SuperStack® II Switch 3900 and 9300 Implementation Guide

Release 3.0

http://www.3com.com/

Part No. 10012709
Published September 1999

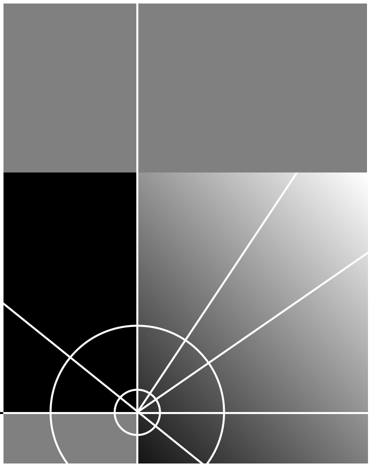

3Com Corporation
5400 Bayfront Plaza
Santa Clara, California
95052-8145

3Com Corporation reserves the right to revise this documentation and to make changes in content from time to time without obligation on the part of 3Com Corporation to provide notification of such revision or change.

3Com Corporation provides this documentation without warranty, term, or condition of any kind, either implied or expressed, including, but not limited to, the implied warranties, terms, or conditions of merchantability, satisfactory quality, and fitness for a particular purpose. 3Com may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

If there is any software on removable media described in this documentation, it is furnished under a license agreement included with the product as a separate document, in the hardcopy documentation, or on the removable media in a directory file named LICENSE.TXT or !LICENSE.TXT. If you are unable to locate a copy, please contact 3Com and a copy will be provided to you.

UNITED STATES GOVERNMENT LEGEND

If you are a United States government agency, then this documentation and the software described herein are provided to you subject to the following:

All technical data and computer software are commercial in nature and developed solely at private expense. Software is delivered as “Commercial Computer Software” as defined in DFARS 252.227-7014 (June 1995) or as a “commercial item” as defined in FAR 2.101(a) and as such is provided with only such rights as are provided in 3Com’s standard commercial license for the Software. Technical data is provided with limited rights only as provided in DFAR 252.227-7015 (Nov 1995) or FAR 52.227-14 (June 1987), whichever is applicable. You agree not to remove or deface any portion of any legend provided on any licensed program or documentation contained in, or delivered to you in conjunction with, this documentation.

Unless otherwise indicated, 3Com registered trademarks are registered in the United States and may or may not be registered in other countries.

3Com, the 3Com logo, CoreBuilder, Dynamic*Access*, NETBuilder II, PACE, SmartAgent, SuperStack, and Transcend are registered trademarks of 3Com Corporation. 3Com Facts is a service mark of 3Com Corporation.

AppleTalk and Macintosh are registered trademarks of Apple Computer, Incorporated. Banyan and VINES are registered trademarks of Banyan Worldwide. DEC, DECnet, and PATHWORKS are registered trademarks of Compaq Computer Corporation. AIX, IBM, and NetView are registered trademarks and NetBIOS is a trademark of International Business Machines Corporation. Internet Explorer, Microsoft, MS-DOS, Windows, and Windows NT are registered trademarks of Microsoft Corporation. Netscape, Netscape Navigator, and the Netscape N and Ship’s Wheel logos are registered trademarks of Netscape Communications Corporation in the United States and other countries. IPX, Novell, and NetWare are registered trademarks of Novell, Inc. Sun and SunNet Manager are trademarks of Sun Microsystems, Inc. Xerox and XNS are trademarks of Xerox Corporation. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

All other company and product names may be trademarks of the respective companies with which they are associated.

CONTENTS

ABOUT THIS GUIDE

1 CONFIGURATION OVERVIEW

2 MANAGEMENT ACCESS

3 PHYSICAL PORT NUMBERING

4 SYSTEM PARAMETERS

5 ETHERNET

6 BRIDGE-WIDE AND BRIDGE PORT PARAMETERS

8 TRUNKING

9 RESILIENT LINKS

10 VIRTUAL LANS

11 INTERNET PROTOCOL (IP)

12 IP MULTICAST FILTERING WITH IGMP

13 DEVICE MONITORING

A TECHNICAL SUPPORT

INDEX

ABOUT THIS GUIDE

The *SuperStack II Switch 3900 and 9300 Implementation Guide* provides information that you need to use features of the SuperStack® II Switch 3900 and Switch 9300 systems after you install and attach either of them to your network. Before you use this guide:

- Verify that your system is installed and set up using the *SuperStack II Switch 3900 Getting Started Guide* or the *SuperStack II Switch 9300 Getting Started Guide*.
- Become familiar with the *Command Reference Guide*, which documents the commands that you use to configure and manage your system from the Administration Console, a menu-driven command line interface that is embedded in the system software.
- If you want to manage your system from a Web browser, become familiar with the *Web Management User Guide for the SuperStack II Switch 3900 and Switch 9300*.
- Read Chapter 1 for an overview of the configuration process.

This guide is intended for the system or network administrator who is responsible for configuring, using, and managing the Switch 3900 and Switch 9300 systems. It assumes a working knowledge of LAN operations and familiarity with LAN communications protocols.

If the information in the release notes that relate to your version of software differs from the information in this guide, follow the release notes.

Conventions

Table 1 and Table 2 list icon and text conventions that are used throughout this guide.

Table 1 Icons

Icon	Type	Description
ℹ	Information note	Information that describes important features or instructions
⚠	Caution	Information that alerts you to potential loss of data or potential damage to an applications, system, or device
⚡	Warning	Information that alerts you to potential personal injury
L2	Layer 2 switch	In figures, a switch that can perform Layer 2 functions
L2/3	Layer 3 switch	In figures, a switch that can perform both Layer 2 and Layer 3 functions

Table 2 Text Conventions

Convention	Description
`Screen displays`	This typeface represents information as it appears on an Administration Console menu or screen display.
Syntax	The word “syntax” means that you evaluate the syntax provided and then supply the appropriate values. Example: To set the system date and time, use the following syntax: `mm/dd/yy hh:mm:ss: xM`
`Commands`	The word “command” means that you enter the Administration Console command exactly as shown in the text and then press Return or Enter. Commands appear in bold. For example: To remove an IP interface, enter the following command: **`ip interface remove`** ℹ *This guide is conceptual in nature and all commands are not addressed. When this guide does reference a command, you see the full form of the command. For valid minimum abbreviations and complete command information, see the* Command Reference Guide.
The words “enter” and “type”	When you see the word “enter” in this guide, you must type something, and then press Return or Enter. Do not press Return or Enter when an instruction simply says “type.”

Table 2 Text Conventions (continued)

Convention	Description
Keyboard key names	If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del.
Words in *italics*	Italics are used to: ■ Emphasize a point ■ Denote a new term when it is defined in text

Documentation

This section describes the documentation delivery formats and lists the titles of documents that pertain to the Switch 3900 and Switch 9300.

Formats

Documents are available in three forms:

- Paper Documents

  The paper documents that are shipped with your system and components are listed in the next section.

- Documents on CD-ROM

  The *Documentation CD-ROM* contains online versions of the paper documents, this *Implementation Guide*, and the *Command Reference Guide*.

- Documents on the Web

  You can view and print most 3Com documentation from the Web:

  `http://www.3com.com`

To order additional copies of the paper documents and the CD-ROM, contact your network supplier.

Paper Documents

These documents are shipped with your system:

- *SuperStack II Switch 3900 and 9300 Unpacking Instructions*

  How to unpack your system. Also, an inventory list of all the items that are shipped with your system.

- *SuperStack II Switch 3900 and 9300 Software Installation and Release Notes*

  Information about the software release, including new features, software corrections, and known problems. It also describes any changes to the documentation.

- *SuperStack II Switch 3900 and 9300 Quick Installation Guide*

  Quick reminders and information for system installation.

- *SuperStack II Switch 3900 Getting Started Guide* and *SuperStack II Switch 9300 Getting Started Guide*

  All of the procedures necessary for getting your system up and running, including information on installing, cabling, powering up, configuring, and troubleshooting the system.

- *SuperStack II Switch 3900 and 9300 Command Quick Reference*

  A list of the Administration Console commands for the systems in a convenient booklet.

- *Web Management User Guide for the SuperStack II Switch 3900 and Switch 9300*

  Overview, installation, and troubleshooting information for the suite of applications that help you manage your system from a Web browser.

In addition, each optional component ships with a guide:

- *1000BASE-SX/1000BASE-LX Gigabit Ethernet Module Installation Guide*

  How to install the optional Gigabit Ethernet module.

- *SuperStack II Switch Advanced RPS User Guide*

  How to install the Advanced Redundant Power Supply (RPS) and how to use it to provide redundant and resilient power supplies.

- *SuperStack II Switch Advanced RPS 'Y' Cable Type 2 User Guide*

  How to install the Y cable with the Advanced Redundant Power Supply (RPS) to provide fully redundant capabilities.

Documents on CD-ROM

The *Documentation CD-ROM* that ships with your system contains:

- Online versions of the paper documents that are shipped with your system and optional components.

The CD-ROM does not include release notes or the SuperStack II Switch 3900 and 9300 Command Quick Reference *booklet.*

- *SuperStack II Switch 3900 and Switch 9300 Implementation Guide* (this guide)
- *Command Reference Guide*

  Information about the commands used to configure the system. This multiplatform guide documents commands for the Switch 3900 and Switch 9300 as well as other 3Com systems.
- Help system for the Web Management suite of applications

  Online Help system for the SuperStack II Switch 3900 and Switch 9300 Web Management software. See the *Web Management User Guide for the SuperStack II Switch 3900 and Switch 9300* for information about Web Management and the related Help system.

Documentation Comments

Your suggestions are very important to us. They help us to make our documentation more useful to you.

Please send e-mail comments about this guide to:

`sdtechpubs_comments@ne.3com.com`

Please include the following information when you comment:

- Document title
- Document part number (found on the front or back page of each document)
- Page number

Example:

SuperStack II Switch 3900 and 9300 Implementation Guide

Part Number 10012709

Page 25

Year 2000 Compliance

For information on Year 2000 compliance and 3Com products, visit the 3Com Year 2000 Web page:

`http://www.3com.com/products/yr2000.html`

1 CONFIGURATION OVERVIEW

This chapter provides an initial configuration procedure for the SuperStack® II Switch 3900 and SuperStack II Switch 9300 with Release 3.0 software installed.

If you are upgrading an existing system to Release 3.0 or later software, be sure to read the release notes for upgrade considerations and procedures first.

System Configuration Procedure

Software is installed on each system at the factory. Because the software boots from flash memory when you power on the system, the system is immediately ready to configure according to your network needs.

See the Getting Started Guide *for your system and the* Quick Installation Guide *for information about the physical installation process and LED diagnostic indicators.*

3Com recommends that you use the following procedure to configure your system. Follow the steps that apply to your network needs and ignore the steps that do not apply:

1 Establish management access and use the Administration Console.

To perform initial system management tasks you must connect to the system using the Console port and a terminal or modem serial connection. The Console port is located on the front panel and is clearly labeled. For information about the required settings for the Console port, see Chapter 2 in this guide.

The first time that you log into the system, you must use the Administration Console as the management interface. It is the system's internal menu-driven command line interface. For detailed information about the menus and commands, see the *Command Reference Guide*.

2 Choose a subsequent management access method.

After you establish management access, you can continue to manage the system through the Console port as described in step 1, or you can configure the system so that you can use the following access methods:

- **Out-of-band management port (Switch 9300 only)** — The 10BASE-T Ethernet port on the front panel of the SuperStack II Switch 9300.

  Before you can use the out-of-band management port, you must create an IP interface in the system. To begin the process, use the Console port connection from step 1 and enter:

  ```
  ip interface define
  ```

  Switch 9300 only: When the system prompts you to enter an interface type, enter:

  ```
  system
  ```

  This selection indicates that it is an out-of-band interface. For more information about IP interfaces, see Chapter 11.

- **In-band management port** — Any Fast Ethernet or Gigabit Ethernet port in the system.

  Before you can use an in-band port for management, you must create an in-band IP interface, which consists of an IP address and a VLAN (such as the default VLAN).

  To begin the process, use the Console port connection from step 1 or, on the Switch 9300, use an out-of-band management connection and enter:

  ```
  ip interface define
  ```

  Switch 9300 only: When the system prompts you to enter an interface type, enter:

  ```
  vlan
  ```

  This selection indicates that it is an in-band interface. The Switch 3900 does not prompt you to specify an interface type because it supports only in-band interfaces.

  When the system prompts you to specify a VLAN interface index number, use the default index (1), unless you already have other appropriate VLANs defined. For more information about IP interfaces, see Chapter 11. For more information about VLANs, see Chapter 10.

Both the Switch 3900 and Switch 9300 allow you to specify a maximum of two unique IP interfaces, as shown in Table 3 and Table 4.

Table 3 IP interface Options for the SuperStack II Switch 9300

Interface Type	Number of IP Interfaces*
`system`	up to 1
`vlan`	up to 2

* The total number of IP interfaces cannot exceed 2.

Table 4 IP interface Options for the SuperStack II Switch 3900

Interface Type	Number of IP Interfaces*
`vlan`	up to 2

* The total number of IP interfaces cannot exceed 2.

For more information about management access, see Chapter 2.

3 Choose a subsequent management interface.

You can continue to use the Administration Console as the management interface, or, after you establish an IP interface, you have two more management interface options:

- **Web Management software** — A suite of HTML-based Help and several monitoring applications are shipped with your system. For more information, see the *Web Management User Guide*.
- **SNMP-based applications** — 3Com Transcend® software products are examples of SNMP-based network management applications that you can use as management interfaces. To manage the system in-band with such applications, set the SNMP parameters through the `snmp` menu. For more information, see Chapter 2 and Chapter 13.

4 Learn how ports are numbered in the system.

For pictures that explain port numbering for these systems, see Chapter 3.

5 Administer system parameters.

To configure Simple Network Time Protocol (SNTP) parameters, administer nonvolatile data (nvData), update system software, display your system configuration, and perform other system activities, see Chapter 4.

6 Set Ethernet parameters.

To label Ethernet ports, set the port mode, enable flow control, and control autonegotiation and other settings, see Chapter 5.

7 Set bridge-wide and bridge port parameters.

To set Spanning Tree Algorithm parameters for the system or specific ports, see Chapter 6.

8 Modify Class of Service priority queue assignments.

Your system includes one high priority queue and one low priority queue per port so that it can prioritize business-critical or time-critical traffic over other network traffic. To receive Class of Service handling, packets must use the IEEE 802.1Q format and carry one of eight possible numeric traffic class values as noted in the IEEE 802.1p specification. To modify the default queue configuration or set a rate limit on the high priority queue, see Chapter 7.

9 Create trunks to increase bandwidth and resiliency.

To increase the bandwidth and resiliency between two switches, you can combine multiple Fast Ethernet or Gigabit Ethernet ports into a single high-speed link called a trunk (also known as an aggregated link). For more information about trunking, see Chapter 8.

3Com recommends that you configure trunks before you define VLANs because trunk configurations affect VLAN configurations. This approach minimizes your administrative tasks.

10 Increase network availability with resilient links.

To use the resilient links feature to ensure network availability between two switches or between a switch and a server, see Chapter 9.

11 Define port-based VLANs.

To create logical workgroups, which are generally equivalent to Layer 2 broadcast domains, you can define port-based VLANs. These VLANs support IEEE 802.1Q tagging on each port. You can also select a VLAN mode that determines whether data can be forwarded between VLANs. For more information about port-based VLANs, including their rules of operation and the maximum number of VLANs that you can configure, see Chapter 10.

To minimize your administrative tasks, 3Com recommends that you configure trunks before you define VLANs.

12 Define a second IP interface.

You defined the first IP interface in step 2. The primary reason to configure a second IP interface is for management access from a different subnetwork.

You must define at least one additional VLAN beyond the default VLAN before you can define a second IP interface. Each IP interface of type VLAN (up to 2 are allowed) must be assigned to a unique VLAN.

13 Enable IP multicast filtering in the system.

The Internet Group Management Protocol (IGMP) allows the system to forward IP multicast packets only to the ports that lead to group members. When IGMP is disabled, the system floods IP multicast packets to all ports in compliance with IEEE 802.1D. IGMP helps to conserve network bandwidth. See Chapter 12 for more information.

14 Monitor devices and validate network paths.

As part of your overall approach to network management, you can use the system's device monitoring features to analyze your network periodically and identify potential network problems before they become serious problems. To identify potential problems in your network, use baselining, roving analysis, and RMON information. To test and validate paths in your network, use IP tools like ping and traceRoute. For more information on these features, see Chapter 13.

2

MANAGEMENT ACCESS

This chapter explains the methods that you can use to configure management access to the system. It describes the types of management applications and the communication and management protocols that deliver data between your management device (UNIX workstation, PC, or Macintosh computer) and the system. It also contains information about port connection options.

This chapter covers the following topics:

- Management Access Overview
- Key Concepts
- Key Guidelines for Implementation
- Administration Console Access
- Web Management Access
- SNMP Access
- Standards, Protocols, and Related Reading

Management Access Overview

The system gives you the flexibility to access and manage your system using several methods. You can administer your system using:

- The Administration Console
- Web Management suite of applications
- An external SNMP-based network management application such as 3Com's Transcend Network Control Services

The Administration Console and most of the Web Management applications are embedded in the software and are available for immediate use.

Administration Console Overview

The Administration Console is an internal, character-oriented, menu-driven, user interface for performing system administration such as displaying statistics or changing option settings.

You can view the Administration Console from a terminal, a PC, a Macintosh, or from a UNIX workstation.

You can access the Administration Console through the Console (serial) port or through an Ethernet port using an IP interface.

Figure 1 shows a sample output of SuperStack® II Switch 3900 menu options that can be viewed from the various devices.

Figure 1 Viewing the Administration Console

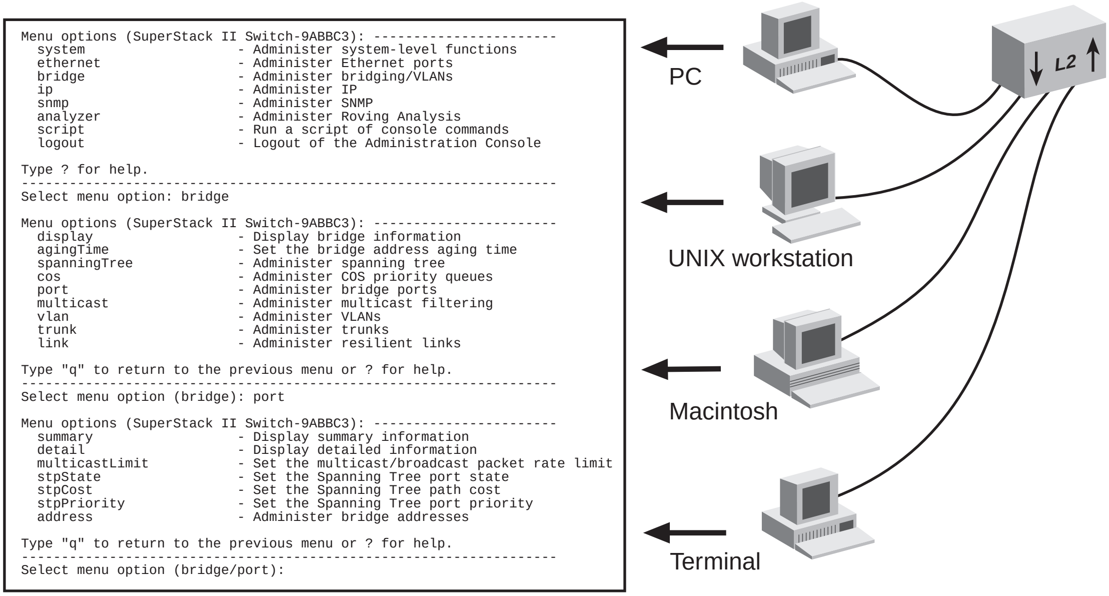

```
Menu options (SuperStack II Switch-9ABBC3): -----------------------
  system                   - Administer system-level functions
  ethernet                 - Administer Ethernet ports
  bridge                   - Administer bridging/VLANs
  ip                       - Administer IP
  snmp                     - Administer SNMP
  analyzer                 - Administer Roving Analysis
  script                   - Run a script of console commands
  logout                   - Logout of the Administration Console

Type ? for help.
-------------------------------------------------------------------
Select menu option: bridge

Menu options (SuperStack II Switch-9ABBC3): -----------------------
  display                  - Display bridge information
  agingTime                - Set the bridge address aging time
  spanningTree             - Administer spanning tree
  cos                      - Administer COS priority queues
  port                     - Administer bridge ports
  multicast                - Administer multicast filtering
  vlan                     - Administer VLANs
  trunk                    - Administer trunks
  link                     - Administer resilient links

Type "q" to return to the previous menu or ? for help.
-------------------------------------------------------------------
Select menu option (bridge): port

Menu options (SuperStack II Switch-9ABBC3): -----------------------
  summary                  - Display summary information
  detail                   - Display detailed information
  multicastLimit           - Set the multicast/broadcast packet rate limit
  stpState                 - Set the Spanning Tree port state
  stpCost                  - Set the Spanning Tree path cost
  stpPriority              - Set the Spanning Tree port priority
  address                  - Administer bridge addresses

Type "q" to return to the previous menu or ? for help.
-------------------------------------------------------------------
Select menu option (bridge/port):
```

Web Management Overview

The Web Management software consists of embedded Web Management applications and installable tools:

- **Embedded Web Management applications** — Use the embedded Web Management applications for most of your device configuration and management tasks. You can manage a single port or device, or, using multiple windows, you can manage multiple devices. This software, which is part of the system software image, contains:
  - **WebConsole** — An HTML-based set of configuration forms.
  - **DeviceView** — A Java-based application that displays a real-time image of the device. You can manage each port, module, or system by clicking the part of the image that you want to manage.
  - **Performance features** — Dynamic monitoring through graphing of QoS statistics and Ethernet interfaces.
  - **Help** — Access to the configuration form on which you set up the installable Help files as well as access to links to support information on the 3Com Web site.
- **Installable tools** — Install these optional tools on your workstation from the 3Com Web site:
  - **DeviceView accessories** — To set up e-mail notification for Status Logging
  - **WebManage Framework** — To group your access links to the devices that you manage
  - **Form-specific Help** — To get more information about WebConsole, DeviceView, and Performance forms

For details about this software, see the *Web Management User Guide.*

SNMP-Based Network Management Overview

For more complete network management, you can use an external SNMP-based application such as 3Com's Transcend Network Control Services or another network management application. You access external applications through an Ethernet port using an IP interface.

Figure 2 shows an example of a Transcend Network Control Services Device View screen.

Figure 2 Transcend Network Control Services Device View Screen

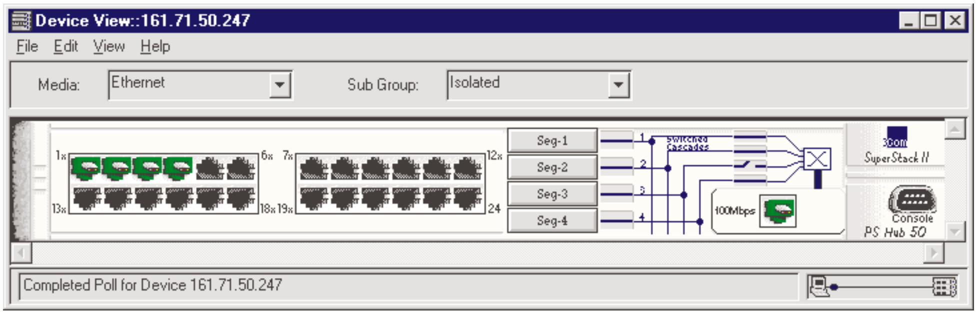

Media: Ethernet

Sub Group: Isolated

Completed Poll for Device 161.71.50.247

Key Concepts

This section describes the relationship between the methods of management access described in the previous sections and how they fit into established networking protocols. It also looks further into the concepts of in-band and out-of-band management using IP.

The SuperStack II Switch 9300 provides in-band or out-of-band management. The SuperStack II Switch 3900 provides in-band management only and does not have an out-of-band Ethernet port.

OSI Protocols

Management and administration on the system occur through the layers of the Open Systems Interconnection (OSI) reference model.

Figure 3 shows how the different management access methods fit into the OSI model.

Figure 3 OSI Protocols

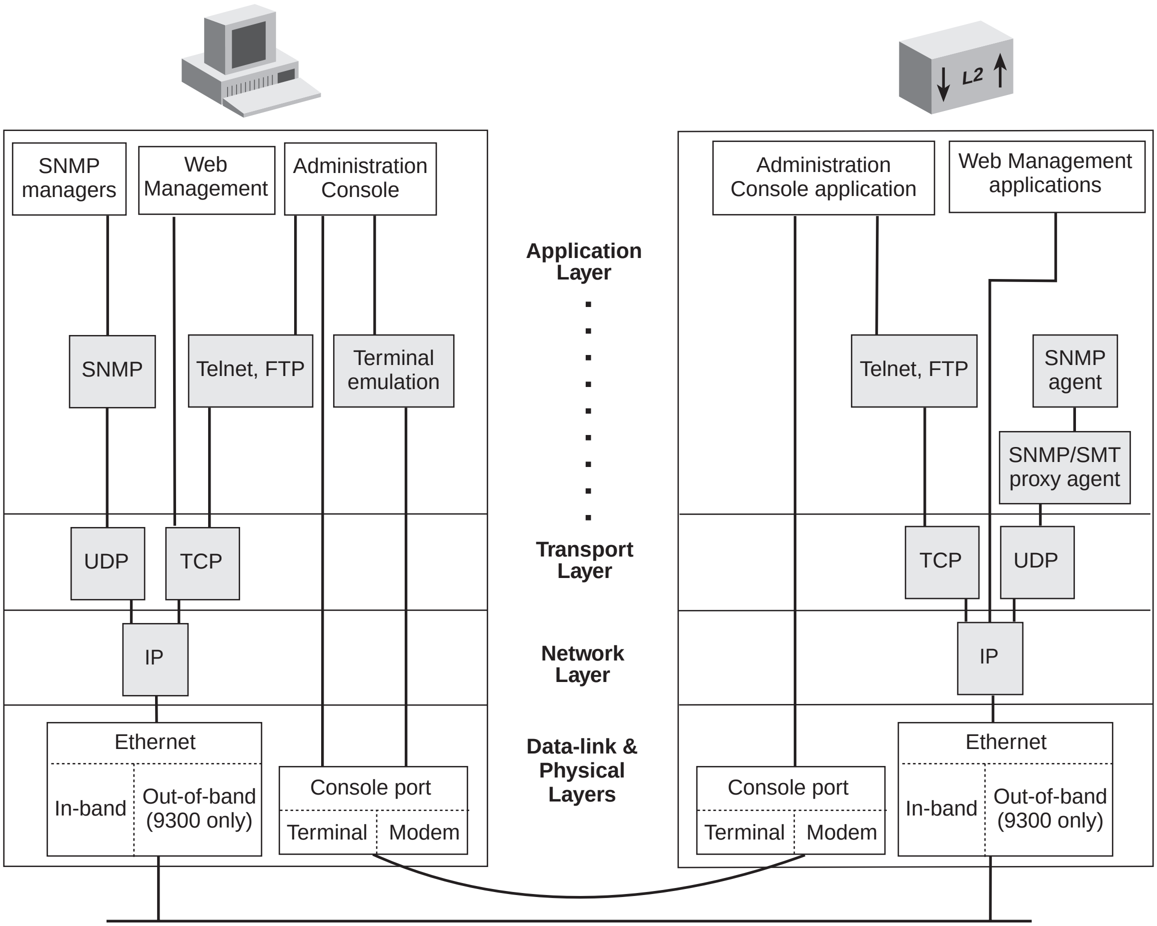

Protocols

The system supports the following protocols:

- Virtual terminal protocols, such as Telnet
- Simple Network Management Protocol (SNMP)

Virtual Terminal Protocols

A *virtual terminal protocol* is a software program, such as Telnet, that allows you to establish a management session from a Macintosh, a PC, or a UNIX workstation. Because Telnet runs over TCP/IP, you must have at least one IP address configured on the system before you can establish access to it with a virtual terminal protocol. Within the Administration Console, you configure an IP address by defining an IP interface. See the *Command Reference Guide* for additional information about defining IP addresses for in-band or out-of-band management.

Terminal emulation differs from a virtual terminal protocol in that you must connect a terminal directly to the Console port.

Figure 4 shows a UNIX workstation connected to the system through a virtual terminal protocol (Telnet), and a terminal connecting directly to the Console port through a null modem cable.

Figure 4 Administration Console Access

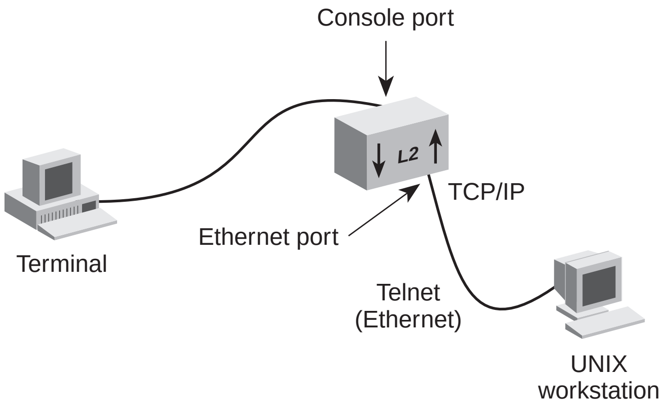

Simple Network Management Protocol

Simple Network Management Protocol (SNMP) is the standard management protocol for multivendor IP networks. SNMP supports transaction-based queries that allow the protocol to format messages and to transmit information between reporting devices and data-collection programs. SNMP runs on top of the User Datagram Protocol (UDP), offering a connectionless-mode service. Figure 5 shows a PC connected to the system through an Ethernet port.

Figure 5 SNMP Manager Access

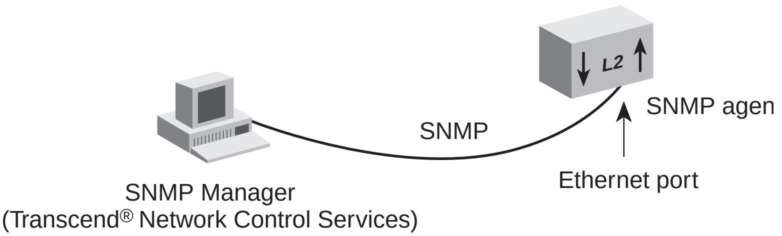

See Chapter 13 for additional information about SNMP.

IP Management Concepts

In-band and out-of-band management each have advantages and disadvantages:

- **In-Band Management** — If you manage your system and its attached LANs over the same network that carries your regular data traffic, then you are managing your network *in band*. This kind of management is often the most convenient and inexpensive way to access your system. The disadvantage is that, if your data network is faulty or congested, you may not be able to diagnose the problem because management requests are sent over the same network.
- **Out-of-Band Management (9300 only)** — If you are using a dedicated network for management data, then you are managing your network *out of band*. Although this is a more expensive way to access your network, you are able to diagnose problems even when your data network is faulty.

Key Guidelines for Implementation

This section describes guidelines for the different ways to access your system.

Access Methods

You can access your management application on the system several ways; locally through a terminal connection or remotely using a modem or an IP connection. Table 5 describes these methods.

Table 5 Management Access Methods

Access Method	Access Description	Interface
Terminal	Connect directly to the Administration Console and stay attached during system reboots.	Console (serial) port
Modem	Access the Administration Console by dialing in from remote sites.	Console (serial) port
IP	■ Access the Administration Console with up to four Telnet sessions. ■ Use an external SNMP management application to communicate with the system's SNMP agent. ■ Use your Internet browser to connect to the embedded Web Management suite of configuration forms.	In-band or out-of-band Ethernet port that is assigned to an IP interface

Setting Up the Terminal Port

Use the Administration Console to change the baud setting to match the speed of your terminal.

Baud setting changes take effect immediately after you confirm the change. You must adjust the baud setting of your terminal or terminal emulator to match your management interface port before you can reestablish communication using the terminal port. When you change the baud setting to something other than 9600, the new setting becomes the new default, even after you enter the `system nvData reset` *command.*

You can use the `system serialPort terminalSpeed` *command through the terminal serial port or through an IP interface. However, if you change the terminal speed while in a Telnet session, you must reboot the system before the change takes effect.*

Setting Up the Modem Port

Use the Administration Console to match your external modem speed. Then configure the external modem by establishing a connection between your current Administration Console session and the modem port.

You must establish a connection to the modem with the `system serialPort connectModem` *command after you change the modem speed and before you dial in. This sequence allows the modem to synchronize its baud rate with the system.*

See the *Getting Started Guide* for your system for Console port pin assignments. For additional information about modem port settings, see the *Command Reference Guide*.

IP Management Interface

An IP management interface allows you to manage the system in-band through an Ethernet port on the system or out-of-band (9300 only) through the out-of-band Ethernet port. You can access the system through an IP interface in one of the following ways:

- Use Telnet to connect up to four concurrent remote sessions to the Administration Console using a terminal program from a host computer.
- Run Web Management to access its management applications to manage and monitor your system.
- Run an SNMP-based network management application to manage and monitor your system.

IP is a standard networking protocol that is used for communications among various networking devices. To gain access to the system using TCP/IP or to manage the system using SNMP, you must set up an IP interface for your system.

How you set up the IP interface depends on whether you plan to manage the system in band (with your regular network traffic) or out of band (with a dedicated network):

For Telnet access, Web Management access, or SNMP access, you must first define an IP interface. You can use either an out-of-band (Switch 9300 only) or in-band port for the IP interface. For the Switch 9300, be careful not to assign the same IP address to both the out-of-band and the in-band ports. Also, be sure not to assign an out-of-band port IP address that is on the same subnetwork as any of the in-band IP interfaces. You can have a maximum of two IP addresses on either system.

- **In-band management** — To manage your network in-band, you need to set up an IP routing interface. An IP interface consists of a unique IP address and an assigned VLAN (such as the default VLAN). See Chapter 10 for information about defining a VLAN and see Chapter 11 for information about setting up an IP routing interface.
- **Out-of-band management (9300 only)** — To manage your system out of band, you need to assign an IP address and subnet mask for the out-of-band Ethernet port on your system. The out-of-band Ethernet port is the 10BASE-TX port located next to the Console port. See Chapter 3 for a diagram that calls out the ports on the system. See Chapter 11 for background information on IP addresses and subnet masks.

Administration Console Access

The first time that you access the Administration Console, access the system at the *administer* level and press Return at the password prompt. The initial password is null. Subsequent access is described next.

Password Levels

The Administration Console supports three password levels, allowing the network administrator to provide different levels of access for a range of users, as described in Table 6.

Table 6 Password Access Levels

Access Level	For the User Who Needs to	Allows the User to
Administer	Perform system setup and management tasks (usually a single network administrator)	Perform system-level administration (such as setting passwords, loading new software, and so on)
Write	Perform active network management	Configure network parameters (such as setting bridge aging time)
Read	Only view system parameters and settings	Access only *display* menu items (display, summary, detail)

Passwords are stored in nonvolatile (NV) memory. You must enter the password correctly before you can continue.

When you access the Administration Console, the top-level menu appears. You manage and monitor your system by selecting options from this menu and from others below it. Each menu option is accompanied by a brief description.

For additional information about using the Administration Console, see the *Command Reference Guide*.

Terminal Port Access

Direct access to the Administration Console through the Console (serial) port is often preferred because you can remain connected and monitor the system during system reboots. Also, certain error messages are sent to the serial port, regardless of the interface through which the associated action was initiated. A Macintosh or PC attachment can use any terminal emulation program for connecting to the terminal serial port. A workstation attachment under UNIX can use an emulator such as TIP.

Modem Port Access

You can access the Administration Console from your PC or Macintosh using an external modem attached to the Console (serial) port.

When you have configured the external modem from the Administration Console, the system transmits characters that you have entered as output on the modem port. The system echoes characters that it receives as input on the modem port to the current Administration Console session. The console appears to be directly connected to the external modem.

Web Management Access

Most Web Management applications are an embedded part of the SuperStack II 3900 and SuperStack II 9300 system software image. They include WebConsole, DeviceView, and Performance monitoring tools. Additional installable applications include Help.

After you have set up your IP address for the CoreBuilder 3500 system, you can access Web Management applications directly in your Web browser by entering the IP address.

In the Web Management interface window, you can list and manage all your devices from one central location. You can easily add and delete devices and group the devices in ways that make sense to you, for example, by location or subnetwork.

For more information, see the *Web Management User Guide*.

Browser Requirements

Web Management requires either Microsoft Internet Explorer 4.01 or later or Netscape Navigator 4.03 or later.

- **Netscape Navigator** — If you are using Netscape Navigator 4.03 or 4.04, be sure to install the Netscape JDK 1.1 Patch. Download the patch from the following location:

  ```
  http://help.netscape.com/filelib.html#smartupdate
  ```

  If you encounter problems accessing Help files when you use Netscape, clear the browser memory cache and disk cache and restart the browser.

- **Internet Explorer** — If you use Internet Explorer, install the latest 4.01 Service Pack 1. This service pack makes Internet Explorer Year 2000 compliant and fixes other product-support issues. Download the 4.01 Service Pack 1 from the following location:

  `http://www.microsoft.com/msdownload/iebuild/ie4sp1_win32/en/ie4sp1_win32.htm`

  If the above link is unavailable, download the service pack from the Microsoft home page:
  `http://www.microsoft.com`

SNMP Access

You can use an external SNMP-based application such as 3Com's Transcend Network Control Services lets you access your system through an Ethernet port using an IP interface. SmartAgent® intelligent agents are the foundation of the Transcend architecture. SmartAgent software and RMON work together to provide automatic network-wide monitoring, analysis, and reporting. For additional information about Transcend Network Control Services, see the 3Com Web page:

`http://www.3com.com`

Standards, Protocols, and Related Reading

The following standards and protocols apply to management access on your system:

- **IP** — Internet Protocol. In TCP/IP, the standard for sending the basic unit of data, an IP datagram, through an internet. Defined in RFC 791.
- **TCP/IP** — Transmission Control Protocol/Internet Protocol. Set of protocols developed by the U.S. Defense Department's Advanced Research Projects Agency (ARPA) during the early 1970s. Its intent was to develop ways to connect different kinds of networks and computers. TCP/IP is the protocol of the network, but does not have the functionality that OSI provides.
- **UDP** — User Datagram Protocol. Internet standard protocol in the TCP/IP suite that allows an application on one machine to send a datagram to an application program on another machine. UDP provides no acknowledgments or guaranteed delivery.

- **Telnet** — A standard terminal emulation protocol, supported by almost every TCP/IP implementation, that allows users (clients) to log on to many different hosts (servers) from a single virtual terminal running at their desktop. The connection is always initiated in the client-to-server direction.
- **SNMP** — Simple Network Management Protocol. Standardized method of managing and monitoring network devices on TCP/IP-based internets.

3 PHYSICAL PORT NUMBERING

This chapter contains illustrations that label the ports on the SuperStack® II Switch 3900 and Switch 9300 systems.

This chapter covers the following topics:

- Switch 3900 Port Numbering
- Switch 9300 Port Numbering

The Switch 3900 provides in-band management only. It has no out-of-band management port.

Switch 3900 Port Numbering

The following illustrations (Figure 6, Figure 7, Figure 8, and Figure 9) identify the ports on the front and rear panels of the 24-port and 36-port Switch 3900 systems.

Figure 6 Front Panel of the 24-Port Switch 3900 System

Ethernet Ports (Ports 1 - 24)
24 10/100BASE-TX
RJ-45 Ethernet ports

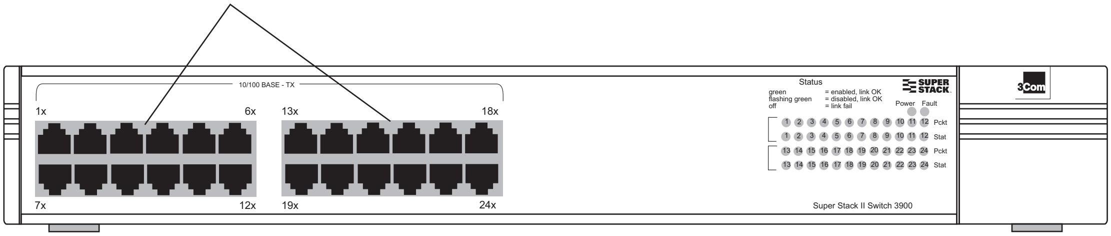

Figure 7 Rear Panel of the 24-Port Switch 3900 System

Optional Gigabit Ethernet Slots (Ports 26 and 27)
Allow you to add additional
Gigabit Ethernet modules

Fixed Gigabit Ethernet Port (Port 25)
Provides a 1000BASE-SX link

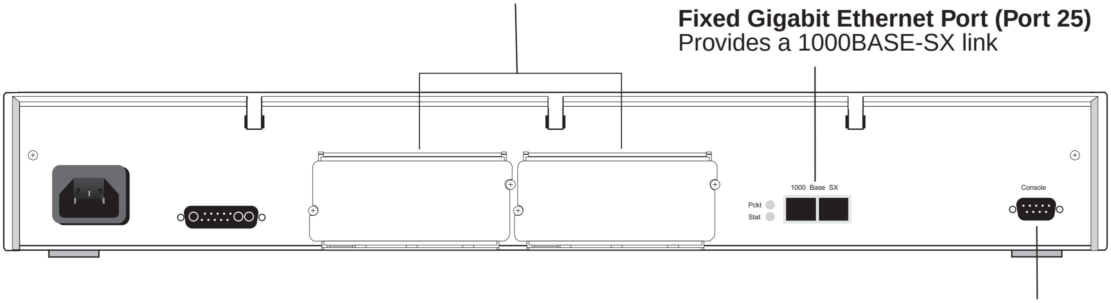

Console Port (Not numbered)
Allows you to connect a workstation or modem to the Switch 3900 for management access

Figure 8 Front Panel of the 36-Port Switch 3900 System

Ethernet Ports (Ports 1 - 36)
36 10/100BASE-TX
RJ-45 Ethernet ports

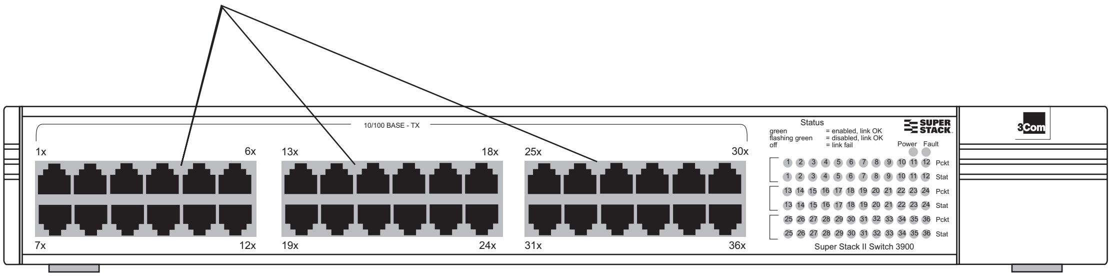

Figure 9 Rear Panel of the 36-Port Switch 3900 System

Optional Gigabit Ethernet Slots (Ports 38 and 39)
Allow you to add additional
Gigabit Ethernet modules

Fixed Gigabit Ethernet Port (Port 37)
Provides a 1000BASE-SX link

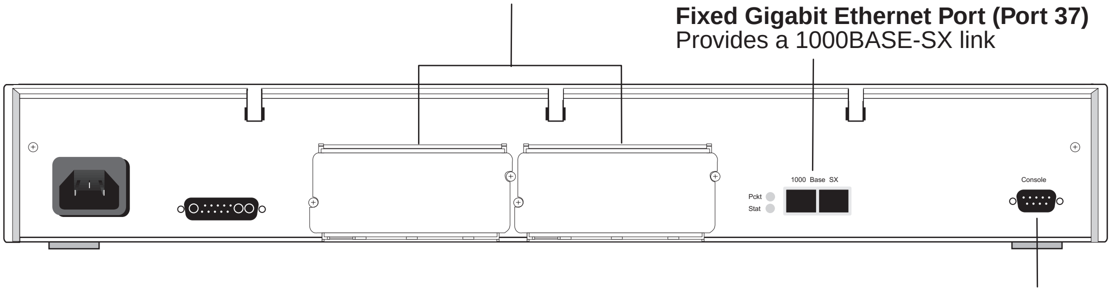

Console Port (Not numbered)
Allows you to connect a workstation or modem to
the Switch 3900 for management access

Switch 9300 Port Numbering

The following illustrations (Figure 10 and Figure 11) label the ports on the front panels of the Switch 9300 SX and Switch 9300 SX/LX systems.

Figure 10 Front Panel Ports of the Switch 9300 SX System

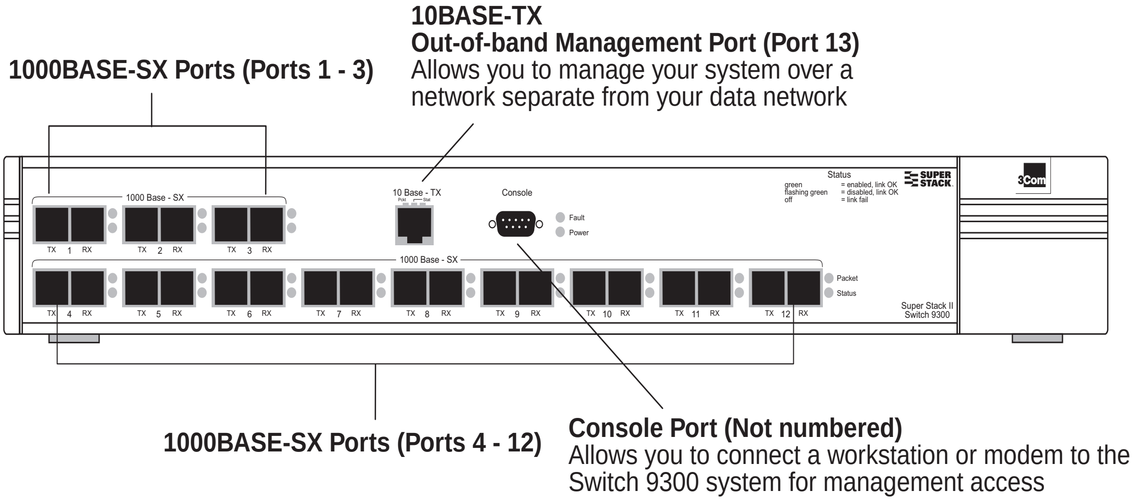

1000BASE-SX Ports (Ports 1 - 3)

10BASE-TX
Out-of-band Management Port (Port 13)
Allows you to manage your system over a network separate from your data network

1000BASE-SX Ports (Ports 4 - 12)

Console Port (Not numbered)
Allows you to connect a workstation or modem to the Switch 9300 system for management access

Figure 11 Front Panel Ports of the Switch 9300 SX/LX System

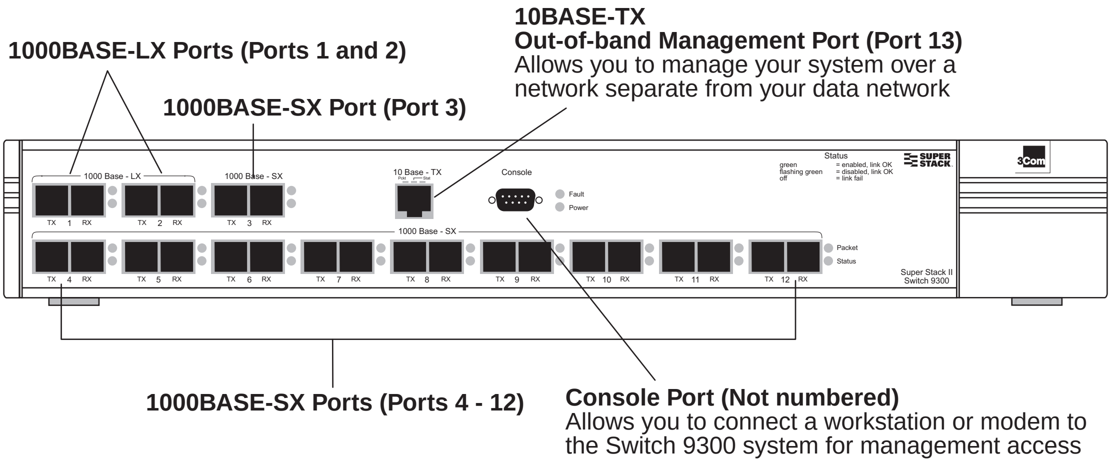

1000BASE-LX Ports (Ports 1 and 2)

1000BASE-SX Port (Port 3)

10BASE-TX
Out-of-band Management Port (Port 13)
Allows you to manage your system over a network separate from your data network

1000BASE-SX Ports (Ports 4 - 12)

Console Port (Not numbered)
Allows you to connect a workstation or modem to the Switch 9300 system for management access

4 SYSTEM PARAMETERS

This chapter provides guidelines and other information about the system parameters that you can configure.

This chapter covers these topics:

- System Parameters Overview
- Key Concepts
- Key Guidelines for Implementation
- Security
- Software Update
- nvData Operations
- Simple Network Time Protocol (SNTP)
- Standards, Protocols, and Related Reading

You can manage system parameters in either of these ways:

- *From the* `system menu` *on the Administration Console. See the* Command Reference Guide*.*
- *From the System folder of the Web Management software. See the* Web Management User Guide*.*

System Parameters Overview

On the Administration Console, you use the `system` menu to set or modify values for system parameters or functions. For many of these parameters, you can also use the configuration forms in the System folder of the Web Management suite of applications software.

Features

You can set or modify the values for when you perform the following tasks:

- Display your system's current configuration
- Take a snapshot of your system's current system configuration and status
- Create and modify passwords
- Create and maintain a statistics baseline

  See Chapter 13 for details.
- Set and administer your system's serial port baud rates

  See Chapter 2 for more details.
- Modify your system's date and time

See the *Command Reference Guide* for descriptions of the commands that you use to set and modify system parameters.

You can also set options for the following, as discussed in these sections later in this chapter:

- Security
- Software Update
- nvData Operations
- Simple Network Time Protocol (SNTP)

Benefits

Using the options on the `system` menu:

- Provides an easy method for setting and modifying system parameters.
- Provides added security by limiting IP and Web Management access to your system.
- Decreases the time and cost of modifying your system configuration. You do not need to make frequent changes from the same source and then reboot your system to apply the changes.
- Reduces the cost of software upgrades by providing an easier process for remote upgrade operations.
- Provides an easy method for changing your system time, changing time zones, and resetting for daylight savings time through SNTP.

Key Concepts

Review these terms and key concepts for system parameters.

- **Save** — Use this option on the `nvData` menu to save nvData to a file on a remote system.
- **Restore** — Use this option on the `nvData` menu to restore data from a file on a network host.
- **Examine** — Use this option on the `nvData` menu to examine a previously saved nvData file header.
- **Simple Network Time Protocol (SNTP)** — SNTP is an adaptation of the Network Time Protocol (NTP). NTP is used to synchronize computer clocks in the global Internet. For more detailed information on NTP and how it is used in your system, see “Simple Network Time Protocol (SNTP)” later in this chapter.
- **Trusted IP Client** — One or more clients that you can allow to have management access to your system. You can configure up to 5 IP addresses or 5 subnet works on this access list.

Key Guidelines for Implementation

This section briefly explains how to set and modify the values for system parameters that you can set.

The system sets most of the parameter values during power-on. To set parameters that are not defined by the system or to modify predefined values, use one of the following methods:

- The `system` menu at the top level of the Administration Console
- The system folder in the Web Management software

Here are the basic steps for setting or modifying system parameter values:

1. Access the menu or form that governs a system parameter.
2. Open the appropriate system-level function.
3. Specify a value.

Security

You can now limit IP management access to your system through the Administration Console or the Web Management software.

- On the Administration Console, you can limit IP management access through the `system console security` menu.
- On the Web Management software, use a security option in the WebManage folder on the Web console.

To limit IP management access, you can use the `system console security` parameter to configure up to 5 IP addresses or 5 subnets, called *trusted IP clients*. If an IP address or subnet is not on the trusted IP client list, that IP address or subnet cannot be used to access the system using the Web Management software, the Administration Console, or SNMP.

Unless you configure trusted IP clients, a user with the appropriate password at a remote device can access the system.

Security Options

To configure trusted IP clients from the Administration Console, you can use the following options:

- **Display** — Shows the IP address and subnet mask of each trusted IP client.
- **Define** — Allows you to supply the IP address and subnet mask of a trusted IP client.
- **Remove** — Removes an IP client from the trusted list.
- **Message** — Controls the message that is displayed when access is denied.
- **Access** — Enables or disables checking for trusted IP clients. By default, checking for trusted IP clients is disabled.

The Web Management software offers these security options:

- **Display** — Displays the trusted IP clients and indicates whether checking for trusted IP clients is enabled or disabled.
- **Configuration** — Allows you to enable or disable checking for trusted IP clients and control the message displayed to a user when access is denied.
- **Add Trusted Client** — Defines a trusted IP client.
- **Remove Trusted Client** — Removes a trusted IP client from the list.

Configuration Procedure

Configure trusted IP clients in this order:

1. Define the trusted IP clients.
2. Display the list of configured trusted IP clients to verify that you have configured them correctly.
3. Enable the checking for trusted IP clients (using the access option on the Administration Console or the Configuration folder in the Web Management software).

CAUTION: *Be careful when you define trusted IP clients. If you specify an incorrect IP address or subnetwork, you can affect your ability to access the system:*

- *For Web Management access, the change is immediate. Therefore, an incorrect IP address or subnet forces you to establish local access via the serialPort.*
- *For Telnet access, the change takes effect at your next login.*

Important Considerations

- If you modify a trusted IP client definition through the Web Management software, the change also affects Telnet and SNMP access to the system. If you modify a trusted IP client definition through Telnet access to the Administration Console, the change also affects SNMP and Web Management access to the system.
- Use the subnet mask to allow all addresses on a particular subnetwork to have trusted access. For example, the IP address 158.101.112.219 with a subnet mask of 255.255.255.0 allows all addresses on the 158.101.112 subnetwork to have trusted access, whereas the same IP address with a subnet mask of 255.255.255.255 only allows only access by 158.101.112.219.
- The trusted IP client information is retained, that is, saved in nvData after a system reboot.

Software Update

You can load a new or updated version of the system software into your system's flash memory with software update option on the System menu through the Administration Console. Depending on your network load, loading software into flash memory can take approximately 10 to 15 minutes to complete.

Important Considerations

Consider the following guidelines *before* you update the system software:

- You can load the system software into flash memory while the system is operating. The system does not have to be powered off.
- Verify that you have defined an IP address on your system.
- To guard against failure during the software upgrade, be sure to save the software to nvData *before* you perform the system software upgrade.

Consider the following points *after* you upgrade the system software:

- If the executable software image that is stored in flash memory is corrupted (for example, if a power failure occurs during the update), contact 3Com Technical Support.
- You can continue to run the old software after you perform a system software upgrade. When it is convenient, reboot your system to use the upgraded software.

nvData Operations

With the nvData features, you can perform these tasks:

- Save and restore your system configuration for backup.
- Examine a saved nvData file header.
- Reset system data to its factory default values, if necessary.

See the *Command Reference Guide* for details on these commands.

Saving nvData

You can use this command remotely to save nvData from your system to a file on another system.

When you save nvdata, the system copies the data, which is located in nonvolatile memory, to a disk file at the location that you specify. The system saves all configurable parameters in nonvolatile memory.

Restoring nvData

Use the nvData restore option on the `system nvData` menu to restore a previous configuration that you have saved to an external file.

Effects and Consequences

Consider the following guidelines before you restore nvData:

- Do not confuse nvData restore with nvData reset. You use `nvData reset` *only* to reset your system configuration values to their factory default settings.
- After you restore nvData, the software presents a proposal for how to restore the data based on the following restoration rules:
  - *Rule 1* **Exact match** — The system IDs and revisions (if applicable) all match between the saved configuration and the configuration of the system on which you are restoring the image.
  - *Rule 2* **System ID mismatch** — System IDs do not match between the saved configuration and the target system. In this case, the system informs you of the mismatch and then prompts you to continue.

  If neither of these rules succeeds, you cannot apply the saved configuration to your system.

- Before you restore a system with mismatched system IDs, consider the following issues that might cause problems after the nvData is restored:
  - Management IP addresses (which are defined in IP interface configurations) are saved as nvData and restored. Restoring management IP addresses can cause duplicate IP address problems. To avoid these problems, change the IP addresses of any defined interfaces before you connect the restored system to the network.
  - Statically configured MAC addresses are saved as nvData. After a successful restore operation, verify that you have no duplicate addresses.

Viewing nvData Information

To verify that you have successfully saved nvData to the file that you specified, view (examine) the header information for that file. The header information shows pertinent product and system information.

Example for the Switch 9300:

```
Select menu option: system nvdata examine
Host IP Address [158.101.100.1]: 158.101.112.34
NV Control file (full pathname): [/tftpboot/mec]
Product ID 4, Product Type 1
System ID 102D00
Saved 1999-05-20T09:45:23 AM Version 2.
```

Resetting nvData

To reset the system settings back to their factory default values use the nvData reset option on the `system nvData` menu.

Effects and Consequences

Consider these points *before* you reset nvdata on your system:

- Resetting nvdata erases all user-configured data, including all passwords, *except* the terminalSpeed and modemSpeed baud settings and the system boot parameters. Therefore, before you reset all affected values, document your configuration so that you can reconfigure the system after you reset it, or save the existing nvData to a file. See “Saving nvData” earlier in this chapter for details.
- You can reset nvData on a system only when it is directly connected through the Administration Console. You cannot reset nvData through a Telnet connection.

Simple Network Time Protocol (SNTP)

This section provides an overview of the Simple Network Time Protocol (SNTP) and implementation guidelines.

SNTP Overview

SNTP is an adaptation of the Network Time Protocol (NTP), which is used to synchronize computer clocks in the global Internet. NTP provides comprehensive mechanisms to access national time and frequency dissemination services, organize the time-synchronization subnetwork, and adjust the local clock in each participating subnetwork peer.

SNTP is a simplified access strategy for servers and clients using NTP version 3. The access paradigm is identical to the User Datagram Protocol (UDP)/TIME protocol, so it is relatively easy to adapt a UDP/TIME client implementation to operate using SNTP. SNTP is designed to operate in a dedicated server configuration with existing NTP and other SNTP clients and servers.

SNTP can operate in either unicast mode (point-to-point), multicast mode (point-to-multipoint), or anycast mode (multipoint-to-point):

- **Unicast client** — Sends a request to a designated server at its unicast address and expects a reply within a specified time frame. From the reply, the unicast client can determine the time and (optional) the round-trip delay and local clock offset relative to the responding server.
- **Multicast server** — Periodically sends a unsolicited message to a designated IP local broadcast address or multicast group address and expects no requests from clients.
- **Anycast client** — Sends a request to a designated IP local broadcast address or multicast group address. One or more anycast servers reply with their individual unicast addresses. The anycast client binds to the first reply that it receives and then continues the operation in unicast mode.

Implementing SNTP

The system software provides an SNTP client, which works with distributed SNTP time servers to synchronize the system clock to international time standards. If you want to use SNTP, you must enable it from the Administration Console, and then select a Daylight Savings Time parameter and a time zone parameter.

Effects and Consequences

The SNTP client operates in unicast mode, which means that the client and server end-system addresses are assigned following the usual IP conventions. Although SNTP in these systems supports one server at a time, you can define up to three servers for backup. Therefore, when the client does not receive a response from the first server within a designated time, it sends a request to the next server on the list.

Standards, Protocols, and Related Reading

See the following Internet Requests for Comments (RFCs) for more information on the protocols that are described in this chapter:

- **RFC 2030** — Simple Network Time Protocol, v4.0, specification
- **RFC 1305** — Network Time Protocol, v3.0, specification, implementation, and analysis
- **RFC 868** — Time Protocol specification

To obtain copies of Internet RFCs and proposed standards, visit the Internet Engineering Task Force (IETF) Web site:

`http://www.ietf.org`

5 ETHERNET

This chapter provides guidelines and other key information about how to implement Ethernet ports.

The chapter covers these topics:

- Ethernet Overview
- Key Concepts
- Key Guidelines for Implementation
- Port Enable and Disable (Port State)
- Port Labels
- Autonegotiation
- Port Mode (Switch 3900)
- Flow Control
- PACE Interactive Access (Switch 3900)
- Port Monitoring (Switch 3900)
- Standards, Protocols, and Related Reading

You can manage Ethernet port features in either of these ways:

- *From the* `ethernet` *menu of the Administration Console. See the* Command Reference Guide.
- *From the Ethernet folder of the Web Management software. See the* Web Management User Guide.

Ethernet Overview

Ethernet is a standardized, packet-based network that supports an exponential hierarchy of three line speeds:

- **10 Mbps** — Ethernet
- **100 Mbps** — Fast Ethernet (Switch 3900)
- **1000 Mbps** — Gigabit Ethernet

All speeds of Ethernet are based on the IEEE 802.3 standard protocol called Carrier Sense Multiple Access with Collision Detection (CSMA/CD), which controls network access. With CSMA/CD, a station that intends to transmit listens for other Ethernet traffic on the network. When the station does not detect network activity, the station transmits.

Features

You can configure these features on Ethernet ports:

- **Port state** — Whether a port is enabled (placed online) or disabled (placed off-line)
- **Port label** — An alphanumeric port identifier
- **Port mode** — Port speed (10 Mbps, 100 Mbps [Switch 3900], or 1000 Mbps) and duplex mode (half-duplex or full-duplex)
- **Autonegotiation** — A feature that allows some ports to automatically identify and negotiate speed and duplex mode with a receiving device
- **Flow control** — A Fast Ethernet and Gigabit Ethernet port mode that pauses and resumes transmissions
- **PACE® Interactive Access (also called PACE Access)** — An algorithm that reduces network jitter, provides reliable timing, and optimizes LAN bandwidth use

In addition, some important Ethernet features depend on which Ethernet equipment you use, how you configure it, and how you connect it:

- **Resilient links** — Protect a network against an individual link or device failure by providing a secondary backup link that is inactive until it is needed
- **Trunking** — Increases bandwidth between switches and servers
- **Trunk Control Message Protocol (TCMP)** — Increases the availability of trunked links by handling physical configuration errors

Benefits

Ethernet, Fast Ethernet, and Gigabit Ethernet technologies allow you to configure and optimize:

- Link bandwidths
- Link availability

Link Bandwidths

As your network needs to support more users and increasingly bandwidth-intensive applications, you can configure Ethernet networks to keep pace with (or exceed) the capacity demands at two locations:

- **To end stations** — Depending on your application needs and network growth, you can migrate workstation connections from shared 10 Mbps to switched 100 Mbps Fast Ethernet. 3Com’s Ethernet network interface cards (NICs) can automatically sense and configure themselves to an upgraded connection speed.
- **Between servers and switches** — Ethernet systems allow you to increase the bandwidth between switches or between servers and switches as your network requires. This increase is accomplished using *trunking* technology (also called *link aggregation*), which works at Layer 2 in the Open Systems Interconnection (OSI) model. For more information about trunking, see Chapter 8.

Link Availability

Ethernet technologies also allow you to design high levels of availability into your network through the use of the following technologies:

- **Resilient links** — If the main link fails, the standby link immediately takes over traffic from the main link. For more information about resilient links, see Chapter 9.
- **Trunking** — The underlying TCMP technology detects and handles physical configuration errors in point-to-point network connections. For more information about trunking, see Chapter 8.

Other Benefits

The hierarchy of Ethernet, Fast Ethernet, and Gigabit Ethernet technologies offers these additional network benefits:

- Easy configuration and expansion of point-to-point links
- Easy implementation and management of workstation moves, adds, changes, and upgrades
- Low-cost expansion of switch-to-switch or switch-to-server bandwidths without having to change devices or cabling
- Improved reliability of network timing and optimization of LAN bandwidth using PACE Interactive Access (Switch 3900)

Key Concepts

These concepts are important to implementing Ethernet:

- **Carrier Sense Multiple Access with Collision Detection (CSMA/CD)** — The standardized Ethernet protocol that controls device access to the network
- **Collision** — When two or more stations attempt to transmit simultaneously
- **Port mode** — An Ethernet port's speed and duplex mode
- **Port speed** — 10 Mbps (Ethernet), 100 Mbps (Fast Ethernet [on the Switch 3900]), and 1000 Mbps (Gigabit Ethernet)
- **Port state** — Whether a port is enabled (placed online) or disabled (placed off-line)
- **Duplex mode** — Whether a port supports one-way (half-duplex) or two-way (full-duplex) transmissions
- **Autonegotiation** — A feature that allows some ports to identify and negotiate speed and duplex mode with a receiving device
- **Flow control** — A Fast Ethernet and Gigabit Ethernet port mode that pauses and resumes transmissions
- **Packet** — The basic unit of communications in Ethernet networks. While packets can vary in size, they have a consistent format.
- **Resilient link**— A technology that protects networks against an individual link or device failure by providing a secondary backup link that is inactive until it is needed
- **Trunking** — A technology that combines multiple Fast Ethernet or Gigabit Ethernet ports into a single high-speed channel, thereby increasing bandwidth among switches and among servers and switches
- **Trunk Control Message Protocol (TCMP)** — A protocol that detects and handles physical configuration errors in a point-to-point or point-to-multipoint configuration, thereby increasing availability of trunked links

- **PACE Interactive Access (also called PACE Access)** — (Switch 3900) An algorithm that controls traffic flow on a point-to-point link with an end station. In a typical half-duplex Ethernet connection, you can never achieve high rates of utilization because of the randomness of collisions. If a switch and end station both try to send data, a collision occurs, forces retransmission, and lowers link utilization.

  PACE Interactive Access enables higher link utilization by altering the switch's *back-off* behavior. Instead of continuing to send data after winning a collision, the switch waits, allows the end station to send a packet, and then retransmits. The result is an interleaving of transmissions between the end station and the switch.

  This feature avoids repetitive collisions and prevents an end station from "capturing" the link. (With conventional Ethernet, a packet collision can cause the last station that transmitted successfully to monopolize Ethernet access and cause delays.)

- **Network areas** — 3Com uses a three-tiered framework to describe the functional areas in a LAN:

  - **Wiring closet** — This area provides connections to user workstations. It also includes downlinks into the data center or campus interconnect area.

  - **Data center** — This area receives connections from wiring closets and campus interconnect areas. Most local server farms reside here.

  - **Campus interconnect** — This area appears as a separate location only in larger networks; smaller networks usually have only wiring closets and data centers. The campus interconnect links campus data centers to each other. It may also include an enterprise server farm and connections to a wide area network.

Ethernet Packet Processing

All packets on an Ethernet network are received promiscuously by an Ethernet port. A port can discard packets for either of the following reasons:

- There is no buffer space available.
- The packet is in error.

Figure 12 shows the order in which packet discard tests are made.

Figure 12 How Packet Processing Affects Ethernet Receive Packet Statistics

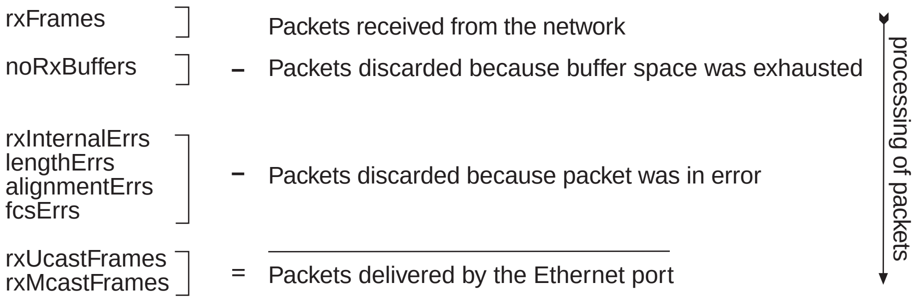
rxFrames
Packets received from the network
noRxBuffers
–
Packets discarded because buffer space was exhausted
rxInternalErrs
lengthErrs
alignmentErrs
fcsErrs
–
Packets discarded because packet was in error
rxUcastFrames
rxMcastFrames
=
Packets delivered by the Ethernet port
processing of packets

Packets also may be delivered directly to an Ethernet port by bridge, router, or management applications. A transmitted packet can be discarded for any of the following reasons:

- The Ethernet port is disabled.
- There is no room on the transmit queue.
- An error occurred during packet transmission.

Figure 13 shows the order in which these discard tests are made.

Figure 13 How Packet Processing Affects Ethernet Transmit Packet Statistics

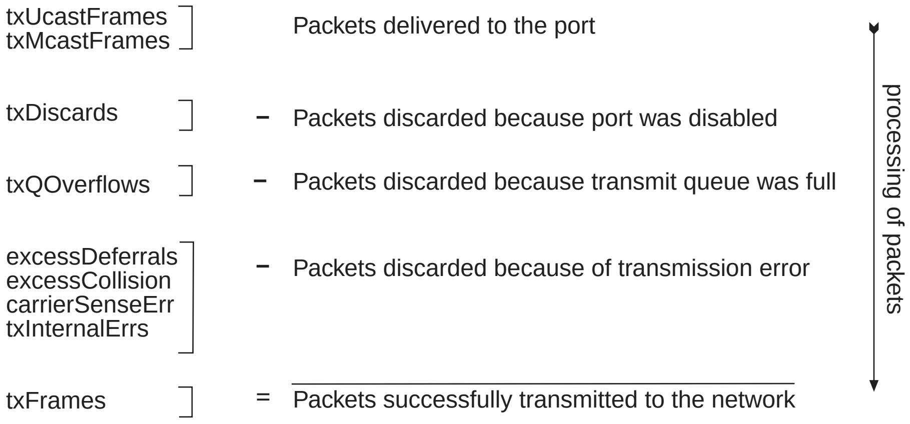
txUcastFrames
txMcastFrames
Packets delivered to the port
txDiscards
– Packets discarded because port was disabled
txQOverflows
– Packets discarded because transmit queue was full
excessDeferrals
excessCollision
carrierSenseErr
txInternalErrs
– Packets discarded because of transmission error
txFrames
= Packets successfully transmitted to the network
processing of packets

Key Guidelines for Implementation

Consider these important factors when you design and configure Ethernet networks.

Link Bandwidths

Recommended link capacities in a network normally depend on the speed requirements of end-user workstations, as shown in Table 7. In areas that may benefit from 1000-Mbps connections, you may be able to substitute trunked Fast Ethernet, subject to the issues raised in Chapter 8.

Table 7 Recommendations for Structuring Bandwidth Across the LAN

	Desktops to Wiring Closet	Wiring Closet to Data Center	Data Center to Campus Interconnect
Mainstream networks	Switched 10 or Shared 10/100	Switched 100	Switched 1000
Power networks	Switched 10/100	Switched 1000	Switched 1000+

Trunking

Consider these important factors when you implement and trunk Fast Ethernet or Gigabit Ethernet links:

- 3Com recommends that you use trunks to increase network availability in the following circumstances:
  - Switch-to-switch connections in the data center and campus interconnect areas
  - Switch-to-server connections in the data center and campus interconnect areas
  - Downlinks from the data center to the campus interconnect
- When multiple links are trunked, it can be difficult to manage and troubleshoot individual port-to-port connections if a connectivity problem occurs. This issue may not be of concern in a server farm room. But if you use trunking extensively between wiring closets and data centers, the large number of connections involved and their distributed nature may make their management and troubleshooting difficult.
- When you work with trunks, be sure that you understand the port numbering for your system. For more information:
  - About port numbering, see Chapter 3.
  - About trunking, see Chapter 8.

Port Enable and Disable (Port State)

You can enable Ethernet ports (place them online) or disable them (place them off-line).

Important Considerations

- Because it stops all network traffic through the port, disabling a port may adversely affect a live network.
- When a port is enabled, the port transmits packets normally. When a port is disabled, the port neither sends nor receives packets.
- The portState is off-line for disabled ports and on-line for enabled ports with an active link.

Port Labels

Port labels serve as useful reference points and as an accurate way for you to identify ports for management applications.

Labeling Ports

- Label Ethernet ports so that you can easily identify the devices that are attached to them (such as LANs, workstations, or servers). For example, you can assign `engineeringserver` as a label.
- The new port label appears in system displays the next time that you display information for that port.
- Port labels can include up to 32 ASCII characters, including the null terminator.

Autonegotiation

This feature enables some ports to identify and negotiate speed and duplex mode with a remote device.

Important Considerations

- In most cases, if autonegotiation does not properly detect the remote port speed, the vendor of the remote device implemented either autonegotiation or a change in port speed in a noncompliant way. If autonegotiation does not properly detect the port speed, you can manually set the port speed and duplex mode.
- Table 8 lists Ethernet port types on your system, whether they support autonegotiation, and which features they negotiate.

Table 8 Port Types and Autonegotiation Attributes

Port Type	Supports Autonegotiation?	Negotiable Attributes	Default Values for Negotiable Attributes
10/100BASE-TX (Switch 3900)	Yes	Port speed	10 Mbps
		Duplex mode	Half-duplex
1000BASE-LX	Yes	Duplex mode	Full-duplex
		Flow control	If autonegotiation is enabled, the system's best effort is `On`
1000BASE-SX	Yes	Duplex mode	Full-duplex
		Flow control	If autonegotiation is enabled, the system's best effort is `On`

- **10/100BASE-TX ports** — Enabling autonegotiation causes both the port speed and duplex mode attributes to be autonegotiated.
- **1000BASE-SX ports** — Both link partners must either enable or disable autonegotiation. As long as autonegotiation is `enabled`, the system's best effort for handling flow control is `On`.
- When you enable autonegotiation, the system ignores your requested portMode information for 10/100BASE-TX ports and your requested flowControl information for 1000BASE-SX ports. When you disable autonegotiation, the system recognizes the requested portMode values for ports that have portMode options and the requested flowControl values for 1000BASE-SX ports.

- On the Switch 3900, you can use the `portMode` option to manually configure or modify the port speed and duplex mode. Use the `flowControl` option to manually configure or modify flow control. (The Switch 9300 ports are 1000-Mbps and full-duplex by default.)
- Autonegotiation is enabled by default on the ports that support it.

Port Mode (Switch 3900)

On the Switch 3900 system, you can change the port speed and duplex mode for the 10/100BASE-TX ports. You cannot change the port speed or duplex mode for Gigabit Ethernet ports.

The Switch 9300 does not support a `portMode` *command. Switch 9300 ports are 1000 Mbps and full-duplex by default.*

Important Considerations

- When you configure duplex mode, configure both sending and receiving ports identically. If the port speeds differ, the link does not come up. If the duplex modes differ, link errors occur.
- Enabling full-duplex mode on a port disables collision detection.
- Autonegotiation must be disabled on a port before a port mode selection can take effect.
- Table 9 lists the duplex port mode options available for 10/100BASE-TX ports.

Table 9 Port Mode Options

Port Type	Duplex Port Mode	Resulting Port Mode	[Default]
10/100BASE-TX	100full	100 Mbps, full-duplex	10half
	100half	100 Mbps, half-duplex	
	10full	10 Mbps, full-duplex	
	10half	10 Mbps, half-duplex	

Flow Control

The flow control mode allows a Fast Ethernet or Gigabit Ethernet port to:

- Decrease the frequency with which it sends packets to a receiving device, if packets are being sent too rapidly.
- Send flow control packets to a sending device, to request that the device slow its speed of transmission.

Important Considerations

- The default setting for flow control is off.
- The system does not count flow control packets in receive or transmit statistics.
- Table 10 lists the effects of flow control options for Fast Ethernet and Gigabit Ethernet ports.

Table 10 Flow Control Options

Flow Control Option	Description	Available on Port Type
on	Port recognizes flow control packets and responds by pausing transmission. The port can generate flow control packets as necessary to slow incoming traffic.	Gigabit Ethernet Fast Ethernet (Switch 3900)
off	Port ignores flow control packets and does not generate flow control packets.	Gigabit Ethernet Fast Ethernet (Switch 3900)
rxOn	Port recognizes flow control packets and responds by halting transmission. The port does not generate flow control packets.	Gigabit Ethernet
txOn	Port ignores flow control packets, but it can generate flow control packets, if necessary.	Gigabit Ethernet

PACE Interactive Access (Switch 3900)

PACE Interactive Access (which is also called PACE Access) prevents excessive network jitter (variation in the timing of packet delivery that can cause garbled sound, jerky images, and delays). PACE technology also provides reliable timing and optimizes LAN bandwidth utilization.

Important Considerations

- Use PACE Interactive Access only on half-duplex Ethernet links between a switch and a single end station. (This setting has no effect on full-duplex links.)
- Do not use PACE Interactive Access when a repeater is connected to a switch port.

Port Monitoring (Switch 3900)

The Ethernet port monitoring feature allows you to:

- Monitor 10/100 Mbps Ethernet ports for excessive collisions, multiple collisions, late collisions, runts, and FCS errors.
- Compare these error counters against user-defined thresholds.
- Disable a port that reaches an error threshold.
- Report the reason that a port is disabled to the Administration Console, MIB databases, and SNMP traps.
- Reenable the port after an initial backoff time interval.
- Continue monitoring.

Standards, Protocols, and Related Reading

The system supports these Ethernet standards:

- **IEEE 802.3** — 10BASE-T Ethernet over unshielded twisted pair (UTP) wiring
- **IEEE 802.3u** — 100BASE-T Fast Ethernet over UTP or fiber-optic cable
- **IEEE 802.3z** — 1000BASE-SX Gigabit Ethernet over multimode fiber-optic cable and 1000BASE-LX Gigabit Ethernet over multimode or single-mode fiber-optic cable
- **IEEE 802.3ad** — Link aggregation (trunking) proposed standard with which 3Com systems will comply

Ethernet Protocols

- **IEEE 802.3** — Carrier Sense Multiple Access with Collision Detection, which controls Ethernet access. A station that intends to transmit listens for network traffic. If it detects none, it transmits.

  If two or more stations transmit at about the same time, their packets experience a *collision,* and the colliding data streams do not reach their destinations. The stations stop transmitting, send an alert to other stations, and wait a random amount of time before trying again.

Media Specifications

Table 11 summarizes the system's Ethernet media options.

Table 11 Ethernet Media Specifications

Type	Speed	Media	Connector	Recommended Distance (max)
10/100BASE-TX	10/100 Mbps	Category 5 UTP	RJ-45	100 m
1000BASE-SX	1000 Mbps	multimode fiber	SC	220 m (62.5 micron @ 160 MHz*km modal bandwidth) 275 m (62.5 micron @ 200 MHz*km modal bandwidth) 500 m 50 micron @ 400 MHz*km modal bandwidth) 550 m (50 micron @ 500 MHz*km modal bandwidth)
1000BASE-LX	1000 Mbps	single-mode fiber	SC	5 km (9 micron) (qualified for up to 10 km)
		multimode fiber	Duplex SC conditioned launch cable	550 m (62.5 and 50 micron @ all modal bandwidths)

Related Reading

For information about Ethernet media options, see the *Getting Started Guide* for your system.

6

Bridge-Wide and Bridge Port Parameters

This chapter provides an overview of bridging concepts and the Spanning Tree Protocol and describes the bridging options and guidelines for your system.

The chapter covers these topics:

- Bridging Overview
- Key Bridging Concepts
- How the Spanning Tree Protocol Works
- Key Guidelines for Implementation
- STP Bridge and Port Parameters
- Frame Processing
- MAC Address Table
- Broadcast and Multicast Limit for Bridge Ports
- Standards, Protocols, and Related Reading

You can manage most bridge-wide and bridge port commands in either of these ways:

- *From the* `bridge` *menu of the Administration Console. See the* Command Reference Guide.
- *From the Bridge folder of the Web Management software. See the* Web Management User Guide.

Bridging Overview

A bridge interconnects two or more LANs and allows them to communicate as if they were one LAN. Bridges examine incoming frames, make forwarding decisions based on the information that the frames contain, and forward the frames toward the destination. Bridges operate at the Layer 2 data link layer of the Open Systems Interconnection (OSI) reference model. Because bridges operate at this layer, they are not required to examine the upper-layer information.

Your system supports transparent bridging, a form of bridging that attaches two or more LANs, listens promiscuously to every packet that is transmitted, and stores each received packet until the packet can be transmitted on to other LANs.

Your system complies with the requirements that are outlined in the IEEE 802.1D Media Access Control (MAC) Bridges base standard. A compliant bridge must, at minimum:

- Learn source addresses from packets that stations on attached LANs transmitted.
- Age addresses of stations (on attached LANs) that have not transmitted a packet for a prolonged period.
- Store and forward packets from one LAN to another.
- Use the Spanning Tree Protocol for loop detection.

Benefits

Bridges provide the following benefits:

- Bridges extend the effective length of a LAN, allowing you to attach distant stations that could not otherwise be connected.
- Bridges can provide a level of separation that prevents some potential damaging errors or undesirable packets from spreading or multiplying on the network.
- Because bridges only forward a percentage of total traffic received, they diminish the traffic that devices on connected segments experience and increase available bandwidth.
- Bridges allow a larger number of devices to communicate than a single LAN can support.

Features

Your system supports several features that are closely related to the bridging process and are therefore categorized under `bridge` on the system interface.

The following bridging topics are covered in this chapter:

- **Spanning Tree Protocol (STP)** — You can configure bridge-wide and bridge port settings to calculate a network topology that reflects a single, loop-free path between any two devices.
- **Multicast and broadcast limits** — You can assign per-port multicast threshold values to limit the per-second forwarding rate of incoming broadcast and multicast traffic from the segment that is attached to that port.

The following bridging topics are covered in other chapters:

- **Class of Service (CoS)** — Your system can process packets through two priority queues. You assign each of the eight priority levels specified in the IEEE 802.1p standard to one of the two queues. For more information, see Chapter 7.
- **Multicast filtering with IGMP** — By understanding the Internet Group Management Protocol (IGMP), your system can direct IP multicast packets only to the ports that require them, instead of flooding to all ports. This process conserves bandwidth at the edge of the network. For more information, see Chapter 12.
- **Virtual LANs (VLANs)** — A VLAN is a logical grouping methodology that allows dispersed users to communicate as if they were physically connected to the same LAN (broadcast domain). For more information about VLANs, see Chapter 10.
- **Trunking** — You can configure your system to aggregate multiple network links into a single point-to-point trunk to increase bandwidth and redundancy without replacing cabling. For more information about trunking, see Chapter 8.
- **Resilient links** — Resilient links protect your network against an individual link or device failure by providing a secondary backup link that is inactive until needed. For more information about resilient links, see Chapter 9.

Key Bridging Concepts

Before you configure bridge-wide or bridge port parameters, review the following key concepts.

Learning Addresses

Bridges *learn* addresses so that they can determine which packets to forward from one bridge port to another. A bridge learns addresses by processing the network traffic that it receives. For a bridge to learn the address of a station on the network (a *source address*), that station must transmit a packet. Addresses that are learned are called *dynamic addresses*.

Each bridge maintains a table, called the *address table*, which lists each learned address and associates it with a port. (The address table also lists manually configured addresses called *static addresses*.)

The system can store up to 16K addresses in its address table.

Aging Addresses

A dynamic address remains in the bridge address table as long as the station to which it relates regularly transmits packets through the bridge. If the station does not transmit within a specified period of time, the address is *aged out* (deleted) from the address table.

Address aging ensures that, if a station moves to a different segment on the network, packets are no longer forwarded to the station's former location. Address aging is necessary because a bridge can learn only a finite number of addresses.

Forwarding, Filtering, and Flooding Packets

A bridge filters, floods, or forwards packets by comparing:

- The packet's destination address to the source addresses in the bridge's address table.
- The destination bridge port (if known) to the port on which the packet was received.

The bridge compares the destination address to the addresses in the address table and does one of the following:

- *If the destination address is known* to the bridge, the bridge identifies the port on which the destination address is located.
  - If the destination bridge port is *different* from the bridge port on which the packet was received, the bridge forwards the packet to the destination bridge port.
  - If the destination bridge port is the *same* as the port on which the packet was received, the bridge filters (discards) the packet.
- *If the destination address is not known* to the bridge, the bridge forwards the packet to all active bridge ports other than the bridge port on which the packet was received. This process is called *flooding*.

Spanning Tree Protocol

A bridge maintains connectivity between LANs with assistance from the Spanning Tree Protocol (STP), which is specified in the IEEE 802.1D MAC Bridges standard.

When a bridge attaches to any single LAN with more than one path, this results in a *loop* in the network topology. Because the bridge receives the same packet from multiple ports within a short period of time, a loop can cause a bridge to continually question where the source of a given packet is located. As a result, the bridge forwards and multiplies the same packet continually, which clogs the LAN bandwidth and eventually affects the bridge’s processing capability.

A backup or redundant path is a valuable concept nevertheless. STP balances both concerns by allowing redundant paths to exist but keeps them inactive until they are needed.

STP uses an algorithm which compares values from a few different parameters to determine all possible paths and then map out a loopless network topology which ensures that only one active path exists between every pair of LANs. STP keeps one bridge port active and puts redundant bridge ports in the *blocking* state. A port in the blocking state neither forwards nor receives data packets. See Figure 14.

After STP logically eliminates the redundant paths, the network configuration stabilizes. Thereafter, if one or more of the bridges or communication paths in the stable topology fail, STP recognizes the changed configuration and, within a few seconds, activates redundant links to ensure network connectivity is maintained.

For more detailed information about Spanning Tree, see “How the Spanning Tree Protocol Works” next in this chapter.

Figure 14 Spanning Tree Algorithm Blocks Redundant Links

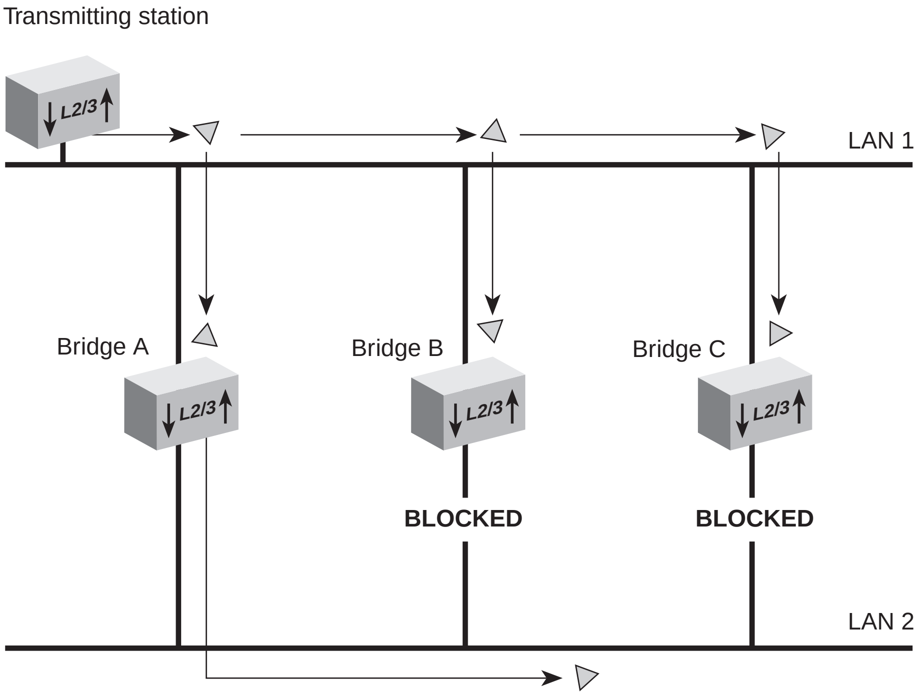

How the Spanning Tree Protocol Works

Using the Spanning Tree Protocol (STP), bridges transmit messages to each other that allow them to calculate the Spanning Tree topology. These messages are special packets called *Configuration Bridge Protocol Data Units* (CBPDUs), or configuration messages.

CBPDUs at Work

CBPDUs do not propagate through the bridge as regular data packets do. Instead, each bridge acts as an end station, receiving and interpreting CBPDUs.

Bridge Hierarchy

The CBPDUs help bridges establish a hierarchy (or a *calling order*) among themselves for the purposes of creating a loopless network.

Based on the information in the CBPDUs, the bridges elect a *root bridge*, which is at the top level of the hierarchy. The bridges then choose the best path on which to transmit information to the root bridge.

The bridges that are chosen as the best path, called *designated bridges*, form the second level of the hierarchy:

- A designated bridge relays network transmissions to the root bridge through its *root port*. Any port that transmits to the root bridge is a root port.
- The designated bridges also have *designated ports* — the ports that are attached to the LANs from which the bridge is receiving information.

Figure 15 shows the hierarchy of the STP bridges and their ports.

Figure 15 Hierarchy of the Root Bridge and the Designated Bridge

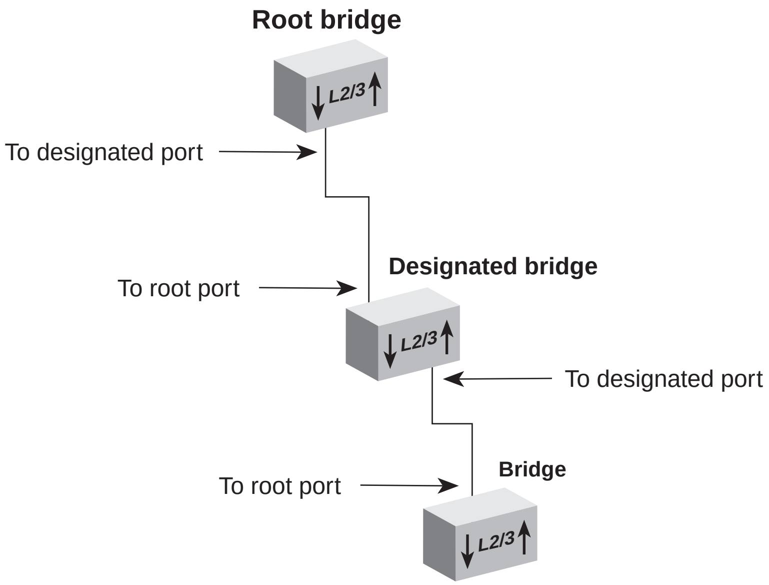

Actions That Result from CBPDU Information

From the information that the CBPDUs provide:

- Bridges elect a single bridge to be the *root bridge*. The root bridge has the lowest bridge ID among all the bridges on the extended network.
- Bridges calculate the best path between themselves and the root bridge.
- Bridges elect as the *designated bridge* on each LAN the bridge with the *least cost path* to the root bridge. The designated bridge forwards packets between that LAN and the path to the root bridge. For this reason, the root bridge is *always* the designated bridge for its attached LANs. The port through which the designated bridge is attached to the LAN is elected the *designated port.*
- Bridges choose a *root port* that gives the best path from themselves to the root bridge.
- Bridges select ports to include in the STP topology. The ports that are selected include the root port plus any designated ports. Data traffic is forwarded to and from ports that have been selected in the STP topology.

Figure 16 shows a bridged network with its STP elements.

Figure 16 STP Root and Designated Bridges and Ports

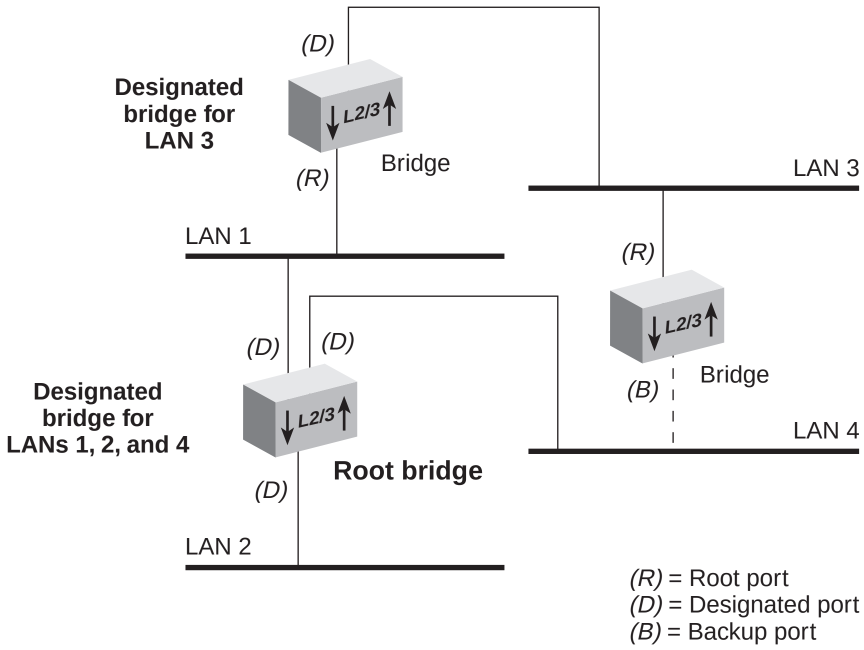

Contents of CBPDUs

Bridges use information in CBPDU to calculate a STP topology. The content of a CBPDU includes:

- **Root ID** — The identification number of the root bridge.
- **Cost** — The cost of the least-cost path to the root from the transmitting bridge. One of the determining factors in cost is the speed of the bridge's network interface; that is, the faster the speed, the lower the cost.
- **Transmitting bridge ID** — The identification of the bridge that transmits the CBPDU, which includes the bridge address and the bridge priority.
- **Port identifier** — Includes the port priority as well as the number of the port from which the transmitting bridge sent the CBPDU.

  The port identifier is used in the STP calculation only if the root IDs, transmitting bridge IDs, and costs (when compared) are equal. In other words, the port identifier is a tiebreaker in which the lowest port identifier takes priority. This identifier is used primarily for selecting the preferred port when two ports of a bridge are attached to the same LAN or when two routes are available from the bridge to the root bridge.

Comparing CBPDUs

Here are three examples that show how the bridge determines the best CBPDU. In every case, the root ID is the most important determining factor. If the root ID fields are equal, then the cost is compared. The last determining factor is the transmitting bridge ID. If the CBPDUs all have the same root ID, cost, and transmitting bridge ID, then the port identifier is used as a tiebreaker.

Example 1. Root ID is lower for Message 1. The bridge saves Message 1.

Message 1			**Message 2**		
root ID	**cost**	**transmitter**	**root ID**	**cost**	**transmitter**
12	15	35	31	12	32

Example 2. Root ID is the same for Message 1 and Message 2, but cost is lower in Message 1. The bridge saves Message 1.

Message 1			Message 2		
root ID	**cost**	**transmitter**	**root ID**	**cost**	**transmitter**
29	15	80	29	18	38

Example 3. Root ID and cost are the same for Message 1 and Message 2, but the transmitting bridge ID is lower in Message 1. The bridge saves Message 1.

Message 1			Message 2		
root ID	**cost**	**transmitter**	**root ID**	**cost**	**transmitter**
35	80	39	35	80	40

How a Single Bridge Interprets CBPDUs

The following case describes how *a single bridge* interprets CBPDUs and contributes to the Spanning Tree configuration.

1. When Spanning Tree is first started on a network, the bridge acts as if it is the root bridge and transmits a CBPDU from each of its ports with the following information:

   - Its own bridge ID as the root ID (for example, `85`)
   - Zero (0) as the cost (because, for the moment, it is the root bridge)
   - Its own bridge ID as the transmitting ID (for example, `85`)

   Thus, its CBPDU looks like this: `85.0.85`.

2. The bridge receives CBPDUs on each of its ports from all other bridges and saves the *best* CBPDU from each port.

   The bridge determines the best CBPDU by comparing the information in each message that arrives at a particular port to the message that is currently stored at that port. In general, the lower the value of the CBPDU, the *better* it is. When the bridge comes across a better CBPDU than it has stored, it replaces the old message with the new one.

3 From the messages that are received, the bridge identifies the root bridge.

For example, if the bridge receives a CPBDU with the contents 52.0.52, then it assumes that the bridge with ID 52 is the root (because 52 is smaller than 85).

4 Because the bridge now knows the root bridge, it can determine its distance to the root and elect a root port.

It examines CBPDUs from all ports to see which port has received a CBPDU with the smallest cost to the root. This port becomes the root port.

5 Now that the bridge knows the contents of its own CBPDU, it can compare this updated CBPDU with the ones that its other ports received:

- If the bridge's message is better than the ones received on any of its ports, then the bridge assumes that it is the designated bridge for the attached LANs.
- If the bridge receives a better CBPDU on a port than the message it would transmit, it no longer transmits CBPDUs on that LAN. When the algorithm stabilizes, only the designated bridge transmits CBPDUs on that LAN.

How Multiple Bridges Interpret CBPDUs

The previous section looked at how a single bridge reviews CBPDUs and makes decisions. The following examples illustrate how STP determines the topology for an entire network.

Figure 17 and Figure 18 shows the same network topology — six bridges that connect six LANs. The topology is designed with redundant links for backup purposes, which create loops in the extended network. Figure 17 shows the network at the start of the STP topology calculation. Figure 18 shows the network after the STP topology has stabilized.

Figure 17 Starting the Spanning Tree Calculation

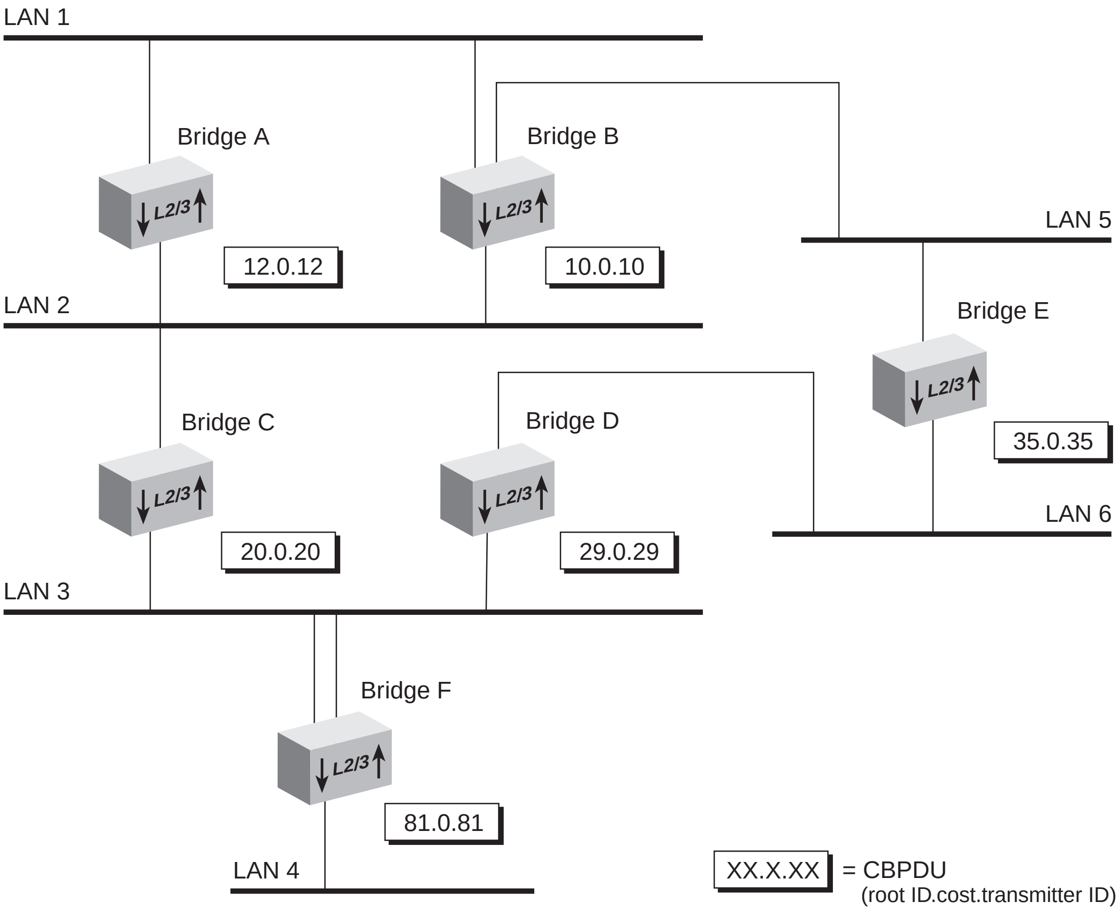

Figure 18 Spanning Tree Topology Calculated

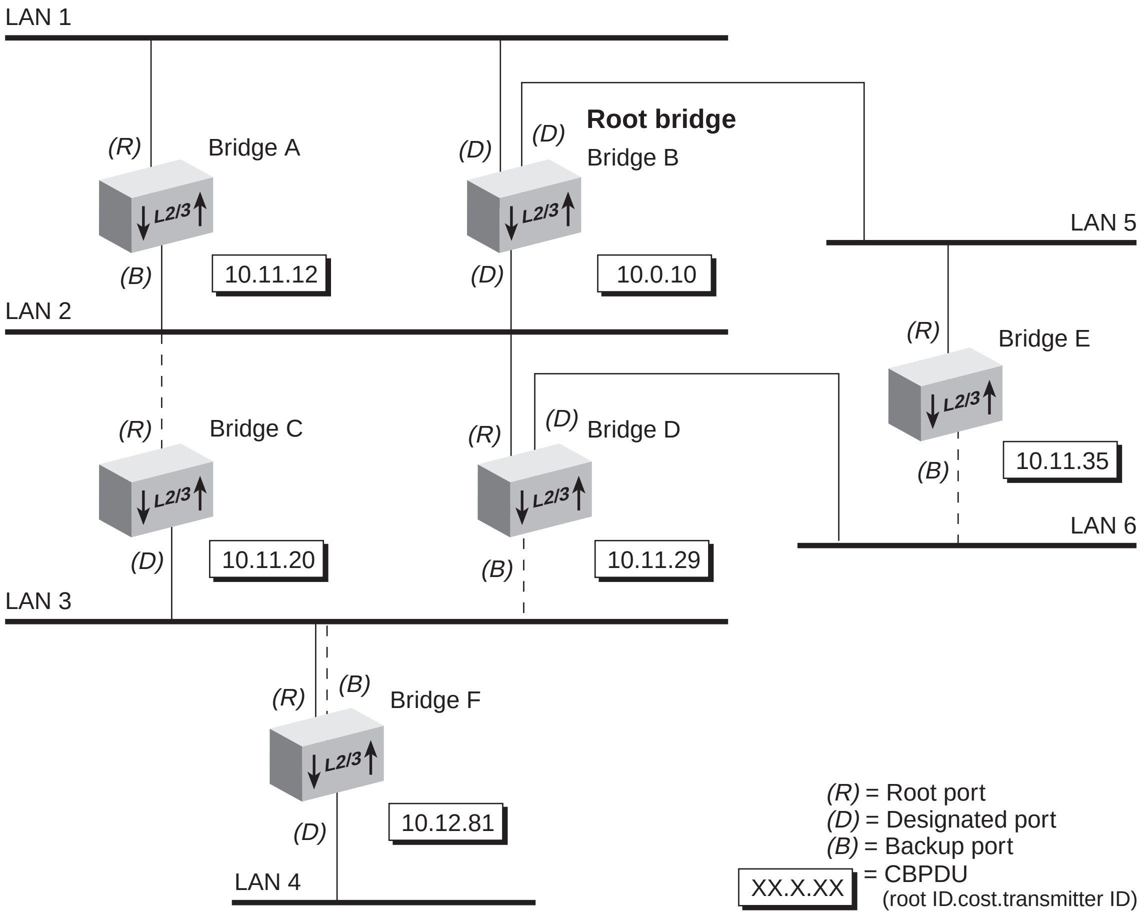

Determining the Root Bridge

The root ID portion of the CBPDU determines which bridge actually becomes the root bridge. In Figure 17, notice how each bridge assumes itself to be the root and transmits a CBPDU that contains its own bridge ID as both the *root ID* and the *transmitting bridge ID*, and zero as the *cost*. In Figure 18, because Bridge B has the lowest root ID of all the bridges, it becomes the root and all other bridges change their root ID to Bridge B's ID (10).

Determining the Root Ports

Next, each bridge (except for the root bridge) must select a root port. To select a root port, each bridge determines the most cost-effective path for packets to travel from each of its ports to the root bridge. The cost depends on:

- The port path cost.
- The root path cost of the designated bridge for the LAN to which this port is attached.

If the bridge has more than one port attachment, the port with the lowest cost becomes the root port, and the other ports become either designated or backup ports. If bridges have redundant links to the same LAN, then the port with the lowest port identifier becomes the root port.

In Figure 18, Bridge F has two links to LAN 3 (through port 1 and port 2). Because the lowest port identifier for Bridge F is port 1, it becomes the root port, and port 2 becomes a backup port to LAN 3.

Determining the Designated Bridge and Designated Ports

For a LAN attached to a single bridge, that bridge is the LAN's designated bridge. For a LAN that is attached to more than one bridge, a designated bridge must be selected from among the attached bridges.

The root bridge functions as the designated bridge for all of its directly attached LANs.

For example, Bridge B, the root bridge in Figure 18, is also the designated bridge for LANs 1, 2, and 5.

A designated bridge must be determined for LANs 3, 4, and 6:

- Because Bridges C, D, and F are all attached to LAN 3, one of them must be the designated bridge for that LAN:
  - The algorithm first compares the root ID of these bridges, which is the same for all.
  - The cost is then compared. Bridge C and Bridge D both have a cost of 11. Bridge F, with a cost of 12, is eliminated as the designated bridge.
  - The transmitting bridge ID is compared between Bridge C and Bridge D. Because Bridge C's ID (20) is smaller than Bridge D's (29), Bridge C becomes the designated bridge for LAN 3.
- The designated bridge for LAN 6 is either Bridge D or Bridge E. Because Bridge D's transmitting bridge ID (29) is lower than Bridge E's (35), Bridge D becomes the designated bridge for that LAN.
- The designated bridge for LAN 4 is Bridge F, the only bridge that is attached to that LAN.

The port that attaches the designated bridge to the LAN determines the designated port. If more than one port is attached to the same LAN, then the port identifier determines the designated port.

Spanning Tree Port States

Because STP determines the network configuration or adjusts it, depending on events that occur, it places bridge ports in one of the following states at all times: listening, learning, forwarding, blocking, or disabled. Table 12 describes these states.

Table 12 Spanning Tree Protocol Port States

Port State	Description
Listening	When STP is configuring, all ports are placed in the listening state. Each port remains in this state until the root bridge is elected. While in the listening state, the bridge continues to run STP and to transmit CBPDUs on the port; however, the bridge discards data packets that are received on that port and does not transmit data packets from that port. The listening state should be long enough for a bridge to hear from all other bridges on the network. After being in the listening state, the bridge ports that are to proceed to the forwarding state go into the learning state. All other bridge ports go into the blocking state.
Learning	The learning state is similar to the listening state except that data packets are received on that port for the purpose of learning which stations are attached to that port. After spending the specified time in this state without receiving information to change the port to the blocking state, the bridge changes the port to the forwarding state. The time that the port spends in each of the listening and learning states is determined by the value of the *forward delay* parameter.
Forwarding	After the port enters the forwarding state, the bridge performs standard bridging functions.
Blocking	When a port is put in the blocking state, the bridge continues to receive CBPDUs on that port (monitoring for network reconfigurations), but it does not transmit them. In addition, the bridge does not receive data packets from the port, learn the locations of station addresses from it, or forward packets onto it.
Disabled	A port is disabled when the STP has been disabled on the port or when the port has failed. In the disabled state, the port does not participate in the Spanning Tree algorithm. The port continues to forward frames only if STP is disabled for the entire bridge and the link is up.

Figure 19 illustrates the factors that cause a port to change from one state to another. The arrows indicate the direction of movement between states. The numbers correspond to the factors that affect the transition.

Figure 19 Factors in Spanning Tree Port State Transitions

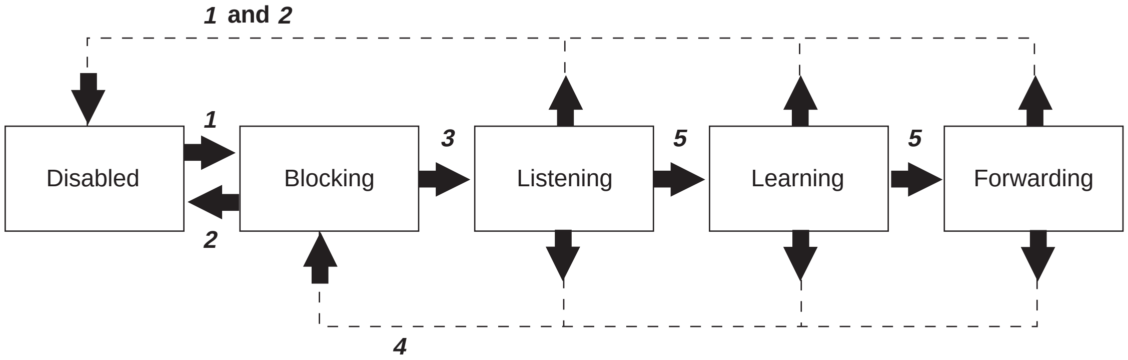

1 Port enabled by either network administrator or initialization

2 Port disabled by either network administrator or failure

3 Spanning Tree algorithm selects port as designated or root

4 Spanning Tree algorithm does not select port as designated or root

5 Forwarding timer (forward delay) expires

As shown in Figure 19, for a port in the blocking state to transition to the listening state, STP must select that port as a designated or root port. After the port enters the listening state, forward delay must expire before the port can transition to the learning state. Then another forward delay period must expire (listening state) before the port can transition to the forwarding state. If you disable a port in the listening, learning, or forwarding state or if port initialization fails, then that port becomes disabled.

Reconfiguring the Bridged Network Topology

STP reconfigures the bridged network topology when any of the following events occur:

- Bridges are added or removed.
- The root bridge fails.
- You change any of the bridging parameters that influence the topology decision.

Resulting Actions

Whenever a designated bridge detects a topology change, it sends a Topology Change Notification Bridge Protocol Data Unit (BPDU) through its root port. This information is eventually relayed to the root bridge.

The root bridge then sets the Topology Change Flag in its CBPDU so that the information is broadcast to all bridges. It transmits this CBPDU for a fixed amount of time to ensure that all bridges are informed of the topology change.

If a port changes from the blocking state to the forwarding state as a result of the topology change, STP sends the topology information to all the ports before that port starts forwarding data. This delay prevents temporary data loops.

When a network reconfiguration occurs, a bridge flushes all dynamic addresses (including STP path information) from its address table. This action ensures that the bridge learns the correct paths and continues to forward packets to the correct LANs.

Key Guidelines for Implementation

Consider the following guidelines when you configure bridging parameters on your system:

- When you disable bridge-wide STP, the bridge cannot participate in the algorithms for loop detection and so forth.
- Table 13 describes the forwarding behavior of a port based on its bridge and port STP states:

Table 13 Port Forwarding Behavior Depends on Bridge and Port STP States

Bridge STP State	Port STP State	Port Participates in STP?	Port Forwards Frames?
Disabled	Disabled	No	Yes, if link state is up.
	Enabled	No	Yes, if link state is up.
	Removed	No	Yes, if link state is up.
Enabled	Disabled	No	No
	Enabled	Yes	Determined by STP, provided that the link state is up.
	Removed	No	Yes, if link state is up.

- When STP is removed from the port but is enabled for the bridge, the port is invisible to STP but can forward frames. Removing the port from STP is useful if you have an edge switch device that is connected to end stations (such as PCs) that are frequently turned on and off.
- The port numbering shown for your ports is always sequential. See Chapter 3 for more information about port numbering.
- When you are prompted to select ports, specify the ? option to see a matrix of information about your bridge ports, including a Selection column, a Port column, and a Label column.
  - *Without trunking,* the Selection and Port columns contain the same port numbers, which indicates that you can select each port.
  - *With trunking,* the Selection column indicates that you can select the anchor port (lowest-numbered port) in the trunk, and the Port column shows each port that is associated with the trunk. The Label column contains the trunk name, if you have assigned one.

- If you want to specify a multicast limit for a trunk, be sure to apply it to the trunk's anchor port (lowest-numbered port) only. However, be aware that the multicast limit applies to *each link* in the trunk (that is, it is not an aggregate).
- You can enable STP with trunks. You may find it useful to configure a backup trunk that STP places in the blocking state. See Chapter 8 for more information about trunking.
- If you want to define one or more resilient link pairs on the system, STP cannot be enabled.
  - If STP is enabled and you define a resilient link pair, the system rejects it toward the end of the definition process.
  - If you have one or more resilient links defined (STP is disabled) and you try to enable STP, the system rejects this request. You cannot enable STP until you remove all resilient link pairs.
- If you have specified `allClosed` as the VLAN mode and you want to administer bridge port address options, you must specify the correct VLAN interface index because each VLAN in `allClosed` mode operates with a unique address table.

STP Bridge and Port Parameters

On a bridge-wide basis, you can enable or disable the Spanning Tree Protocol (STP) and set STP bridge parameters. On a bridge-port basis, you can enable, disable, or remove STP and set STP bridge port parameters.

Administering Bridge-wide STP Parameters

You can modify the following STP bridge-wide parameters:

- **STP state on a bridge** — When STP is disabled on the system, the bridge does not participate in the Spanning Tree algorithm and other STP settings have no effect on bridge operation or network topology calculations. If other devices on the network are running STP, then these packets are bridged.
- **Bridge priority** — The *bridge priority* influences the choice of the root bridge and the designated bridge. The *lower* the bridge's priority number, the *more likely* it is that the bridge is chosen as the root bridge or a designated bridge. The bridge priority value (`0x0-0xffff`) is appended as the most significant portion of a bridge identifier (for example: `8000 00803e003dc0`). It is a 2-octet value.
- **Bridge maximum age** — The *bridge maximum age* determines when the stored configuration message information is judged to be too old and is discarded from the bridge's memory. If the value is too small, then STP may reconfigure the topology too often, causing temporary loss of connectivity in the network. If the value is too large, the network may take longer than necessary to adjust to a new STP configuration after a topology change such as the restarting of a bridge. A conservative value assumes a delay variance of 2 seconds per hop. The recommended value is 20 seconds.

  The value that you set for bridge maximum age is only used if the system is selected as the root bridge. Otherwise, the system uses the value that is assigned to it by the root bridge.
- **Bridge hello time** — *Hello time* is the period between the configuration messages that a root bridge generates. If the probability of losing configuration messages is high, shorten the time to make the protocol more robust. Alternatively, to lower the overhead of the algorithm, lengthen the time. The recommended value is 2 seconds.

  The value that you set for bridge hello time is only used if the system is selected as the root bridge. Otherwise, the system uses the value that is assigned to it by the root bridge.

- **Bridge forward delay** — The *forward delay* value specifies the amount of time that a bridge spends in each of the listening and the learning states. This value temporarily prevents a bridge from starting to forward data packets to and from a link until news of a topology change has spread to all parts of a bridged network. The delay gives enough time to turn off to all links that need to be turned off in the new topology before new links are turned on.

  Setting the value too low can result in temporary loops while the Spanning Tree algorithm reconfigures the topology. Setting the value too high can lead to a longer wait while the STP reconfigures the topology. The recommended value is 15 seconds.

  The value that you set for bridge forward delay is only used if the system is selected as the root bridge. Otherwise, the system uses the value that is assigned to it by the root bridge.

- **STP group address** — The STP group address is a single address to which a bridge listens when it receives STP information. Each bridge on the network sends STP packets to the group address. Every bridge on the network receives STP packets that were sent to the group address, regardless of which bridge sent the packets.

  You may run separate STP domains in your network by configuring different STP group addresses. A bridge only acts on STP frames that are sent to the group address for which it is configured. Frames with a different group address are ignored.

  Because there is no industry standard about group address, bridges from different vendors may respond to different group addresses. If STP does not seem to be working in a mixed-vendor environment, verify that the group addresses are identical on all devices.

Administering STP Parameters on Bridge Ports

You can enable, disable, or remove STP for one or more ports on the system. This setting affects the operation of a port only if STP is enabled for the system. You can also set the following STP port parameters:

- **Port path cost** — The STP algorithm adds the path cost to the root cost field in a configuration message that is received on this port. The system uses this value to determine the path cost to the root through this port. You can set this value individually on each port. The range is 1 through 65535.

  A higher path cost value makes the LAN that is reached through the port more likely to be low in the Spanning Tree topology. The lower the LAN is in the topology, the less through traffic it carries. For this reason, assign a high path cost to a LAN that has a lower bandwidth or to one on which you want to minimize traffic.

- **Port priority** — The STP port priority influences the choice of port when the bridge has two ports connected to the same LAN, which creates a loop. The port with the lowest port priority is selected by STP. Port priority is a 1-octet value. The range for the port priority is 0x0 through 0xff hexadecimal. The default is `0x80`.

Frame Processing

All frames that are received on a physical interface and not explicitly directed to the system or discarded are delivered to the corresponding bridge port. The bridge port either forwards each frame to another bridge port or discards it.

The system can discard an incoming frame for the following reasons:

- The destination station is on the same segment as the source station.
- The receive bridge port is blocked.
- There is a problem with the frame.

  The physical interface does not deliver frames with errors to the bridge port. Thus, the `rxFrames` fields in the Ethernet statistics display and bridge statistics display often report different values — that is, the latter value is lower because it does not count frames in error.

A frame that is forwarded from a physical interface to a bridge port is then transmitted onto a physical interface unless it is discarded. The system can discard a frame at this point for the following reasons:

- The transmit bridge port is blocked.
- The frame is too large for the corresponding physical interface.

MAC Address Table

Your system includes several options for managing MAC addresses on bridge ports. The system recognizes two different kinds of addresses:

- **Static MAC addresses** — Addresses that you manually add to the bridge address table using menu options. These addresses never age; you must add and remove them manually.
- **Dynamic MAC addresses** — Addresses that the bridge learns by receiving and processing packets and ages. In the bridge address table, each dynamic address is associated with a specific port and is assigned an age so that it can be cleared from the table if the station is inactive.

Your system can store up to 16K addresses.

Aging Time

The bridge aging time is the maximum period (in seconds) that dynamically learned forwarding information (addresses) is held in the bridge address table before it is aged out. Use this parameter to configure the system to age addresses in a timely manner, without increasing packet flooding beyond acceptable levels.

Important Considerations

- All dynamic addresses are flushed from the bridge address table whenever you cycle power to the system or reboot the system. All dynamic addresses are also flushed when STP reconfigures the topology. Both dynamic and static addresses are flushed when you reset nonvolatile data.
- If you have multiple ports associated with a trunk, the addresses that are defined for the anchor port apply to all ports in the trunk.
- You can remove individual MAC addresses from selected ports. Typically, this action is only applied to static addresses because the system can quickly relearn dynamic addresses that you remove.
- A statically configured address is never aged and it cannot be learned dynamically on a different port until it is removed from the port on which it is configured.

- The number of static MAC addresses that you can configure depends on the availability of system resources.
- If a station whose address is statically-configured on one port is moved to a different port, the system discards all received packets as a security measure and increments a statistical counter. (From the `bridge display` of the Administration Console, see the `rxSecurityDiscs` field. From the Bridge Display option on the Web Management interface, see the Received Security Discards column.)

Broadcast and Multicast Limit for Bridge Ports

You can assign a rate limit to any bridge port in the system to control the per-second forwarding rate of incoming multicast and broadcast packets. If the limit is reached, all remaining packets received in that second of time are dropped. You can configure the limit to affect either broadcast packets only or both broadcast and multicast packets. This feature is useful for suppressing potential multicast or broadcast storms.

Important Considerations

When you set a limit, consider the following:

- A value of zero means that there is no limit set on the port. The system default is zero on all ports.
- You specify the limit in frames per second. To determine an appropriate limit, measure the normal amount of broadcast or multicast traffic on your network.
- If you want to specify a limit for a trunk, you only need to specify the trunk's anchor port (lowest-numbered port) when you configure the limit for the entire trunk. However, be aware that the multicast limit operates on *each link* in the trunk.
- If you have IP multicast application traffic on your network, be sure that any limits that you configure do not constrain these traffic flows.

Standards, Protocols, and Related Reading

Refer to the following standard for more information about the bridging methodology described in this chapter:

- *IEEE 802.1D: Media Access Control (MAC) Bridges*

  This base standard specifies requirements for transparent bridging. To obtain a copy of this standard, register for an on-line subscription at the Institute of Electrical and Electronics Engineers (IEEE) Web site:

  `http://www.ieee.org`

7

CLASS OF SERVICE (COS)

The *IEEE 802.1D Media Access Control (MAC) Bridges* base standard has been amended in recent years to include various supplements. One such supplement standard is *IEEE 802.1p: Traffic Class Expediting and Dynamic Multicast Filtering*. This chapter describes the traffic prioritization portion of this standard and how it is implemented in your system.

This chapter covers these topics:

- Overview
- Key Concepts
- CoS in Your System
  - CoS Architecture
  - Configuring Priority Levels
  - Configuring a Rate Limit on Queue 1
  - Handling Tagged and Untagged Packets
- Standards, Protocols, and Related Reading

You can administer Class of Service (CoS) commands from the `bridge cos` *menu of the Administration Console. See the* Command Reference Guide *for details.*

Overview

Many network technologies, such as Ethernet and Fiber Distributed Data Interface (FDDI), have no inherent ability to distinguish between different types of traffic such as data, voice, and video, or even perhaps between different data applications. Thus, all traffic competes for the same bandwidth and is processed in a single queue by network devices. This approach to network service is described as "best effort" because there is no way to prioritize certain traffic ahead of other traffic.

If the network load is high and network devices become congested, certain bandwidth-intensive applications may receive a poor *quality of service (QoS)*. A jittery video conference display that does not reflect real-time movement or a crisp picture is an example of poor quality of service.

To overcome this limitation in Ethernet and FDDI, switch and router vendors developed a variety of QoS-oriented features that work at higher levels in the Open Systems Interconnection (OSI) model. Users can configure these features to better control how different types of traffic are processed and forwarded through the system and ultimately the network as whole. QoS techniques are designed to address the different latency and throughput needs of time-sensitive applications, as well as to address the desire to prioritize business-critical information over non-critical information.

While QoS features clearly benefit a network with bandwidth constraints, they can also add complexity and cost into network equipment and administration activities. Thus, the practical aim of the IEEE 802.1p standard is to outline a simplified version or subset of QoS techniques that preserves the high speed, low cost nature of traditional LAN bridging. Because the IEEE 802.1p standard addresses queuing and prioritization based on a numeric traffic class but it does not address bandwidth reservation or other approaches to QoS, the approach is often distinguished with the term *Class of Service (CoS)*.

Key Concepts

Before you configure CoS options in your system, review the following key concepts.

Basic Elements of the Standard

The two basic elements of the IEEE 802.1p standard are:

- Multiple processing queues in devices.

  The standard does not require a specific number of queues, but rather how different types of traffic could be allocated across up to eight queues.

- Priority levels carried in packets.

  Packets that include priority levels are processed through the device queues in a way that is configured by the network administrator. The standard identifies eight different priority levels using numbers 0 through 7.

  Table 14 outlines the different types of traffic that the standards body envisioned carrying different priority levels. However, you can apply the eight numbers in any way you choose to identify your network application traffic.

Table 14 Priority Levels and Traffic Types Envisioned by Standards Committee

Priority Level	Traffic Type
1	Background
2	(Spare)
0 (default)	Best Effort
3	Excellent Effort
4	Controlled Load
5	Video, less than 100 milliseconds latency and jitter
6	Voice, less than 10 milliseconds latency and jitter
7	Network Control

Format of Prioritized Packets

Priority level information can only be carried inside packets that are formatted according to the IEEE 802.1Q standard; such packets carry an extra 2 octets of data called a *tag*. The priority level information occupies 3 bits of this tag and VLAN information occupies 12 bits.

The following definitions summarize the difference between tagged and untagged packets and clarify two types of tagged packets:

- **Untagged packet** — Does not include an IEEE 802.1Q tag.
- **Tagged packet** — Includes an IEEE 802.1Q tag. There are two types:
  - **Priority-tagged packet** — Carries priority level information but no VLAN information.
  - **VLAN-tagged packet** — Carries priority level information and VLAN information.

Queues and Priority Levels

Compliance with the IEEE 802.1p standard means that a device must recognize eight priority levels (0 through 7), however the number of queues in a given device can vary. (Eight queues are not required.)

When there are fewer than eight device queues, a packet's priority level does not always indicate how it will be processed relative to other packets, because more than one priority level will be assigned to at least one of the queues. When multiple priority levels are assigned to the same queue, all packets in that queue are processed in the same manner, regardless of their priority level.

The characteristics of a given queue as well as overall product design determine how the packets in that queue are processed relative to packets in other queues. The device vendor identifies these characteristics.

CoS in Your System

Using the Administration Console, you can:

- Enable or disable CoS (the setting affects all ports), which changes the architecture from one to two queues per port.
- Modify how the eight priority levels are assigned between the two queues.

  The priority levels are initially assigned according to recommendations in the IEEE 802.1p standard. To modify queue assignments, see "Configuring Priority Levels" later in this section.
- Set a rate limit on the high priority queue.

  See "Configuring a Rate Limit on Queue 1" later in this section.
- Display a summary CoS configuration.

CoS Architecture

When CoS is enabled, your system uses two CoS queues per port:

- Queue 1 is always the high priority queue.
  - Each Fast Ethernet port has a queue-specific buffer of 64 KB.
  - Each Gigabit Ethernet port has a queue-specific buffer of 128 KB.
  - You can affect the flow of queue 1 traffic by configuring a rate limit. See "Configuring a Rate Limit on Queue 1" later in this chapter.
- Queue 2 is always the low priority queue.
  - Each Fast Ethernet port has a queue-specific buffer of 256 KB.
  - Each Gigabit Ethernet port has queue-specific buffer of 512 KB.

When CoS is disabled, the high priority queues and associated buffers are shut off; all traffic flows through the low priority queues.

CoS settings are stored in non-volatile memory. Thus, in the event of a power cycle or reboot, user-configured settings are retained.

Important Considerations

- In nonblocking situations, CoS settings have no impact on traffic flow through the system.
- In blocking situations, queue 1 (high priority) traffic on a given port is processed ahead of queue 2 traffic on that same port. Traffic in queue 2 on that port is either delayed (buffered) or dropped (if buffers become full) as needed to allow the queue 1 traffic to be forwarded.
- Queue 1 traffic on a given port is not necessarily processed ahead of queue 2 traffic from other ports. This is because the switch selects traffic from each port in an approximate round-robin fashion.

Configuring Priority Levels

By default, CoS is enabled with priorities 4, 5, 6, and 7 assigned to queue 1 and priorities 0, 1, 2, and 3 assigned to queue 2. This arrangement conforms with IEEE recommendations, but you can change it at any time.

Although there are two physical queues per port, the priority levels (traffic classes) that you assign to each queue actually apply to all ports in the system.

When you assign one or more priority levels to one of the queues, the system automatically assigns the remaining priority levels to the other queue.

If CoS is disabled, you can still modify the priority assignments to each queue; they simply do not effect traffic until you enable CoS.

Configuring a Rate Limit on Queue 1

You can configure a rate limit for queue 1. The rate limit is configured as a percentage of the number of packets received on each port.

The percentage refers to the number of packets that are processed from queue 1 out of every 8 packets received on the port. This *n of 8 packets* formula means that, in real terms, there are eight supported rate limit percentages: 12.5, 25, 37.5, 50, 62.5, 75, 87.5, and 100. The rate limit operates as a threshold, not as a bandwidth reservation technique.

Considering that the system does not accept decimal values, you must enter a whole number in the range that corresponds to one of the eight percentages. Depending on the number you enter, the system rounds down to the nearest *n of 8* value, although the summary display retains the number that you enter.

For example, if you enter any whole number between 88 and 99 as the rate limit, the operating rate limit percentage is 87.5; that is, for every 8 packets received on a given port, 7 packets are selected from queue 1 and 1 packet is selected from queue 2. Table 15 provides a reference chart:

Table 15 Implementation Guidelines for the CoS Rate Limit

Operating Percentage	Range of Possible Values	For every 8 packets received on the port, this number of packets is selected from queue 1
12.5	1–24	1
25	25–37	2
37.5	38–49	3
50	50–62	4
62.5	63–74	5
75	75–87	6
87.5	88–99	7
100	100	8

Important Considerations

- If the number of packets in queue 1 exceeds the rate limit, packets are held in the queue 1 buffer. When this buffer becomes full, the system begins to drop packets from that queue.
- The default rate limit of 100 percent means that queue 1 can “starve” queue 2 under the right conditions. That is, on a given port, packets in queue 2 are always buffered if there are any packets in queue 1 on that same port. Queue 2 packets are processed only after all packets from queue 1 have been processed. If the queue 2 buffers become full, the system begins to drop packets in that queue.

Handling Tagged and Untagged Packets

Consider the following points about how the system processes tagged and untagged packets with respect to CoS information:

- If CoS is enabled and an untagged packet enters a port, the packet is always processed through the low priority queue.
- As described earlier, the CoS priority level information is carried in the IEEE 802.1Q tag. After the packet enters the system, its format is subject to change according to VLAN configurations and ingress and egress rules.
- If an untagged packet enters the system and the VLAN settings modify the packet to become a tagged packet, the system can insert VLAN information but cannot set a priority level other than 0 (whether CoS is enabled or not).
- If a tagged packet enters the system and the tag is retained upon forwarding the packet, the system leaves the priority level as-is (whether CoS is enabled or not).
- If a tagged packet enters the system and VLAN rules cause the tag to be stripped prior to forwarding, the CoS priority information is lost thereafter unless the packet is later processed by a device that can insert tags and priority levels other than 0.

Standards, Protocols, and Related Reading

The following standards provide more information about Class of Service:

- *IEEE 802.1p Traffic Class Expediting and Dynamic Multicast Filtering*

  A supplement to the IEEE 802.1D MAC Bridges base standard that addresses traffic prioritization in local area networks.
- *IEEE 802.1Q Virtual Bridged LANs*

  A base standard that specifies requirements for virtual LANs and a packet format that includes VLAN and CoS priority information.

To obtain copies of these standards, register for an on-line subscription with the Institute of Electrical and Electronics Engineers (IEEE) Web site:

`http://www.ieee.org`

8

TRUNKING

This chapter provides guidelines and other important information about how to implement the trunking function for your SuperStack® II Switch 3900 and Switch 9300.

This chapter covers the following topics:

- Trunking Overview
- Key Concepts
- Key Guidelines for Implementation
- Defining Trunks
- Modifying Trunks
- Removing Trunks
- Standards, Protocols, and Related Reading

You can manage Ethernet trunking in either of these ways:

- *From the* `bridge trunk` *menu of the Administration Console. See the* Command Reference Guide.
- *From the Define Wizard in the Bridge Trunk folder of the Web Management software. See the* Web Management User Guide*.*

Trunking Overview

A *trunk* (also known as an *aggregated link*) works at Layer 2 of the Open Systems Interconnection (OSI) model and allows you to combine multiple Fast Ethernet (Switch 3900 only) or Gigabit Ethernet ports into a single high-speed link between two switches. See Figure 20.

Figure 20 Example of a Trunk

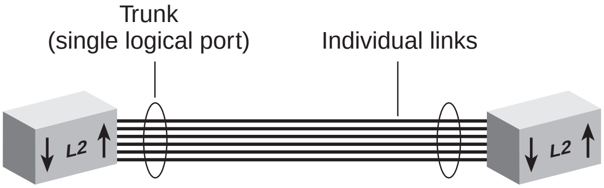

The system treats trunked ports in the same way that it treats individual ports. Also, all higher-level network functions — including Spanning Tree algorithms, virtual LANs (VLANs), and Simple Network Management Protocol (SNMP) management — do not distinguish a trunk from any other network port.

Features

You can configure the following trunking features:

- **Define** — You specify ports and characteristics associated with the trunk.
- **Modify** — You modify a trunk's characteristics or add or remove a port from the trunk.
- **Remove** — You remove a trunk definition from the system.

Benefits

Trunking can help you meet your network capacity and availability needs. With trunks, you can cost-effectively increase the bandwidth between switches or between servers and switches as your network requires. With trunking, you combine multiple Fast Ethernet (Switch 3900 only) or Gigabit Ethernet ports into a single high-speed channel.

If Gigabit Ethernet is not available, you can use trunked Fast Ethernet on the Switch 3900 to increase network capacity. After Gigabit Ethernet is in place and the time comes to scale links beyond 1000 Mbps, you can use trunking to create multigigabit connections.

Trunks also enhance network availability because the Trunk Control Message Protocol (TCMP) detects and handles physical configuration errors in the point-to-point configuration. The system automatically distributes traffic across the ports that are associated with the trunk. If any of the trunk's ports go down or up, the system automatically redistributes traffic across the new arrangement of operational ports.

Key Concepts

Before you configure trunking on your system, become familiar with the key concepts in this section.

Port Numbering in a Trunk

When you combine ports on a trunk, the system logically groups the physical ports that you specify into a single bridge port, identified by a single bridge port number in bridge statistics. For example, Figure 21 shows that Ethernet ports 2, 3, and 4 are represented by bridge port 2 after trunking. The lowest numbered port in the trunk, called the *anchor port*, represents the entire trunk. After trunking, you can select bridge port 2 when you specify bridge port or VLAN information, but you cannot select bridge ports 3 or 4 because they are part of the trunk.

Figure 21 Bridge Port Numbering After Trunking

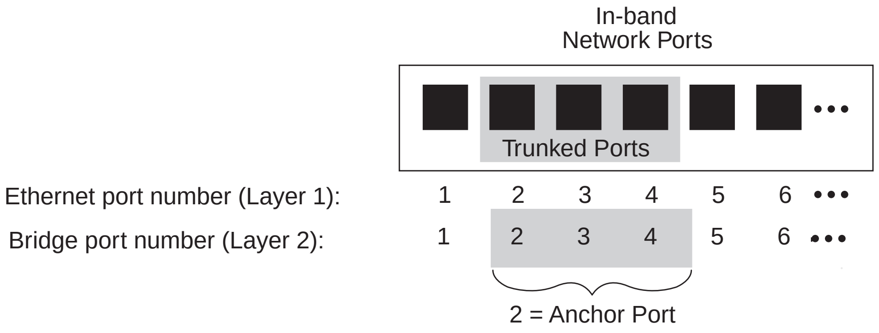

Regardless of whether you define trunking, the physical port numbering on your system remains the same.

It is important to understand the relationships between Ethernet, bridge, and VLAN port-related information:

- **Ethernet port information** — Each physical port is always listed individually, regardless of whether it is part of a trunk.
- **Bridge port information** — This information uses the concept of bridge ports. When you perform bridge port operations, you specify the trunk's anchor port, not the other ports in the trunk, as the representative bridge port. In the bridge port displays, each selectable bridge port has a port field that contains multiple port numbers if the bridge port represents a trunk (for example, `3,5` or `6-8`).
- **VLAN information** — When you define VLANs (as described in Chapter 10), you must specify the bridge ports that you want to be part of the VLAN. If you have a trunk, you specify its anchor port as the bridge port. The VLAN that you create then includes all of the physical ports in the trunk.

Trunk Control Message Protocol (TCMP)

The Trunk Control Message Protocol (TCMP) performs the following functions:

- Detects and corrects trunks that violate trunk configuration rules
- Ensures orderly activation and deactivation of trunk ports

The system runs a separate TCMP agent for each trunk. If TCMP detects an invalid configuration, the protocol restricts the trunk to the largest subset of ports that is a valid configuration.

By default, TCMP is enabled. Keeping TCMP enabled is optional, but recommended. If you disable TCMP, the network still functions, but without automatic trunk validation and reconfiguration.

Each TCMP agent:

- Periodically transmits a TCMP helloMessage through every trunk port.
- Continuously listens for helloMessages from other trunk ports.
- Builds a list of ports that TCMP has detected.
- Uses this list to activate or deactivate trunk ports to maintain valid trunk configurations.

TCMP uses three trunk port states to control port activation and deactivation:

- **notInUse** — A trunk port in this state has not been *selected* to participate in the trunk.
- **selected** — TCMP has *selected* the trunk port to participate in the trunk, but the port has not yet become *active*.
- **inUse** — A trunk port is fully *active* on the trunk.

Terminology

3Com uses a three-tiered framework to describe the different functional areas in a local area network (LAN):

- **Wiring closet** — This area provides connections to user workstations. It also includes downlinks into the data center or campus interconnect.
- **Data center** — This area receives connections from wiring closets and campus interconnect areas. Most local server farms reside here.
- **Campus interconnect** — This area only appears as a separate location in larger networks; smaller networks usually have just wiring closets and data centers. The campus interconnect links campus data centers to each other and may also include an enterprise server farm and connections to a wide area network.

Key Guidelines for Implementation

Consider the following important factors when you implement and configure trunks.

General Guidelines

- To minimize your administrative tasks, define trunks before you define VLANs.
- A system supports up to four point-to-point trunks, each built from up to six ports. All channels in a trunk must be *parallel* and must connect:
  - Correctly configured ports.
  - Identical types of ports (with no two ports on a trunk connected to the same network).
  - Identical types of network nodes (switches or servers).
- You cannot mix Fast Ethernet (Switch 3900 only) and Gigabit Ethernet links in a trunk. All links to be trunked must be homogeneous.
- When multiple links are trunked, it can be difficult to manage and troubleshoot individual port-to-port connections if a connectivity problem occurs. This issue may not be of concern in a server farm room. But if you use trunking extensively between wiring closets and data centers, the large number of connections involved and their distributed nature may make their management and troubleshooting difficult. 3Com recommends that you apply trunking only *within* data center and campus interconnect areas.
- 3Com recommends that you use trunks to increase network availability in the following circumstances:
  - Switch-to-switch connections in the data center and campus interconnect areas
  - Switch-to-server connections in the data center and campus interconnect areas
  - Downlinks from the data center to the campus interconnect
- The trunking feature in 3Com switches is a proprietary implementation. No *de facto* standards exist.
- 3Com trunking technology interoperates with similar technology from other vendors, including Sun Microsystems and Cisco Systems.

Trunk Capacity Guidelines

- The device-to-device burst-transmission rate across a trunk is limited to the speed of just *one* of the port-to-port links within the trunk. For example, the maximum burst rate over a 400 Mbps pipeline with four trunked Fast Ethernet links on the Switch 3900 is 100 Mbps. This limitation preserves frame ordering between devices, usually by moving all traffic between two specific MAC addresses across *only one port-to-port link*. Therefore, trunking provides no direct benefit for some one-way applications, such as server-to-server backups. This limit exists for most vendor implementations.
- The total throughput of a trunk is typically less than the bandwidth that is obtained by adding the theoretical capacity of its individual links. For example, four 1000 Mbps links do not yield a 4000 Mbps trunk. This is true with all vendor implementations.
- In the Switch 3900, a trunked Fast Ethernet pipeline may seem to offer comparable bandwidth to a single Gigabit Ethernet link, and trunked Fast Ethernet may seem like a good way to buy some time before you upgrade connections to Gigabit Ethernet. Table 16 shows that, given a choice, trunking Fast Ethernet may not be an effective strategy.

  If you cannot upgrade to Gigabit Ethernet, then trunking Fast Ethernet in switch-to-switch or switch-to-server links can help you fine-tune or expand network capacity. After Gigabit Ethernet is in place, you can use trunking to further expand switch-to-switch or server-to-switch links.

Table 16 Comparing Gigabit Ethernet with Trunked Fast Ethernet

Comparison Point	Gigabit Ethernet	Trunked Fast Ethernet
Max burst rate	1000 Mbps	100 Mbps
Max aggregate rate	1000 Mbps (2000 Mbps full-duplex)	600 Mbps (over 6 links) (1200 Mbps full-duplex)
Standards compliance	IEEE 802.3z	In progress

Defining Trunks

To define a trunk, you specify the ports that you want to be in the trunk.

Important Considerations

- If you have already defined other trunks on your system, you cannot select ports that are part of an existing trunk.
- Devices that you use in a trunking configuration must have the hardware to support the trunking algorithm.
- You can define more than one trunk at a time, which saves having to reboot the system after each trunk definition.
- When you define a trunk, you specify ports and characteristics associated with the trunk (including Gigabit Ethernet flow control). You can specify them all in one `define` operation.
- When you create the trunk, the entire trunk assumes the current port characteristics.
- Trunk names cannot be longer than 32 characters.
- 3Com recommends that the TCMP state be `enabled`. But devices can operate without TCMP. When TCMP is not in effect on a point-to-point link, its configuration validation is simply absent.
- If your system has more than one media type (for example, Fast Ethernet on the Switch 3900 and Gigabit Ethernet), you are prompted for a media type before you are prompted for the trunk information.
- Trunk names become the port labels when you display information on the trunks.
- All ports in the trunk are set to the selected operating mode (half-duplex mode or full-duplex mode).
- When you create a virtual LAN (VLAN) that includes ports that are part of a trunk, specify the *anchor* port (lowest-numbered port) that is associated with the trunk. For example, if ports 1 through 3 are associated with a trunk, specifying port 1 defines the VLAN to include all of the physical ports in the trunk. If you have not defined trunks, specify one or more port numbers, or specify `all` to assign all ports to the VLAN interface.

- When you create a trunk that includes ports that are part of a VLAN, those ports are removed from the VLAN. You must modify the VLAN and add the new bridge port to the appropriate VLAN. This situation does not apply to the default VLAN (all ports are part of the default VLAN, including the trunk's anchor port).
- Performing an `nvData reset` operation erases all previous trunk information.

Modifying Trunks

You can modify a trunk in two ways:

- You can modify a trunk's characteristics (for example, the operating mode or the TCMP state).
- You can add or remove a port from the trunk.

Important Considerations

- You must keep at least one port that you defined in the original trunk. To completely redefine a trunk configuration, remove the trunk and define a new one.
- You cannot modify, add, or remove ports that are part of different trunks from the one that you are modifying.
- If you have more than one media type on your system (for example, Fast Ethernet on the Switch 3900 and Gigabit Ethernet), you are prompted for a media type before you are prompted for the trunk information.
- Any changes that you make to the trunk's characteristics take effect immediately and do not interrupt trunk operations. If you add or remove a port, however, you must reboot the system to implement the change.

Removing Trunks

You can remove one, several, or all trunks using a single `remove` operation. This feature saves having to reboot the system after each trunk remove.

Important Consideration

- If you remove a Gigabit Ethernet module that has trunks defined, NVRAM is not cleaned up, but the trunk ports are available for use by a replacement module of the same type.

Standards, Protocols, and Related Reading

The Switch 3900 and Switch 9300 systems support these Ethernet standards:

- **IEEE 802.3** — 10BASE-T Ethernet over unshielded twisted pair (UTP) wiring
- **IEEE 802.3u** — 100BASE-T Fast Ethernet over UTP or fiber-optic cable
- **IEEE 802.3z** — 1000BASE-SX Gigabit Ethernet over multimode fiber-optic cable and 1000BASE-LX Gigabit Ethernet over multimode or single-mode fiber-optic cable

Ethernet is a standardized, packet-based network that supports an exponential hierarchy of three line speeds:

- **10 Mbps** — Ethernet
- **100 Mbps** — Fast Ethernet (Switch 3900 only)
- **1000 Mbps** — Gigabit Ethernet

All speeds of Ethernet are based on the IEEE 802.3 standard protocol called Carrier Sense Multiple Access with Collision Detection (CSMA/CD), which controls network access. With CSMA/CD, a station that intends to transmit listens for other Ethernet traffic on the network. When the station does not detect network activity, the station transmits.

9 RESILIENT LINKS

This chapter provides an overview, guidelines, and other important information about how to implement resilient links on your system. The chapter covers these topics:

- Resilient Links Overview
- Key Concepts
- Key Guidelines for Implementation
- Defining and Modifying Resilient Links
- Removing Resilient Links
- Resilient Link State
- Resilient Link Active Port

You can manage resilient links in either of these ways:

- *From the* `bridge link` *menu of the Administration Console. See the* Command Reference Guide*.*
- *From the Bridge folder of the Web Management software. See the* Web Management User Guide*.*

Resilient Links Overview

Resilient links protect your network against an individual link or device failure by providing a secondary backup link that is inactive until it is needed. A resilient link consists of a pair of links — one main link and one standby link. If the main link fails, the standby link immediately takes over the task of the main link. Figure 22 shows a resilient link pair.

Figure 22 Resilient Link Pair

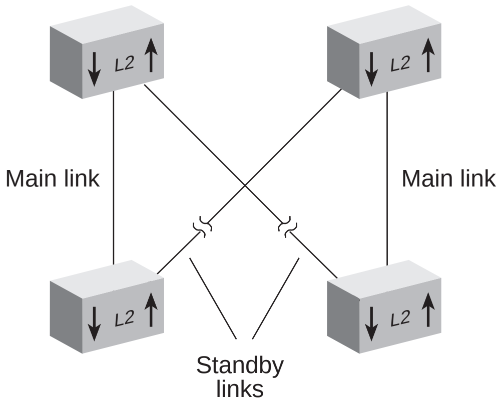

Resilient Links in Operation

Under normal network conditions, the main link carries your network traffic. If a signal loss is detected, the device immediately enables the standby link so that it carries the data and sends a trap to the network management station to alert you of the signal loss. The standby port assumes the profile and carries the network traffic of the main port.

If the main link has a higher bandwidth than its standby link, traffic is switched back to the main link, provided that no loss of link is detected for 2 minutes. Otherwise, you must manually switch traffic back to the main link.

Switchover time to the backup link takes less than 1 second, ensuring no session timeouts and therefore seamless operation.

To be informed about network activity, you can configure the system to generate a Simple Network Management Protocol (SNMP) trap whenever a switchover from one link to the other occurs or whenever the link state (up or down) of either link in the resilient pair changes. For more information about how to configure SNMP traps, see Chapter 13.

Benefits

- Resilient links enable you to protect critical links and prevent network downtime if those links fail.
- 3Com recommends that you implement resilient links to increase network availability in these configurations:
  - Switch-to-switch downlinks from the wiring closet to the data center. The resilient link pair must terminate on a Layer 2 data-center switch.
  - Server-to-switch connections in the data center and campus interconnect areas.

Key Concepts

You can perform three operations with resilient links:

- **Define** — Specify a name and the ports you want to include in the link.

  When you define a resilient link pair, you define:

  - **The main port** — The port on which network traffic runs under normal operation.
  - **The standby port** — The port to which network traffic shifts if the main port fails.

  You can define either the main port or the standby port as the *active* port, that is, the port that is carrying network traffic. The main port is usually defined as the active port.
- **Modify** — Modify the name and ports that are associated with an existing resilient link.
- **Remove** — Remove one or more resilient links from the system.

You can configure these features with resilient links:

- **Link state** — Enable or disable a resilient link pair.
- **Active port** — Specify either port in the pair as the port that carries network traffic.

Key Guidelines for Implementation

Consider these guidelines when you configure resilient links:

- Create resilient links before you define your Virtual LANs (VLANs). If you plan to create resilient links to be part of a VLAN, create the resilient links before you create the VLAN.
- You must reboot the system when you finish defining resilient links. (You can define multiple links in one define operation.)
- If you perform an `nvData reset` operation, the system erases all existing resilient link settings.
- Resilient links are a point-to-point technology; they do not react to *downstream* network failures.
- Inactive links take up ports but do not add to network capacity.
- You cannot set up resilient link pairs if the Spanning Tree Protocol (STP) is enabled.
- You cannot disable ports that are part of a resilient link pair unless a link failure occurs.
- You need to define a resilient link only at one end of the link.
- If an active standby port fails and you have defined a link on the main port, the ports toggle and the main port becomes active.

Defining and Modifying Resilient Links

When you define or modify a resilient link, you specify the name of the link and the ports that you want to associate with the link. You can specify all ports in one define operation.

Important Considerations

- After you define or modify one or more resilient links, you must reboot the system.
- You can define more than one resilient link at a time so that you reboot the system only once.
- When you create a resilient link that includes ports that are part of a VLAN, those ports are removed from the VLAN. You must modify the VLAN to add the new bridge port. This consideration does not apply to the Default VLAN. All ports are part of the Default VLAN.
- If you have already defined other resilient links on your system, you cannot select ports that are part of an existing resilient link to be part of an a different resilient link pair.
- You cannot select a trunked port as part of a resilient link; nor can you select the trunk itself as part of a resilient link.
- The name of a resilient link can be up to 32 characters long.

See the *Command Reference Guide* for a complete description of the commands for defining and modifying resilient links.

Removing Resilient Links

You can remove one or more resilient links with a single remove command.

Important Consideration

After you remove a resilient link pair, you must reboot the system.

Resilient Link State

You can enable or disable one or more resilient link pairs with a single command.

Important Considerations

- When the link state is *enabled*, the resilient link can transmit and receive traffic.
- When the link state is *disabled*, the resilient link no longer transmits or receives frames.

Resilient Link Active Port

The active port is the port that carries traffic. You can designate either the main port or the standby port as the active port.

Important Considerations

- Only one port in a resilient link pair is active at a time.
- By default, the system defines the main port in a resilient pair as the active port when you reboot, unless the main link is down.
- If the main link is of equal or lesser bandwidth than the active standby link, the switchover back to main link is not automatic. If you want the main link to be active, you must configure it as the active port.

10 VIRTUAL LANS

This chapter provides guidelines and other key information about how to use virtual LANs (VLANs) on your system. The chapter covers these topics:

- VLAN Overview
- Key Concepts
- Key Guidelines for Implementation
- VLAN allOpen and allClosed Mode
- Port-based VLANs
  - The Default VLAN
  - User-Configured Port-based VLANs
- Rules of VLAN Operation
- Modifying and Removing VLANs

You can manage VLANs in either of these ways:

- *From the* `bridge vlan` *menu of the Administration Console. See the* Command Reference Guide*.*
- *From the Bridge VLAN folder of the Web Management software. See the* Web Management User Guide*.*

VLAN Overview

A *virtual LAN (VLAN)* is a logical grouping that allows users to communicate as if they were physically connected to a single LAN, independent of the physical configuration of the network. A VLAN is generally considered equivalent to a Layer 2 broadcast domain.

Your system's point of attachment to a given VLAN is called a *VLAN interface*. A VLAN interface exists entirely within a single switch; you control the configuration of the VLAN interfaces on the switch. A VLAN and a VLAN interface are analogous to an IP subnetwork and an IP interface on a router.

Need for VLANs

If a bridge port in a LAN switch receives a frame with a broadcast, multicast, or unknown destination address, it forwards the data to all bridge ports in the VLAN that are associated with the frame except the port on which it was received. This process is referred to as bridge *flooding*. As networks grow and the amount and types of traffic increase, bridge flooding may create unnecessary traffic problems that can clog the LAN.

To help control the flow of traffic through a switch and meet the demands of growing networks, vendors have responded by:

- Using customized packet filtering to further control which packets are forwarded through the bridge. These filters can be complex to configure.
- Using more and more routers as broadcast firewalls to divide the network into broadcast domains. As the number of legacy routers increase, latency begins to degrade network performance, administration overhead increases, and operating costs rise.
- Using the Spanning Tree algorithm in switches to control the flow of traffic among LANs (for redundant links). This mechanism works best only in certain types of LAN topologies.

VLANs provide a high-performance and easy-to-implement alternative to routers for broadcast containment. Using switches with VLANs:

- Each network segment can contain as few as one user (approaching private port LAN switching), while broadcast domains can be as large as 1,000 users or even more.
- VLANs can help you track workstation movements to new locations without manual reconfiguration of IP addresses.
- VLANs can be used to isolate unicast traffic to a single broadcast domain, providing a form of network security.

Benefits

You can use VLANs to:

- Reduce the cost of equipment moves, upgrades, and other changes and simplify network administration.
- Create virtual workgroups in which members of the same department or section appear to share the same LAN, with most of the network traffic staying in the same VLAN broadcast domain. They can isolate broadcast and multicast traffic to a single broadcast domain, as well as unicast traffic.
- Help avoid flooding and minimize broadcast and multicast traffic.
- Reduce the need for routing to achieve higher network performance and reduced costs.
- Control or filter communication among broadcast domains.

Features

The system supports the following VLAN features:

- **Settable modes** — For the entire system, you can establish a less-restrictive VLAN environment with allOpen mode or a more secure VLAN environment with allClosed mode. The VLAN mode dictates the requirements for the port-based VLANs. See “Terminology” for more information about the VLAN modes.
- **Port-based VLANs** — Determines VLAN membership by means of a VLAN ID (VID) that is assigned to a group of ports. Your system provides a special port-based VLAN by default that contains all ports. This special VLAN is called the *default VLAN*. The system also supports *static* (user-defined) port-based VLAN configuration. See “VLAN allOpen and allClosed Mode” for information on the default VLAN. See “User-Configured Port-based VLANs” later in this chapter for information on static VLAN configuration.

- **Per-port IEEE 802.1Q tagging** — Selecting IEEE 802.1Q tagging for each port in a VLAN dictates that frames be encapsulated and tagged as specified in the IEEE 802.1Q standard. See "User-Configured Port-based VLANs" later in this chapter for more information on per-port tagging.
- **Ethernet Links** — Depending on the configuration of your system, you can configure VLANs that incorporate Gigabit Ethernet links as well as 10/100 Mbps Ethernet. See Chapter 5 for Ethernet information.

Key Concepts

Before you configure VLANs, review the following key concepts.

IEEE 802.1Q and Per-port Tagging

IEEE 802.1Q is a standard for VLANs. It aims to:

- Define an architecture to logically partition bridged LANs and provide services to defined user groups, independent of physical location
- Allow interoperability between multivendor equipment

IEEE 802.1Q defines the bridging rules for VLANs (called ingress and egress rules). See "Rules of VLAN Operation" later in this chapter for detailed descriptions of these rules.

IEEE 802.1Q also specifies a tag format that embeds explicit VLAN membership information within each frame in a 12-bit VLAN ID (VID), providing 4094 possible VLANs. (IEEE 802.1D, which incorporates 802.1p, uses this same frame format but takes advantage of an additional 3 bits to specify the priority levels used for Class of Service differentiation.)

The system supports per-port tagging (that is, you can select IEEE 802.1Q tagging or no tagging on each port in the VLAN). Tagged and nontagged ports can coexist in the same VLAN group. There are two modes:

- **802.1Q tagging mode** — With this form of tagging, VLAN frames are encapsulated and tagged as specified in the IEEE 802.1Q standard. In frame tagging mode, an explicit header that identifies to which VLAN the frame belongs is inserted into each frame of interswitch data. Frames in the same VLAN can be tagged or untagged. An untagged port in a VLAN cannot insert a tag, but it can recognize a tagged frame. Use this mode for VLANs in an IEEE 802.1Q environment.
- **Nontagging mode** — The default tagging mode. Use this mode for ports if your environment includes end stations that do not support IEEE 802.1Q VLANs. Nontagged VLAN ports accept tagged frames; however, any traffic that is transmitted from an untagged port on a VLAN is untagged.

Devices (end stations, routers, switches, and so forth) that are connected to an explicitly tagged port must be capable of supporting 802.1Q tagging. If the port is untagged in the VLAN to which these devices belong, however, they do not have to support 802.1Q tagging.

VLAN IDs

Each VLAN is identified by its VLAN ID (VID). For the VLANs that you create, the system keeps track of its used VLAN ID numbers to help you to select the next available VLAN ID. If tagging is enabled on the transmit port for that VLAN, data frames that are sent by the system are tagged according to IEEE 802.1Q (the tag contains the VID). Tagged IEEE 802.1Q data frames that are received on the system are assigned to the VLAN that corresponds to the VID contained in the tag. The default VLAN always uses the VID of 1.

Before you assign any VIDs, review the information in Table 17.

Table 17 Assigning ID Numbers to VLANs

VLAN ID Number	Description
VID 1	Default VLAN assigned by IEEE and 3Com Corporation
VID 4095	Reserved
VID 2 – 4094	Numbers that you assign when you create VLANs

Terminology

Review the following terms:

- **Default VLAN** — The predefined port-based VLAN interface on your system that always uses VID 1 and the name Default. The default VLAN initially includes all of the bridge ports without any tagging, but you can modify the bridge ports and tag status of the default VLAN. See "The Default VLAN" later in this chapter for more information.
- **VLAN origin** — The method used to create the VLAN. Because this system supports only static (user-configured) VLAN configuration, the VLAN display always shows an origin of `static` for your VLANs.
- **VLAN mode** — A system-wide mode that determines whether data with a unicast MAC address can be forwarded between configured VLANs (`allOpen`). In `allClosed mode`, each VLAN has its own address table and data cannot be forwarded between VLANs. The default VLAN mode is `allOpen`. See "VLAN allOpen and allClosed Mode" for more information about how to select the VLAN mode.
- **Tagging type** — On each port in the VLAN, whether there is explicit VLAN membership information in each frame (the IEEE 802.1Q header and the VID). Types: no tagging or IEEE 802.1Q tagging.
- **Port membership** — The bridge ports that you assign to be part of the VLAN. All bridge ports are initially part of the default VLAN. If you have created trunks, you must specify the anchor port (that is, the lowest-numbered port) in the trunk to add the entire trunk to the VLAN.
- **VLAN name** — A name that you assign to the VLAN. It can contain up to 32 ASCII characters. If the name includes spaces, enclose the name in quotation marks. The default VLAN uses the name Default.
- **Ingress and egress rules** — *Ingress* rules determine the VLAN to which an incoming frame belongs. If a frame cannot be assigned to any VLAN, it is assigned to the null VLAN, which contains no ports and has no associated address table in allClosed mode. *Egress* rules determine whether the frame is forwarded, flooded, or filtered, as well as the tag status of the transmitted frame. For more information on ingress and egress rules, see "Rules of VLAN Operation" later in this chapter.

Key Guidelines for Implementation

Consider all of the following guidelines before you configure port-based VLANs on your system.

- Your system supports a maximum of 127 port-based VLANs.
- The VLAN mode of allOpen or allClosed applies to *all* VLANs that are associated with the system. Configure the VLAN mode *before* you define any VLANs.

If you change the VLAN mode after you have defined VLANs, the system deletes all VLANs and redefines the default VLAN. See "Modifying the VLAN Mode" later in this chapter.

- Consider maintaining the system's default VLAN. The default VLAN preserves the flooding of unspecified traffic because it initially contains all of the system's bridge ports with no tagging.
- Define trunks *before* you define your VLANs. (If you define a VLAN with certain ports and subsequently configure some of those ports to be part of a trunk, the system removes those ports from that VLAN.) See "Trunking and the Default VLAN" for information about how trunking actions affect the default VLAN. When you define a VLAN that includes trunk ports, you must specify the trunk's anchor port (lowest-numbered port). For trunking information, see Chapter 8.
- You can configure overlapping VLANs if the VLANs have some distinguishing characteristic. For example, a bridge port can be shared by multiple VLANs as long as there is a distinguishing characteristic for the shared port, such as the tagging type.
- When the system receives a frame, the frame is assigned to a VLAN using the ingress rules. See "Ingress Rules" later in this chapter. When it transmits the frame, the system determines the tag status (none or IEEE 802.1Q tagging) by referring to the tag status of the transmit port in the frame's assigned VLAN. In allOpen mode, if a frame is transmitted on a port that does not belong to the assigned VLAN, it is transmitted untagged.

VLAN allOpen and allClosed Mode

You can select allOpen or allClosed as the VLAN mode for your entire system. The default is `allOpen`.

The 3Com terms "allOpen" and "allClosed" are equivalent to the respective terms "Shared VLAN Learning" (SVL) and "Independent VLAN Learning" (IVL) that are used in the IEEE 802.1Q standard. 3Com imposes the restriction of choosing one VLAN mode for the entire system; more complex logic for assigning SVL and IVL to individual ports is described in the IEEE 802.1Q standard.

Important Considerations

- In general, select your VLAN mode *before* you define your VLANs.
- Select a VLAN mode as follows:
  - **allOpen** — Use this less-restrictive mode if you have no security issues concerning the forwarding of data between VLANs. The allOpen mode is the default VLAN mode for all VLANs that you create. The allOpen mode implies that the system uses a single bridge address table for all of the VLANs on the system (the default configuration).

    This mode permits *tagged* data with a unicast MAC address to be forwarded between VLANs. For example, tagged data received on VLAN 2 with a destination of VLAN 3 is forwarded to VLAN3.
  - **allClosed** — Use this more-restrictive mode if you are concerned about security between VLANs. Data cannot be forwarded between VLANs. The allClosed mode implies that each VLAN that you create has its own address table.
- Your selection of a VLAN mode affects how you manipulate bridge port address options (using the Administration Console or the Web). For example:
  - If you select allClosed mode, you *must* specify a VLAN interface index to identify the appropriate bridge address table.
  - If you select allOpen mode (the default), the entire system has only one address table, so you can manipulate the bridge port address options without specifying a VLAN interface index.

Modifying the VLAN Mode

To change your VLAN mode, follow these steps:

1. Using the Administration Console or Web Management, modify the VLAN mode to specify the new VLAN mode.

   When you change the mode, the system deletes all of your existing configured VLANs and reverts to the default VLAN.

2. Reconfigure your VLANs.

Using allOpen Mode

Here are the requirements for defining port-based VLANs in allOpen mode:

- For bridge ports that are not shared, the distinguishing characteristic is a VLAN's ownership of specific ports. (For example, VLAN 1 has ports 1 through 6 and VLAN 2 has ports 7 through 12.)
- For VLANs with shared (overlapped) ports, you *must* use IEEE 802.1Q tagging for those ports. When tagging is the distinguishing characteristic, only one of the shared ports can be set to tagging none.

Using allClosed Mode

Here are the requirements for defining port-based VLANs in allClosed mode:

- For bridge ports that are not shared, the distinguishing characteristic is a VLAN's ownership of specific ports. (For example, VLAN 1 has ports 1 through 6 and VLAN 2 has ports 7 through 12.)
- For VLANs with shared (overlapped) ports, you *must* use IEEE 802.1Q tagging for those ports. When tagging is the distinguishing characteristic, only one of the shared ports can be set to the tagging mode of *none*.

Port-based VLANs

Port-based VLANs logically group together one or more bridge ports on the system and implicitly use the generic protocol type *unspecified*. Each arbitrary collection of bridge ports is designated as a *VLAN interface.* The VLAN interface belongs to a given VLAN. Flooding of all frames that are received on bridge ports in a VLAN interface is constrained to that VLAN interface.

Your system supports the following types of port-based VLANs:

- The default VLAN, a special VLAN predefined on the system
- User-configured port-based VLANs (static VLANs)

The Default VLAN

The system predefines a port-based VLAN with a VID of 1 and the name Default to initially include all of the system's bridge ports without any tagging.

The default VLAN is the flood domain when:

- The system receives data for a protocol that is not supported by any VLAN in the system
- The system receives data for a protocol that is supported by defined VLANs, but these VLANs do not contain the port that is receiving the data

See "Rules of VLAN Operation" later in this chapter for more details.

Modifying the Default VLAN

The default VLAN always uses a VID of 1 and the name Default. It initially has no tagging on any of the ports. Keeping the default VLAN intact, without any modifications, ensures that the system has a VLAN for all of its ports. If necessary, you can modify the default VLAN to remove certain ports or change the tag status of a port. However, you cannot change the VID or name of the default VLAN.

You can delete the default VLAN as long as you do not have an IP interface associated with it. (Layer 2 systems support one IP interface per VLAN; you can define up to two IP interfaces that use the interface type `VLAN`.) If you try to remove the default VLAN and it is associated with an IP interface, the system displays the following error message:

```
Request failed- Cannot delete vlan 1  -  interface in use by
client
```

If you remove the default VLAN completely, you can redefine it by specifying the VID of 1. (The system indicates that 1 is associated with the default VLAN only.)

To ensure that data can be forwarded, associate a bridge port with a VLAN. This association is mandatory in allClosed mode. If you remove the default VLAN (and you have no other VLANs defined for the system), your ports may not forward data until you create a VLAN for them.

Trunking and the Default VLAN

Another benefit of maintaining the default VLAN (with any number of ports) involves trunking. 3Com strongly recommends that you define your trunks *before* you define your VLANs.

Trunking with the default VLAN intact

Trunking actions (including trunking used with MultiPoint Link Aggregation) affect the default VLAN in these ways:

- If you have only the default VLAN with all ports and you define a trunk (or subsequently remove a trunk), the ports that are listed in the VLAN summary for the default VLAN do not change. In this case, maintaining the default VLAN with all ports ensures that trunks can come and go without causing any VLAN changes.
- If you have the default VLAN as well as additional VLANs and you then define a trunk for ports in one of the other VLANs, the system removes those ports from that VLAN and places them in the default VLAN. The same action occurs when you remove an existing trunk from a VLAN that you created after the trunk. For example:

Ports Before Action	Trunking Action	Ports After Action
default VLAN: ports 1–4 vlan2: ports 5–11	Define a trunk with ports 7 and 8	default VLAN: ports 1–4, 7–8 vlan2: ports 5–6, 9–11

- If you have the default VLAN as well as other VLANs and you then modify an existing trunk that has ports in one of the VLANs, any port that you removed from the trunk is removed from the VLAN and placed in the default VLAN. For example:

Ports Before Action	Trunking Action	Ports After Action
default VLAN: ports 1–4 vlan2: ports 5–11 (ports 5–8 are trunk ports)	Modify existing trunk to have ports 6–8 (remove port 5, the anchor port)	default VLAN: ports 1–5 vlan2: ports 6–11 (port 6 is new anchor port)

Trunking with the default VLAN removed

If you remove the default VLAN, there is no place to which to return ports that are altered by trunking, as discussed in these examples:

- If you have VLANs (but no default VLAN) and you then define a trunk for ports in one of the VLANs, those ports are removed from that VLAN and are not assigned to any other VLAN. If you later remove the trunk, these ports are not reassigned to the VLAN; they no longer have a VLAN associated with them. For example:

Ports Before Action	Trunking Action	Ports After Action
vlan2: ports 1–11	Define trunk with ports 5–8	vlan2: ports 1–4, 9–11

- If you have VLANs (but no default VLAN) and you then modify an existing trunk that has ports in one VLAN, any port that is removed from the trunk is removed from the VLAN and no longer has a VLAN. For example:

Ports Before Action	Trunking Action	Ports After Action
vlan2: ports 1–11 (ports 5–8 are trunk ports)	Modify existing trunk to have ports 6–8 (remove port 5, the anchor port)	vlan2: ports 1–4, 6–11 (port 6 is new anchor port)

See Chapter 8 for more information on using trunks.

User-Configured Port-based VLANs

You can explicitly configure port-based VLAN interfaces on your system.

Important Considerations

When you create this type of VLAN interface, review these guidelines:

- When you select the bridge ports that you want to be part of the VLAN, the bridge ports that you specify as part of the VLAN are the same as your physical ports, unless you have created trunks.
- If you define trunks, a single bridge port called the *anchor port* (the lowest-numbered port in the trunk) represents all ports that are part of the trunk. Only the anchor bridge port for the trunk is selectable when you create VLANs; the other bridge ports in the trunk are not selectable. For more information on trunking, see Chapter 8.
- Decide whether you want the ports that you are specifying for the VLAN interface to be shared by any other VLAN interface on the system. Shared ports produce *overlapped* VLANs; ports that are not shared produce *nonoverlapped* VLANs.
- The per-port tagging options are IEEE 802.1Q tagging or no tagging. The IEEE 802.1Q tagging option embeds explicit VLAN membership information in each frame.
- Overlapped VLANs require tagging; that is, two port-based VLAN interfaces may contain the same bridge port if one of the VLAN interfaces defines the shared port to use IEEE 802.1Q tagging. This rule is true for either allOpen or allClosed mode. For example, a shared bridge port is set to tagging *none* for one VLAN and *IEEE 802.1Q* tagging for the other VLAN, or to *IEEE 802.1Q* tagging for each VLAN.
- To define a port-based VLAN interface, specify this information:
  - The VID, or accept the next-available VID.
  - The bridge ports that are part of the VLAN. (If you have trunk ports, specify the anchor port for the trunk.)
  - Tag status (none or IEEE 802.1Q).
  - The unique name of the VLAN interface.

Example 1: Nonoverlapped VLANs

Figure 23 shows two systems that have nonoverlapping port-based VLANs and no port tagging. Ports 1–3 on Layer 2 Switch Device 1 make up VLANA, and ports 4–6 make up VLANB. All frames that are received on a port are assigned to the VLAN that is associated with that port. For instance, all frames that are received on port 2 in VLANA are assigned to VLANA, regardless of the data that is contained in the frames.

Figure 23 Port-based VLANs without Overlapped Ports

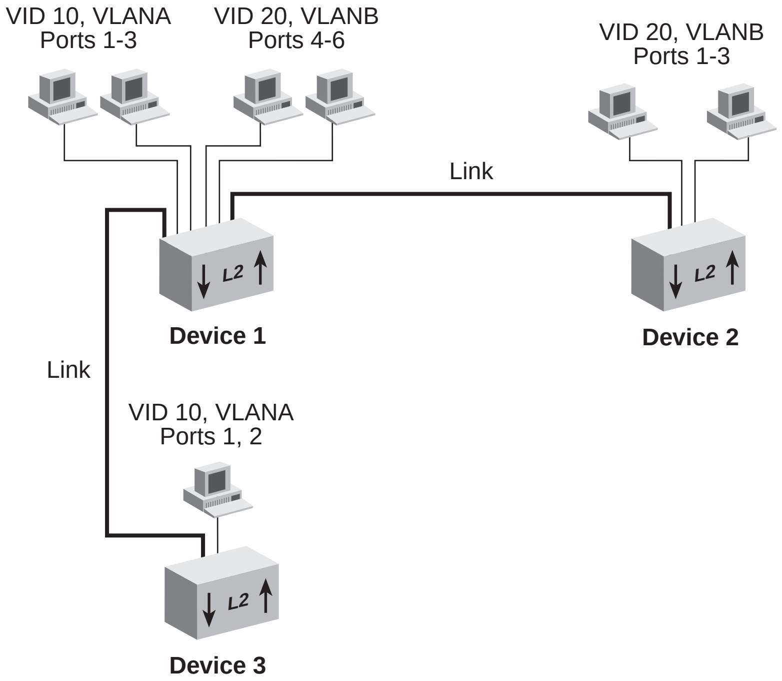

After an incoming frame is assigned to a VLAN, the frame is forwarded, filtered, or flooded within its VLAN, based on the standard bridging rules.

This situation causes different behavior for allOpen VLANs versus allClosed VLANs. For example, for allClosed VLANs, if an untagged frame is received on a port in VLANA with a destination address that is known in the address table of VLANB, the frame is flooded throughout VLANA because it is an unknown address for VLANA.

For allOpen VLANs, however, there is one address table; therefore, the frame is forwarded to the port that corresponds to the known destination address. If the transmit port is not a member port of VLANA, the frame is transmitted according to that port's tag status on VLANB.

Table 18 shows the information that can be used to configure these VLANs *without* overlapped ports on Device 1 (the device at the upper left):

Table 18 Port-based VLAN Definitions without Overlapped Ports for Device 1

VLANA	**VLANB**
VLAN Index 2	VLAN Index 3
VID 10	VID 20
Bridge ports *1–3*	Bridge ports *4–6*
Tagging *none* for ports 1–3:	Tagging *none* for ports 4–6
VLAN name *VLANA*	VLAN name *VLANB*

Example 2: Overlapped VLANs

Figure 24 shows port-based VLANs that *overlap* on bridge port 6. This bridge port is tagged in one VLAN (VLAND) but not in the other VLAN (VLANC).

Figure 24 Port-based VLAN Definitions with Overlapped Port

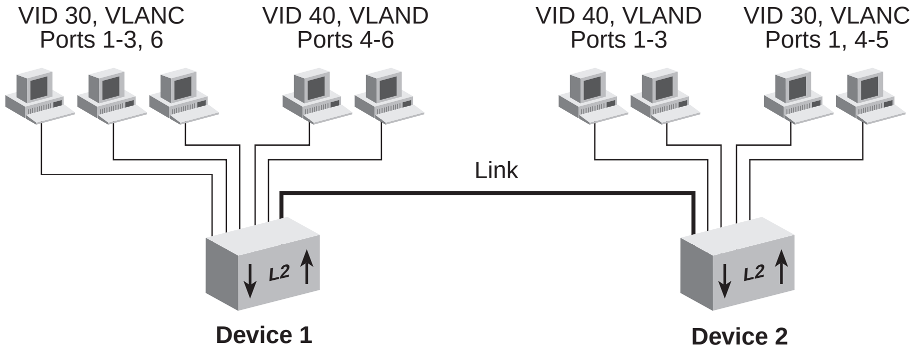

Table 19 shows the information that you use to configure these VLANs *with* overlapped ports on Device 1.

Table 19 Port-based VLAN Definitions with Overlapped Ports for Device 1

VLANC	VLAND
VLAN Index 3	VLAN Index 4
VID 30	VID 40
Bridge ports *1–3,6*	Bridge ports *4–6*
Tagging *none* for ports 1–3, 6	Per-port tagging: Port 4 — *none* Port 5 — *none* Ports 6 — *IEEE 802.1Q*
VLAN name *unspecA*	VLAN name *unspecB*

If your VLAN includes trunk ports, specify the anchor port (lowest-numbered port) of the trunk. For example, if ports 1 through 3 in VLANC are associated with a trunk, specify only bridge port *1* to define the VLAN to include all of the physical ports in the trunk (ports 1 through 3). The tagging type for port 1 applies to all ports in the trunk.

Rules of VLAN Operation

After you select a VLAN mode for the system and create VLAN interfaces with VLAN characteristics such as IEEE 802.1Q or no tagging and port membership, the system determines the details of VLAN operation by observing two main types of rules:

- **Ingress rules** — Assign an incoming frame to a specific VLAN.
- **Egress rules** — Use standard bridging rules to determine whether the frame is forwarded, flooded, or filtered. These rules also determine the tag status of the transmitted frame.

These rules are classified in the IEEE 802.1Q standard. In addition, the system relies on some system-specific rules.

Ingress Rules

These rules determine the VLAN to which an *incoming* frame belongs. The frame is assigned to the VLAN that matches most closely. A protocol match hierarchy is used to find the most specific match.

The ingress rules, which are classified according to your VLAN mode, use the following process to determine the most specific match:

1. IEEE 802.1Q tag VID value
2. The default VLAN (an untagged VLAN with all ports and a VID of 1), or any port-based VLAN

Ingress Rules for allClosed VLANs

- If the incoming frame is an IEEE 802.1Q tagged frame, the frame is assigned to the VLAN if the receive port and the VID of the frame match those of the VLAN. If there is no match, the frame is dropped.
- If the incoming frame is not tagged, the frame is assigned to the VLAN if the receive port is untagged (that is, if tagging is set to none) and if the receive port of the frame matches that of the VLAN. If there is no match, the frame is dropped.

Ingress Rules for allOpen VLANs

- If the frame is an IEEE 802.1Q tagged frame, the frame is assigned to the VLAN if the VID of the frame matches that of the VLAN. If there is no VID match, the frame is dropped.
- If the frame is not tagged, the frame is assigned to the VLAN if the receive port is untagged (that is, if tagging is set to none) and if the receive port of the frame matches that of the VLAN. If there is no match, the frame is dropped.

Egress Rules

These rules determine whether the *outgoing* frame is forwarded, filtered (dropped), or flooded; they also determine the frame's tag status. The same standard bridging rules apply to both open and closed VLANs, but they result in different behavior based on the allOpen mode (one address table for the system) versus allClosed mode (one address table for each VLAN). In allClosed mode, if an untagged frame is associated with a VLAN using VID 1 and has a destination address associated with a VLAN using VID 2, the frame is flooded over the VID 1 VLAN. In allOpen mode, the frame is forwarded out of the port in the VID 2 VLAN (where the address is known) and with the tag status of that port.

Standard Bridging Rules for Outgoing Frames

The frame is handled according to these bridging rules:

- If the transmit port is tagged and is not a member of the assigned VLAN, the frame is dropped.
- If the frame's destination address matches an address that was learned on the receive port, it is *filtered* (dropped).
- If the frame's destination address matches an address that was learned on a port other than the receive port, it is *forwarded* to that port.
- If a frame with an unknown, multicast, or broadcast destination address is received, then it is *flooded* (that is, forwarded to all ports on the VLAN that is associated with the frame, except the port on which it was received).
- If the frame's destination address matches a MAC address of one of the bridge's ports, it is further processed, not forwarded immediately. This type of frame is a management/configuration frame, such as a RIP update, SNMP get/set PDU, Administration Console Telnet packet, or a Web Management Interface http packet.

Tag Status Rules

After the VLAN and the transmit ports are determined for the frame, the tag status rules determine whether the frame is transmitted with an IEEE 802.1Q tag:

- For each port on which a frame is to be transmitted, if that port is tagged for the VLAN that is associated with the frame, transmit the frame as a tagged frame.
- For each port on which a frame is to be transmitted, if that port is *not* tagged for the VLAN that is associated with the frame, transmit the frame as an untagged frame.

Modifying and Removing VLANs

You can modify or remove any VLANs on your system. Review the following guidelines before you modify or remove VLANs:

- When you modify VLAN information for a VLAN interface on your system, you have the option to change VLAN characteristics such as the member bridge ports and the form of explicit tagging.
- When you modify or remove a VLAN interface, you must specify a VLAN interface index to identify the VLAN interface. The default VLAN uses the VLAN interface index of 1. You cannot modify the VID or name of the default VLAN.
- You cannot delete a VLAN if you have an IP interface associated with it. (Layer 2 systems support one IP interface per VLAN; you can define up to two IP interfaces that use the interface type `vlan`.) If you try to remove a VLAN associated with an IP interface, the system displays an error message stating that the interface is in use.
- If you add ports to a specific VLAN, you are permitting additional traffic through that port. If you remove ports from a specific VLAN and the default VLAN is intact, those ports come under jurisdiction of the default VLAN and therefore have an unspecified protocol type and no explicit or implicit tagging.
- In general, verify that each bridge port is associated with at least one VLAN in order to handle traffic.

- If you modify the default VLAN to remove certain ports, verify that those ports are included in another VLAN. See "Modifying the Default VLAN" earlier in this chapter for more information about the default VLAN.
- If you remove the default VLAN and you have no other VLANs defined for the system, your ports may not be able to forward data until you create a VLAN for them (for example, if you are using allClosed mode). If you redefine the default VLAN, it must use the VID of 1. (VID 1 is reserved for the default VLAN.)

11 INTERNET PROTOCOL (IP)

This chapter provides guidelines and other key information about how to configure interfaces on your system to use IP.

The chapter covers these topics:

- Overview
- Key Concepts
- Configuring IP Interfaces
- Address Resolution Protocol (ARP)
- Routing Information Protocol (RIP)
- Domain Name System (DNS)
- Standards, Protocols, and Related Reading

You can manage IP interface features in either of these ways:

- *From the* `ip` *menu of the Administration Console. See the* Command Reference Guide.
- *From the IP folder of the Web Management software. See the* Web Management User Guide.

Overview

The Internet Protocol (IP) is a standard network protocol that is used for communications among various networking devices. Your system is a Layer 2 forwarding device that has the ability to advertise IP (Layer 3) addresses.

You must set up an IP interface for your system:

- To gain access to the system using TCP/IP or to manage the system using SNMP.
- To establish a routing table.

System Management: In-Band Versus Out-Of-Band

You can manage Switch 3900 and Switch 9300 systems in-band. For the Switch 9300, you can also manage it out-of-band.

- In-band management (3900 and 9300)

  In-band management means that you use one of the ports that carries data traffic to manage the system. Bandwidth use on the port is divided between data traffic and management traffic.

  If you are managing your network in-band, you need to set up an IP routing interface and at least one VLAN. For information about setting up an IP routing interface, see “Establish an IP Interface” later in this chapter. For information about setting up VLANs, see Chapter 10.

- Out-of-band management (9300 only)

  Out-of-band management means that you use a 10BASE-T port dedicated to management traffic so that you do not have to use the resources of a data traffic port for management functions.

  If you are managing your system out-of-band, you need to assign an IP address and subnet mask for the out-of-band Ethernet port on your system. For information about configuring these IP attributes for the out-of-band Ethernet port, see “Establish an IP Interface” later in this chapter.

The `ip interface define` (in-band) and `management ip interface define` (out-of-band) options on the system console are documented in the *Command Reference Guide*. You can also use Web Management to configure an IP interface.

Routing Table

Even though your system is a Layer 2 device and therefore not capable of routing, it can advertise IP addresses. Setting up an IP interface establishes a routing table, which contains information such as IP addresses that serve as the next hop in forwarding packets to their ultimate destinations, whether the routes are static or dynamic, and how many hops it takes packets to get to their destinations. See "Routing Table Elements" later in this chapter for more information.

Key Concepts

Review these topics before you establish an IP interface:

- IP Addresses
- IP Interfaces
- Routing Table Elements

IP Addresses

IP addresses are 32-bit addresses that consist of a *network part* (the address of the network where the host is located) and a *host part* (the address of the host on that network). See Figure 25.

Figure 25 IP Address: Network Part and Host Part

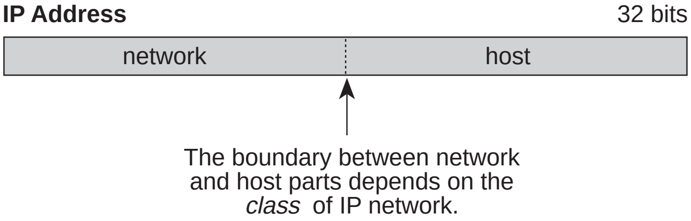
IP Address
32 bits
network
host
The boundary between network and host parts depends on the *class* of IP network.

IP addresses differ from Ethernet and Fiber Distributed Data Interface (FDDI) MAC addresses, which are unique hardware-configured 48-bit addresses. A central agency assigns the network part of the IP address, and you assign the host part. All devices that are connected to the same network share the same network part (also called the *prefix*).

Dotted Decimal Notation

The actual IP address is a 32-bit number that is stored in binary format. These 32 bits are segmented into 4 groups of 8 bits — each group is referred to as a *field* or an *octet*. Decimal notation converts the value of each field into a decimal number, and the fields are separated by dots.

Figure 26 Dotted Decimal Notation for IP Addresses

10011110.01100101.00001010.0010000 = Binary notation

158.101.10.32 = Decimal notation

The decimal value of an octet whose bits are all 1s is 255.

Network Portion

The location of the boundary between the network part and the host part depends on the class that the central agency assigns to your network. The three primary classes of IP addresses are A, B, and C:

- **Class A address** — Uses 8 bits for the network part and 24 bits for the host part. Although only a few Class A networks can be created, each can contain a very large number of hosts.
- **Class B address** — Uses 16 bits for the network part and 16 bits for the host part.
- **Class C address** — Uses 24 bits for the network part and 8 bits for the host part. Each Class C network can contain only 254 hosts, but many such networks can be created.

The high-order bits of the network part of the address designate the IP network class. See Table 20.

Table 20 How Address Class Corresponds to the Address Number

Address Class	High-order Bits	Address Number (Decimal)
A	0nnnnnnn	0 – 127
B	10nnnnnn	128 – 191
C	11nnnnnn	192 – 254

Subnetwork Portion

The IP address can also contain a *subnetwork part* at the beginning of the host part of the IP address. Thus, you can divide a single Class A, B, or C network internally, allowing the network to appear as a single network to other external networks. The subnetwork part of the IP address is visible only to hosts and gateways on the subnetwork.

When an IP address contains a subnetwork part, a *subnet mask* identifies the bits that constitute the subnetwork address and the bits that constitute the host address. A subnet mask is a 32-bit number in the IP address format. The *1* bits in the subnet mask indicate the network and subnetwork part of the address. The *0* bits in the subnet mask indicate the host part of the IP address, as shown in Figure 27.

Figure 27 Subnet Masking

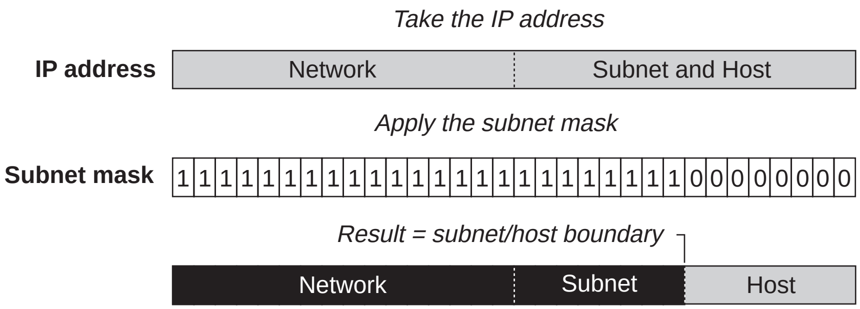

Figure 28 shows an example of an IP address that includes network, subnetwork, and host parts. Suppose the IP address is *158.101.230.52* with a subnet mask of *255.255.255.0*. Because this is a Class B address, this address is divided as follows:

- *158.101* is the network part
- *230* is the subnetwork part
- *52* is the host part

As shown in this example, the 32 bits of an IP address and subnet mask are usually written using an integer shorthand. This notation translates four consecutive 8-bit groups (octets) into four integers that range from 0 through 255. The subnet mask in the example is written as 255.255.255.0.

Traditionally, subnet masks were applied to octets in their entirety. However, one octet in the subnet mask can be further subdivided so that part of the octet indicates an *extension* of the network number, and the rest of the same octet indicates the host number, as shown in Figure 28.

Figure 28 Extending the Network Prefix

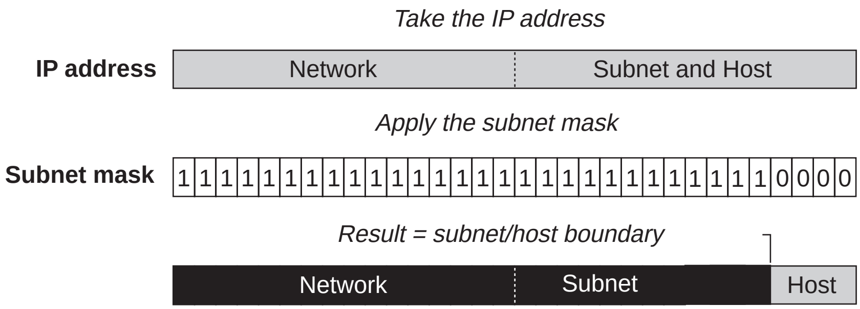

Using the Class B IP address from Figure 27 (158.101.230.52), the subnet mask is 255.255.255.240.

The number that includes both the Class B natural network mask (255.255) and the subnet mask (255.240) is sometimes called the *extended network prefix*.

Continuing with the previous example, the subnetwork part of the mask uses 12 bits, and the host part uses the remaining 4 bits. Because the octets are actually binary numbers, the number of subnetworks that are possible with this mask is 4096 (2^{12}), and the number of hosts that are possible in each subnetwork is 16 (2^4).

Subnet Mask Numbering

An alternate method to represent the subnet mask numbers is based on the number of bits that signify the network portion of the mask. Many Internet Service Providers (ISPs) now use this notation to denote the subnet mask. See Table 21.

Table 21 Subnet Mask Notation

Standard Mask Notation	Network Prefix Notation
100.100.100.100 (255.0.0.0)	100.100.100.100/8
100.100.100.100 (255.255.0.0)	100.100.100.100/16
100.100.100.100 (255.255.255.0)	100.100.100.100/24

The subnet mask 255.255.255.255 is reserved as the default broadcast address.

IP Interfaces

An IP interface connects the system to a subnetwork. On your system, more than one port can connect to the same subnetwork.

Each IP interface has an IP address and a subnet mask. This IP interface address defines both the number of the network to which the IP interface is attached and its host number on that network. An interface IP address serves two functions:

- For sending IP packets to or from the system
- For defining the network and subnetwork numbers of the segment that is connected to that interface

Routing Table Elements

With a routing table, a device determines how to send a packet toward its ultimate destination. The routing table contains an entry for every learned and locally defined network. The size of the routing table is dynamic and can hold thousands of entries; the actual number depends upon what other protocols are being routed.

A device uses the routing table when the destination IP address of the packet is not on a network or subnetwork to which it is directly connected. The routing table provides the IP address of a device that can forward the packet toward its destination.

The routing table consists of the following elements:

- **Destination IP address** — The destination network, subnetwork, or host
- **Subnet mask** — The subnet mask for the destination network
- **Metric** — A measure of the distance to the destination. In the Routing Information Protocol (RIP), the metric is the number of hops through devices
- **Gateway** — The IP address of the interface through which the packet travels on its next hop
- **Status** — Information that the routing protocol has about the route, such as how the route was put into the routing table

Figure 29 shows a sample routing table.

Figure 29 Sample Routing Table

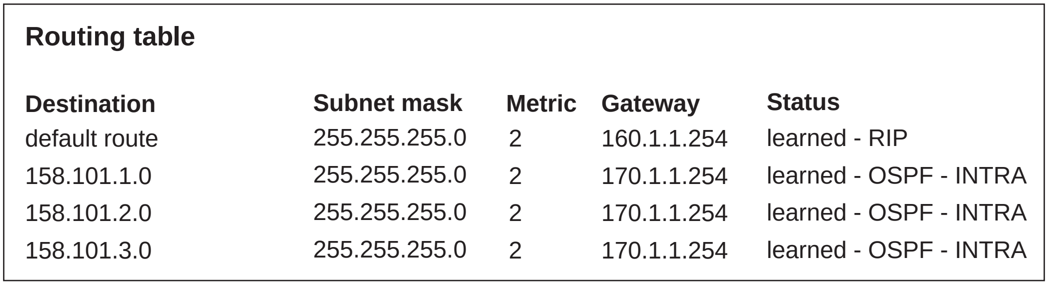

Routing table

Destination	Subnet mask	Metric	Gateway	Status
default route	255.255.255.0	2	160.1.1.254	learned - RIP
158.101.1.0	255.255.255.0	2	170.1.1.254	learned - OSPF - INTRA
158.101.2.0	255.255.255.0	2	170.1.1.254	learned - OSPF - INTRA
158.101.3.0	255.255.255.0	2	170.1.1.254	learned - OSPF - INTRA

Routing table data is updated statically or dynamically:

- **Statically** — You manually enter static routes in the routing table. Static routes are useful in environments where no routing protocol is used or where you want to override some of the routes that are generated with a routing protocol. Because static routes do not automatically change in response to network topology changes, manually configure only a small number of reasonably stable routes. Static routes do not time out, but they can be learned.
- **Dynamically** — Switches use RIP to automatically exchange routing data and to configure their routing tables dynamically. Routes are recalculated at regular intervals. This process helps you to keep up with network changes and allows the system to reconfigure routes quickly and reliably. Interior Gateway Protocols (IGPs), which operate within networks, provide this automated method.

Default Route

In addition to the routes to specific destinations, a routing table can contain a *default route*. The switch uses the default route to forward packets that do not match any other routing table entry. A default route is often used in place of static routes to numerous destinations that all have the same gateway IP address and interface number. The default route can be configured statically, or it can be learned dynamically.

Configuring IP Interfaces

To set up and use an IP interface, you must perform these tasks in the following order:

1. Configure trunks (optional).
2. Configure VLANs.
3. Establish an IP interface.
4. Administer the IP interface.

Configure Trunks (Optional)

Trunks (also known as aggregated links) allow you to combine multiple Fast Ethernet or Gigabit Ethernet links into a single high-speed link between two switches.

If you intend to use trunking on your system, configure the trunks *before* you define VLANs and IP interfaces. You must specify the anchor port (the lowest-numbered port) to associate with the trunk. For example, if ports 7 through 12 are associated with a trunk, specifying 7 to 12 defines the VLAN to include all of the physical ports in the trunk (ports 7 through 12). For more information about trunking, see Chapter 8.

Configure VLANs

Before you set up an IP interface, you must first configure a *port-based VLAN* to associate with the IP interface. For more information about VLANs, see Chapter 10.

Establish an IP Interface

To establish an IP interface, follow these steps:

1. Define your trunks (see Chapter 8).
2. Determine your interface parameters.
3. Define the IP interfaces.

Interface Parameters

Each IP routing interface has these standard characteristics:

- **IP address** — An address from the range of addresses that the Internet Engineering Task Force (IETF) assigns to your organization. This address is specific to your network and system.
- **Subnet mask** — The 32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number/subnetwork number and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnetwork part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.
- **State** — The status of the IP interface. It indicates whether the interface is available for communications (`up`) or unavailable (`down`).
- **Interface type (9300 only)** — The type of IP interface you are configuring. When you define an in-band interface, use `vlan` as the interface type. When you define the out-of-band interface, use `system` as the interface type. See Chapter 2 for information about in-band versus out-of-band management.
- **VLAN interface index (for in-band management)** — The number of the VLAN that is associated with the IP interface. When the switch prompts you for this option, the menu identifies the available VLAN indexes.

Important Consideration

Consider the following issue before you establish an IP interface:

- Before you assign IP addresses, understand how those addresses fit into the IP addressing scheme in your network. Plan for future expansion of address numbers as well.

Defining an IP Interface

After you determine the VLAN index, IP address, and subnet mask for each IP interface, you can define each interface. Use the Administration Console (`ip interface define`) or the Web Management application to define an IP interface.

You must define a VLAN that you want to associate with the IP interface before you define the interface. VLANs are described in Chapter 10.

To administer your IP interface, use the following IP-related protocols:

- ARP
- RIP
- DNS

These protocols are discussed later in this chapter.

Administering IP Interfaces

Keep these points in mind as you administer the IP interface:

- Flush the ARP cache regularly if you set the age time to 0.
- Set up a default route.

  The system uses the default route to forward packets that do not match any other routing table entry. You may want to use the default route in place of routes to numerous destinations that all have the same gateway IP address. If you do not use a default route, the system is more likely to return an `address not found` error.
- Before you can define static routes, you must define at least one IP interface. See "Defining an IP Interface" earlier in this chapter for more information. Remember the following guidelines:
  - Static routes remain in the routing table until you remove them or the corresponding interface.
  - Static routes take precedence over dynamically learned routes to the same destination.
  - Static routes are included in periodic RIP updates sent by your system.

Address Resolution Protocol (ARP)

ARP is a low-level protocol that locates the MAC address that corresponds to a given IP address. This protocol allows a device to use IP addresses to make routing decisions while it uses MAC addresses to forward packets from one hop to the next.

You do not need to implement ARP — the system has ARP capability built in, but you can manipulate and display the contents of the ARP cache.

When the device knows the IP address of the *next* hop towards the packet destination, the device translates that IP address into a MAC address before sending the packet. To perform this translation, the device first searches its *ARP cache*, which is a table of IP addresses with their corresponding MAC addresses. Each device that participates in IP routing maintains an ARP cache. See Figure 30.

Figure 30 Example of an ARP Cache

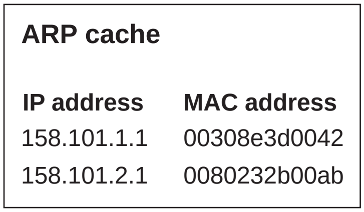

ARP cache

IP address	MAC address
158.101.1.1	00308e3d0042
158.101.2.1	0080232b00ab

If the IP address does not have a corresponding MAC address, the device broadcasts an *ARP request* packet to all the devices on the network. The ARP request contains information about the target and source addresses for the protocol (IP addresses). See Figure 31.

Figure 31 Example of an ARP Request Packet

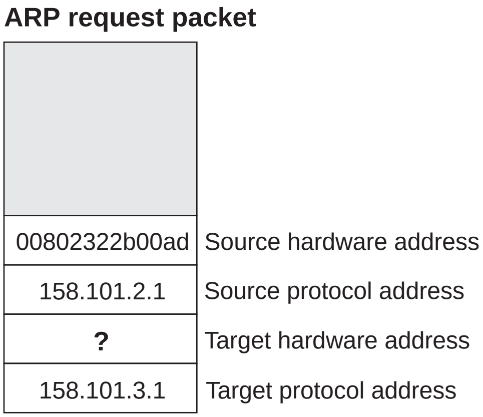

ARP request packet

00802322b00ad	Source hardware address
158.101.2.1	Source protocol address
?	Target hardware address
158.101.3.1	Target protocol address

When devices on the network receive this packet, they examine it. If their address is not the target protocol address, they discard the packet. When a device receives the packet and confirms that its IP address matches the target protocol address, the receiving device places its MAC address in the target hardware address field and sends the packet back to the source hardware address.

When the originating device receives this *ARP reply*, it places the new MAC address in its ARP cache next to the corresponding IP address. See Figure 32.

Figure 32 Example of ARP Cache Updated with ARP Reply

ARP cache

IP address	MAC address
158.101.1.1	00308e3d0042
158.101.2.1	0080232b00ab
158.101.3.1	0134650f3000

After the MAC address is known, the device can send the packet directly to the next hop.

Important Considerations

Keep the following things in mind about this protocol:

- Enter a static ARP entry when the ARP resolution does not result in an ARP entry in the cache. For example, some applications do not respond to ARP requests and, consequently, specific network operations may time out for lack of address resolution.
- Enter a static ARP entry in a test environment if your test analyzer cannot respond to an ARP request.
- Setting an ARP cache age time of zero (no aging) is useful in the middle of lengthy tests so that ARP requests do not have to be issued.

  If you do set an ARP cache age time of zero, be aware that the ARP cache can quickly grow in size and consume system resources. In this case, be sure to flush the ARP cache after your tests are complete.
- You can keep ARP cache entries if you refresh the ARP cache; otherwise, the system removes the entries after they reach their defined age time.

Routing Information Protocol (RIP)

RIP is the protocol that implements routing by using Distance Vector Algorithms (DVAs) to calculate the route with the fewest number of hops to the destination of a route request. Although your system is not a router, it assists routing by keeping known routes in its routing table. RIP is an Interior Gateway Protocol (IGP) for TCP/IP networks.

A device using RIP can operate in active or passive mode:

- **Active devices** — usually routers, broadcast RIP messages to all devices in a network or subnet and update their internal routing tables when they receive a RIP message.
- **Passive devices** — usually hosts, listen for RIP messages and update their internal routing tables, but do not send RIP messages.

On your system, RIP acts as a passive device.

Basic RIP Parameters

RIP has several parameters that you must consider when you set up RIP to use in your network. When you configure an IP interface, the system already has the RIP parameters set to the following defaults:

- **RIP-1 Mode** — learn
- **RIP-2 Mode** — learn

RIP Mode

The two available settings for both RIP modes are as follows:

- **Disabled** — The system ignores all incoming RIP packets and does not generate any RIP packets of its own.
- **Learn** — The system processes all incoming RIP packets, but it does not transmit RIP updates.

RIP-1 Versus RIP-2

Like RIP-1, RIP-2 allows the system to dynamically configure its own routing table. RIP-2 is much more flexible and efficient than RIP-1, however, because RIP-2 recognizes the multicast method, in which a message is advertised to a subset of the network. (RIP-1 uses the broadcast method, which advertises to the whole network.) RIP-2 includes a subnet mask in its header. (See Figure 33.)

If your system receives a RIP-2 packet, your system puts the route into the routing table with the subnet mask.

Figure 33 RIP-1 Versus RIP-2

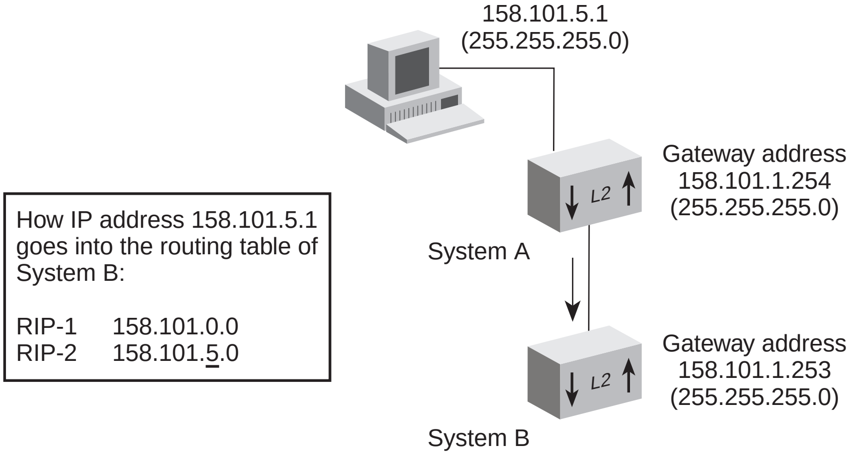

In this way, the system keeps track of the RIP-1 and RIP-2 address routes in its routing table.

Domain Name System (DNS)

The Domain Name System (DNS) client allows you to specify a hostname rather than an IP address when you perform various operations (for example, when you use `ping` or `traceRoute` to contact an IP station).

With DNS, you can specify one or more name servers that are associated with a domain name. Each name server maintains a list of IP addresses and their associated host names. When you use `ping` or `traceRoute` with a hostname, the DNS client attempts to locate the name on the name servers that you specify. When the DNS client locates the name, it resolves it to the associated IP address.

You can resolve an IP address to a host name or a host name to an IP address on a name server. Enter either the host name or the IP address; the DNS client displays the pair.

Important Considerations

When you set up DNS servers on your LAN, remember the following:

- Always set up more than one DNS name server (a primary and secondary server) so that the lookup service does not have a single point of failure.
- If your Internet service provider (ISP) changes the Classes of Internetwork Service, change the DNS settings on each host that the ISP serve.

See UNIX Network File System (NFS) documentation for information about how to create and maintain lists of domain names and IP addresses on the name servers.

See Chapter 13 and the Command Reference Guide *for information about how to use* `ping` *and* `traceRoute`*.*

Standards, Protocols, and Related Reading

This section describes how to obtain more technical information about IP.

Requests for Comments (RFCs)

Documents called Requests for Comments (RFCs) contain information about the entire set of protocols that comprise IP. Some of the RFCs that pertain to the discussions in this chapter are:

- **RFC 791** — Internet Protocol
- **RFC 1219** — Subnetwork Numbers
- **RFC 1058** — RIP
- **RFC 1723** — RIP Version 2
- **RFC 2400** — Internet Official Protocol Standards

To obtain copies of Internet RFCs and proposed standards, visit the Internet Engineering Task Force (IETF) Web site:

`http://www.ietf.org`

Standards Organizations

Standards organizations ensure interoperability, create reports, and recommend solutions for communications technology. The most important standards groups are:

- International Telecommunications Union (ITU)
- American National Standards Institute (ANSI)
- International Standards Organization (ISO)
- Institute of Electrical and Electronic Engineers (IEEE)
- Internet Engineering Task Force (IETF)
- National Institute of Standards and Technology (NIST)

Related Reading

For more information about the IP suite, see the following books:

- *High Speed Networks: TCP/IP and ATM Design Principles* by William Stallings, Prentice Hall, 1998
- *Local Area Networks: Architectures and Implementations* by James Martin, Prentice Hall, 1994
- *Internetworking with TCP/IP: Principles, Protocols, and Architecture* by Douglas Comer, Prentice Hall, 1995

12

IP MULTICAST FILTERING WITH IGMP

The Internet Group Management Protocol (IGMP) gives your system a way to forward IP multicast application traffic only to ports that require it and filter it on other ports to increase bandwidth efficiency in the network. This chapter provides an overview, guidelines, and other key information about IGMP functions and their effect on IP multicast traffic.

The chapter covers these topics:

- Overview
- Key Concepts
- Configuring IGMP in Your System
- Key Implementation Guidelines
- Processing IP Multicast Packets
- Effects of MAC Address Aliasing
- Operating as the Querier
- Locating Multicast Routers
- Aging the IGMP Tables
- Standards, Protocols, and Related Reading

You can manage IGMP commands in either of these ways:

- *From the* `bridge multicast igmp` *menu of the Administration Console. See the* Command Reference Guide *for details*.
- *From an SNMP management application (not included with the system) using a private 3Com IGMP Snooping MIB, which is available from the 3Com Software Library on the Web.*

Overview

To transport their content to a community of network users, bandwidth-hungry applications often generate packets in the IP multicast format. Many network standards and protocols have been designed to create an efficient delivery system for IP multicast traffic from the source (usually a server) to the destinations (users). IGMP is one of these protocols.

If IGMP functions are disabled or not present in a Layer 2 switch, the switch floods all IP multicast packets to all ports — that is, it operates in compliance with the *IEEE 802.1D MAC Bridges* base standard. If IGMP functions are present and enabled, a switch can forward IP multicast traffic only to ports that require it and filter it on other ports.

Defined in Internet RFC 1112 and RFC 2236, IGMP performs two main functions in a Layer 2 switch: *snooping* and *querying*. Descriptions of these functions and how they work together are explained later in this chapter.

Benefits

Support for IGMP in Layer 2 switches benefits your network in many ways:

- IGMP reduces the amount of bandwidth that an IP multicast stream would otherwise occupy at the edge of an IP multicast delivery tree; it is especially useful in flat network designs (large broadcast domains with cascading switches).
- IGMP reduces the amount of unwanted traffic that a host encounters on its Ethernet segment, which may be critical for users with low-speed connections to the network.
- IGMP requires minimal configuration in network devices and hosts. For example:
  - Snooping and querying functions can be easily enabled. The querying function requires that the system have at least one IP interface configured.
  - IP-capable end stations do not usually require any special configuration because IGMP is already part of the IP protocol stack.
- Because more IP multicast applications are available each year, support for IGMP in switches helps prolong the life span of existing network topologies and available bandwidth.

To understand the fundamental benefit that IGMP provides for users attached to a switch, see Figure 34.

Figure 34 IP Multicast Traffic Flow Before and After IGMP Snooping

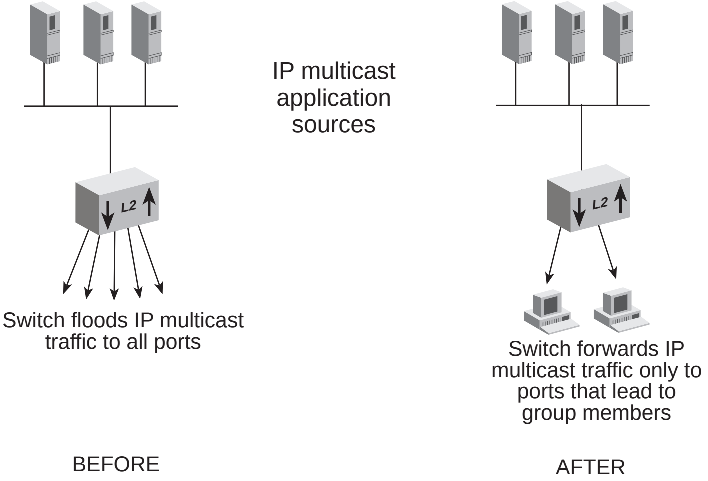
IP multicast application sources
L2
L2
Switch floods IP multicast traffic to all ports
Switch forwards IP multicast traffic only to ports that lead to group members
BEFORE
AFTER

Key Concepts

IGMP plays a specific role in the overall delivery process for IP multicast traffic. Before you modify IGMP parameters in your system, review the following key concepts about IP multicast packets.

Devices That Generate IP Multicast Packets

Application sources (usually servers) generate IP multicast packets as the way to deliver their information (such as a video stream) to interested hosts (such as PC end stations).

Network devices generate IP multicast packets as the way to communicate with each other to establish a delivery path. These packets are issued by specific supporting protocols, such as IGMP.

Group Addresses and Group Members

An IP multicast packet differs from a unicast packet by the presence of a *multicast group address* in the destination address field of the IP header. Each application uses a unique group address, and hosts refer to these group addresses when they tell network devices which IP multicast transmissions they want to receive. In doing so, hosts become *group members*. Hosts can join and leave one or more groups at any time.

Communication Protocols

To coordinate an efficient, loopfree delivery path for IP multicast packets, certain protocols are used among network devices and hosts:

- A multicast routing protocol such as Distance-Vector Multicast Routing Protocol (DVMRP) is used between routers (if routers exist between sources and group members).
- A protocol such as IGMP is used between routers, switches, and hosts in each subnetwork or broadcast domain.

Figure 35 Protocols That Coordinate the Delivery of IP Multicast Traffic

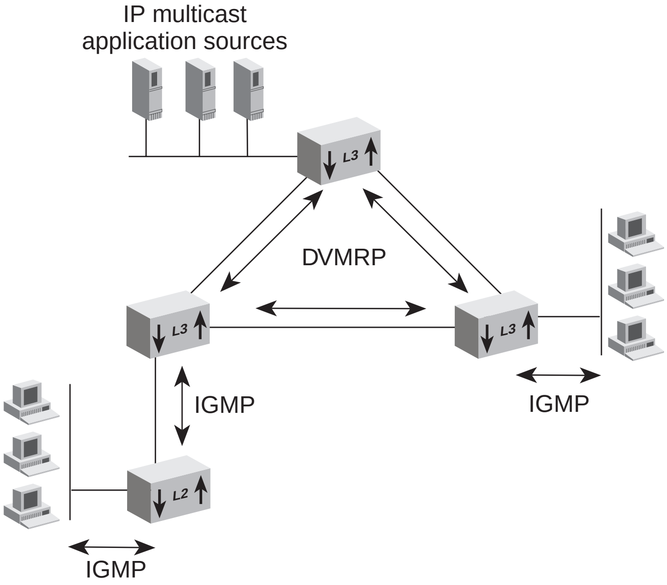

Routers are not required for transmission of IP multicast packets between sources and group members. Compare Figure 34 and Figure 35; both represent valid designs in which IGMP can help conserve bandwidth.

IP Multicast Delivery Process

Even though there may be several, perhaps thousands of, intended recipients for a given IP multicast transmission, only one copy of a given packet is generated at the source. (This contrasts with a unicast approach, which would generate one copy per recipient.) The single copy of each IP multicast packet travels until the path to reach group members diverges, at which point the packet is replicated to ensure that one copy of the packet continues on each branch in the delivery tree. Thus, a significant benefit of the IP multicast delivery process is bandwidth conservation.

How Routers and Switches Use IGMP

Routers and switches use IGMP in similar ways:

- A router uses IGMP to determine which routing interfaces lead to group members.
- A switch uses IGMP to determine which port segments lead to group members.

Routers and switches both construct filters on ports that do not require group traffic to be forwarded. On each device, one group's traffic may be forwarded to one set of ports and another group's traffic may be forwarded to a different set of ports.

Tracking Group Member Locations

The ability to detect the location of group members is a product of two IGMP functions — querying and snooping — working together to construct individual delivery trees for each IP multicast group.

In each subnetwork or broadcast domain (VLAN), all switches and routers perform snooping (for their own connections), but only one of the devices needs to perform the querying. If there are multiple devices (routers and switches) in the subnetwork or broadcast domain that support querying, the one with the lowest IP address is elected as the *querier*.

The querier periodically sends a *query message* to all hosts on the subnetwork or broadcast domain and requests that they reply with the IP multicast groups for which they want to receive traffic.

A host responds to a query by sending an *IGMP report.* The querier as well as IGMP-capable devices between hosts and the querier *snoop* on the report's content to track which hosts (and the associated ports) belong to which groups.

If, after a certain period of time, a router or switch does not receive responses from any host on a given interface or port for a particular IP multicast group, the router or switch *prunes* the interface or port from the delivery tree to conserve network bandwidth.

If a switch has downstream group members, an upstream router, and upstream sources, and it is elected as the querier, it must forward host reports to the upstream router to ensure that it continues to receive IP multicast traffic from the router.

How Hosts Use IGMP

Each host uses IGMP to communicate with the querier in a few different ways.

Host Membership Reports

Hosts transmit *Host Membership Reports* (hereafter called *IGMP reports*) in response to queries. A host sends a separate report for each group that it wants to join or to which it currently belongs. Hosts do not send reports if they are not or do not want to be group members.

Join Message

Rather than wait for a query, a host can also send an IGMP report on its own initiative to inform the querier that it wants to begin receiving a transmission for a specific group. This is called a *join* message. The benefit is faster transmission linkages.

Leave-Group Messages

Leave-group messages are a type of host message defined in IGMP version 2. If a host wants to leave an IP multicast group, it issues a leave-group message. Upon receiving such a message, the querier determines whether that host is the last group member on the subnetwork by issuing a *group-specific query*.

Leave-group messages lower *leave latency* — that is, the time between when the last group member on a given subnetwork or segment sends a report and when a router or switch stops forwarding traffic for that group. This process conserves bandwidth. The alternative is for the router or switch to wait for an aging period to expire before it ceases to forward the group traffic.

Report Suppression and Effect on Switch Activity

If host A hears an IGMP report from host B for the same group, host A suppresses (does not transmit) its own report for that group.

This approach was designed to conserve bandwidth, and it works well with routers because a routing interface only needs to know is that there is at least one group member in a subnetwork for it to forward group traffic onto that subnetwork.

However, because a switch that is not operating as the querier must forward IGMP reports to the querier, a switch must be able to track which ports lead to the querier, to routers, and to group members, so that it can forward IGMP reports only to the querier.

If the switch flooded IGMP reports, hosts on other segments would suppress their own reports for identical groups, which would cause the switch to set overly restrictive filters. Restricted forwarding of IGMP reports is necessary to allow the switch to receive IGMP reports from at least one host per group on each of its ports.

Configuring IGMP in Your System

Your system supports IGMP version 1 (RFC 1112) and version 2 (RFC 2236). You can manage the following IGMP parameters:

- Enable or disable the snooping function and the querying function
  - Both settings apply to all ports.
  - Both settings are disabled by default.
  - If enabled, the querying requires that you configure an IP interface. See "Operating as the Querier" later in this chapter.
- Display a command configuration summary
- Display VLANs with active snooping activity
- Display group and port information per VLAN
- Display the designated querier per VLAN
- Display IP multicast router ports per VLAN
- Display the port in the VLAN that last received a query

With snooping and querying enabled, your switch:

- Builds for each VLAN a table of member ports per IP multicast group and adjusts the table over time through an aging mechanism as well as by processing IGMP messages.
- Sets per-port filters that allow targeted forwarding of IP multicast group traffic.
- Determines which ports lead to routers. See "Locating Multicast Routers" later in this chapter.
- Determines which port leads to the querier (unless the system is the querier). See "Operating as the Querier" later in this chapter.
- Supports a private IGMP Snooping MIB (`3cigmpSnoop.mib`).

Key Implementation Guidelines

Consider these points when you configure IGMP options in your system:

- IGMP snooping and querying works for *IP* multicast packets only. The system floods other protocol-based multicast packets to all ports in compliance with the IEEE 802.1D standard.
- All IGMP-capable devices in a subnetwork or broadcast domain must perform snooping to track group membership activity on their own respective ports, but only one device needs to be the querier for all. Thus, in a single switch, you can enable snooping and disable querying as long as another device in the subnetwork or broadcast domain can be the querier.
- To maximize the effectiveness of IGMP in a flat network design or large broadcast domain that includes IP multicast sources, the querier should be positioned as close to the source of IP multicast traffic (usually servers) as possible. You can accomplish this by enabling the query function only on certain devices.
- Because some IP multicast applications transmit a large number of unsolicited packets or may require security protection, 3Com recommends that you place IP multicast sources upstream from a Layer 3 switch or router.
- If you have configured Open VLAN mode and IP multicast packets are tagged (IEEE 802.1Q format), then the IGMP tables in each VLAN share information with each other. VLANs do not form barriers in the flow of IP multicast traffic, even though they form separate broadcast domains. IGMP operates as if there was a single broadcast domain. However, if you have configured Open VLAN mode and the packets are not tagged, then VLANs do form barriers and the IP multicast traffic is restricted.
- If you have configured Closed VLAN mode, then the IGMP tables do not share information and forwarding of IP multicast packets is restricted to the ports in the VLAN on which the IP multicast packet arrives. If all VLANs contain potential IP multicast group members, then you would need to have a router connection to each VLAN to ensure IP multicast packet delivery to those group members.
- Ensure that any rate limit that you implement with the `bridge port multicastLimit` command does not interfere with the flow of IP multicast traffic. If you select the `BcastOnly` option, the rate limit does not affect multicast packets. For more information about the bridge port multicast limit feature, see Chapter 6.

- To reduce the impact of MAC address aliasing, verify that your IP multicast applications do not use binary group addresses in the range [224 – 239]. [0,128].0.x, where x equals 0 – 255. See "Effects of MAC Address Aliasing" later in this chapter for more information.
- Your switch supports both IGMP version 1 and 2. For maximum benefit, verify that the IP stack in your host endstations also supports both IGMP version 1 and 2.
- If a resilient link pair exists between two switches that are downstream from the IP multicast source and IGMP querier and the main link in the pair fails over to its standby link, the IP multicast transmission to group members is temporarily interrupted. The transmission resumes after hosts send IGMP join messages or after the next query travels down the [now active] standby link and hosts respond with IGMP reports. Queries are sent every 125 seconds. Depending on when the last query was sent prior to the link failover, the interruption can last up to 2 minutes.

Processing IP Multicast Packets

Table 22 summaries how your system processes various types of IGMP and other IP multicast packets.

Table 22 How the System Processes IP Multicast Packets

Packet Type	Is Forwarded To*
IGMP Membership Reports	Ports in the broadcast domain that lead to multicast routers Port that leads to the querier (if another device is the querier)
IGMP Queries	All ports in the broadcast domain
IGMP Leave-Group Messages	All ports in the broadcast domain
Packets with addresses [224 – 139].[0,128].0.x where x = 0 – 255	All ports in the broadcast domain (See "Effects of MAC Address Aliasing" later in this chapter.)
Packets addressed to known (registered) IP multicast group	Ports on which IGMP membership reports for that group have been heard Ports that lead to multicast routers Port that leads to the querier (if another device is the querier)
Packets addressed to unknown IP multicast group (group for which no host has registered)	No ports, except those that lead to multicast routers and the querier (if another device is the querier).

* Except for the port on which the packet originated.

Some ports may not be available for carrying traffic. Two examples are ports that have been administratively disabled or ports that the Spanning Tree Protocol has prevented from being in the forwarding state.

Effects of MAC Address Aliasing

Operating as a Layer 2 device, your system filters IP multicast traffic by referring to hexidecimal MAC addresses that correspond to binary IP multicast group addresses.

A multicast MAC address is created by selecting only the low order 23-bits in the Class D binary IP address, translating that portion to hexidecimal format, and attaching it to a standard set of bits that signify it is a multicast packet (01-00-5E). For example, IP address 224.10.8.5 becomes MAC address 01-00-5E-0A-08-05.

Because only a portion of the binary IP address is translated, several different binary addresses can map to the same hexidecimal MAC address. This situation is called *MAC address aliasing.*

Because the system cannot distinguish such packets, MAC address aliasing has two main implications:

- Some packets are forwarded to more ports than actually require it.

  For example, if requests for multicast group 226.1.2.3 are registered on port 1 and requests for group 227.1.2.3 are registered on port 2, these IP addresses map to the same MAC address and the system forwards traffic for both groups to both ports.

- Packets with certain addresses can never be filtered by IGMP.

  In most cases, such packets are routing protocol advertisements that use addresses from the block of permanent reserved addresses that is administered by the Internet Assigned Numbers Authority (IANA). To ensure these advertisements make their way through the network, it is important that these packets do not get filtered by IGMP. However, if an IP multicast application uses a group address that maps (due to MAC address aliasing) to one of these permanent addresses, these packets cannot be filtered by IGMP either.

Important Considerations

- To reduce the effects of MAC address aliasing, verify that your IP multicast applications do not use binary group addresses in the range [224 – 239]. [0,128].0.x, where x equals 0 – 255.
- See Table 23 for several examples of permanent reserved addresses. For a complete and current list, visit the Web site of the Internet Assigned Numbers Authority (IANA) at:

  `http://www.iana.org`

Table 23 Examples of Class D Permanent Address Assignments

Address	Meaning
224.0.0.0	Base Address (Reserved)
224.0.0.1	All systems on this subnet
224.0.0.2	All routers on this subnet
224.0.0.4	All DVMRP routers
224.0.0.5	All OSPF routers
224.0.0.6	All OSPF designated routers
224.0.0.7	All ST routers
224.0.0.8	All ST hosts
224.0.0.9	All RIP version 2 routers
224.0.0.11	Mobile agents
224.0.0.12	DHCP server/relay agent
224.0.0.13	All PIM routers
224.0.0.14	RSVP, Encapsulation
224.0.0.15	All CBT routers

Operating as the Querier

For your system to offer itself as a potential IGMP querier for its subnetwork or broadcast domains (VLANs), you must:

1 Enable the IGMP snooping option.

The switch cannot perform as the querier if snooping is disabled.

2 Enable the IGMP querying option.

3 Configure an IP interface (IP address) that the system can insert as the source address of query packets. You can do *either* of the following:

- Configure an in-band IP interface. On the Switch 9300, in-band interfaces are labeled with `vlan` as the *type* of IP interface. In either the Switch 3900 or Switch 9300, if you have two in-band interfaces, the system uses the IP address of index 1.
- Use the `bridge multicast igmp queryIpAddress` command to configure a unique IP address that is only used in IGMP queries.

The destination address in query packets is 224.0.0.1, the “all hosts on this subnetwork” IP multicast address.

As the querier, the switch sends general queries every 125 seconds and group-specific queries when prompted by host leave-group messages.

Locating Multicast Routers

A switch must be able to identify which of its ports are connected to multicast routers so that it can forward appropriate IP multicast traffic to them. For example:

- If the switch is operating as the querier, any upstream router (the router that leads back toward an IP multicast source) also needs to see IGMP reports so it can decide whether to begin, stop, or continue forwarding group traffic on the subnetwork that includes the switch.
- Downstream routers need to receive all IP multicast traffic because other group members may be attached to them.

Your system records which of its ports lead to IP multicast routers by snooping on protocol advertisements that are sent by the routers. Your system can recognize the advertisements of the following multicast routing protocols:

- Distance Vector Multicast Routing Protocol (DVMRP)
- Multicast Open Shortest Path First (MOSPF)
- Protocol Independent Multicast (PIM) version 1
- Protocol Independent Multicast (PIM) version 2

The system uses a 100-second interval to age the records it keeps for multicast router locations. Protocol advertisements are sent much more frequently than this.

Aging the IGMP Tables

If your system receives no host reports for a given group on a given port within a certain period of time (the aging interval), it ages that entry in its IGMP tables and sets a filter for that group on that port.

The aging interval is the period from which the last query is sent until the time when the system stops forwarding traffic on that port. The aging interval is determined in the following way: multiply the query interval (125 seconds) by two and then add one query response interval (10 seconds). Thus, it is approximately 4.5 minutes.

Standards, Protocols, and Related Reading

The following standards apply to IGMP and how your system processes multicast packets:

- *IEEE 802.1D Media Access Control (MAC) Bridges*

  A base standard that specifies requirements for transparent bridging. To obtain copies of standards, register for an on-line subscription the Institute of Electrical and Electronics Engineers (IEEE) Web site:

  `http://www.ieee.org`

- RFC 1112: *Host Extensions for IP Multicasting*. S. Deering, July 1986.
- RFC 2236: *IGMP version 2*. W. Fenner, November 1997.

  To obtain copies of Internet RFCs and proposed standards, visit the Internet Engineering Task Force (IETF) Web site:

  `http://www.ietf.org`

13 DEVICE MONITORING

This chapter provides descriptions and key operational information about device monitoring features and tools of your SuperStack® II Switch 3900 and Switch 9300 systems. The chapter covers these topics.

- Device Monitoring Overview
- Key Concepts and Tools
- Statistical Baselines
- Roving Analysis
- Ping
- traceRoute
- Simple Network Management Protocol (SNMP)
- Remote Monitoring (RMON)
- Management Information Base (MIB)

You can manage baselining, roving analysis, and SNMP in either of these ways:

- *From the appropriate menus of the Administration Console. See the* Command Reference Guide*.*
- *From the appropriate folders of the Web Management software. See the* Web Management User Guide*.*

RMON MIBs are accessible only through applications that implement SNMP.

Device Monitoring Overview

You can use the device monitoring features and tools described in this chapter to analyze your network periodically and to identify potential network problems before they become serious. To identify potential problems in your network, use:

- Baselining
- Roving analysis
- RMON information

To test and validate paths in your network, use:

- Ping
- traceRoute

The SNMP protocol and the MIB are also described in this chapter to give you some background on how performance data is collected about the network.

Key Concepts and Tools

This section describes key concepts and tools for monitoring your device.

Administration Console

The Administration Console provides you with access to all the features of your system. It also provides you access to some of the device monitoring tools, such as:

- Statistical Baselines
- Roving Analysis
- Ping
- traceRoute
- snapshot

You access the Administration Console locally through the serial terminal or modem port on the system or remotely via a Telnet connection. See Chapter 2 for more information.

Web Management Tools

The Web Management suite of monitoring and configuration tools provides you access to the system remotely via the Internet. See the *Web Management User Guide* for more information.

Network Management Platform

Use the network management platform to view the health of your overall network. With the platform, you can understand the logical configuration of your network and configure views of your network to see how devices work together and watch how traffic changes over time. The network management platform that supports your Transcend® Network Control Services software installation can provide valuable troubleshooting tools.

SmartAgent Embedded Software

Traditional SNMP management places the burden of collecting network management information on the management station. In this traditional model, software agents collect information about throughput, record errors or packet overflows, and measure performance based on established thresholds. Through a polling process, agents pass this information to a centralized network management station whenever they receive an SNMP query. Management applications then make the data useful and alert the user if there are problems on the device.

For more information about traditional SNMP management, see "Simple Network Management Protocol (SNMP)" later in this chapter.

SmartAgent® software uses remote monitoring and is self-monitoring — that is, it collects and analyzes its own statistical, analytical, and diagnostic data. Use it to conduct network *management by exception* — that is, management in which you are notified only if a problem occurs. Management by exception is unlike traditional SNMP management, in which the management software collects *all* data from the device through polling.

Other Commonly Used Tools

These commonly used tools can help you troubleshoot your network:

- Built-in system features such as baselining, remote monitoring (RMON), and creating snapshots.
- Network software tools, such as ping and traceroute, that can help you to troubleshoot and test your system.
- Analyzers connected to your system's roving analysis ports to monitor devices.
- Network utility software, such as Telnet, FTP, and TFTP, to troubleshoot, configure, and upgrade your system.
- Tools such as cable testers for working on physical network problems.

Statistical Baselines

Normally, the system starts to compile statistics for MACs and ports at power on. With baselining, you can view statistics that are compiled over the period of time since a baseline was set. By viewing statistics relative to a baseline, you can more easily evaluate recent activity in your system or on your network.

Important Considerations

- Baselining is maintained across Administration Console sessions. Statistics that you view after setting the baseline indicate that they are relative to the baseline. To view statistics as they relate only to the most recent power-on, disable the baseline.
- Baselining affects the statistics that are displayed for Ethernet ports and bridges.

Setting a Baseline

You can reset the baseline counters to zero (0). The system maintains the accumulated totals since power-on. The baseline is time-stamped.

Enabling or Disabling Baselines

When you reenable a baseline, the counters return to the values that accumulated since the most recent baseline that you set. Disabling a baseline returns the counters to the total accumulated values since the last power-on.

Displaying the Current Baseline

You can display the current baseline to verify when the baseline was last set and to determine if you need a newer baseline for viewing statistics.

Roving Analysis

Roving analysis is the mirroring of Fast Ethernet and Gigabit Ethernet port traffic to another port within the system. This second port has an external RMON-1 probe or network analyzer attached such as the 3Com Transcend Enterprise Monitor. Through the probe, you can monitor traffic on any switched segment. Figure 36 shows a sample configuration.

- The port with the analyzer attached is called the *analyzer port*.
- The port that is being monitored is called the *monitor port*.

Figure 36 Connecting an Analyzer to the System

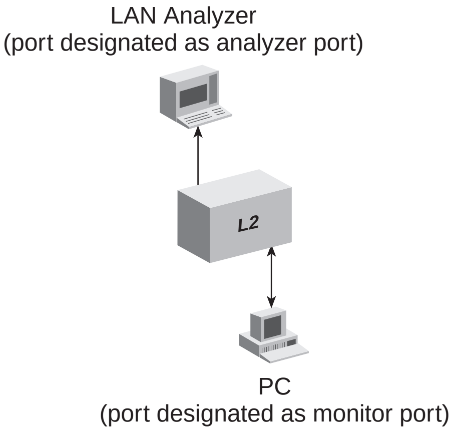
LAN Analyzer
(port designated as analyzer port)
L2
PC
(port designated as monitor port)

The purpose of roving analysis is to:

- Analyze traffic loads on each segment so that you can continually optimize your network loads by moving network segments
- Troubleshoot switched network problems (for example, to find out why a particular segment has so much traffic)

When you set up a roving analysis configuration, the system copies both transmit and receive port data and forwards it to the port on which the network analyzer is attached — without disrupting the regular processing of the packets.

You can configure one analyzer port and one monitor port on a system at a time.

Key Guidelines for Implementation

To enable the monitoring of ports on a system, follow these general steps:

1. Add the port on which you want to attach the network analyzer.
2. Start roving analysis.
   a. Select the port that you want to monitor.
   b. Enter the analyzer port's MAC address.

The system provides commands to add and remove (define and undefine) the analyzer port, to display the current analyzer and monitor ports, and to start and stop analysis.

See the *Command Reference Guide* for details.

Important Considerations

- The network analyzer cannot be located on the same bridge port as the port that you want to monitor.
- For more accurate analysis, attach the analyzer to a dedicated port instead of through a repeater.
- When the analyzer port is set, it cannot receive or transmit any other data. Instead, it receives only the data from the port(s) to be monitored.
- If Spanning Tree Protocol was enabled on the analyzer port, it is automatically disabled. When the analyzer port is undefined, the port returns to its configured Spanning Tree state and functions as it did before it was set as an analyzer port.
- When you configure a port that is part of a virtual LAN (VLAN) as an analyzer port, the port is removed from all VLANs of which it is a member. When you remove the analyzer port, it becomes a member of the default VLAN. You have to manually add it back to its original VLANs.
- You cannot use roving analysis to monitor trunk ports or resilient link ports.

Ping

The ping feature is a useful tool for network testing, performance measurement, and management. It uses the Internet Control Message Protocol (ICMP) echo facility to send ICMP echo request packets to the IP destination that you specify.

When a router sends an echo request packet to an IP station using ping, the router waits for an ICMP echo reply packet. The response indicates whether the remote IP is available, unreachable, or not responding.

Important Consideration

When you specify a host name with ping, the host name and its associated IP address must be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See "Domain Name System (DNS)" in Chapter 11.

Using ping

The system provides two ping functions:

- **ping** — Uses the hostname or IP address to ping a host with default options
- **advancedPing** — Uses the hostname or IP address to ping a host with the advanced ping options that you specify

Ping Responses

This list gives the possible responses to a ping:

- If the host is reachable, the system displays information about the ICMP reply packets and the response time to the ping. The amount of information depends on whether the quiet option is enabled or disabled.
- If the host does not respond, the system displays the ICMP packet information and this message: `Host is Not Responding.` (You may see this message if you have not configured your gateway IP address.)
- If the packets cannot reach the host, the system displays the ICMP packet information and this message: `Host is Unreachable.` A host is unreachable when there is no route to that host.

Strategies for Using Ping

Follow these strategies for using ping:

- Ping devices when your network is operating normally so that you have a performance baseline for comparison.
- Ping by *IP address* when:
  - You want to test devices on different subnetworks. This method allows you to ping your network segments in an organized way, rather than having to remember all the hostnames and locations.
  - Your DNS server is down and your system cannot look up host names properly. You can ping with IP addresses even if you cannot access hostname information.
- Ping by *hostname* when you want to identify DNS server problems.
- To troubleshoot problems involving large packet sizes, ping the remote host repeatedly, increasing the packet size each time.

traceRoute

Use the traceRoute feature to track the route of an IP packet through the network. TraceRoute information includes all of the nodes in the network through which a packet passes to get from its origin to its destination. The traceRoute feature uses the IP time-to-live (TTL) field in User Datagram Protocol (UDP) probe packets to elicit an ICMP Time Exceeded message from each gateway to a particular host.

Using traceRoute

The system provides two traceRoute functions:

- **traceRoute** — Uses the hostname or IP address to trace a route to a host with default options
- **advancedTraceRoute** — Uses the hostname or IP address to trace a route to a host with the advanced traceRoute options that you specify

traceRoute Operation

To track the route of an IP packet, the traceRoute feature launches UDP probe packets with a small TTL value and then listens for an ICMP Time Exceeded reply from a gateway. Probes start with a small TTL of 1 and increase the value by 1 until one of the following events occurs:

- The system receives a Port Unreachable message, indicating that the packet reached the host.
- The probe exceeds the maximum number of hops. The default is `30`.

At each TTL setting, the system launches three UDP probe packets, and the traceRoute display shows a line with the TTL value, the address of the gateway, and the round-trip time of each probe. If a probe answers from different gateways, the traceRoute feature prints the address of each responding system. If no response occurs in the 3-second time-out interval, traceRoute displays an asterisk (*) for that probe. Other characters that can be displayed include the following:

- `!N` — Network is unreachable
- `!H` — Host is unreachable
- `!P` — Protocol is unreachable
- `!F` — Fragmentation is needed
- `!<n>` — Unknown packet type

Simple Network Management Protocol (SNMP)

The Simple Network Management Protocol (SNMP), one of the most widely used management protocols, provides the means for management communication between network devices and your management workstation across TCP/IP networks.

See Figure 3 in Chapter 2 to review where SNMP fits in the Open System Interconnection (OSI) reference model for the network environment.

Most management applications, including 3Com Transcend Status Watch applications, require SNMP to perform their management functions.

Manager and Agent Operations

SNMP communication requires a *manager* (the station that is managing network devices) and an *agent* (the software in the devices that communicates with the management station). SNMP provides the language and the rules that the manager and agent use to communicate.

Managers can discover agents:

- Through autodiscovery tools on Network Management Platforms (such as HP OpenView Network Node Manager)
- When you manually enter IP addresses of the devices that you want to manage

For agents to discover their managers, you must provide the agent with the IP address of the management station or stations.

Managers send requests to agents (either to send information or to set a parameter), and agents provide the requested data or set the parameter. Agents also communicate with managers (without being requested by the managers) through trap messages, which inform the manager that certain events have occurred.

You can use SNMP to manage the Switch 9300 system either in-band or out-of band. You can use SNMP to manage the Switch 3900 system in-band only. See Chapter 2 for more information.

IP Address Assignment

For the manager and agent to be able to communicate with one another, you need to assign an IP address. How you do so depends on how the management station is attached. On the Switch 3900, you can use an in-band Ethernet port. On the Switch 9300, you can use the out-of-band Ethernet port or an in-band Ethernet port.

- **In-band (both)** — Use the `ip interface define` command to assign the IP address for the in-band Ethernet port. On the Switch 9300, you must supply `vlan` as the interface type.
- **Out-of-Band (Switch 9300 only)** — Use the `ip interface define` command to assign an IP address to the out-of band Ethernet port. On the Switch 9300, you must supply `system` as the interface type.

SNMP Messages

SNMP supports queries (called *messages*) that allow the protocol to transmit information between the managers and the agents. Types of SNMP messages:

- **Get** and **Get-next** — The management station requests an agent to report information.
- **Set** — The management station requests an agent to change one of its parameters.
- **Get Responses** — The agent responds to a Get, Get-next, or Set operation.
- **Trap** — The agent sends an unsolicited message informing the management station that an event has occurred.

Management Information Bases (MIBs) define what can be monitored and controlled within a device (that is, what the manager can Get and Set). An agent can implement one or more groups from one or more MIBs. See “Management Information Base (MIB)” later in this chapter for more information.

Trap Reporting

Traps are events that devices generate to indicate status changes. Every agent supports some trap reporting. You must configure trap reporting at the devices so that these events are reported to your management station to be used by the network management platforms (such as HP OpenView Network Node Manager or SunNet Manager). To do this, you configure the system to send traps to one or more IP destination addresses. See "Administering SNMP Trap Reporting" later in this chapter.

You do not need to enable all traps to effectively manage a switch. To decrease the burden on the management station and on your network, you can limit the traps reported to the management station.

Table 24 lists the MIBs that are supported by the Switch 3900 and the Switch 9300. More traps may be defined in vendors' private MIBs.

Table 24 Traps Supported by SNMP

Trap No.	Trap Name	Source	Indication
1	Cold Start	MIB II	The agent has started or been restarted.
2	Link Down	MIB II	The status of an attached communication interface has changed from *up* to *down*.
3	Link Up	MIB II	The status of an attached communication interface has changed from *down* to *up*.
4	Authentication Failure	MIB II	The agent received a request from an unauthorized manager.
5	New Root	Bridge MIB	The sending agent has become the new root of the Spanning Tree.
6	Topology Change	Bridge MIB	Any transitions from the Learning state to the Forwarding state or from the Forwarding state to the Blocking state of the bridge-configured ports.
7	System Overtemperature	3C System MIB	The system temperature exceeds a certain threshold.
8	Power Supply Failure	3C System MIB	The trap that is generated when a power supply unit fails.
26	Rising Alarm Trap	RMON MIB	An alarm entry crosses its rising threshold.
27	Falling Alarm Trap	RMON MIB	An alarm entry crosses its falling threshold.
28	Response Received Trap	POLL MIB	A disabled device begins responding.

Table 24 Traps Supported by SNMP (continued)

Trap No.	Trap Name	Source	Indication
29	Response Not Received Trap	POLL MIB	An enabled device stops responding.
30	Resilient Link Switch Trap	3C Resilient link MIB	This trap is generated in response to either of these conditions: ■ If one of the ports in a resilient link pair changes state, which causes a switchover of the active port. ■ If there was no active port and a port has become active.
31	Resilient Link No Switch Trap	3C Resilient link MIB	This trap is generated when one of the ports in a resilient link pair changes state but does *not* cause a switchover of the active port. If such a switchover occurs, trap 30 is generated.
34	Port Monitor Trap	3C System MIB, Port Monitor Table	This trap is generated when the system has exceeded the excessive collision, multiple collision, late collision, runt packet, or FCS error thresholds. This could be due to a duplex mismatch or a malfunctioning device on the port. See Chapter 5 for details. Applies to Switch 3900 only.

To minimize SNMP traffic on your network, you can implement trap-based polling, which allows the management station to start polling only when it receives certain traps. Your management applications must support trap-based polling for you to take advantage of this feature.

Security

SNMP uses community strings as a form of management security. To enable management communication, the manager must use the same community strings that are configured on the agent. You can define both read and read/write community strings.

Because community strings are included unencoded in the header of a User Datagram Protocol (UDP) packet, packet capture tools can easily access this information. As with any password, change the community strings frequently.

Setting Up SNMP on Your System

To manage your system from an external management application, you must configure SNMP community strings and set up trap reporting, as described in this section.

You must also assign an IP address to either the system out-of-band Ethernet port or an in-band Ethernet port, depending on where the management station is attached. See Chapter 2 for more information.

You can manage the system using an SNMP-based external management application. This application (called the SNMP manager) sends requests to the system, where they are processed by the internal SNMP agent.

You can gain access to the Remote Monitoring (RMON) capabilities of your system through SNMP applications such as Transcend® Network Control Services software. See "Remote Monitoring (RMON)" later in this chapter for information about the RMON capabilities of your system.

The SNMP agent provides access to the collection of information about your system. (You can view many system-specific settings.) Your views of MIB information differ depending on the system SNMP management method that you choose. In addition, you can configure a system SNMP agent to send traps to an SNMP manager to report significant events.

Access to system information through SNMP is controlled by community string, discussed next.

Displaying Community Strings

You can display the currently assigned SNMP community strings.

Configuring Community Strings

A community string is an octet string, included in each SNMP message, that controls access to system information. The system SNMP agents internally maintain two community strings that you can configure:

- *Read-only* community strings with the default *public*
- *Read-write* community strings with the default *private*

When an SNMP agent receives an SNMP request, the agent compares the community string in the request with the community strings that are configured for the agent:

- SNMP *get*, *get-next*, and *set* requests are valid if the community string in the request matches the agent's *read-write* community.
- SNMP *get* and *get-next* requests are valid if the community string in the request matches the agent's *read-only* community string or read-write community string.

Community string length

When you set a community string, you can specify any value up to 48 characters long.

Administering SNMP Trap Reporting

For network management applications, you can use the Administration Console to manually administer the trap reporting address information. (See the *Command Reference Guide* for descriptions of the `snmp trap` commands.)

- **Displaying Trap Reporting Information** — When you display the trap reporting information, the system displays the various SNMP traps and their currently configured destinations.
- **Configuring Trap Reporting** — You can add new trap reporting destination configurations and modify existing configurations. You can define up to 10 destination addresses and the set of traps that are sent to each destination address.

The trap numbers that you enter direct the system to send the corresponding traps to the destination address when the events occur. No unlisted traps are transmitted.

- **Removing Trap Destinations** — When you remove a destination, no SNMP traps are reported to that destination.
- **Flushing All SNMP Trap Destinations** — When you flush the SNMP trap reporting destinations, you remove all trap destination address information for the SNMP agent.

Controlling SNMP Write Requests

You can `enable` or `disable` SNMP write requests.

Remote Monitoring (RMON)

This section provides information about Remote Monitoring (RMON) and the RMON-1 Management Information Base (MIB) groups that are implemented in your system. The following topics are included:

- Overview of RMON
- RMON Benefits
- 3Com Transcend RMON Agents
- RMON-1 Groups

You can gain access to the RMON capabilities of the system through SNMP applications such as Transcend Network Control Services software, not through the serial interface or Telnet. For more information about the details of managing 3Com devices using RMON and Transcend tools, see the user documentation for the 3Com Transcend Network Control Services for Windows suite of applications.

Overview of RMON

RMON provides a way to monitor and analyze a local area network (LAN) from a remote location. The Internet Engineering Task Force (IETF) defines RMON-1 (RMON Version 1) in documents RFC 1271 and RFC 1757.

A typical RMON implementation has two components:

- **Your system** — Your system's built-in probe functionality examines all the LAN traffic on its segments, and keeps a summary of statistics (including historical data) in its local memory.
- **Management station** — Communicates with your system and collects the summarized data from it. The station can be on a different network from the system and can manage the system's probe function through either in-band or out-of-band connections.

The RMON specification consists almost entirely of the definition of the MIB. The RMON MIB contains standard MIB variables that are defined to collect comprehensive network statistics that alert you to significant network events. If the embedded RMON agent operates full time, it collects data on the correct port when the relevant network event occurs.

RMON Benefits

From a network management station, traditional network management applications poll network devices such as switches, bridges, and routers at regular intervals. The console gathers statistics, identifies trends, and highlights network events. The console polls network devices constantly to determine if the network is within its normal operating conditions.

As network size and traffic levels grow, however, the network management station can become overburdened by the amount of data it must collect. Frequent console polling also generates significant network traffic that itself can create problems for the network.

The RMON implementation in your system offers solutions to both of these problems:

- The system examines the network without affecting the characteristics and performance of the network.
- The system can report by exception rather than by reporting constant or frequent information. That is, the system informs the network management station directly if the network enters an abnormal state. The station can then use more information gathered by the system, such as historical information, to diagnose the abnormal condition.

3Com Transcend RMON Agents

RMON requires one probe per LAN segment. Because a segment is a portion of the LAN that is separated by a bridge or router, the cost of implementing many probes in a large network can be high.

To solve this problem, 3Com has built an inexpensive RMON probe into the Transcend SmartAgent software in each system. With this probe you deploy RMON widely around the network at a cost of no more than the cost of traditional network monitors.

Placing probe functionality inside the system has these advantages:

- You can integrate RMON with normal device management.
- The system can manage conditions proactively.

The system associates statistics with individual ports and then takes action based on these statistics. For example, the system can generate a log event and send an RMON trap if errors on a port exceed a threshold set by the user.

Figure 37 shows an example of the RMON implementation.

Figure 37 Embedded RMON Implemented on the System

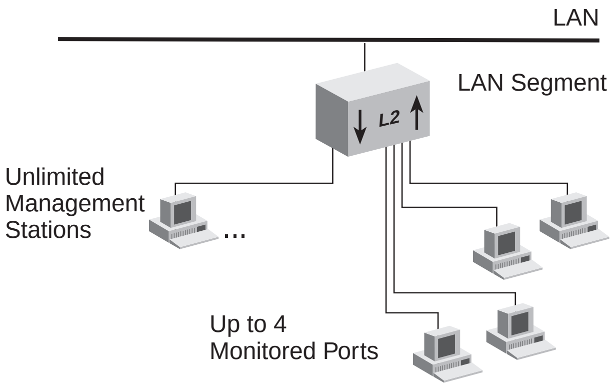

Important Considerations

- To manage RMON, you must assign an IP address to the system. See Chapter 11 for information about managing IP interfaces.
- The system will keep RMON Statistics (group 1) data on all ports.
- The system can be configured to keep all other RMON group data — History (group 2), Alarm (group 3), and Event (group 9) — on up to four ports across the entire system.
- There is no limit to the number of network management stations monitoring the data.

RMON-1 Groups

Your system software offers full-time embedded RMON support using SNMP for four RMON-1 groups. (RMON-1 defines 10 groups.) Table 25 briefly describes these groups.

Table 25 RMON-1 Groups Supported in the System

RMON-1 Group	Group Number	Purpose
Statistics	1	Maintains utilization and error statistics for the segment being monitored
History	2	Gathers and stores periodic statistical samples from the statistics group
Alarm	3	Allows you to define thresholds for any MIB variable and trigger alarms
Event	9	Allows you to define actions (generate traps, log alarms, or both) based on alarms

Statistics Group

The statistics group records frame statistics for Ethernet interfaces. The information available per interface segment includes:

- Number of received octets
- Number of received packets
- Number of received broadcast packets
- Number of received multicast packets
- Number of received packets with cyclic redundancy check (CRC) or alignment errors
- Number of received packets that are undersized but otherwise well-formed
- Number of received packets that are oversized but otherwise well-formed
- Number of received undersized packets with either a CRC or an alignment error
- Number of detected transmit collisions

Byte sizes include the 4-byte FCS but exclude the framing bits. The following Ethernet packet length counters are implemented in the RMON-1 statistics group to keep track of the frame sizes that are encountered:

- 64 Bytes
- 65 – 127 Bytes
- 128 – 511 Bytes
- 512 – 1023 Bytes
- 1024 – 1518 Bytes (1024 – 1522 Bytes when tagging is enabled)

History Group

The history group records periodic statistical samples for Ethernet interfaces and store them for later retrieval. The information available per interface for each time interval includes:

- Number of received octets
- Number of received packets
- Number of received broadcast packets
- Number of received multicast packets
- Number of received packets with CRC or alignment errors
- Number of received packets that are undersized but otherwise well-formed
- Number of received packets that are oversized but otherwise well-formed
- Number of received undersized packets with either a CRC or an alignment error
- Number of detected transmit collisions
- Estimate of the mean physical layer network utilization

Alarm Group

The system supports the following RMON alarm mechanisms:

- Counters
- Gauges
- Integers
- Timeticks

These RMON MIB objects yield alarms when the network exceeds predefined limits. The most frequently used objects are *counters*, although the other objects may be used in much the same way. The balance of this chapter illustrates RMON functions using counters.

Counters hold and update the number of times an event occurs on a port, or switch. *Alarms* monitor the counters and report when counters exceed their set threshold.

Counters are useful when you compare their values at specific time intervals to determine rates of change. The time intervals can be short or long, depending on what you measure.

Occasionally, counters can produce misleading results. Because counters are finite, they are useful for comparing rates. When counters reach a predetermined limit, they *roll over* (that is, return to 0). A single low counter value may accurately represent a condition on the network. On the other hand, the same value may simply indicate a rollover.

When you disable a port, the application may not update some of its associated statistics counters.

An alarm calculates the difference in counter values over a set time interval and remembers the high and low values. When the value of a counter exceeds a preset threshold, the alarm reports this occurrence.

Using Transcend Network Control Services or any other SNMP network management application, you can assign alarms to monitor any counter, gauge, timetick, or integer. See the documentation for your management application for details about setting up alarms.

Setting Alarm Thresholds

Thresholds determine when an alarm reports that a counter has exceeded a certain value. You can set alarm thresholds manually through the network, choosing any value for them that is appropriate for your application. The network management software monitors the counters and thresholds continually during normal operations to provide data for later calibration.

Figure 38 shows a counter with thresholds set manually.

Figure 38 Manually Set Thresholds

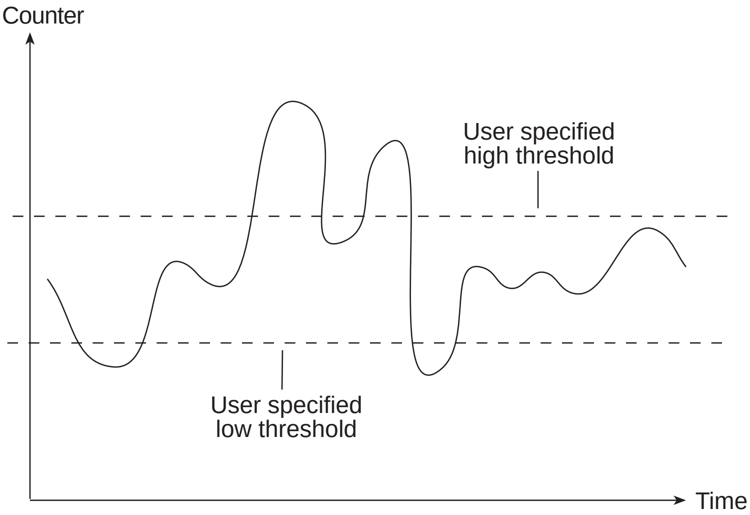

You can associate an alarm with the high threshold, the low threshold, or both. The actions that occur because of an alarm depend on the network management application.

RMON Hysteresis Mechanism

The RMON hysteresis mechanism prevents small fluctuations in counter values from causing alarms. Alarms occur only when either:

- The counter value exceeds the high threshold after previously falling below the low threshold. (An alarm does not occur if the value has not fallen below the low threshold before rising above the high threshold.)
- The counter value falls below the low threshold after previously exceeding the high threshold. (An alarm does not occur if the value has not risen above the high threshold before falling below the low threshold.)

For example, in Figure 38, an alarm occurs the first time that the counter exceeds the high threshold, but not the second time. At the first instance, the counter is rising from below the low threshold. In the second instance, the counter is not rising from below the low threshold,

Event Group

The event group logs alarms or traps network event descriptions. Although alarm group thresholds trigger most events, other RMON groups may define event conditions.

Management Information Base (MIB)

This section provides information on the Management Information Base (MIB). A MIB is a structured set of data that describes the way that the network is functioning. The management software, known as the *agent*, gains access to this set of data and extracts the information it needs. The agent can also store data in the MIB. The following topics are covered:

- MIB Files
- Compiler Support
- MIB Objects
- MIB Tree
- MIB-II
- RMON-1 MIB
- 3Com Enterprise MIBs

MIB Files

The organization of a MIB allows a Simple Network Management Protocol (SNMP) network management package, such as the Transcend Network Control Services application suite, to manage a network device without having a specific description of that device. 3Com ships the following MIB files as ASN.1 files.

- **BRIDGE-MIB.mib** — Bridge MIB, RFC 1493

  Unsupported groups and tables in this MIB:

  - dot1dSr group
  - dot1dStatic group
- **ETHERNET-MIB.mib** — Ethernet MIB, RFC 1398
- **IANAifType-MIB-V1SMI.mib** — Internet Assigned Numbers Authority MIB, SMI Version 1, RFC 1573

- **IF-MIB-V1SMI.mib** — Interface MIB, RFC 1573

  Unsupported tables in this MIB:

  - ifTestTable
  - ifRcvAddressTable
  - ifHC 64-bit counters
- **MIB2-MIB.mib** — MIB-II MIB, RFC 1213

  Unsupported groups and tables in this MIB:

  - egp group
- **RMON-MIB.mib** — RMON MIB, RFC 1757

  Supported groups in this MIB:

  - statistics
  - history
  - alarm
  - event
- **SNMPv2-MIB.mib** — Used by other MIBs, RFC 1907
- **3Com Enterprise MIBs** — See "3Com Enterprise MIBs" later in this chapter.

Compiler Support

ASN.1 MIB files are provided for each of these MIB compilers:

- SunNet Manager (version 2.0)
- SMICng (version 2.2.06)

MIB Objects

The data in the MIB consists of objects that represent features of the equipment that an agent can control and manage. Examples of objects in the MIB include a port that you can enable or disable and a counter that you can read.

A counter is a common type of MIB object used by RMON. A counter object may record the number of frames transmitted onto the network. The MIB may contain an entry for the counter object similar to the one in Figure 39.

Figure 39 Example of an RMON MIB Counter Object

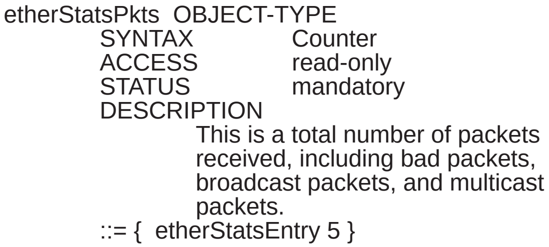

```
etherStatsPkts  OBJECT-TYPE
          SYNTAX              Counter
          ACCESS              read-only
          STATUS              mandatory
          DESCRIPTION
                    This is a total number of packets
                    received, including bad packets,
                    broadcast packets, and multicast
                    packets.
          ::= {  etherStatsEntry 5 }
```

The counter object information includes these items:

- The name of the counter. In Figure 39, the counter is called *etherStatsPkts* (Ethernet, Statistics, Packets).
- Access level. In Figure 39, access is read-only.
- The number of the counter's column in the table. In Figure 39, the counter is in column 5 of the *etherStatsEntry* table.

The name of the table where the counter resides is *3CometherStatTable,* although this name does not appear in the display.

To manage a network, you do not need to know the contents of every MIB object. Most network management applications, including Transcend Network Control Services, make the MIB transparent. However, by knowing how different management features are derived from the MIB you can better understand how to use the information they provide.

MIBs include MIB-II, other standard MIBs (such as the RMON MIB), and vendors' private MIBs (such as enterprise MIBs from 3Com Corporation). These MIBs and their objects are part of the MIB tree.

MIB Tree

The MIB tree is a structure that groups MIB objects in a hierarchy and uses an abstract syntax notation (ASN.1) to define manageable objects. Each item on the tree is assigned a number (shown in parentheses after each item), which creates the path to objects in the MIB. See Figure 40. This path of numbers is called the object identifier (OID). Each object is uniquely and unambiguously identified by the path of numeric values.

When the system software performs an SNMP Get operation, the management application sends the OID to the agent, which in turn determines if the OID is supported. If the OID is supported, the agent returns information about the object.

For example, to retrieve an object from the RMON MIB, the software uses this OID:

```
1.3.6.1.2.1.16
```

which indicates this path:

```
iso(1).indent-org(3).dod(6).internet(1).mgmt(2).mib(1).RMON(
16)
```

Figure 40 MIB Tree Showing Key MIBs

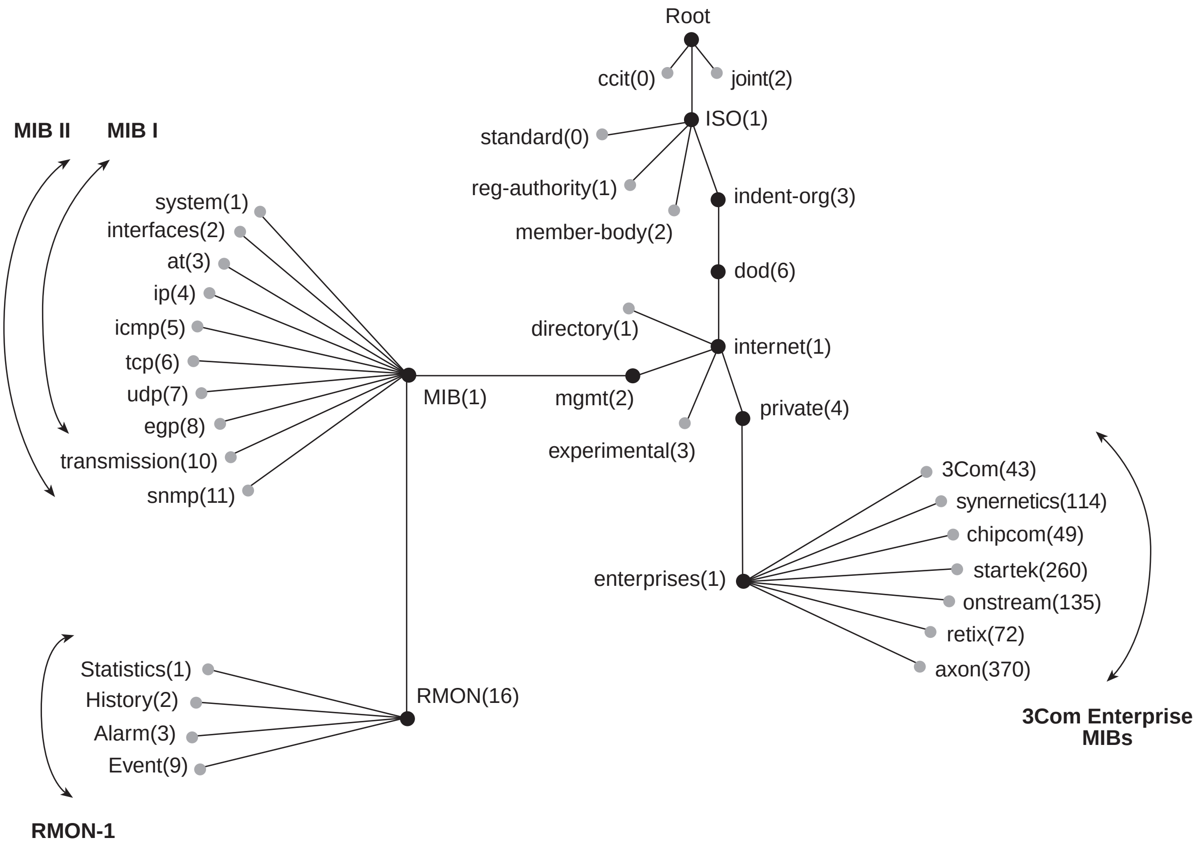

MIB-II

MIB-II defines various groups of manageable objects that contain device statistics as well as information about the device, device status, and the number and status of interfaces.

The MIB-II data is collected from network devices using SNMP. As collected, this data is in its raw form. To be useful, data must be interpreted by a management application, such as Status Watch.

MIB-II, the only MIB that has reached Internet Engineering Task Force (IETF) standard status, is the one MIB that all SNMP agents are likely to support.

Table 26 lists the MIB-II object groups. The number following each group indicates the group's branch in the MIB subtree.

MIB-I supports groups 1 through 8; MIB-II supports groups 1 through 8, plus two additional groups.

Table 26 MIB-II Group Descriptions

MIB-II Group	Purpose
system(1)	Operates on the managed node
interfaces(2)	Operates on the network interface (for example, a port or MAC) that attaches the device to the network
at(3)	Were used for address translation in MIB-I but are no longer needed in MIB-II
ip(4)	Operates on the Internet Protocol (IP)
icmp(5)	Operates on the Internet Control Message Protocol (ICMP)
tcp(6)	Operates on the Transmission Control Protocol (TCP)
udp(7)	Operates on the User Datagram Protocol (UDP)
egp(8)	Operates on the Exterior Gateway Protocol (EGP)
transmission(10)	Applies to media-specific information (implemented in MIB-II only)
snmp(11)	Operates on SNMP (implemented in MIB-II only)

RMON-1 MIB

RMON-1 is a MIB that enables the collection of data about the network itself, rather than about devices on the network.

The IETF definition for the RMON-1 MIB specifies several groups of information. These groups are described in Table 27.

Table 27 RMON-1 Group Descriptions

RMON-1 Group	Description
Statistics(1)	Total LAN statistics
History(2)	Time-based statistics for trend analysis
Alarm(3)	Notices that are triggered when statistics reach predefined thresholds
Event(9)	Reporting mechanisms for alarms

3Com Enterprise MIBs

3Com enterprise MIBs allow you to manage unique and advanced functionality of 3Com devices. These MIBs are shipped with your system. Figure 40 shows some of the 3Com enterprise MIB names and numbers. The following MIBs are included in 3Com (43):

- **3cigmpSnoop.mib** — 3Com IGMP Snooping MIB (43.10.37.1)
- **3com0304.mib** — 3Com Resilient Links MIB (43.10.15)
- **3cPoll.mib** — 3Com Remote Polling MIB (43.29.4.22)
- **3cProd.mib** — 3Com Transcend Product Management MIB
- **3cSys.mib** — 3Com System MIB (43.29.4)

  Unsupported groups in this MIB:

  - a3ComSysSlot
  - a3ComSysControlPanel
  - a3ComSysSnmp
- **3cSysBridge.mib** — 3Com Bridging MIB (43.29.4.10)
- **3cSysFt.mib** — 3Com File Transfer MIB (43.29.4.14)
- **3cTrunk.mib** — 3Com Port Trunking MIB (43.10.1.15.1)
- **3cVlan.mib** — 3Com VLAN MIB (43.10.1.14.1)
- **3cWeb.mib** — 3Com Web Management MIB (43.29.4.24)

MIB names and numbers are usually retained when organizations restructure their businesses; therefore, some of the 3Com enterprise MIB names do not contain the word “3Com.”

A

TECHNICAL SUPPORT

3Com provides easy access to technical support information through a variety of services. This appendix describes these services.

Information contained in this appendix is correct at time of publication. For the most recent information, 3Com recommends that you access the 3Com Corporation World Wide Web site.

Online Technical Services

3Com offers worldwide product support 24 hours a day, 7 days a week, through the following online systems:

- World Wide Web site
- 3Com Knowledgebase Web Services
- 3Com FTP site
- 3Com Bulletin Board Service (3Com BBS)
- 3Com Facts^{SM} Automated Fax Service

World Wide Web Site

To access the latest networking information on the 3Com Corporation World Wide Web site, enter this URL into your Internet browser:

`http://www.3com.com/`

This service provides access to online support information such as technical documentation and software, as well as support options that range from technical education to maintenance and professional services.

3Com Knowledgebase Web Services

This interactive tool contains technical product information compiled by 3Com expert technical engineers around the globe. Located on the World Wide Web at **`http://knowledgebase.3com.com`**, this service gives all 3Com customers and partners complementary, round-the-clock access to technical information on most 3Com products.

3Com FTP Site

Download drivers, patches, software, and MIBs across the Internet from the 3Com public FTP site. This service is available 24 hours a day, 7 days a week.

To connect to the 3Com FTP site, enter the following information into your FTP client:

- Hostname: **`ftp.3com.com`**
- Username: **`anonymous`**
- Password: **`<your Internet e-mail address>`**

You do not need a user name and password with Web browser software such as Netscape Navigator and Internet Explorer.

3Com Bulletin Board Service

The 3Com BBS contains patches, software, and drivers for 3Com products. This service is available through analog modem or digital modem (ISDN) 24 hours a day, 7 days a week.

Access by Analog Modem

To reach the service by modem, set your modem to 8 data bits, no parity, and 1 stop bit. Call the telephone number nearest you:

Country	Data Rate	Telephone Number
Australia	Up to 14,400 bps	61 2 9955 2073
Brazil	Up to 28,800 bps	55 11 5181 9666
France	Up to 14,400 bps	33 1 6986 6954
Germany	Up to 28,800 bps	4989 62732 188
Hong Kong	Up to 14,400 bps	852 2537 5601
Italy	Up to 14,400 bps	39 2 27300680
Japan	Up to 14,400 bps	81 3 5977 7977
Mexico	Up to 28,800 bps	52 5 520 7835
P.R. of China	Up to 14,400 bps	86 10 684 92351
Taiwan, R.O.C.	Up to 14,400 bps	886 2 377 5840
U.K.	Up to 28,800 bps	44 1442 438278
U.S.A.	Up to 53,333 bps	1 847 262 6000

Access by Digital Modem

ISDN users can dial in to the 3Com BBS using a digital modem for fast access up to 64 Kbps. To access the 3Com BBS using ISDN, call the following number:

1 847 262 6000

3Com Facts Automated Fax Service

The 3Com Facts automated fax service provides technical articles, diagrams, and troubleshooting instructions on 3Com products 24 hours a day, 7 days a week.

Call 3Com Facts using your Touch-Tone telephone:

1 408 727 7021

Support from Your Network Supplier

If you require additional assistance, contact your network supplier. Many suppliers are authorized 3Com service partners who are qualified to provide a variety of services, including network planning, installation, hardware maintenance, application training, and support services.

When you contact your network supplier for assistance, have the following information ready:

- Product model name, part number, and serial number
- A list of system hardware and software, including revision levels
- Diagnostic error messages
- Details about recent configuration changes, if applicable

If you are unable to contact your network supplier, see the following section on how to contact 3Com.

Support from 3Com

If you are unable to obtain assistance from the 3Com online technical resources or from your network supplier, 3Com offers technical telephone support services. To find out more about your support options, call the 3Com technical telephone support phone number at the location nearest you.

When you contact 3Com for assistance, have the following information ready:

- Product model name, part number, and serial number
- A list of system hardware and software, including revision levels
- Diagnostic error messages
- Details about recent configuration changes, if applicable

Here is a list of worldwide technical telephone support numbers:

Country	Telephone Number	Country	Telephone Number
Asia, Pacific Rim			
Australia	1 800 678 515	P.R. of China	10800 61 00137 or
Hong Kong	800 933 486		021 6350 1590
India	+61 2 9937 5085	Singapore	800 6161 463
Indonesia	001 800 61 009	S. Korea	
Japan	0031 61 6439	From anywhere in S. Korea:	00798 611 2230
Malaysia	1800 801 777	From Seoul:	(0)2 3455 6455
New Zealand	0800 446 398	Taiwan, R.O.C.	0080 611 261
Pakistan	+61 2 9937 5085	Thailand	001 800 611 2000
Philippines	1235 61 266 2602		
Europe			
From anywhere in Europe, call:	+31 (0)30 6029900 phone +31 (0)30 6029999 fax		
Europe, South Africa, and Middle East			
From the following countries, you may use the toll-free numbers:			
Austria	0800 297468	Netherlands	0800 0227788
Belgium	0800 71429	Norway	800 11376
Denmark	800 17309	Poland	00800 3111206
Finland	0800 113153	Portugal	0800 831416
France	0800 917959	South Africa	0800 995014
Germany	0800 1821502	Spain	900 983125
Hungary	00800 12813	Sweden	020 795482
Ireland	1800 553117	Switzerland	0800 55 3072
Israel	1800 9453794	U.K.	0800 966197
Italy	1678 79489		
Latin America			
Argentina	AT&T +800 666 5065	Mexico	01 800 CARE (01 800 2273)
Brazil	0800 13 3266	Peru	AT&T +800 666 5065
Chile	1230 020 0645	Puerto Rico	800 666 5065
Colombia	98012 2127	Venezuela	AT&T +800 666 5065
North America	1 800 NET 3Com (1 800 638 3266) Enterprise Customers: 1 800 876-3266		

Returning Products for Repair

Before you send a product directly to 3Com for repair, you must first obtain an authorization number. Products sent to 3Com without authorization numbers will be returned to the sender unopened, at the sender's expense.

To obtain an authorization number, call or fax:

Country	Telephone Number	Fax Number
Asia, Pacific Rim	+ 65 543 6500	+ 65 543 6348
Europe, South Africa, and Middle East	+ 31 30 6029900	+ 31 30 6029999
Latin America	1 408 326 2927	1 408 326 3355
From the following countries, you may call the toll-free numbers; select option 2 and then option 2:		
Austria	0800 297468	
Belgium	0800 71429	
Denmark	800 17309	
Finland	0800 113153	
France	0800 917959	
Germany	0800 1821502	
Hungary	00800 12813	
Ireland	1800553117	
Israel	1800 9453794	
Italy	1678 79489	
Netherlands	0800 0227788	
Norway	800 11376	
Poland	00800 3111206	
Portugal	0800 831416	
South Africa	0800 995014	
Spain	900 983125	
Sweden	020 795482	
Switzerland	0800 55 3072	
U.K.	0800 966197	
U.S.A. and Canada	1 800 NET 3Com (1 800 638 3266)	1 408 326 7120 (not toll-free)
	Enterprise Customers: 1 800 876 3266	

INDEX