

Intel® Blade Server Ethernet Switch Module IXM5414E: Installation and User's Guide

A Guide for Technically Qualified Assemblers of Intel® Identified Subassemblies/Products

C66107-004



Contents

Safety	v
1 Introducing the Intel® Blade Server Ethernet Switch Module IXM5414E.....	1
Related publications.....	2
Notices and statements used in this book.....	3
Major components of the IXM5414E switch module.....	3
Specifications and features.....	4
2 Installing and Removing the Intel® Blade Server Ethernet Switch Module IXM5414E.....	9
Ethernet interface requirements.....	9
Installation guidelines.....	10
Installing the IXM5414E switch module.....	11
Removing the IXM5414E switch module.....	14
3 Information Panel LEDs and External Ports.....	17
Information panel.....	17
LEDs.....	17
4 Switch Management and Operating Concepts.....	19
Intel® Blade Server Ethernet Switch Module IXM5414E overview.....	19
Switch module management and control.....	20
Switching concepts.....	24
Security.....	35
Quality of Service (QoS).....	38
5 Web-Based Network Management.....	41
Introduction.....	41
Remotely managing the switch module.....	41
Getting started.....	42
System.....	45
Switching.....	91
Class of service.....	115
Security.....	116
QoS.....	137
Logout.....	147
6 Updating the Ethernet Switch Software.....	149
Determining the software version.....	149
Upgrading the switch software.....	150
Resetting and restarting the switch module.....	152
7 Command Line Interface Management.....	155
Command Line Interface (CLI) conventions.....	155
Remotely managing the IXM5414E switch module.....	158
IXM5414E switch module system commands.....	160
Switching configuration commands.....	187
Class of Service commands.....	205
Security configuration commands.....	206
Quality of Service (QoS) commands.....	219

Appendix A RJ-45 Pin Specifications	227
Appendix B Cable Lengths	229
Appendix C Run-time Switching Software Default Settings	231
Appendix D CLI Command Tree	239
Appendix E CLI Configuration Examples	249
IEEE 802.1w configuration example	251
VLAN configuration example	252
Link aggregation configuration example	253
IGMP snooping configuration example	254
Access Control List configuration example	255
Appendix F Troubleshooting the Spanning Tree Protocol	257
Appendix G Getting Help and Technical Assistance	275
Appendix H Notices	277

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 ασφάλειας
(safety information).

לפני שתתקינו מוצר זה, קראו את הנראות 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.

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 2:



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



Class 1 Laser Product
Laser Klasse 1
Laser Klass 1
Luokan 1 Laserlaite
Appareil À Laser de Classe 1

Statement 3:



≥ 18 kg (39.7 lb)



≥ 32 kg (70.5 lb)



≥ 55 kg (121.2 lb)

CAUTION:

Use safe practices when lifting.

Statement 4:



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 5:



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 6:



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

Statement 7:



⚠️ 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.



1 Introducing the Intel® Blade Server Ethernet Switch Module IXM5414E

Thank you for purchasing an Intel® Blade Server Ethernet Switch Module IXM5414E. This *Installation and User's Guide* contains information about:

- Setting up and installing your switch module
- Configuring your switch module

For installation details, see Chapter 2 “Installing and Removing the Intel® Blade Server Ethernet Switch Module IXM5414E” on page 9. For additional information, see the instructions in your appropriate server board chassis publications.

Your IXM5414E switch module is one of up to four switch modules that can be installed in the SBCE configuration of the blade chassis.

This high-performance IXM5414E switch module is ideally suited for networking environments that require superior microprocessor performance, efficient memory management, flexibility and reliable data storage.

Performance, reliability and expansion capabilities were key considerations in the design of your switch module. These design features make it possible for you to customize the system hardware to meet your needs today, while providing flexible expansion capabilities for the future.

The product name, machine type and serial number are located on the identification label on the side of the IXM5414E switch module. The Media Access Control (MAC) address also is located on the identification label. See “Major components of the IXM5414E switch module” on page 3 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.

Record your product information in this table.	
Product name	Intel® Blade Server Ethernet Switch Module IXM5414E
Type	_____
Model number	_____
Serial number	_____
Media access control (MAC) address	_____

Verify that the shipping carton contains an Intel® Blade Server Ethernet Switch Module IXM5414E. If the switch module is missing or damaged, contact your local reseller for replacement. Otherwise, return the switch module to its static-protective package.

/ NOTE

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

Related publications

This *Installation and User's Guide* contains setup and installation instructions for your IXM5414E switch module. This publication also provides general information about your switch module, including getting started and how to configure the switch module.

In addition to this *Installation and User's Guide*, the *Intel® Server Boards and Server Chassis Safety Information* is included with your switch module. This multilingual publication is provided in PDF on the *Resource* CD. It contains translated versions of the caution and danger statements that appear in the documentation.

Depending on your switch model, additional publications might be included on the *Resource* CD.

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 *Resource 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 IXM5414E switch module

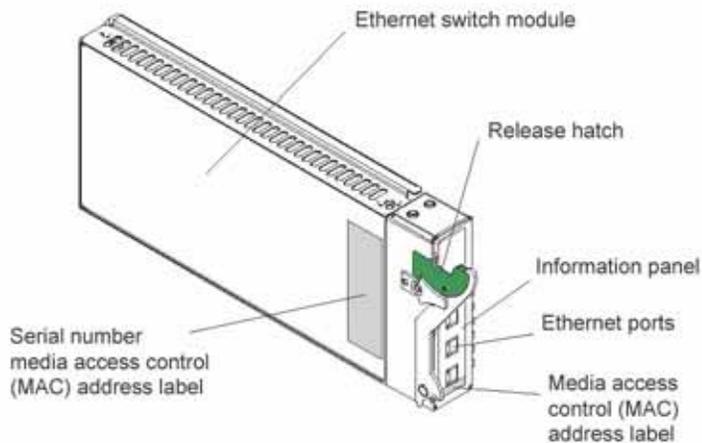
The green on components and labels on your IXM5414E switch module and on the platform identifies hot-swap or hot-plug components. You can install or remove these components while the system is running, provided that your system is configured to support this function.

The blue color on components and labels indicates touch points where a component can be gripped, a latch can be moved, and so on.

The following illustration shows the major components of your 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 3 “Information Panel LEDs and External Ports” on page 17. For more information about the MAC address, see “IP addresses and SNMP community names” on page 21.

Specifications and features

The following section provides a summary of the specifications and features for your IXM5414E switch module.

- **Ports**
 - Four external 1000BASE-T ports for making 10/100/1000 Mbps connections to a backbone, end stations, and servers
 - Fourteen internal full-duplex gigabit ports, one connected to each of the blade servers
 - Two internal full-duplex 100 Mbps ports connected to the management modules
- **Performance features**
 - Transmission method: Store-and-forward
 - Packet filtering/forwarding rate
 - Full-wire speed for all connections
 - 148k packets per second per port (for 100 Mbps)
 - 1.48m packets per second (pps) per port (for 1000 Mbps)
 - Media Access Control (MAC) address learning: Automatic update. Supports 3584 MAC address.
 - Forwarding table age time: Maximum age: 10 to 1,000,000 seconds. Default is 300 seconds
 - Support for 128 concurrent VLANs
 - Switch Topology: Star

- **Standards**

The following standards apply to the IXM5414E 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
- 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

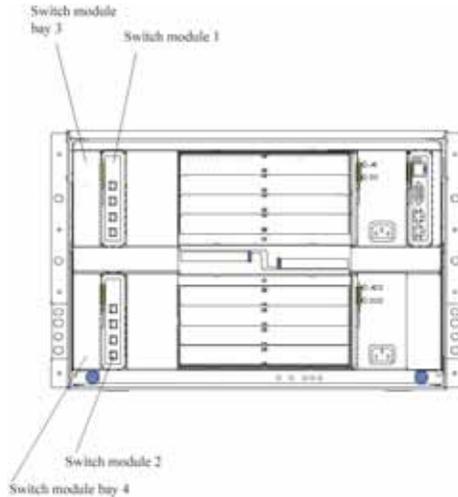
- Independent VLAN Learning (IVL) support
- Port Mirroring
- IGMP Snooping
- Static MAC Filtering
- System Facilities
 - Event and Error Logging Facility
 - Run-time and Configuration Download Capability
 - PING Utility
- Quality of Service (QOS) Support
 - Bandwidth Provisioning
 - Maximum Burst Rate (MBR)
 - Per Port (Interface)
 - Per VLAN
 - Access Control Lists
 - Source IP
 - Destination IP
 - Source L4 Port
 - Destination L4 Port
- **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-ierf-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; RSVA/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
- QOS / SNMP Support in Enterprise MIBs
 - Available through Management Module
 - Private MIBs for full configuration of ACL and Bandwidth Provisioning functionality
- **Network Cable Support**
 - 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)

2 Installing and Removing the Intel® Blade Server Ethernet Switch Module IXM5414E

The following illustration shows the I/O module bay locations in the SBCE platform.



Attention: To maintain proper system cooling, each I/O module bay must contain either a module or a filler module; each blade bay must contain either a blade or a filler blade.

Ethernet interface requirements

The SBCE platform supports a minimum of one hot-swap Ethernet switch module in I/O module bay 1. This switch module is a fully functional four-connector Ethernet switch that provides a network connection to Ethernet Link 1 in all the blade servers in the SBCE. To provide a network connection for Ethernet Link 2 in each blade server, install an Ethernet switch module in I/O module bay 2.

If you install an interface option on any blade server, you must install a hot-swap switch module of the same interface type in I/O module bay 3 to obtain connection 1 for the interface option. To provide connection 2 for the interface option, install a switch module of that interface type in I/O module bay 4. The switch modules in I/O module bays 3 and 4 provide connections to all the interface options in the SBCE.

Important: The switch modules in I/O module bays 3 and 4 and all blade server interface options in the SBCE must use the same interface type. For example: if you install an Ethernet interface option on a blade server, the switch modules that you install in I/O module bays 3 and 4 must be Ethernet. All other interface options in the SBCE must also be Ethernet interface options.

The following table summarizes the application for each switch module.

I/O module bay	Switch-module function
1	Connection 1 (Ethernet Link 1) for all blade servers in the SBCE

I/O module bay	Switch-module function
2	Connection 2 (Ethernet Link 2) for all blade servers in the SBCE
3	Connection 3 (from all blade server interface options in the SBCE)
4	Connection 4 (from all blade server interface options in the SBCE)

For additional information, see the *Intel® Blade Server Chassis SBCE : Installation and User's Guide* on the Resource CD.

Installation guidelines

Before you begin installing the IXM5414E switch module in your SBCE, read the following information:

- Become familiar with the safety and handling guidelines specified under Appendix H “Notices” on page 277 and “Handling static-sensitive devices”, and read the safety statements in the SBCE option publications.
- The green color on components and labels in your SBCE identifies hot-swap or hot-plug components. You can install or remove hot-swap modules while the SBCE is running. For complete details about installing or removing a hot-swap or hot-plug component, see the detailed information in this chapter.
- The blue color on components and labels identifies touch points where you can grip a component, move a latch, and so on.
- You do not need to turn off the SBCE to install or replace any of the hot-swap modules on the rear of the SBCE.

System reliability considerations

Attention: To help ensure proper cooling and system reliability, make sure that:

- Each of the I/O module bays on the rear of the SBCE has either a module or filler module installed.
- A removed hot-swap module is replaced with an identical module or filler module within 1 minute of removal.
- Cables for the optional modules are routed according to the illustrations and instructions in this document.

Handling static-sensitive devices

Attention: Static electricity can damage electronic devices and your system. To avoid damage, keep static-sensitive devices in their static-protective packages until you are ready to install them.

To reduce the possibility of electrostatic discharge, observe the following precautions:

- Limit your movement. Movement can cause static electricity to build up around you.
- Handle the device carefully, holding it by its edges or its frame.
- Do not touch solder joints, pins, or exposed printed circuitry.

- Do not leave the device where others can handle and possibly damage it.
- While the device is still in its static-protective package, touch it to an unpainted metal part of the SBCE platform for at least two seconds. (This drains static electricity from the package and from your body.)
- Remove the device from its package and install it directly into your SBCE without setting it down. If it is necessary to set the device down, place it in its static-protective package. Do not place the device on your SBCE platform or on a metal table.
- Take additional care when handling devices during cold weather because heating reduces indoor humidity and increases static electricity.

Installing the IXM5414E switch module

Statement 8:



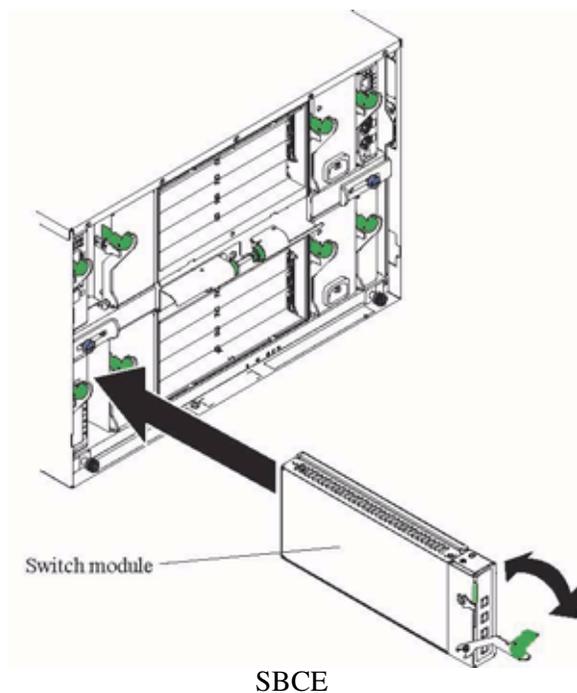
xxCAUTION:

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



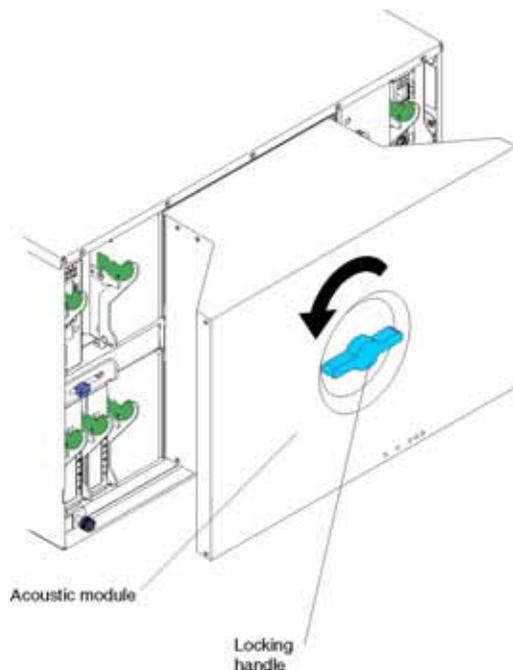
Hazardous voltage, current, and energy levels are present inside any component that has this label attached. There are no serviceable parts inside these components. If you suspect a problem with one of these parts, contact a service technician.

The following illustrations show how to install a switch module in the rear of the SBCE platform.

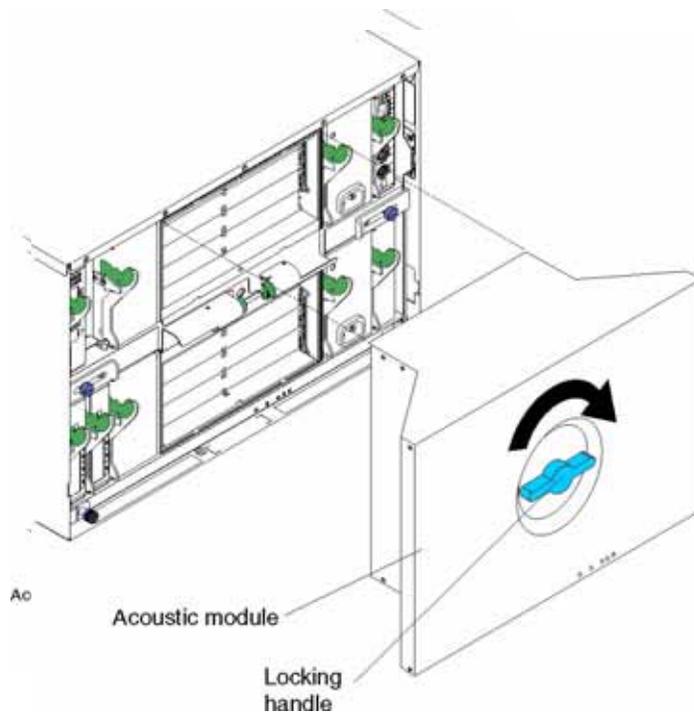


Complete the following steps to install the IXM5414E switch module.

1. Review the information in “Safety” on page v and in “Installation guidelines” on page 10.
2. Remove the acoustic attenuation module, if installed, from the rear of the SBCE platform. The following illustrations show how to remove the module from the SBCE platform.



3. Select an I/O module bay in which to install the switch module, in accordance with the instructions in “Ethernet interface requirements” on page 9.
4. Remove the filler module from the selected I/O module bay. Store the filler module for future use.
5. If you have not already done so, touch the static-protective package that contains the switch module to an unpainted metal part of the SBCE platform for at least two seconds.
6. Remove the switch module from its static-protective package.
7. Ensure that the release latch on the switch module is in the open position (perpendicular to the module).
8. Slide the switch module into the appropriate I/O module bay until it stops.
9. Push the release latch on the front of the switch module to the closed position.
10. Make sure that the LEDs on the switch module indicate that it is operating properly. Verify that:
 - The DC power LED and the ac power LED on each power module are lit.
 - The OK LED on each management module is lit.
 - The OK LED on each switch module is lit.
11. If you have other switch modules to install, do so now; otherwise, continue with the next step.
12. Attach any cables required by the switch module. For the location of the connectors on the SBCE platform, see *Intel® Server Chassis SBCE Installation and User’s Guide* on the *Resource CD*.
13. Replace the acoustic attenuation module if you removed it in Step 2. The following illustration shows how to replace the acoustic attenuation module in the SBCE platform.



Removing the IXM5414E switch module

Statement 8:



xxCAUTION:

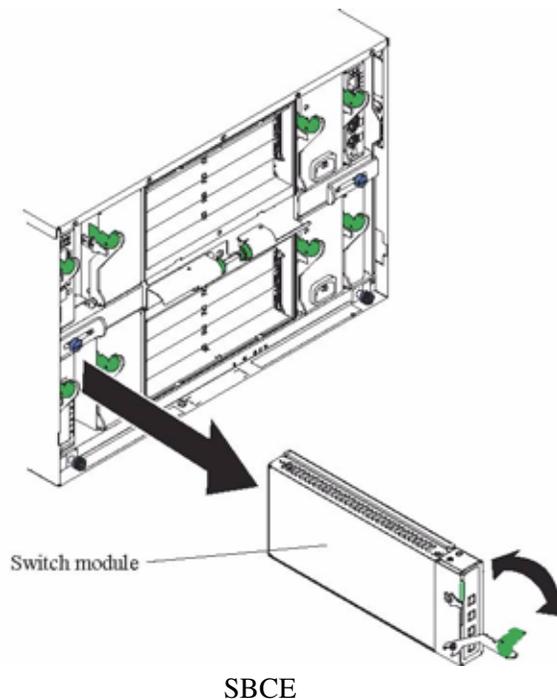
Never remove the cover on a power supply or any part that has the following label attached.



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.

Complete the following steps to remove the IXM5414E switch module.

1. Select an appropriate I/O module bay from which to remove a switch module, in accordance with the instructions in “Ethernet interface requirements” on page 9.
2. Unplug any cables from the selected switch module.
3. For the SBCE platform, pull the release latch toward the side of the switch module as shown in the illustration below. The module moves out of the I/O module bay about 0.64 cm (0.25 inch).



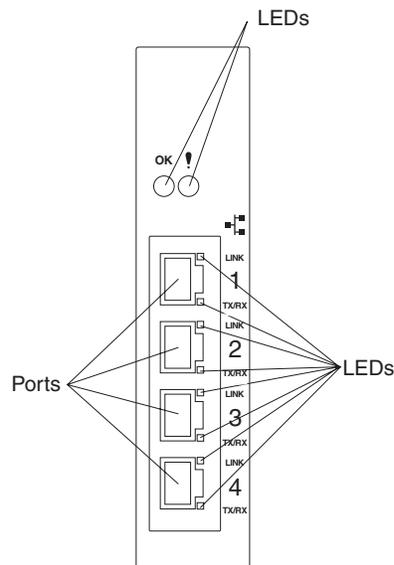
4. Slide the switch module out of the I/O module bay and set it aside.
5. Place either another switch module or a filler module in the I/O module bay within 1 minute.
6. If you placed another switch module in the I/O module bay, reconnect any cables that you unplugged in Step 2.
7. Replace the acoustic attenuation module option if you removed it in step 1.

3 Information Panel LEDs and External Ports

This chapter describes the information panel and LEDs (also known as indicators) on the Intel® Blade Server Ethernet Switch Module IXM5414E. This chapter also identifies the external ports on the information panel.

Information panel

The information panel of the IXM5414E switch module consists of LEDs and four external 1000BASE-T ports, as shown in the following illustration.

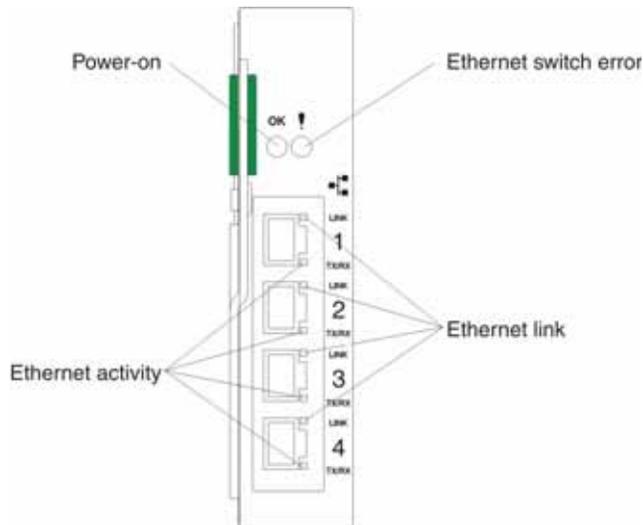


The Intel® Blade Server Ethernet Switch Module IXM5414E contains:

- Comprehensive LEDs, which display the status of the switch module and the network (see “LEDs”).
- Fourteen internal ports, one connected to each of the processor blades.
- Two internal full-duplex 10/100 Mbps ports connected to the management module.
- Four external 1000BASE-T Ethernet ports for 10/100/1000 Mbps connections to external Ethernet devices such as backbones, end stations and servers. These ports are identified as Ext1, Ext2, Ext3 and Ext4 in the switch configuration menus and are labeled 1 through 4 (from top to bottom) on the switch module, as shown in the preceding illustration.

LEDs

The LEDs on the information panel of the IXM5414E switch module include OK, !, Ethernet link, and Ethernet activity. The following illustration shows the LEDs on the switch module. A description of each LED follows the illustration.



Notes:

1. The illustrations in this document may differ slightly from your hardware.
2. An amber LED illuminates when a system error or event has occurred. To identify the error or event, check the LEDs on the information panel of the switch module.

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

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

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

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

4 Switch Management and Operating Concepts

This chapter discusses many of the concepts and features used to manage the Intel® Blade Server Ethernet Switch Module IXM5414E and the concepts necessary to understand how it functions. In addition, this chapter explains many important points regarding these features.

Configuring the switch module to implement these concepts and use its many features is discussed in detail in the following chapters.

Intel® Blade Server Ethernet Switch Module IXM5414E 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 used in this document. However, it is important that you also understand the operating environment of the SB-HE platform with regard to the internal switches.

IXM5414E switch modules are hot-swappable subsystems that provide Ethernet switching capabilities within the chassis of the SB-HE platform. The primary purpose of the switch module is to provide Ethernet interconnectivity among the processor blades, management modules and the external network infrastructure.

The SB-HE platform may be configured with up to four independent switch modules, supporting up to fourteen server blades. Ports 1 through 14 on the switch module correspond to server blades 1 through 14, respectively (numbered left to right when viewed from the front of the chassis). Each switch module has four external 10/100/1000 Mbps Ethernet ports for connection to the external network infrastructure. These ports are identified as Ext.1, Ext.2, Ext.3 and Ext.4 in the switch module configuration menus and are labeled 1 through 4 on the switch module (see Chapter 3 “Information Panel LEDs and External Ports” on page 17 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 IXM5414E switch module has been pre-configured with default parameter settings that can be used with some typical installations. Most installations will need some configuration of parameters. Information on initial software configuration can be found in “Remotely managing the IXM5414E switch module” on page 158 and “IXM5414E switch module system commands” on page 160.

Chassis configuration and operation

Each IXM5414E switch module is an integral subsystem within an overall SB-HE platform. For additional platform level information, see the applicable *Installation and User's Guide* publications on the *Resource CD*. Each chassis includes one or two management modules (MM) as the central element for overall chassis management and control. The switch module includes 100-Mbps internal Ethernet ports that can only be accessed by the management modules. To prevent inadvertent changes, this management port is “hidden” and does not appear in the port configuration and status screens. The factory default settings will *only* permit management and control access to the switch module through the 10/100 Mbps Ethernet port on the management module. You can use the four external 10/100/1000 Mbps Ethernet ports on the 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 and User's Guide* publications on the *Resource CD* for more information).

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 IXM5414E switch module. Complete the following initial configuration steps:

1. Connect the Ethernet port of the management module to a 10/100 Mbps network (with access to a management station) or directly to a management station.
2. Initially configure the management module with the appropriate IP addresses for network access (see the applicable *SB-HE Installation and User's Guide* publications on the *Resource CD* for more information).
3. From the management-module Web interface, click I/O Module Tasks, click Management; then, click the bay in which the switch module is installed.
4. Click **Advanced Management** under the selected bay and make sure that the following Ethernet switch module features are enabled:
 - External ports
 - External management over all ports
 - Preserve new IP configuration on all resets
5. Click **Save**.

Note: When management of the Ethernet switch module is enabled through the four external ports, the switch module will acquire its IP address from a Dynamic Host Configuration Protocol (DHCP) server when the switch module is turned on or reset.

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

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

The switch module supports three primary management and control user interfaces. A built-in Web browser interface is the primary interface (see Chapter 5 "Web-Based Network Management" on page 41 for detailed information). The Web browser interface can be invoked from the management and configuration utility program, along with the Telnet interface that provides a Command Line Interface

(CLI) (see Chapter 7 “Command Line Interface Management” on page 155 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 private MIB objects through SNMP protocols.

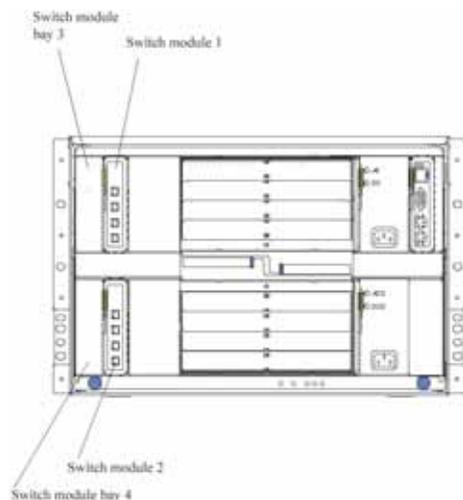
IP addresses and SNMP community names

Each 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 transmission control protocol/Internet protocol (TCP/IP) application. The switch module default IP address is 10.90.90.9 x , where x depends on the number of the I/O module bay into which you have installed the switch module, as shown in Table 1.

Table 1. Default IP addresses based on I/O module bay numbers

I/O module bay number	Default IP address
Switch Module Bay 1	10.90.90.91
Switch Module Bay 2	10.90.90.92
Switch Module Bay 3	10.90.90.94
Switch Module Bay 4	10.90.90.97

The following illustration shows the I/O module bay locations.

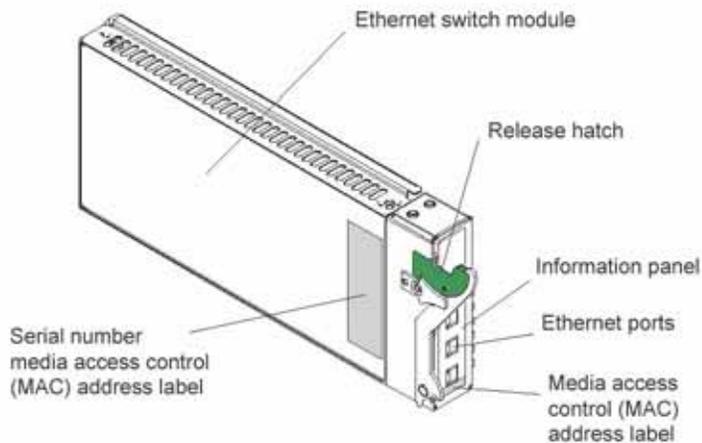


You can change the default switch module IP address to meet the requirements of your networking address scheme.

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 switch MAC address can also be displayed using CLI command **show inventory** or from the Web Interface.

In addition, you can also set an IP address for a gateway router. This becomes necessary when the network management station and switch modules are located on different IP networks, requiring management packets to go through a router to reach the network manager.

For security, you can specify the IP addresses of the network managers that are permitted to manage the switch module using the **config snmpcommunity ipaddr** CLI command or the Web Interface equivalent. 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, someone accidentally turned off the switch module) or less serious, such as a port-status change. The switch module generates traps and sends them to the network manager (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 sent from the switch module; they may then need to 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 like passwords in that the community string entered for a given IP address must be used in the management station software, otherwise a trap will be sent.

The following are trap types that the switch module can send to a trap recipient:

Cold start This trap indicates that the switch module has been turned on and initialized such that software settings are reconfigured and hardware systems are restarted. A cold start is different from a factory reset in that configuration settings saved to

nonvolatile random-access memory (NVRAM) are used to reconfigure the switch module.

Warm start This trap indicates that the switch module has been restarted; however, the power-on self-test (POST) is skipped.

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 has 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 are stored in the switch module in the management information base (MIB). The switch module uses the standard MIB-II management information base module. Consequently, values for MIB objects can be retrieved using any SNMP-based network management software. In addition to the standard MIB-II module, the switch module also supports its own proprietary enterprise MIB as an extended management information base. This MIB can also be retrieved by specifying the object identifier (OID) of the MIB as the network manager. MIB values can be either Read-only or Read/Write.

Read-only MIB variables 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 kilobytes) has been received and forwarded through a port.

Read/Write MIBs variables 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 a third-party vendor's SNMP software to manage the switch module, a diskette listing the switch module proprietary enterprise MIBs can be obtained by request. If your software provides functions to browse or modify MIBs, you can also get the MIB values and change them (if the attributes of the MIBs permit the write operation). However, this process can become complicated, because you must know the MIB OIDs and retrieve them one by one.

Port mirroring

The IXM5414E switch module enables you to copy packets that were transmitted and received on a source port and to redirect the copies to another target port. The source port can be either one of the four 10/100/1000 Mbps external ports, or one of the fourteen internal blade server ports. The target port is where you will connect a monitoring/troubleshooting device, such as a sniffer or an RMON probe. The target port must be one of the four 10/100/1000 Mbps external ports.

You can attach a monitoring device to the mirrored port, such as a sniffer or an RMON probe, to view details about the packets that pass through the first port. This is useful for network monitoring and troubleshooting purposes.

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. SNMP can be used to perform many of the same functions as a directly connected console, or can be used within an integrated network management software package such as IBM® NetView or Hewlett Packard OpenView. SNMP performs the following functions:

- Sending and receiving SNMP packets using the IP protocol
- 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 management station (a designated computer on the network). The SNMP agent and the user program both employ 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 which function like passwords. The remote user SNMP application and the switch module's SNMP agent must use the same community string. SNMP community strings of up to 20 characters can be entered using the CLI **snmp community** commands described in Chapter 7 "Command Line Interface Management" on page 155.

Switching concepts

This section introduces the concepts and protocols relevant to the switching functionality of the Intel® Blade Server Ethernet Switch Module IXM5414E.

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 number through which they can be reached. Packets sent to known addresses are therefore transmitted only through relevant destination ports, thus reducing network traffic. For example, if port 1 receives a packet 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, those learned by the switch by observing network traffic, are deleted from the table if they are not accessed within the aging time. Static entries, those entered by a network administrator, are not subject to the aging process.

The aging time can be from 10 to 1,000,000 seconds, with a default value of 300 seconds. Setting the value too high could mean that some entries in the table become out of date, causing the switch module to make incorrect packet-forwarding decisions. If the aging time is too short, however, entries may 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 ports on the switch, thus unnecessarily increasing 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 enables the duplicate links to be used in the event of a failure of the primary link. 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 also accomplished automatically, without operator intervention.

This automatic network reconfiguration provides maximum uptime to network users. However, the concepts of the Spanning Tree Algorithm and Protocol are complicated and complex subjects and must be fully researched and understood. 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

Improper configuration of the switch module's external ports or improper 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 created by the Spanning Tree Algorithm is referred to as the Common Spanning Tree (CST) in some of the commands described in this document.

The original Spanning Tree Algorithm 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 which are full-duplex and ports which 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 running certain protocols (e.g. LLC2 and NETBEUI) that are sensitive to the slightly increased probability of frame misordering. The IXM5414E switch module defaults to IEEE 802.1D operation, but can be configured to use the algorithm and protocols 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 IXM5414E switch module defaults to IEEE 802.1D compatibility mode operation, but can be configured to use the algorithm and protocols defined in IEEE 802.1w instead. Where this document refers to IEEE 802.1D, you should be aware that the reference is to IEEE 802.1D compatibility mode.

For additional information about both forms of the Spanning Tree Protocol, see Appendix H on page 277.

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 chassis. 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.

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

Notes about VLANs on the IXM5414E 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 IXM5414E switch module implements IEEE 802.1Q VLANs, which require 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:

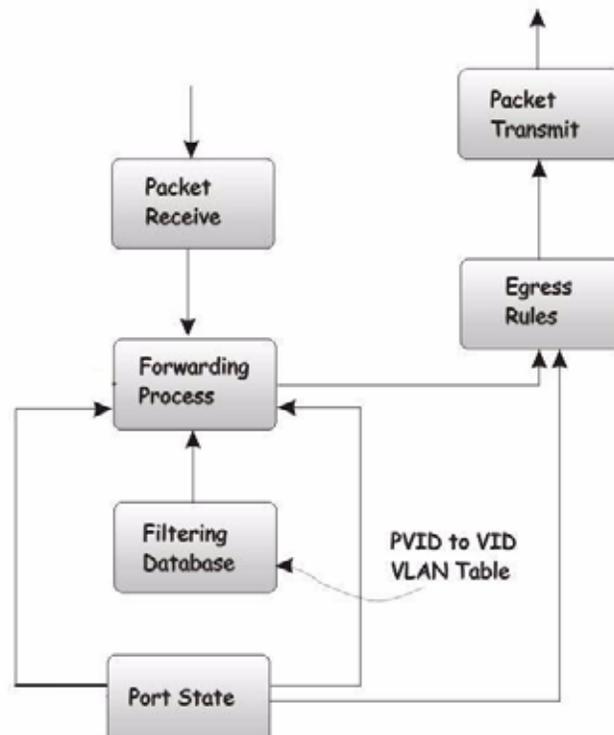
Forwarding rules between ports

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 24. For more information about port VLAN IDs (PVIDs), see “Port VLAN ID” on page 29. For more information about tagging and untagging, see “Tagging and untagging” on page 29. For more information about port states, see “IEEE 802.1D STP port states” on page 259 and “IEEE 802.1w STP port states” on page 260.

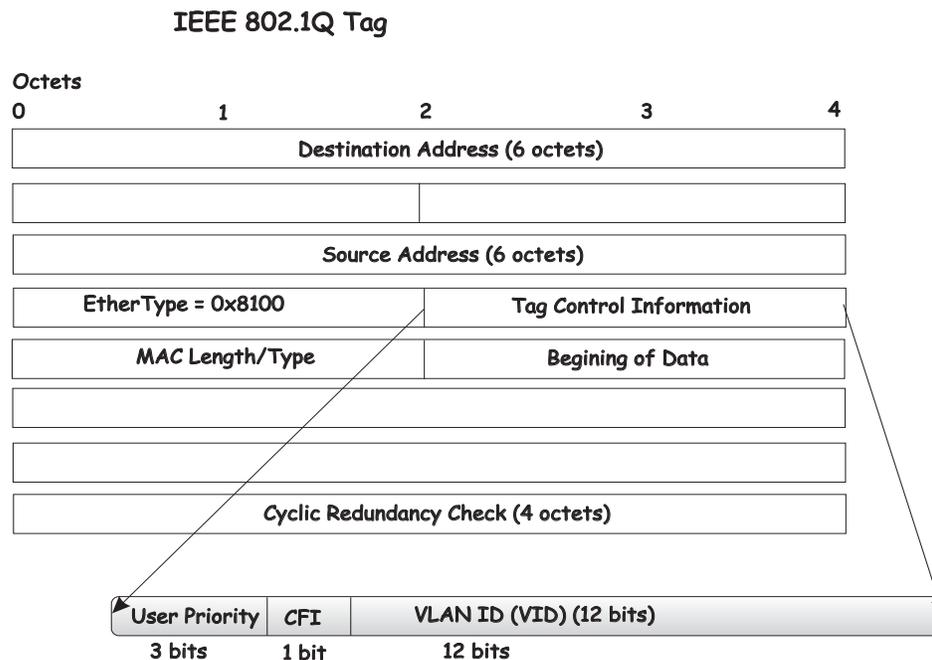
802.1Q Packet Forwarding



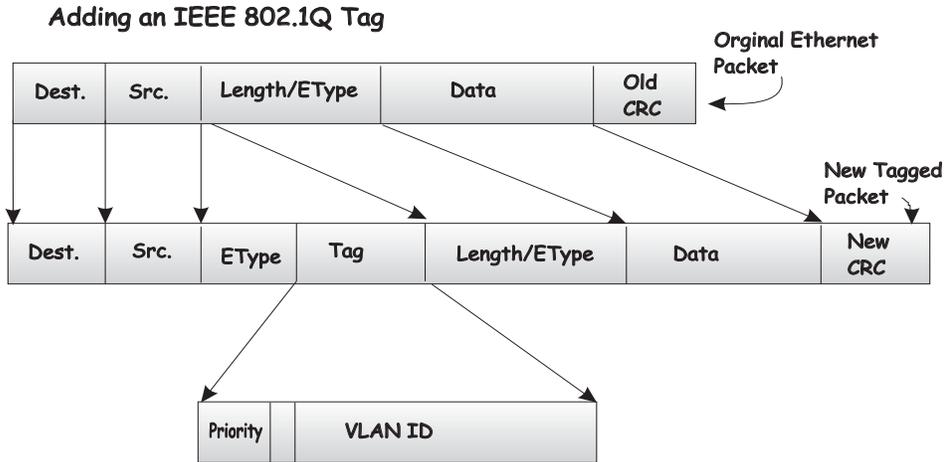
IEEE 802.1Q VLAN tags

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) and 12 bits of VLAN ID (VID). The 3 bits of user priority are used according to the protocols 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.



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.



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 given 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 can have as many VIDs 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.

Egress rules

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.
- 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.

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 defined in the filter.

Static entries are never aged and can only be removed 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 a given 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 Intel® Blade Server Ethernet Switch Module IXM5414E 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 GPDU. 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.

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 given 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 may collapse IP multicast group memberships into a single Ethernet multicast group.

A traditional Ethernet network may 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 may be removed from an IGMP group in response to an IGMP Leave Group message.

Link aggregation (LAG)

The Intel® Blade Server Ethernet Switch Module IXM5414E 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 SB-HEs 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 only be assigned to one trunk group, whether 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 Web device panel display.

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 LACPDU.

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.

Security

IEEE 802.1X

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 may 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 IXM5414E 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.

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 IXM5414E 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 may 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 appropriate 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

Table 2. Secure Shell Feature Details

SSH Feature	Component Type
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 Intel® Blade Server Ethernet Switch Module IXM5414E 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 may 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 may be assigned to multiple TCs. Once 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
- 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 an interface. Filtering is done on inbound traffic.

An ACL rule may 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.

5 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 your Intel® Blade Server Ethernet Switch Module IXM5414E, be sure that the management modules in your SBCE platform are properly 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 *Resource CD* for more information.

Introduction

The Intel® Blade Server Ethernet Switch Module IXM5414E 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 Chapter 7 “Command Line Interface Management” on page 155 for additional information.

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

Remotely managing the switch module

The IXM5414E 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 SBCE 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 IXM5414E 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 *Resource 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 IXM5414E 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 only be enabled through the management module configuration interface. Once this mode is enabled, the external Ethernet ports will support both management traffic and SBCE application data traffic. Also, the IXM5414E 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 SBCE subsystems through the management module. However, management access to the IXM5414E 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 IXM5414E switch module for this mode of operation.

The two previously described modes are only applicable to the IXM5414E switch module. The management module can only be remotely accessed through the 10/100 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.

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.

Note: When the switch module is attached to the 100 Mbps port on the management module, use the statically assigned switch module IP address. When management of the Ethernet switch module is enabled through the four external ports, use the IP address the switch module acquired from a Dynamic Host Configuration Protocol (DHCP) server when the switch module was turned on or reset.

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



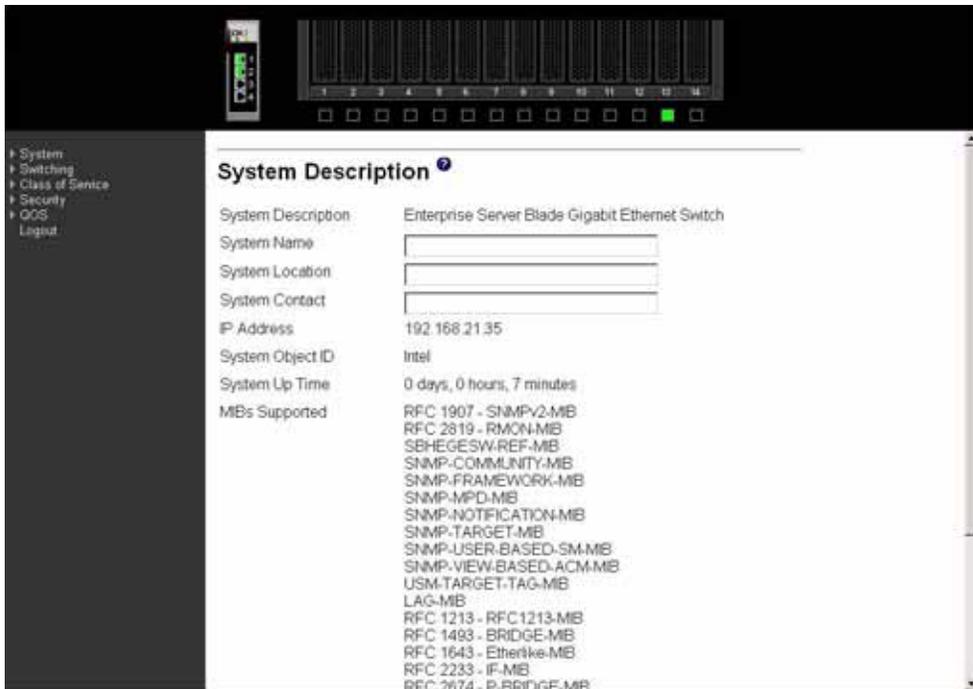
Click on Login, and a dialog box similar to the following will open:



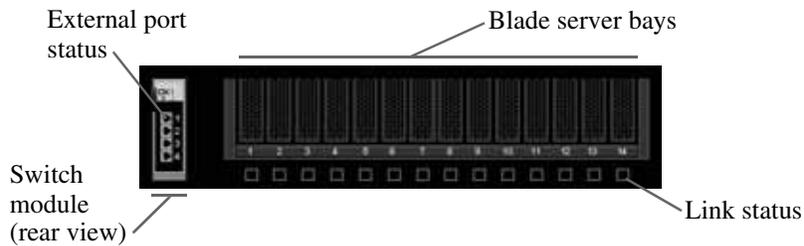
Enter “USERID” in the User name field and enter “PASSWORD” (with a zero in place of the O) in the Password field. Click the OK button. This opens the main page in the management module.

/ NOTE

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



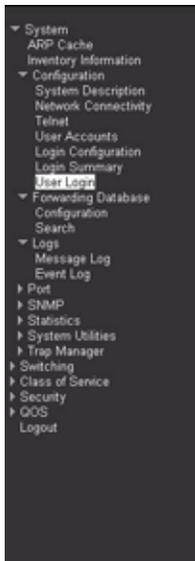
If java mode is enabled for the switch (the default is enabled) the top panel shows a real-time information-panel display of the switch module, as shown below. You can change the java mode on the Network Connectivity Configuration menu (See “Network connectivity” on page 49).



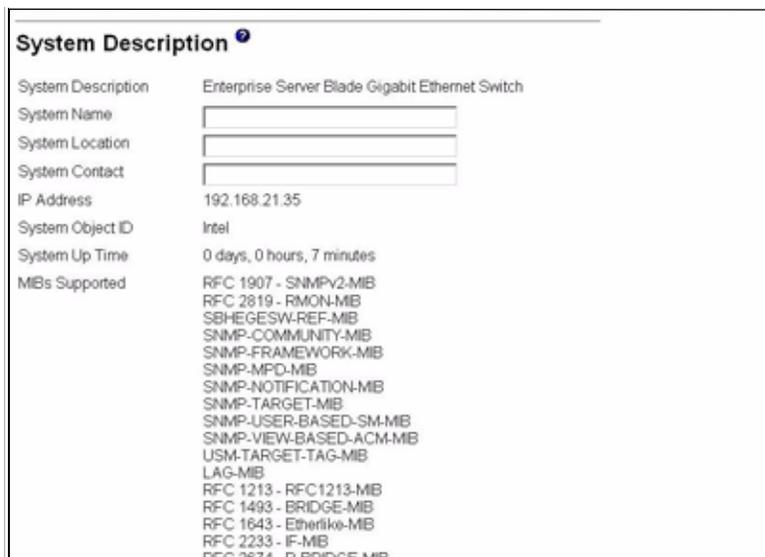
The panel on the left side of the screen displays the main menu. The main menu contains:

- System
- Switching
- Class of service
- Security
- QOS
- 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 “System description” on page 48.



System

The 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.

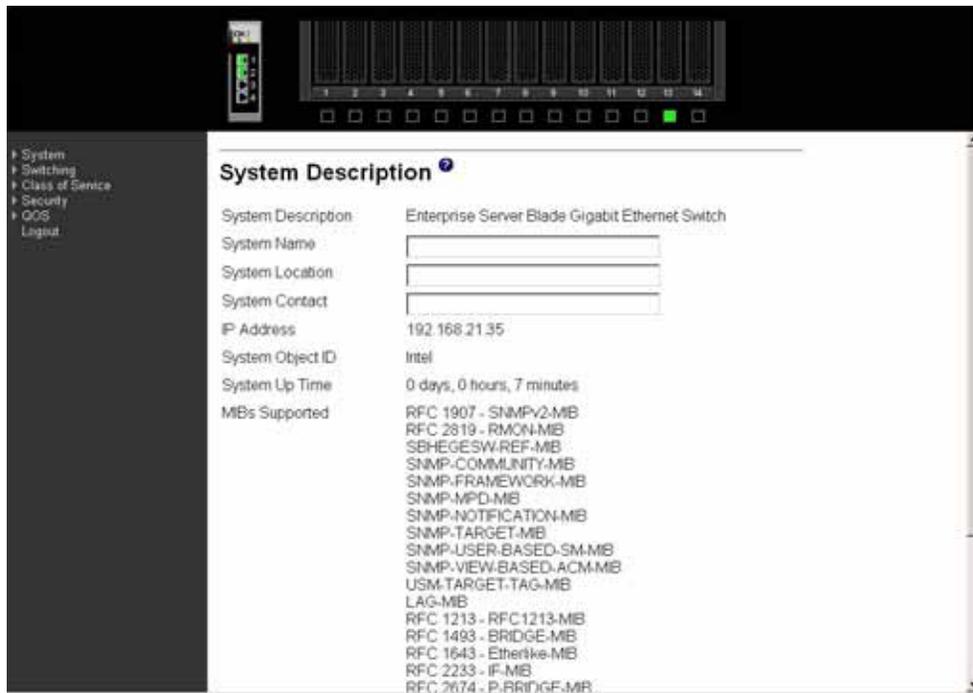
Configuration

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 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. May be up to 31 alphanumeric characters. The factory default is blank.

System Contact

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

IP Address

The IP address of the interface. The factory default value 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. (See Table 1. “Default IP addresses based on I/O module bay numbers” on page 21)

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.

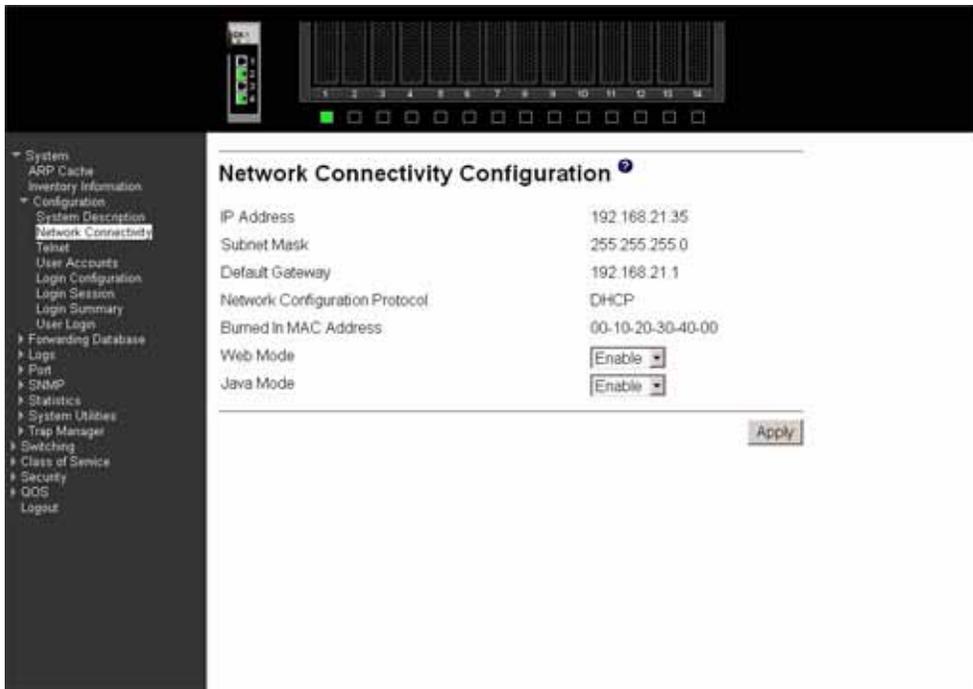
Network connectivity

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, the switch must first be configured with its IP information (IP address, subnet mask and default gateway).

Once 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, where x is determined by the number of the I/O-module bay into which you have installed the Ethernet switch module (see Table 1. “Default IP addresses based on I/O module bay numbers” on page 21).

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 configuration methods are DHCP, BootP and none. The factory default method is none. When management of the Ethernet switch module is enabled through the four external ports, the switch module will acquire its IP address from a Dynamic Host Configuration Protocol (DHCP) server when the switch module is turned on or reset.

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

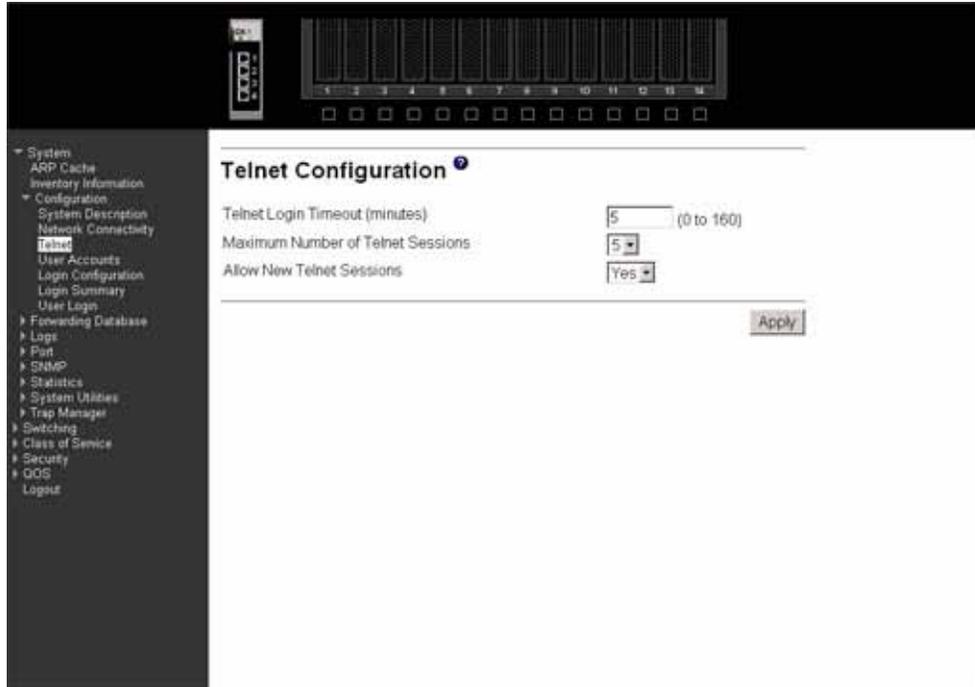
Web Mode Specify whether the switch may be accessed from a web browser through TCP port 80. If you choose to Enable web mode you will be able to manage the switch from a web browser. The factory default is Enabled.

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.



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 may 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.

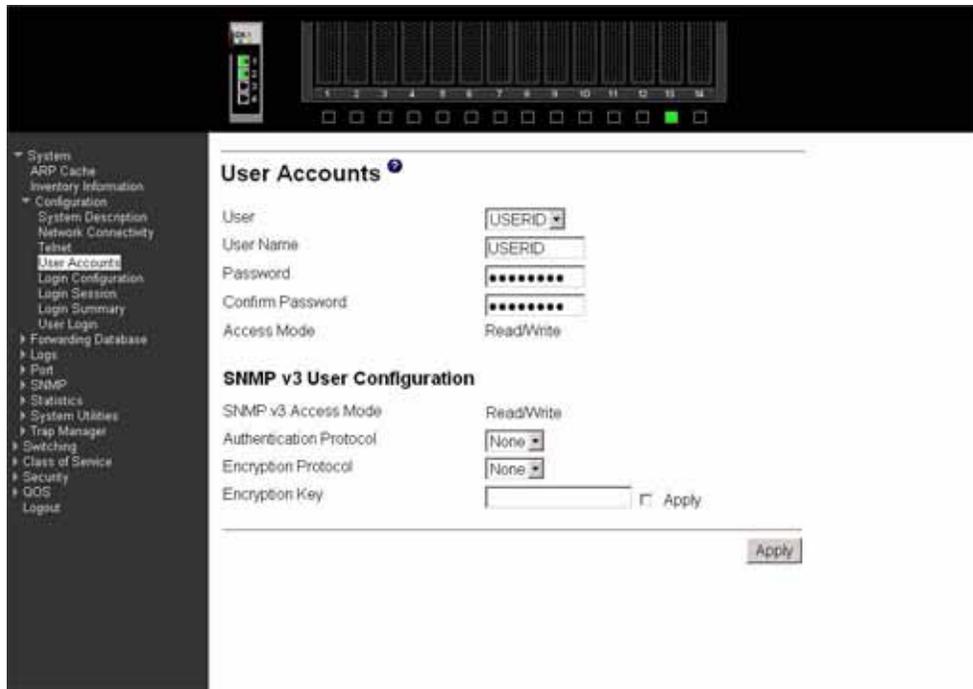
Allow New Telnet Sessions

Indicates whether new Telnet sessions are allowed. If you set this to no, new Telnet and SSH 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 only 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 may 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 may be up to 16 characters long. If None is specified for the protocol, the Encryption Key is ignored.

Encryption Key

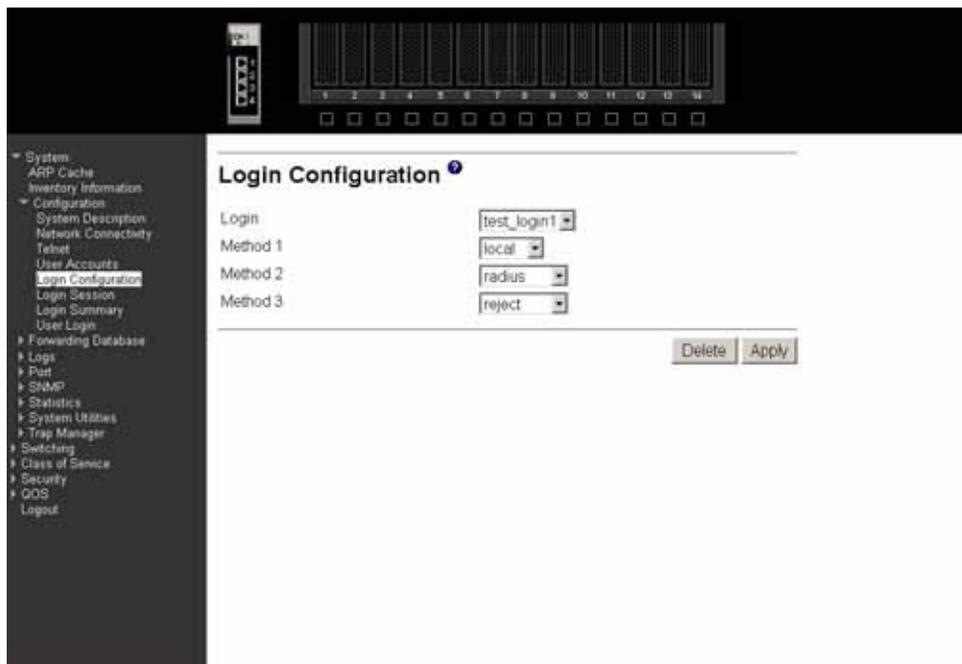
If you selected DES in the Encryption Protocol field, enter the SNMPv3 Encryption Key here. Otherwise this field is ignored. Valid keys are 0 to 15 characters long. The Apply checkbox must be checked in order to change the Encryption Protocol and Encryption Key.

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 may not delete. All newly created users are also assigned to the defaultList until you specifically assign them to a different list.



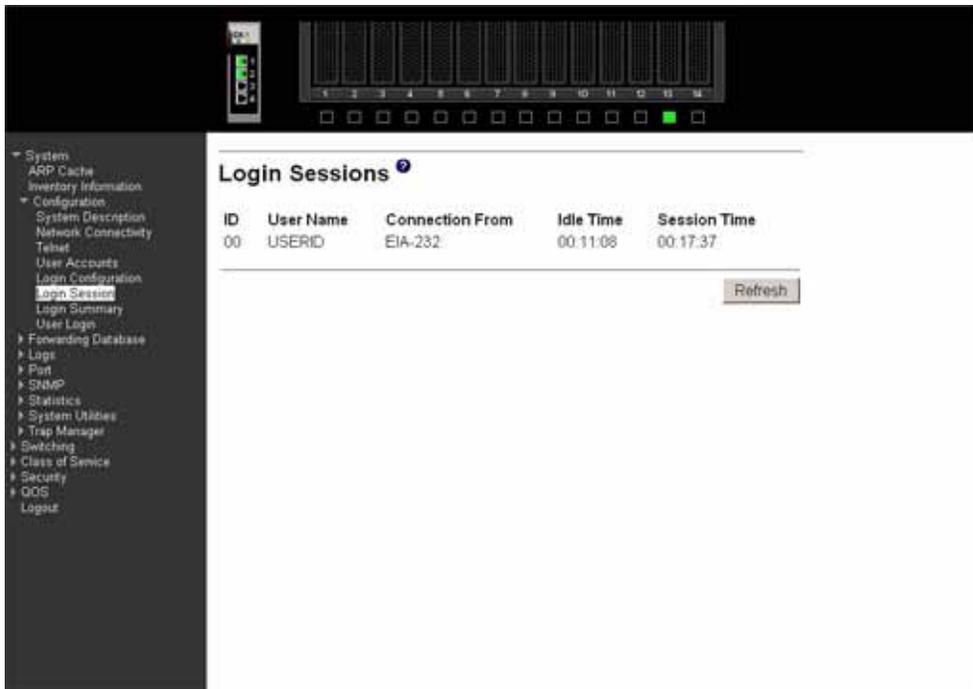
- Login** Select the authentication login list you want to configure. Select Create to define a new login list. When you create a new login list, Local is set as the initial authentication method.
- Login Name** If you are creating a new login list, enter the name you want to assign. It can be up to 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 may 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 only use this button if you have Read/Write access.

Login session

This panel displays the details for all user login sessions.



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.

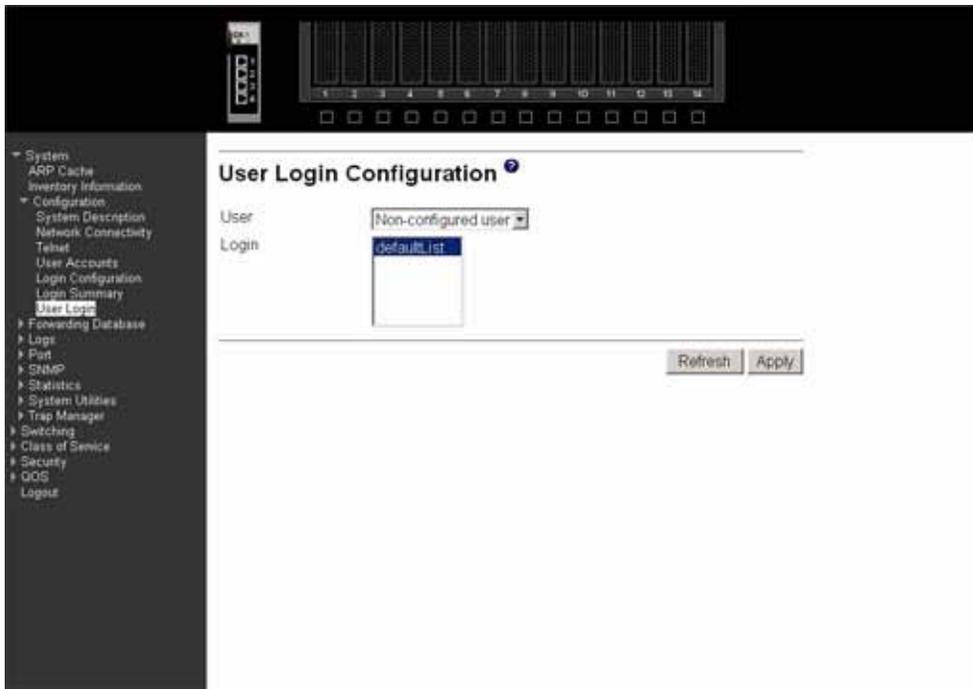


- Login** Identifies the authentication login list summarized in this row.
- Method List** The ordered list of methods configured for this login list.
- Login Users** The users you assigned to this login list on the User Login Configuration screen. This list is used to authenticate the users for system 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 may not 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.

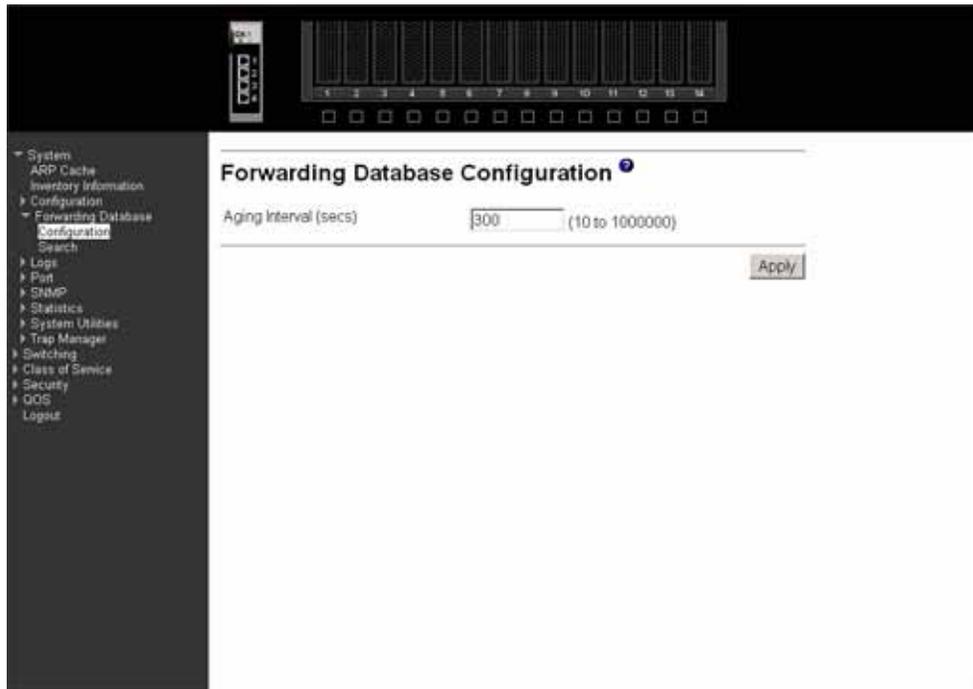
Click the Apply button to cause the changes made on this screen to take effect on the switch click. If you want the switch to retain the new values across a power cycle, you must perform a save.

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.



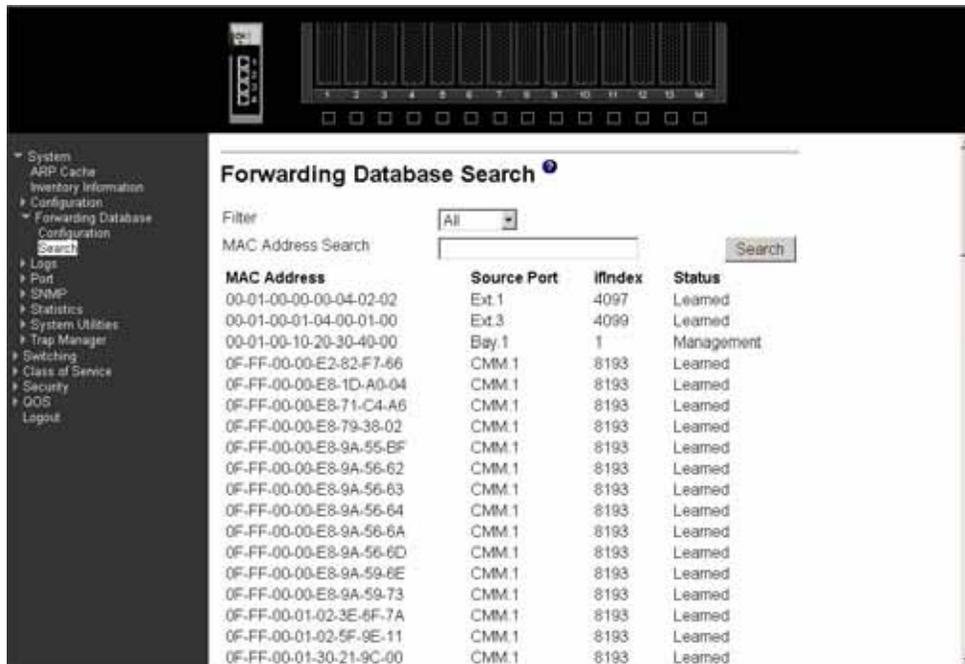
Aging Interval (secs)

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 given 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. Once 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 may 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, CMM refers to Chassis 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

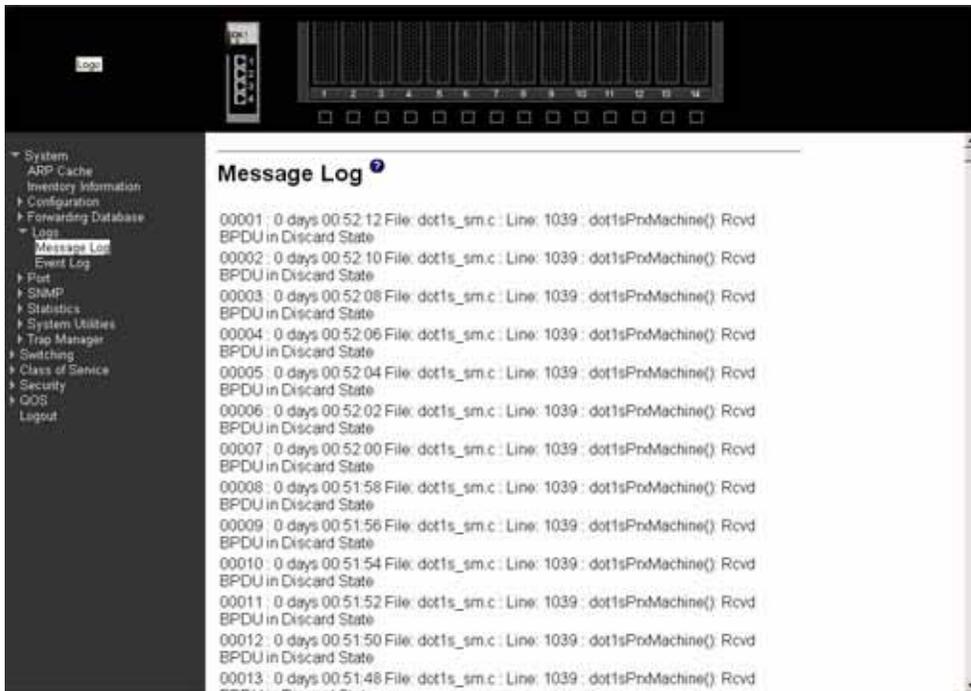
This menu provides access to the following two logs:

- 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 system 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.

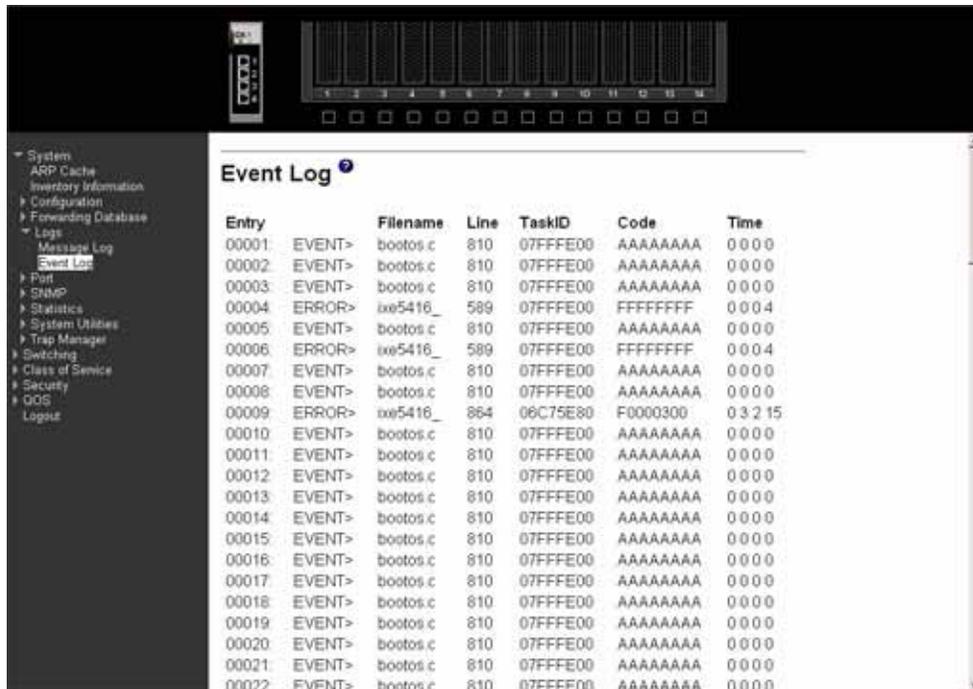
Line The line number within the source file of the code that detected the event.

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 (the actual number depends on the platform and OS), and is erased when an attempt is made to add an entry after it is full. The event log is preserved across system resets.



- 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.

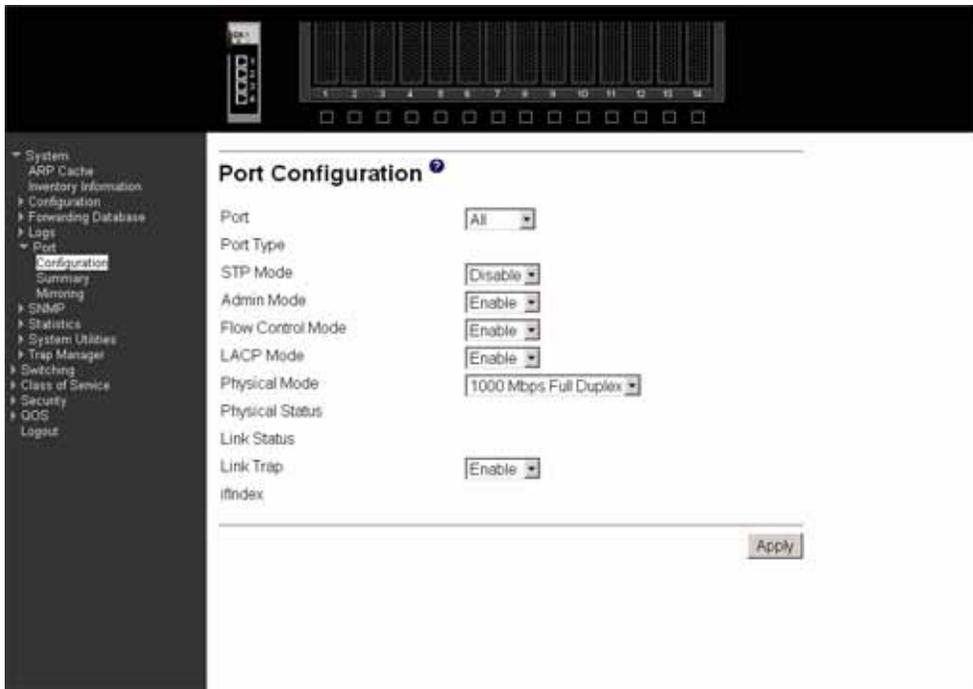
Port

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

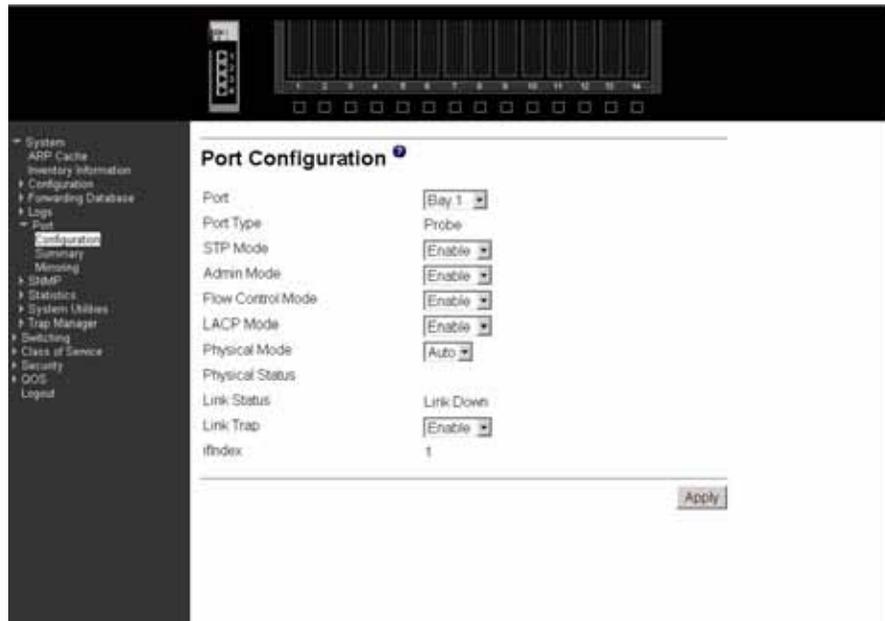
Selects the interface for which data is to be displayed or configured.

Port Type

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

Probe

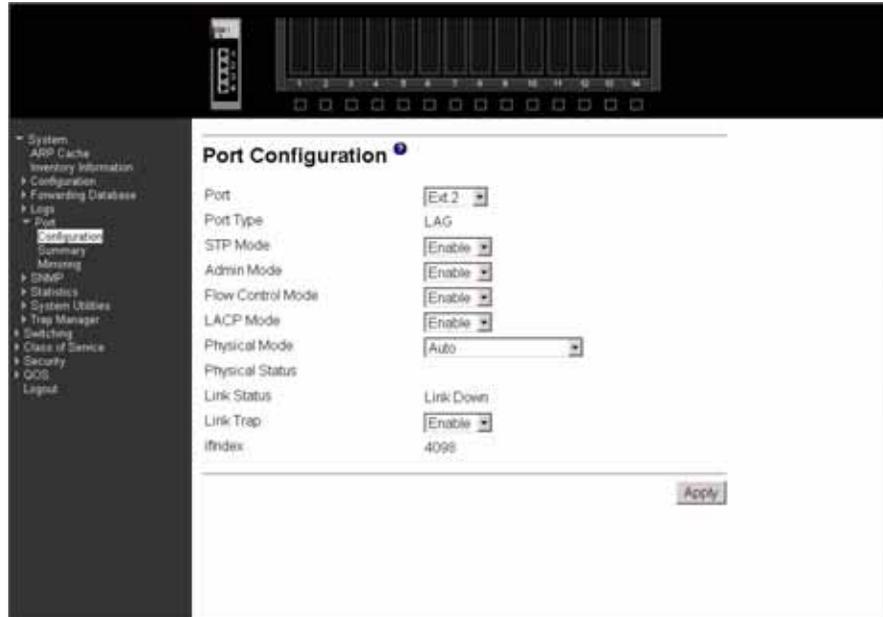
Monitoring port, participating in Port Mirroring. Following is how this panel displays when the port type is Probe.



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.

LACP 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 may be Enabled or Disabled by selecting the corresponding line on the pull-down entry field. The factory default is Enabled.

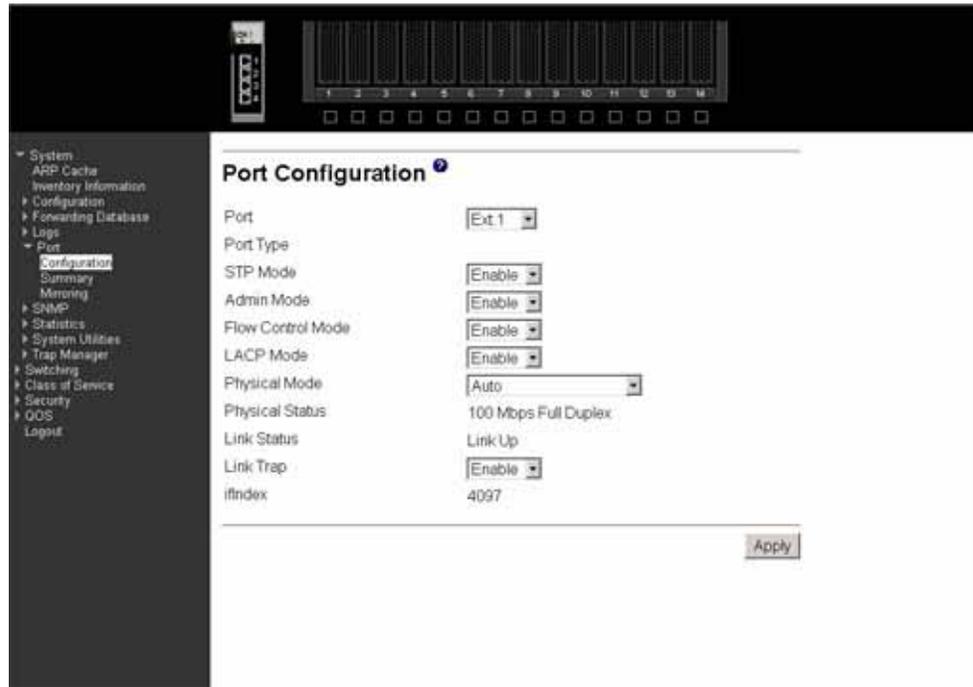
Physical Mode

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 only use this menu 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	Control Mode
Bay 1	Probe	Enabled	Disabled	Disabled Port	Enable	Enable
Bay 2	Mirrored	Enabled	Disabled	Disabled Port	Enable	Enable
Bay 3		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 4		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 5		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 6		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 7		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 8		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 9		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 10		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 11		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 12		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 13		Enabled	Disabled	Disabled Port	Enable	Enable
Bay 14		Enabled	Disabled	Disabled Port	Enable	Enable
Ext 1		Enabled	Forwarding	Designated Port	Enable	Enable
Ext 2	LAG	Enabled	Disabled	Disabled Port	Enable	Enable
Ext 3		Disabled	Manual forwarding	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.

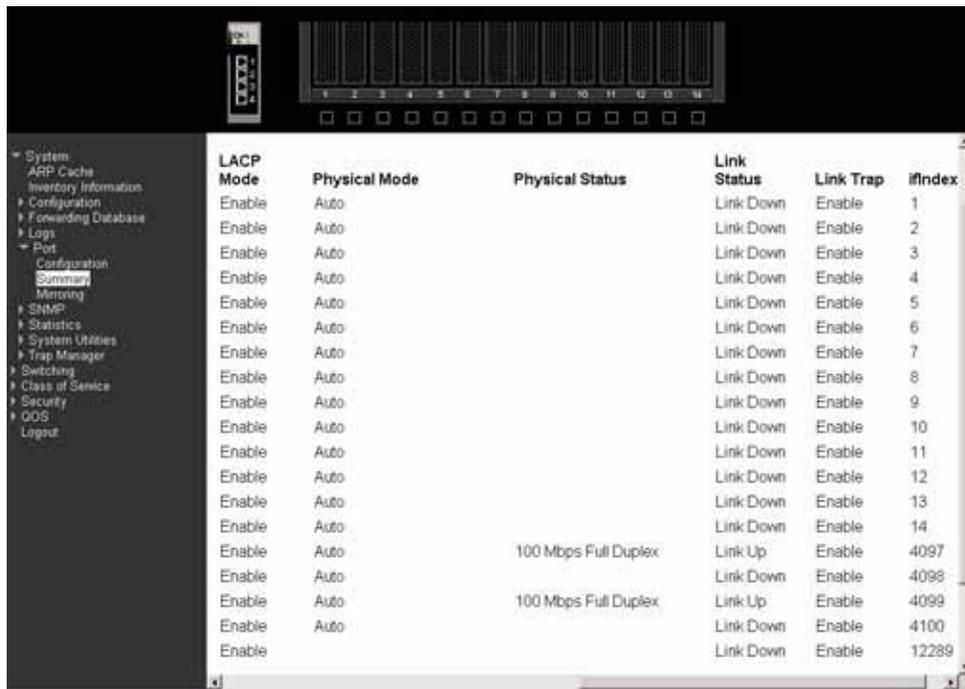
Admin Mode

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	9
Enable	Auto		Link Down	Enable	10
Enable	Auto		Link Down	Enable	11
Enable	Auto		Link Down	Enable	12
Enable	Auto		Link Down	Enable	13
Enable	Auto		Link Down	Enable	14
Enable	Auto	100 Mbps Full Duplex	Link Up	Enable	4097
Enable	Auto		Link Down	Enable	4098
Enable	Auto	100 Mbps Full Duplex	Link Up	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

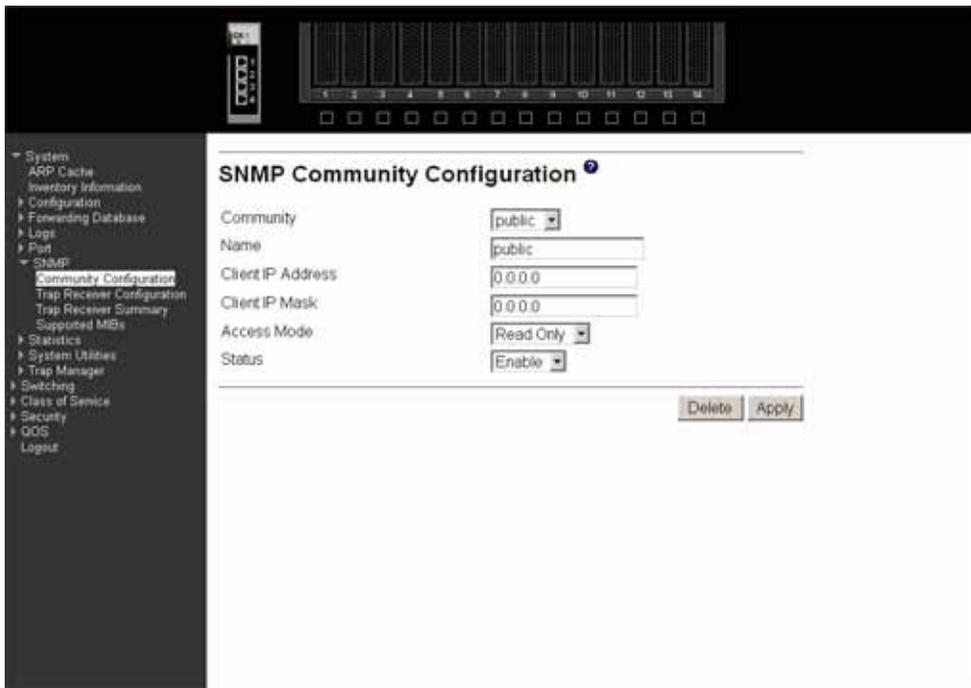
Community configuration

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.

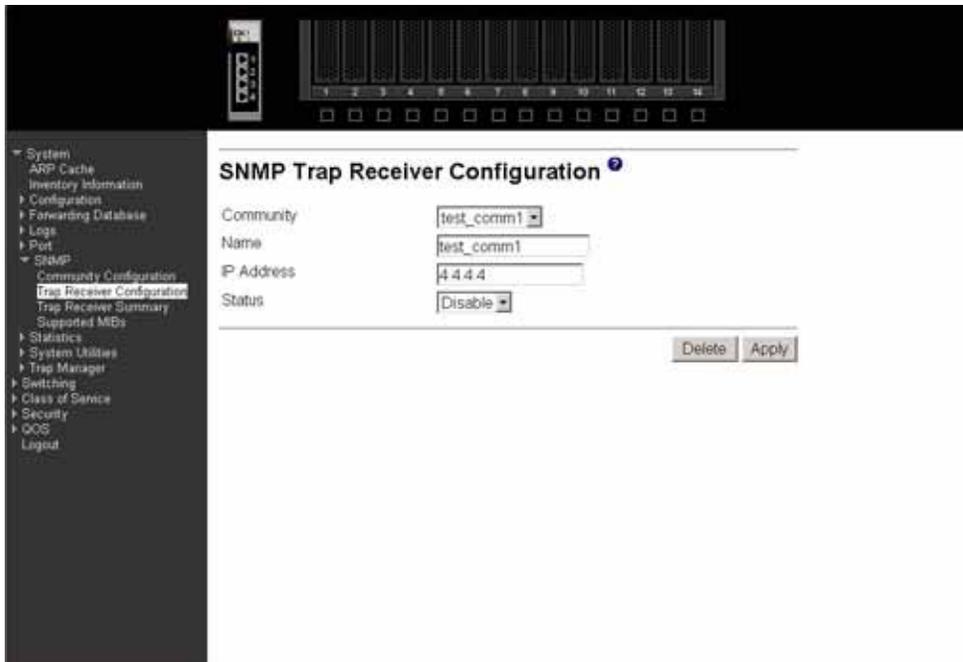
Click the Delete button to delete the currently selected Community Name. If you want the switch to retain the new values across a power cycle, you must perform a save.

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.



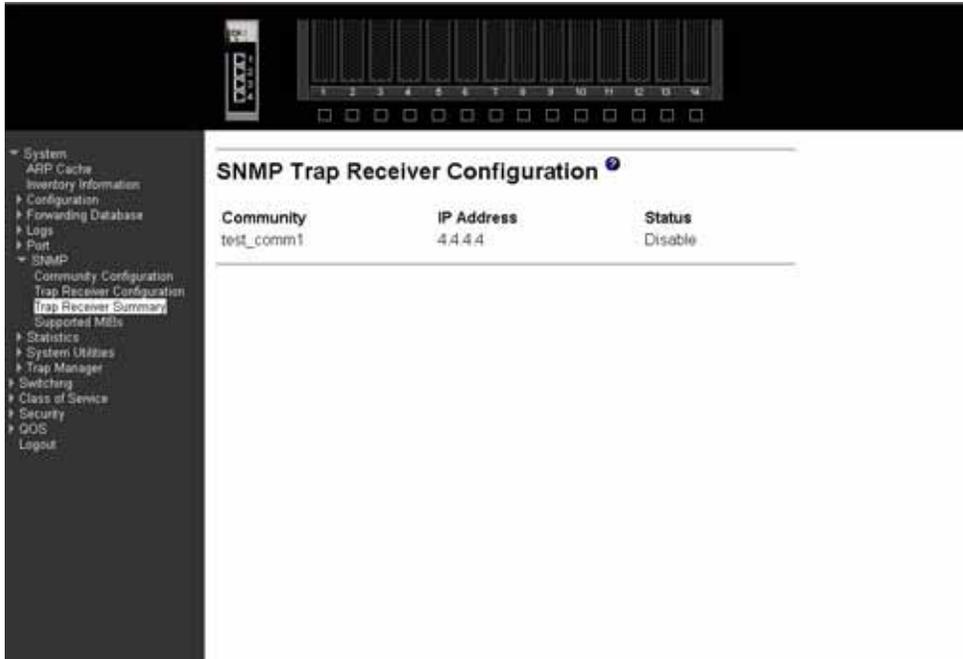
- Community** This field adds an SNMP trap receiver community name and associated IP address.
- Name** Enter the community string for the SNMP trap packet to be sent to the trap manager. This may be up to 16 characters and is case sensitive.
- IP Address** Enter the IP address to receive SNMP traps from 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.

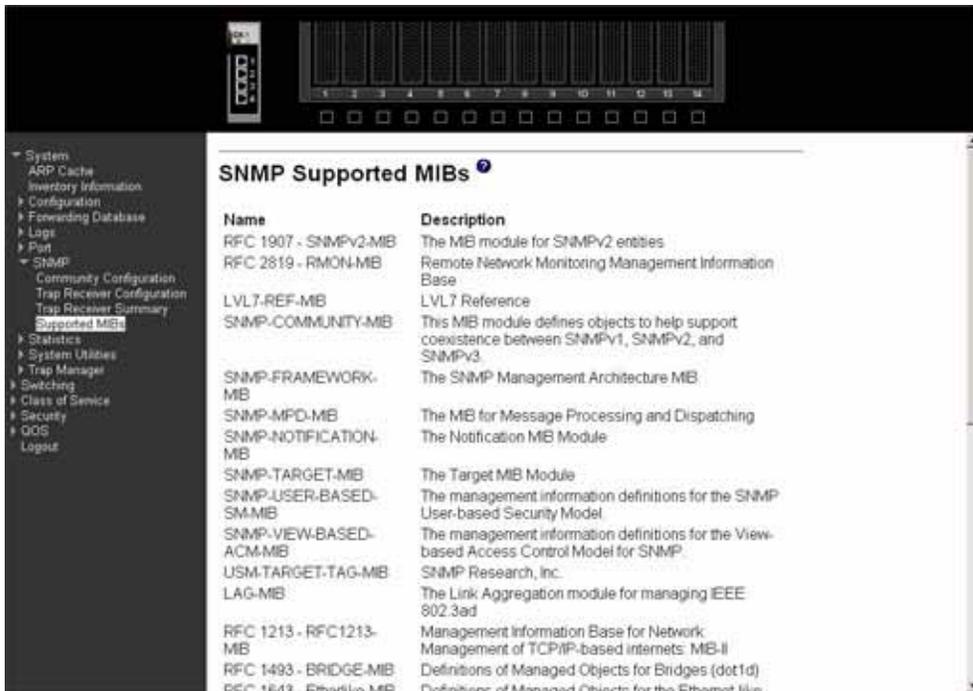
Status Indicates whether traps are currently Enabled for this community:

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 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.

The screenshot shows a network switch management interface. On the left is a navigation menu with categories like System, Configuration, Logs, Port, SNMP, and Statistics. The 'Statistics' section is expanded to show 'Switch Detailed Statistics'. The main area displays a table of statistics for a specific switch.

Statistic	Value
ifIndex	24577
Octets Received	2501120
Packets Received Without Error	17603
Unicast Packets Received	15613
Multicast Packets Received	1941
Broadcast Packets Received	51
Receive Packets Discarded	0
Octets Transmitted	2788335
Packets Transmitted Without Errors	15787
Unicast Packets Transmitted	15768
Multicast Packets Transmitted	0
Broadcast Packets Transmitted	19
Transmit Packets Discarded	0
Most Address Entries Ever Used	42
Address Entries in Use	25
Maximum VLAN Entries	3584
Most VLAN Entries Ever Used	1

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:

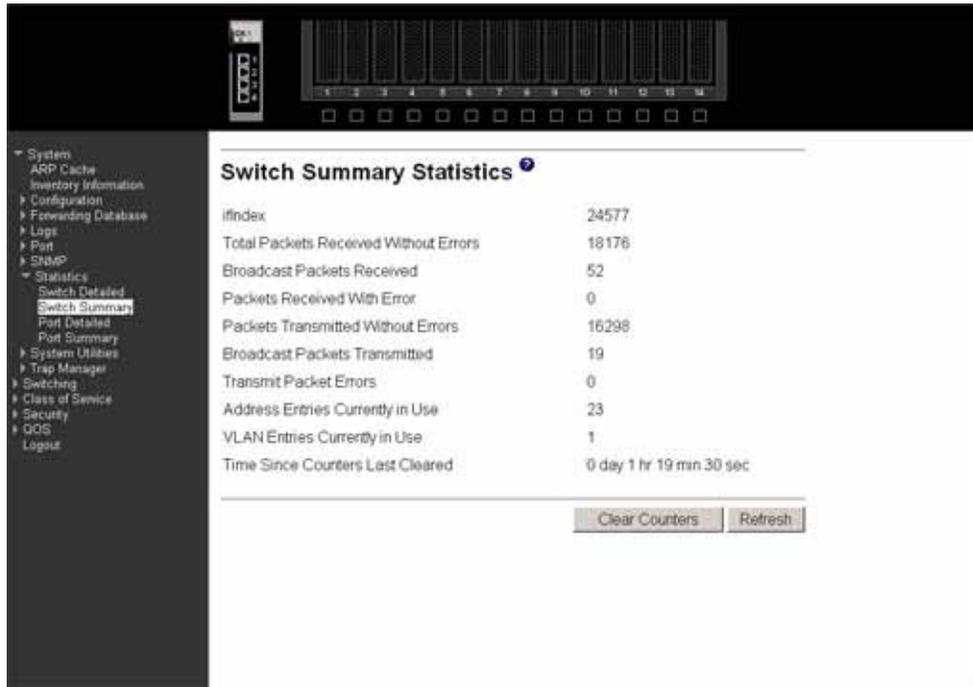
The elapsed time in days, hours, minutes and seconds since the statistics for this port were last cleared.

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.



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

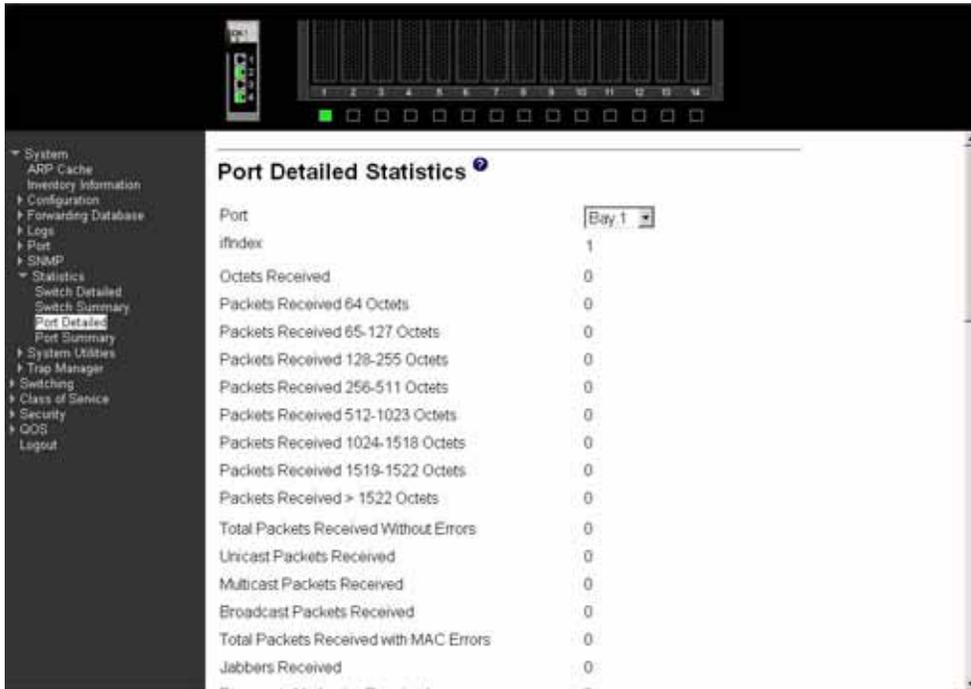
The elapsed time in days, hours, minutes and seconds since the statistics for this port were last cleared.

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 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.



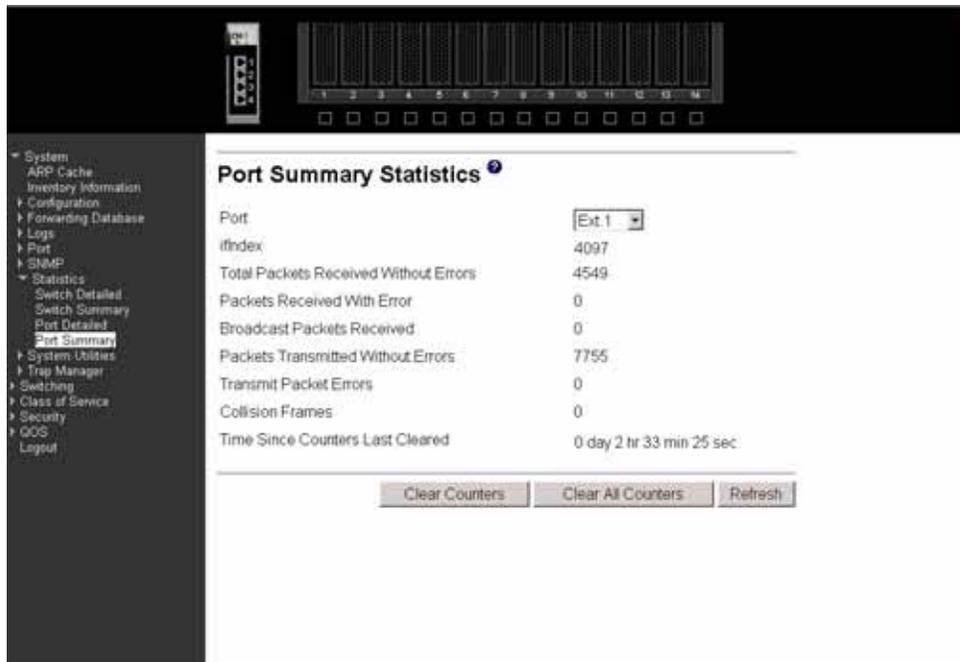
Click the Clear Counters button to clear all the counters, resetting all statistics for this port to default values.

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.

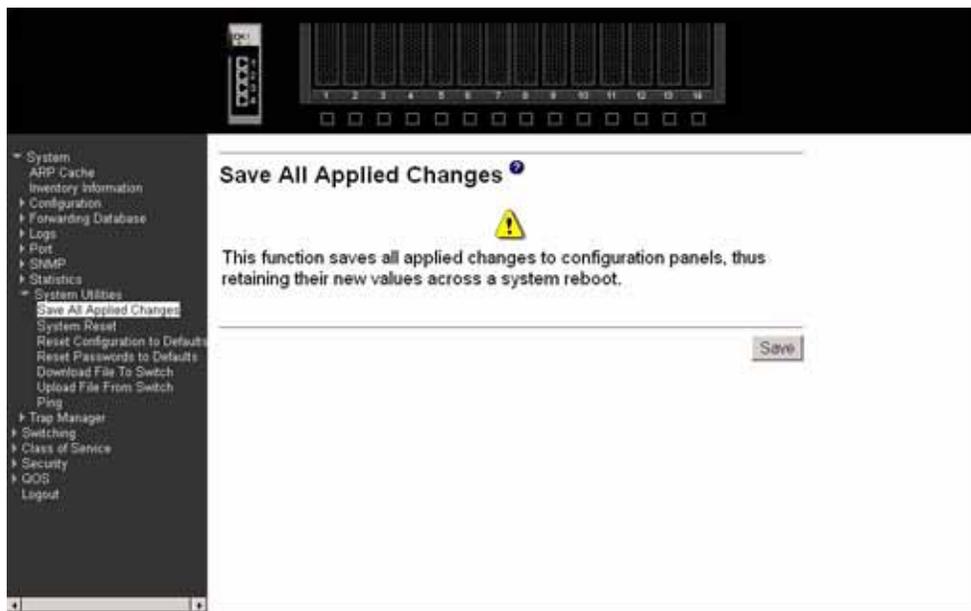
System utilities

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

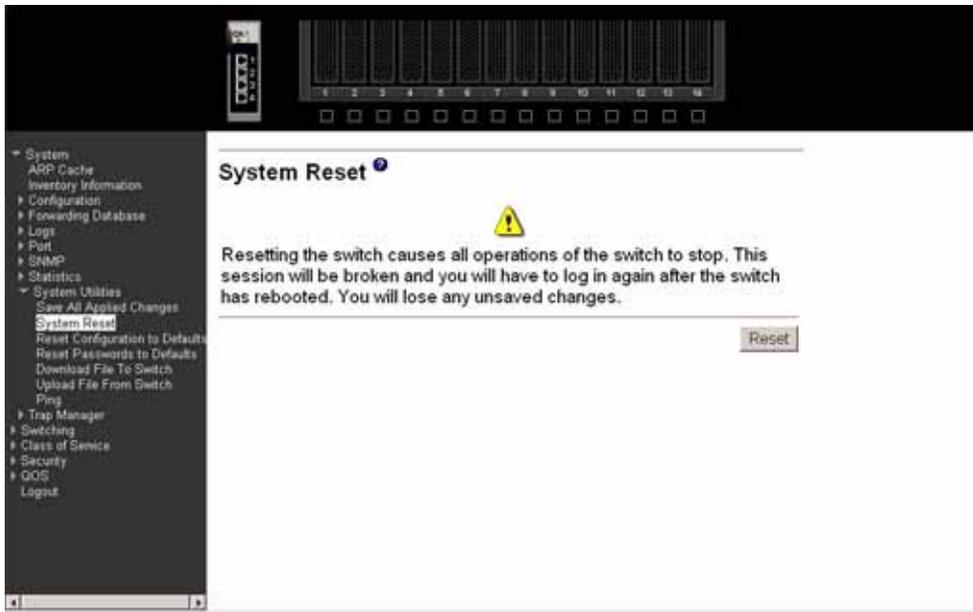
Save all applied changes

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.



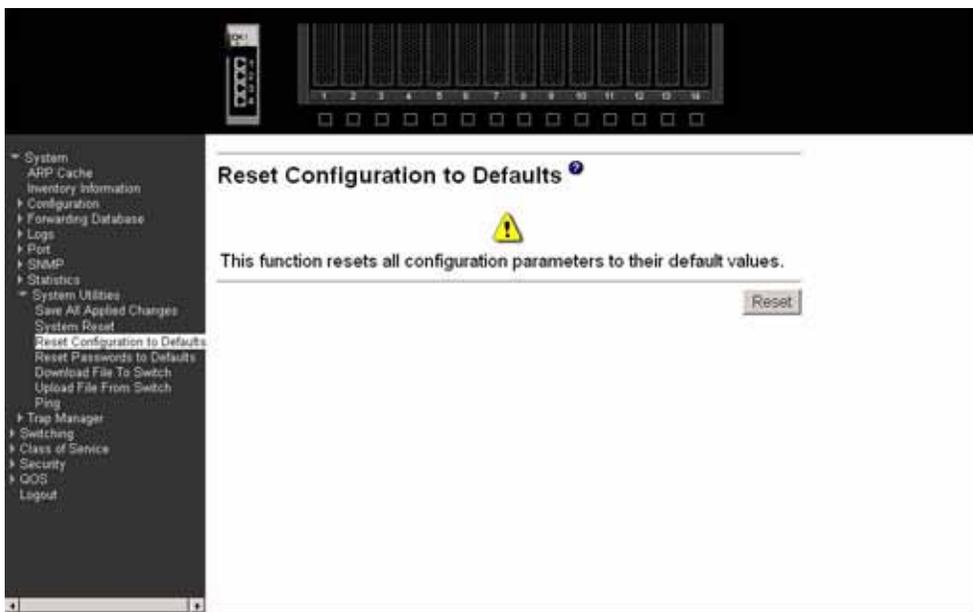
System reset

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.



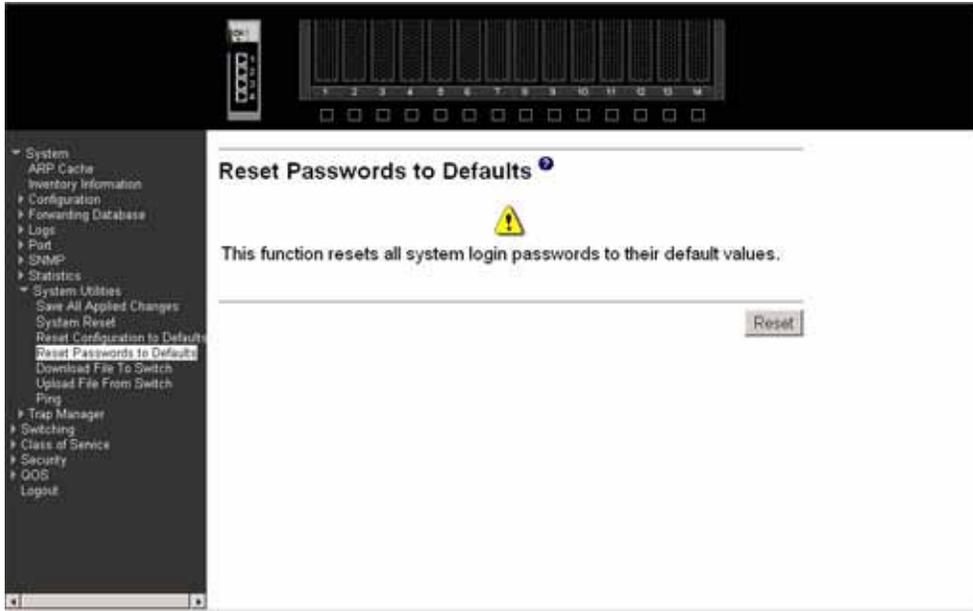
Reset configuration to defaults

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.



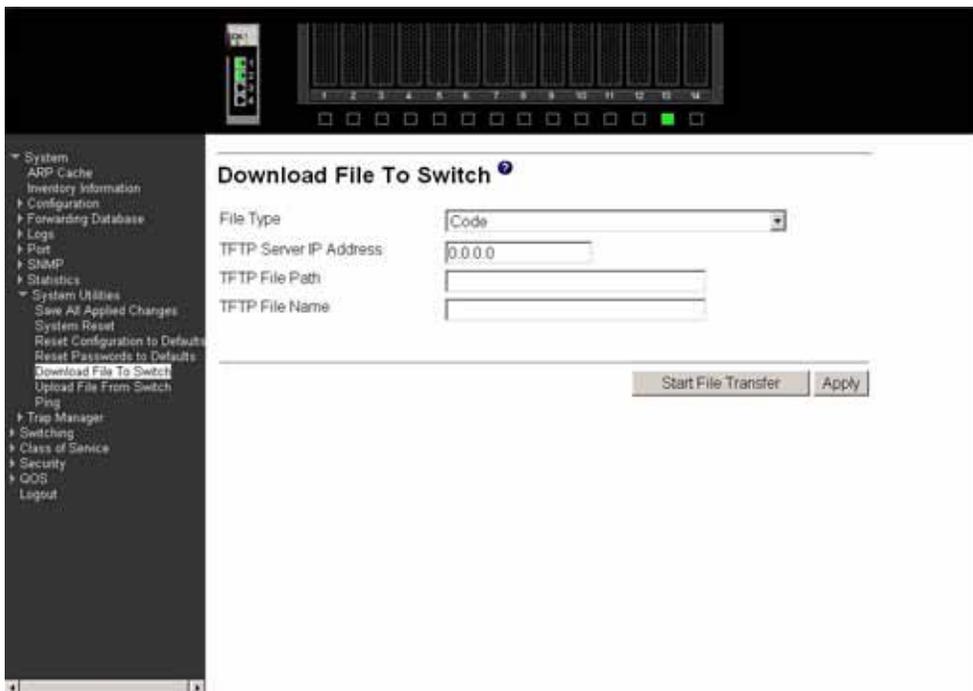
Reset passwords to defaults

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.



File Type

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

- Code
- 8051 MCU Code
- Config
- 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
- 8051 MCU Code

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 Intel® Blade Server Ethernet Switch Module IXM5414E 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 may 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 4. TFTP Upload Scenarios

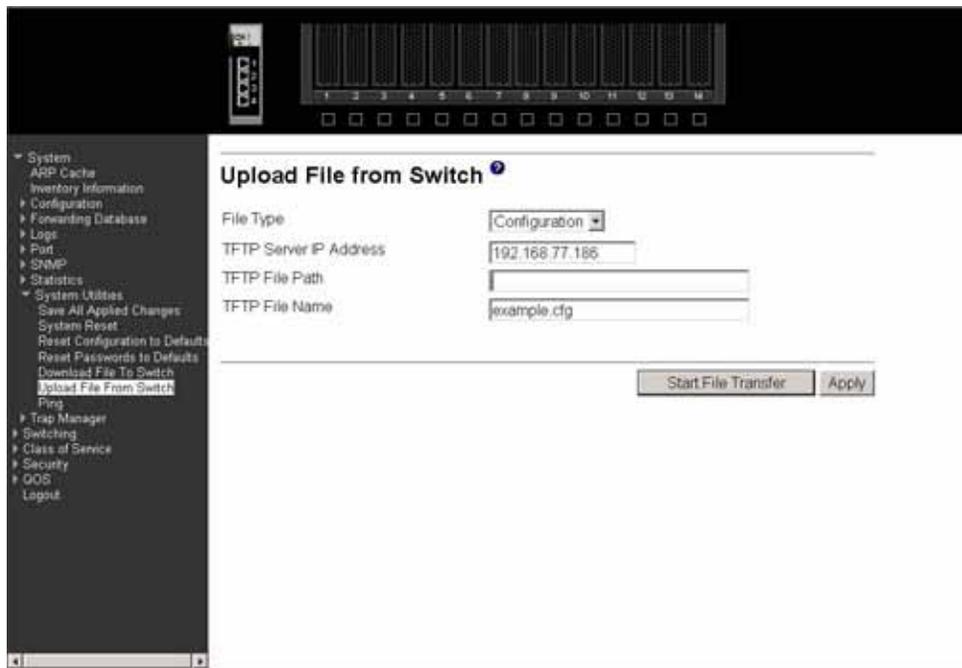
TFTP Server path	TFTP Client path
c:\tftp\	blank
c:\	tftp\
c:	\tftp\

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 86 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:

config	Configuration file
errorlog	Error log
msglog	Message log

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. Once 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.

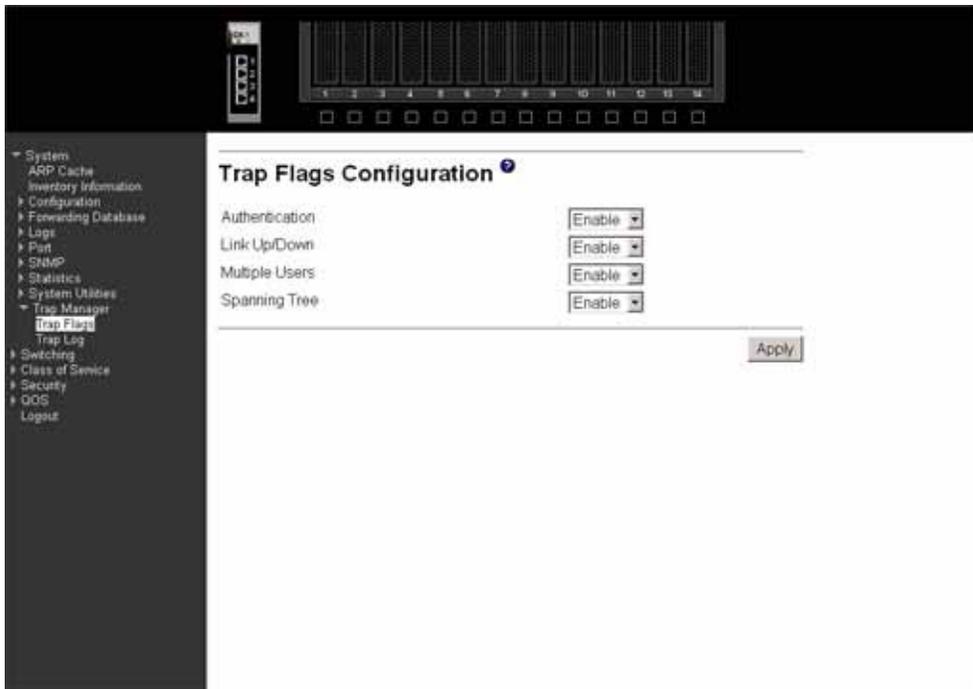
Trap manager

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.

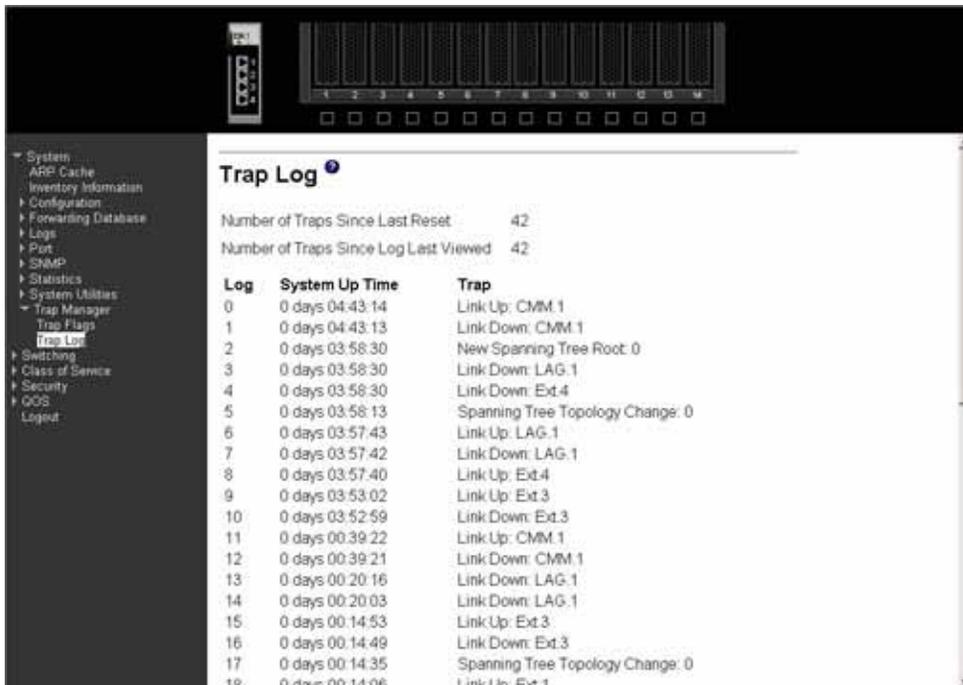
Spanning Tree

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
- Filters
- GARP
- IGMP snooping
- Link aggregation
- Multicast forwarding database

- Spanning tree

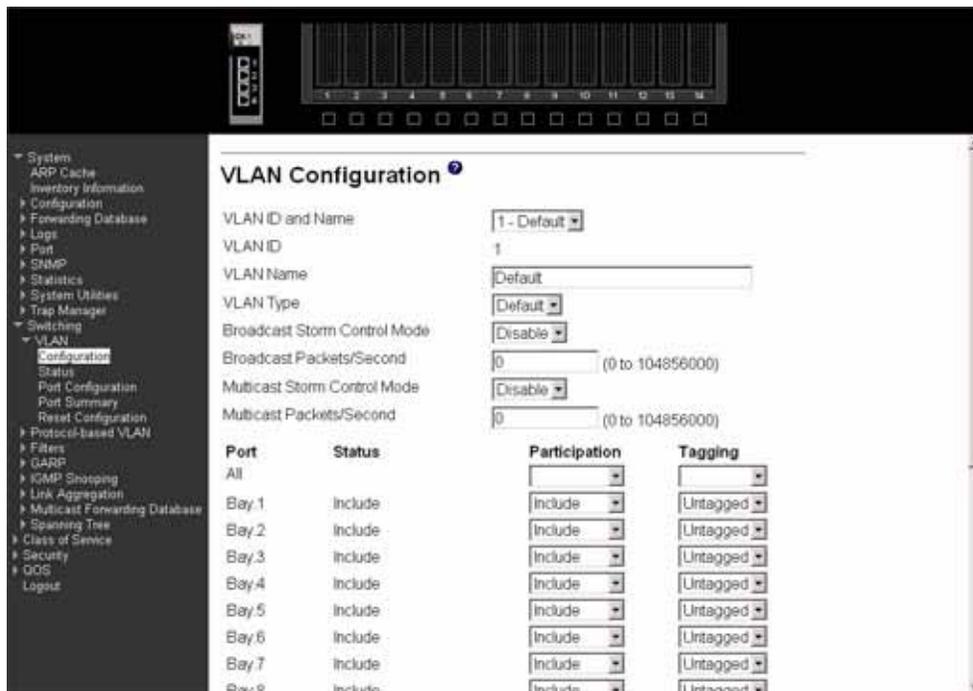
VLAN

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.

In order 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	

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.

In order 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

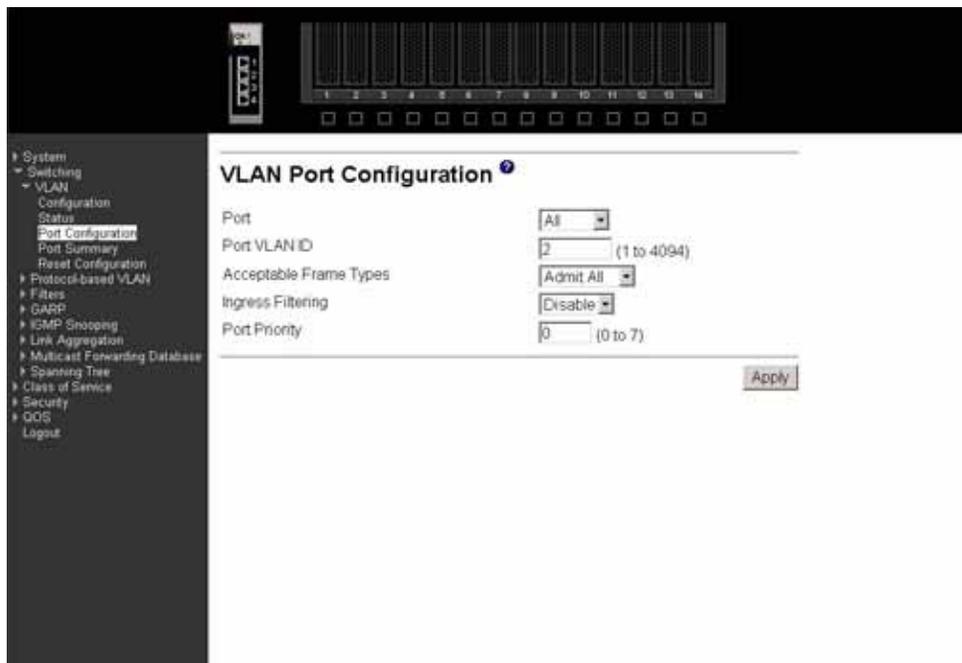
This field shows the mode of multicast storm control on the VLAN. If storm control is Enabled, storms are controlled by counting the number of multicast packets within a certain time period. If a count limit is exceeded, the packets are discarded.

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.

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	2	Admit All	Disabled	0
Bay.2	2	VLAN Only	Disabled	0
Bay.3	2	VLAN Only	Enabled	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
Bay.9	1	Admit All	Disabled	0
Bay.10	1	Admit All	Disabled	0
Bay.11	1	Admit All	Disabled	0
Bay.12	1	Admit All	Disabled	0
Bay.13	1	Admit All	Disabled	0
Bay.14	1	Admit All	Disabled	0
Ext.1	1	Admit All	Disabled	0
Ext.2	1	Admit All	Disabled	0
Ext.3	1	Admit All	Disabled	0
Ext.4	1	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 may 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.

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 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.

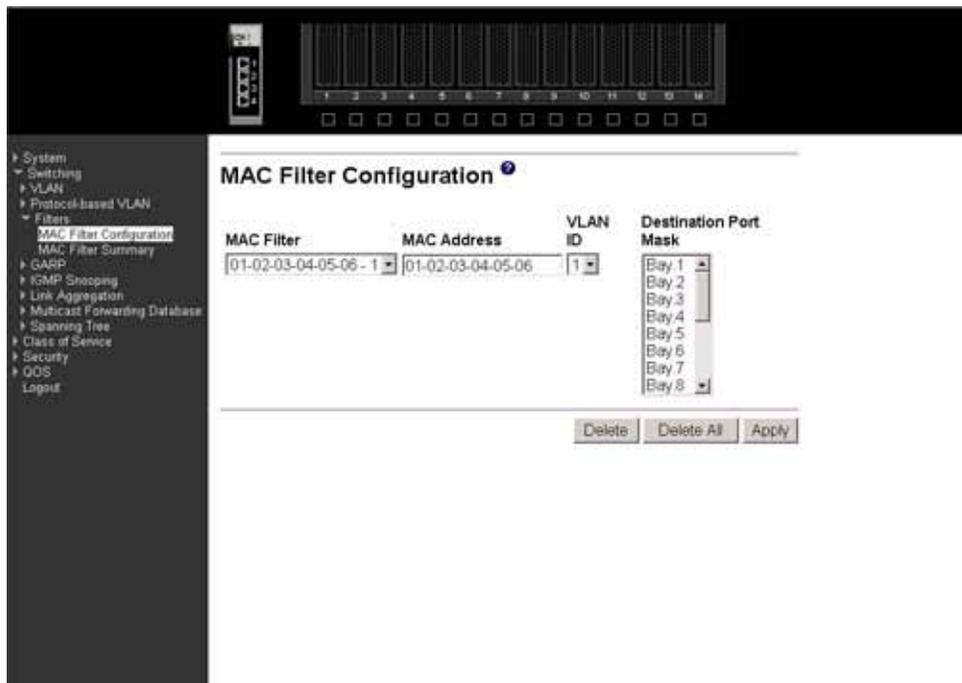
Filters

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 may be created.

MAC Address

The MAC address of the filter in the format 00-01-1A-B2-53-4D. You can only change this field 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 only change this field when you have selected the Create Filter option and you can only select 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.

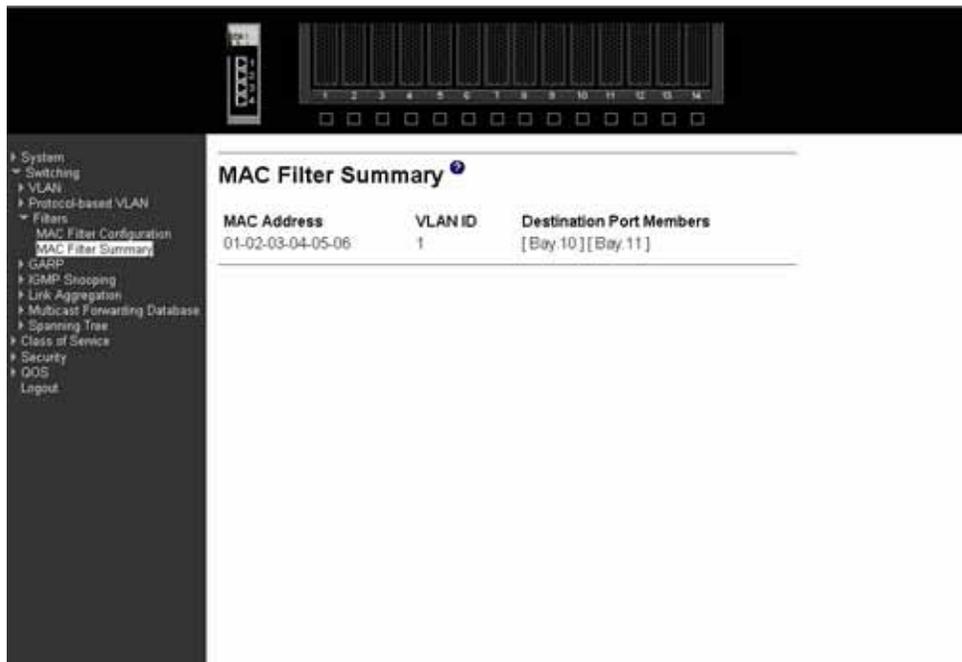
Click the Delete button to remove the currently selected filter.

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 may 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.

Port	Port GVRP Mode	Port GMRP Mode	Join Timer (centiseecs)	Leave Timer (centiseecs)	Leave All Timer (centiseecs)
Bay 1	Enabled	Enabled	20	60	1000
Bay 2	Enabled	Enabled	20	60	1000
Bay 3	Enabled	Enabled	20	60	1000
Bay 4	Enabled	Enabled	20	60	1000
Bay 5	Enabled	Enabled	20	60	1000
Bay 6	Enabled	Enabled	20	60	1000
Bay 7	Enabled	Enabled	20	60	1000
Bay 8	Enabled	Enabled	20	60	1000
Bay 9	Enabled	Enabled	20	60	1000
Bay 10	Enabled	Enabled	20	60	1000
Bay 11	Enabled	Enabled	20	60	1000
Bay 12	Enabled	Enabled	20	60	1000
Bay 13	Enabled	Enabled	20	60	1000
Bay 14	Enabled	Enabled	20	60	1000
Ext 1	Enabled	Enabled	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 (centiseocs)

Specifies the time between the transmission of GARP PDUs registering (or re-registering) membership for a VLAN or multicast group in centiseocs. An instance of this timer exists for each GARP participant for each port. Permissible values are 10 to 100 centiseocs (0.1 to 1.0 seconds). The factory default is 20 centiseocs (0.2 seconds).

Leave Timer (centiseocs)

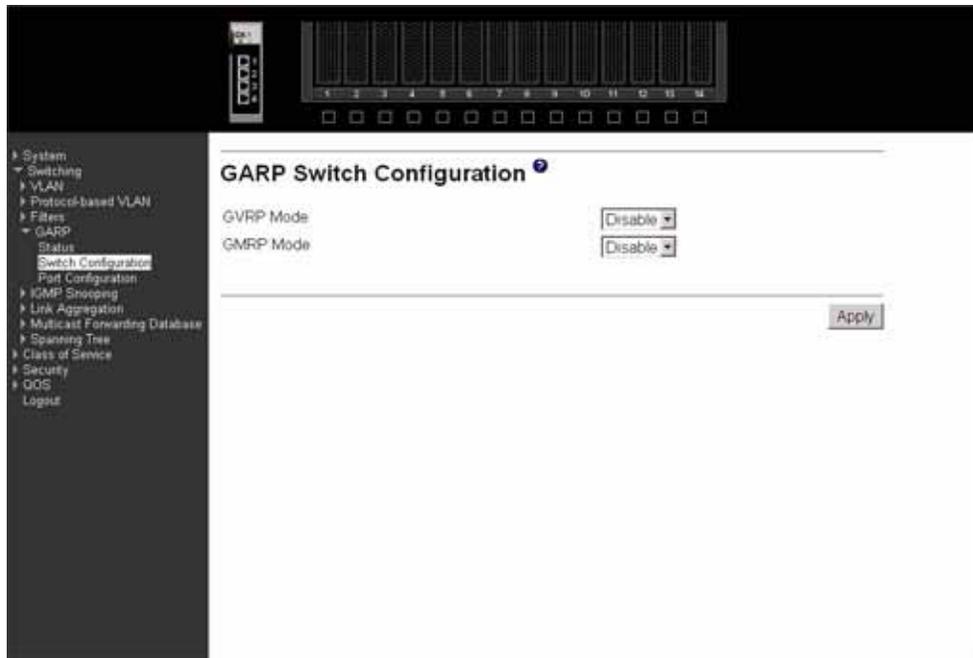
Specifies the time to wait after receiving an unregister request for a VLAN or multicast group before deleting the associated entry, in centiseocs. This allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service. An instance of this timer exists for each GARP participant for each port. Permissible values are 20 to 600 centiseocs (0.2 to 6.0 seconds). The factory default is 60 centiseocs (0.6 seconds).

Leave All Timer (centiseocs)

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*LeaveAllTime. Permissible values are 200 to 6000 centiseocs (2 to 60 seconds). The factory default is 1000 centiseocs (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.

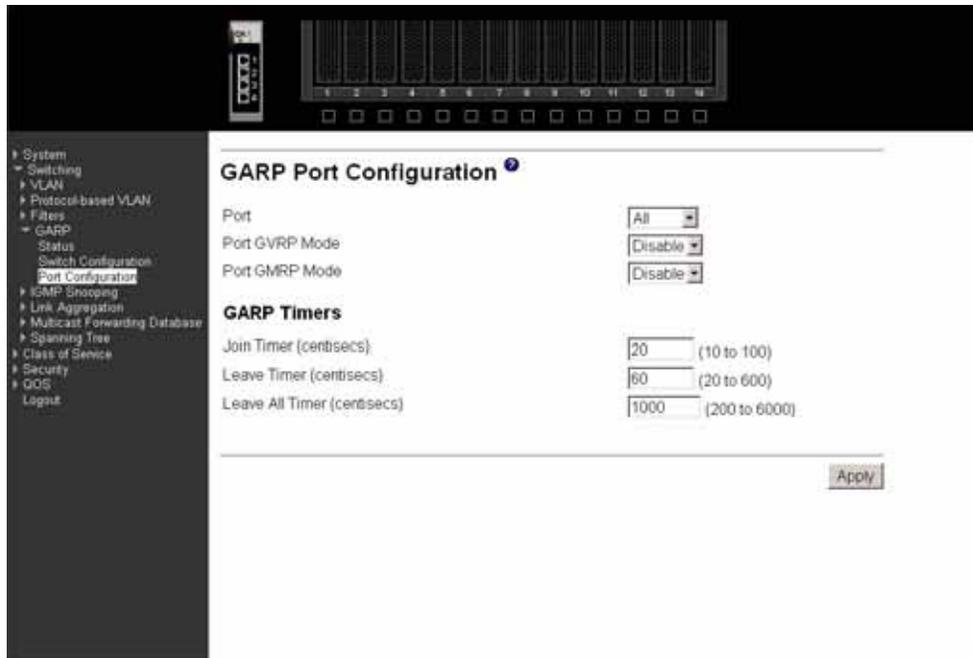
GMRP Mode

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.



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 centiseconds. Enter a number between 10 and 100 (0.1 to 1.0 seconds). The factory default is 20 centiseconds (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 centiseconds. 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 centiseconds (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 centiseconds. 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.

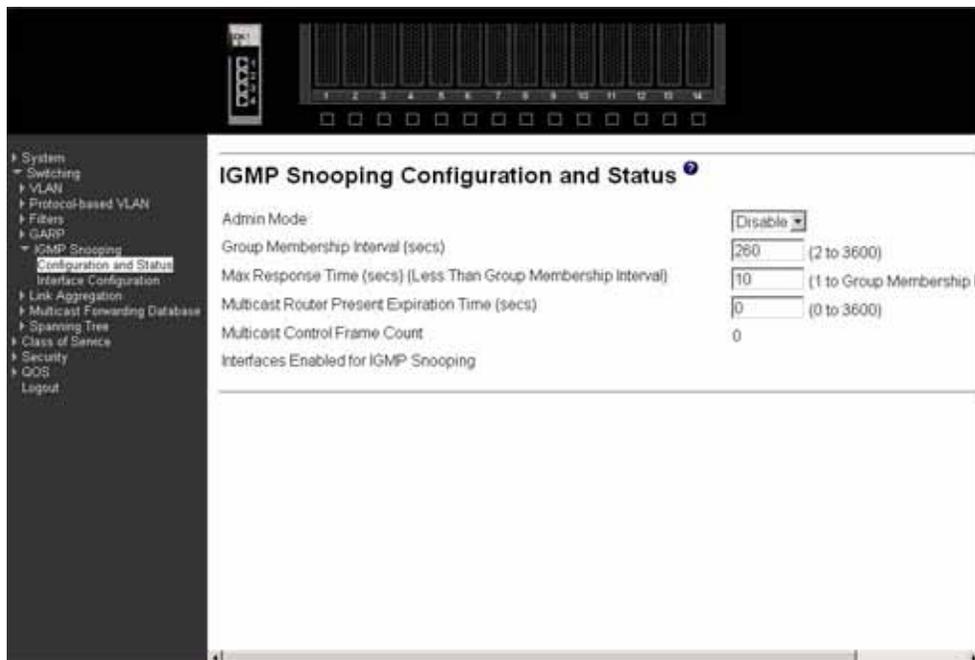
IGMP snooping

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.

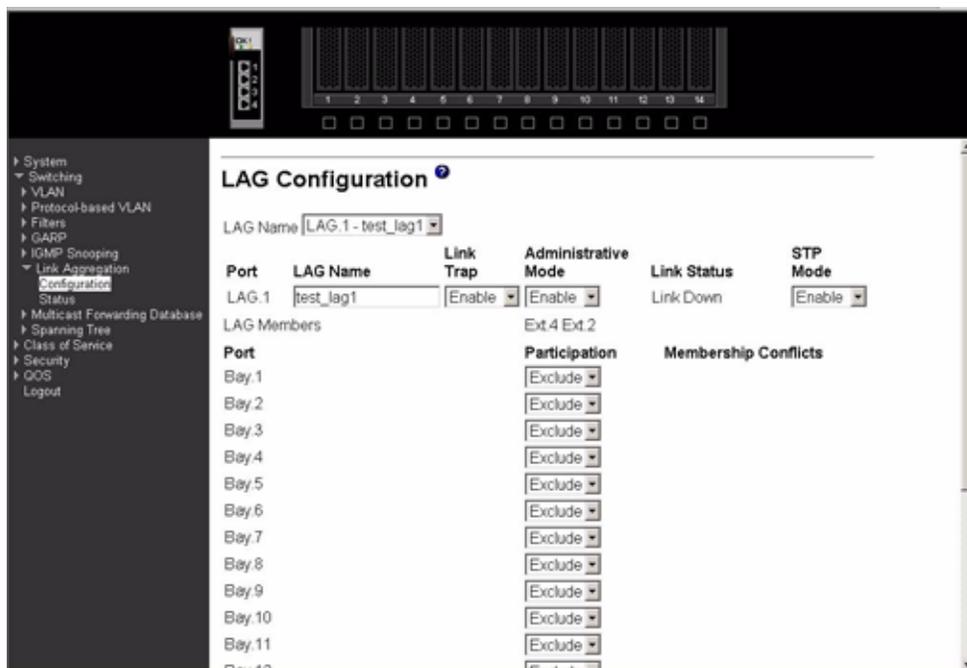
LAG

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 9 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 may only be a member of one LAG at a time. If the entry is blank, it is not currently a member of any LAG.

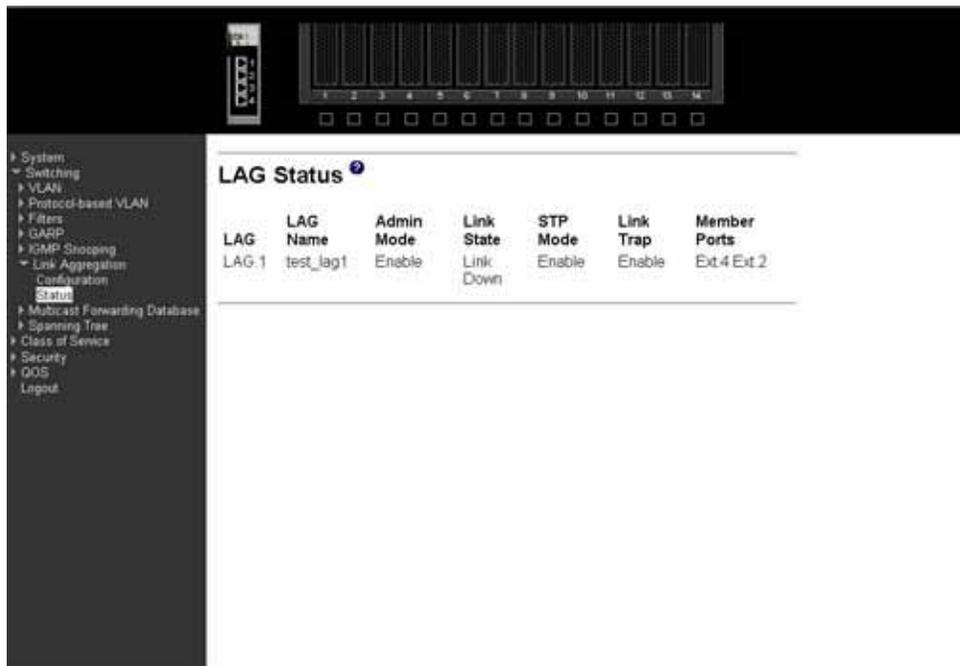
Click the Refresh button to refresh the data on the screen with the present state of the data in the switch.

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 given LAG.

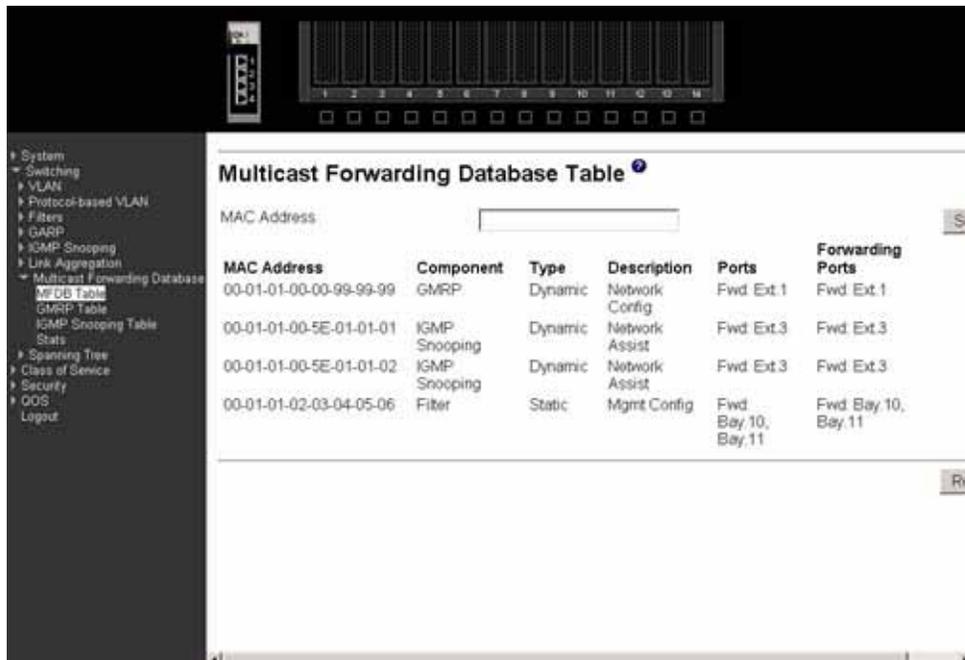
MFDB

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 may 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.

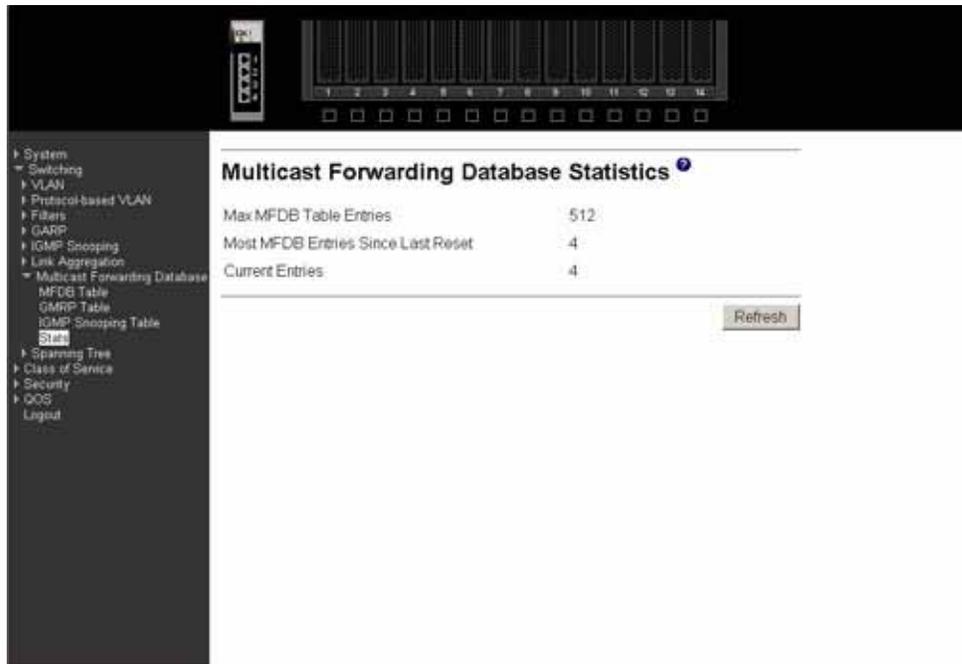
Ports The list of interfaces that are designated for forwarding (Fwd:) and filtering (Flt:).

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.



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

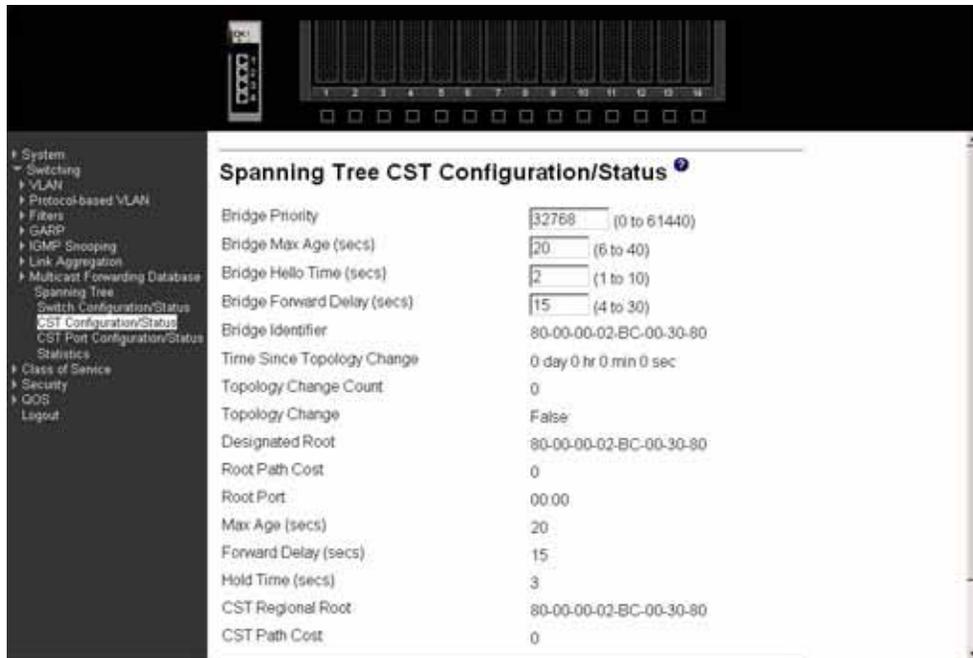
A derived value identifying the configuration.

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.



Bridge Priority

Specifies the bridge priority. The value may 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 $(2 * 4096 - 1)$ it will be set to 4096. The default priority is 32768.

Bridge Max Age (secs)

Specifies the bridge maximum age timeout value. The value may be between 1 and 40, and should be less than or equal to $((2 * \text{Bridge Forward Delay}) - 1)$ and greater than or equal to $(2 * (\text{Bridge Hello Time} + 1))$. 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

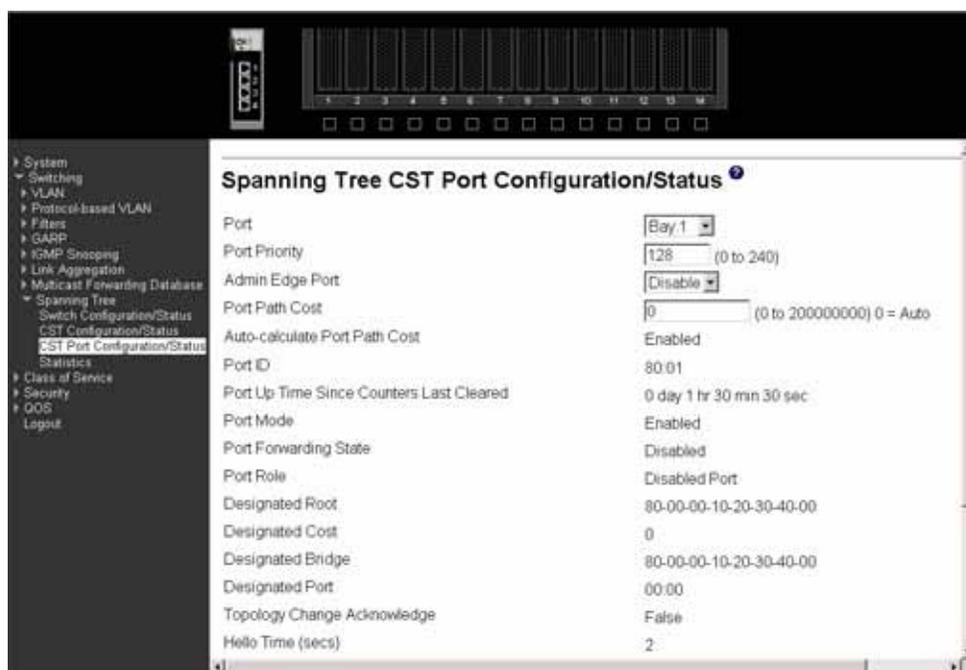
Path Cost to the CST tree Regional Root.

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.



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.

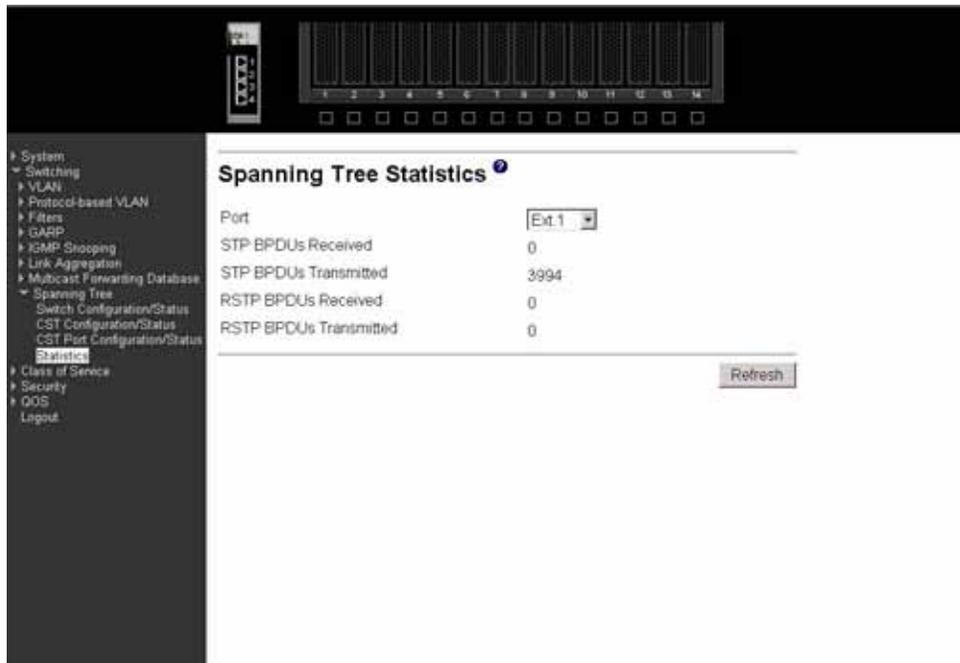
Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

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.

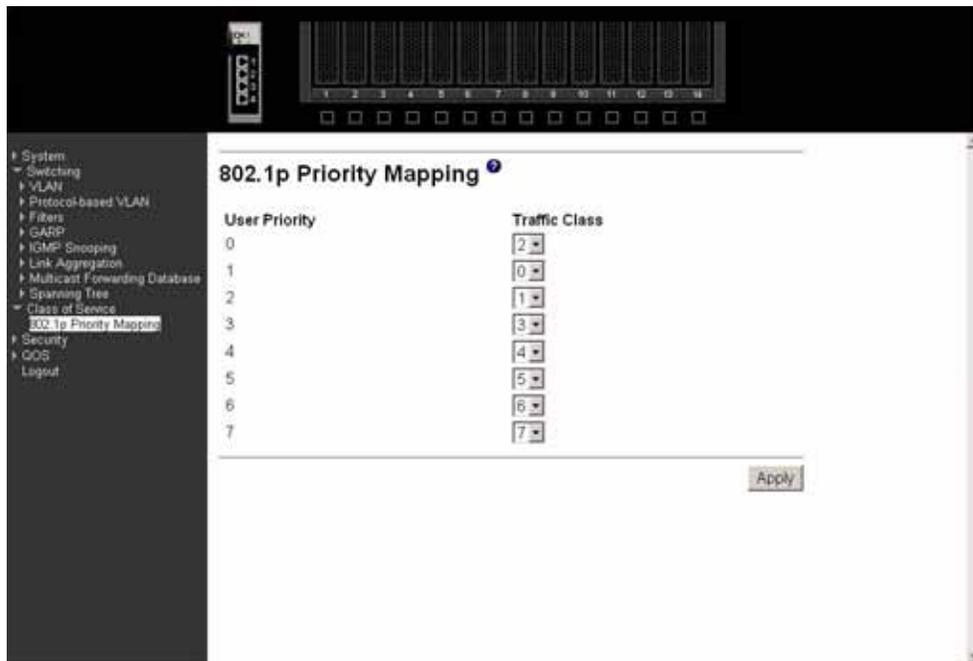
Click the Refresh button to update the screen with the most recent data.

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

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 Intel® Blade Server Ethernet Switch Module IXM5414E. 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

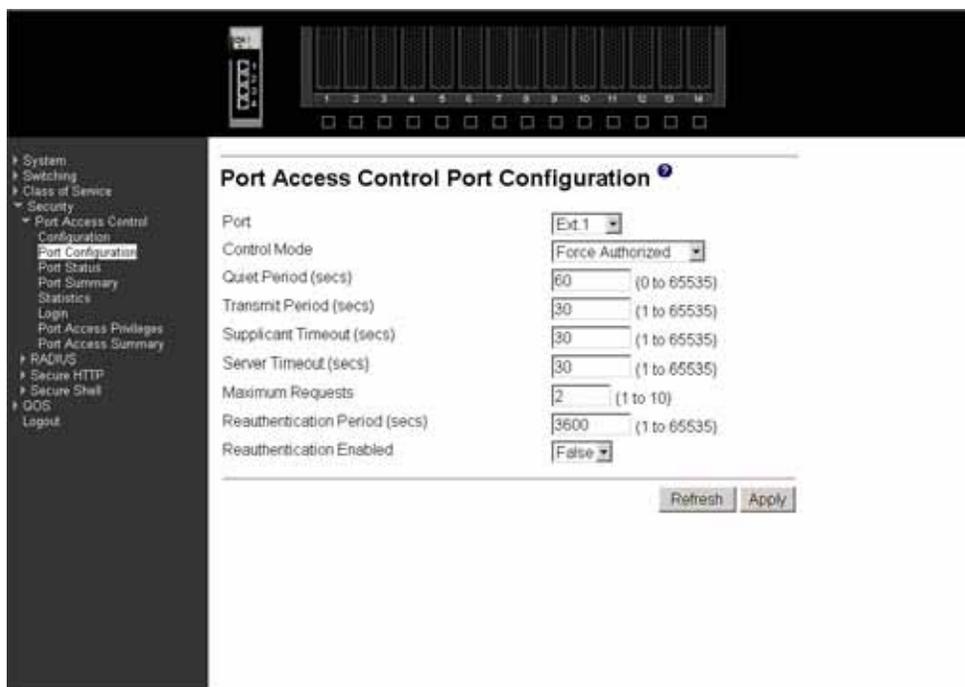
Lists the two options for administrative mode: Enable and Disable. The default value is Disable.

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. Once 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. Once 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	true	false	Unauthorized
Bay 2	forceunauthorized	forceunauthorized	false	false	Unauthorized
Bay 3	auto	auto	false	false	Unauthorized
Bay 4	auto	auto	false	false	Unauthorized
Bay 5	auto	auto	false	false	Unauthorized
Bay 6	auto	auto	false	false	Unauthorized
Bay 7	auto	auto	false	false	Unauthorized
Bay 8	auto	auto	false	false	Unauthorized
Bay 9	auto	auto	false	false	Unauthorized
Bay 10	auto	auto	false	false	Unauthorized
Bay 11	auto	auto	false	false	Unauthorized
Bay 12	auto	auto	false	false	Unauthorized
Bay 13	auto	auto	false	false	Unauthorized
Bay 14	auto	auto	false	false	Unauthorized
Ext 1	forceauthorized	forceauthorized	false	false	Authorized
Ext 2	auto	auto	false	false	Unauthorized
Ext 3	auto	auto	false	false	Unauthorized
Ext 4	auto	auto	false	false	Unauthorized

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 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.

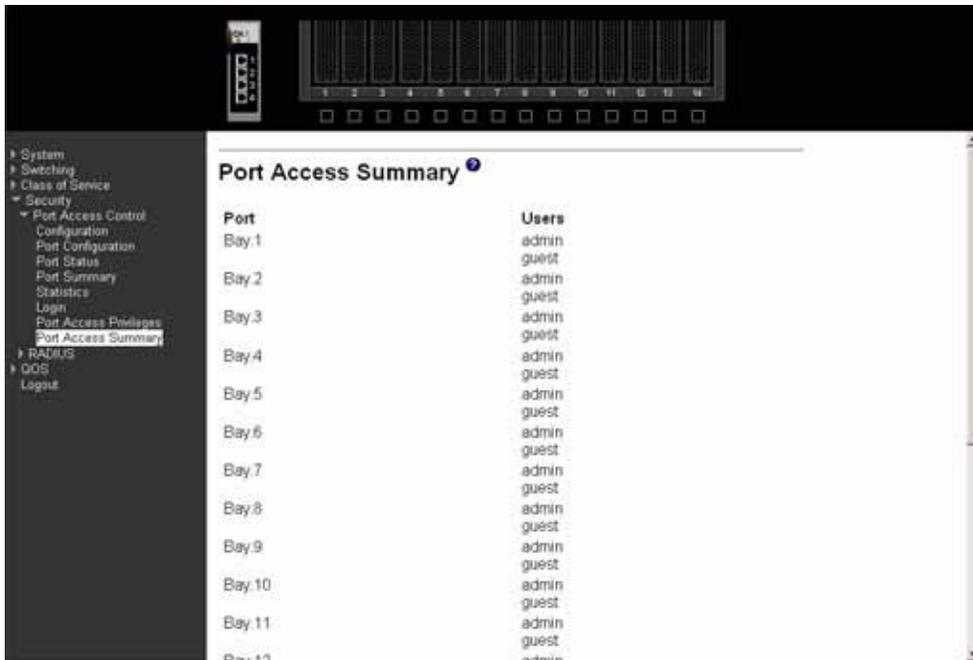
Users Select the users that may have access to the selected port or ports.

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:

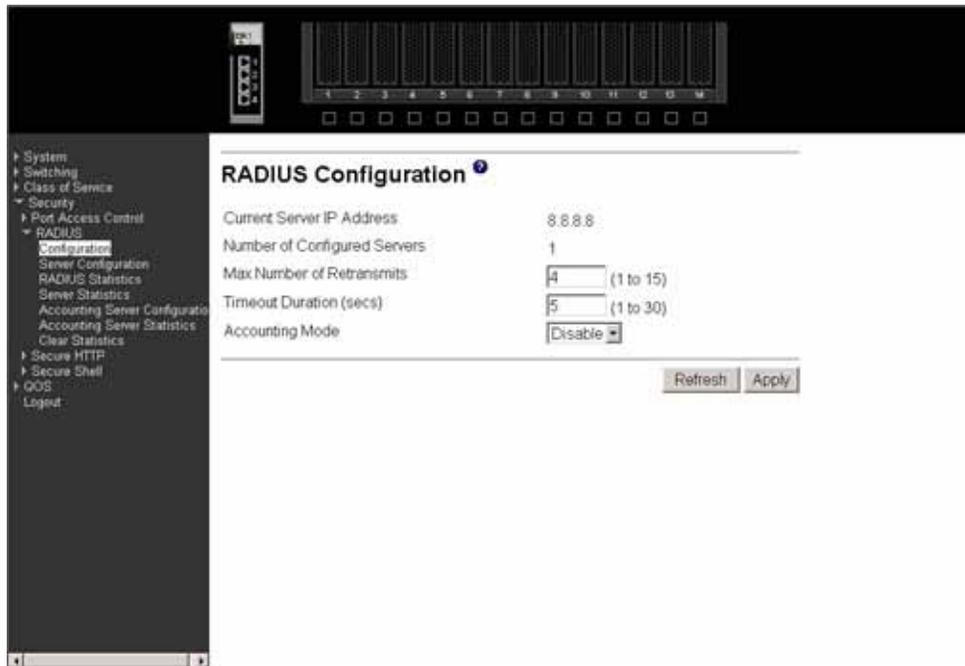
- Configuration
- Server configuration
- RADIUS statistics
- Server statistics
- Accounting server configuration
- Accounting server statistics
- Clear statistics

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

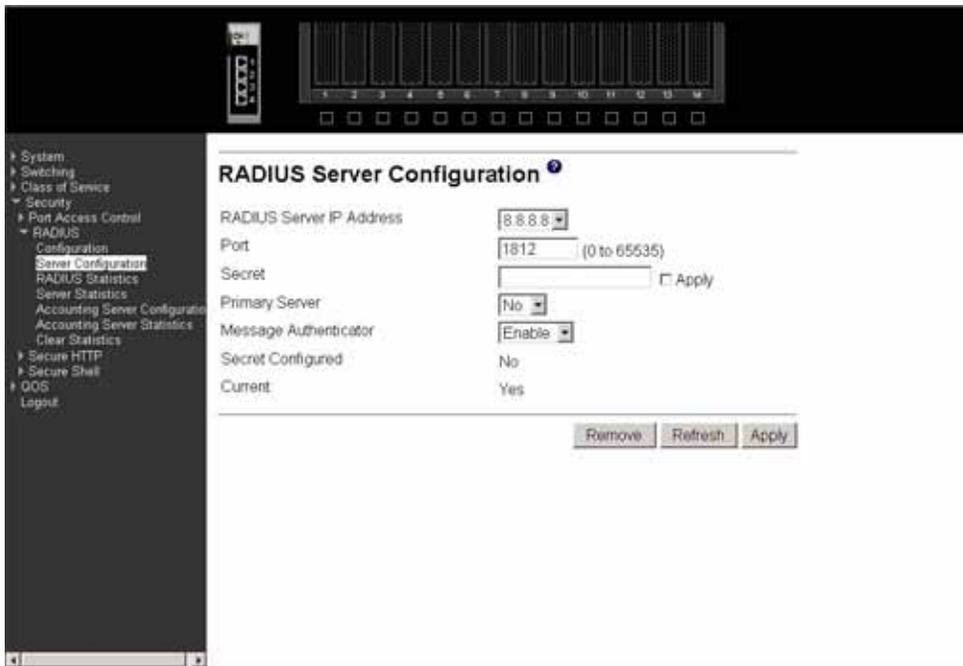
Select whether the RADIUS accounting mode is Enabled or Disabled.

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.



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.

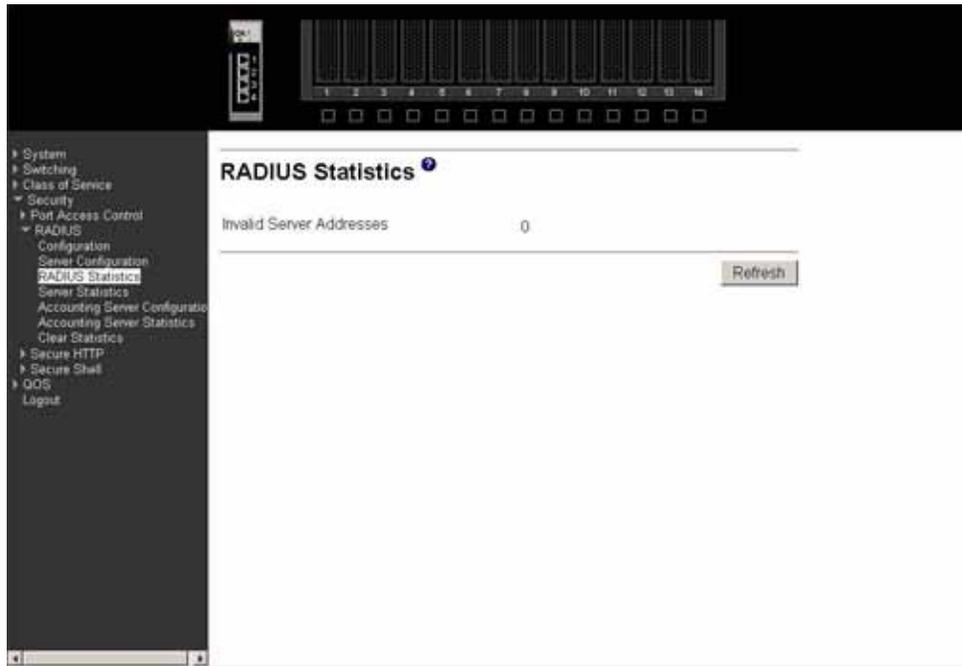
Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

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.



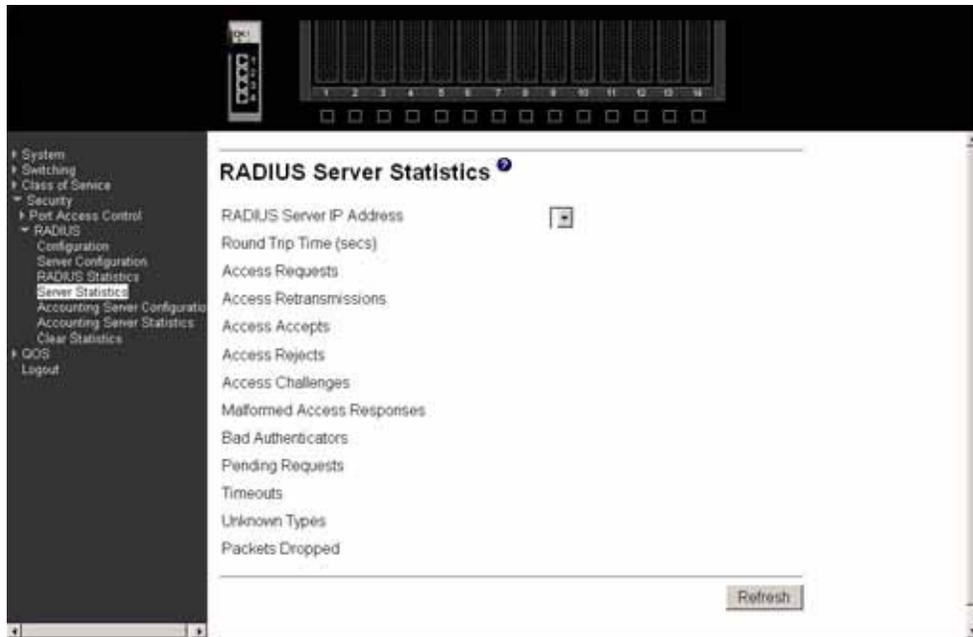
Invalid Server Addresses

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.

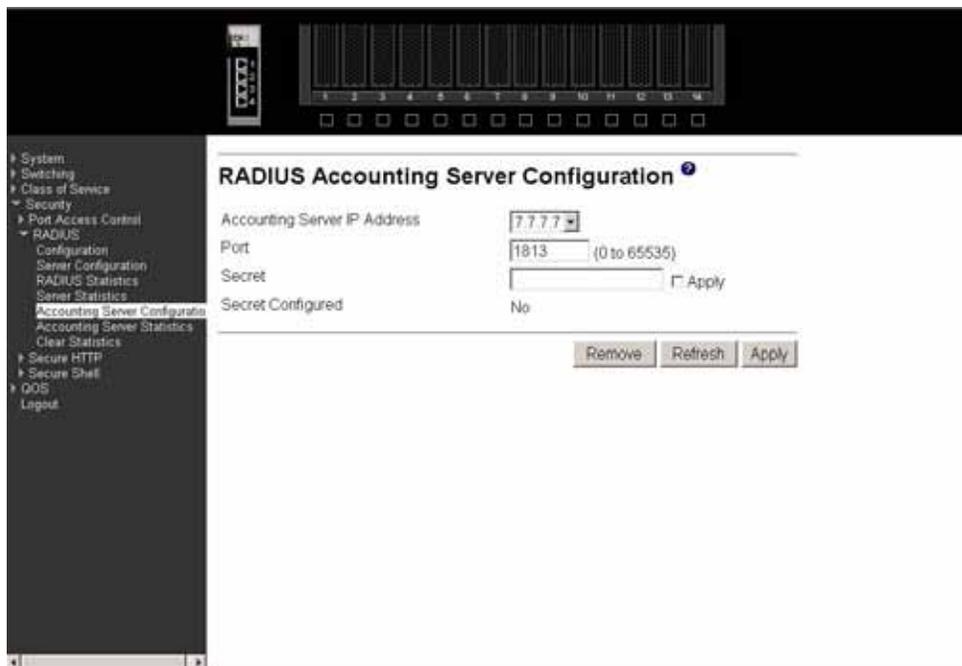
Packets Dropped

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.



The screenshot shows a web interface for configuring a RADIUS Accounting Server. The title is "RADIUS Accounting Server Configuration". On the left is a navigation menu with categories like System, Switching, Class of Service, Security, Port Access Control, RADIUS, and Accounting Server Configuration. The main content area contains the following fields:

- Accounting Server IP Address: A dropdown menu showing "7.7.7.7".
- Port: A text input field containing "1813" with a note "(0 to 65535)".
- Secret: A text input field that is currently empty.
- Secret Configured: A checkbox labeled "Apply" which is currently unchecked.

At the bottom of the configuration area are three buttons: "Remove", "Refresh", and "Apply".

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 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.

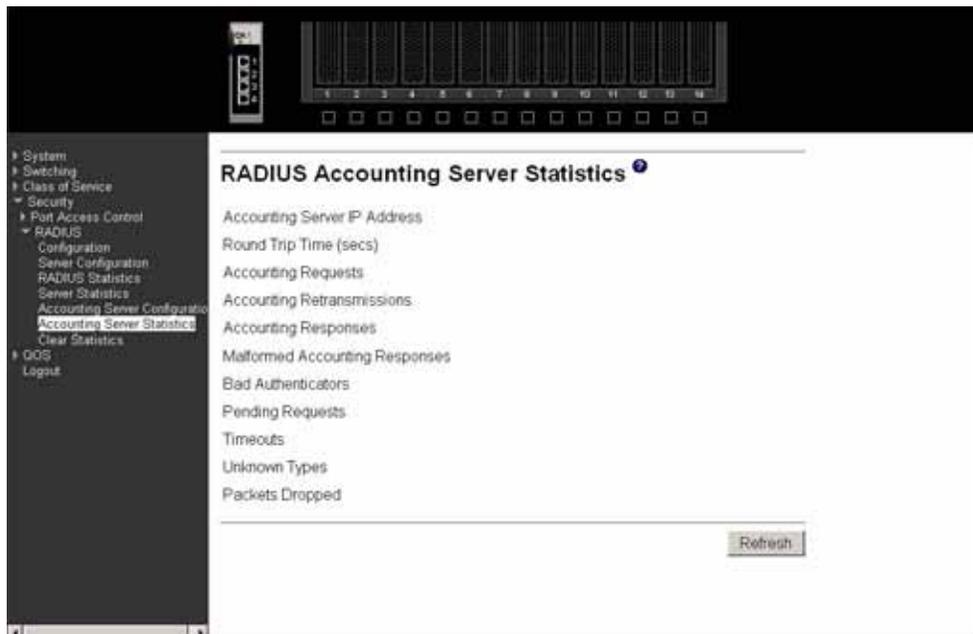
Click the Apply button to send the updated screen to the switch and cause the changes to take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

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.



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

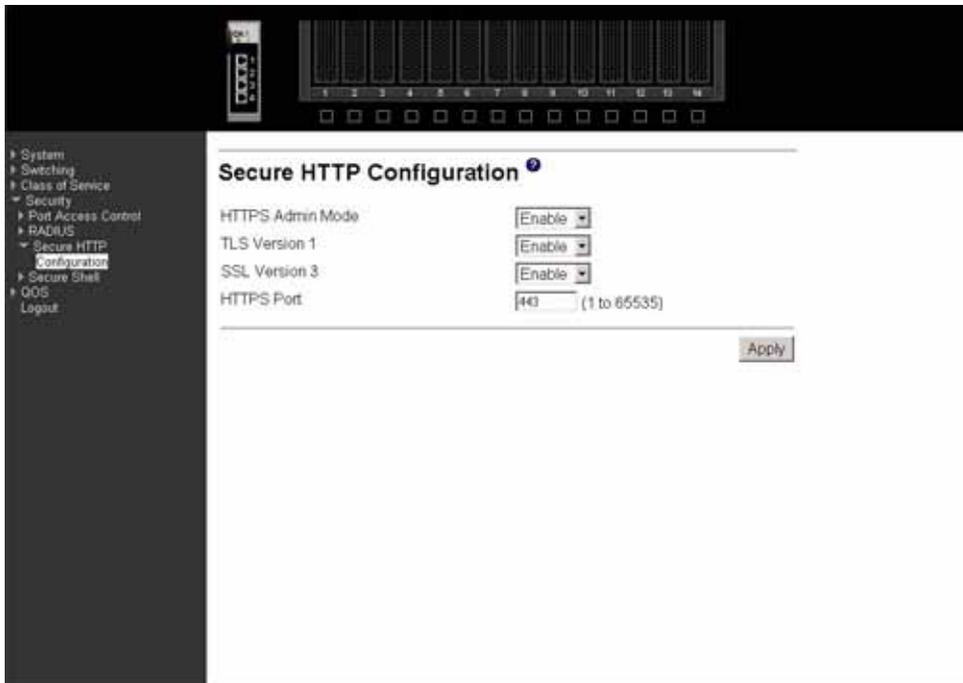
Use this panel to reset all RADIUS statistics for the switch. Click the Clear button to clear the accounting server, authentication server and RADIUS statistics.

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.

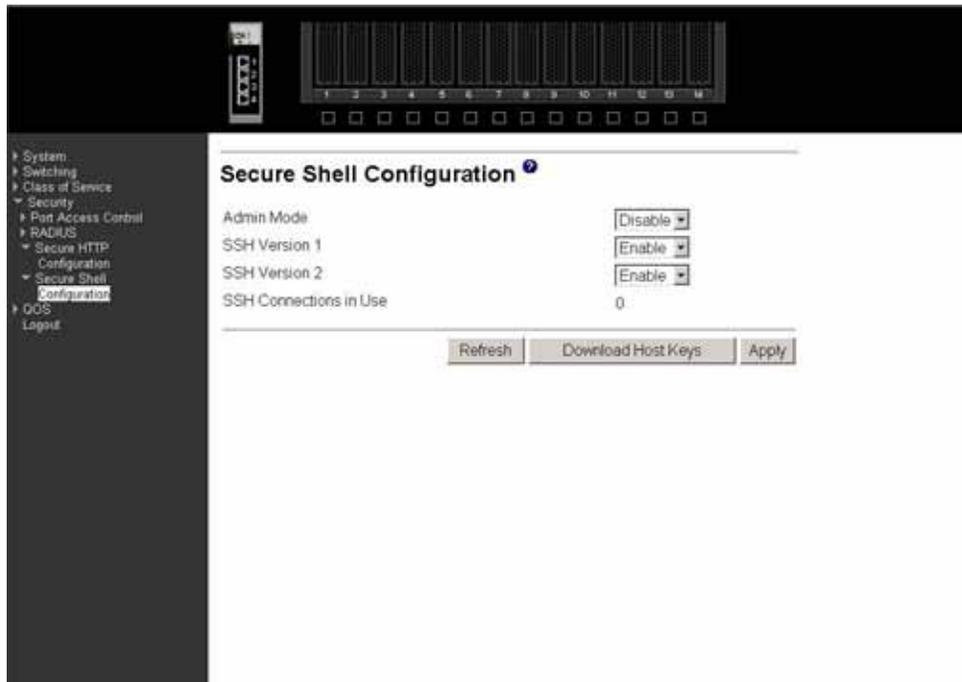
Click the Apply button to send the updated screen to the switch and have the changes take effect on the switch. If you want the switch to retain the new values across a power cycle you must perform a save.

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

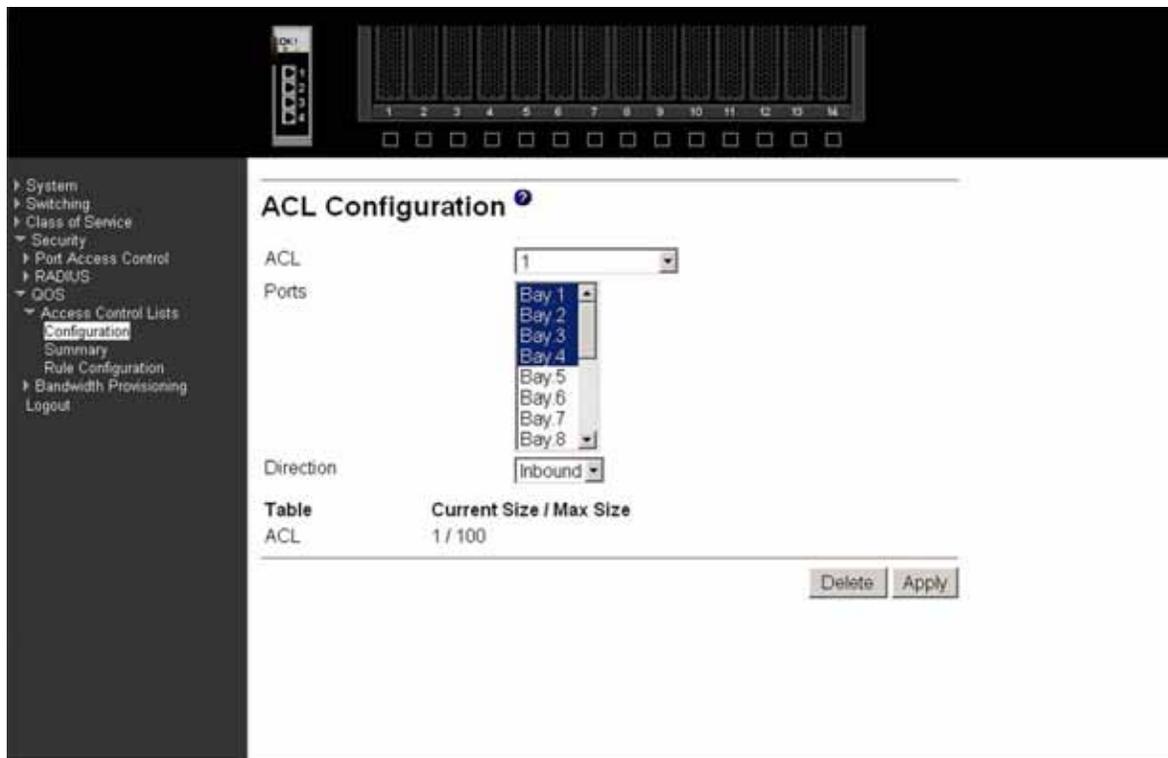
Access Control Lists

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.



ACL Make a selection from the pull-down menu. You may 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.

Direction

Select the packet filtering direction for the ACL from the pull-down menu. Currently the only choice is Inbound. The packet direction for a given 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.

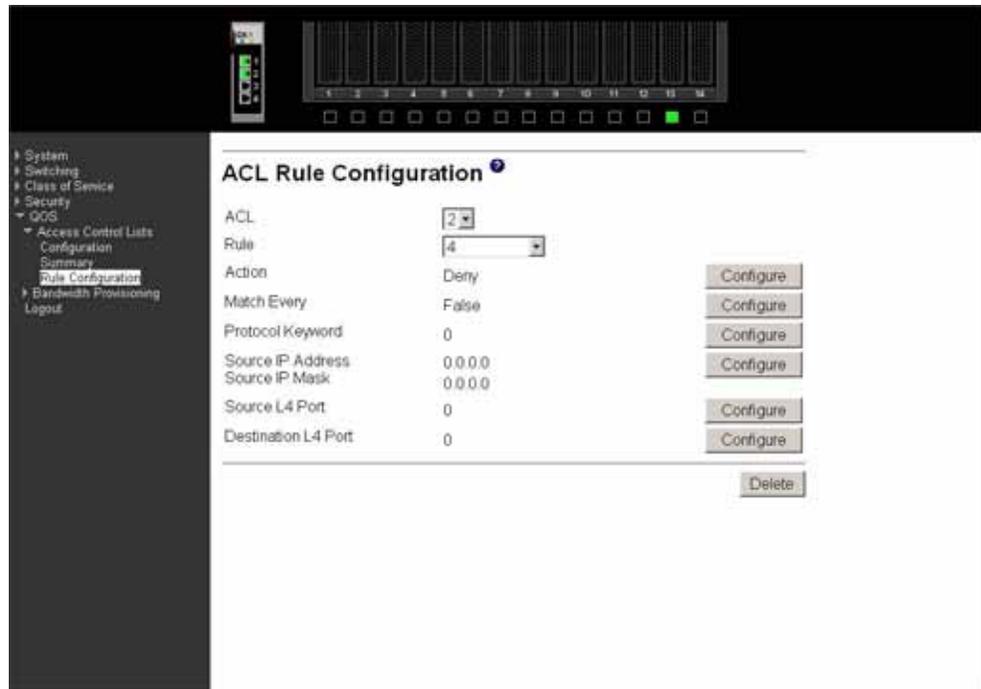
Click the Refresh button to update the screen with the latest information.

Rule configuration

This panel configures the rules associated with an ACL. When the screen first displays, you will see the first four fields 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 may 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.

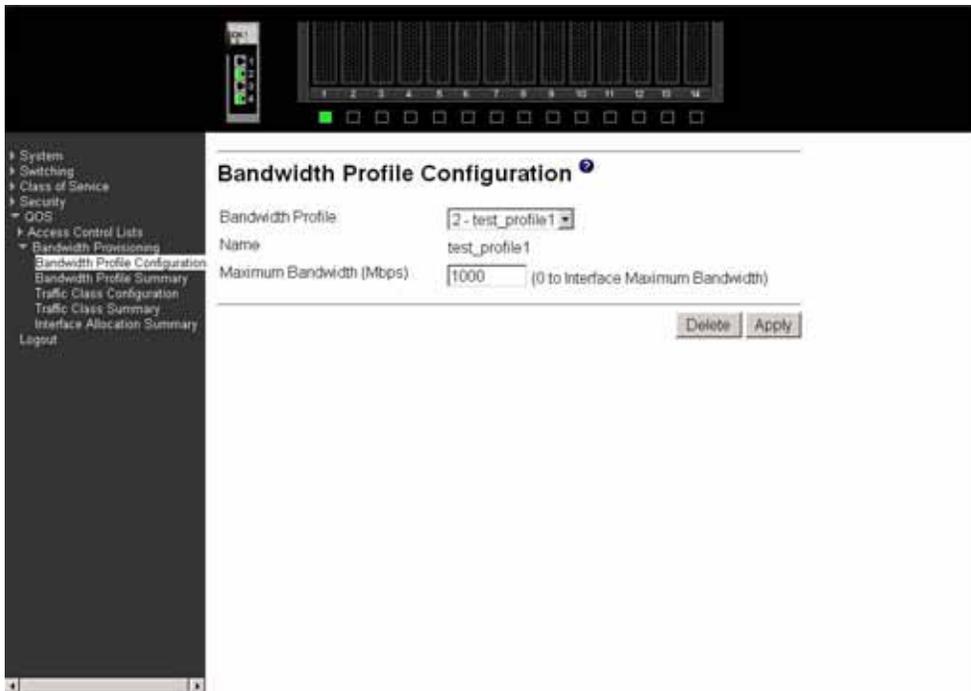
Bandwidth provisioning

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.



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 may enter up to 15 alphanumeric characters and may include the underscore _ or the dash -. You cannot change the name after the initial configuration.

Maximum Bandwidth

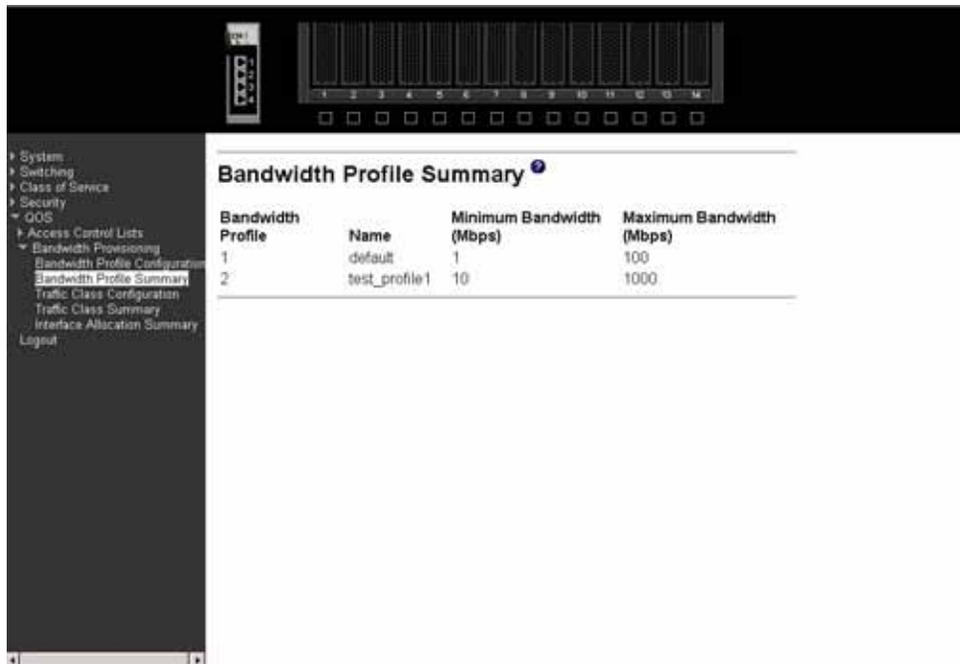
Enter the maximum allowable bandwidth for this bandwidth allocation profile.

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.



Bandwidth Profile

Displays the number associated with the bandwidth profile.

Name

Displays the name of the bandwidth profile.

Allocated Minimum Bandwidth

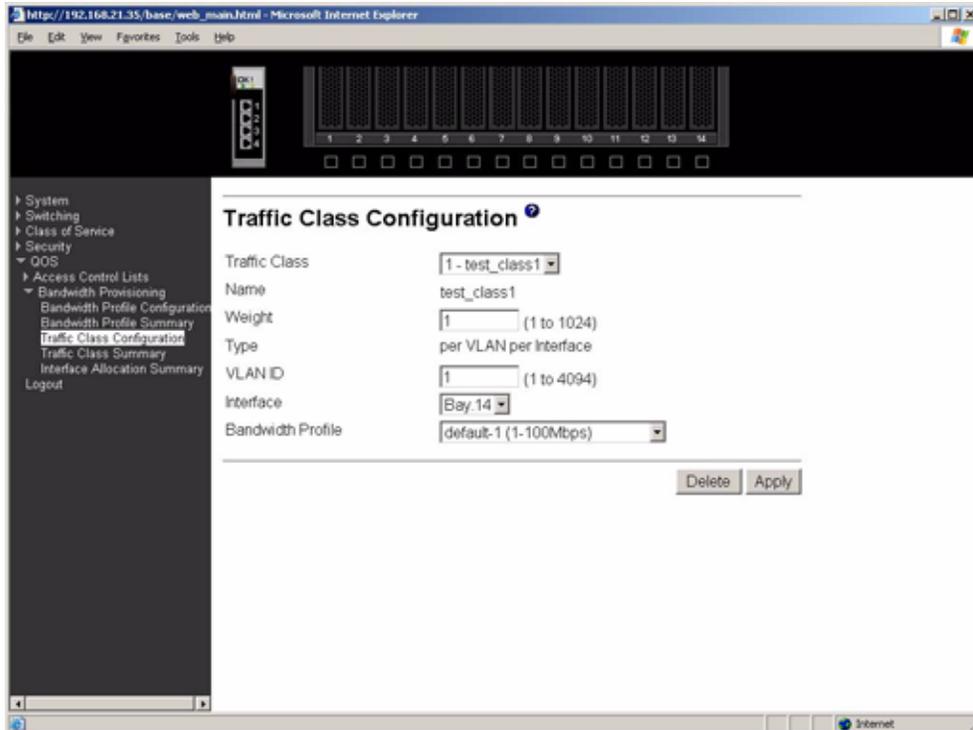
Displays the sum of the minimum guaranteed bandwidth for all bandwidth profiles configured on this interface.

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.



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 may enter up to 15 alpha-numeric characters and may 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 may 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	test_class1	1	0	per VLAN per Interface	1	Bay 14	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.

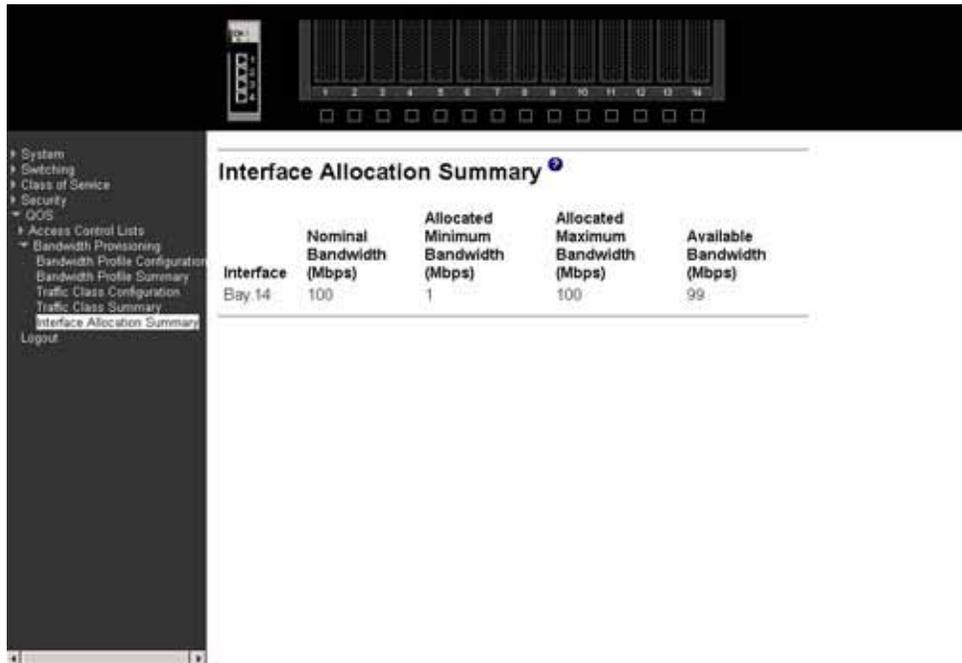
VLAN ID The VLAN ID with which this Traffic Class is associated.

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 14	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.

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."



6 Updating the Ethernet Switch Software

Two types of software run on the Ethernet switch module: the software image and the Micro-Controller Unit (MCU) code. You can update both the software image and the code, using either the switch module's CLI commands through a Telnet session, or by using the switch module's web interface.

This chapter describes how to:

- determine the software version that is currently installed on the switch module
- obtain the latest version of the switch software and upgrade the switch
- reset the switch module to activate the software upgrade.

Determining the software version

After you install the switch module in your SBCE unit, ensure that the latest software is installed on the switch module. To determine the version of the software that is installed on the switch module, complete the following steps:

1. Log on to the management module web interface. (Obtain the IP address of the management module from your system administrator.) The login window opens.
2. From the Monitors option on the left bar, choose Firmware VPD. The Firmware VPD window opens.
3. Under the I/O Module Firmware VPD section, locate the I/O module-bay number that displays the corresponding software version of the Ethernet switch module installed.

The screenshot shows the Management Module web interface. On the left is a navigation menu with options like Monitors, Blade Tasks, I/O Module Tasks, and MM Control. The main content area is titled 'Management Module' and contains three sections:

- Blade Server Firmware VPD:** A table with columns: Bay(s), Name, Firmware Type, Build ID, Released, Revision. It lists data for Bay 14, including BIOS, Diagnostics, and Blade sys. mgmt. proc.
- I/O Module Firmware VPD:** A table with columns: Bay, Type, Firmware Type, Build ID, Released, Revision. It lists data for Bay 1, including Ethernet SM, Boot ROM, Main Application 1, and Main Application 2.
- Management Module Firmware VPD:** A table with columns: Bay, Name, Firmware Type, Build ID, File Name, Released, Revision. This section is currently empty.

Below the Blade Server table, there is a note: "To reread firmware VPD for a blade, select the blade, and click 'Reload VPD'. This process may take a while." Below this note is a 'Target' dropdown menu set to 'All Blades' and a 'Reload VPD' button.

Obtaining the latest version

To determine the latest version of the switch module software, available from Intel, complete the following steps:

1. Go to <http://downloadfinder.intel.com>.
2. Enter "IXM5414E" in the download search box and click "Go".
A Results window opens, displaying a list of links to the current software update.
3. Compare the software version that you had noted from the I/O Module Firmware VPD section with the version of the latest software update. If the two software versions do not match, download the latest version to a TFTP server .

Upgrading the switch software

You upgrade the switch software using a TFTP server application that runs under the operating system in your server. Make sure that this software is installed on your server; then, download the software image or the MCU code from the Intel Web site into a directory on your TFTP server. Enable the TFTP server and set the directory that contains the image as the default directory.

Upgrading the image using Telnet

To upgrade the switch software using a Telnet session, complete the following steps:

1. Start a Telnet session to connect to the switch module.
2. Enter your user ID and password.
After successful login, the CLI prompt displays.
3. Enter *transfer download* and press Enter.
4. Enter *serverip ipaddress* where *ipaddress* is the IP address of your TFTP server, and press Enter.
5. Enter *datatype Code* and press Enter.
6. Enter *filename IXM54_nnn.opr* where *nnn* is the software sequence number of the new switch software, and press Enter.
7. Enter *start* and press Enter to begin the software upgrade process.
8. Review the information on the screen and enter 'y' when prompted to confirm the correctness of the entries.

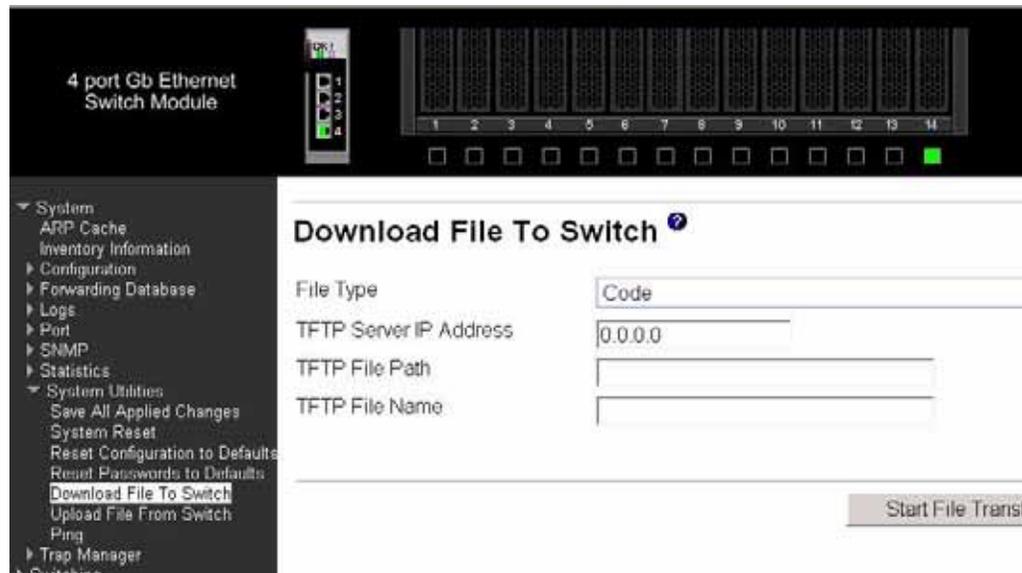
Upgrading the MCU code using Telnet

To upgrade the MCU code using a Telnet session, complete the following steps:

1. Start a Telnet session to connect to the switch module.
2. Enter your user ID and password.
After successful login, the CLI prompt displays.
3. Enter *transfer download* and press Enter.
4. Enter *serverip ipaddress* where *ipaddress* is the IP address of your TFTP server, and press Enter.
5. Enter *datatype 8051-mcu-code* and press Enter.
6. Enter *filename IXM54_MCU_{nnn}.hex*, where *nnn* is the software sequence number of the new switch MCU Code, and press Enter.
7. Enter *start* and press Enter to begin the software upgrade process.

- Review the information on the screen and enter 'y' when prompted to confirm the correctness of the entries.
After confirmation, the MCU code is transferred to the switch from the TFTP server. After successful update, the switch module is automatically powered off.
- To activate the new MCU code, turn on the Ethernet switch module through the management-module interface (see "Resetting and restarting the Ethernet switch module" Section)

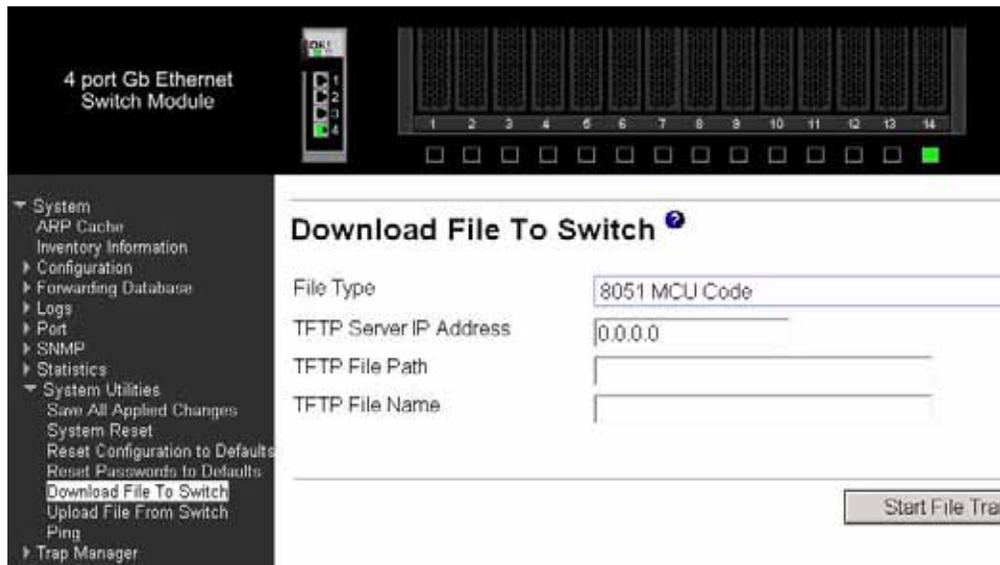
Upgrading the image using web interface



To upgrade the switch software using the web interface, complete the following steps:

- Log on to the management module web interface.
- From the I/O Module Tasks menu, click Management / Advanced Management.
- Click Start Web Session; then, log on to the Ethernet switch module.
- From the System Utilities menu, click Download File to Switch. The "Download File to Switch" window opens.
- Enter the following information in the "Download file to Switch" window.
 - In the File type field, select Code from the drop-down list.
 - In the TFTP Server IP Address field, enter the IP address of your TFTP server.
 - In the TFTP File Name field, enter filename *IXM54_nnn.opr* where *nnn* is the software sequence number of the new switch software.
- Click Start File Transfer to download the new switch software

Upgrading the MCU code using web interface



To upgrade the switch MCU code using the Web interface, complete the following steps:

1. Log on to the management module web interface.
2. From the I/O Module Tasks menu, click Management / Advanced Management.
3. Click Start Web Session; then, log on to the Ethernet switch module.
4. From the System Utilities menu, click Download File to Switch. The "Download File to Switch" window opens.
5. Enter the following information in the "Download file to Switch" window.
 - In the File type field, select 8051-MCU-Code from the drop-down list.
 - In the TFTP Server IP Address field, enter the IP address of your TFTP server.
 - In the TFTP File Name field, enter filename *IXM54_MCU nnn .hex* where *nnn* is the software sequence number of the new switch MCU code.
6. Click Start File Transfer to download the new switch software
After confirmation, the MCU code is transferred to the switch from the TFTP server. After a successful update, the switch module is then automatically powered off.
7. To activate the new MCU code, turn on the Ethernet switch module through the management-module interface.

Resetting and restarting the switch module

To activate the new image, you must restart the switch module through the management-module interface. Complete the following steps to reset the switch module:

1. From the I/O Module Tasks menu, click Management / Advanced Management.
2. Select the I/O-module bay on which the software update is installed.
3. From the I/O Module Tasks menu, click Power/Restart.

4. Click Power Off Module(s).
5. Select the I/O-module bay on which the software update is installed, again.
6. Click Power On Module(s). Wait 70 seconds for POST to be completed
7. Make sure that the latest switch operating-system software is correctly installed on the Ethernet switch module. From the Monitors menu, click Firmware VPD. The Firmware VPD window opens.
8. In the Firmware VPD window, locate the I/O Module Firmware VPD section. Scroll down to the number of the I/O-module bay that contains the Ethernet switch module that you just installed; then, note the corresponding level of the software for the switch module. Make sure that the latest switch operating-system software is correctly installed on the Ethernet switch module.

Management Module

Bay 1: WMN189288729

- Monitors:
 - System Status
 - Event Log
 - LEDs
 - Fuel Gauge
 - Hardware VPD
 - Firmware VPD**
- Blade Tasks:
 - Power/Restart
 - Remote Control
 - Firmware Update
 - Configuration
 - Serial Over LAN
- I/O Module Tasks:
 - Power/Restart
 - Management
 - Firmware Update
- MM Control:
 - General Settings

Blade Server Firmware VPD

Bay(s)	Name	Firmware Type	Build ID	Released
14	SN#ZJ1TRL3AW10J	BIOS	BSEC05DUS	07/30/2003
		Diagnostics	BSOT08AUS	07/02/2003
		Blade sys. mgmt. proc.	BR8T17A	n/a

To reread firmware VPD for a blade, select the blade, and click "Reload VPD". This process may take a while.

Target:

I/O Module Firmware VPD

Bay	Type	Firmware Type	Build ID	Released	Revision
1	Ethernet SM	Boot ROM	BRISM02	07/02/2004	3.04

7 Command Line Interface Management

Your Intel® Blade Server Ethernet Switch Module IXM5414E supports a management interface that you can use to set up and control your device over the network using the TCP/IP Telnet protocol. You can use this facility to perform the same network management functions that you can perform using the Web Interface. You can also use the Telnet interface to configure the switch module for management using an SNMP-based network management system. This chapter describes how to use the CLI to access the IXM5414E switch module, change its settings, and monitor its operation.

Important: Before you configure your IXM5414E switch module, be sure that the management modules in your SBCE unit are properly 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 *Resource* CD for more information.

Command Line Interface (CLI) conventions

The Command Line Interface (CLI) syntax, conventions and terminology are described in this section. Each CLI command referenced in this document is illustrated using the structure outlined below.

Format

Some commands, such as **show inventory**, do not require parameters. Other commands, such as **config lag deleteport**, have parameters for which you must supply a value. Parameters are positional — you must type the values in the correct order. Optional parameters will follow required parameters. For example:

config vlan mcaststorm <1-4094> <enable/disable> [*packets per second*]

- **config vlan mcaststorm** is the command name.
- <1-4094> <enable/disable> are the required values for the command.
- [*packets per second*] is the optional value for the command.

config lag deleteport <logical port> <port/listofports/all>

- **config lag deleteport** is the command name.
- <logical port> <port/listofports/all> are the required values for the command. Please note that usually the actual value of the parameter as seen in the CLI, e.g. <1-4094>, is used in the documentation. In some instances a generic term(s) such as <port/listofports/all> must be used since listing all possible choices is not possible.

Command name

The following conventions apply to the command name:

- The command name is displayed in this document in bold font and must be typed exactly as shown.
- Once you have entered enough letters of a command name to uniquely identify the command, hitting the space bar or Tab key will cause the system to complete the word.

- Entering Ctrl-Z will return you to the root level command prompt.

Parameters

The following conventions apply to the parameters:

- Parameters are order dependent.
- Parameters are displayed in this document in bold italic font, which must be replaced with a name or number.
- To use spaces as part of a name parameter, enclose it in double quotes, for example, “System Name with Space”.
- Parameters may be required or optional, and may have a list of choices.
 - *<parameter>* The angle brackets indicate that the parameter is required and you must enter a value in place of the brackets and text.
 - *[parameter]* The square brackets indicate that the parameter is optional and you may choose to enter a value in place of the brackets and text.
 - *choice1/choice2* Enter one and only one of the values listed.

Values

Some parameters are used frequently. This section explains the format you should use when providing values for them.

ipaddr	Enter a valid IP address made up of four decimal digits ranging from 0 to 255. The default for all IP addresses consists of zeros (that is, 0.0.0.0). The interface IP address of 0.0.0.0 is invalid. In some cases, you can also enter the IP address as a 32-bit number.
macAddr	The MAC address format is six hexadecimal numbers separated by hyphens, for example 00-06-29-32-81-40.
port	This is used to identify a physical interface, in the form of bay.port for an I/O module bay and ext.port for an external port. You enter a name and number separated by a period, for example: <ul style="list-style-type: none"> bay.1 identifies I/O module bay 1 ext.4 identifies external port 4
listofports	This is a comma-delimited list of valid ports, in the form of bay.port,bay.port or ext.port,ext.port. Port lists must NOT contain spaces and each interface must have its prefix specified (for example: bay.10,ext.2,bay.1)
logical port	This is used to identify a logical interface – a Link Aggregation Group or a VLAN. You enter a name and number separated by a period, for example: <ul style="list-style-type: none"> lag.3 identifies LAG 3 vlan.2 identifies VLAN 2

character strings

Use double quotation marks to identify character strings, for example, “System Name with Spaces”. An empty string (“”) is not valid.

Comments

When you are writing a test or configuration script you may add comments by using the “#” character to flag the beginning of a comment. The comment flag character can begin a word anywhere on the command line, and all input following this character will be ignored. Any command line that begins with the character “#” is recognized as a comment line and is ignored by the parser.

For example:

```
#Script file for displaying the ip interface  
#Display information about interfaces  
show ip interface ext.1 #Displays information about the first external interface  
#Display information about the next interface  
show ip interface ext.2  
#End of the script file
```

Special characters

Certain special key combinations speed up use of the CLI. They are listed in this section. Also, help is available for the CLI by typing **HELP**:

DEL, BS	delete previous character
Ctrl-A	go to beginning of line
Ctrl-E	go to end of line
Ctrl-F	go forward one character
Ctrl-B	go backward one character
Ctrl-D	delete current character
Ctrl-H	display command history or retrieve a command.
Ctrl-U, X	delete to beginning of line
Ctrl-K	delete to end of line
Ctrl-W	delete previous word
Ctrl-T	transpose previous character
Ctrl-P	go to previous line in history buffer
Ctrl-N	go to next line in history buffer
Ctrl-Z	return to root command prompt
Tab, <SPACE>	command-line completion
Exit	go to next lower command prompt
!!	execute the most recent command
!-n	execute the nth most recent command

!n	execute the nth command in history buffer
!str	execute the most recent command that starts with the string “str”.
!*str	execute the most recent command that contains the string “str”.
?	list choices

Remotely managing the IXM5414E switch module

The IXM5414E switch module supports two remote-access modes for management over Ethernet connections. You can select the mode that is best suited for your environment. The switch module has an internal Ethernet path to the management module and its four external Ethernet ports.

- 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 10/100 Mbps Ethernet port on the management module. With this mode, the IP addresses and SNMP parameters of the Ethernet switch modules can be manually assigned through the *SBCE Management and Configuration Program*. This mode allows you to provide a secure LAN for management of the platform’s subsystems separately from the data network.

Important: In this mode, the IXM5414E switch module does not respond to remote management commands from the four external Ethernet ports on the switch module.

See the applicable *Installation and User’s Guide* publications on the *Resource CD* for additional instructions for configuring the Intel® Blade Server Ethernet Switch Module IXM5414E for this mode of operation.

- You can choose to enable remote management of the IXM5414E 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 only be enabled through the management module configuration interface. Once this mode is enabled, the external Ethernet ports will support both management traffic and data traffic. Also, the Ethernet switch module will be able to transmit DHCP request frames through the external Ethernet ports.

This mode allows the switch module 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 chassis subsystems through the management module. However, management access to the IXM5414E switch module link will be lost if the switch module IP address is not on the same subnet as the management module. This chapter contains additional instructions for configuring the switch module for this mode of operation.

The two previously described modes are only applicable to the Intel® Blade Server Ethernet Switch Module IXM5414E. The management module can only be remotely accessed through the 10/100 Mbps Ethernet port on the management module.

Connecting to the IXM5414E switch module

When you know the IP address for your switch module and have an existing network connection, you can use the Telnet program (in VT-100 compatible terminal mode) to access and control the switch module. If you need to obtain the IP address for your switch module or establish a network connection, consult your system or network administrator. Be sure to use the correct IP address in the required command, as specified in this section.

The IXM5414E switch module supports user-based security that you can use to prevent unauthorized users from accessing the switch module or changing its settings. This section tells you how to log on to the switch module for the first time.

Complete the following steps to connect to the switch module through the Telnet interface:

1. Display a window that contains a DOS prompt command line; for example, **C:\>**.
2. Type the following command on the DOS prompt command line and press Enter: **telnet x.x.x.x** where **x.x.x.x** is the IP address for your switch module

When you first connect to the switch module, you will be prompted to enter a user ID followed by a password. Enter **USERID** in response to the prompt for a user ID and enter **PASSWORD** in response to the prompt for a password (notice the use of the zero and not the “O”). This will give you Read/write access to the switch module. By default, the switch module has one Read-only account named “**GUEST.**” The password for the Read-only GUEST account is left blank, just press Enter. For security you should change these default passwords after you log onto the system for the first time.

/ NOTE

All user IDs and passwords are CASE SENSITIVE.

Only a user with Read/write privileges can add new user accounts or make changes to existing user accounts. Another function available with a Read/write account is updating firmware and configuration files.

Changing configuration settings

The IXM5414E switch module has two levels of memory: normal random-access memory (RAM) and non-volatile RAM (NVRAM). When you enter a configuration change, the new settings will be immediately applied to the switching software in RAM. The new settings will remain in effect until the switch is restarted or you make another change. To make the changes permanent you need to issue the **save config** command which stores the current configuration in NVRAM. When the switch configuration settings have been saved to NVRAM, they become the default settings for the switch. These settings will be used every time the switch module is restarted.

/ NOTE

Some settings require you to restart the switch before they will take effect. Make sure you save the new configuration to NVRAM first.

There are two ways to change the configuration stored in NVRAM:

- Save a new configuration using the **save config** command.
- Reset all configuration values to the initial settings listed in Appendix Appendix C “Run-time Switching Software Default Settings” on page 231 by issuing the **clear config** command. This restores the configuration settings that were entered at the factory and causes a reboot. Loading the factory default configuration will erase any user accounts (and all other configuration settings) that you might have entered and return the switch module to its original state at the time of purchase.

Managing user accounts

Access to the IXM5414E switch module is controlled through an authorized user ID and password. The switch supports a maximum of six user accounts, only one of which can have Read/write privileges. The interface does not permit deletion of the currently logged-in user in order to prevent accidentally deleting all the users with Root privileges.

To log in after you have created a registered user, enter **login** at a command line prompt:

1. Type your user ID when prompted and press Enter.
2. Type your password when prompted and press Enter.

/ NOTE

The passwords used to access the switch module ARE case-sensitive.

Only the user with Read/write privileges can add new user accounts or make changes to existing user accounts. Before you can update a user account, you must also enter the password (if any) for that user account.

Complete the following steps to update a user account:

1. Enter the **config users passwd** command with the name of the account and the new password as parameters
2. Enter the old password when prompted, or just press enter if the account did not have a password

To delete a user account simply enter the **config users delete** command with the name of the account.

Initial configuration

Some settings must be entered to enable the IXM5414E switch module to be managed from an SNMP-based Network Management System such as SNMP version 1 or to be able to access the switch module using the Telnet protocol.

If the management of the Ethernet switch module is enabled through the four external ports, the switch module will acquire its IP address from a Dynamic Host Configuration Protocol (DHCP) server when the switch module is turned on or reset. The DHCP protocol enables IP addresses, network masks, and default gateways to be assigned by a DHCP server.

IXM5414E switch module system commands

This section describes the commands that you use to configure and manage the switch. These commands include:

- System information and statistics commands
- System configuration commands
- System description commands
- System utility commands
- Trap management commands

Later sections describe the commands that you use to configure and manage the various protocols running on the switch.

System commands

These commands display and configure system information and statistics.

Address Resolution Protocol (ARP) cache

show arp switch

Use this command to display the connectivity between the switch and other devices. The Address Resolution Protocol (ARP) cache identifies the MAC addresses of the IP stations communicating with the switch.

Format **show arp switch**

MAC Address

A unicast MAC address of a device on a subnet attached to one of the switch's routing 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.

Forwarding DB

config forwardingdb agetime

Use this command to configure the forwarding database address aging timeout.

Default 300

Format **config forwardingdb agetime <seconds>**

Seconds The <seconds> parameter must be within the range of 10 to 1,000,000 seconds.

show forwardingdb agetime

Use this command to display the address aging timeout for the forwarding database.

Format **show forwardingdb agetime**

Agetime The address aging timeout for the forwarding database in seconds.

show forwardingdb learned

Use this command to display forwarding database entries for learned addresses.

Format **show forwardingdb learned**

show forwardingdb table

Use this command to display the forwarding database entries. If the command is entered with no parameter, the entire table is displayed. This is the same as entering the optional <all> parameter. Alternatively, you can enter a MAC address to display the table entry for that address and all entries following it.

Format **show forwardingdb table**

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.

Port The physical interface on which the MAC address was learned.

ifIndex	The ifIndex of the MIB interface table entry associated with the port.
Status	The status of the entry. The possible values are:
	Static
	The value of the corresponding instance was added by the system or a user and cannot be relearned.
	Learned
	The entry was learned by observing the source MAC addresses of incoming traffic, and is currently in use.
	Management
	The system MAC address, identified with Bay.1.
	Self
	The MAC address of one of the switch's physical interfaces.

Inventory information

show inventory

Use this command to display inventory information for the switch.

Format **show inventory**

Switch 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 two-octet code that identifies the manufacturer.

Burnedin MAC Address

The burned-in universally administered MAC address of this switch.

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 Element

Identifies the network processor hardware.

Additional Packages

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

Logs

show eventlog

Use this command to display the event log, which is used to hold error messages for catastrophic events. After the event is logged and the updated log is saved in FLASH memory, the switch will be reset. The log can hold at least 2,000 entries (the actual number depends on the platform and OS), and is erased when an attempt is made to add an entry after it is full.

Format	show eventlog
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.
Task Id	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.

/ NOTE

Event log information is retained across a switch module reset.

show msglog

Use this command to display the message log. The message log contains system trace information that records non-critical problems.

Format	show msglog
Time	The time the event occurred, calculated from the time the switch was last reset.
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.

/ NOTE

Message log information is not retained across a switch module reset and wraps after 512 entries.

Port commands

System and configuration

config port adminmode

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

Default	enable
Format	config port adminmode <port/listofports/all> <enable/disable>

config port autoneg

Use this command to enable or disable automatic negotiation on one or more ports.

Default	enable
----------------	--------

Format **config port autoneg** <port/listofports/all> <enable/disable>

config port flowcontrol

Use this command to enable or disable IEEE 802.3x flow control for one or more ports.

Default disable

Format **config port flowcontrol** <port/listofports/all> <enable/disable>

config port lacpmode

Use this command to enable or disable the Link Aggregation Control Protocol (LACP) on one or more ports.

Default disable

Format **config port lacpmode** <port/listofports/all> <enable/disable>

config port linktrap

Use this command to enable or disable link status traps for one or more ports.

/ NOTE

This command is valid only when the Link Up/Down Flag is enabled (see “config trapflags linkmode” on page 186).

Format **config port linktrap** <port/listofports/all> <enable/disable>

config port physicalmode

Use this command to configure the speed and duplex mode for one or more ports. For this configuration to take effect, auto negotiation must be disabled.

Format **config port physicalmode** <port/listofports/all> <1000f/100f/100h/10f/10h>

Acceptable values are:

1000f **1000BASE-T full duplex**

100f **100BASE-T full duplex**

100h **100BASE-T half-duplex**

10f **10BASE-T full duplex**

10h **10BASE-T half duplex**

show port

Use this command to display port information.

Format **show port** <port/listofports/all>

Port The interface number of the physical port or LAG whose information is displayed on the line.

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

Mon Monitoring port, participating in Port Mirroring.

Probe

Probe port, participating in Port Mirroring.

LAG Member of a LAG.

Admin Mode

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

Physical Mode

Displays the port speed and duplex mode. If auto-negotiation is specified for the port, then the duplex mode and speed will be set by the auto-negotiation process. Note that the port's maximum capability (full duplex -100M) will be advertised. The factory default is auto.

Physical Status

Indicates the 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.

LACP Mode Displays whether Link Aggregation Control Protocol is enabled or disabled on this port.

FlowControl Mode

Displays whether flow control is enabled or disabled on this port.

Mirroring commands***config mirroring create***

Use this command to configure a probe port and a mirrored port for port mirroring. The first port is the probe port and the second port is the mirrored port. If this command is executed while port mirroring is enabled, it will have the effect of changing the probe and mirrored port values. The probe port will be removed from all VLANs.

Format **config mirroring create** *<port>* *<port>*

config mirroring delete

Use this command to remove the port mirroring designation from both the probe port and the mirrored port. The probe port must be manually re-added to any desired VLANs.

Format **config mirroring delete**

config mirroring mode

Use this command to configure the port mirroring mode. The possible values are enable and disable. The probe and mirrored ports must be configured before port mirroring can be enabled. If enabled, the probe port will mirror all traffic received and transmitted on the physical mirrored port. It is not necessary to disable port mirroring before modifying the probe and mirrored ports.

Default disable

Format **config mirroring mode** *<enable/disable>*

show mirroring

Use this command to display the port mirroring information for the switch module.

Format **show mirroring**

Port Mirroring Mode

Indicates whether the port mirroring feature is enabled or disabled.

Probe Port The port that is configured as the probe port. If this value has not been configured, 'Not Configured' will be displayed.

Mirrored Port

The port that is configured as the mirrored port. If this value has not been configured, 'Not Configured' will be displayed.

Simple Network Management Protocol (SNMP)

SNMP community commands

config snmpcommunity accessmode

Use this command to configure SNMP access to switch information for a specific community name. The access mode can be Read-only (also called public) or Read/write (also called private).

Format **config snmpcommunity accessmode <readonly/readwrite> <name>**

config snmpcommunity create

Use this command to add (and name) a new SNMP community. A community name associates the switch with a set of SNMP managers with a specified privileged level. The name can be up to 16 case-sensitive characters long.

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.

Default 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.

Format **config snmpcommunity create <name>**

config snmpcommunity delete

Use this command to remove a name from the SNMP community table.

Format **config snmpcommunity delete <name>**

config snmpcommunity ipaddr

Use this command to specify 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 IP mask before being compared to this IP address. Note that if the 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". The parameter <name> is the applicable community name, and may be up to 16 alphanumeric characters.

Default 0.0.0.0

Format **config snmpcommunity ipaddr <ipaddr> <name>**

config snmpcommunity ipmask

Specify the mask to be ANDed with the requesting entity's IP address before comparison with the SNMP community IP address associated with the same community name. If the result matches the SNMP community IP address then the address is an authenticated IP address. For example, if the IP address = 9.47.128.0 and the corresponding IP 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". The parameter <name> is the applicable community name, and may be up to 16 alphanumeric characters.

Default 0.0.0.0

Format **config snmpcommunity ipmask <ipmask> <name>**

config snmpcommunity mode

Use this command to activate or deactivate 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.

Default The default private and public communities are enabled by default. The four undefined communities are disabled by default.

Format **config snmpcommunity mode <enable/disable> <name>**

show snmpcommunity

Use this command to display SNMP community information.

Up to 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 SNMP Version 1 (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).

Format **show snmpcommunity**

SNMP Community Name

The community name of this row of the table.

Client IP Address

An 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

The mask that will 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 Client IP 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 The access level for this community. Either Read/write or Read-only.
Status The status of this community. Either enable or disable.

SNMP trap commands

config snmptrap create

Use this command to add an SNMP trap receiver community name and associated IP address. The maximum length of name is 16 case-sensitive alphanumeric characters.

Format **config snmptrap create <name> <ipaddr>**

config snmptrap delete

Use this command to delete a trap receiver from a community.

Format **config snmptrap delete <name> <ipaddr>**

config snmptrap ipaddr

Use this command 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.

Format **config snmptrap ipaddr <ipaddrold> <name> <ipaddrnew>**

config snmptrap mode

Use this command to enable or disable an 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).

Format **config snmptrap mode <enable/disable> <name> <ipaddr>**

show snmptrap

Use this command to display information about SNMP trap receivers. Trap messages are sent across the 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.

Format **show snmptrap**

SNMP Trap Name

The community string of the SNMP trap packet sent to the trap manager. Note that trap receiver communities and SNMP communities are separate and distinct.

IP Address The IP address that receives SNMP traps from the switch for this trap receiver community.

Status Indicates whether traps are currently enabled for this community

- Enable -**
traps will be sent
- Disable -**
traps will not be sent.

System configuration

Network connectivity

config network javamode

Use this command to enable or disable the java applet that displays a picture of the switch module at the top right of the screen when you are using the Web interface. 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.

- Default** enable
- Format** **config network javamode** *<enable/disable>*

config network parms

Use this command to set the IP Address, subnet mask and gateway of the router. The IP Address and the gateway must be on the same subnet.

- Default** 10.90.90.9x 255.255.255.0 0.0.0.0
- Format** **config network parms** *IPaddress netmask [gateway]*

config network protocol

Use this command to specify the network configuration protocol. Select DHCP, BootP, or None. The change is effective immediately.

- Default** none
- Format** **config network protocol** *none/bootp/dhcp*
- Values**
 - none** The network interface is manually configured with static IP information.
 - bootp** The network interface is configured to acquire an IP configuration using the Bootstrap Protocol (BootP).
 - dhcp** The network interface is configured to acquire an IP configuration using the Dynamic Host Configuration Protocol (DHCP). If you select DHCP and management of the Ethernet switch module through the four ports is enabled, the switch module will acquire its IP address, network mask, and default gateway from a DHCP server when the switch module is turned on or reset.

config network webmode

Use this command to enable or disable access to the switch module via the Web interface. When access is enabled a user can login to the switch from a web browser through TCP port 80. Disabling access takes effect immediately on all interfaces.

- Default** enable
- Format** **config network webmode** *<enable/disable>*

show network

Use this command to display network configuration settings that are necessary for in-band connectivity.

Format **show network**

IP Address The IP address of the interface. The factory default value is 10.90.90.9x, where x is determined by the the number of the I/O-module bay in which the Ethernet switch module is installed. See Table 1. “Default IP addresses based on I/O module bay numbers” on page 21. The IP address may have been acquired from a DHCP server or may be the static IP address acquired from the management module. See the Network Configuration Protocol item for more information.

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.

Burned In MAC Address

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

Network Configuration Protocol

Indicates what network protocol was used on the last or current power-up cycle, if any. The configuration methods are DHCP, BootP and none. The factory default method is none.

Note: When management of the Ethernet Switch Module is enabled through the four external ports, and Dynamic Host Configuration Protocol (DHCP) is enabled, the switch module acquires its IP address from a DHCP server when the switch module is turned on or reset; otherwise, the switch module acquires a static IP address from the management module.

Web Mode Indicates whether the switch may be accessed from a web browser. If web mode is enabled you can manage the switch from a web browser. The factory default is enabled.

Java Mode Indicates whether the java applet that displays a picture of the switch at the top right of the screen is enabled or disabled. If the applet is enabled 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.

Telnet

config telnet maxsessions

Use this command to configure the number of simultaneous Telnet and Secure Shell (SSH) sessions that can be established. A value of 0 indicates that no Telnet session can be established. The range is 0 to 5.

Default 5

Format **config telnet maxsessions <0-5>**

config telnet mode

Use this command to allow or disallow new Telnet and SSH sessions. If sessions are enabled, new Telnet sessions can be established until there are no more sessions available. If sessions are disabled, no new Telnet sessions are established but an established session will remain active until the session is terminated or an abnormal network error ends it.

Default enable

Format **config telnet mode** *<enable/disable>*

config telnet timeout

Use this command to specify the number of minutes of inactivity that will occur on a Telnet or SSH session before the switch logs off. A value of 0 indicates there will be no timeout and the session will remain active indefinitely. The time is a decimal value from 0 to 160.

Changing the timeout value does not affect an active session until the session is reaccessed. Any keystroke will also activate the new timeout duration.

Default 5

Format **config telnet timeout** *<0-160>*

show telnet

Use this command to display Telnet settings.

Format **show telnet**

Telnet Login Timeout (minutes)

The number of minutes of inactivity that will occur on a Telnet or SSH session before the switch logs off. A value of zero means there will be no timeout.

Maximum Number of Telnet Sessions

The number of simultaneous Telnet and SSH sessions allowed.

Allow New Telnet Sessions

Indicates whether new Telnet and SSH sessions are allowed.

User accounts

config users add

Use this command to add a new user account if the maximum number of users has not been reached. The *<name>* can be up to eight alphanumeric characters and is case-sensitive. A maximum of six user IDs can be defined.

Format **config users add** *<name>*

config users delete

Use this command to remove a user account.

Format **config users delete** *<name>*

/ NOTE

The admin user account cannot be deleted.

config users passwd

Use this command to change the password of an existing user. The password is up to eight alphanumeric characters and is case-sensitive.

After you enter this command you will be prompted for the user's current password. If none, press enter.

Default Blank (indicating no password) for users with Read-only access. For those with Read/write access the factory standard password is "PASSWORD." Please note the use of zero instead of the letter "O."

Format **config users passwd <user>**

config users snmpv3 accessmode

Use this command to specify the SNMPv3 access privileges for the specified user account. The valid accessmode values are <readonly> or <readwrite>. The <user> is the login user name for which the specified access mode will apply.

Default readwrite for admin user; readonly for all other users

Format **config users snmpv3 accessmode <user> <readonly/readwrite>**

config users snmpv3 authentication

Use this command to specify 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. The <user> is the user account for which the specified authentication protocol will be used.

Default no authentication

Format **config users snmpv3 authentication <user> <none/md5/sha>**

config users snmpv3 encryption

Use this command to specify the encryption protocol and key to be used to authenticate a user account. The valid encryption protocols are none or DES. The DES protocol requires a key, which can be specified on the command line. The key may be up to 16 characters long. If the DES protocol is specified but a key is not provided, you will be prompted for the key. If none is specified as the protocol, you may not enter a key. The <user> is the user account for which the specified encryption protocol will be used.

Default no encryption

Format **config users snmpv3 encryption <user> <none/des [key]>**

show users info

Use this command to display the configured user names and their settings. This command is only available for the user with Read/write privileges.

Format **show users info**

User Name The name the user will use to login using the serial port, Telnet or Web.

User Access Mode

Shows 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 only be one Read/write user and up to five Read-only users.

SNMPv3 Access Mode

Displays the SNMPv3 Access Mode. 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 does not have to be the same as the CLI and Web access modes.

SNMPv3 Authentication

The protocol (if any) that will be used to authenticate the user.

SNMPv3 Encryption

The encryption protocol (if any) that will be used for the authentication process.

Login

config loginsession close

Use this command to close a specified Telnet session.

Format **config loginsession close <sessionid/all>**

show loginsession

Use this command to display currently active Telnet and serial port connections to the switch.

Format show loginsession

ID Login Session ID

User Name The account name used to login via the serial port or Telnet.

Connection From

The IP address of the Telnet client machine or EIA-232 for the serial port connection.

Idle Time Time this session has been idle.

Session Time

Total time this session has been connected.

System description

config prompt

Use this command to change the prompt that is displayed when you use the CLI. You may enter up to 64 alphanumeric characters.

Format **config prompt <system prompt>**

config syscontact

Use this command to configure the name of the person or organization responsible for the switch. The range for name is from 1 to 31 alphanumeric characters.

Format **config syscontact <contact>**

config syslocation

Use this command to configure the physical location assigned to the switch. The range for name is from 1 to 31 alphanumeric characters.

Format **config syslocation <location>**

config sysname

Use this command to configure the name assigned to the switch. The range for name is from 1 to 31 alphanumeric characters.

Format **config sysname <name>**

show stats port detailed

Use this command to display detailed statistics for a specified port.

Format **show stats port detailed <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).

Packets Received Successfully**Total Packets Received Without Error**

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.

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/Undersized 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.

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.

Packets Transmitted

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 1522 octets in length (excluding framing bits but including FCS octets).

Max Info

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

Packets Transmitted Successfully

Total Packets Transmitted Successfully

The total number of packets that have been transmitted by this port to its segment.

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 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.

Transmit Discards**Total Transmit Packet 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.

Protocol Statistics**BPDUs Received**

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

BPDUs Transmitted

The number of BPDUs (Bridge Protocol Data Units) transmitted from the spanning tree layer.

802.3x Pause Frames Transmitted

The number 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.

GVRP PDUs Received

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

GVRP PDUs Transmitted

The number of GVRP PDUs transmitted by the GARP layer.

GVRP PDUs Failed Registrations

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

GMRP PDUs Received

The number of GMRP PDUs received.

GMRP PDUs Transmitted

The number of GMRP PDUs transmitted.

GMRP PDUs 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.

show stats port summary

Use this command to display a summary of the statistics for a specified port.

Format **show stats port summary <port>**

Packets Received Without Error

The total number of packets (including multicast and broadcast packets) received on this port.

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 Error

The total number of packets transmitted from the interface.

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.

show stats switch detailed

Use this command to display detailed statistics for all CPU traffic.

Format **show stats switch detailed**

Received**Octets Received**

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

Packets Received Without Errors

Total number of packets received on the network

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.

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.

VLAN Entries**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 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.

show stats switch summary

Use this command to display a summary of the statistics for all switch traffic.

Format **show stats switch summary**

Packets Received Without Error

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

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.

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 the switch were last cleared.

show sysinfo

Use this command to display switch information.

Format **show sysinfo**

Switch Description

The product name of the switch.

System Name

The name used to identify the switch.

System Location

Text used to identify the location of the switch. May be up to 31 alphanumeric characters. The factory default is blank.

System Contact

Text used to identify a contact person for the switch. May be up to 31 alphanumeric characters. The factory default is blank.

System ObjectID

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 the switch.

System utilities

System utility commands

The commands in this section allow you to fine tune your systems performance and functionality.

clear config

Use this command 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 will be prompted to confirm that the reset should proceed.

Format **clear config**

clear igmpsnooping

Use this command to clear the tables managed by the Internet Group Management Protocol (IGMP) Snooping function. The switch will attempt to delete these entries from the Multicast Forwarding Database (MFDB). You will be prompted to confirm that you want to issue this command

Format **clear igmpsnooping**

clear lag

Use this command to clear all LAGs. You will be prompted to confirm that you want to issue this command.

Format **clear lag**

clear pass

Use this command to reset all user passwords to the factory defaults. You will be prompted to confirm that the password reset should proceed.

Format **clear pass**

clear stats port

Use this command to clear the statistics for a specified port. You will be prompted to confirm that you want to issue this command.

Format **clear stats port <port/listofports/all>**

clear stats switch

Use this command to clear the statistics for the switch. You will be prompted to confirm that you want to issue this command.

Format **clear stats switch**

clear transfer

Use this command to reset the file transfer parameters to the factory defaults. You will be prompted to confirm that you want to issue this command.

Format **clear transfer**

clear traplog

Use this command to clear the trap log. You will be prompted to confirm that you want to issue this command.

Format **clear traplog**

clear vlan

Use this command to reset the VLAN configuration parameters to the factory defaults. You will be prompted to confirm that you want to issue this command.

Format **clear vlan**

logout

Use this command to close the current Telnet connection or reset the current serial connection. If you have any saved configuration changes, you will be prompted to save them.

If you logout without issuing a **save config** command any configuration changes you have made will be lost.

Format **logout**

ping

Use this command 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. The switch will send three Ping requests and display the results. The switch can be pinged from any IP workstation with which it is connected through the default VLAN (VLAN 1), as long as there is a physical path between the switch and the workstation.

Format **ping <ipaddr>**

reset system

Use this command to reset the switch without powering it 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 will be prompted to confirm that the reset should proceed. A successful reset is indicated by the LEDs on the switch.

Format **reset system**

save config

Use this command to permanently save configuration changes made since the previous save or reboot to Non-Volatile Random Access Memory (NVRAM). You are prompted to verify your choice.

Format **save config**

show history

Use this command to show the contents of the command history buffer. The output will display the oldest command in the history buffer first and the **show history** command (the newest command) last.

Format **show history**

Transfer download commands

transfer download datatype

Use this command to configure the type of file to be downloaded to the switch.

Default **code**

Format **transfer download datatype <code/config>**

transfer download filename

Use this command to specify the name of the file that is to be downloaded to the switch. The switch will remember the last file name used.

You may specify the file path as part of the file name if the string is less than 31 characters. Otherwise, use the transfer download path command.

This command is valid only when the Transfer Mode is TFTP. See transfer download mode.

Format **transfer download filename <name>**

transfer download path

Use this command to specify the directory path on the TFTP server where the file to be downloaded to the switch is located. The switch will remember the last file path used.

This command is valid only when the Transfer Mode is TFTP. See **transfer download mode**. Details of the TFTP path are explained under the command **transfer upload path**.

Format **transfer download path <path>**

transfer download serverip

Use this command to configure the IP address of the server on which a file to be downloaded is located.

This command is valid only when the transfer mode is TFTP. See **transfer download mode**.

Default 0.0.0.0

Format **transfer download serverip <ipaddr>**

transfer download start

Use this command to start a download transfer. After the current settings are displayed you will be prompted to confirm your decision. This command will close your connection to the host.

Format **transfer download start**

The following information fields are displayed:

TFTP Server IP

The IP address of the server where the file is to be downloaded.

TFTP Path

The directory path specification for the file to be downloaded.

TFTP Filename

The name of the file to be downloaded.

Data Type

The type of file to be downloaded: config, error log, message log or trap log.

Transfer upload commands

TFTP upload example

This example shows three ways to specify the same 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 shown below:

Table 5. TFTP Upload Scenarios

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

The directory path statement can be cleared by issuing the **clear config** command.

Format **transfer upload path <path>**

transfer upload datatype

Use this command to specify the type of file to be uploaded from the switch.

Format **transfer upload datatype** <config/errorlog/msglog/traplog>

The datatype is one of the following:

config	Configuration file
errorlog	Error log
msglog	Message log
traplog	Trap log (the default)

transfer upload filename

Use this command to specify the name of the file to be uploaded from the switch. The switch will remember the last file name used.

You may specify the file path as part of the file name if the string is less than 31 characters. Otherwise, use the **transfer upload path** command to specify the directory path.

This command is valid only when the Transfer Mode is TFTP. See **transfer upload mode**.

Format **transfer upload filename** <name>

transfer upload path

Use this command to specify the directory path on the TFTP server where you want to save a file uploaded from the switch. The switch will remember the last file path used.

/ NOTE

This command is valid only when the transfer mode is TFTP. See the command, **transfer upload mode**.

The Intel® Blade Server Ethernet Switch Module IXM5414E 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 may remain blank. See the following path setup example.

transfer upload serverip

Use this command to configure the IP address of the server on which a file to be uploaded is to be located.

It is valid only when the transfer mode is TFTP. See “transfer upload mode”.

Default 0.0.0.0

Format **transfer upload serverip** <ipaddr>

transfer upload start

Use this command to start an upload transfer. After the current settings are displayed you will be prompted to confirm your decision. Note that issuing this command will close your connection to the host.

Format **transfer upload start**

The following information fields are displayed:

TFTP Server IP Address

The Internet Protocol (IP) address of the server where the file is to be uploaded.

TFTP File Path

The directory path specification for the file to be uploaded.

TFTP File Name

The name to be given to the file after it has been uploaded.

File Type

The type of file to be uploaded: config, error log, message log or trap log.

Trap manager

config trapflags authentication

Use this command to enable or disable the Authentication Flag, which determines whether a trap message is sent when the switch detects an authentication failure.

Default enable

Format **config trapflags authentication <enable/disable>**

config trapflags linkmode

Use this command to enable or disable 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.

Default enable

Format **config trapflags linkmode <enable/disable>**

config trapflags multiusers

Use this command to enable or disable 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.

Default enable

Format **config trapflags multiusers <enable/disable>**

config trapflags stpmode

Use this command to enable or disable STP traps. When enabled, topology change notification trap messages will be sent.

Default enable

Format **config trapflags stpmode <enable/disable>**

show trapflags

Use this command to display 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.

Format **show trapflags**

Authentication Flag

Indicates whether authentication failure traps will be sent (enable) or not (disable).

Link Up/Down Flag

Indicates whether a trap will be sent when the link status changes from up to down or vice versa.

Multiple Users Flag

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 serial port).

Spanning Tree Flag

Indicates whether spanning tree traps will be sent.

show traplog

Use this command to display the trap log.

Format **show traplog****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.

/ NOTE

Trap log information is not retained across a switch module reset.

Switching configuration commands

This section describes the commands you use to manage the switch and to show the current management settings.

This section also provides detailed explanations of said switching commands. The commands are divided into nine groups:

- Generic Attributes Registration Protocol (GARP) commands
- IGMP snooping commands
- Link Aggregation (LAG) commands
- MAC filter commands
- Mirroring commands
- Multicast Forwarding Database (MFDB) commands
- Spanning tree commands
- Virtual Local Area Network (VLAN) commands

Generic Attribute Registration Protocol (GARP) commands

config garp gmrp adminmode

Use this command to enable or disable the GARP Multicast Registration Protocol (GMRP) on the switch module.

Default disable

Format **config garp gmrp adminmode** *<enable/disable>*

config garp gmrp interfacemode

Use this command to enable or disable the GMRP on one, some or all interfaces. If an interface which has GARP enabled is enabled for routing or is made a member of a LAG, GARP functionality will be disabled on that interface. GARP functionality will subsequently be re-enabled if routing is disabled or LAG membership is removed from an interface that previously had GARP enabled.

Default disable

Format **config garp gmrp interfacemode** *<port/listofports/all><enable/disable>*

config garp gvrp adminmode

Use this command to enable or disable GVRP on the switch module.

Default disable

Format **config garp gvrp adminmode** *<enable/disable>*

config garp gvrp interfacemode

Use this command to enable or disable GVRP for one, some or all interfaces. If GVRP is disabled, Join Time, Leave Time and LeaveAll Time have no effect.

Default disable

Format **config garp gvrp interfacemode** *<port/listofports/all> <enable/disable>*

config garp jointimer

Use this command to configure the GARP Join Time for the specified port(s). Join Time is the interval between the transmission of GARP Protocol Data Units (PDUs) registering (or re-registering) membership for a VLAN or multicast group.

This command has an effect only when GVRP is enabled. The time may range from 10 to 100 centiseconds.

Default 20 centiseconds (0.2 seconds)

Format **config garp jointimer** *<port/listofports/all> <10-100>*

config garp leavealltimer

Use this command to configure how frequently LeaveAll PDUs are generated for the specified port(s). A LeaveAll PDU indicates that all registrations will be unregistered. Participants would need to rejoin in order to maintain registration. The value applies per port and per GARP participation. The time may range from 200 to 6000 centiseconds.

This command has an effect only when GVRP is enabled.

Default 1000 centiseconds (10 seconds)

Format **config garp leavealltimer** *<port/listofports/all> <200-6000>*

config garp leavetimer

Use this command to configure the GARP Leave Time for the specified port(s). Leave Time is the time to wait after receiving an unregister request for a VLAN or a multicast group before deleting the VLAN entry or group. This allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service. The time may range from 20 to 600 centiseconds.

This command has an effect only when GVRP is enabled.

Default 60 centiseconds (0.6 seconds)

Format `config garp leavetimer <port/listofports/all> <20-600>`

show garp info

Use this command to display GARP information for the IXM5414E switch module.

Format `show garp info`

GMRP Admin Mode

This displays the administrative mode of GMRP for the switch module. The default is disable.

GVRP Admin Mode

This displays the administrative mode of GVRP for the IXM5414E switch module. The default is disable.

show garp interface

Use this command to display GARP information for one, some or all interfaces.

Format `show garp interface <port/listofports/all>`

Port This displays the identification of the interface that this row in the table describes.

Join Timer Displays the interval between the transmission of GARP PDUs registering (or re-registering) membership for an attribute. Current attributes are a VLAN or a multicast group. There is an instance of this timer on a per-port, per-GARP participant basis. Permissible values are 10 to 100 centiseconds (0.1 to 1.0 seconds) in increments of 1 centisecond (0.01 seconds). The factory default is 20 centiseconds (0.2 seconds).

Leave Timer Displays the period of time to wait after receiving an unregister request for an attribute before deleting the attribute. Current attributes are a VLAN or a multicast group. This allows time for another station to assert registration for the same attribute in order to maintain uninterrupted service. There is an instance of this timer on a per-port, per-GARP participant basis. Permissible values are 20 to 600 centiseconds (0.2 to 6.0 seconds) in increments of 1 centisecond (0.01 seconds). The factory default is 60 centiseconds (0.6 seconds).

LeaveAll Timer

Shows 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. There is an instance of this timer on a per-port, per-GARP participant basis. The LeaveAll Period Time is set to a random value in the range of LeaveAll Time to (1.5*LeaveAll Time). Permissible values are 200 to 6000 centiseconds (2 to 60 seconds) in increments of 1 centisecond (0.01 seconds). The factory default is 1000 centiseconds (10 seconds).

Port GMRP Mode

Indicates the GMRP administrative mode for the port. It may be enabled or disabled. If this parameter is disabled, Join Time, Leave Time and LeaveAll Time have no effect. The factory default is disabled.

Port GVRP Mode

Indicates the GVRP administrative mode for the port. It may be enabled or disabled. If this parameter is disabled, Join Time, Leave Time and LeaveAll Time have no effect. The factory default is disabled.

IGMP snooping commands

config igmpsnooping adminmode

Use this command to enable or disable IGMP Snooping on the switch module.

Default disable

Format **config igmpsnooping adminmode <enable/disable>**

config igmpsnooping groupmembershipinterval

Use this command to configure the IGMP Group Membership Interval time on the IXM5414E switch module. The group membership interval time is the amount of time in seconds that a switch will wait for a report from a particular group on a particular interface before deleting the interface from the entry. This value must be greater than the IGMP maximum response time value. The range is 2 to 3600 seconds.

Default 260 seconds

Format **config igmpsnooping groupmembershipinterval <2-3600>**

config igmpsnooping interfacemode

Use this command to enable or disable IGMP Snooping on a selected interface. The <port/listofports/all> parameter identifies the interface(s) on which to enable or disable IGMP Snooping. If an interface which has IGMP Snooping enabled is enabled for routing or becomes a member of a LAG, IGMP Snooping functionality will be disabled on that interface. IGMP Snooping functionality will subsequently be re-enabled if routing is disabled, or the interface is deleted from the LAG.

Default disable

Format **config igmpsnooping interfacemode <port/listofports/all> <enable/disable>**

config igmpsnooping maxresponse

Use this command to configure the IGMP Maximum Response time on the IXM5414E switch module. The maximum response time is the amount of time in seconds that a switch will wait after sending a query on an interface because it did not receive a report for a particular group in that interface. This value must be less than the IGMP query interval time value. The range is 1 to 3599 seconds.

Default 10 seconds

Format **config igmpsnooping maxresponse <1-3599>**

config igmpsnooping mcrtrexpiretime

Use this command to configure the Multicast Router Present Expiration time on the switch module. This is the amount of time in seconds that a switch will wait for a query to be received on an interface before the interface is removed from the list of interfaces with multicast routers attached. The range is 0 to 3600 seconds. A value of 0 indicates an infinite timeout, i.e. the time never expires.

Default 0

Format `config igmpsnooping mcrtrexpiretime <0-3600>`

show igmpsnooping

Use this command to display IGMP Snooping information for the IXM5414E switch module. Configuration information is displayed whether or not IGMP Snooping is enabled. Status information is only displayed when IGMP Snooping is enabled.

Format `show igmpsnooping`

Admin Mode

This indicates whether or not IGMP Snooping is enabled on the switch.

Group Membership Interval (secs)

This displays the IGMP Query Interval Time. This is the amount of time the switch will wait for a report for a particular group on a particular interface before it sends a query on that interface.

Max Response Time (secs)

This displays the amount of time the switch will wait after sending a query on an interface because it did not receive a report for a particular group on that interface.

Multicast Router Present Expiration Time (secs)

If a query is not received on an interface within this amount of time, the interface is removed from the list of interfaces with multicast routers attached.

Interfaces Enabled for IGMP Snooping

This is the list of interfaces on which IGMP Snooping is enabled.

The following status value is only displayed when IGMP Snooping is enabled.

Multicast Control Frame Count

This displays the number of multicast control packets that have been processed by the CPU.

Link Aggregation (LAG) commands

config lag addport

Use this command to add a physical port to a LAG. The first interface parameter designation is of a configured LAG and the second identifies the port to be added. There can be a maximum of 8 member ports.

Format `config lag addport <logical port> <port>`

config lag adminmode

Use this command to enable or disable the specified LAG(s). The option <all> sets every configured LAG to the same administrative mode setting.

Format `config lag adminmode <logical port/listofports/all> <enable/disable>`

config lag create

Use this command to configure a new LAG, assign a name and generate a logical port number for it. To display the assigned logical port number use the **show lag** command. The <name> parameter is a string of up to 15 alphanumeric characters.

Format **config lag create <name>**

config lag deletelag

Use this command to delete the specified LAG(s). The <all> option removes all configured LAGs.

Format **config lags deletelag <logical port/listofports/all>**

config lag deleteport

Use this command to delete one or more ports from a LAG. The first interface parameter designates a configured LAG. The second interface number designates a port that is a member of the LAG. Use <all> to delete all ports in the specified LAG.

Format **config lag deleteport <logical port> <port/listofports/all>**

config lag linktrap

Use this command to enable or disable link trap notifications for the specified LAG. The option <all> sets every configured LAG to the same administrative mode setting.

Default enable

Format **config lag linktrap <logical port/listofports/all> <enable/disable>**

config lag name

Use this command to define a name for the specified LAG. Name is an alphanumeric string up to 15 characters. Use this command to modify the name that was associated with the LAG when it was created.

Format **config lag name <logical port> <name>**

show lag

Use this command to display an overview of all link aggregation groups (LAGs) on the switch.

Format **show lag <logical port/listofports/all>**

Logical Port The logical port identifying the LAG, in the format lag.port.

LAG Name The name of this LAG.

Link State Indicates whether the link is up or down.

Admin Mode

The administrative mode. The factory default is enabled.

Link Trap Mode

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

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.

Mbr 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 given LAG.
Port Speed	The speed of the LAG. A LAG is always full-duplex.

MAC filter commands

config macfilter adddest

Use this command to add the <port> to the destination filter set for the MAC filter with the MAC address of <macaddr> and VLAN of <vlan>. The <macaddr> parameter must be specified as a 6-byte hexadecimal number in the format of 00-12-34-56-78-90. The <vlan> parameter must identify a valid VLAN.

The <port> parameter identifies the destination port(s) to be added to the destination port filter set for the MAC filter. If <all> is selected, all ports will be added to the destination port filter set. Packets for the specified MAC address and VLAN ID will only be transmitted out of ports that are in the filter set.

Format **config macfilter adddest <macaddr> <vlan> <port/listofports/all>**

config macfilter create

Use this command to add a static MAC filter entry for a MAC address and VLAN pair. The <macaddr> parameter must be specified as a 6-byte hexadecimal number in the format of 00-12-34-56-78-90.

Filters may not be defined for 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

The <vlan> parameter must identify a valid VLAN.

Up to 100 static MAC filters may be created.

Format **config macfilter create <macaddr> <vlan>**

config macfilter deldest

Use this command to remove one or more ports from the destination filter set for the MAC filter with the MAC address of <macaddr> and VLAN of <vlan>. The <macaddr> parameter must be specified as a 6-byte hexadecimal number in the format of 00-12-34-56-78-90. The <vlan> parameter must identify a valid VLAN.

The <port> parameter identifies the destination port(s) to be removed from the destination port filter set for the MAC filter. If <all> is selected, all ports will be removed from the destination port filter set.

Format **config macfilter deldest <macaddr> <vlan> <port/listofports/all>**

config macfilter remove

Use this command to remove the static MAC filter entry for the given MAC address on the VLAN. The <macaddr> parameter must be specified as a 6-byte hexadecimal number in the format of 00-12-34-56-78-90. The <vlan> parameter must identify a valid VLAN.

Format **config macfilter remove <macaddr> <vlan>**

show macfilter

Use this command to display the Static MAC Filtering information. If <all> is selected as the first parameter, all the Static MAC Filters in the switch module are displayed. If a <macaddr> is entered, a VLAN ID must also be entered and the Static MAC Filter information will be displayed only for that MAC address and VLAN ID pair.

Format **show macfilter <all/macaddr <all/vlan>>**

MAC Address

The MAC address of the static MAC filter entry.

VLAN ID The VLAN ID of the static MAC filter entry.

Destination Port(s)

The port(s) in the destination filter. Packets with the associated MAC address and VLAN ID will only be transmitted out of ports in the list.

Multicast Forwarding Database (MFDB) commands

show mfdb gmrp

Use this command to display the GMRP entries in the Multicast Forwarding Database (MFDB) table.

Format **show mfdb gmrp**

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.

Interfaces The list of interfaces that are designated for forwarding (Fwd:) and filtering (Flt:).

show mfdb igmpsnooping

Use this command to display the IGMP Snooping entries in the MFDB.

Format **show mfdb igmpsnooping**

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.

Interfaces The list of interfaces that are designated for forwarding (Fwd:) and filtering (Flt:).

show mfdb staticfiltering

Use this command to display the Static Filtering entries in the MFDB.

Format **show mfdb staticfiltering**

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.

Interfaces The list of interfaces that are designated for forwarding (Fwd:) and filtering (Flt:).

show mfdb stats

Use this command to display the MFDB statistics.

Format **show mfdb stats**

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 the switch was reset. This value is also known as the MFDB high-water mark.

Current Entries

Displays the current number of entries in the MFDB table.

show mfdb table

Use this command to display the MFDB information. If the command is entered with no parameter, the entire table is displayed. This is the same as entering the optional <all> parameter. The user can display the table entry for one MAC address by specifying the MAC address as an optional parameter.

Format **show mfdb table [macaddr/all]**

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 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.

Component The component that is responsible for this entry in the MFDB. Possible values are IGMP Snooping, GMRP, and Static Filtering.

Description The text description of this multicast table entry.

Interfaces The list of interfaces that are designated for forwarding (Fwd:) and filtering (Flt:).

Forwarding Interfaces

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.

Spanning tree commands

Spanning tree bridge commands

config spanningtree bridge forwarddelay

Use this command to configure the Bridge Forward Delay parameter to a new value. Forwarddelay is used by bridges to ensure that a new network topology has stabilized before leaving the blocking state. The forwarddelay value is in whole seconds within a range of 4 to 30, with the value being greater than or equal to $((\text{Bridge Max Age} / 2) + 1)$.

Default 15

Format `config spanningtree bridge forwarddelay <4-30>`

config spanningtree bridge hellotime

Use this command to configure the Hello Time parameter to a new value. Hellotime determines how often a hello message is broadcast; it cannot be longer than MaxAge but should be longer than forwarddelay. The hellotime value is in whole seconds within a range of 1 to 10 with the value being less than or equal to $((\text{Bridge Max Age} / 2) - 1)$.

Default 2

Format `config spanningtree bridge hellotime <1-10>`

config spanningtree bridge maxage

Use this command to configure the Bridge Max Age parameter to a new value. This is the value that all bridges use for maxage when this bridge is acting as the root: A BPDU will be discarded when its age exceeds maxage. The maxage value is in whole seconds within a range of 6 to 40, with the value being less than or equal to $(2 \text{ times } (\text{Bridge Forward Delay} - 1))$.

Default 6

Format `config spanningtree bridge maxage <6-40>`

config spanningtree bridge priority

Use this command to configure the Bridge Priority parameter to a new value. The bridge priority value is the first two octets of the eight octet Bridge ID. This value is a number between 0 and 61440. The lower the number the higher the priority. The twelve least significant bits will be masked according to the IEEE 802.1s specification. This will cause the priority to be rounded down to the next lower valid priority.

Default 32768

Format `config spanningtree bridge priority <0-61440>`

show spanningtree bridge

Use this command to display the STP settings for the bridge.

Format `show spanningtree bridge`

Bridge Priority

The priority component of the bridge identifier. Valid values range from 0-61440, in increments of 4096. The lower the number the higher the priority. The factory default is 32768.

Bridge Identifier

The unique identifier associated with this bridge instance. It consists of the bridge priority and the bridge's base MAC address.

Bridge Max Age

The value that all bridges use for Max Age when this bridge is acting as the root: a BPDU will be discarded when its age exceeds maxage.

Bridge Hello Time

The value that all bridges use for HelloTime when this bridge is acting as the root. HelloTime determines how often a hello message is broadcast; it cannot be longer than maxage but should be longer than forwarddelay.

Bridge Forward Delay

The value that all bridges use for Forward Delay when this bridge is acting as the root. Forwarddelay is used by bridges to ensure that a new network topology has stabilized before leaving the blocking state. Note that IEEE 802.1D specifies that the range for this parameter is related to the value of STP Bridge Maximum Age.

Bridge Hold Time

Minimum time between transmission of Configuration Bridge Protocol Data Units (BPDUs).

Spanning tree Common Spanning Tree (CST) commands

config spanningtree cst port edgeport

Use this command to specify whether a port is an edge port within the Common Spanning Tree (CST). This will allow the port to transition to Forwarding State without delay. The <port> is the port to be affected. The edgeport value can either be "true" or "false".

Default false

Format config spanningtree cst port edgeport <port> <true/false>

config spanningtree cst port pathcost

Use this command to configure the path cost to a new value for the specified port in the CST. The <port> is the port to be affected. The pathcost value can be specified as a number in the range of 1 to 200000000 or auto. If <auto> is specified, the pathcost value will be set based on Link Speed.

Default auto

Format config spanningtree cst port pathcost <port> <1-200000000/auto>

config spanningtree cst port priority

Use this command to configure the port priority to a new value for use within the CST. The <port> is the port to be affected. The priority value is a number in the range of 0 to 240 in increments of 16.

Default 128

Format config spanningtree cst port priority <port> <0-240>

show spanningtree cst detailed

Use this command to display STP settings for the CST.

Format show spanningtree cst detailed

Bridge Priority

The value of the first two octets of the eight octet Bridge ID. Valid values are 0 to 61440. Factory default is 32768.

Bridge Identifier

The unique identifier associated with this bridge instance.

Time Since Topology Change

The time (in seconds) since the last time a topology change was detected by the bridge entity.

Topology Change Count

The total number of topology changes detected by this bridge since the management entity was last reset or initialized.

Topology Change in progress

Boolean value of the topology change parameter for the switch indicating whether a topology change is in progress on any port assigned to the CST.

Designated Root

The identifier of the bridge currently assumed to be the root of the spanning tree.

Root Path Cost

The cost of the path to the root as seen from this bridge.

Root Port Identifier

The port number of the port which offers the lowest cost path from this bridge to the root bridge.

Root Port Max Age

The maximum age of STP information learned from the network on any port before it is discarded.

Root Port Bridge Forward Delay

The value that all bridges use for forwarddelay when this bridge is acting as the root. Values range from 4 to 30. The Factory default is 15 seconds.

Hello Time

The amount of time between the transmission of Configuration BPDUs by this node or any port when it is the root of the spanning tree or trying to become the root.

Bridge Hold Time

Minimum time between transmission of Configuration BPDUs.

CST Regional Root

The regional root bridge.

Regional Root Path Cost

The cost of the path to the regional root as seen from this bridge.

Associated FIDs

List of forwarding database identifiers currently associated with this bridge instance.

Associated VLANs

List of VLAN IDs currently associated with this bridge instance.

show spanningtree cst port detailed

Use this command to display the settings and parameters for a specific switch port within the CST. The <port> is the port to be affected.

Format **show spanningtree cst port detailed <port>**

Port Identifier

The port identifier for this port within the CST.

Port Priority

The priority of the port within the CST.

Port Forwarding State

The forwarding state of the port within the CST.

Port Role The role of the specified interface within the CST.

Auto-calculate Port Path Cost

Indicates whether automatic calculation of the port path cost is enabled.

Port Path Cost

The configured path cost for the specified interface.

Designated Port Cost

Path Cost offered to the LAN by the designated port.

Designated Bridge

The bridge containing the designated port.

Designated Port Identifier

Port used to forward frames towards the root bridge for this CST on this LAN. It is the port with the lowest cost path to the bridge and the highest port priority.

Topology Change Acknowledgement

Value of flag in next Configuration BPDU transmission indicating if a topology change is in progress for this port.

Hello Time The hello time in use for this port.

Edge Port The configured value indicating if this port is an edge port.

Edge Port Status

The derived value of the edge port status. True if operating as an edge port; false otherwise.

Point To Point MAC Status

Derived value indicating if this port is part of a point to point link.

CST Regional Root

The regional root identifier in use for this port.

CST Path Cost

The configured path cost for this port.

show spanningtree cst port summary

Use this command to display the status of one, some or all ports within the CST. The parameter <port/listofports/all> indicates the port or ports to be affected.

Format **show spanningtree cst port summary** <port/listofports/all>

Port The interface being displayed.

STP Mode Whether the STP is enabled or disabled on the port.

STP State The port's current spanning tree state. This state controls what action a port takes on receipt of a frame. Possible states are: disabled, blocking, listening, learning, forwarding and broken.

Port Role The role of the specified port within the spanning tree.

Link Status The operational status of the link. Possible values are "Up" or "Down".

Link Trap The link trap configuration for the specified interface.

Spanning tree port commands

config spanningtree port migrationcheck

Use this command to force the specified port to transmit RST BPDUs. The <port> parameter specifies the port(s) to be affected. To set the migration check for all ports with a single command, <all> can be specified. Note that the forceversion parameter for the switch must be set to 802.1w for this command to work.

Default disable

Format **config spanningtree port migrationcheck** <port/listofports/all> <enable/disable>

config spanningtree port mode

Use this command to configure the Administrative Switch Port State to a new value for the specified port. The <port> parameter specifies the port(s) to be affected. To enable or disable all ports with a single command, <all> can be specified. Note that a maximum of 4095 ports can be enabled.

Default disable

Format **config spanningtree port mode** <port/listofports/all> <enable/disable>

show spanningtree port

Use this command to display the STP statistics for a specific switch port.

Format **show spanningtree port** <port>

Port mode Enabled or disabled.

Port Up Time Since Counters Last Cleared

The time in days, hours, minutes, and seconds since the counters were last reset.

STP BPDUs Transmitted

The number of STP BPDUs sent by this port.

STP BPDUs Received

The number of STP BPDUs received by this port.

RSTP BPDUs Transmitted

The number of Rapid Reconfiguration STP BPDUs sent by this port.

RSTP BPDUs Received

The number of Rapid Reconfiguration STP BPDUs received by this port.

Spanning tree summary commands

config spanningtree adminmode

Use this command to configure the STP operational mode. While the operational mode is disabled, the spanning tree configuration is retained and can be changed, but it is not activated.

Default disable

Format **config spanningtree adminmode** <enable/disable>

config spanningtree forceversion

Use this command to select which version of the STP will be used. The <version> can be one of the following:

- 802.1D - IEEE 802.1D functionality supported: STP BPDUs are transmitted rather than R(Rapid)STP BPDUs
- 802.1w - IEEE 802.1w functionality supported: RSTP BPDUs are transmitted rather than STP BPDUs

Default IEEE 802.1D
Format `config spanningtree forceversion <802.1D/802.1w>`

show spanningtree summary

Use this command to display STP settings and parameters for the switch.

Format `show spanningtree summary`

Spanning Tree Adminmode
Enabled or disabled.

Spanning Tree Version
Indicates which version of the STP is being run. Possible values are IEEE 802.1w, or IEEE 802.1D.

Configuration Digest Key
Calculated value used as part of the configuration identifier.

Configuration Format Selector
Identifies the level of the IEEE 802.1 standard in use by the switch.

Virtual Local Area Network (VLAN) commands

config vlan bcstorm

Use this command to enable or disable broadcast storm control for a particular Virtual Local Area Network (VLAN). If broadcast storm control is enabled, storms are controlled by counting the number of broadcast packets within a certain time period. If the [packets per second] count limit is exceeded, the packets are discarded.

Default disable

Format `config vlan bcstorm <1-4094> <enable/disable> [packets per second]`

config vlan create

Use this command to create a new VLAN and assign it an ID. The ID is a VLAN identification number in the range of 2-4094 (ID 1 is reserved for the default VLAN).

Format `config vlan create <2-4094>`

config vlan delete

Use this command to delete an existing VLAN. The ID is a valid VLAN identification number. The default VLAN cannot be deleted.

Format `config vlan delete <2-4094>`

config vlan makestatic

Use this command to change a dynamically created VLAN (one that is created by GVRP registration) to a static VLAN (one that is permanently configured and defined) The number identifies an existing VLAN.

Format `config vlan makestatic <2-4094>`

config vlan mcstorm

Use this command to enable or disable multicast storm control for a particular VLAN. If multicast storm control is enabled, storms are controlled by counting the number of multicast packets within a certain time period. If the [packets per second] count limit is exceeded, the packets are discarded.

Default disable
Format **config vlan mcaststorm** <1-4094> <enable/disable> [packets per second]

config vlan name

Use this command to change the name of a VLAN. The name is an alphanumeric string of up to 16 characters, and the number identifies an existing VLAN.

Default The name for VLAN ID 1 is always Default. The default name for other VLANs is a blank string.
Format **config vlan name** <name> <2-4094>

config vlan participation

Use this command to configure the degree of participation for a specific interface in a VLAN. The number identifies an existing VLAN, and the parameter <port/listofports/all> indicates the port or ports to be affected.

Format **config vlan participation** <exclude/include/auto> <1-4094> <port/listofports/all>

Participation options are:

include The interface is always a member of this VLAN. This is equivalent to registration fixed.
exclude The interface is never a member of this VLAN. This is equivalent to registration forbidden.
auto The interface is dynamically registered in this VLAN by GVRP. The interface will not participate in this VLAN unless a join request is received on this interface. This is equivalent to registration normal.

config vlan port acceptframe

Use this command to configure the frame acceptance mode for the specified port(s). Possible values are:

all Both tagged and untagged frames are accepted. Untagged frames will be assigned the PVID and default priority configured for the port(s) for this VLAN.
vlan Untagged frames are discarded.

With either option, VLAN tagged packets are forwarded in accordance with the IEEE 802.1Q VLAN Specification.

Default all
Format **config vlan port acceptframe** <all/vlanonly> <port/listofports/all>

config vlan port priority

Use this command to change the default IEEE 802.1p port priority assigned to untagged frames received on the specified port(s) for the specified VLAN.

Default 0
Format **config vlan port priority** <0-7> <port/listofports/all>

config vlan port pvid

Use this command to change the VLAN ID that the specified port(s) will assign to untagged frames if untagged frames are accepted.

Default 1
Format `config vlan port pvid <1-4094> <port/listofports/all>`

config vlan port tagging

Use this command to configure the tagging behavior for a specific interface in a VLAN. If tagging is enabled, all traffic is transmitted as tagged frames. If tagging is disabled, all traffic is transmitted as untagged frames. The parameter <port/listofports/all> indicates the port or ports to be affected.

Format `config vlan port tagging <enable/disable> <port/listofports/all>`

show vlan detailed

Use this command to display detailed information, including interface information, for a specific VLAN.

Format `show vlan detailed`

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 The type of VLAN. A VLAN can be:

- the Default VLAN (VLAN ID = 1)
- a static VLAN, one that is created using the config vlan create command
- a Dynamic VLAN, one that is created by GVRP registration

In order to change a VLAN from Dynamic to Static, use the config vlan makestatic command.

Broadcast Storm Control

Displays the administrative mode of broadcast storm control for this VLAN. The threshold value for broadcast storm control is in packets per second.

Multicast Storm Control

Displays the administrative mode of multicast storm control for this VLAN. The threshold value for broadcast storm control in packets per second.

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

Current Displays 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.

Configured Displays the configured 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	Displays the tagging behavior for 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.

show vlan port

Use this command to display VLAN port information.

Format `show vlan port <port/listofports/all>`

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 may 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.

GVRP Indicates whether GVRP is enabled or disabled on the port.

Default Priority

The IEEE 802.1p priority that will be assigned to untagged frames accepted on this port for this VLAN.

show vlan summary

Use this command to display information about all configured VLANs.

Format `show vlan summary`

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 is created using the **config vlan create** command
- a Dynamic VLAN, one that is created by GVRP registration

In order to change a VLAN from dynamic to static, use the **config vlan makestatic** command.

BcastStorm This displays the administrative mode of broadcast storm control for this 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.

McastStorm This displays the administrative mode of multicast storm control for this 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.

Class of Service commands

config classofservice 802.1pmapping

Use this command to map an User priority to a Traffic Class priority queue.

Default = see table below

Table 6. Classofservice 802.1p Mapping

IEEE 802.1p priority	IXE5416 priority queue
0	2
1	1
2	0
3	3
4	4
5	5
6	6
7	7

Format `config classofservice 802.1pmapping <0-7> <0-7>`

show classofservice 802.1pmapping

Use this command to show the current mapping of IEEE 802.1p priority values to traffic class priority queues.

Format `show classofservice 802.1pmapping`

User Priority

The IEEE 802.1p priority number. The range is 0 to 7.

Traffic Class Priority Queue

The priority queue number. The range is 0 to 7.

Security configuration commands

This section describes the commands used to configure and manage the security features of the Intel® Blade Server Ethernet Switch Module IXM5414E. These features include:

- Authentication commands
- IEEE 802.1X Port-based network access control
- Remote Authentication Dial-In User Service (RADIUS)
- Secure Shell (SSH) commands
- Secure Socket Layer (SSL) commands

Authentication commands

config authentication login create

Use this command to create an authentication login list. The <listname> is up to 15 alphanumeric characters and is case sensitive. Up to 10 authentication login lists can be configured on the switch. When a list is created, the authentication method “local” is set as the first method. Authentication methods can be changed using the config authentication login set command.

Format **config authentication login create <listname>**

config authentication login delete

Use this command to delete the specified authentication login list. The command will fail if any of the following conditions are true:

- The login list name is invalid or does not identify an existing login list
- The specified login list is currently assigned to a user or to the nonconfigured user
- The specified login list is the default login list included with the default configuration and was not created using the **config authentication login set** command.

Format **config authentication login delete <listname>**

config authentication login set

Use this command to configure an ordered list of methods for the specified authentication login list. You may specify up to three methods. The possible methods are local, radius, and reject.

The value of local indicates that the user’s locally stored ID and password should be used for authentication. The value of radius indicates that the user’s ID and password will be authenticated using the RADIUS server. The value of reject indicates that the user is never authenticated.

To authenticate a user, the authentication methods in the user’s login list will be attempted in order until an authentication attempt succeeds or fails.

Note that the default login list included with the default configuration can not be changed.

Format **config authentication login set <listname> <local/radius/reject>**
[local/radius/reject] [local/radius/reject]

config users defaultlogin

Use this command to assign the authentication login list to be used when a non-configured user attempts to log in to the system. This setting is overridden by the authentication login list assigned to

a specific user if the user is configured locally. If this value is not configured, users will be authenticated using local authentication only.

Format **config users defaultlogin <listname>**

config users login

Use this command to assign the specified authentication login list to the specified user for system login. The <user> must be a configured user and <listname> must be a configured login list.

If the user is assigned a login list that requires remote authentication, all access to the interface from CLI, web, and Telnet sessions will be blocked until the authentication is complete. Refer to the discussion of maximum delay in the config radius maxretransmit and config radius timeout commands.

Note that the login list associated with the user with Read/write privileges cannot be changed, to prevent accidental lockout from the switch.

Format **config users login <user> <listname>**

show authentication login info

Use this command to display the ordered authentication methods for all authentication login lists.

Format **show authentication login info**

Authentication Login List

The login list whose information is displayed on this line.

Method 1 The first method in the login list, if any.

Method 2 The second method in the login list, if any.

Method 3 The third method in the login list, if any.

show authentication login users

Use this command to display information about the users assigned to the specified login list. If the login list is assigned to non-configured users, the word “default” will appear as the user name.

Format **show authentication login users <listname>**

User The user assigned to the specified login list.

Component The component, either user or 802.1X, for which the login list is assigned.

show users authentication

Use this command to display all user and authentication login information for the switch, including the login list assigned to the default user.

Format **show users authentication**

User A list of all users with an assigned login list.

System login

The authentication login list assigned to the user for system login.

802.1X The authentication login list assigned to the user for IEEE 802.1X port security.

IEEE 802.1X commands

clear dot1x port stats

Use this command to reset the IEEE 802.1X statistics for the specified port(s).

Format **clear dot1x port stats <port/all>**

config dot1x adminmode

Use this command to enable or disable authentication support on the switch. The default value is disable. In disabled mode, the dot1x configuration is retained and can be changed, but it is not activated.

Default disable

Format **config dot1x adminmode <enable/disable>**

config dot1x defaultlogin

Use this command to assign the authentication login list to use for non-configured users for IEEE 802.1X port security. This setting is overridden by the authentication login list assigned to a specific user if the user is configured locally. If this value is not configured, users will be authenticated using local authentication only.

Format **config dot1x defaultlogin <listname>**

config dot1x login

Use this command to assign the specified authentication login list to the specified user for port security. The <user> must be a configured user and the <listname> must be a configured login list.

Format **config dot1x login <listname>**

config dot1x port controlmode

Use this command to configure the authentication mode to be used on the specified port or ports. The control mode may be one of the following:

forceunauthorized

The authenticator Port Access Entity (PAE) unconditionally sets the controlled port(s) to unauthorized mode

forceauthorized

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.

Default auto

Format **config dot1x port controlmode <port/listofports/all>**
<forceunauthorized/forceauthorized/auto>

config dot1x port initialize

Use this command to begin the initialization sequence on the specified port. This command is only valid if dot1x is enabled and the control mode for the specified port is “auto”.

Default disable

Format **config dot1x port initialize <port>**

config dot1x port maxrequests

Use this command to configure the maximum number of times the authenticator state machine on the specified port will retransmit an Extensible Authentication Protocol Over LANs (EAPOL) EAP

Request/Identity before timing out the supplicant. The maximum requests value must be in the range of 1 and 10.

Default 2

Format `config dot1x port maxrequests <port> <1-10>`

config dot1x port quietperiod

Use this command to configure the value, in seconds, of the timer used by the authenticator state machine on the specified 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 must be a value in the range of 0 and 65535.

Default 60

Format `config dot1x port quietperiod <port> <0-65535>`

config dot1x port reauthenabled

Use this command to enable or disable reauthentication of the supplicant for the specified port. The reauthenabled value must be true or false. If the value is true reauthentication will occur. Otherwise, reauthentication will not be allowed.

Default false

Format `config dot1x port reauthenabled <port> <true/false>`

config dot1x port reauthenticate

Use this command to begin the reauthentication sequence on the specified port. This command is only valid if dot1x is enabled and the control mode for the specified port is “auto”.

Default disable

Format `config dot1x port reauthenticate <port>`

config dot1x port reauthperiod

Use this command to configure 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 reauthperiod must be between 1 and 65535.

Default 3600

Format `config dot1x port reauthperiod <port> <1-65535>`

config dot1x port servertimeout

Use this command to configure the value, in seconds, of the timer used by the authenticator on the specified port to timeout the authentication server. The server timeout must be between 1 and 65535.

Default 30

Format `config dot1x port servertimeout <port> <1-65535>`

config dot1x port supptimeout

Use this command to configure the value, in seconds, of the timer used by the authenticator state machine on the specified port to timeout the supplicant. The supplicant timeout must be between 1 and 6553.

Default 30

Format `config dot1x port supptimeout <port> <1-65535>`

config dot1x port transmitperiod

Use this command to configure 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 must be a value in the range of 1 and 65535.

Default 30

Format `config dot1x port transmitperiod <port> <1-65535>`

config dot1x port users add

Use this command to add the specified user to the list of users with access to the specified port(s). The user must be a configured user and the port must be a valid port. By default, a user is given access to all ports.

Default all

Format `config dot1x port users add <user> <port/all>`

config dot1x port users remove

Use this command to remove the specified user from the list of users with access to the specified port(s).

Format `config dot1x port users remove <user> <port/all>`

show dot1x port detailed

Use this command to display the details of the IEEE 802.1X configuration parameters for the specified port.

Format `show dot1x port detailed <port>`

Port The interface whose configuration is displayed on this row.

Protocol Version

The version of IEEE 802.1X active on the port. Currently this is always 1.

PAE Capabilities

The port access entity state of the port. Either authenticator or supplicant.

Authenticator PAE State

The current state of the authenticator state machine. Possible values are initialize, disconnected, connecting, authenticating, authenticated, aborting, held, forceauthorized, and forceunauthorized.

Backend Authentication State

The current state of the back-end authentication state machine. Possible values are request, response, success, fail, timeout, idle, and initialize.

Quiet Period (secs)

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 value is expressed in seconds and will be in the range 0 and 65535.

Transmit Period (secs)

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 value is expressed in seconds and will be between 1 and 65535.

Supplicant Timeout (secs)

The timer used by the authenticator state machine on this port to timeout the supplicant. The value is expressed in seconds and will be between 1 and 65535.

Server Timeout (secs)

The timer used by the authenticator on this port to timeout the authentication server. The value is expressed in seconds and will be in the range of 1 and 65535.

Maximum Requests

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 value will be in the range of 1 and 10.

Reauthentication Period (secs)

The timer used by the authenticator state machine on this port to determine when reauthentication of the supplicant takes place. The value is expressed in seconds and will be between 1 and 65535.

Reauthentication Enabled

Indicates whether reauthentication is enabled for the port.

Key Transmission Enabled

Indicates whether a key is transmitted to the supplicant from the port.

Control Direction

Indicates the control direction for the port. Possible values are both and in.

show dot1x port stats

Use this command to display the IEEE 802.1X statistics for the specified port.

Format **show dot1x port stats <port>**

Port The interface whose statistics are displayed on this row.

EAPOL Frames Received

The number of valid Extensible Authentication Protocol over LANs (EAPOL) frames of any type that have been received by the authenticator port.

EAPOL Frames Transmitted

The number of valid EAPOL frames of any type that have been transmitted by the authenticator port.

EAPOL Start Frames Received

The number of EAPOL start frames that have been received by the authenticator port.

EAPOL Logoff Frames Received

The number of EAPOL logoff frames that have been received by the authenticator port.

Last EAPOL Frame Version

The protocol version number in the most recently received EAPOL frame.

Last EAPOL Frame Source

The source MAC address in the most recently received EAPOL frame.

EAP Response/ID Frames Received

The number of EAP response/identity frames that have been received by the authenticator port.

EAP Response Frames Received

The number of EAP response frames (other than response/identity frames) that have been received by the authenticator port.

EAP Request/ID Frames Transmitted

The number of EAP response/identity frames that have been transmitted by the authenticator port.

EAP Response Frames Transmitted

The number of EAP response frames (other than response/identity frames) that have been transmitted by the authenticator port.

Invalid EAPOL Frames Received

The number of EAPOL frames that have been received by the authenticator port with an unrecognized frame type.

EAP Length Error Frames Received

The number of EAPOL frames that have been received by the authenticator port with an incorrect length.

show dot1x port summary

Use this command to display a summary of the IEEE 802.1x configuration parameters for the specified port(s).

Format **show dot1x port summary** *<port/listofports/all>*

Port The interface whose configuration is displayed on this row.

Control Mode

The configured control mode: forceunauthorized, forceauthorized or auto.

Operating Control Mode

The active control mode.

Reauthentication Enabled

Indicates whether reauthentication is enabled for the port.

Transmission Enabled

Indicates whether a key is transmitted to the supplicant from the port.

Port Status Indicates whether a port is authorized.

show dot1x port users

Use this command to display IEEE 802.1X port security information about locally configured users.

Format **show dot1x port users** *<port>*

User The locally configured users with access to the specified port.

show dot1x summary

Use this command to display a summary of the IEEE 802.1X configuration parameters for the switch.

Format **show dot1x summary**

Administrative mode

Indicates whether authentication control is enabled on the switch.

Remote Authentication Dial-In User Service (RADIUS) commands

RADIUS accounting commands

config radius accounting mode

Use this command to enable or disable the RADIUS accounting function.

Default disable

Format **config radius accounting mode** *<enable/disable>*

config radius accounting server add

Use this command to configure the IP address to be used to access the accounting server. Only a single accounting server can be configured. If an accounting server is currently configured it must be removed using the config radius accounting server remove command before this command will succeed.

Format **config radius accounting server add** *<ipaddr>*

config radius accounting server port

Use this command to configure which User Datagram Protocol (UDP) port will be used to access the accounting server. The IP address specified must match that of the previously configured accounting server. If a port is already configured for the accounting server, the new port will replace the previously configured value.

Default 1813

Format **config radius accounting server port** *<ipaddr>* *<0-65535>*

config radius accounting server remove

Use this command to remove a configured accounting server. The IP address specified must match that of the previously configured accounting server. Since only a single accounting server is supported, issuing this command will cause future accounting attempts to fail.

Format **config radius accounting server remove** *<ipaddr>*

config radius accounting server secret

Use this command to configure the secret shared between the RADIUS client and accounting server. The IP address specified must match that of the previously configured accounting server. When you enter this command, you will be prompted to enter the secret, which must be an alphanumeric value of 20 characters or less.

Format **config radius accounting server secret** *<ipaddr>*

show radius accounting stats

Use this command to display the RADIUS statistics for the accounting server.

Format **show radius accounting stats** *<ipaddr>*

Accounting Server IP Address

The IP address of the server whose statistics are displayed on this row.

Round Trip Time

The time, in hundredths of a second, between the most recent RADIUS accounting response and the matching accounting request from this RADIUS accounting server.

Accounting Requests

The number of RADIUS accounting request packets sent to this accounting server, not including retransmissions.

Accounting Retransmissions

The number of RADIUS accounting request packets retransmitted to this accounting server.

Accounting Responses

The number of RADIUS packets received from this accounting server.

Malformed Accounting Responses

The number of malformed RADIUS accounting response packets received from this accounting server, including packets with invalid length but not including packets with bad authenticators or unknown types.

Bad Authenticators

The number of RADIUS accounting response packets received from this accounting server, including packets with invalid authenticators.

Pending Requests

The number of RADIUS accounting request packets sent to this accounting server that have not yet timed out or received a response.

Timeouts

The number of RADIUS packets sent to this accounting server that have timed out.

Unknown Types

The number of RADIUS packets of unknown type received from this accounting server.

Packets Dropped

The number of RADIUS packets received from this accounting server dropped for a reason not otherwise included in this list.

show radius accounting summary

Use this command to display a summary of the RADIUS accounting configuration parameters for the switch.

Format **show radius accounting summary**

Accounting Mode

Indicates whether accounting mode is enabled or disabled.

IP Address

The IP address of the RADIUS accounting server currently in use.

Port

The port used to access the accounting server.

Secret configured

Indicates whether a secret has been configured for the accounting server.

RADIUS configuration / summary commands**clear radius stats**

Use this command to reset all RADIUS statistics for the switch. You will be prompted to confirm this choice.

Format **clear radius stats**

config radius maxretransmit

Use this command to configure the maximum number of times a request packet is retransmitted when no response is received from the RADIUS server. The maxretransmit value is an integer in the range of 1 and 15.

Consideration should be given to the maximum delay time when configuring RADIUS maxretransmit and timeout values. If multiple RADIUS servers are configured, the maxretransmit

value on each will be exhausted before the next server is attempted. A retransmit will not occur until the configured timeout value on that server has passed without a response from the RADIUS server. Therefore, the maximum delay in receiving a response from the RADIUS application equals the sum of maxretransmit times timeout for all configured servers. If the RADIUS request was generated by a user login attempt, all user interfaces will be blocked until the RADIUS application returns a response.

Default 4

Format `config radius maxretransmit <1-15>`

config radius timeout

Use this command to configure the timeout value (in seconds) after which a request must be retransmitted to the radius server if no response is received.

Consideration should be given to the maximum delay time when configuring RADIUS maxretransmit and timeout values. If multiple RADIUS servers are configured, the maxretransmit value on each will be exhausted before the next server is attempted. A retransmit will not occur until the configured timeout value on that server has passed without a response from the RADIUS server. Therefore, the maximum delay in receiving a response from the RADIUS application equals the sum of maxretransmit times timeout for all configured servers. If the RADIUS request was generated by a user login attempt, all user interfaces will be blocked until the RADIUS application returns a response.

Default 5

Format `config radius timeout <1-30>`

show radius stats

Use this command to display RADIUS statistics for the switch that are not associated with a specific server or accounting server.

Format `show radius stats`

Invalid Server Address

The number of RADIUS access response packets received from an unknown address.

show radius summary

Use this command to display a summary of the RADIUS configuration parameters for the switch.

Format `show radius summary`

Current Server IP Address

The IP address of the RADIUS server currently used for authentication.

Number of Configured Servers

The number of RADIUS servers that have been configured.

Max Number of Retransmits

The maximum number of times a request packet will be retransmitted.

Timeout Duration (secs)

The timeout value, in seconds, for request retransmissions.

Accounting Mode

Indicates whether accounting is currently enabled.

RADIUS server commands

config radius server add

Use this command to configure the IP address used to connect to a RADIUS server. Up to three servers can be configured for each RADIUS client. If three servers are currently configured, one must be removed using the **config radius server remove** command before the add command will succeed. Once a server has been added it will be identified in future commands by its IP address.

Format **config radius server add <ipaddr>**

config radius server msgauth

Use this command to enable or disable the message authenticator attribute for the specified RADIUS server. Enabling the message authenticator attribute provides additional security for the connection between the RADIUS client and server. Some RADIUS servers require that the message authenticator attribute be enabled before authentication requests from the RADIUS client will be accepted. The IP address specified must match that of a configured server.

Format **config radius server msgauth <ipaddr> <enable/disable>**

config radius server port

Use this command to configure which UDP port will be used to access the specified RADIUS server. The IP address specified must match that of the previously configured RADIUS server.

Default 1812

Format **config radius server port <ipaddr> <0-65535>**

config radius server primary

Use this command to specify which configured server should be the primary server for this RADIUS client. The primary is the server that is used by default for handling RADIUS requests. The remaining configured servers are used only if the primary server cannot be reached. A maximum of three servers can be configured on each client. Only one server can be configured as the primary server. If a primary server is currently configured and this command is issued, the server specified by the IP address used in this command will become the new primary server. The IP address specified must match that of a configured server.

Format **config radius server primary <ipaddr>**

config radius server remove

Use this command to remove a configured RADIUS server. The IP address specified must match that of the previously configured RADIUS server. When a server is removed all configuration for the server is erased including the shared secret. If the removed server was the primary server, one of the remaining configured servers will be used as the RADIUS server for future RADIUS requests.

Format **config radius server remove <ipaddr>**

config radius server secret

Use this command to configure the secret shared between the RADIUS client and server. A secret must be configured for each RADIUS server. The IP address specified must match that of a previously configured RADIUS server. When you enter this command, you will be prompted to enter the secret, which must be an alphanumeric value of 20 characters or less.

Format **config radius server secret <ipaddr>**

show radius server stats

Use this command to display the statistics for a configured RADIUS server.

Format **show radius server stats <ipaddr>**

Server IP Address

The IP address of the server whose information is displayed on this row.

Round Trip Time

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 to this server.

Packets Dropped

The number of RADIUS packets received from this server dropped for a reason not otherwise included in this list.

show radius server summary

Use this command to display a summary of the configured RADIUS servers.

Format **show radius server summary**

Current Indicates the server currently in use for authentication.

IP Address The IP address of the authentication server.

Port The port used to access the authentication server.

Type Indicates whether the server is primary or secondary.

Secret configured

Indicates whether a secret has been configured for the authentication server.

Secure Shell (SSH) commands

config ssh adminmode

Use this command to enable or disable SSH.

Default Disabled

Format `config ssh adminmode <enable/disable>`

config ssh protocol

Use this command to set or remove protocol levels (or versions) for SSH. Either SSH1 (1), SSH2 (2) or both (1 and 2) can be set.

Default both

Format `config ssh protocol <ssh1/ssh2/both>`

show ssh info

Displays the SSH settings.

Format `show ssh info`

Administrative Mode

Indicates whether the administrative mode of SSH is enabled or disabled.

Protocol Level

The protocol level may have the values of version 1, version 2 or both versions 1 and 2.

Connections Specifies the current SSH connections.

Secure Socket Layer (SSL) commands

config http secureport

Use this command to configure the SSL port where port is between 1 and 65535.

Default 443

Format `config http secureport <port>`

config http secureprotocol

Use this command to enable or disable SSL and set protocol levels (versions). The protocol level can be set to TLS1, SSL3 or to both TLS1 and SSL3.

Default both

Format `config ip http secure-protocol <ssl3/tls1/both> <add/remove>`

config http secureserver adminmode

Command is used to enable/disable the SSL for secure HTTP.

Default disable

Format `config http secureserver adminmode <enable/disable>`

show http info

Displays the http settings for the switch.

Format `show http info`

Mode	Privileged EXEC
Secure-Server Administrative Mode	Indicates whether the administrative mode of secure HTTP is enabled or disabled.
Secure Protocol Level	The protocol level may have the values of SSL3, TSL1 or both.
Secure Port	Specifies the port configured for SSL.
HTTP Mode	Indicates whether the HTTP mode is enabled or disabled.

Quality of Service (QoS) commands

This section describes the commands used to configure and manage the Quality of Service (QoS) features of the Intel® Blade Server Ethernet Switch Module IXM5414E. These features include:

- Access Control Lists (ACLs)
- Bandwidth provisioning

Access Control List (ACL) commands

An 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. This section describes the commands you use to specify the interfaces to which an ACL applies, whether it applies to inbound or outbound traffic and its match criteria.

config acl create

Use this command to create an ACL identified by the parameter <aclid>. The ACL number is an integer from 1 to 100.

Format **config acl create <aclid>**

config acl delete

Use this command to delete an ACL identified by the parameter <aclid> from the system.

Format **config acl delete <aclid>**

config acl interface add

Use this command to associate an ACL with an interface and specifies whether it affects inbound or outbound traffic. The <direction> parameter can have the values of in or out. The <aclid> parameter specifies the ACL to add.

Format **config acl interface add <port> <direction> <aclid>**

config acl interface remove

Use this command to disassociate an ACL from an interface for the specified direction. The <direction> parameter can have the values of in or out. The <aclid> parameter specifies the ACL to remove.

Format **config acl interface remove <port> <direction> <aclid>**

config acl rule action

Use this command to specify the action for the ACL and rule referenced by the parameters <aclid> and <rulenum>. The values of permit or deny indicate how this rule is applied.

Format **config acl rule action <aclid> <rulenum> <permit/deny>**

config acl rule create

Use this command to create a rule within the ACL referenced by the parameter <aclid>. The rule is identified by the <rulenum> parameter. An ACL may have up to 10 user-specified rules, whose <rulenum> ranges from 1 to 10. Rules are created with a default action of deny.

Default deny

Format **config acl rule create <aclid> <rulenum>**

config acl rule delete

Use this command to remove a rule from the ACL referenced by the parameter <aclid>. The rule is identified by the <rulenum> parameter.

Format **config acl rule delete <aclid> <rulenum>**

config acl rule match dstip

Use this command to specify a destination IP address and mask match condition for the ACL rule referenced by the <aclid> and <rulenum> parameters. The <ipaddr> and <ipmask> parameters are 4-digit dotted-decimal numbers which represent the destination IP address and IP mask, respectively.

Format **config acl rule match dstip <aclid> <rulenum> <ipaddr> <ipmask>**

config acl rule match dstl4port keyword

Use this command to specify a destination layer 4 port match condition for the ACL rule referenced by the <aclid> and <rulenum> parameters. The <portkey> parameter uses a single keyword notation and currently has the values of 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 a port range.

This command and the **config acl match destl4port number** command are two methods of specifying the destination layer 4 port range as a match condition. Either command can be used to configure or modify the destination layer 4 port range.

Format **config acl rule match dstl4port keyword <aclid> <rulenum> <portkey>**

config acl rule match dstl4port number

Use this command to specify a destination layer 4 port match condition for the ACL rule referenced by the <aclid> and <rulenum> parameters. The <startport> and <endport> parameters identify the first and last ports that are part of the port range. They have values from 0 to 65535. The ending port must have a value equal to or greater than the starting port. The starting port, ending port and all ports in between will be part of the destination port range.

Either this command or the **config acl match destl4port keyword** command may be used to specify a destination layer 4 port range as a match condition.

Format **config acl rule match dstl4port range <aclid> <rulenum> <startport> <endport>**

config acl rule match every

Use this command to specify a match condition in which all packets will be considered to match the ACL rule referenced by the <aclid> and <rulenum> parameter. If the parameter <true/false> is set to <true>, all packets will be either permitted or denied based on the action setting for the rule and no other match conditions may be specified. Specifying <false> allows other match conditions to be specified.

Format **config acl rule match every <aclid> <rulenum> <true/false>**

config acl rule match protocol keyword

Use this command to specify the IP protocol of a packet as a match condition for the ACL rule referenced by the <aclid> and <rulenum> parameters. The <protocolkey> parameter identifies the protocol using a single keyword notation and has the possible values of ICMP, IGMP, IP, TCP and UDP. A protocol keyword of ip is interpreted to match all protocol number values.

Either this command or the **config acl match protocol number** command can be used to specify an IP protocol value as a match criterion.

Format **config acl rule match protocol keyword <aclid> <rulenum> <protocolkey>**

config acl rule match protocol number

Use this command to specify a protocol number as a match condition for the ACL rule referenced by the <aclid> and <rulenum> parameters. The <protocolnum> parameter identifies the protocol by number. The protocol number is a standard value assigned by IANA and is an integer from 0 to 255.

Either this command or the **config acl match protocol keyword** command can be used to specify an IP protocol value as a match criterion.

Format **config acl rule match protocol number <aclid> <rulenum> <protocolnum>
<protocolmask>**

config acl rule match srcip

Use this command to specify a packet's source IP address and Mask as a match condition for the ACL rule referenced by the <aclid> and <rulenum> parameters. The <ipaddr> and <ipmask> parameters are 4-digit dotted-decimal numbers which represent the source IP address and IP mask, respectively.

Format **config acl rule match srcip <aclid> <rulenum> <ipaddr> <ipmask>**

config acl rule match srcl4port keyword

Use this command to specify a source layer 4 port match condition for the ACL rule referenced by the <aclid> and <rulenum> parameters. The <portkey> uses a single keyword notation and has the possible values of 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.

This command and the **config acl match srcl4port number** command are two methods of specifying the source layer 4 port range as a match condition. Either command can be used to configure or modify the source layer 4 port range.

Format **config acl rule match srcl4port keyword <aclid> <rulenum> <portkey>**

config acl rule match srcl4port number

Use this command to specify a packet's source layer 4 port match condition for the ACL rule referenced by the <aclid> and <rulenum> parameters. The <startport> and <endport> parameters identify the first and last ports that are part of the port range and have values from 0 to 65535. The ending port must have a value equal or greater than the starting port. The starting port, ending port and all ports in between will be part of the contiguous source port range.

Either this command or **config acl match srcl4port keyword** can be used to specify a source layer 4 port range as a match criterion.

Format **config acl rule match srcl4port range <aclid> <rulenum> <startport> <endport>**

show acl detailed

Use this command to display an ACL and all of the rules that are defined for the ACL. The <aclid> is the number used to identify the ACL.

Format **show acl detailed <aclid>**

Rule Number

Displays the number identifier for each rule that is defined for the ACL.

Action Displays the action that will be taken if a packet matches the rule's criteria. The choices are permit or deny.

Protocol Displays which IP protocol (if any) is a match condition for the rule. The possible values are ICMP, IGMP, IP, TCP, and UDP.

Source IP Address

Displays the source IP address (if any) that is a match condition for this rule.

Source IP Mask

Displays the source IP mask (if any) that is a match condition for this rule.

Source Ports Displays the source port range (if any) that is a match condition for this rule.

Service Type Field Match

Indicates whether an IP DSCP, IP Precedence or IP TOS match condition is specified for this rule.

Service Type Field Value

Indicates the value specified for the Service Type Field Match (IP DSCP, IP Precedence or IP TOS) if it a match condition for this rule.

show acl summary

Use this command to display a summary of the ACLs associated with interfaces in the system.

Format **show acl summary**

ACL ID Displays the ACL identifier.

Rules Displays the number of rules that are associated with this ACL.

Interface(s) Displays the interfaces associated with this ACL.

Direction Displays the packet filtering direction for the ACL on the interface. The possible values displayed are inbound and outbound.

Bandwidth provisioning commands

BW provisioning BW allocation commands

config bwprovisioning bwallocation create

Use this command to create a bandwidth allocation profile. The <name> field is an alphanumeric string up to 15 characters.

Format **config bwprovisioning bwallocation create <name>**

config bwprovisioning bwallocation delete

Use this command to delete a bandwidth allocation profile from the system. The <name> field is the user supplied name associated with the bandwidth allocation profile. A bandwidth allocation profile may not be deleted while it is associated with a traffic class.

Format **config bwprovisioning bwallocation delete <name>**

config bwprovisioning bwallocation maxbandwidth

This commands configures the maximum allowable bandwidth for this bandwidth allocation profile. The <maxbandwidth> parameter is a value from 0 to the maximum bandwidth of the interface to be associated with this profile. The bandwidth allocation profile maximum bandwidth must be greater than or equal to the minimum bandwidth. If this value is set to 0, it will not allow any traffic for this bandwidth allocation profile.

Default 100 Mbps

Format **config bwprovisioning bwallocation maxbandwidth <name> <maxbandwidth>**

show bwprovisioning bwallocation detailed

Use this command to display detailed bandwidth allocation information for the specified bandwidth allocation profile.

Format **show bwprovisioning bwallocation detailed <name>**

Bandwidth Allocation Profile Name

Displays the user-defined name of this bandwidth allocation profile.

Minimum Bandwidth

Displays the minimum guaranteed bandwidth of this bandwidth allocation profile in Mbps.

Maximum Bandwidth

Displays the maximum allowable bandwidth of this bandwidth allocation profile in Mbps.

Associated Traffic Class(es)

Displays the traffic classes that have been associated with this bandwidth allocation profile. This field is blank if there are no traffic classes associated with this bandwidth allocation profile.

show bwprovisioning bwallocation summary

Use this command to display the bandwidth allocation information for all bandwidth allocation profiles in the system.

Format **show bwprovisioning bwallocation summary**

Bandwidth Allocation Profile Name

Displays the user-defined name of this bandwidth allocation profile.

Minimum Bandwidth

Displays the minimum guaranteed bandwidth of this bandwidth allocation profile in Mbps.

Maximum Bandwidth

Displays the maximum allowable bandwidth of this bandwidth allocation profile in Mbps.

BW provisioning traffic class commands**config bwprovisioning trafficclass bwallocation**

Use this command to associate a bandwidth allocation profile with a traffic class. The <bwprofile> parameter must represent a valid bandwidth allocation profile.

Format **config bwprovisioning trafficclass bwallocation <name> <bwprofile>**

config bwprovisioning trafficclass create

Use this command to create a traffic class. The <type> field indicates the type of traffic class. The only supported value for type is vlan. The <name> field is an alphanumeric string up to 15 characters.

Format **config bwprovisioning trafficclass create <type> <name>**

config bwprovisioning trafficclass delete

Use this command to delete a traffic class from the system. The <name> field identifies the traffic class to be deleted. When a traffic class is deleted, its association with a bandwidth allocation profile is automatically removed.

Format **config bwprovisioning trafficclass delete <name>**

config bwprovisioning trafficclass port

Use this command to attach a traffic class to a specific interface. The <port> interface must indicate a valid physical or logical interface. The sum of the minimum bandwidth allocations 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 attached to the same port. 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.

Format **config bwprovisioning trafficclass port <name> <port>**

config bwprovisioning trafficclass vlan

Use this command to associate a VLAN with a traffic class. The <vlanid> field is the VLAN ID for the traffic class within the range of 1 to 4094.

The VLAN parameter can identify an invalid VLAN (the VLAN does not need to exist in the system.)

Format **config bwprovisioning trafficclass vlan <name> <vlanid>**

config bwprovisioning trafficclass weight

Use this command to configure the priority for this traffic class. The <weight> parameter will be a value between 1 and 1024.

Default 1

Format **config bwprovisioning trafficclass weight <name> <weight>**

show bwprovisioning trafficclass allocatedbw

Use this command to display the bandwidth allocated. The allocated minimum bandwidth should not exceed the interface bandwidth unless the interface is a LAG interface.

Format **show bwprovisioning trafficclass allocatedbw <port>**

Port The specified interface.

Allocated Minimum Bandwidth

Displays the sum of the minimum guaranteed bandwidth for all traffic classes configured on this interface.

Allocated Maximum Bandwidth

Displays the sum of the maximum allowable bandwidth for all traffic classes configured on this interface.

show bwprovisioning trafficclass detailed

Use this command to display the traffic class information for the specified traffic class.

Format **show bwprovisioning trafficclass detailed <name>**

Traffic Class Name

Displays the name of this traffic class.

Port Displays the port to which this traffic class is attached.

VLAN ID Displays the VLAN ID with which this traffic class is associated.

Weight Displays the weight of this traffic class.

Accept Byte Count

Displays the number of bytes accepted.

Bandwidth Allocation Profile

Displays the bandwidth allocation profile associated with this traffic class. This field is blank when there is no bandwidth allocation profile associated with this traffic class.

The following attributes are only displayed when there is a bandwidth allocation profile associated with this traffic class.

Minimum Bandwidth

Displays the minimum bandwidth defined for this traffic class.

Maximum Bandwidth

Displays the maximum bandwidth defined for this traffic class.

show bwprovisioning trafficclass summary

Use this command to display the traffic class information for all traffic classes in the system.

Format **show bwprovisioning trafficclass summary**

Traffic Class Name

Displays the user-defined name of this traffic class.

Port Displays the interface to which this traffic class is attached.

VLAN ID Displays the Virtual Local Area Network (VLAN) ID with which this traffic class is associated.

Weight Displays the weight of this traffic class.

Bandwidth Allocation Profile

Displays the bandwidth allocation profile associated with this traffic class. This field is blank when there is no bandwidth allocation profile associated with this traffic class.

Appendix A RJ-45 Pin Specifications

The four external Ethernet ports of this switch module are auto-configuring and will work with straight-through or crossover cables when connected to other Ethernet equipment. Review the documentation that comes with the product you are connecting to for matching cable pin assignments.

The following illustration and table show the standard RJ-45 receptacle/connector and their corresponding pin assignments.

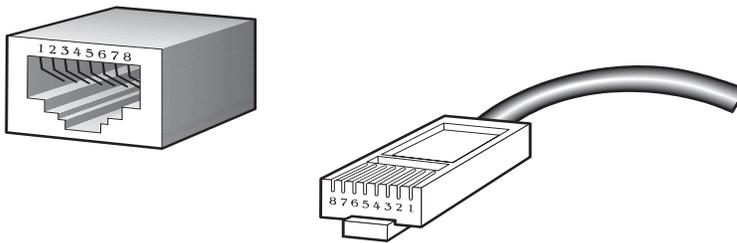


Table 7. Standard Ethernet cable, RJ-45 pin assignment

Contact (pin number)	Label	Media direct interface signal
1	TPO+	Tx + (transmit)
2	TPO-	Tx - (transmit)
3	TP1+	Rx + (receive)
4	TP2+	Not used
5	TP2-	Not used
6	TP1-	Rx - (receive)
7	TP3+	Not used
8	TP3-	Not used

Appendix B Cable Lengths

Use the following table as a guide for the maximum cable lengths:

Table 8. Maximum cable lengths

Standard	Data transmission rate	Media type	Maximum distance
1000BASE-T	1000 Mbps	Category 5e UTP cable	100 meters (328.1 ft)
		Category 5 UTP cable	
100BASE-TX	100 Mbps	Category 5 UTP cable	100 meters (328.1 ft)
10BASE-T	10 Mbps	Category 3 UTP cable	100 meters (328.1 ft)

Appendix C Run-time Switching Software Default Settings

The following table contains the default settings for the run-time switching software variables. Variables are separated by category and further by sub-headings (listed alphabetically within category). "Default value" is self-explanatory while "Command" lists the CLI command used to change the default setting.

Table 9. Default settings for run-time switching software variables

Heading	Sub-heading	Variable	Default value	Command	
Quality of Service					
	ACL				
		ACL Rule	None	config acl rule create	
	Bandwidth Provisioning				
		Bandwidth Allocation Maximum	100 mbps	config bwprovisioning bwallocation maximum	
		Traffic Class Weight	1	config bwprovisioning trafficclass weight	
Security					
	IEEE 802.1X				
			Add users	All	config dot1x port users add
			Control Mode	Auto	config dot1x port controlmode
			Initialization	Disable	config dot1x port initialize
			Maximum # of requests	2	config dot1x port maxrequests
			Mode	Disable	config dot1x adminmode
			Port initialize	Disable	config dot1x port initialize
			Quiet Period	60 seconds	config dot1x port quietperiod
			Reauthentication Enabled	False	config dot1x port reauthenabled
			Reauthentication Period	3600 seconds	config dot1x port reauthperiod

Table 9. Default settings for run-time switching software variables (continued)

Heading	Sub-heading	Variable	Default value	Command
		Reauthentication Sequence	Disable	config dot1x port reauthenticate
		Server Timeout	30 seconds	config dot1x port servertimeout
		Supplicant Time Out	30 seconds	config dot1x port supptimeout
		Transmit Period	30 seconds	config dot1x port transmitperiod
	Remote Authentication Dial-in User Service (RADIUS)			
	<i>Accounting</i>			
		Accounting Server Port	1813	config radius accounting server port
		Mode	Disable	config radius accounting mode
	<i>Configuration</i>			
		Maximum Retransmits	4	config radius maxretransmits
		Timeout	5 minutes	config radius timeout
	<i>Server</i>			
		Server Port	1812	config radius accounting server port
	Secure Shell (SSH)			
		Mode	Disable	config ssh adminmode
		Protocol	Both (SSH1 and SSH2)	config ssh protocol
	Secure Socket Layer (SSL)			
		Secure port	443	config http secureport
		Secure Protocol	Both (SSL3 and TLS1)	config http secureprotocol
		Secure Server Mode	Disable	config http secureserver adminmode
Switching				
	VLAN Switching			

Table 9. Default settings for run-time switching software variables (continued)

Heading	Sub-heading	Variable	Default value	Command
		Accept frame	all	config vlan port acceptframe
		Broadcast Storm	disable	config vlan bcaststorm
		Default port VID	1	config vlan port pvid
		Multicast Storm	disable	config vlan mcaststorm
		Name	VLAN1 = Default	config vlan name
		Port priority	0	config vlan port priority
	GARP			
		GARP administration	disable	config garp gmrp adminmode
		GARP interface	disable	config garp gmrp interfacemode
		GARP join timer	20 centiseconds	config garp jointimer
		GARP leave all timer	1000 centiseconds	config garp leavealltimer
		GARP leave timer	60 centiseconds	config garp leavetimer
	<i>GVRP</i>			
		GVRP administration	disable	config garp gvrp adminmode
		GVRP interface	disable	config gvrp gmrp interfacemode
		GVRP join timer	20 centiseconds	config gvrp jointimer
		GVRP leave all timer	1000 centiseconds	config gvrp leavealltimer
		GVRP leave timer	60 centiseconds	config gvrp leavetimer
	IGMP Snooping			
		Group Membership Interval	260 seconds	config igmpsnooping groupmembershipinterval
		Interface	disable	config igmpsnooping interfacemode
		Maximum response time	10 seconds	config igmpsnooping maxresponse
		MCRT Expiration Time	0 seconds	config igmpsnooping mcrtexpiretime
		Mode	Disable	config igmpsnooping adminmode

Table 9. Default settings for run-time switching software variables (continued)

Heading	Sub-heading	Variable	Default value	Command
	Link Aggregation			
		LAG linktrap	enable	config lag linktrap
	Spanning Tree Protocol (STP)			
	<i>Bridge</i>			
		Forward Delay	15 secs	config spanningtree bridge forwarddelay
		Hello Time	2 secs	config spanningtree bridge hellotime
		Max Age	6 secs	config spanningtree bridge maxage
		Priority	32768	config spanningtree bridge priority
	<i>Configuration</i>			
		Admin Mode	Disable	config spanningtree adminmode
		Configuration name	The base MAC address displayed using hexadecimal notation	config spanningtree configuration name
		Forced Version	IEEE 802.1D	config spanningtree forceversion
		Revision level	0	config spanningtree configuration revision
	<i>CST</i>			
		Edgeport	False	config spanningtree cst port edgeport
		Pathcost	Auto	config spanningtree cst port pathcost
		Priority	128	config spanningtree cst port priority
	<i>Port</i>			
		Migration Check	Disable	config spanningtree port migrationcheck
		Port Mode	Disable	config spanningtree port mode
System				
		Auto log-out	10 min	

Table 9. Default settings for run-time switching software variables (continued)

Heading	Sub-heading	Variable	Default value	Command
		Configuration update	Disable	
		Default gateway	0.0.0.0	
		IP address	10.90.90.9x, where x depends on the number of the bay into which you have installed the switch module.	
		Subnet mask	255.255.255.0	
	Configuration			
		System Contact	Blank	config syscontact
		System Location	Blank	config syslocation
		System Name	Blank	config sysname
	Forwarding Database			
		Forwarding Database aging time	300 seconds	config forwardingdb agetime
	Port			
	<i>Configuration</i>			
		Auto Negotiation	Enable	config port autoneg
		Flow control	Disable	config port flowcontrol
		LACP mode	Disable	config port lacpmode
		Port Enable	Enable	config port adminmode
	<i>Mirroring</i>			
		Mirroring Mode	Disable	config mirroring mode
	Network Connectivity			
		IP connectivity parameters	10.90.90.9x 255.255.255.0 0.0.0.0	config network parms
		IP connectivity protocol	None	config network protocol
		Java enable status	Enable	config network javamode
		Web enable status	Enable	config network webmode
	SNMPcommunity			

Table 9. Default settings for run-time switching software variables (continued)

Heading	Sub-heading	Variable	Default value	Command
		IP address	0.0.0.0	config snmpcommunity ipaddr
		IP Mask	0.0.0.0	config snmpcommunity ipmask
		Mode	Default private and public communities are enabled by default. The four undefined communities are disabled by default	config snmpcommunity mode
		Type	Public/Private	config snmp community create
	Telnet			
		Max Number of Sessions	5	config telnet maxsessions
		Status	Enable	config telnet mode
		Ttimeout	5	config telnet timeout
	User Accounts			
		Password	Blank	config users passwd
		SNMPv3 Access Mode	R/W for admin, ReadOnly for others	
		SNMPv3 Authentication	No authorization	config users snmpv3 authentication
		SNMPv3 Encryption	No encryption	config users snmpv3 encryption
	Utilities			
	<i>Transfer</i>			
		Transfer Upload/Download Datatype	Code	transfer upload/download datatype
		Transfer upload/download Filename	Blank	transfer upload/download filename
		Transfer Upload\Download IP Address	0.0.0.0	transfer upload/download serverip
		Transfer Upload/download Path	Blank	transfer upload/download path
	Trap Management			

Table 9. Default settings for run-time switching software variables (continued)

Heading	Sub-heading	Variable	Default value	Command
		Authenticate Trapflags	Enable	config trapflags authentication
		Trapflags Linkmode	Enable	config trapflags linkmode
		Trapflags Multiusers	Enable	config trapflags multiusers
		Trapflags STP	Enable	config trapflags stpmode

Appendix D CLI Command Tree

This appendix presents the CLI command tree used in conjunction with the Intel® Blade Server Ethernet Switch Module IXM5414E.

SWITCHING					
clear					
	config				
	igmpsnooping				
	lag				
	dot1x	port	stats		
	pass				
	radius	stats			
	stats	port			
		switch			
	transfer				
	traplog				
	vlan				
config					
	acl	create			
		delete			
		interface	add		
			remove		
		rule	action		
			create		
			delete		
			match	dstip	
				dstl4port	keyword
					number
				every	
				protocol	keyword
					number

				srcip		
				src14port	keyword	
					number	
	authentication	login	create			
			delete			
			set			
	bwprovisioning	bwallocation	create			
			delete			
			maxbandwidth			
		trafficclass	bwallocation			
			create			
			delete			
			port			
			vlan			
			weight			
	classofservice	802.1mapping				
	dot1x	adminmode				
		defaultlogin				
		login				
		port	controlmode			
			initialize			
			maxrequests			
			quietperiod			
			reauthenable			
			reauthenticate			
			reauthperiod			
			servertimeout			
			supptimeout			
			transmitperiod			
			users	add		
				remove		

	forwardingdb	agetime			
	garp	gmrp	adminmode		
			interfacemode		
		gvrp	adminmode		
			interfacemode		
		jointimer			
		leavealltimer			
		leavetimer			
	http	secureport			
		secureprotocol			
		secureserver	adminmode		
	igmpsnooping	adminmode			
		groupmember- shipinterval			
		interfacemode			
		maxresponse			
		mcrtextpiretime			
	lag	addport			
		adminmode			
		create			
		deletelag			
		deleteport			
		linktrap			
		name			
	loginsession	close			
	macfilter	adddest			
		create			
		deldest			
		remove			
	mirroring	create			
		delete			
		mode			

	network	javamode				
		webmode				
	port	adminmode				
		autoneg				
		flowcontrol				
		lacpmode				
		linktrap				
		physicalmode				
	prompt					
	protocol	create				
		delete				
		interface	add			
			remove			
		protocol	add			
			remove			
		vlan	add			
			remove			
	radius	accounting	mode			
			server	add		
				port		
				remove		
				secret		
		maxretransmit				
		server	add			
			msgauth			
			port			
			primary			
			remove			
			secret			
		timeout				
	snmpcommunity	accessmode				

		create				
		delete				
		ipaddr				
		ipmask				
		mode				
	snmptrap	create				
		delete				
		ipaddr				
		mode				
	spanningtree	adminmode				
		bridge	forwarddelay			
			hellotime			
			maxage			
			priority			
		cst	port	edgeport		
				pathcost		
				priority		
		forceversion				
		port	migrationcheck			
			mode			
	ssh	adminmode				
		protocol				
	syscontact					
	syslocation					
	sysname					
	telnet	maxsessions				
		mode				
		timeout				
	trapflags	authentication				
		linkmode				
		multiusers				

		stpmode				
	users	add				
		defaultlogin				
		delete				
		login				
		passwd				
		snmpv3	accessmode			
			authentication			
			encryption			
	vlan	bcaststorm				
		create				
		delete				
		makestatic				
		mcaststorm				
		name				
		participation				
		port	acceptframe			
			priority			
			pvid			
			tagging			
help						
logout						
ping						
reset	system					
save	config					
show	acl	detailed				
		summary				
	arp	switch				
	authentication	login	info			
			users			
	bwprovisionin g	bwallocation	detailed			

			summary			
		trafficclass	allocatedbw			
			detailed			
			summary			
	classofservice	802.1pmapping				
	dot1x	port	detailed			
			stats			
			summary			
			user			
		summary				
	eventlog					
	forwardingdb	agetime				
		learned				
		table				
	garp	info				
		interface				
	history					
	http	info				
	igmpsnooping					
	inventory					
	lag					
	loginsession					
	macfilter					
	mfdb	gmrp				
		igmpsnooping				
		staticfiltering				
		stats				
		table				
	mirroring					
	msglog					
	network					

	port					
	protocol					
	radius	accounting	stats			
			summary			
		server	stats			
			summary			
		stats				
		summary				
	snmpcommunity					
	snmptrap					
	spanningtree	bridge				
		cst	detailed			
			port	detailed		
				summary		
		port				
		summary				
	ssh	info				
	stats	port	detailed			
			summary			
		switch	detailed			
			summary			
	sysinfo					
	telnet					
	trapflags					
	traplog					
	users	authentication				
		info				
	vlan	detailed				
		port				
		summary				
transfer						

	download	datatype				
		filename				
		path				
		serverip				
		start				
	upload	datatype				
		filename				
		path				
		serverip				
		start				

Appendix E CLI Configuration Examples

This appendix provides examples of using the CLI to configure the Intel® Blade Server Ethernet Switch Module IXM5414E for some key functions.

Bridging configuration example

This section provides sample CLI commands showing how to configure the Intel® Blade Server Ethernet Switch Module IXM5414E for basic bridging support. Bridging support, conforming to the IEEE 802.1D compatibility mode specified in IEEE 802.1s, is enabled for the switch and for all ports by default. All ports are enabled by default, and defaults are also provided for timers and protocol parameters.

Although the switch will operate correctly as a bridge implementing the base Spanning Tree Protocol (STP) as configured at the factory, the configuration script in this section will show you how to override the defaults. Before you do so, make sure that you fully understand the protocol and that the values you provide are consistent with each other.

Set a new bridge priority level. Setting the priority level affects the likelihood of the bridge being elected as the root of the spanning tree (the lower the number the greater the probability). It is the only way to change the bridge identifier, which consists of the bridge priority concatenated with the switch's base MAC address. The default value is 32768. If all bridges retain their default priority values, the bridge with the lowest MAC address will become the root bridge.

```
config spanningtree bridge priority 7680
```

Set new port priority levels. Setting the priority level affects the likelihood of the port being elected as the root port of the spanning tree (the lower the number the greater the probability). It is the only way to change the port identifier, which consists of the port priority concatenated with the port's interface number. The default value is 128.

```
config spanningtree port priority ext.1 16
```

```
config spanningtree port priority ext.2 32
```

Set new timer values. The timer values will only take effect if the bridge becomes the root bridge, in which case they will take effect for all bridges in the network.

```
config spanningtree bridge maxage 30
```

```
config spanningtree bridge forwarddelay 16
```

```
config spanningtree bridge hellotime 14
```

Assign new path cost values to the ports whose priority values were changed. The lower the path cost the more likely that a port will be elected as the root port.

```
config spanningtree port pathcost ext.1 8
```

```
config spanningtree port pathcost ext.2 16
```

In addition to the parameters that affect the Spanning Tree Protocol, other parameters and protocols are defined in IEEE 802.1D which you may also change. For example, IEEE 802.1p has been included in the latest version of 802.1D. Use the following commands to change the

default priority mapping provided by the switch. These commands affect all of the interfaces on the switch and leave the defaults unchanged for priority levels 3-7.

```
config classofservice 802.1p mapping 0 0
```

```
config classofservice 802.1p mapping 1 2
```

```
config classofservice 802.1p mapping 2 1
```

The switch supports two protocols based on the Generic Attribute Registration Protocol (GARP) defined in IEEE 802.1D: GARP Multicast Registration Protocol (GMRP) and GARP VLAN Registration Protocol (GVRP). These protocols are disabled by default.

```
config garp gmrp adminmode enable
```

```
config garp gmrp interfacemode all
```

```
config garp gvrp adminmode enable
```

```
config garp gvrp interfacemode all
```

While the Spanning Tree Protocol is needed to maintain the network topology, forwarding of frames also requires that the switch learn the location of end stations. The switch does this by recording the port on which packets from a source MAC address are received. The forwarding database is used to hold this information. You can control how long an address will remain in the database if no traffic is seen from it (the aging timer).

```
config forwardingdb agetime 500
```

IEEE 802.1w configuration example

This section shows you how to configure the Intel® Blade Server Ethernet Switch Module IXM5414E to support rapid reconfiguration of the spanning tree topology. The IEEE 802.1w support specified in IEEE 802.1s defines a new configuration algorithm and protocol that provide significantly faster reconfiguration of the spanning tree than the original algorithm and protocol defined in the base IEEE 802.1D standard. While the old and new protocols will successfully interoperate, the IEEE 802.1 standards committee recommends the use of the new protocol.

Configuration of the switch to support IEEE 802.1w is simple. In normal operation, the bridge timers are not used to control reconfiguration, and the default values should be adequate. Bridge and port priorities and path costs are still required, and are configured as shown for IEEE 802.1D.

Configure the switch to use rapid reconfiguration.

```
config spanningtree forceversion 802.1w
```

To disable support for rapid reconfiguration.

```
config spanningtree forceversion 802.1d
```

VLAN configuration example

This section provides sample CLI commands showing how to configure the Intel® Blade Server Ethernet Switch Module IXM5414E to support IEEE 802.1Q VLANs. Configuring VLANs allows you to partition your network on a logical rather than physical basis. The only physical restriction is that both ends of a point-to-point link must be in the same VLAN. There are many possible logical partitions – one common one being department membership.

The script in the following example shows you how to create and configure VLANs on your switch.

Create and name two VLANs (the names are optional).

```
config vlan create 1  
config vlan name 1 vlan_one  
config vlan create 2  
config vlan name 2 vlan_two
```

Assign the ports that will belong to `vlan_one`. This will be a tagged VLAN – only tagged packets will be accepted by member ports, and all packets transmitted from member ports will be tagged.

```
config vlan participation include 1 bay.1,bay.2  
config vlan port tagging enable 1 bay.1,bay.2  
config vlan port acceptframe vlanonly 1 bay.1,bay.2
```

Assign the ports that will belong to `vlan_two`. Untagged packets will be accepted by member ports `bay.3` and `bay.4` and assigned the default PVID of 2, and all packets transmitted from member ports will be untagged. Note that `bay.2` is a member of both `vlan_one` and `vlan_two`, and that `ext.1` and `ext.2` will never be members.

```
config vlan participation include 2 bay.2,bay.3,bay.4  
config vlan participation exclude 2 ext.1,ext.2  
config vlan port acceptframe all 2 bay.3,bay.4
```

Assign the same default PVID to ports `bay.3` and `bay.4`.

```
config vlan port pvid 2 bay.3,bay.4
```

Link aggregation configuration example

This section provides sample CLI commands showing how to configure the Intel® Blade Server Ethernet Switch Module IXM5414E to support IEEE 802.3ad aggregated links. By defining a Link Aggregation Group (LAG) you can treat multiple physical links between two end-points as one logical link. The LAG will also be seen by management functions as a single link.

LAGs are used to increase both link bandwidth and reliability: they are often used for links to the Internet or to shared servers. The script in the following example shows you how to configure and enable two LAGs on the same switch.

Create and name two LAGs.

```
config lag create lag_internet
```

```
config lag create lag_server
```

When the switch creates the LAGs, it will assign logical interface IDs that you will use to identify them in subsequent commands. Use the following command to find out what IDs have been assigned:

```
show lag all
```

Add the physical ports to the LAGs. (Assume that lag_internet was assigned ID lag.1 and lag_server was assigned ID lag.2.)

```
config lag addport lag.1 ext.1
```

```
config lag addport lag.1 ext.2
```

```
config lag addport lag.2 ext.3
```

```
config lag addport lag.2 ext.4
```

Enable both LAGs.

```
config lag adminmode lag.1,lag.2 enable
```

The previous command could have been issued instead as:

```
config lag adminmode all enable
```

IGMP snooping configuration example

This section provides sample CLI commands showing how to configure the Intel® Blade Server Ethernet Switch Module IXM5414E to support IGMP Snooping. Activating IGMP Snooping allows you to restrict the forwarding of multicast packets to network segments that need to see the packets. The switch uses information gained from examining IGMP packets to decide how to forward multicast packets.

You can activate IGMP Snooping for both individual and aggregated physical interfaces. The script in the following example show you how to configure IGMP Snooping.

Enable IGMP Snooping on the switch.

config igmpsnooping adminmode enable

IGMP Snooping will be enabled with default values for the group membership interval, maximum response and multicast router present expiration timers. This command overrides the default for the multicast router present expiration timer.

config igmpsnooping mcrtrexpiretime 2400

Enable IGMP Snooping for a set of physical ports and for a LAG.

config igmpsnooping interfacemode bay.1,bay.2,bay.3,bay.4 enable

config igmpsnooping interfacemode lag.1 enable

To display information about the IGMP Snooping configuration issue:

show igmpsnooping

To display information about all multicast addresses issue:

show mfdb table all

Access Control List configuration example

This section provides sample CLI commands showing how to configure the Intel® Blade Server Ethernet Switch Module IXM5414E to support Access Control Lists (ACLs). ACLs offer one way of adding Quality of Service support to your network.

You define an ACL to control who can use your network or network resources by allowing or prohibiting access. The ACL specifies one or more match criteria that will be used to determine whether a given packet will be admitted to the network. The first match criteria met by a packet determines whether the packet is admitted. If the packet matches none of the criteria, it will be dropped.

An ACL consists of up to ten rules, each applied to one or more of the following fields:

- Source IP address
- Destination IP address
- Source Layer-4 port
- Destination Layer-4 port
- Type of Service byte
- Internet Protocol number

The script in the following example restricts access to the network to UDP and TCP traffic from a defined set of IP source addresses.

Create Access Control List 1.

```
config acl create 1
```

Create Rule 1 for ACL 1.

```
config acl rule create 1 1
```

Define the content of ACL 1 Rule 1. Packets will be accepted only if they are TCP packets from the source IP address set defined by the specified IP address and mask.

```
config acl rule action 1 1 permit
```

```
config acl rule match protocol keyword 1 1 tcp
```

```
config acl rule match dstip 1 1 192.168.50.0 255.255.255.0
```

Create Rule 2 for ACL 1.

```
config acl rule create 1 2
```

Define the content of ACL 1 Rule 2. Packets will be accepted only if they are UDP packets from the source IP address set defined by the specified IP address and mask. This is the same source IP address set defined for TCP traffic.

```
config acl rule action 1 2 permit
```

```
config acl rule match protocol keyword 1 2 udp
```

```
config acl rule match dstip 1 2 192.168.50.0 255.255.255.0
```

Apply ACL 1 to inbound traffic received on external ports 1-4. Packets that do not match the criteria specified in Rules 1 or 2 will be dropped.

```
config acl interface add ext.1 inbound 1
```

```
config acl interface add ext.2 inbound 1  
config acl interface add ext.3 inbound 1  
config acl interface add ext.4 inbound 1
```

Appendix F 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 Intel® Blade Server Ethernet Switch Module IXM5414E, 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 a bridging device's ability 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 10. 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

Table 10. STP parameters – bridge

Parameter	Description	Default value
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 11. 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's view of the spanning tree topology
- Topology Change Notification (TCN) messages

Each BPDU includes 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 that bridge's root port (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 given 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 264. |
| Learning | The port is adding addresses to its forwarding database but not yet forwarding packets. For additional information, see “Forwarding state” on page 266. |

- Forwarding** The port is forwarding packets. For additional information, see “Forwarding state” on page 266.
- 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 268. Note that the STP port state of disabled applies only to the port’s role within the spanning tree, and should not be confused with the port’s 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 LACPDU) 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 may be configured to transition directly to the forwarding state when the bridge is reinitialized – this is appropriate 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 12. 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

Table 12. 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
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 13. 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:

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

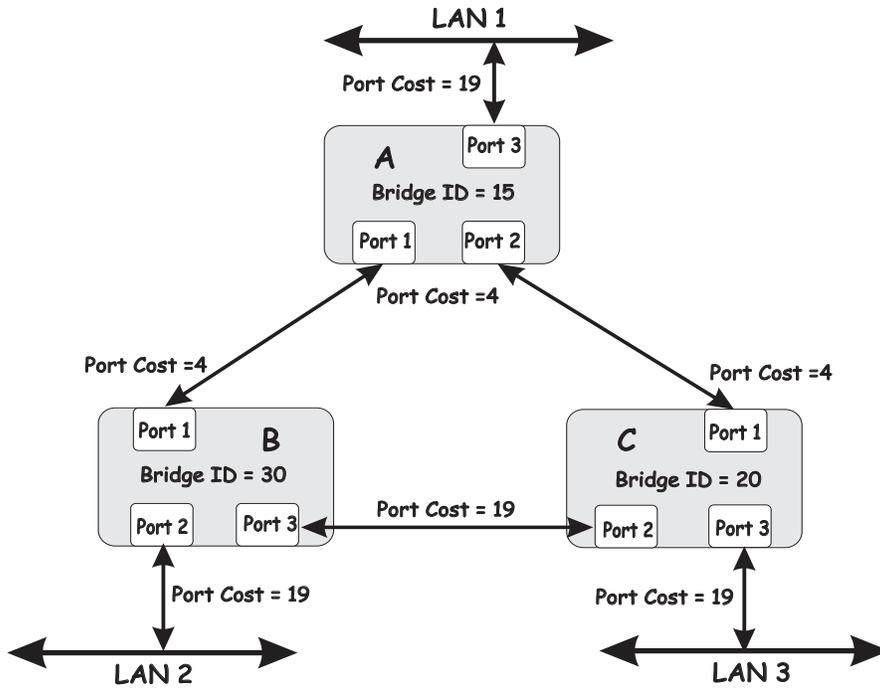
You can set a port cost from 1 to 200000000, or specify auto. The lower the number, the greater the probability that the port will be chosen to forward packets. If you specify auto the switch will assign the port cost based on the link speed.

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.

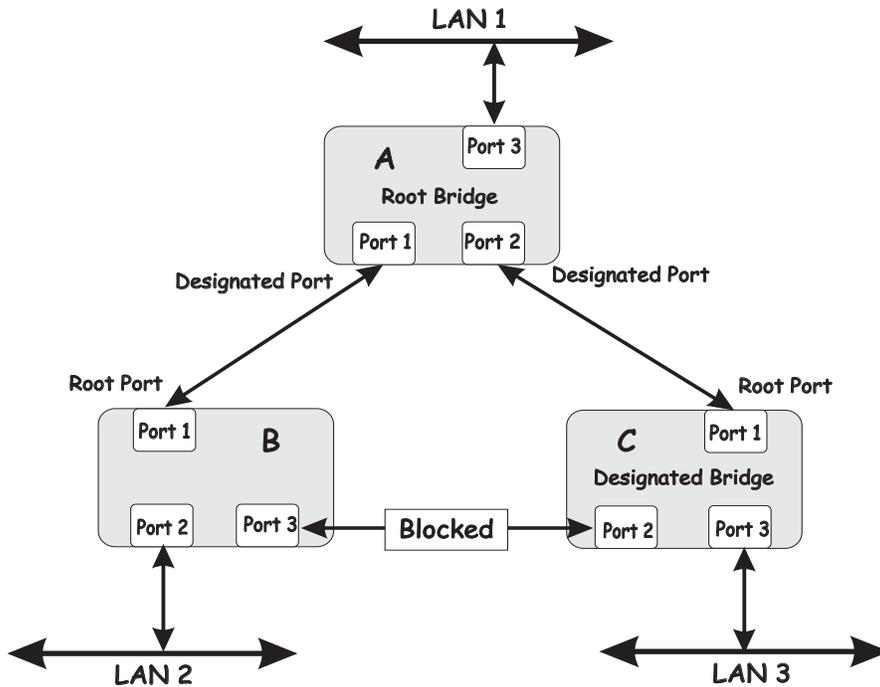
STP can be applied as shown in the following illustration. In this example, STP breaks the loop by blocking the connection between bridges B and C. The decision to block a particular connection is based on the STP calculation of the most current bridge and port settings. If bridge A broadcasts a packet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

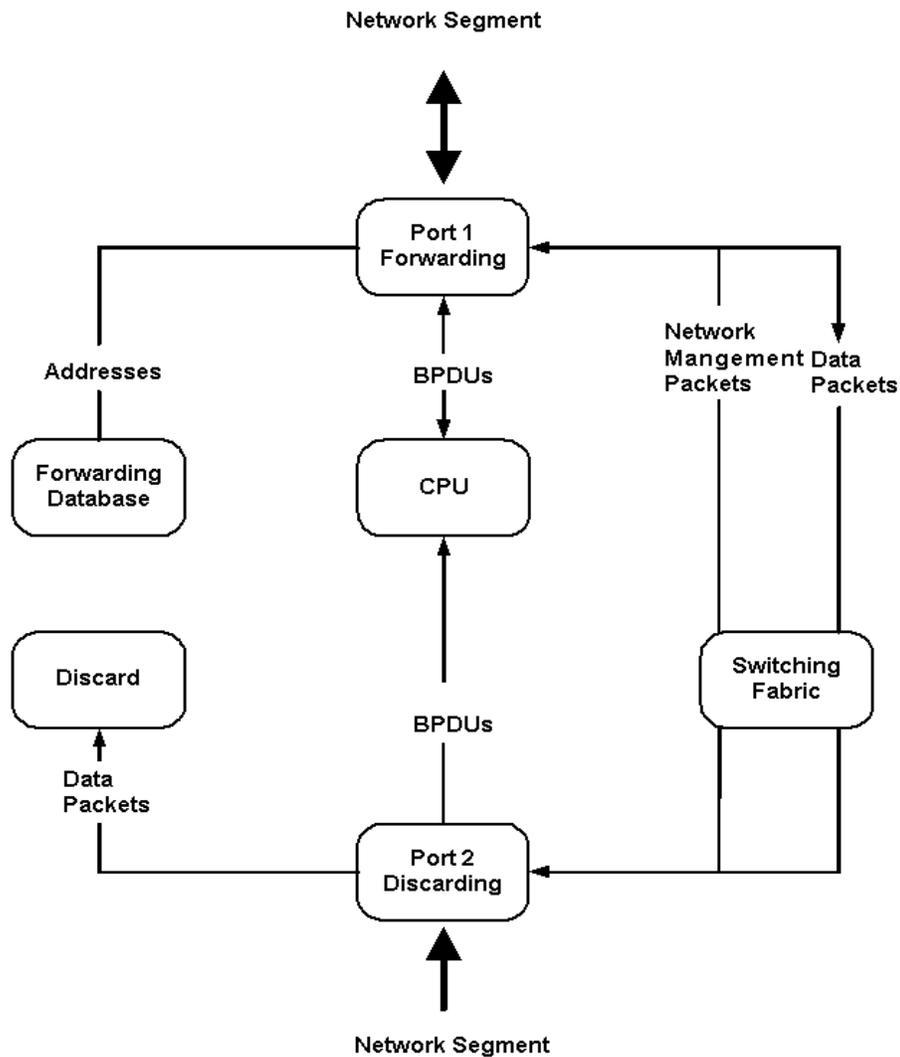
Discarding state

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.



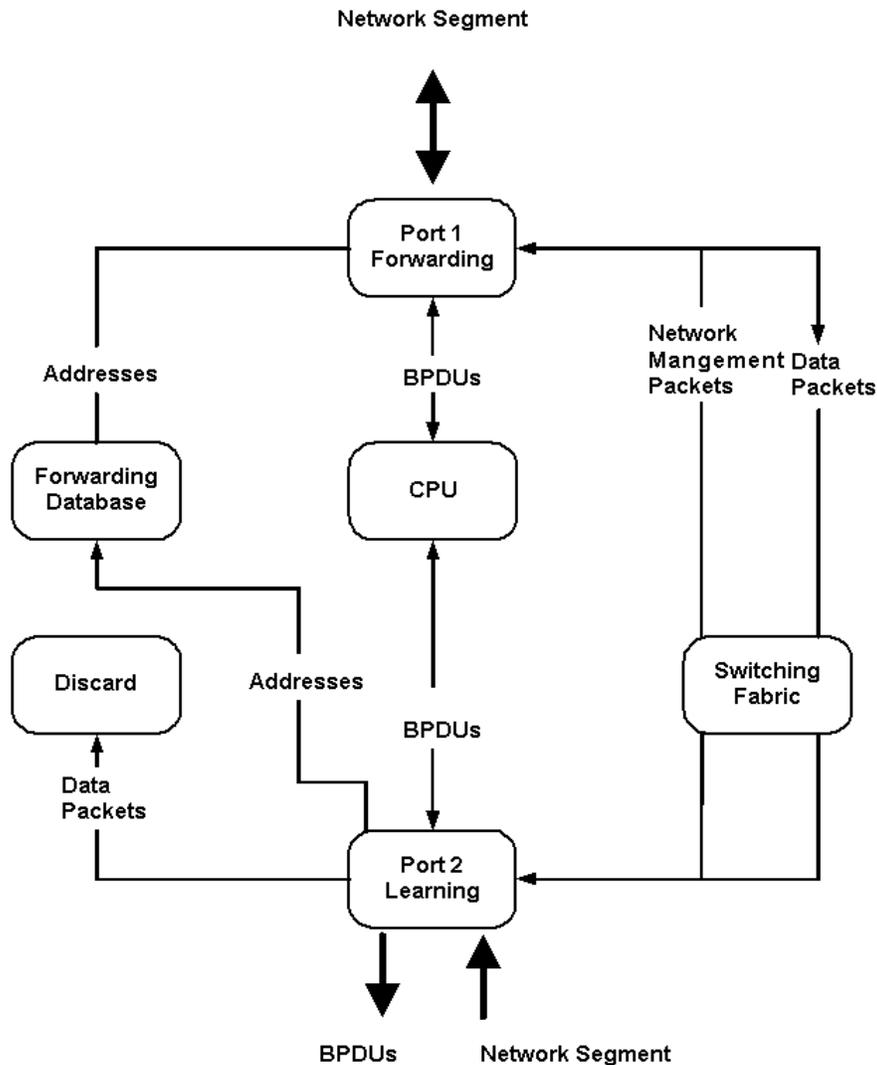
Learning state

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.



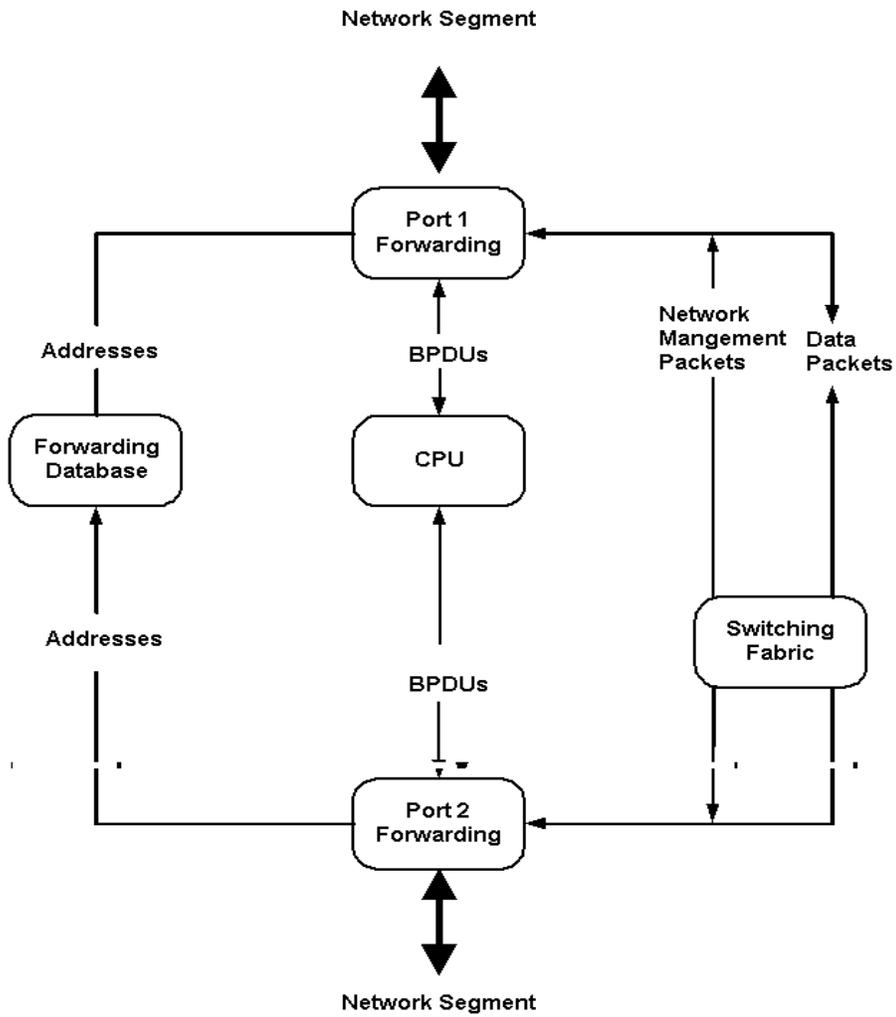
Forwarding state

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.



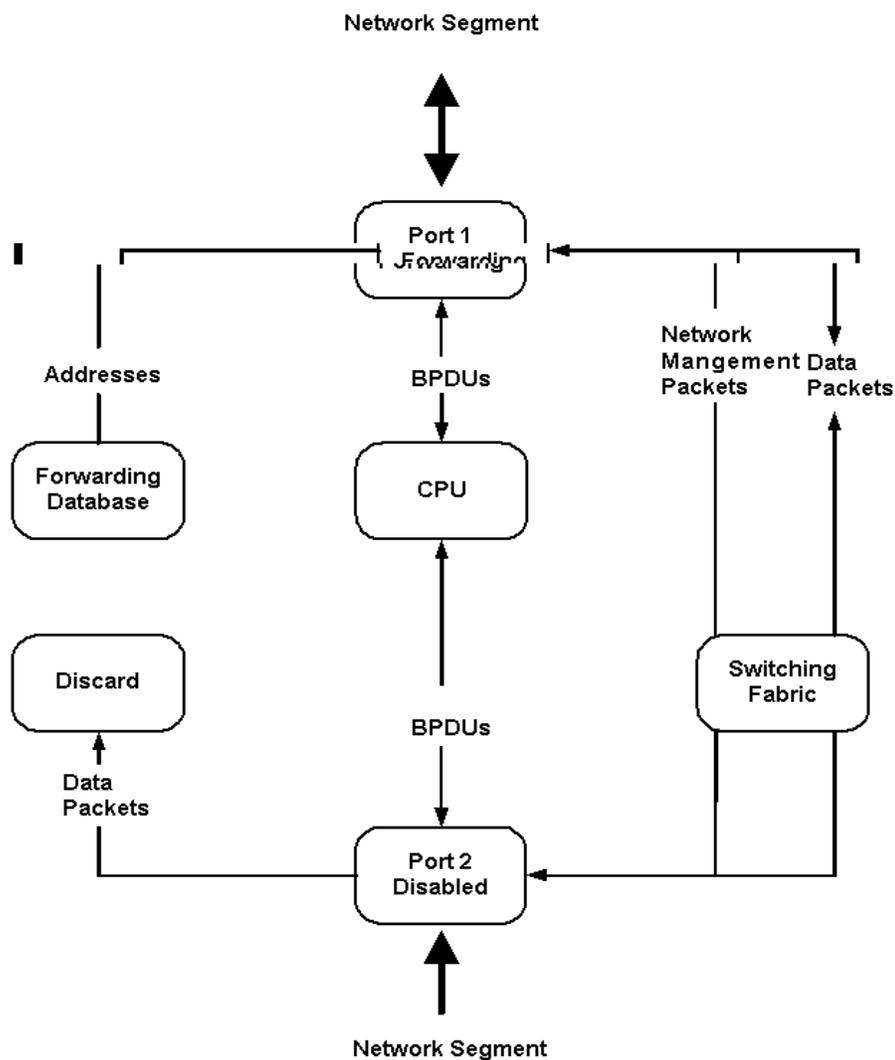
Disabled state

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's 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.

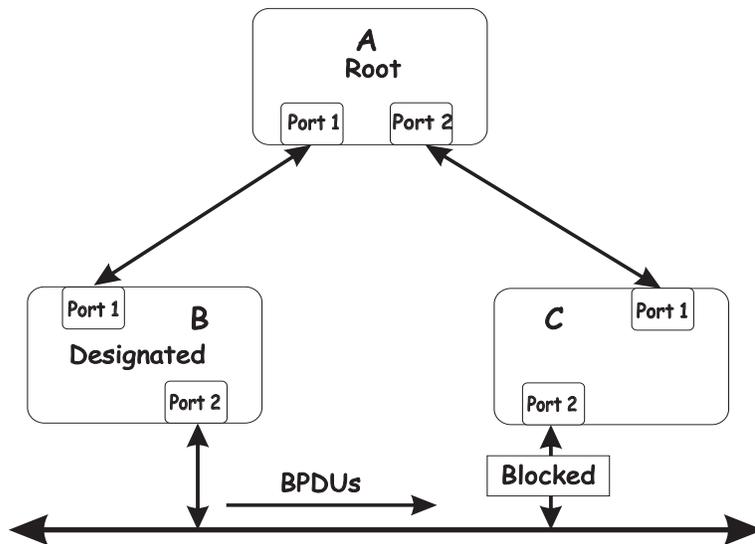


Troubleshooting STP

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.

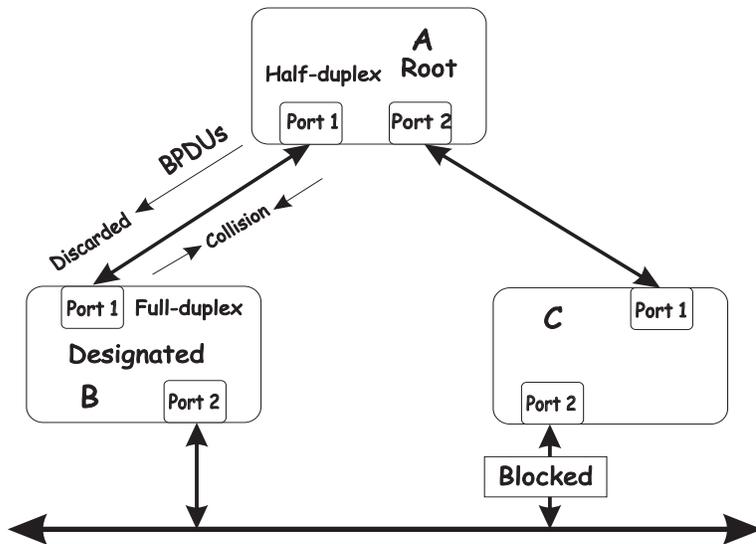
/ NOTE

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

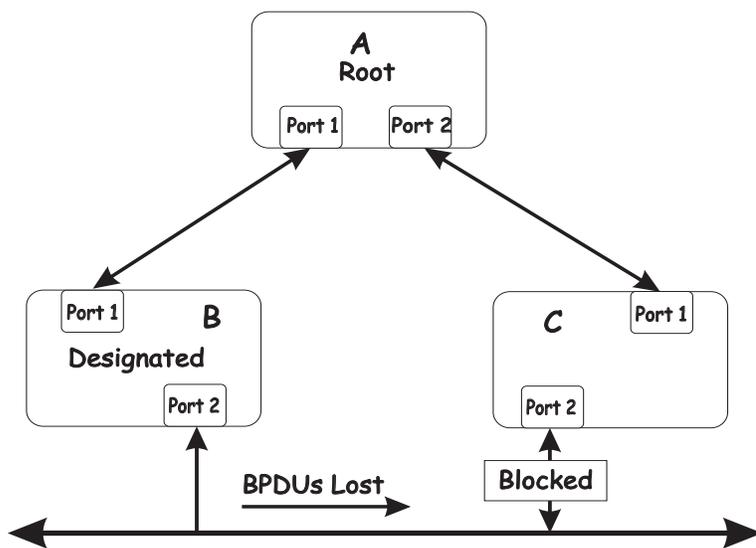
A mismatch in the duplex state of two ports is a very common configuration error for a point-to-point link. If one port is configured as full duplex and the other port is left in auto-negotiation mode, the second port will end up in half-duplex because ports explicitly configured as half- or full-duplex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

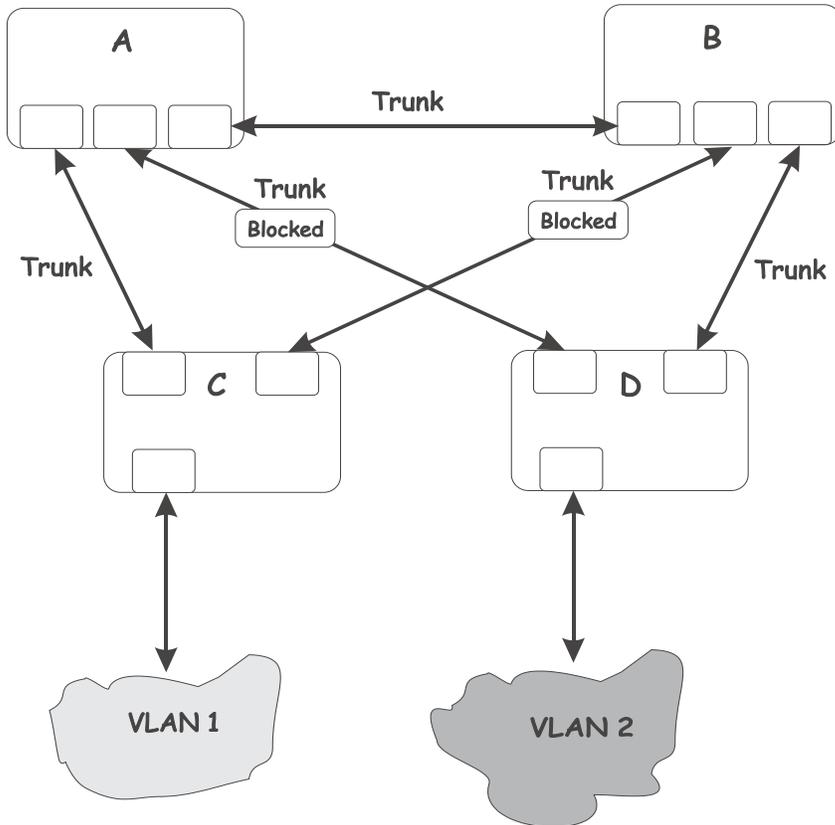
- Know which links are redundant.

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

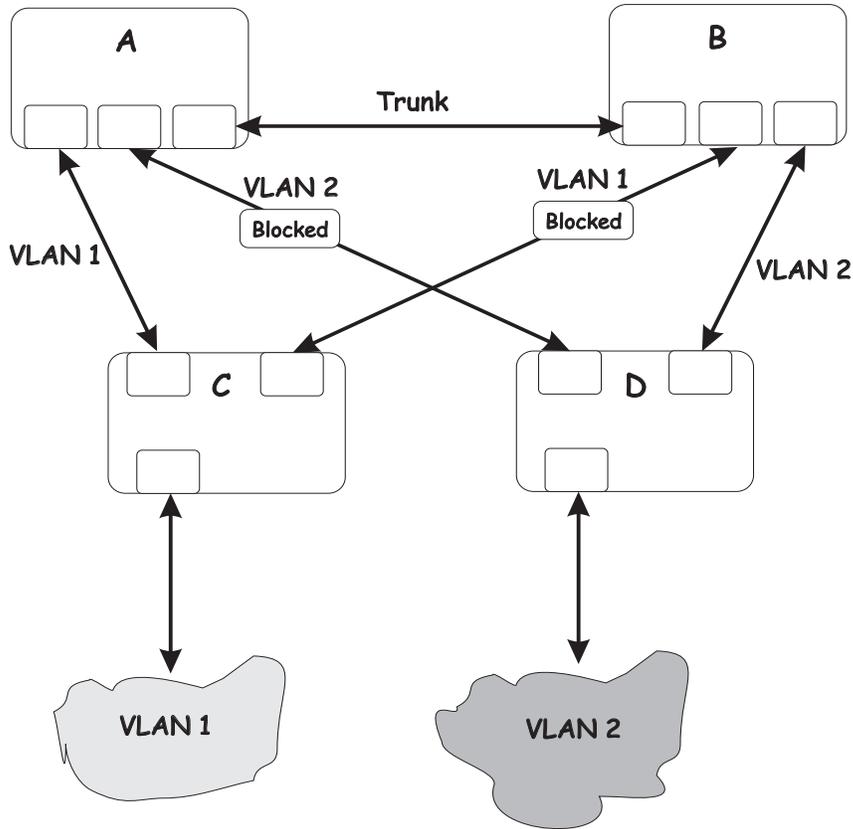
For each VLAN, know which ports should be discarding in a stable network. A network illustration that shows each physical loop in the network and which ports break which loops is extremely helpful.

- 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.

Appendix G Getting Help and Technical Assistance

This appendix contains information about where to go for additional information on Intel products, what to do if you experience a problem with your server platform, 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.

You can solve many problems without outside assistance by following the troubleshooting procedures that Intel provides in the publications that are provided on the Resource CD that ships with your system and software. The documentation also describes the diagnostic tests that you can perform. Most 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 server platform and pre installed software, if any, is available on the Resource CD that comes with your system. The Resource CD includes user manuals, maintenance manuals and troubleshooting guides. 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. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software.

Hardware and software service and support

Contact your Intel Support Representative for hardware and software service and support.

Appendix H Notices

Safety and regulatory information

/ NOTE

These service procedures are designed to help you isolate problems. They are written with the assumption that you have model-specific training on all computers, or that you are familiar with the computers, functions, terminology, and service information provided in this manual.

Important Safety Instructions

Read all caution and safety statements in this document before performing any of the instructions. See *Intel Server Boards and Server Chassis Safety Information* on the Resource CD and/or at <http://support.intel.com>.

Wichtige Sicherheitshinweise

Lesen Sie zunächst sämtliche Warn- und Sicherheitshinweise in diesem Dokument, bevor Sie eine der Anweisungen ausführen. Beachten Sie hierzu auch die Sicherheitshinweise zu Intel-Serverplatinen und -Servergehäusen auf der Ressourcen-CD oder unter <http://support.intel.com>.

重要安全指导

在执行任何指令之前，请阅读本文档中的所有注意事项及安全声明。参见 Resource CD（资源光盘）和/或 <http://support.intel.com> 上的 *Intel Server Boards and Server Chassis Safety Information*（《Intel 服务器主板与服务器机箱安全信息》）。

Consignes de sécurité

Lisez attention toutes les consignes de sécurité et les mises en garde indiquées dans ce document avant de suivre toute instruction. Consultez *Intel Server Boards and Server Chassis Safety Information* sur le CD Resource CD ou bien rendez-vous sur le site <http://support.intel.com>.

Instrucciones de seguridad importantes

Lea todas las declaraciones de seguridad y precaución de este documento antes de realizar cualquiera de las instrucciones. Vea *Intel Server Boards and Server Chassis Safety Information* en el CD Resource y/o en <http://support.intel.com>.

General Safety

Follow these rules to ensure general safety:

- Observe good housekeeping in the area of the machines during and after maintenance.
- Do not perform any action that causes hazards to the customer, or that makes the equipment unsafe.
- Place removed covers and other parts in a safe place, away from all personnel, while you are servicing the machine.
- Keep your tool case away from walk areas so that other people will not trip over it.
- Do not wear loose clothing that can be trapped in the moving parts of a machine. Ensure that your sleeves are fastened or rolled up above your elbows. If your hair is long, fasten it.
- Insert the ends of your necktie or scarf inside clothing or fasten it with a nonconductive clip, approximately 8 centimeters (3 inches) from the end.
- Do not wear jewelry, chains, metal-frame eyeglasses, or metal fasteners for your clothing.
Remember: Metal objects are good electrical conductors.
- Wear safety glasses when you are: hammering, drilling soldering, cutting wire, attaching springs, using solvents, or working in any other conditions that might be hazardous to your eyes.
- After service, reinstall all safety shields, guards, labels, and ground wires. Replace any safety device that is worn or defective.
- Reinstall all covers correctly before returning the machine to the customer.

Electrical Safety

CAUTION:

Electrical current from power, telephone, and communication cables can be hazardous. To avoid personal injury or equipment damage, disconnect the server system power cords, telecommunication systems, networks, and modems before you open the server covers.

Important: Observe the following rules when working on electrical equipment:

- Disconnect all power before performing a mechanical inspection.
- Before you start to work on the machine, unplug the power cord. or power-off the wall box that supplies power to the machine and to lock the wall box in the off position.
- Regularly inspect and maintain your electrical hand tools for safe operational condition.
- Do not use worn or broken tools and testers.
- Never assume that power has been disconnected from a circuit. First, check that it has been powered-off.
- Always look carefully for possible hazards in your work area. Examples of these hazards are moist floors, nongrounded power extension cables, power surges, and missing safety grounds.
- Do not touch live electrical circuits with the reflective surface of an inspection mirror. The surface is conductive; such touching can cause personal injury and machine damage.

Handling electrostatic discharge-sensitive devices

Any computer part containing transistors or integrated circuits (IC) should be considered sensitive to electrostatic discharge (ESD). ESD damage can occur when there is a difference in charge between objects. Protect against ESD damage by equalizing the charge so that the server, the part, the work mat, and the person handling the part are all at the same charge.

/ NOTE

Use product-specific ESD procedures when they exceed the requirements noted here.

Make sure that the ESD-protective devices you use have been certified (ISO 9000) as fully effective.

When handling ESD-sensitive parts:

- Keep the parts in protective packages until they are inserted into the product.
- Avoid contact with other people.
- Wear a grounded wrist strap against your skin to eliminate static on your body.
- Prevent the part from touching your clothing. Most clothing is insulative and retains a charge even when you are wearing a wrist strap.
- Use the black side of a grounded work mat to provide a static-free work surface. The mat is especially useful when handling ESD-sensitive devices.
- Select a grounding system, such as those in the following list, to provide protection that meets the specific service requirement.
 - Attach the ESD ground clip to any frame ground, ground braid, or green-wire ground.
 - Use an ESD common ground or reference point when working on a double-insulated or battery-operated system. You can use coax or connector-outside shells on these systems.
 - Use the round ground-prong of the AC plug on AC-operated computers.

/ NOTE

The use of a grounding system is desirable but not required to protect against ESD damage.

CAUTION:

If your system has a module containing a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of. Do not:

- Throw or immerse into water
- Heat to more than 100°C (212°F)
- Repair or disassemble
- Dispose of the battery as required by local ordinances or regulations.

CAUTION:

When laser products (such as CD-ROMs, DVD-ROM 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.

 **CAUTION:**

Hazardous energy is present when the blade is connected to the power source. Always replace the blade cover before installing the blade.

Regulatory specifications and disclaimers

Safety compliance

USA:	UL 60950 - 3rd Edition/CSA 22.2. No. 60950
Canada:	cUL certified - 3rd Edition/CSA 22.2. No. 60950- for Canada (product bears the single cUL mark for U.S. and Canada)
Europe:	Low Voltage Directive, 73/23/EEC UL/CB to EN60950 3rd Edition
International:	UL/CB to IEC 60950 3rd Edition UL/CB - EN60 950 3rd Edition UL/CB - EMKO-TSE (74-SEC) 207/94
Australia/New Zealand:	CB Report to IEC 60950, 3rd Edition plus international deviations

Electromagnetic compatibility (EMC)

USA:	FCC CFR 47 Part 2 and 15, Verified Class A Limit
Canada:	IC ICES-003 Class A Limit
Europe:	EMC Directive, 89/336/EEC EN55022, Class A Limit, Radiated & Conducted Emissions EN55024 ITE Specific Immunity Standard EN61000-4-2 ESD Immunity (Level 2 Contact Discharge, Level 3 Air Discharge) EN61000-4-3 Radiated Immunity (Level 2) EN61000-4-4 Electrical Fast Transient (Level 2) EN61000-4-5 AC Surge EN61000-4-6 Conducted RF EN61000-4-8 Power Frequency Magnetic Fields EN61000-4-11 Voltage Dips and Interrupts
Japan:	VCCI Class A ITE (CISPR 22, Class A Limit)
Australia/New Zealand:	AS/NZS 3548, Class A Limit
Taiwan:	BSMI Approval
Korea:	RRL Approval
Russia:	GOST Approval
International:	CISPR 22, Class A Limit

Electromagnetic compatibility notice (USA)

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.

Electromagnetic compatibility notices (International)

Europe (CE Declaration of Conformity): This product has been tested in accordance to, and complies with the Low Voltage Directive (73/23/EEC) and EMC Directive (89/336/EEC). The product has been marked with the CE Mark to illustrate its compliance.

Japan EMC Compatibility:

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。

English translation of the notice above: This is a Class A product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may arise. When such trouble occurs, the user may be required to take corrective actions.

ICES-003 (Canada): Cet appareil numérique respecte les limites bruits radioélectriques applicables aux appareils numériques de Classe A prescrites dans la norme sur le matériel brouilleur: “Appareils Numériques”, NBM-003 édictée par le Ministre Canadian des Communications.

English translation of the notice above: This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus set out in the interference-causing equipment standard entitled “Digital Apparatus”, ICES-003 of the Canadian Department of Communications.

BSMI (Taiwan): The BSMI Certification number and the following warning is located on the product safety label which is located visibly on the external chassis.

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。

RRL Korea:

기종별	주요특징
A급 기기	이 기기는 업무용으로 전자파 적합평가를 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바랍니다. 판매와 구매시 또는 구입하였을 때에는 가정용으로 표창하시기 바랍니다.
B급 기기	이 기기는 가정용으로 전자파 적합평가를 한 기기로서 주거지역에서는 물론 모든 지역에서 사용할 수 있습니다.

※ 비고
 A급 기기 : 업무용 정보통신기기를 말한다.
 B급 기기 : 가정용 정보통신기기를 말한다.

English translation of the previous notice:

Device	User's Information
Class A device	This device complies with RRL EMC and is operated in commercial environment so that distributors or users pay attention to this point. If the product is sold or purchased improperly, please exchange this product to what can be used at home.
Class B device	This device complies with RRL EMC and is operated in a residential area so that it can be used at all other location as well as residential area.
Remarks: Class A device - operated in a commercial area. Class B device - operated in a residential area.	

