)

Specify the reauthentication period for the selected port. The reauthentication period is the value, in seconds, of the timer used by the authenticator state machine on this port to determine when

reauthentication of the supplicant takes place. The reauthentication period range is 1 to 65535. The default value is 3600.

Reauthentication Enabled

Enable or Disable the reauthentication of the supplicant for the specified port. If the value true is selected reauthentication will occur. Otherwise, reauthentication will not be allowed. The default value is false.

Click the Initialize button to begin the initialization sequence on the selected port. This button is only selectable if the control mode is auto. If the button is not selectable, it will be grayed out. When you click this button the action is immediate and you will not need to press the Apply button for the action to occur.

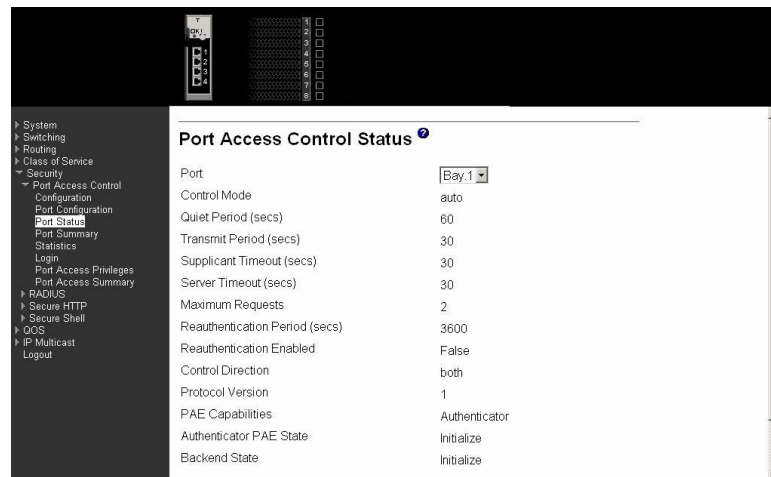
Click the Reauthenticate button to begin the reauthentication sequence on the selected port. This button is only selectable if the control mode is auto. If the button is not selectable, it will be grayed out. When you click this button the action is immediate and you will not need to press the Apply button for the action to occur.

Click the Refresh button to update the information on the page.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Port status

This panel displays the details of the IEEE 802.1X configuration parameters for the specified port.



Port

Select the port whose information will be displayed. When the selection is changed, a screen refresh will occur causing all fields to be updated for the newly selected port. All physical interfaces are valid.

Control Mode

Displays the configured control mode for the specified port. Options are:

force unauthorized

The authenticator port access entity (PAE) unconditionally sets the controlled port to unauthorized.

force authorized

The authenticator PAE unconditionally sets the controlled port(s) to authorized mode.

auto

The authenticator PAE sets the controlled port(s) mode to reflect the result of the authentication exchanges between the supplicant, authenticator and authentication server.

Quiet Period (secs)

This field displays the configured quiet period for the selected port. This quiet period is the value, in seconds, of the timer used by the authenticator state machine on this port to define periods of time in which it will not attempt to acquire a supplicant. The quiet period is the period for which the authenticator does not attempt to acquire a supplicant after a failed authentication exchange with the supplicant. The quiet period range is 0 to 65535.

Transmit Period (secs)

Displays the configured transmit period for the selected port. The transmit period is the value, in seconds, of the timer used by the authenticator state machine on the specified port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant. The transmit period range is 1 to 65535.

Supplicant Timeout (secs)

Displays the configured supplicant timeout for the selected port. The supplicant timeout is the value, in seconds, of the timer used by the authenticator state machine on this port to timeout the supplicant. The supplicant timeout range is 1 to 65535.

Server Timeout (secs)

Displays the configured server timeout for the selected port. The server timeout is the value, in seconds, of the timer used by the authenticator on this port to timeout the authentication server. The server timeout range is 1 to 65535.

Maximum Requests

Displays the configured maximum requests for the selected port. The maximum requests value is the maximum number of times the authenticator state machine on this port will retransmit an EAPOL EAP Request/Identity before timing out the supplicant. The maximum requests value range is 1 to 10.

Reauthentication Period (secs)

Displays the configured reauthentication period for the selected port. The reauthentication period is the value, in seconds, of the timer used by the authenticator state machine on this port to determine when reauthentication of the supplicant takes place. The reauthentication period value range is 1 to 65535.

Reauthentication Enabled

Indicates whether reauthentication is enabled on the selected port. If you select the value true reauthentication will occur. Otherwise, reauthentication will not be allowed.

Control Direction

Displays the control direction for the specified port. The control direction dictates the degree to which protocol exchanges take place between supplicant and authenticator. This affects whether the controlled port exerts control over communication in both

directions (disabling both incoming and outgoing frames) or just incoming (disabling only the reception of incoming frames). This field is not configurable on some platforms.

Protocol Version

Displays the protocol version associated with the selected port. The only possible value is 1, corresponding to the first version of the IEE 802.1X specification.

PAE Capabilities

Displays the PAE functionality of the selected port. Possible values are Authenticator or Supplicant.

Authenticator PAE State

Displays the current state of the authenticator PAE state machine. Possible values are:

- Initialize
- Disconnected
- Connecting
- Authenticating
- Authenticated
- Aborting
- Held
- Force Authorized
- Force Unauthorized

Backend State

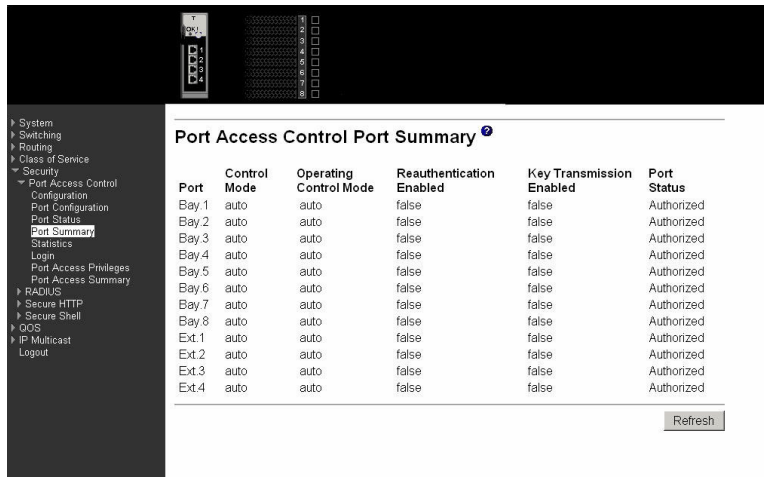
Displays the current state of the backend authentication state machine. Possible values are:

- Request
- Response
- Success
- Fail
- Timeout
- Initialize
- Idle

Click the Refresh button to update the information on the page.

Port summary

This panel displays a summary of the IEEE 802.1X configuration parameters for all switch ports.



Port	Control Mode	Operating Control Mode	Reauthentication Enabled	Key Transmission Enabled	Port Status
Bay.1	auto	auto	false	false	Authorized
Bay.2	auto	auto	false	false	Authorized
Bay.3	auto	auto	false	false	Authorized
Bay.4	auto	auto	false	false	Authorized
Bay.5	auto	auto	false	false	Authorized
Bay.6	auto	auto	false	false	Authorized
Bay.7	auto	auto	false	false	Authorized
Bay.8	auto	auto	false	false	Authorized
Ext.1	auto	auto	false	false	Authorized
Ext.2	auto	auto	false	false	Authorized
Ext.3	auto	auto	false	false	Authorized
Ext.4	auto	auto	false	false	Authorized

Port The port whose settings are displayed in the associated table row.

Control Mode Displays the configured control mode for the port. Possible values are:

Force Unauthorized

The authenticator port access entity (PAE) unconditionally sets the controlled port to unauthorized.

Force Authorized

The authenticator PAE unconditionally sets the controlled port(s) to authorized mode.

Auto The authenticator PAE sets the controlled port(s) mode to reflect the result of the authentication exchanges between the supplicant, authenticator and authentication server.

Operating Control Mode

Displays the control mode under which the port is actually operating. Possible values are:

Force Unauthorized

The authenticator PAE unconditionally sets the controlled port to unauthorized.

Force Authorized

The authenticator PAE unconditionally sets the controlled port(s) to authorized mode.

Auto The authenticator PAE sets the controlled port(s) mode to reflect the result of the authentication exchanges between the supplicant, authenticator and authentication server.

Reauthentication Enabled

Displays whether reauthentication of the supplicant for the specified port is allowed. The possible values are true and false. If the value is true reauthentication will occur. Otherwise, reauthentication will not be allowed.

Key Transmission Enabled

Displays whether key transmission is enabled on the selected port.

The possible values are true and false. If the value is true, keys will be transmitted to the supplicant. Otherwise, keys will not be transmitted.

Port Status Displays the authorization status of the specified port. The possible values are Authorized and Unauthorized.

Click the Refresh button to update the information on the page.

Statistics

This panel displays the IEEE 802.1X statistics for the specified port.

Port	Bay 1
EAPOL Frames Received	0
EAPOL Frames Transmitted	0
EAPOL Start Frames Received	0
EAPOL Logoff Frames Received	0
Last EAPOL Frame Version	0
Last EAPOL Frame Source	00-00-00-00-00-00
EAP Response/ID Frames Received	0
EAP Response Frames Received	0
EAP Request/ID Frames Transmitted	0
EAP Request Frames Transmitted	0
Invalid EAPOL Frames Received	0
EAPOL Length Error Frames Received	0

Port Select the port whose information is to be displayed. When the selection is changed, a screen refresh occurs causing all fields to be updated for the newly selected port. All physical interfaces are valid.

EAPOL Frames Received

The number of valid EAPOL frames of any type that have been received by this authenticator.

EAPOL Frames Transmitted

The number of EAPOL frames of any type that have been transmitted by this authenticator.

EAPOL Start Frames Received

The number of EAPOL start frames that have been received by this authenticator.

EAPOL Logoff Frames Received

The number of EAPOL logoff frames that have been received by this authenticator.

Last EAPOL Frame Version

The protocol version number carried in the most recently received EAPOL frame.

Last EAPOL Frame Source

The source MAC address carried in the most recently received EAPOL frame.

EAP Response/ID Frames Received

The number of EAP response/identity frames that have been received by this authenticator.

EAP Response Frames Received

The number of valid EAP response frames (other than response/identity frames) that have been received by this authenticator.

EAP Request/ID Frames Transmitted

The number of EAP request/identity frames that have been transmitted by this authenticator.

EAP Request Frames Transmitted

The number of EAP request frames (other than request/identity frames) that have been transmitted by this authenticator.

Invalid EAPOL Frames Received

The number of EAPOL frames that have been received by this authenticator with an invalid length.

EAP Length Error Frames Received

The number of EAPOL frames that have been received by this authenticator in which the frame type is not recognized.

Click the Refresh button to update the information on the page.

Click the Clear All button to reset all statistics for all ports to 0. There is no confirmation prompt. When this button is clicked, the statistics are immediately cleared.

Click the Clear button to reset the statistics for the selected port. There is no confirmation prompt. When this button is clicked, the statistics are immediately cleared.

Login

Use this panel to assign a selected authentication login list to a selected user for port security. Both user and the login list must already be configured.



Users Select the user name to be configured.

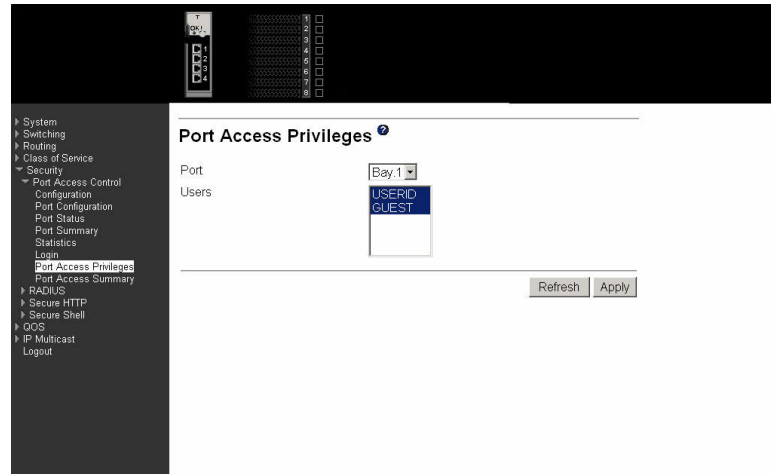
Login Selects the login list to be associated with the selected user. All configured login lists are displayed.

Click the Refresh button to update the information on the page.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch.

Port access privileges

Use this panel to add the specified user to the list of users with access to the specified port(s). By default, a user is given access to all ports.



Port Select a port from the pull-down menu. All physical ports are available for this selection.

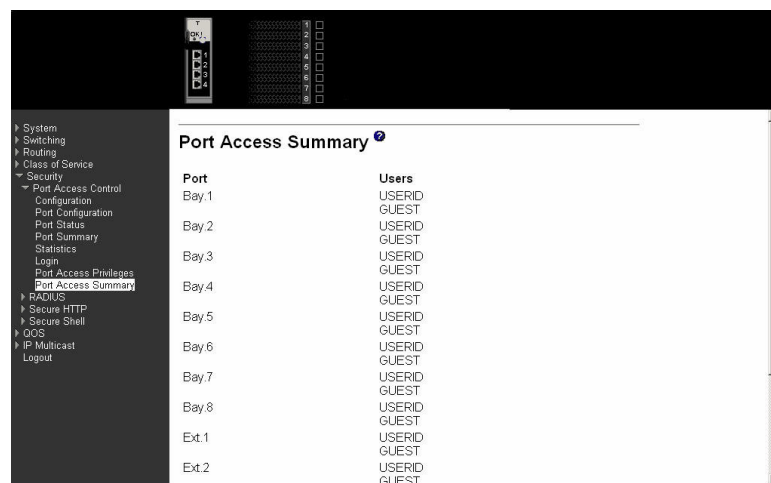
Users Select the users that can have access to the selected port or ports.

Click the Refresh button to update the information on the page.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch.

Port access summary

This panel displays IEEE 802.1X port security information about locally configured users.



Port The port whose information is displayed on this line.

Users The locally configured users with access to the specified port.

Click the Refresh button to update the information on the page.

RADIUS

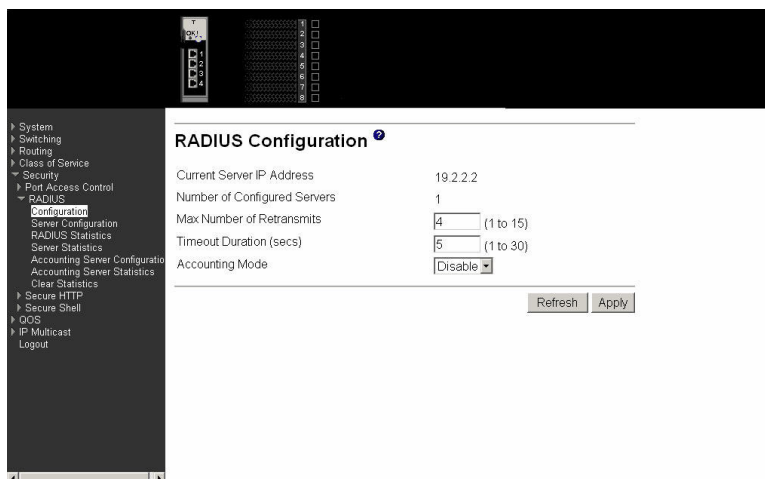
The Remote Authentication Dial-in User Service (RADIUS) menu provides access to the following panels:

- Configuration
- Server configuration
- RADIUS statistics
- Server statistics
- Accounting server configuration
- Accounting server statistics
- Clear statistics

Configuration

Use this panel to configure RADIUS parameters for the switch.

Consideration should be given to the maximum delay time when configuring RADIUS maximum retransmit and timeout values. If multiple RADIUS servers are configured, the maximum retransmit value on each is exhausted before the next server is attempted. A retransmit does not occur until the configured timeout value on that server has passed without a response. Therefore, the maximum delay in receiving a response from the RADIUS application equals the sum of maximum retransmit times the timeout for all configured servers. If the RADIUS request was generated by a user login attempt, all user interfaces are blocked until the RADIUS application returns a response.



Current Server IP Address

The IP address of the current server. This field is blank if no servers are configured.

Number of Configured Servers

The number of RADIUS servers that have been configured. The range for this value is 0 to 3.

Max Number of Retransmits

The value of the maximum number of times a request packet is retransmitted. The valid range is 1 - 15.

Timeout Duration (secs)

The timeout value, in seconds, for request retransmissions. The valid range is 1 - 30.

Accounting Mode

Select whether the RADIUS accounting mode is Enabled or Disabled.

Click the Refresh button to update the information on the page.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch.

Server configuration

Use this panel to configure the IP address of a RADIUS server. Up to three servers can be configured for each RADIUS client.

The screenshot displays the 'RADIUS Server Configuration' web page. On the left is a dark navigation menu with a tree structure. The main area is white and contains the following configuration fields:

- RADIUS Server IP Address:** 19.2.2.2
- Port:** 1812 (0 to 65535)
- Secret:** [Empty text box] Apply
- Primary Server:** No
- Message Authenticator:** Enable
- Secret Configured:** No
- Current:** Yes

At the bottom right of the configuration area are three buttons: 'Remove', 'Refresh', and 'Apply'.

RADIUS Server IP Address

Select the RADIUS Server to be configured. Select Add to add a new server.

Port The User Datagram Protocol (UDP) port used by this server. The valid range is 0 - 65535.

Secret

The shared secret for this server. The data entered in this field will not be displayed.

Apply The Secret is applied only if this box is checked. If the box is not checked, anything entered in the Secret field has no affect and is not retained. This field is only displayed if the user has Read/Write access.

Primary Server

Sets the selected server to be the Primary or Secondary server.

Message Authenticator

Enable or Disable the message authenticator attribute for the selected server.

Secret Configured

Indicates whether the shared secret for this server has been configured.

Current

Indicates whether this server is currently in use as the authentication server.

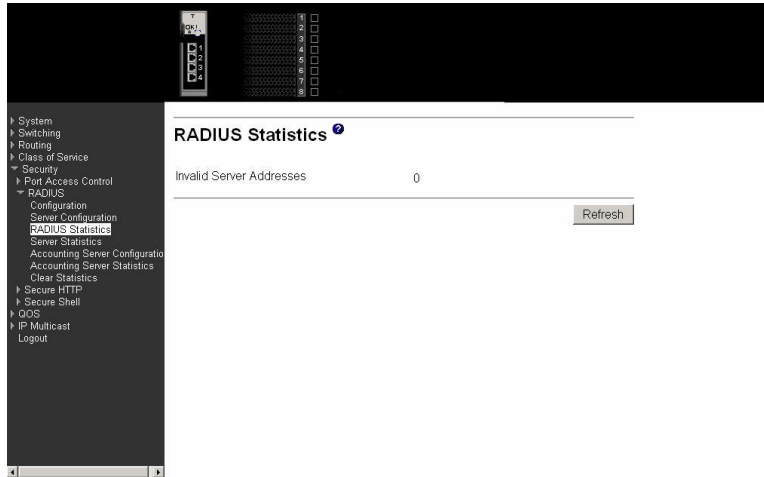
Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Remove button to remove the selected server from the configuration. This button is only available to Read/Write users. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Refresh button to update the information on the page.

RADIUS statistics

This panel displays RADIUS statistics for the switch that are not associated with a specific server or accounting server.



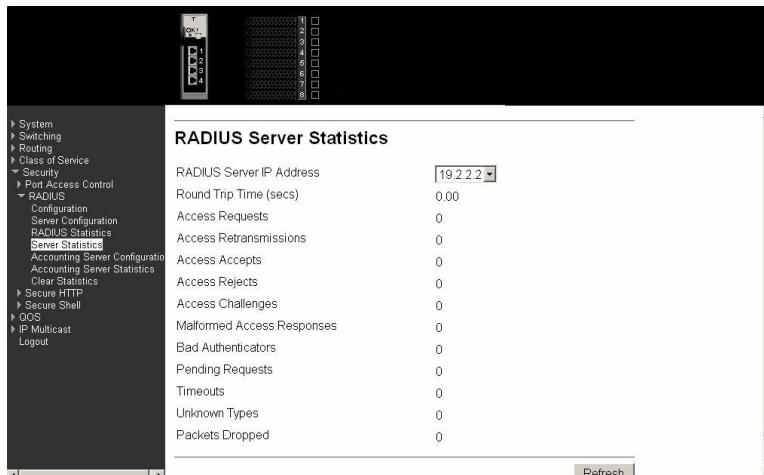
Invalid Server Addresses

The number of RADIUS Access-Response packets received from unknown addresses.

Click the Refresh button to update the information on the page.

Server statistics

This panel displays the statistics for a configured RADIUS server.



RADIUS Server IP Address

Select the IP address of the server whose information is to be displayed.

Round Trip Time (secs)

The time, in seconds, between the most recent RADIUS Access-Reply/Access-Challenge and the matching Access-Request from this RADIUS server.

Access Requests

The number of RADIUS Access-Request packets sent to this server, not including retransmissions.

Access Retransmissions

The number of RADIUS Access-Request packets retransmitted to this server.

Access Accepts

The number of RADIUS Access-Accept packets, both valid and invalid, received from this server.

Access Rejects

The number of RADIUS Access-Reject packets, both valid and invalid, received from this server.

Access Challenges

The number of RADIUS Access-Challenge packets, both valid and invalid, received from this server.

Malformed Access Responses

The number of malformed RADIUS Access-Response packets received from this server, including packets with invalid length but not including packets with bad authenticators, bad signature attributes or unknown types.

Bad Authenticators

The number of RADIUS Access-Response packets received from this server, including packets with invalid authenticators or signature attributes.

Pending Requests

The number of RADIUS Access-Request packets sent to this server that have not yet timed out or received a response.

Timeouts

The number of RADIUS packets sent to this server that have timed out.

Unknown Types

The number of RADIUS packets of unknown type received from this server.

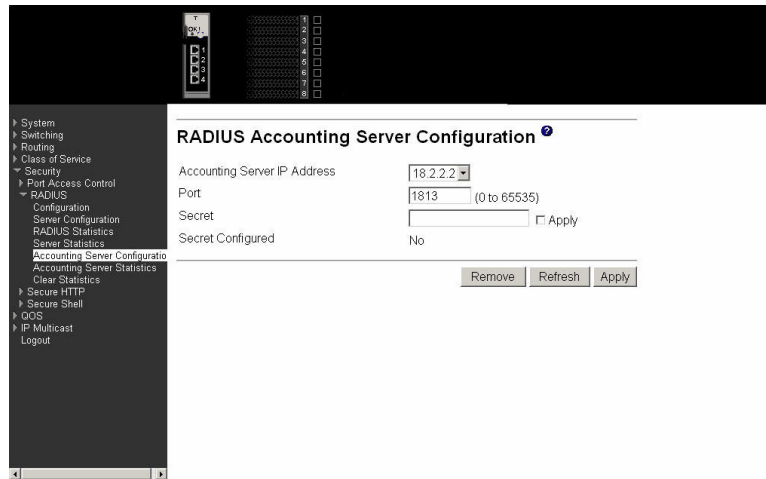
Packets Dropped

The number of RADIUS packets received from this server dropped for a reason not otherwise included in this list.

Click the Refresh button to update the information on the page.

Accounting server configuration

Use this panel to configure the IP address of the accounting server. Only a single accounting server can be configured.



Accounting Server IP Address

Select Add to configure an accounting server or the address of an already configured server.

Port

Specifies the UDP Port to be used by the accounting server. The valid range is 0 - 65535. If the user has Read-only access, the value is displayed but it cannot be changed.

Secret

Specifies the shared secret to use with the specified accounting server. This field is only displayed if the user has Read/Write access.

Apply

The Secret is applied only if this box is checked. If the box is not checked, anything entered in the Secret field has no affect and is not retained. This field is only displayed if the user has Read/Write access.

Secret Configured

Indicates whether the shared secret for this accounting server has been configured.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Remove button to remove the selected accounting server from the configuration. This button is only available to Read/Write users. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Refresh button to update the information on the page.

Accounting server statistics

This panel displays the RADIUS statistics for the accounting server.

RADIUS Accounting Server Statistics	
Accounting Server IP Address	18.2.2.2
Round Trip Time (secs)	0.00
Accounting Requests	0
Accounting Retransmissions	0
Accounting Responses	0
Malformed Accounting Responses	0
Bad Authenticators	0
Pending Requests	0
Timeouts	0
Unknown Types	0
Packets Dropped	0

Refresh

Accounting Server IP Address

Identifies the accounting server associated with the statistics.

Round Trip Time (secs)

Displays the time interval, in hundredths of a second, between the most recent Accounting-Response and the Accounting-Request that matched it from this RADIUS accounting server.

Accounting Requests

Displays the number of RADIUS Accounting-Request packets sent not including retransmissions.

Accounting Retransmissions

Displays the number of RADIUS Accounting-Request packets retransmitted to this RADIUS accounting server.

Accounting Responses

Displays the number of RADIUS packets received on the accounting port from this server.

Malformed Accounting Responses

Displays the number of malformed RADIUS Accounting-Response packets received from this server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.

Bad Authenticators

Displays the number of RADIUS Accounting-Response packets that contained invalid authenticators received from this accounting server.

Pending Requests

Displays the number of RADIUS Accounting-Request packets sent to this server that have not yet timed out or received a response.

Timeouts

Displays the number of accounting timeouts involving this server.

Unknown Types

Displays the number of RADIUS packets of unknown type that were received from this server on the accounting port.

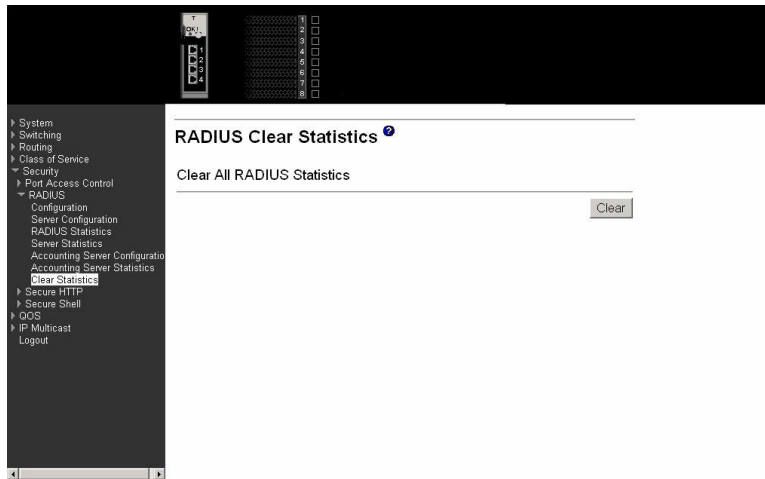
Packets Dropped

Displays the number of RADIUS packets that were received from this server on the accounting port and dropped for some other reason.

Click the Refresh button to update the information on the page.

Clear statistics

Use this panel to reset all RADIUS statistics for the switch. Click the Clear button to clear the accounting server, authentication server and RADIUS statistics.

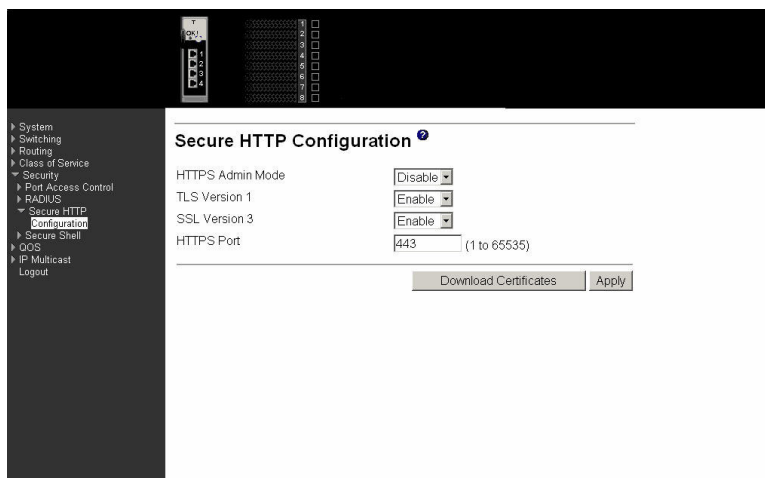


Secure HTTP

The Secure Sockets Layer (SSL) encryption protocol provides a means of abstracting an encrypted connection between two stations, allowing HTTP to operate securely on an open network. This menu provides access to the Secure HTTP configuration panel.

Configuration

Use this panel to configure Secure HTTP variables.



HTTPS Admin Mode

Select Enable or Disable to turn the Administrative Mode of Secure HTTP on or off. The currently configured value is shown when the web page is displayed. The default value is Disable.

TLS Version 1

Select Enable or Disable to turn Transport Layer Security (TLS) Version 1.0 on or off. The currently configured value is shown when the web page is displayed. This field cannot be changed while HTTPS Admin Mode is enabled. The default value is Enable.

SSL Version 3

Select Enable or Disable to turn SSL Version 3.0 on or off. The currently configured value is shown when the web page is displayed. This field cannot be changed while HTTPS Admin Mode is enabled. The default value is Enable.

HTTPS Port

Specify the HTTPS Port Number. The value must be in the range of 1 to 65535. Port 443 is the default value. The currently configured value is shown when the web page is displayed.

Click the Download Certificates button to link to the File Transfer page to download SSL Certificate(s). Download is through the System Utilities menu.

Note: To download SSL Certificate files SSL must be administratively Disabled.

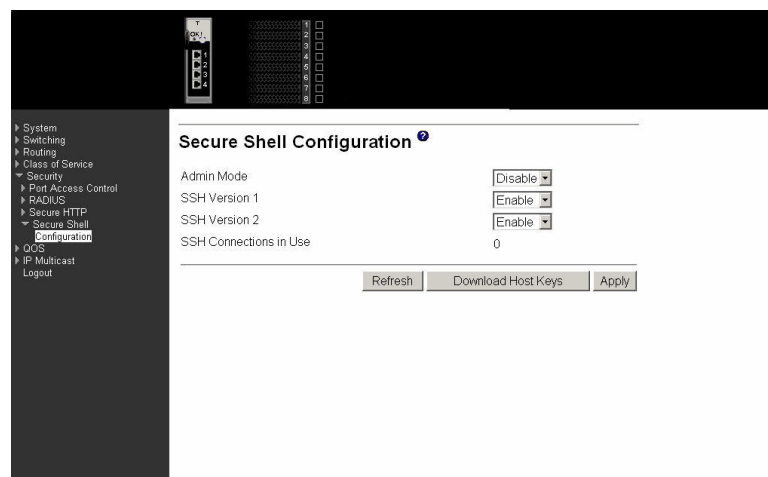
Click the Apply button to send the updated screen to the switch and have the changes take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Secure Shell

Secure Shell (SSH) is the standard encryption protocol used to provide a secure interactive login over a network. This Secure Shell menu provides access to the SSH configuration panel.

Configuration

Use this panel to configure SSH variables.



Admin Mode

Select Enable or Disable to turn the Administrative Mode of SSH on or off. The currently configured value is shown when the web page is displayed. The default value is Disable.

SSH Version 1

Select Enable or Disable to turn Protocol Level 1 for SSH on or off. The

currently configured value is shown when the web page is displayed. The default value is Enable. Either SSH Version 1 or Version 2 must be Enabled at all times.

SSH Version 2

Select Enable or Disable to turn Protocol Level 2 for SSH on or off. The currently configured value is shown when the web page is displayed. The default value is Enable. Either SSH Version 1 or Version 2 must be Enabled at all times.

SSH Connections in Use

Displays the number of SSH connections currently in use in the system.

Click the Download Host Keys button to link to the File Transfer page to download the Host Key(s).

Note: To download SSH key files SSH must be administratively Disabled and there can be no active SSH sessions

Click the Submit button to send the updated screen to the switch and have the changes take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Refresh button to display the current page with the latest settings and status.

QoS

This menu provides access to two Quality of Service (QoS) menus:

- Access Control Lists (ACLs)
- Bandwidth provisioning

Access Control Lists

An Access Control List (ACL) consists of a set of rules which are matched sequentially against a packet. When a packet meets the match criteria of a rule, the specified rule action (Permit/Deny) is taken and the additional rules are not checked for a match. You can specify the interfaces to which an ACL applies using the Configuration screen. You specify the rules for the ACL using the ACL Rule Configuration screen. ACL menu options are:

- Configuration
- Summary
- Rule configuration

Configuration

Use this panel to create an ACL.

The screenshot shows the 'ACL Configuration' web interface. On the left is a navigation menu with categories: System, Switching, Class of Service, Security, and OOS. Under OOS, there are sub-items: Access Control Lists, Configuration (highlighted), Summary, Rule Configuration, Bandwidth Provisioning, and Logout. The main panel is titled 'ACL Configuration' and contains the following elements:

- ACL:** A dropdown menu labeled 'Create New ACL'.
- ACL ID:** An input field containing '0' with '(1 to 100)' next to it.
- Ports:** A multi-selector list containing 'Bay.1', 'Bay.2', 'Bay.3', 'Bay.4', 'Bay.5', 'Bay.6', 'Bay.7', and 'Bay.8'.
- Direction:** A dropdown menu.
- Table:** A table with two columns: 'Current Size / Max Size'. The row shows 'ACL' with the value '0 / 100'.
- Apply:** A button at the bottom right.

ACL Make a selection from the pull-down menu. You can create a new ACL or update the configuration of an existing ACL.

ACL ID

ACL ID must be a whole number between 1 and 100.

Ports

This dynamic multi-selector lists all available valid interfaces for ACL mapping. All non-routing physical interfaces and interfaces participating in LAGs that are not already assigned to an ACL are listed. You can map an interface to one and only one ACL, but multiple interfaces can be assigned to one ACL. To apply an ACL to a LAG interface, the ACL must be applied to all member ports in a LAG.

Direction

Select the packet filtering direction for the ACL from the pull-down menu. Currently the only choice is Inbound. The packet direction for an ACL is the same for all affected interfaces.

Table Displays the current and maximum number of ACLs.

Current Size/Max Size

Displays the number of existing ACLs and the maximum number of configurable ACLs.

Click the Apply button to send the updated configuration to the switch. Configuration changes take effect immediately. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Delete button to remove the currently selected ACL from the switch configuration.

Summary

This panel displays a summary of all ACLs on the switch.

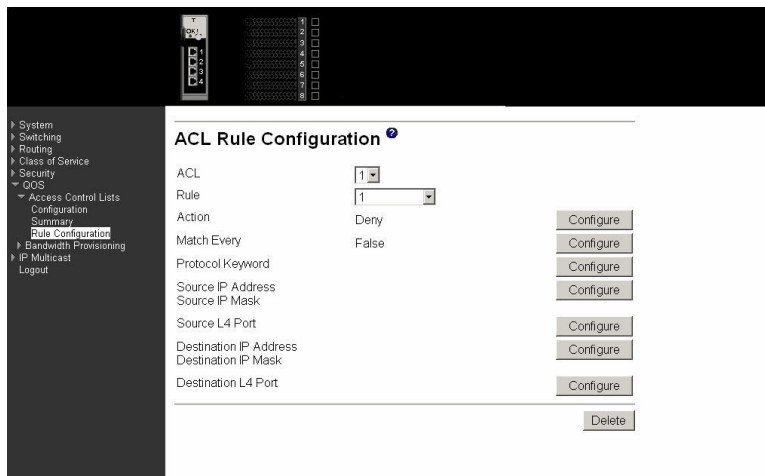


- ACL** The ACL identifier.
- Rules** The number of rules that are associated with this ACL.
- Ports** The interfaces that are associated with this ACL.
- Direction** The packet filtering direction for the ACL on the interface.

Click the Refresh button to update the screen with the latest information.

Rule configuration

This panel configures the rules associated with an ACL. When the screen first displays you will see the first four fields that are described below. If you select False as the Match Entry criteria and click Apply, the screen will be refreshed and you will see the remaining fields. Clicking one of the configure buttons shown on that screen will display a third screen allowing you to configure the match criterion you selected.



- ACL** Use the pull-down menu to select the ACL for which you want to create or update a rule.
- Rule** Enter a whole number in the range of 1 to 10 that will be used to identify the rule. An ACL can have up to 10 user-specified rules.

Action Specify what action should be taken if a packet matches the rule's criteria. Permit means that matching traffic will be accepted, Deny means that it will be excluded.

Match Every Select True or False from the pull-down menu. If you select true you are specifying that all packets will match the selected ACL and Rule and will be either permitted or denied. In this case, since all packets match the rule, you will not be offered the option of configuring other match criteria. To configure specific match criteria for the rule, remove the rule and re-create it, or re-configure Match Every to False for the other match criteria to be visible. Click the Apply button to save your choice and return to the main screen, or click the Cancel button to exit without saving a change.

Protocol Keyword

Specify that a packet's IP protocol is a match condition for the selected ACL rule. If you click Configure on this line you will be shown a new screen where you can select the protocol to be used as the match condition. The possible values are ICMP, IGMP, IP, TCP, and UDP. Either the Protocol Keyword field or the Protocol Number field can be used to specify an IP protocol value as a match criterion.

Protocol Number

Specify that a packet's IP protocol is a match condition for the selected ACL rule and identify the protocol by number. If you click Configure on this line you will be shown a new screen where you can select the protocol to be used as the match condition. The protocol number is a standard value assigned by IANA and is interpreted as an integer from 1 to 255. Either the Protocol Number field or the Protocol Keyword field can be used to specify an IP protocol value as a match criterion.

Source IP Address

Specify that a packet's source IP address is a match condition for the selected ACL rule. If you click Configure on this line you will be shown a new screen where you can select the IP address and mask to be used as the match condition. On that screen you can enter an IP address using dotted-decimal notation.

Destination IP Address

Enter an IP address using dotted-decimal notation to be compared to a packet's destination IP address as a match criteria for the selected ACL rule.

Source IP Mask

Enter the IP Mask in dotted-decimal notation to be used with the Source IP address value.

Source L4 Port Keyword

Specify that a packet's source Layer 4 port is a match condition for the selected ACL rule. If you click Configure on this line you will be shown a new screen where you can select the port to be used as the match condition. The possible values are domain, echo, FTP, ftpdata, HTTP, SMTP, SNMP, Telnet, TFTP, and www. Each of these values translates into its equivalent port number, which is used as both the start and end of the port range.

Source L4 Port Number

Specify a packet's source Layer 4 port number as a match condition for the selected ACL rule.

Destination L4 Port Keyword

Specify that a packet's destination Layer 4 port is a match condition for the selected ACL rule. If you click Configure on this line you will be shown a new screen where you can select the protocol to be used as the match condition. The possible values are domain, echo, FTP, ftpdata, HTTP, SMTP, SNMP, Telnet, TFTP, and www. Each of these values translates into its equivalent port number, which is used as both the start and end of the port range.

Destination L4 Port Number

Specify a packet's destination Layer 4 port number match condition for the selected ACL rule.

Click the Configure button to configure the corresponding match criteria for the selected rule. Click the Delete button to remove the currently selected Rule from the selected ACL. If you want the switch to retain the new values across a power cycle you must perform a save.

Bandwidth provisioning

This menu provides access to the following bandwidth provisioning configuration and summary screens:

- Bandwidth profile configuration
- Bandwidth profile summary
- Traffic class configuration
- Traffic class summary
- Interface allocation summary

Bandwidth profile configuration

Use this panel to create a bandwidth allocation profile.

The screenshot shows a web interface for configuring a bandwidth profile. On the left is a navigation tree with the following items: System, Switching, Routing, Class of Service, Security, QoS, Access Control Lists, Bandwidth Provisioning (expanded), Bandwidth Profile Configuration (selected), Bandwidth Profile Summary, Traffic Class Configuration, Traffic Class Summary, Interface Allocation Summary, IP Multicast, and Logout. The main content area is titled "Bandwidth Profile Configuration" and contains the following fields: "Bandwidth Profile" with a dropdown menu set to "Create", "Name" with an empty text input field, and "Maximum Bandwidth (Mbps)" with a text input field containing "0" and a note "(0 to Interface Maximum Bandwidth)". An "Apply" button is located at the bottom right of the form.

Bandwidth Profile

Select Create from the pull-down menu to configure a new bandwidth

profile, or select one of the existing profiles to display and update its configuration. Bandwidth profile 1, named default, always exists and you cannot change or delete it.

Name Enter the name you want to give to the bandwidth profile. You can enter up to 15 alpha-numeric characters and can include the underscore _ or the dash -. You cannot change the name after the initial configuration.

Maximum Bandwidth


Enter the maximum allowable bandwidth for this bandwidth allocation profile.

Click the Apply button to send the updated configuration to the switch. Configuration changes take effect immediately. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Delete button to delete the selected bandwidth allocation profile from the system.

Bandwidth profile summary

This panel displays the bandwidth allocation information for all bandwidth profiles on the switch.



The screenshot shows a network management interface with a sidebar on the left containing a navigation menu. The main content area displays a table titled "Bandwidth Profile Summary". The table has four columns: "Bandwidth Profile", "Name", "Minimum Bandwidth (Mbps)", and "Maximum Bandwidth (Mbps)". There are two rows of data: profile 1 named "default" with a minimum bandwidth of 1 Mbps and a maximum of 100 Mbps; and profile 2 named "test_profile1" with a minimum bandwidth of 0 Mbps and a maximum of 90000 Mbps.

Bandwidth Profile	Name	Minimum Bandwidth (Mbps)	Maximum Bandwidth (Mbps)
1	default	1	100
2	test_profile1	0	90000

Bandwidth Profile

Displays the number associated with the bandwidth profile.

Name

Displays the name of the bandwidth profile.

Allocated Minimum Bandwidth

Displays the sum of the minimum guaranteed bandwidth for all bandwidth profiles configured on this interface.

Maximum Bandwidth

Displays the sum of the maximum allowable bandwidth for all bandwidth profiles configured on this interface.

Traffic class configuration

Use this panel to create a traffic class.

The screenshot shows the 'Traffic Class Configuration' web interface. On the left is a navigation sidebar with a tree view containing: System, Switching, Routing, Class of Service, Security, QoS, Access Control Lists, Bandwidth Provisioning, Bandwidth Profile Configuration, Bandwidth Profile Summary, Traffic Class Configuration (highlighted), Traffic Class Summary, Interface Allocation Summary, IP Multicast, and Logout. The main panel is titled 'Traffic Class Configuration' and contains the following fields:

- Traffic Class: Create (dropdown)
- Name: [text input]
- Weight: 1 (range 1 to 1024)
- Type: per VLAN per Interface
- VLAN ID: 1 (range 1 to 4094)
- Interface: Bay 1 (dropdown)
- Bandwidth Profile: default-1 (1-100Mbps) (dropdown)

An 'Apply' button is located at the bottom right of the configuration area.

Traffic Class

Select Create from the pull-down menu to configure a new Traffic Class, or select one of the existing classes to display and update its configuration.

Name Enter the name to be given to the Traffic Class. You can enter up to 15 alpha-numeric characters and can include the underscore _ or the dash -. You cannot change the name after the initial configuration.

Weight

Enter the weight to be assigned to the Traffic Class. The weight must be a decimal number from 1 to 1024.

Type The only supported type is per VLAN per Interface.

VLAN ID

Enter the ID of the VLAN to be associated with the traffic class. This is a value between 2 and 4094.

Interface

Select the interface to which the Traffic Class will be applied. The pull-down menu contains the port identification of all interfaces for which a traffic class can be configured.

Bandwidth Profile

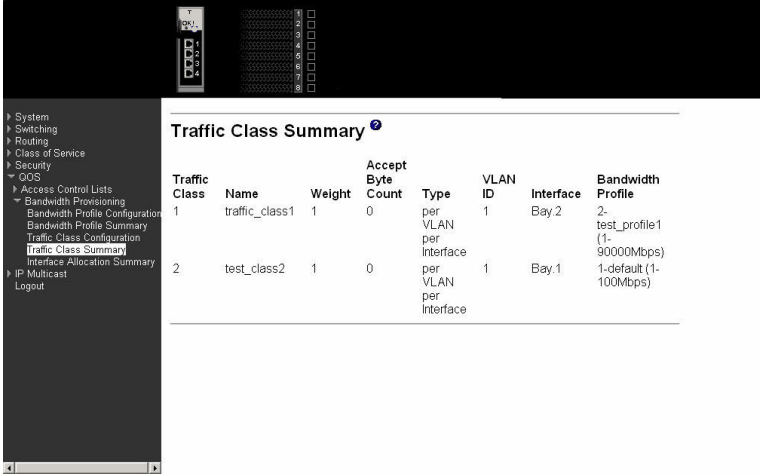
Select the Bandwidth Profile for the Traffic Class from the pull-down menu. The list contains the identification of all Bandwidth Profiles in the form "name-id (min-max Mbps)". If you have not configured any Bandwidth Profiles the list will contain only the default profile. This field associates a bandwidth allocation profile with a Traffic Class. The sum of the bandwidth allocation profile minimum bandwidth of all Traffic Classes associated with the same interface should not exceed the total bandwidth of the interface. There is no restriction on the sum of the maximum bandwidth of all Traffic Classes associated with the same interface. When a Traffic Class is attached to a LAG interface, the bandwidth allocation profile minimum bandwidth parameter will not be applicable to the Traffic Class.

Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

Click the Delete button to remove the currently selected Traffic Class.

Traffic class summary

This panel displays the traffic class information for all Traffic Classes in the system.



Traffic Class	Name	Weight	Accept Byte Count	Type	VLAN ID	Interface	Bandwidth Profile
1	traffic_class1	1	0	per VLAN per Interface	1	Bay.2	2-test_profile1 (1-90000Mbps)
2	test_class2	1	0	per VLAN per Interface	1	Bay.1	1-default (1-100Mbps)

Traffic Class The number of the Traffic Class whose data is displayed in the rest of the line.

Name The user-defined name of this Traffic Class.

Weight The weight of this Traffic Class.

Accept Byte Count The number of bytes accepted for the Traffic Class.

Type The only supported type is per VLAN per Interface.

VLAN ID The VLAN ID with which this Traffic Class is associated.

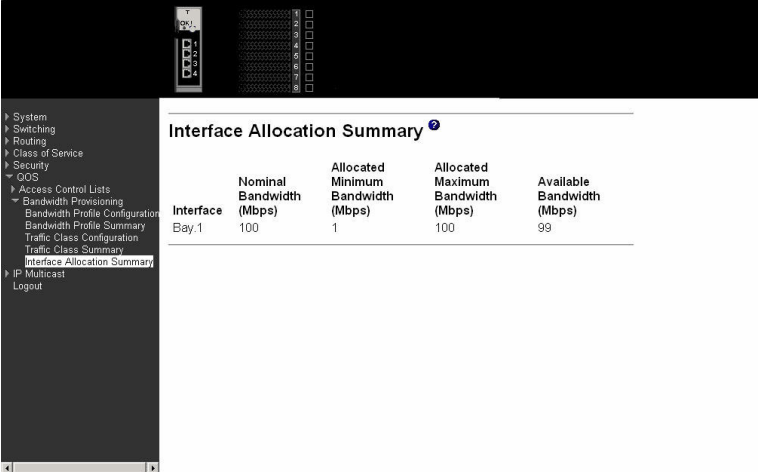
Interface The interface to which the Traffic Class is applied.

Bandwidth Profile The bandwidth allocation profile associated with this Traffic Class in the form “name-id (min-max Mbps)”. This field is blank when there is no bandwidth allocation profile associated with this traffic class.

Interface allocation summary

This panel displays the bandwidth allocated to the listed interfaces. The allocated minimum bandwidth does not exceed the capability of the interface unless the

interface is a LAG.



Interface	Nominal Bandwidth (Mbps)	Allocated Minimum Bandwidth (Mbps)	Allocated Maximum Bandwidth (Mbps)	Available Bandwidth (Mbps)
Bay 1	100	1	100	99

Interface The Port designation of an interface for which you have configured one or more traffic classes.

Nominal Bandwidth (Mbps)

The interface's nominal bandwidth in Mbps. This number is only known for physical interfaces.

Allocated Minimum Bandwidth (Mbps)

The sum of the minimum guaranteed bandwidth for all traffic classes configured on this interface.

Allocated Maximum Bandwidth (Mbps)

The sum of the maximum allowable bandwidth for all traffic classes configured on this interface.

Available Bandwidth (Mbps)

The difference between the Nominal and Allocated Minimum Bandwidths. This number is only known for physical interfaces.

IP multicast

This menu provides access to the following multicast-related submenus:

- Distance vector multicast routing protocol (DVMRP)
- IGMP
- Multicast (Mcast)
- Mdebug
- Protocol independent multicast - dense mode (PIM-DM)
- Protocol independent multicast - sparse mode (PIM-SM)

DVMRP

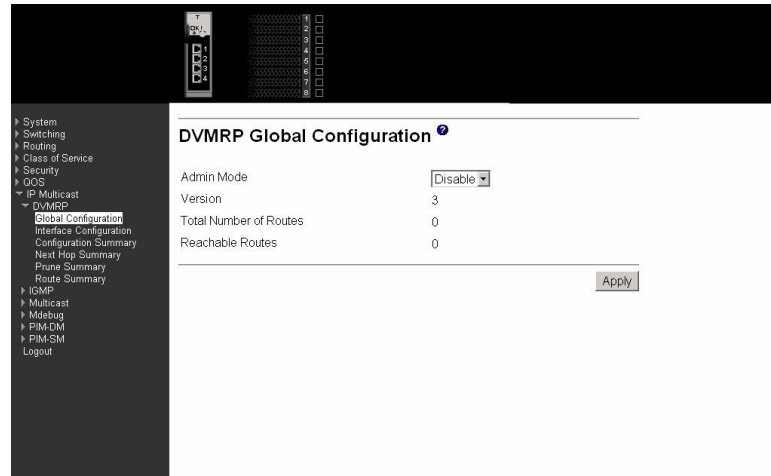
This menu provides access to the following Distance Vector Multicast Routing Protocol (DVMRP) entry panels:

- Global configuration
- Interface configuration
- Configuration summary
- Next hop summary
- Prune summary

- Route summary

Global configuration

Use this panel to enable or disable DVMRP on the switch, and to display the number of routes in the DVMRP routing table.



Admin Mode Select Enable or Disable from the drop-down menu. This sets the administrative status of DVMRP to active or inactive. The default is Disable.

Version The current value of the DVMRP version string. Total

Total Number of Routes

The number of routes in the DVMRP routing table.

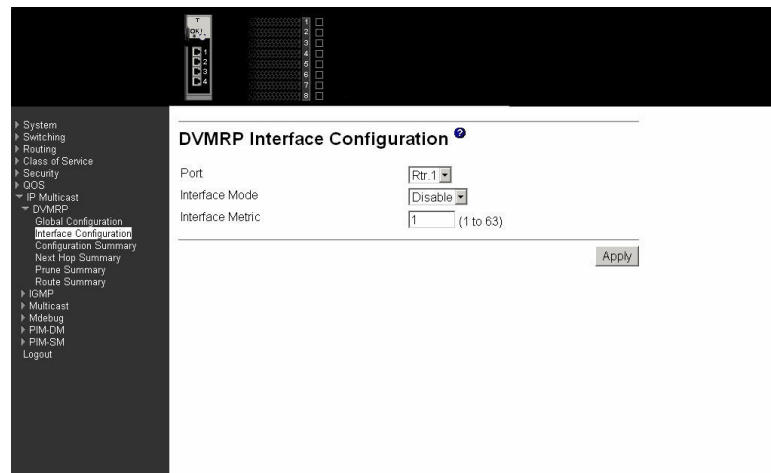
Reachable Routes

The number of routes in the DVMRP routing table that have a non-infinite metric.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Interface configuration

Use this panel to configure a router interface as a DVMRP interface.



Port Select the interface for which data is to be configured. You must

configure at least one router interface before you configure a DVMRP interface. Otherwise you will see a message telling you that no router interfaces are available, and the configuration screen will not be displayed.

Interface Mode

Select Enable or Disable from the pull-down menu to set the administrative mode of the selected DVMRP routing interface.

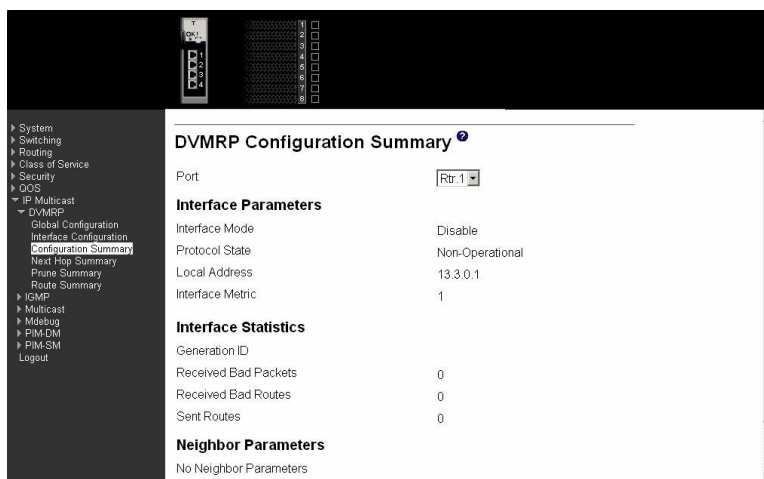
Interface Metric

Enter the DVMRP metric for the selected interface. This value is sent in DVMRP messages as the cost to reach this network. Valid values are from.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Configuration summary

Use this panel to display DVMRP parameters and statistics for the selected DVMRP interface.



Port Select the interface for which data is to be displayed. You must configure at least one router interface before you can display data for a DVMRP interface. Otherwise you will see a message telling you that no router interfaces are available, and the configuration summary screen will not be displayed.

Interface Mode

The administrative mode of the selected DVMRP routing interface, either Enable or Disable.

Protocol State

The operational state of the DVMRP protocol on the selected interface, either operational or non-operational.

Local Address

The IP address used as a source address in packets sent from the selected interface.

Interface Metric

The metric used to calculate distance vectors for the selected interface.

Generation ID

The DVMRP generation ID used by the router for the selected interface. This value is reset every time an interface is (re)started and is placed in prune messages. A change in generation ID informs the neighbor routers that any previous information about this router should be discarded.

Received Bad Packets

The number of invalid packets received on the selected interface.

Received Bad Routes

The number of invalid routes received on the selected interface.

Sent Routes

The number of routes sent on the selected interface.

Neighbor IP

The IP address of the neighbor whose information is displayed.

State The state of the specified neighbor router on the selected interface, either active or down.

The following screens are present on summary screens but are not configurable.

Neighbor Uptime

The DVMRP uptime for the specified neighbor on the selected interface. This is the time since the neighbor entry was learned.

Neighbor Expiry Time

The DVMRP expiry time for the specified neighbor on the selected interface. This is the time left before this neighbor entry will age out, and is not applicable if the neighbor router's state is down.

Generation ID The DVMRP generation ID for the specified neighbor on the selected interface.

Major Version The DVMRP Major Version for the specified neighbor on the selected interface.

Minor Version

The DVMRP Minor Version for the specified neighbor on the selected interface.

Capabilities The DVMRP capabilities of the specified neighbor on the selected interface.

Received Routers

The number of routes received for the specified neighbor on the selected interface.

Received Bad Packets

The number of invalid packets received for the specified neighbor on the selected interface.

Received Bad Routes

The number of invalid routes received for the specified neighbor on the selected interface.

Click the Refresh button to refresh the screen with the new data.

Next hop summary

Use this panel to display next hop information for an entry in the DVMRP routing table.

DVMRP Next Hop Summary

Source IP	Source Mask	Next Hop Interface	Type
9.0.0.0	255.0.0.0	Rtr.1	Leaf

Refresh

Source IP The IP address used with the source mask to identify the source network for this table entry.

Source Mask The network mask used with the source IP address.

Next Hop Interface The outgoing interface for this next hop.

Type The next hop type. 'Leaf' means that no downstream dependent neighbors exist on the outgoing interface. Otherwise, the type is 'branch'.

Click the Refresh button to refresh the screen with the new data.

Prune summary

Use this panel to display information about multicast addresses that have been removed from the DVMRP routing table because of receipt of prune messages.

DVMRP Prune Summary

Group IP	Source IP	Source Mask	Expiry Time (secs)
224.0.0.0	9.1.1.1	255.0.0.0	40

Refresh

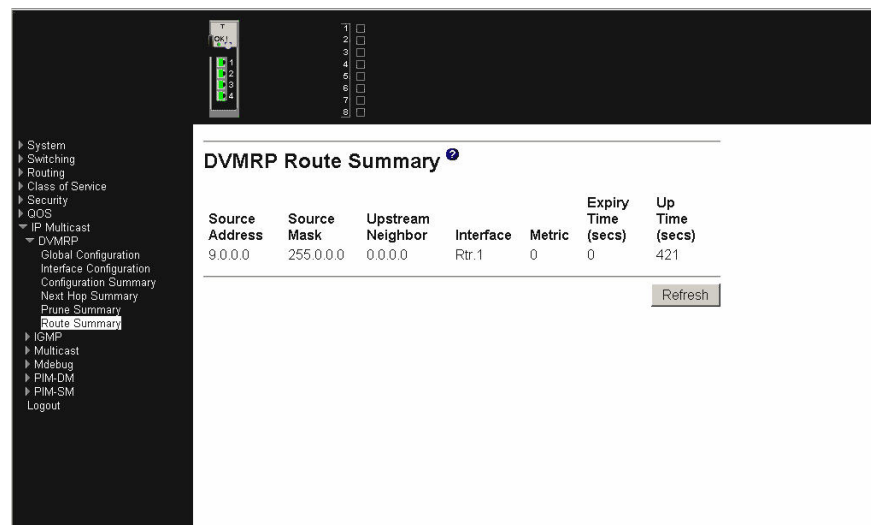
Group IP The group address which has been pruned.

- Source IP** The address of the source or source network which has been pruned.
- Source Mask** The subnet mask to be combined with the source IP address to identify the source or source network which has been pruned.
- Expiry Time** The amount of time remaining before this prune should expire at the upstream neighbor. If no prune messages have been received from downstream neighbors, this is set to value of the default prune lifetime timer, otherwise it is set to the smallest received value or the default timer, whichever is less.

Click the Refresh button to refresh the screen with the new data.

Route summary

Use this panel to display information about sources in the DVMRP routing table.



Source Address

The network address that is combined with the source mask to identify the sources for this entry.

Source Mask The subnet mask to be combined with the source address to identify the sources for this entry.

Upstream Neighbor

The address of the upstream neighbor (e.g., RPF neighbor) from which IP datagrams from these sources are received.

Interface The interface on which IP datagrams sent by these sources are received. A value of 0 typically means the route is an aggregate for which no next-hop interface exists.

Metric The distance in hops to the source subnet.

Expiry Time (secs)

The minimum amount of time remaining before this entry will be aged out.

Up Time (secs)

The time since the route represented by this entry was learned by the router.

Click the Refresh button to refresh the screen with the new data.

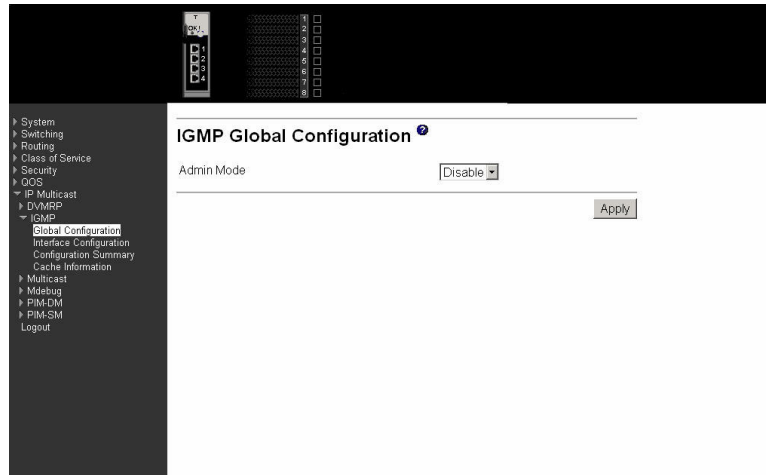
IGMP

This menu provides access to the following IGMP entry panels:

- Global configuration
- Interface configuration
- Configuration summary
- Cache information

Global configuration

Use this panel to enable or disable IGMP for the switch. IGMP must be enabled before any other multicast routing protocols can be enabled.

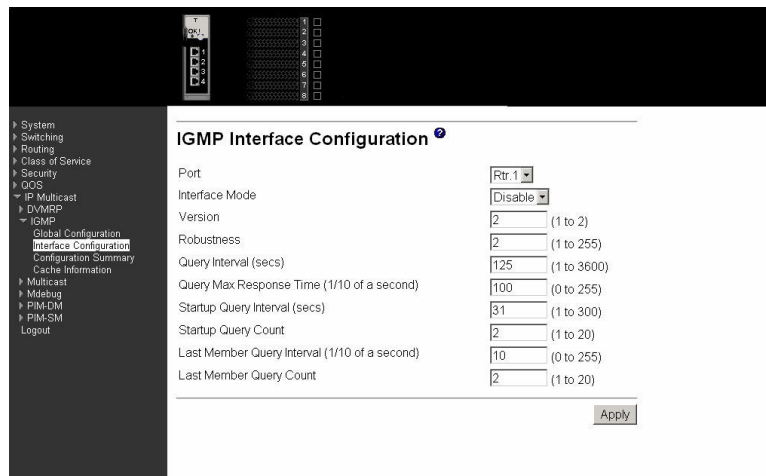


Admin Mode Select Enable or Disable from the pull-down menu to set the administrative status of IGMP in the router to active or inactive. The default is Disable.

Click the Apply button to update the router with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Interface configuration

Use this panel to configure a routing interface as an IGMP interface, or to display the parameters for an existing IGMP interface.



Port Select the physical interface for which data is to be displayed or

configured from the pull-down menu. You must have configured at least one router interface before configuring or displaying data for an IGMP interface, otherwise an error message will be displayed.

Interface Mode

Select Enable or Disable from the pull-down menu to set the administrative status of IGMP on the selected interface. The default is Disable.

Version

Enter the version of IGMP you want to configure on the selected interface. Valid values are 1 to 2 and the default value is 2.

Robustness

Enter the robustness value. This variable allows tuning for the expected packet loss on a subnet. If you expect the subnet to be open to loss, you should enter a higher number for this parameter. IGMP is robust to (robustness variable-1) packet losses. Valid values are from 1 to 255. The default value is 2.

Query Interval (secs)

Enter the frequency in seconds at which IGMP host-query packets are to be transmitted on this interface. Valid values are from 1 to 3600. The default value is 125.

Query Max Response Time (1/10 of a sec)

Enter the maximum query response time to be advertised in IGMPv2 queries on this interface, in tenths of a second. The default value is 10. Valid values are from.

Startup Query Interval (secs)

Enter the number of seconds between the transmission of startup queries on the selected interface. The valid values are from 1 to 300. The default value is 31.

Startup Query Count

Enter the number of queries to be sent on startup. The valid values are from 1 to 20. The default value is 2.

Last Member Query Interval (1/10 of a sec)

- Enter the last member query interval in tenths of a second. This the maximum response time to be inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Valid values are from 0 to 255. The default value is 1. This value is not used for IGMP version 1.

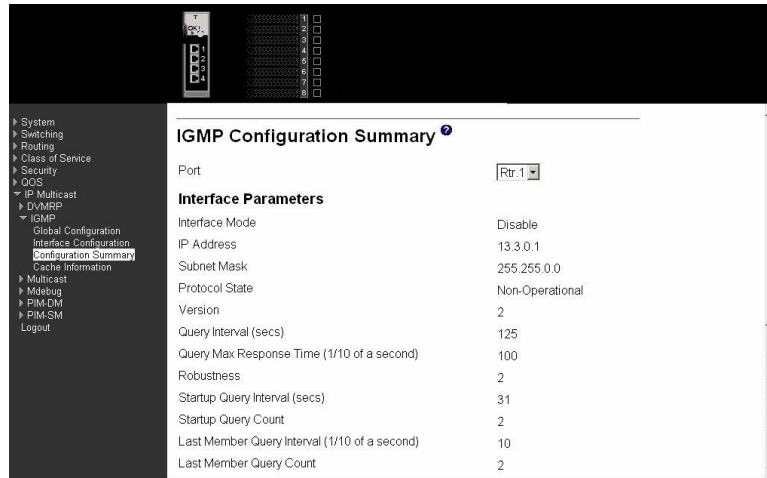
Last Member Query Count

- Enter the number of queries to be sent on receiving a leave group report. Valid values are from 1 to 20. The default value is 2.

Click the Apply button to update the router with the values on this screen. If you want the router to retain the new values across a power cycle, you must perform a save.

Configuration summary

Use this panel to display IGMP parameters and statistics for an IGMP interface.



Port Select the physical interface for which data is to be displayed.

Interface Mode

The administrative status of IGMP on the selected interface.

IP Address

The IP address of the selected interface.

Subnet Mask

The subnet mask for the IP address of the selected interface.

Protocol State

The operational state of IGMP on the selected interface.

Version

The version of IGMP configured on the selected interface.

Query Interval (secs)

The frequency at which IGMP host-query packets are transmitted on the selected interface.

Query Max Response Time (1/10 of a sec)

The maximum query response time advertised in IGMPv2 queries sent from the selected interface.

Robustness

The robustness parameter for the selected interface. This variable allows tuning for the expected packet loss on a subnet. If a subnet is expected to be open to loss, the robustness variable can be increased. IGMP is robust to (robustness variable-1) packet losses.

Startup Query Interval (secs)

The interval at which startup queries are sent on the selected interface.

Startup Query Count

The number of queries to be sent on startup.

Last Member Query Interval (1/10 of a sec)

The last member query interval. The last member query interval is the maximum response time inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. This value can be tuned to modify the

leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This value is not used for IGMP version 1.

Last Member Query Count

The number of queries to be sent on receiving a leave group report.

Querier

The address of the IGMP querier on the IP subnet to which the selected interface is attached.

Querier Status

Indicates whether the selected interface is in querier or non querier mode.

The following fields are not available unless the protocol is enabled and valid entries exist.

Querier Up Time (secs)

The time in seconds since the IGMP interface querier was last changed.

Querier Expiry Time (secs)

The time in seconds remaining before the other querier present timer expires. If the local system is the querier, this will be zero.

Wrong Version Queries

The number of queries that have been received on the selected interface with an IGMP version that does not match the IGMP version configured for the interface, over the lifetime of the entry. IGMP requires that all routers on a LAN be configured to run the same version of IGMP. Therefore, a configuration error is indicated if any queries are received with the wrong version number.

Number of Joins

The number of times a group membership has been added on the selected interface; that is, the number of times an entry for this interface has been added to the cache table. This gives an indication of the amount of IGMP activity on the interface.

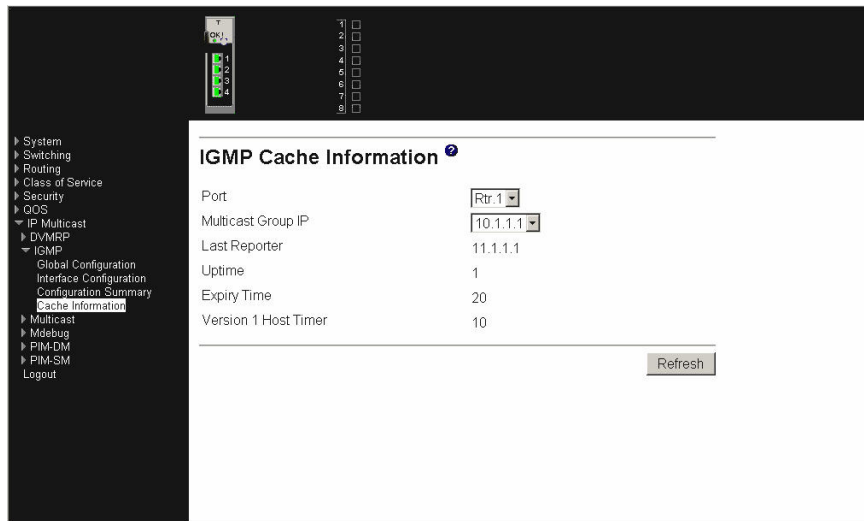
Number of Groups

The current number of entries for the selected interface in the cache table.

Click the Refresh button to refresh the screen with the new data.

Cache information

Use this panel to display information about a specified multicast address on a specified IGMP interface.



Port Select the physical interface for which data is to be displayed.

Multicast Group IP

Select the IP multicast group address for which data is to be displayed. If no group membership reports have been received on the selected interface you will not be able to make this selection, and none of the non-configurable data will be displayed.

Last Reporter The IP address of the source of the last membership report received for the IP Multicast group address on the selected interface.

Uptime The time elapsed since this entry was created.

Expiry Time The minimum amount of time remaining before this entry will be aged out.

Version 1 Host Timer

The time remaining until the local router will assume that there are no longer any IGMP version 1 members on the IP subnet attached to this interface. When an IGMPv1 membership report is received, this timer is reset to the group membership timer. While this timer is non-zero, the local router ignores any IGMPv2 leave messages for this group that it receives on the selected interface. This field is displayed only if the interface is configured for IGMP version 1.

Click the Refresh button to refresh the screen with the new data.

Multicast

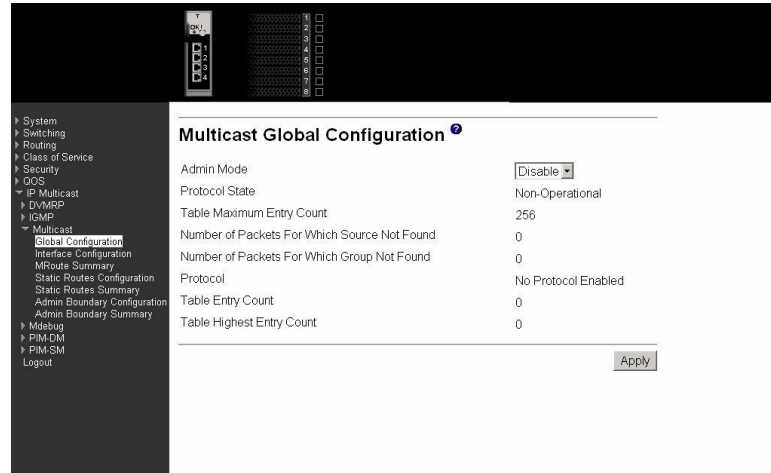
This menu provides access to the following multicast entry panels:

- Global configuration
- Interface configuration
- MRoute summary
- Static routes configuration
- Static routes summary

- Admin boundary configuration
- Admin boundary summary

Global configuration

Use this panel to enable or disable support for multicast forwarding and to display parameters and statistics associated with multicast forwarding.



Admin Mode Select Enable or Disable to set the administrative status of Multicast Forwarding in the router. The default is Disable.

Protocol State The operational state of the multicast forwarding module.

Table Maximum Entry Count The maximum number of entries in the IP Multicast routing table.

Number Of Packets For Which Source Not Found The number of multicast packets that were supposed to be routed but which failed the RPF check.

Number Of Packets For Which Group Not Found The number of multicast packets that were supposed to be routed but for which no multicast route was found.

Protocol The multicast routing protocol presently activated on the router, if any.

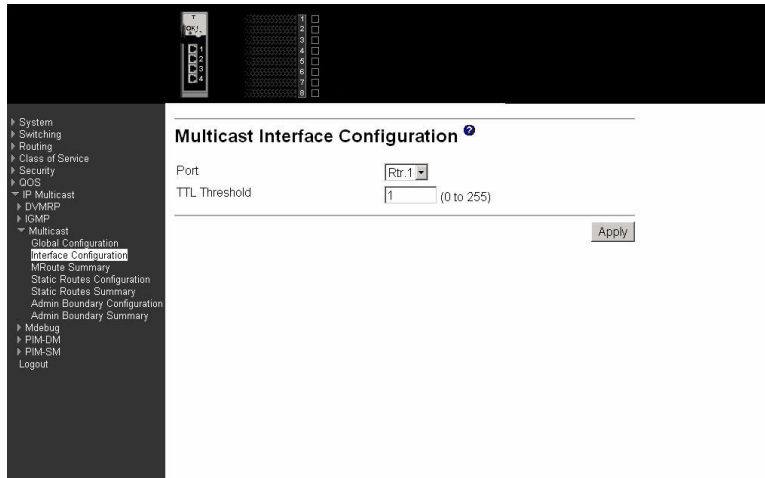
Table Entry Count The number of multicast route entries currently present in the Multicast route table.

Table Highest Entry Count The highest number of multicast route entries that have been present in the Multicast route table.

Click the Apply button to update the router with the values on this screen. If you want the router to retain the new values across a power cycle, you must perform a save.

Interface configuration

This panel displays multicast information for the specified router interface.



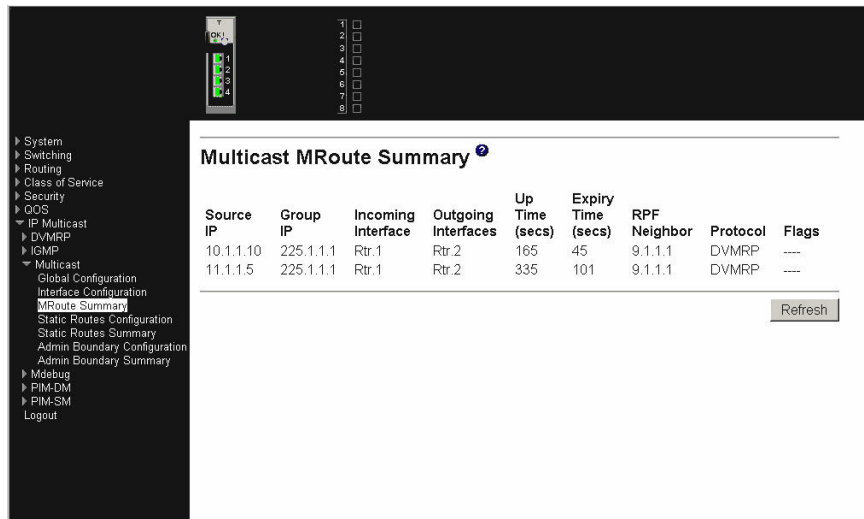
Port The interface to be configured.

TTL Threshold The time to live value to be used in packets transmitted from this interface

Click the Apply button to update the router with the values on this screen. If you want the router to retain the new values across a power cycle, you must perform a save.

MRoute summary

Use this panel to display or configure multicast route information for a specified source and destination address pair.



Source IP Enter the IP address of the multicast packet source to be combined with the Group IP to fully identify a single route whose Mroute table entry you want to display or clear. This field blank can be left blank.

Group IP Enter the destination group IP address whose multicast route(s) you want to display or clear.

Incoming Interface

The incoming interface on which multicast packets for this source/group arrive.

Outgoing Interface(s)

The list of outgoing interfaces on which multicast packets for this source/group are forwarded.

Up Time

The time in seconds since the entry was created.

Expiry Time

The time in seconds before this entry will age out and be removed from the table.

RPF Neighbor

The IP address of the Reverse Path Forwarding neighbor.

Protocol

The multicast routing protocol which created this entry. The possibilities are:

- PIM-DM
- PIM-SM
- DVMRP

Flags

The value displayed in this field is valid if the multicast routing protocol running is PIMSM. The possible values are RPT or SPT. For other protocols, “----” is displayed.

Click the Delete button to delete the entry. If you want the switch to retain the deletion across a power cycle, you must perform a save.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Static routes configuration

Use this panel to configure the parameters for a multicast source to be added to the multicast routing table.

The screenshot shows a web-based configuration interface for a switch. On the left is a dark sidebar with a navigation tree containing items like System, Switching, Routing, and Multicast. The main content area is titled "Multicast Static Routes Configuration" and contains several input fields: "Source" (a dropdown menu with "Create Static Route" selected), "Source IP", "Source Mask", "RPF Neighbor", "Metric", and "Port" (a dropdown menu with "Rtr.1" selected). An "Apply" button is located at the bottom right of the configuration area.

Source

Select Create Static Route to configure a new static entry in the Mroute table, or select one of the existing entries from the pull-down menu.

Source IP

Enter the IP address that identifies the multicast packet source for the entry you are creating.

Source Mask Enter the subnet mask to be applied to the Source IP address.

RPF Neighbor

Enter the IP address of the neighbor router on the path to the source.

Metric

Enter the link state cost of the path to the multicast source. The range is 0 - 255 and the default is one. You can change the metric for a configured route by selecting the static route and editing this field.

Port

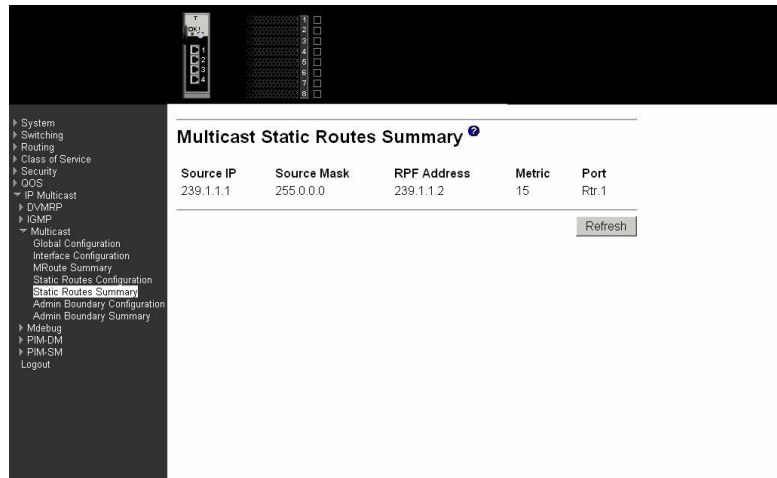
Select the interface number from the drop-down menu. This is interface that connects to the neighbor router for the given source IP address.

Click the Delete button to delete the static entry with the selected Source IP address from the Mroute table.

Click the Apply button to update the router with the values on this screen. If you want the router to retain the new values across a power cycle, you must perform a save.

Static routes summary

Use this panel to display summary information for a multicast source.



Source IP

The IP address that identifies the multicast packet source for this route.

Source Mask

The subnet mask applied to the Source IP address.

RPF Address

The IP address of the RPF neighbor.

Metric

The number of the incoming interface whose IP address is used as RPF for the given source IP address.

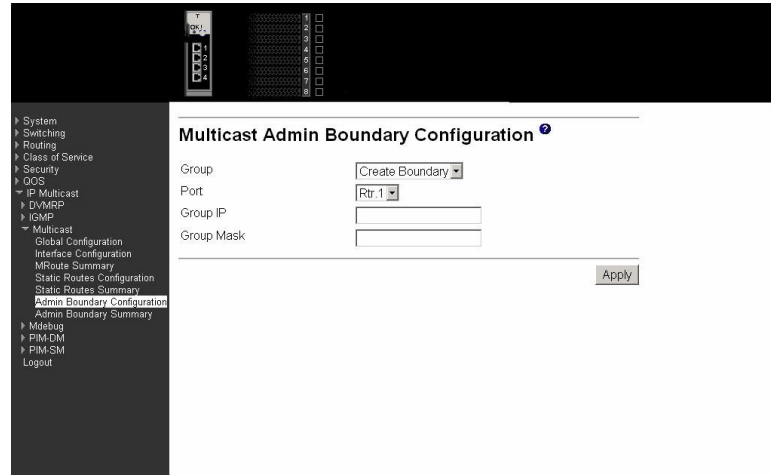
Port

The number of the incoming interface whose IP address is used as RPF for the given source IP address.

Click the Refresh button to refresh the data on the screen with the present state of the data in the router.

Admin boundary configuration

Use this panel to configure or update a range of addresses for which multicast packets will not be forwarded.



The screenshot shows a web-based configuration interface for a network device. On the left is a dark sidebar with a tree view of configuration categories: System, Switching, Routing, Class of Service, Security, QoS, IP Multicast (expanded), IPv6, ICMP, and Multicast (expanded). Under Multicast, sub-items include Global Configuration, Interface Configuration, MRoute Summary, Static Routes Configuration, Static Routes Summary, Admin Boundary Configuration (highlighted), and Admin Boundary Summary. The main content area is titled 'Multicast Admin Boundary Configuration' and contains four fields: 'Group' with a dropdown menu set to 'Create Boundary', 'Port' with a dropdown menu set to 'Rtr. 1', 'Group IP' with an empty text input, and 'Group Mask' with an empty text input. An 'Apply' button is located at the bottom right of the configuration area.

Group Select Create Boundary from the pull-down menu to create a new admin scope boundary, or select one of the existing boundary specifications to display or update its configuration.

Port Select the router interface for which the administratively scoped boundary is to be configured.

Group IP Enter the multicast group address for the start of the range of addresses to be excluded. The address must be in the range of 239.0.0.0 through 239.255.255.255.

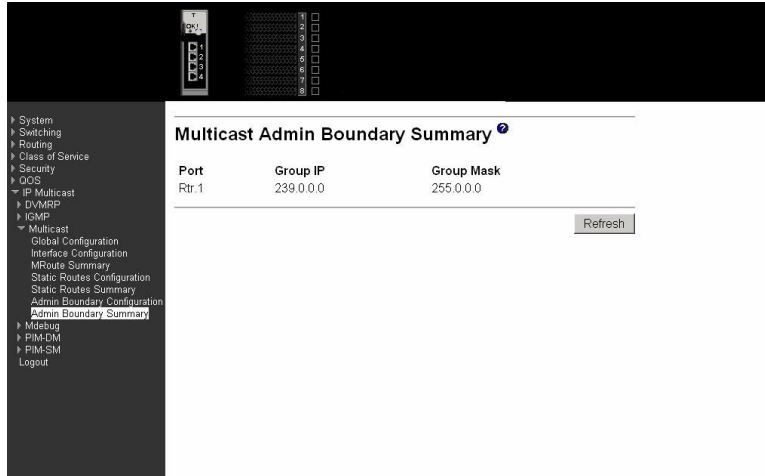
Group Mask Enter the mask to be applied to the multicast group address. The combination of the mask and the Group IP gives the range of administratively scoped addresses for the selected interface.

Click the Delete button to delete the selected administrative scoped boundary.

Click the Apply button to send the updated configuration to the router. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

Admin boundary summary

Use this panel to display information about excluded multicast addresses.



Port The router interface to which the administratively scoped address range is applied.

Group IP The multicast group address for the start of the range of addresses to be excluded.

Group Mask The mask that is applied to the multicast group address. The combination of the mask and the Group IP gives the range of administratively scoped addresses for the selected interface.

Click the Refresh button to refresh the data on the screen with the present state of the data in the router.

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Mdebug

This menu provides access to the following Mdebug entry panels:

- Mrinfo run
- Mrinfo show
- Mstat run
- Mstat show
- Mtrace config
- Mtrace run
- Mtrace show

Mrinfo run

Use this panel to request data about a router interface be displayed on the Mrinfo Show panel.



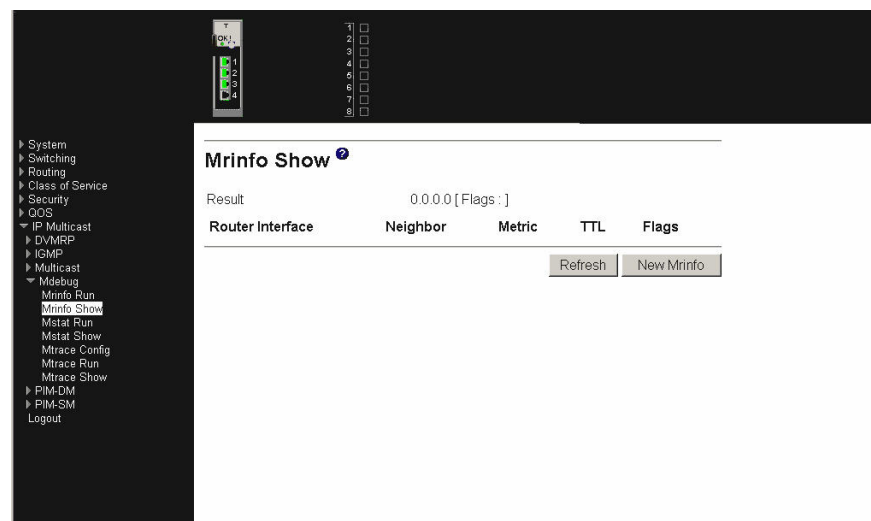
Router Interface

Enter the IP address of the router interface for which you want to see the neighbor router information. If you do not enter an address the router will query itself.

Click the Apply button to initiate the mrinfo command on the router. If the mrinfo command completes successfully the browser will display the Mrinfo Show screen. If the mrinfo command fails, you will see the Mrinfo Run screen again.

Mrinfo show

Use this panel to display information about a router interface. You must use the Mrinfo Run panel to initiate collection of the data displayed on this panel.



Router Interface

The IP address of the router interface for which configuration information was requested.

Neighbor

The IP address of the neighboring router.

Metric

The routing metric for this router.

- TTL** The time-to-live threshold on this hop.
- Flags** The flags indicating whether the router is an IGMP querier or whether or not it has neighbors (leaf router).

Click the New Mrinfo button to redirect the web browser to the Mrinfo Run screen so that you can initiate another mrinfo command.

Click the Refresh button to refresh the content of the screen with the latest data available on the router. Typically, it takes around 20 seconds to process the results after you have initiated the mrinfo command. The contents of the screen have to be refreshed to display the latest results.

Mstat run

Use this panel to initiate a trace of a multicast path. The results will be displayed on the Mstat Show panel.



Source IP Enter the IP address of the multicast-capable source. This is the unicast address of the beginning of the path to be traced.

Receiver IP Enter the IP address of the host to which the mtrace response will be sent by the last hop router. If a value is not entered, the IP address of the router interface through which the mtrace will be sent is used.

Group IP Enter the multicast address of the group to be traced. If you leave this field blank, the multicast address 224.2.0.1 will be used. Valid addresses are 224.0.0.0 through 239.255.255.255.

Click the Apply button to initiate the mstat command on the router. If the mstat command completes successfully the browser will display the Mstat Show screen. If the mstat command fails, you will see the Mstat Run screen again.

Mstat show

This panel displays the path taken by multicast traffic between the specified IP addresses. Forward data flow is indicated by arrows pointing downward and the query path is indicated by arrows pointing upward. For each hop, both the entry and exit addresses of the router are shown if different, along with the initial TTL required for packets to be forwarded at this hop and the propagation delay across the hop. The right half of the screen displays statistics for the path in two groups. Within each group, the columns are the number of packets lost, the number of packets sent, the percentage lost, and the average packet rate at each hop. These statistics

are calculated from differences between traces and from hop to hop. The first group shows the statistics for all traffic flowing out the interface at one hop and in the interface at the next hop. The second group shows the statistics only for traffic forwarded from the specified source to the specified group.

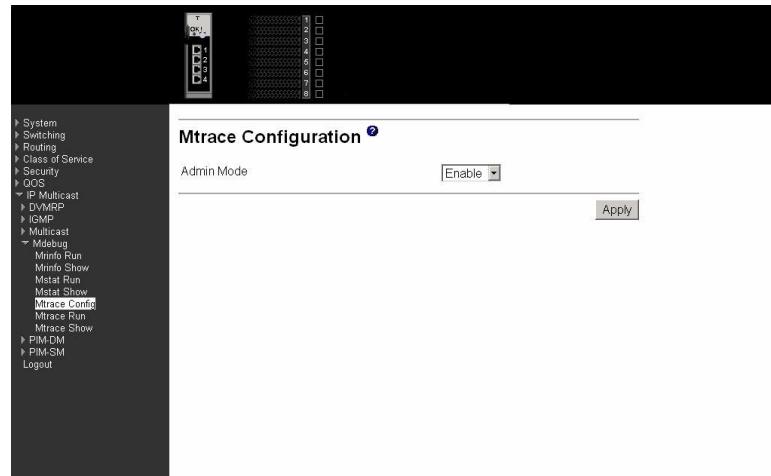


Click the New Mstat button to redirect the web browser to the Mstat Run screen so that you can initiate another mstat command.

Click the Refresh button to refresh the content of the screen with the latest data available on the router. Typically, it takes around 20 seconds to process the results after initiating mstat command. You must refresh the screen to display the latest results.

Mtrace config

Use this panel to enable or disable processing of mtrace requests.



Admin Mode Select Enable or Disable from the pull-down menu. If you select Enable the router will process and forward mtrace requests received from other routers, otherwise received mtrace requests will be discarded.

Click the Apply button to send the updated configuration to the router. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

Mtrace run

Use this panel to initiate an mtrace command.

The screenshot shows the 'Mtrace Run' configuration panel. On the left is a sidebar with a tree view containing the following items: System, Switching, Routing, Class of Service, Security, QoS, IP Multicast (with sub-items DVNRP, IGMP, Multicast, Mdebug), Mirrro Run, Mirrro Show, Metat Run, Metat Show, Mirrro Config, Mtrace Run, Mtrace Show, PIM-DM, PIM-SM, and Logout. The main panel is titled 'Mtrace Run' and contains three input fields: 'Source IP', 'Receiver IP', and 'Group IP'. An 'Apply' button is located at the bottom right of the main panel.

Source IP Enter the IP address of a multicast-capable source. This is the unicast address of the beginning of the path to be traced.

Receiver IP Enter the IP address of the host to which the mtrace response will be sent by the last hop router. If you leave this field blank, mtrace will use the IP address of the router interface through which the mtrace will be sent.

Group IP Enter the Multicast address of the group to be traced. If you do not enter a valid address, multicast address 224.2.0.1 will be used. Valid addresses are 224.0.0.0 through 239.255.255.255.

Click the Apply button to initiate the mtrace command on the router. If the mtrace command completes successfully the browser will display the Mtrace Show screen. If the mtrace command fails, you will see the Mtrace Run screen again.

Mtrace show

Use this panel to display the results of an mtrace command. If the command complete successfully this panel will be displayed automatically.

The screenshot shows the 'Mtrace Show' results panel. The title is 'Mtrace Show'. Below the title, it says 'Result Mtrace for 0.0.0.0 to 0.0.0.0 via 0.0.0.0'. There is a table with the following columns: 'Number of Hops Away from Destination', 'IP Address of Intermediate Router', 'Multicast Protocol in Use', 'TTL Threshold', and 'Time Taken to Forward Between Hops (milliseconds)'. The table contains one row with the following values: 0, 0.0.0.0, (blank), (blank), and (blank). At the bottom right of the table area, there are two buttons: 'Refresh' and 'New Mtrace'.

Number of Hops Away from Destination

The number of hops away from the destination.

IP Address of Intermediate Router

The IP address of the intermediate router in the path being traced between source and destination for the hop number in the previous field.

Multicast Protocol in Use

The multicast protocol in use on this hop.

TTL Threshold

The time-to-live threshold on this hop.

Time Taken to Forward Between Hops (milliseconds)

The time taken for the trace request to be forwarded from the previous hop to this hop.

Click the Refresh button to refresh the content of the screen with the latest data available on the router. Typically, it takes around 20 seconds to process the results after initiating mtrace command. You must refresh the screen to display the latest results.

Click the New Mtrace button to redirect the web browser to the Mtrace Run screen so that you can initiate another mtrace command.

PIM-DM

This menu provides access to the following Protocol Independent Multicast - Dense Mode (PIM-DM) entry panels:

- Global configuration
- Interface configuration
- Interface summary

Global configuration

Use this panel to enable or disable PIM-DM.

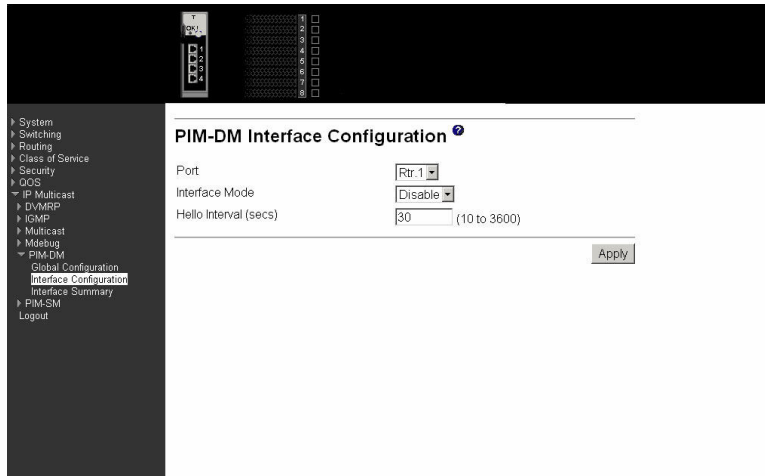


Admin Mode Select Enable or Disable from the pull-down menu to set the administrative status of PIM-DM in the router. The default is Disable.

Click the Apply button to send the updated configuration to the router. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

Interface configuration

Use this panel to configure PIM-DM support for a router interface.



Port

Select the physical interface for which data is to be displayed or configured. You must have configured at least one router interface before configuring or displaying data for a PIM-DM interface, otherwise an error message will be displayed.

Interface Mode

Select Enable or Disable from the pull-down menu to set the administrative status of PIM-DM for the selected interface. The default is Disable.

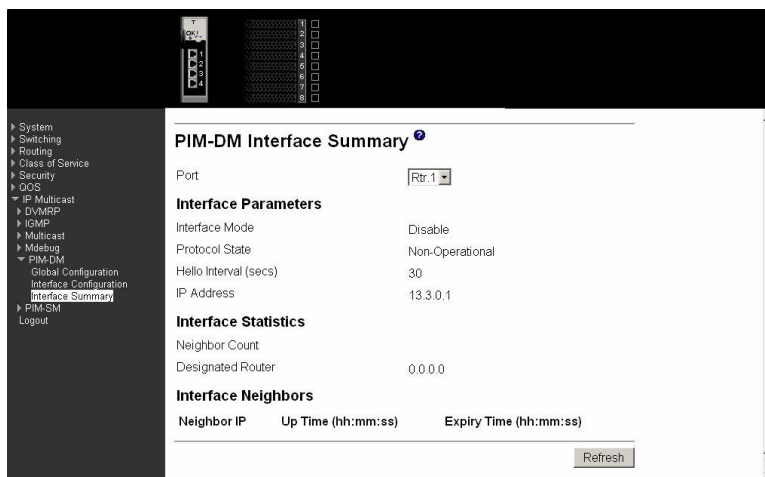
Hello Interval (secs)

Enter the number of seconds between PIM hello messages transmitted from the selected interface. The default value is 30. Valid values are from (10 to 3600).

Click the Apply button to update the switch with the values on this screen. If you want the switch to retain the new values across a power cycle, you must perform a save.

Interface summary

Use this panel to display parameters and statistics for a specified PIM-DM interface.



Port Select the physical interface for which data is to be displayed. There must

be configured at least one router interface before displaying data for a PIM-DM interface, otherwise a message will be displayed.

Interface Parameters

Interface Mode

Displays the administrative status of PIM-DM for the selected interface. The default is Disable.

Protocol State

The operational state of the PIM-DM protocol on this interface.

Hello Interval (secs)

The frequency at which PIM hello messages are transmitted on the selected interface.

IP Address

The IP address of the selected interface.

Interface Statistics

Interface Statistics

Neighbor Count

The number of PIM neighbors on the selected interface.

Designated Router

The designated router on the selected PIM interface. For point- to-point interfaces, this will be 0.0.0.0.

Interface Neighbors

Neighbor IP The IP address of the PIM neighbor for which this entry contains information.

Up Time (hh:mm:ss)

The time since this PIM neighbor (last) became a neighbor of the local router

Expiry Time (hh:mm:ss)

The minimum time remaining before this PIM neighbor will be aged out.

Click the Refresh button to refresh the data on the screen with the present state of the data in the router.

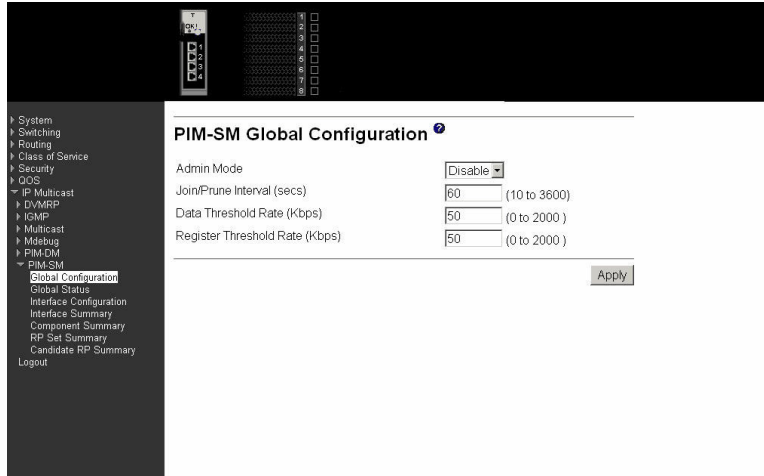
PIM-SM

This menu provides access to the following Protocol Independent Multicast - Sparse Mode (PIM-SM) entry panels:

- Global configuration
- Global status
- Interface configuration
- Interface summary
- Component summary
- RP set summary
- Candidate RP summary

Global configuration

Use this panel to configure global PIM-SM parameters.



Admin Mode The administrative status of PIM-SM in the router: either Enable or Disable.

Join/Prune Interval (secs)

The interval between the transmission of PIM-SM Join/Prune messages.

Data Threshold Rate (Kbps)

The minimum source data rate in K bits/second above which the last-hop router will switch to a source-specific shortest path tree.

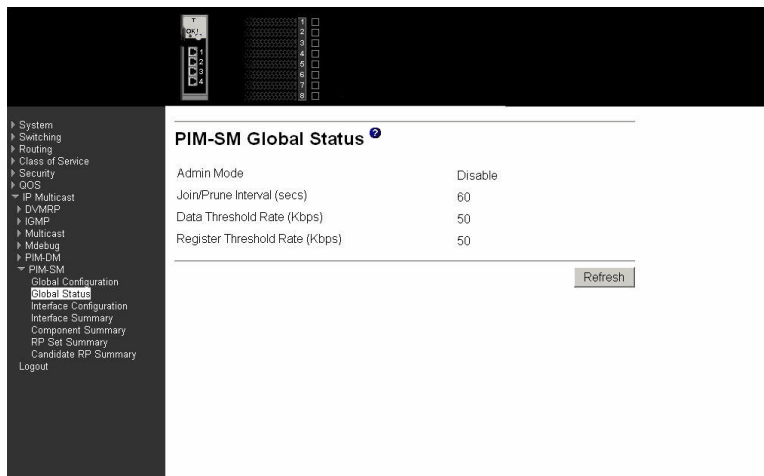
Register Threshold Rate (Kbps)

The minimum source data rate in K bits/second above which the Rendezvous Point router will switch to a source-specific shortest path tree.

Click the Refresh button to refresh the data on the screen with the present state of the data in the router.

Global status

Use this panel to display global PIM-SM parameters.



Admin Mode The administrative status of PIM-SM in the router: either Enable or Disable.

Join/Prune Interval (secs)

The interval between the transmission of PIM-SM Join/Prune messages.

Data Threshold Rate

The minimum source data rate in K bits/second above which the last-hop router will switch to a source-specific shortest path tree.

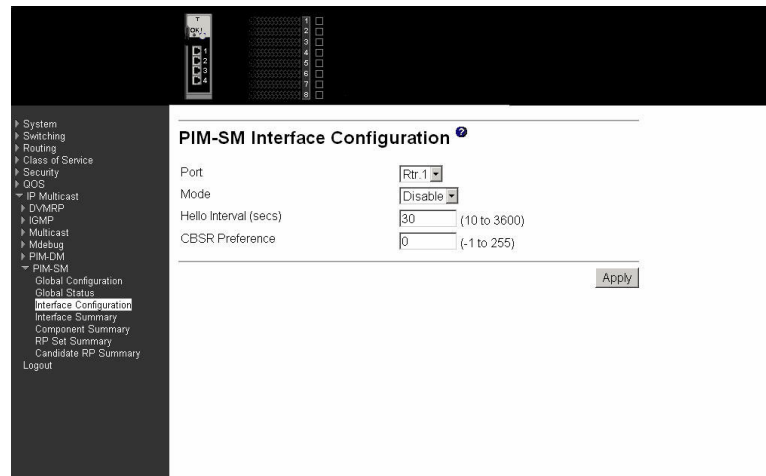
Register Threshold Rate (Kbps)

The minimum source data rate in K bits/second above which the Rendezvous Point router will switch to a source-specific shortest path tree.

Click the Refresh button to refresh the data on the screen with the present state of the data in the router.

Interface configuration

Use this panel to configure PIM-SM parameters for a specified router interface.



Port Select the physical interface for which data is to be displayed or configured.

Mode Select Enable or Disable from the pull-down menu to set the administrative status of PIM-SM in the router. The default is Disable.

Hello Interval (secs)

Enter the time in seconds between the transmission of which PIM Hello messages on this interface. The valid values are from (10 to 3600 secs). The default value is 30.

CBSR Preference

Enter the preference value for the local interface as a candidate bootstrap router. The value of -1 is used to indicate that the local interface is not a candidate BSR interface. The valid values are from (-1 to 255). The default value is 0.

Click the Apply button to update the router with the values on this screen. If you want the router to retain the new values across a power cycle, you must perform a save.

Interface summary

Use this panel to display parameters and statistics for a specified PIM-SM interface.



Port Select the physical interface for which data is to be displayed.

Mode The administrative status of PIM-SM in the router: either Enable or Disable.

Protocol State

The operational status (as opposed to configured status) of the interface. In other words, even though the PIM-SM mode can be Enable for that interface, the status is not operational until all the other prerequisites (link up, routing enabled, etc.) are satisfied for PIM-SM to operate on this interface.

IP Address The IP address of the selected PIM interface.

Net Mask The network mask for the IP address of the selected PIM interface.

Hello Interval (secs)

The frequency at which PIM Hello messages are transmitted on the selected interface.

CBSR Preference

The preference value for the local interface as a candidate bootstrap router. The value of -1 is used to indicate that the local interface is not a candidate BSR interface.

Neighbor Count

The number of PIM neighbors on the selected interface.

Designated Router

The Designated Router on the selected PIM interface. For point-to-point interfaces, this object has the value 0.0.0.0.

IP Address The IP address of the PIM neighbor for this entry.

Up Time (hh:mm:ss)

The time since this PIM neighbor (last) became a neighbor of the local router.


Expiry Time (hh:mm:ss)

The minimum time remaining before this PIM neighbor will be aged out.

Click the Refresh button to refresh the data on the screen with the present state of the data in the router.

Component summary

This panel summarizes the details of the PIM-SM components.



The screenshot shows a web-based network management interface. On the left is a dark sidebar with a navigation menu. The main content area is white and titled "PIM-SM Component Summary". It contains a table with four columns: "Component Index", "Component BSR Address", "Component BSR Expiry Time (hh:mm:ss)", and "Component CRP Hold Time (hh:mm:ss)". The table has one row with the following data: Index 1, BSR Address 9.1.1.2, Expiry Time 00:01:45, and CRP Hold Time 00:00:40. Below the table is a "Refresh" button. The sidebar menu includes items like System, Switching, Routing, Class of Service, Security, QoS, IP Multicast, DMVRP, IGMP, Multicast, Mdebug, PIM-DM, and PIM-SM (expanded to show Global Configuration, Global Status, Interface Configuration, Interface Summary, Component Summary, RP Set Summary, and Candidate RP Summary). At the top of the sidebar is a small status indicator and a vertical list of numbers 1 through 8.

Component Index	Component BSR Address	Component BSR Expiry Time (hh:mm:ss)	Component CRP Hold Time (hh:mm:ss)
1	9.1.1.2	00:01:45	00:00:40

Component Index

Unique number identifying the component index.

Component BSR Address

Displays the IP address of the bootstrap router (BSR) for the local PIM region.

Component BSR Expiry Time

Displays the minimum time remaining before the bootstrap router in the local domain will be declared.

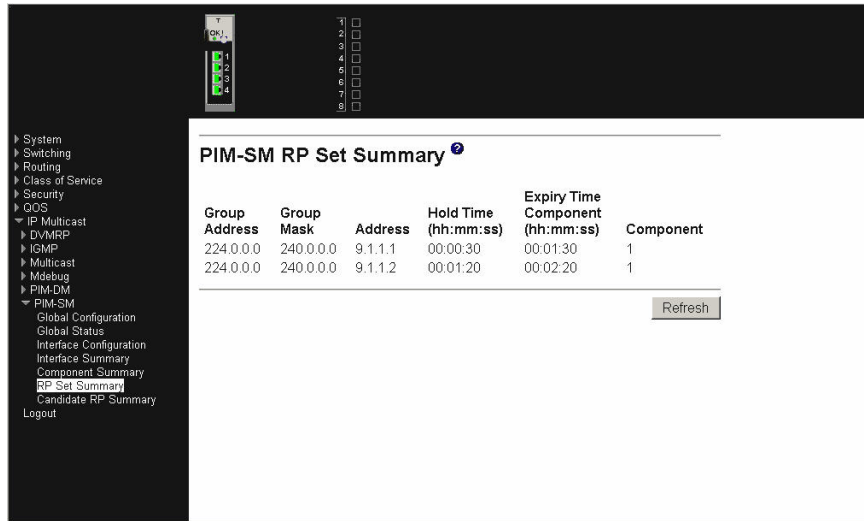
Component CRP Hold Time

The hold time of the component when it is a candidate Rendezvous Point in the local domain.

Click the Refresh button to refresh the data on the screen with the present state of the data in the router.

RP set summary

This panel displays PIM information for candidate Rendezvous Points (RPs) for IP multicast groups in a PIM-SM network.



PIM-SM RP Set Summary

Group Address	Group Mask	Address	Hold Time (hh:mm:ss)	Expiry Time Component (hh:mm:ss)	Component
224.0.0.0	240.0.0.0	9.1.1.1	00:00:30	00:01:30	1
224.0.0.0	240.0.0.0	9.1.1.2	00:01:20	00:02:20	1

Refresh

Group Address

Displays IP multicast group address.

Group Mask Displays Multicast group address mask.

Address Displays IP address of the Candidate-RP.

Hold Time The hold time of a Candidate-RP. If the local router is not the BSR, this value is 0.

Expiry Time Component

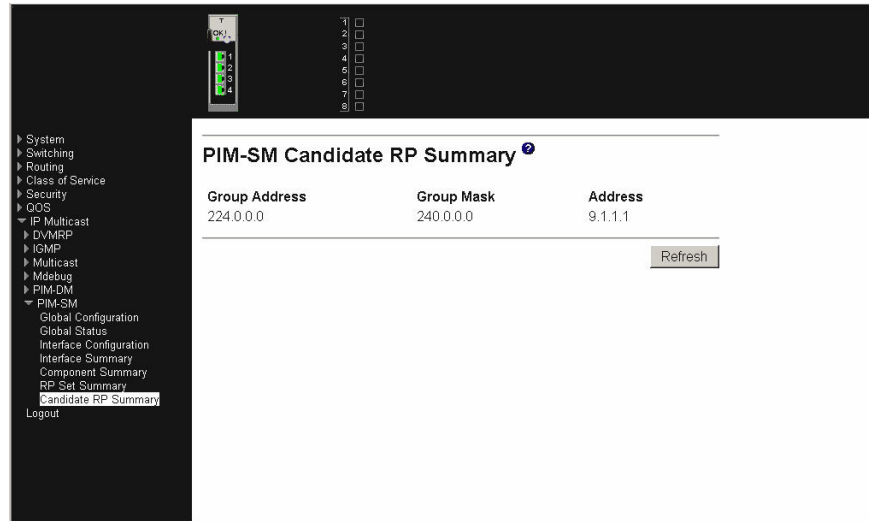
The minimum time remaining before the Candidate-RP will be declared.

Component A number which uniquely identifies the PIM-SM domain to which the router is connected.

Click the Refresh button to refresh the data on the screen with the present state of the data in the router.

Candidate RP summary

Use this panel to display the information transmitted in candidate RP advertisements.



Group Address

The group address transmitted in Candidate-RP-Advertisements.

Group Mask

The group address mask transmitted in Candidate-RP-Advertisements to fully identify the scope of the group which the router will support if elected as a Rendezvous Point.

Address

The unicast address of the interface which will be transmitted in Candidate RP advertisements.

Click the Refresh button to refresh the data on the screen with the present state of the data in the router.

Logout

When you're finished and want to exit the program simply close your browser. If you click the Logout option on the main menu you will get the message, "Please close your browser to logout."



Appendix A. Getting help and technical assistance

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This appendix contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your BladeCenter T unit, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Hardware Maintenance Manual and Troubleshooting Guide* on the IBM *BladeCenter T Documentation* CD or in the *Hardware Maintenance Manual* at the IBM Support Web site.
- Go to the IBM Support Web site at <http://www.ibm.com/pc/support/> to check for technical information, hints, tips, and new device drivers or to submit a request for information.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the online help or in the publications that are provided with your system and software. The information that comes with your system also describes the diagnostic tests that you can perform. Most xSeries and IntelliStation® systems, operating systems, and programs come with information that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the information for the operating system or program.

Using the documentation

Information about your IBM BladeCenter T unit and preinstalled software, if any, is available in the documentation that comes with your system. That documentation includes printed books, online books, readme files, and help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to <http://www.ibm.com/pc/support/> and follow the instructions. Also, you can order publications through the IBM Publications Ordering System at <http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>.

Getting help and information from the World Wide Web

On the World Wide Web, the IBM Web site has up-to-date information about IBM xSeries and IntelliStation products, services, and support. The address for IBM xSeries information is <http://www.ibm.com/eserver/xseries/>. The address for IBM IntelliStation information is <http://www.ibm.com/pc/intellistation/>.

You can find service information for your IBM products, including supported options, at <http://www.ibm.com/pc/support/>.

Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with xSeries servers, IntelliStation workstations, and appliances. For information about which products are supported by Support Line in your country or region, go to <http://www.ibm.com/services/sl/products/>.

For more information about Support Line and other IBM services, go to <http://www.ibm.com/services/>, or go to <http://www.ibm.com/planetwide/> for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

Hardware service and support

You can receive hardware service through IBM Services or through your IBM reseller, if your reseller is authorized by IBM to provide warranty service. Go to <http://www.ibm.com/planetwide/> for support telephone numbers, or in the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

Appendix B. 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.*

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.

Edition notice

© Copyright International Business Machines Corporation 2004. All rights reserved.

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

Trademarks

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

Active Memory	Predictive Failure Analysis
Active PCI	PS/2
Active PCI-X	ServeRAID
Alert on LAN	ServerGuide
BladeCenter	ServerProven
C2T Interconnect	TechConnect
Chipkill	ThinkPad
EtherJet	Tivoli
e-business logo	Tivoli Enterprise
@server	Update Connector
FlashCopy	Wake on LAN
IBM	XA-32
IBM (logo)	XA-64
IntelliStation	X-Architecture
NetBAY	Xcel4
Netfinity	XpandOnDemand
NetView	xSeries
OS/2 WARP	

Intel, MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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.

Adaptec and HostRAID are trademarks of Adaptec, Inc., in the United States, other countries, or both.

Red Hat, the Red Hat “Shadow Man” logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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

Important notes

Processor speeds indicate the internal clock speed of the microprocessor; other factors also affect application performance.

CD-ROM drive speeds list the variable read rate. Actual speeds vary and are often less than the maximum possible.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for approximately 1000 bytes, MB stands for approximately 1 000 000 bytes, and GB stands for approximately 1 000 000 000 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity may vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives available from IBM.

Maximum memory may require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven[®], including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software may differ from its retail version (if available), and may not include user manuals or all program functionality.

Product recycling and disposal


This unit contains materials such as circuit boards, cables, electromagnetic compatibility gaskets, and connectors which may contain lead and copper/beryllium alloys that require special handling and disposal at end of life. Before this unit is disposed of, these materials must be removed and recycled or discarded according to applicable regulations. IBM offers product-return programs in several countries. Information on product recycling offerings can be found on IBM's Internet site at <http://www.ibm.com/ibm/environment/products/prp.shtml>.

Battery return program

This product may contain a sealed lead acid, nickel cadmium, nickel metal hydride, lithium, or lithium ion battery. Consult your user manual or service manual for specific battery information. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to <http://www.ibm.com/ibm/environment/products/batteryrecycle.shtml> or contact your local waste disposal facility.

In the United States, IBM has established a collection process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and battery packs from IBM equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Have the IBM part number listed on the battery available prior to your call.

In the Netherlands, the following applies.



Electronic emission notices

Federal Communications Commission (FCC) statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Industry Canada Class A emission compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Australia and New Zealand Class A statement

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

United Kingdom telecommunications safety requirement

Notice to Customers

This apparatus is approved under approval number NS/G/1234/J/100003 for indirect connection to public telecommunication systems in the United Kingdom.

European Union EMC Directive conformance statement

This product is in conformity with the protection requirements of EU Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to CISPR 22/European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Taiwanese Class A warning statement

警告使用者：
這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Chinese Class A warning statement

聲 明
此為 A 級產品。在生活環境中，該產品可能會造成無線電干擾。在這種情況下，可能需要用戶對其干擾採取切實可行的措施。

Japanese Voluntary Control Council for Interference (VCCI) statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Index

A

- access control lists (ACL) 39
- authentication
 - failure 15
 - protocol 16
- auto-negotiation 57

B

- bandwidth allocation profile 39
- bay locations 13
- BOOTP/DHCP 33
- bridge protocol data unit (BPDU) 46

C

- cable
 - specifications 5
- chassis configuration and operation 12
- Class A electronic emission notice 238
- cold start 15
- configuration
 - chassis 12
 - initial 12
 - IP addresses 13
 - switch module management 68

D

- data loop 60
- data transmission rates 2
- default IP address, Ethernet switch module 13
- default remote-management mode 63
- default settings
 - spanning-tree configuration 47
- DVMRP 40

E

- efficient network operation 61
- electronic emission Class A notice 238
- Ethernet
 - activity LED 9
 - cable specifications 5
 - link LED 9
 - port locations 2
 - switch error LED 9
- Ethernet switch module, default IP address 13
- external components 10
- external-port remote management 64

F

- FCC Class A notice 238
- features 2
- full/half duplex mismatch 57

G

- GARP 25
- general requirements 2
- GMRP 26
- GVRP 25

I

- identification labels 6
- IEEE 802.1Q VLAN.
 - See virtual local area network (VLAN)
- IGMP 27, 40
- indicators 9
- information panel 10
- ingress filtering 23
- IP addresses 13
- IP mapping 30

L

- labels 6
- LAG.
 - See link aggregation
- learning the network topology 17
- LEDs
 - Ethernet activity 9
 - Ethernet link 9
 - Ethernet switch error 9
 - front view 9
 - OK 9
- link aggregation
 - introduction 28
 - static LAGs 28
- link down 15
- link up 15
- location
 - LEDs 10
 - media access control (MAC) address 6
 - ports 10
 - serial number 6

M

- major components 6
- management information base (MIB) 15
- media access control (MAC) address 6, 14

N

- negotiation 57
- network management
 - Web-based 64
- notes, important 236
- notices 6
 - electronic emission 238
 - FCC, Class A 238

O

OK LED 9
OSPF 31

P

performance features 2, 3
performance requirements 2
PIM-DM 42
PIM-SM 43
port
 locations 2, 10
 priority 49
 specifications 2
 states 47
 transition 47
 trunking 28
protocol-based VLANs 24

Q

QoS 38

R

related documentation 5
remote access 63
remote switch management 63
RIP 30
Router Discovery 34

S

serial number 6
specifications
 data transmission rates 2
 network cable 5
 ports 2
 protocols 4
 technical 2
stable topology 46
static MAC filtering 25
STP (spanning tree protocol)
 algorithm (STA) 15, 18
 illustration
 after applying STA rules 51
 before applying STA rules 50
 introduction 18
 parameters 45, 49
 port parameters 45
 port states 47
 topology change 15
switch
 bay numbers 13
 information menu 14
 management operation 11
switch information menu 14
system configuration screens 66

T

tagging 19, 23
technical specifications 2
trademarks 236
traffic class 39
transmission rates 2
traps 14

U

United States electronic emission Class A notice 238
United States FCC Class A notice 238
untagging 23

V

virtual local area network (VLAN)
 configuration 24
 introduction 19
 packet forwarding 20
 tags 21
VLAN routing 32
VRRP 33

W

warm start 15



Part Number: 13N0330

Printed in USA

(1P) P/N: 13N0330

