CoreBuilder® 3500 Implementation Guide

Release 3.0

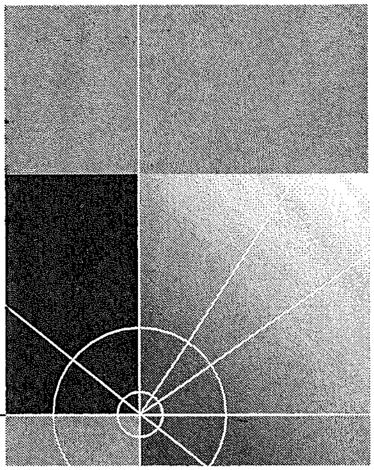

http://www.3com.com/

Part No. 10013506
Published November 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 User Guide.

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.

Postscript is a registered trademark of Adobe Systems, Inc. AppleTalk is a registered trademark of Apple Computer, Incorporated. Banyan and VINES are registered trademarks of Banyan Worldwide. DEC, DECnet, and PATHWORKS are registered trademarks of Compaq Computer Corporation. OpenView is a registered trademark of Hewlett-Packard Company. 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. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd. Xerox and XNS are trademarks of Xerox Corporation.

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

Guide written by 3Com MSD Technical Publications, Marlborough, MA, USA.
We welcome your comments on this guide. Send e-mail to: sdtechpubs_comments@ne.3com.com

CONTENTS

ABOUT THIS GUIDE

3 SYSTEM PARAMETERS

4 PHYSICAL PORT NUMBERING

5 ETHERNET

6 FIBER DISTRIBUTED DATA INTERFACE (FDDI)

7 BRIDGE-WIDE AND BRIDGE PORT PARAMETERS

10 PACKET FILTERING

11 INTERNET PROTOCOL (IP)

14 Open Shortest Path First (OSPF)

15 IPX ROUTING

18 DEVICE MONITORING

A Technical Support

Index

About This Guide

This guide describes information that you need to use features of the CoreBuilder® 3500 system after you install it and attach the system to your network. Before you use this guide:

- Verify that your system is installed and set up using the *CoreBuilder 3500 Getting Started Guide*.
- Become familiar with the *Command Reference Guide*, which documents the commands that are used to configure and manage your system.
- 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 CoreBuilder 3500 system. It assumes a working knowledge of local area network (LAN) operations and familiarity with communications protocols that are used on interconnected LANs.

If the information in the release notes differs from the information in this guide, follow the instructions in the release notes.

The most current versions of guides and release notes are available in Adobe Acrobat Reader Portable Document Format (PDF) or HTML from the 3Com World Wide Web site:

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

Conventions

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

Table 1 Notice Icons

Icon	Notice 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 application, 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 the screen.
`Syntax`	The word "syntax" means that you evaluate the syntax provided and then supply the appropriate values for the placeholders that appear in angle brackets. Example: To set the system date and time, use the following syntax: `CCYY-MM-DDThh:mm:ss`
`Commands`	The word "command" means that you must enter the command exactly as shown and then press Return or Enter. Commands appear in bold. Example: To remove an IP interface, enter the following command: **`ip interface remove`**
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."
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.

Table 2 Text Conventions (continued)

Convention	Description
Words in *italics*	Italics are used to: - Emphasize a point - Denote a new term at the place where it is defined in the text - Identify menu names, menu commands, and software button names. Examples: From the *Help* menu, select *Contents*. Click *OK*.

CoreBuilder 3500 Documentation

The following documents comprise the CoreBuilder 3500 documentation set. Documents are available in one of two forms:

- Paper Documents

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

- Software and Documents on CD-ROM

  The System Software and Documentation CD contains online versions of the paper documents, this *Implementation Guide*, and the *Command Reference Guide*, as well as the CoreBuilder 3500 system software.

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

Paper Documents

These documents are shipped with your system:

- *CoreBuilder 3500 Unpacking Instructions*

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

- *CoreBuilder 3500 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.

- *CoreBuilder 3500 Quick Installation Guide*

  Quick reminders and information for system installation. For greater detail on installation procedures, see the *CoreBuilder 3500 Getting Started Guide*.

- *CoreBuilder 3500 Getting Started Guide*

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

- *CoreBuilder 3500 Command Quick Reference*

  All of the Administration Console commands for the system.

- *Web Management User Guide for the CoreBuilder 3500*

  Overview, installation, and troubleshooting information for the suite of applications that help you manage your system over the Internet.

In addition, each module and field-replaceable component contains a guide:

- *CoreBuilder 3500 System Processor Removal and Replacement Guide*

  Provides overview information and removal and replacement instructions for the CoreBuilder system processor.

- *Module Installation Guides*

  An overview, LED status information, and installation instructions for each module.

- *GBIC Transceiver Installation Guide*

  Installation instructions for the Gigabit Ethernet Interface Converter transceiver.

- *CoreBuilder 3500 Power Supply Assembly Removal and Replacement Guide*

  Overview information and removal and replacement instructions for the CoreBuilder power supplies.

- *CoreBuilder 3500 Fan Tray Removal and Replacement Guide*

  Overview information and removal and replacement instructions for the fan tray.

- *PCMCIA Flash Card User Guide*

  Information on using the PCMCIA card to save and restore system configuration settings.

- *Blank Faceplate Installation Guide*

  Instructions for covering empty slots with the blank faceplate.

Software and Documents on CD-ROM

The compact disc that comes with your system contains:

- System software
- Online versions of the paper guides that are shipped with your system, modules, and field-replaceable components
- *CoreBuilder 3500 Implementation Guide* (this guide)
- Multiplatform *Command Reference Guide*

  Information about the commands used to configure the system. This guide documents commands for the CoreBuilder 3500 as well as other 3Com systems.
- Help system for the Web Management suite of applications

  Online Help system for the CoreBuilder 3500 Web Management software. See the *Web Management User Guide* 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 commenting:

- Document title
- Document part number (found on the inside title page of this guide)
- Page number

Example:

CoreBuilder 3500 Implementation Guide

Part Number 10013506

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 the configuration procedure for the first time that you install a CoreBuilder® 3500 Layer 3 High-Function Switch.

To upgrade the software on an existing switch, see the Software Installation and Release Notes *for configuration information.*

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.

3Com recommends that you use the following procedures the first time that you set up your system and every time that you modify its configuration.

Procedure Summary

These steps are described in detail in the next section:

1. Establish management access.
2. Choose a subsequent management access method.
3. Choose a subsequent management interface.
4. Configure parameters related to the network infrastructure. These include system, bridge-wide and bridge-port, Ethernet, FDDI, and trunking parameters.
5. Define all VLANs.
6. Define routing protocol interfaces and set related parameters.
7. Configure more advanced traffic control features, such as packet filters and Quality of Service (QoS).
8. Monitor the system and analyze network activity.

These steps are described in detail in the next section.

Configuration Procedure

Follow the steps that apply to your network needs and ignore the steps that do not apply.

1 **Establish management access.**

To perform configuration or management tasks, you must initially:

a Connect to the system through its terminal serial port or modem serial port.

For information about the required settings for the serial ports, see Chapter 2 in this guide.

b Use the Administration Console as the management interface.

The Administration Console is a menu-driven command line interface that is embedded in the system software. For specific menu and command information, see the *Command Reference Guide*.

2 **Choose a subsequent management access method.**

You can continue to access your system through a local serial connection, or you can use one of two other local access methods — any in-band port on a media module or the out-of-band 10BASE-T port on the system processor module. To manage the system through either access method, you must first configure an IP address:

- **To configure an IP address for an out-of-band port** — Using the serial port connection from step 1, configure an IP address through the `management ip interface` menu. For more information, see Chapter 2.
- **To configure an IP address for an in-band port** — Using the serial port connection from step 1, configure an IP address by defining an IP VLAN (through the `bridge vlan` menu; see Chapter 9) and an associated IP interface (through the `ip interface` menu; see Chapter 11).

3 Choose a subsequent management interface.

After you configure an IP address, you have additional management interface options:

- **Administration Console** — You can now access this interface from a remote Telnet connection.
- **Web Management software** — From your Web browser, you can access a suite of HTML-based applications that are embedded in the software. For more information, see the *Web Management User Guide*.
- **SNMP-based applications** — One example is 3Com Transcend® Network Control Services software. To manage the system in-band from SNMP-based applications, set the SNMP parameters through the `snmp` menu. For more information, see Chapter 2 and Chapter 18 in this guide, as well as application-specific documentation.

4 Configure parameters related to the network infrastructure.

One or more of the following topics may apply to your system, depending on your network requirements:

- **System parameters** — To choose the file transfer protocol, administer nonvolatile data (nvData), perform system software updates, and display your system configuration, see Chapter 3.
- **Physical port numbering** — To learn the port numbering rules and understand the effects of adding or removing modules, see Chapter 4.
- **Ethernet** — To label Ethernet ports, set the port mode, enable flow control, and control autonegotiation and other settings, see Chapter 5.
- **FDDI** — To configure stations, paths, MACs, and ports, see Chapter 6.
- **Bridge-wide and bridge port parameters** — To set parameters for Spanning Tree Protocol (STP), GARP VLAN Registration Protocol (GVRP), IPX SNAP translation, and IP fragmentation, see Chapter 7.
- **Trunks** — To increase the bandwidth and resiliency between two points, you can aggregate many individual links into a single logical link called a trunk. You must configure trunks before you define VLANs. For more information, see Chapter 8.

5 **Define all VLANs.**

To create logical workgroups, which are generally equivalent to Layer 2 broadcast domains or Layer 3 networks, you can define port-based, protocol-based, and network-based VLANs, and set related modes in the system. You must define VLANs after you define trunks and before you define routing interfaces. For more information, see Chapter 9.

6 **Configure routing interfaces and set related parameters.**

You can use the following protocols to configure routing interfaces and set related parameters:

- **IP** — See Chapter 11.
- **VRRP** — See Chapter 12.
- **IP multicast** — See Chapter 13.
- **OSPF** — See Chapter 14.
- **IPX** — See Chapter 15.
- **AppleTalk** — See Chapter 16.

7 **Configure more advanced traffic control features:**

- **Packet filters** — To improve LAN performance, shape traffic flows, or implement security controls with standard, custom, predefined, and port group packet filters, see Chapter 10.
- **Quality of Service (QoS) and the Resource Reservation Protocol (RSVP)** — To classify, control, and prioritize traffic where available bandwidth is low and your network is carrying time-sensitive or business-critical information, use the QoS and RSVP features. For more information, see Chapter 17.

8 **Monitor the system and analyze network activity.**

You can use the system's device monitoring features such as event logging, baselining, roving analysis, and RMON to record and analyze your network periodically and identify potential network problems before they become serious problems. To test and validate paths in your network, use tools like ping and traceRoute. SNMP and MIBs provide ways to collect performance data on your network. For more information on these features, see Chapter 18.

2

Management Access

This chapter explains the different methods used to configure management access to the system. It describes the different types of applications and the underlying communication protocols that are used to deliver data between your end-station device and the system. It also contains information about connecting to the system directly through one of two serial connections, or through an Ethernet port to an IP (Internet Protocol) interface to run network management applications.

This chapter covers the following topics:

- Management Access Overview
- Key Concepts
- Key Guidelines for Implementation
- Administration Console Access
- Web Management Access
- SNMP Access

Management Access Overview

The system provides you with the flexibility to access and manage your system using several different 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 Web Management are embedded parts of the software and are available for immediate use on your system.

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 a terminal or modem serial port, or through an Ethernet port using an Internet Protocol (IP) interface.

Figure 1 shows a sample output of menu options that can be viewed from the various devices.

Figure 1 Viewing the Administration Console

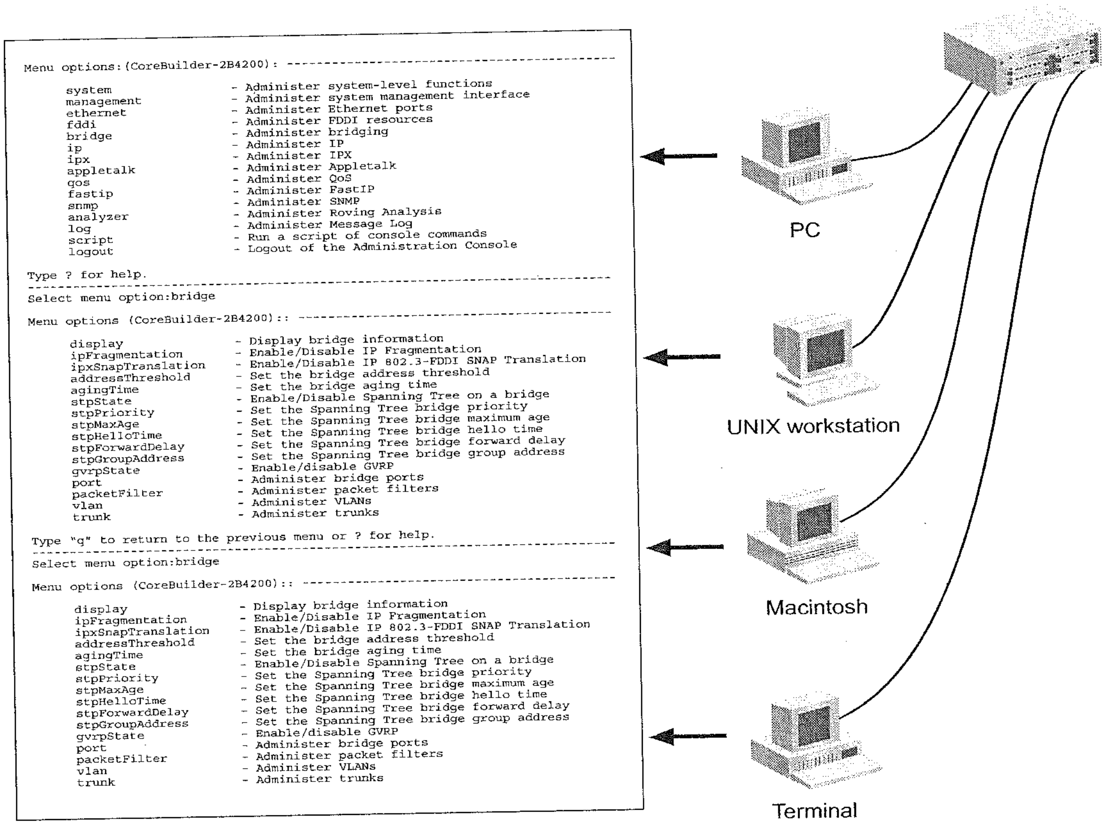

```
Menu options:(CoreBuilder-2B4200): -------------------------------------------

      system                   - Administer system-level functions
      management               - Administer system management interface
      ethernet                 - Administer Ethernet ports
      fddi                     - Administer FDDI resources
      bridge                   - Administer bridging
      ip                       - Administer IP
      ipx                      - Administer IPX
      appletalk                - Administer Appletalk
      qos                      - Administer QoS
      fastip                   - Administer FastIP
      snmp                     - Administer SNMP
      analyzer                 - Administer Roving Analysis
      log                      - Administer Message Log
      script                   - Run a script of console commands
      logout                   - Logout of the Administration Console

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

Menu options (CoreBuilder-2B4200):: ------------------------------------------

      display                  - Display bridge information
      ipFragmentation          - Enable/Disable IP Fragmentation
      ipxSnapTranslation       - Enable/Disable IP 802.3-FDDI SNAP Translation
      addressThreshold         - Set the bridge address threshold
      agingTime                - Set the bridge aging time
      stpState                 - Enable/Disable Spanning Tree on a bridge
      stpPriority              - Set the Spanning Tree bridge priority
      stpMaxAge                - Set the Spanning Tree bridge maximum age
      stpHelloTime             - Set the Spanning Tree bridge hello time
      stpForwardDelay          - Set the Spanning Tree bridge forward delay
      stpGroupAddress          - Set the Spanning Tree bridge group address
      gvrpState                - Enable/disable GVRP
      port                     - Administer bridge ports
      packetFilter             - Administer packet filters
      vlan                     - Administer VLANs
      trunk                    - Administer trunks

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

Menu options (CoreBuilder-2B4200):: ------------------------------------------

      display                  - Display bridge information
      ipFragmentation          - Enable/Disable IP Fragmentation
      ipxSnapTranslation       - Enable/Disable IP 802.3-FDDI SNAP Translation
      addressThreshold         - Set the bridge address threshold
      agingTime                - Set the bridge aging time
      stpState                 - Enable/Disable Spanning Tree on a bridge
      stpPriority              - Set the Spanning Tree bridge priority
      stpMaxAge                - Set the Spanning Tree bridge maximum age
      stpHelloTime             - Set the Spanning Tree bridge hello time
      stpForwardDelay          - Set the Spanning Tree bridge forward delay
      stpGroupAddress          - Set the Spanning Tree bridge group address
      gvrpState                - Enable/disable GVRP
      port                     - Administer bridge ports
      packetFilter             - Administer packet filters
      vlan                     - Administer VLANs
      trunk                    - Administer trunks
```

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 *Software and Documentation CD-ROM* or 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
  - **Filter Builder** — To create and test filters for packets on your switch
  - **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 Device View screen.

Figure 2 Sample Transcend Network Control Services Device View

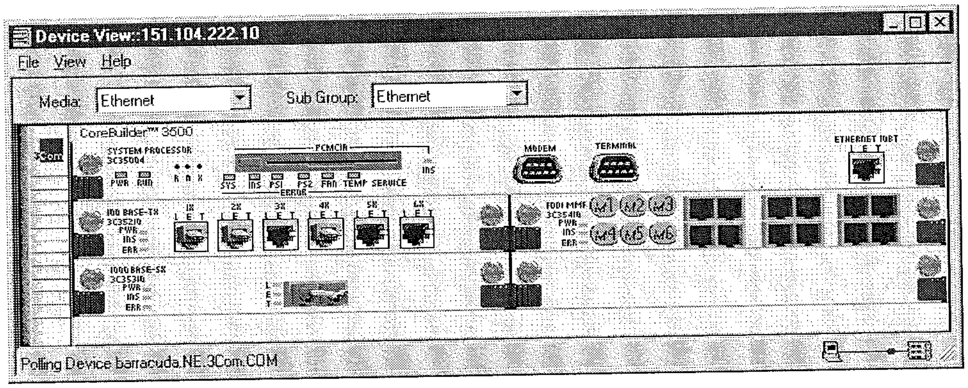

Device View::151.104.222.10

Media: Ethernet

Sub Group: Ethernet

Polling Device barracuda.NE.3Com.COM

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 introduces the concepts of in-band and out-of-band management using IP.

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 for the CoreBuilder 3500

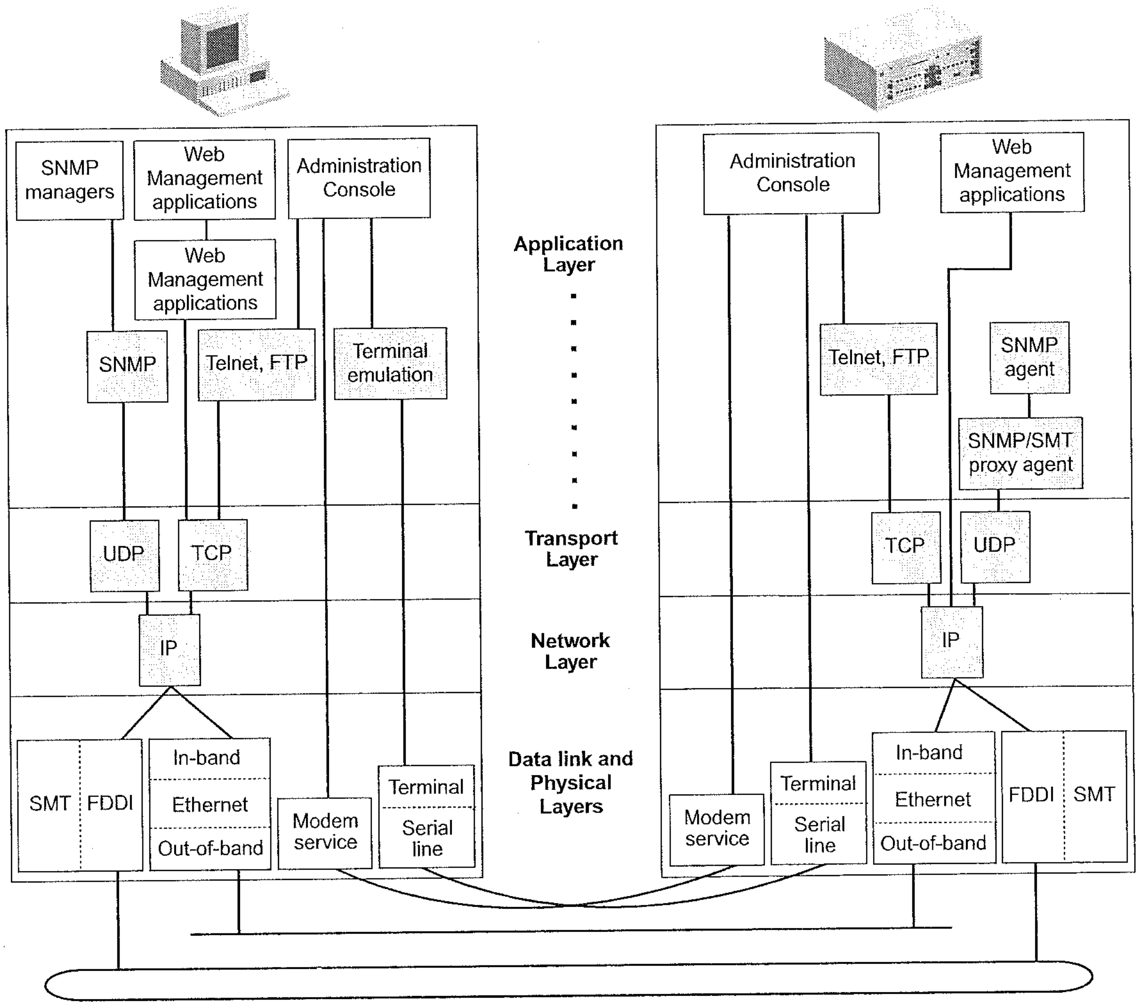

Protocols

The system supports the following protocols:

- Virtual terminal protocols, such as Telnet
- Simple Network Management Protocol (SNMP)
- FDDI Station Management (SMT) protocol

Virtual Terminal Protocols

A *virtual terminal protocol* is a software program, such as Telnet, that allow 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 serial port.

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

Figure 4 Administration Console Access

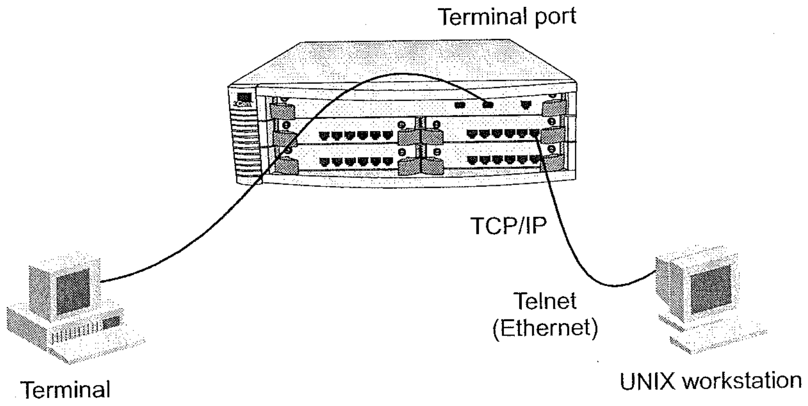

Simple Network Management Protocol

Simple Network Management Protocol (SNMP) is the standard management protocol for multi-vendor 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

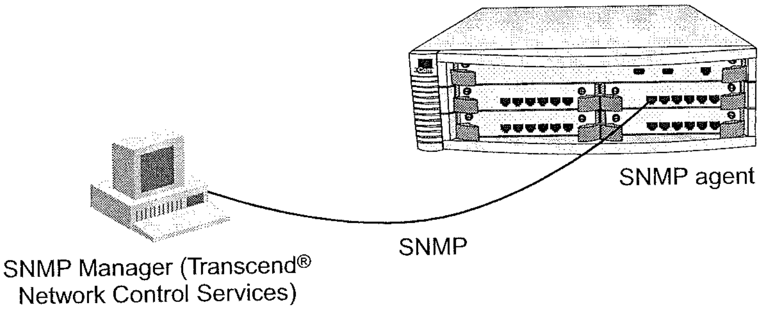

See Chapter 18 for additional information about SNMP.

IP Management Concepts

Both in-band and out-of-band management 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 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

There are several ways you can access your management application on the system; locally through a terminal connection, or remotely using a modem or an IP connection.

Table 3 describes these different methods.

Table 3 Management Access Methods

Access Method	Access Description	Interface
Terminal	Connect directly to the Administration Console and stay attached during system reboots.	Terminal serial port (see "Terminal Port Access").
Modem	Access the Administration Console by dialing in from remote sites.	Modem serial port (see "Modem Port Access").
IP	Access the Administration Console with up to four Telnet sessions. Use Web Management or an external SNMP management application to communicate with the CoreBuilder SNMP agent.	In-band or out-of-band Ethernet port assigned to an IP interface (see "In-Band Management" and "Out-of-Band Management").

Setting Up the Terminal Port

Use the Administration Console to set the baud rate 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 rate to something other than 9600, the new setting becomes the new default, even after you issue a `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 for the change to take 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 by issuing the `system serialPort connectModem` *command after you change the modem speed and before dialing in. This sequence allows the modem to synchronize its baud rate with the system.*

See the *CoreBuilder 3500 Getting Started Guide* for terminal port and modem port pin-outs. For additional information about modem port settings, see the *Command Reference Guide*.

IP Management Interface

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

- You can use Telnet to connect up to four concurrent remote sessions to the Administration Console using a terminal program from a host computer.
- You can run Web Management to access its management applications to manage and monitor your system.
- You can 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, Web Management, or SNMP access, you must first define an IP interface. You can use either an out-of-band or in-band port for the IP interface, but do not assign the same IP address to both the out-of-band and in-band ports. Also, be sure not to assign an out-of-band port IP address that is on the same subnet as any of the in-band IP interfaces.

- **In-Band Management** — If you are managing your network in-band, you need to set up an IP routing interface and at least one VLAN. See Chapter 9 for information about defining a VLAN, and Chapter 11 for information about setting up an IP routing interface. See "In-Band Management" for additional information about in-band management.
- **Out-of-Band Management** — 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 through the `management` menu. The out-of-band Ethernet port is the 10BASE-T port on the system processor module and is not associated with a port number. See Chapter 11 for background information on IP addresses and subnet masks. See "Out-of-Band Management" for additional information about out-of-band management.

Administration Console Access

The first time that you access the Administration Console, access the system at the *administer* level and press the Return key 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 4.

Table 4 Password Access Levels

Access Level	For Users Who Need to	Allows Users 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	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 terminal serial port is often preferred because you can remain on the system and monitor it during system reboots, and 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 modem 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

Web Management applications are an embedded part of the CoreBuilder 3500 Enterprise Switch. 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 of the system.

In the installable WebManage Framework 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. You can 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 are using 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. You can 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, you can download the service pack from the Microsoft home page:

  `http://www.microsoft.com`

See the *Web Management User Guide* for additional information about Web Management.

SNMP Access

You can use an external SNMP-based application such as 3Com Transcend Network Control Services to 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 at:

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

3 System Parameters

This chapter 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
- File Transfer
- 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 software applications.

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 18 for details.
- Set and administer your system's serial port baud rates

  See Chapter 2 for 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 these system parameters.

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

- File Transfer
- 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.
- Provides an easy method of communicating remotely through the fileTransfer option.
- 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:

- **FTP** — File Transfer Protocol. You can send files from one system to another with this protocol.
- **TFTP** — Trivial File Transfer Protocol. Designed to function over the User Datagram Protocol (UDP), this protocol reads and writes files to and from a remote server. It is smaller and easier to operate than FTP, but it lacks most of the FTP features.
- **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 Timing 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 subnetworks 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 on the Administration Console's top-level menu
- The System folder of Web Management software

To set or modify system parameter values, follow these basic steps:

1 Access the menu or form that governs a system parameter.

2 Specify a value.

File Transfer

From the system menu or folder, you can select which protocol you want the system to use to transfer data between systems. Choose either File Transfer Protocol (FTP) or the Trivial File Transfer Protocol (TFTP), which is the default.

Implementing FTP

FTP meets the following file transfer objectives:

- Transfers data reliably and efficiently through an IP connection
- Provides security by ensuring that the person who attempts to use FTP has a valid username and password combination

Important Consideration

- All file transfers using FTP are sent over an IP connection. Before you use FTP, you must configure an IP address for the system. For more information on IP, see Chapter 11.

Implementing TFTP

The Trivial File Transfer Protocol (TFTP) is simpler to use than FTP but has less functionality. TFTP uses UDP as its transport protocol, with a simple stop-and-wait acknowledgment system. Because TFTP has an effective window of only one 512-octet segment, its performance cannot match that of FTP. The most common application for TFTP is bootstrapping a host over a local network.

Important Considerations

Consider the following guidelines before you select TFTP:

- TFTP does not provide access control or security, so use TFTP only when authentication and directory visibility are not required.
- Because TFTP provides no user authentication, you must give *loose* permission to files that are located on your system, that is, make files publicly readable and writable. Otherwise, the TFTP server does not grant requests for file access.
- You must create two files when you are using the save nvData option over TFTP. See "Saving nvData" in this chapter.

For more information on TFTP, see your TFTP server documentation.

Security

You can limit IP management access to your system through the Administration Console or the Web Management software as follows:

- 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` option to configure up to 5 IP addresses or 5 subnetworks, called *trusted IP clients*. If an IP address or subnet is not on the trusted IP client list, the IP address or subnet cannot be used to access the system using the Web Management software, Administration Console, or SNMP.

If you do not configure trusted IP clients on the system, 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, 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.

Important Considerations

Consider the following guidelines *before* you configure trusted IP clients on your system.

Procedures

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 System Configuration form 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, as follows:*

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

Additional 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 or to a PCMCIA flash memory card, with softwareUpdate 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:

- Before you attempt to install the system software, verify that you have extended memory installed on your system. For information on how to verify your system's memory see the *Getting Started Guide.*
- 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

All of the system's configurable parameters are saved in nonvolatile memory. When you work with nonvolatile data (nvData), you can:

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

Saving nvData

When you enter commands to save nvData, the system copies data that is stored in nonvolatile memory to a disk file location that you specify. You can use the `system nvdata save` option to save nvData from your system to a:

- File on another system remotely through FTP or TFTP.
- PCMCIA flash card indirectly.

See the PCMCIA Flash Card User Guide *for a detailed explanation about how to save nvData to a PCMCIA flash memory card.*

Important Considerations

Consider the following guidelines before you perform an `nvData save` operation:

- When you use TFTP, *before* you save data to the file, you have to create two files on the TFTP server. The screen display appears as follows:

```
Select menu option: system nvdata save
Host IP Address [158.101.100.1]: 158.101.112.34
NV Control file (full pathname): [/tftpboot/mecca]
Enter an optional file label {?}: mecca2
Control File:mecca
Data File: mecca.nvd
Saving system ...
```

- You must supply the host IP address and specify the file where you want to save the data according to requirements of your TFTP and FTP implementation.
- Some TFTP implementations require that you store the file in the same directory where the TFTP daemon (server) is running on a remote host.
- Because TFTP does not provide user authentication, give the file loose permissions to make it both readable and writable. TFTP does not grant requests for file access.

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.

Resetting nvData

To reset the system settings back to their factory default values, use the nvData reset option.

Important Considerations

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.

Viewing nvData

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

Example:

```
Select menu option: system nvdata examine
Host IP Address [158.101.100.1]: 158.101.112.34
NV Control file (full pathname): systemdata
Product ID 4, Product Type 1
System ID 102D00
Saved 1999-05-20T09:24:43 AM Version 2.
Labelled: LabSwitch
```

Simple Network Time Protocol (SNTP)

This section covers:

- SNTP Overview
- Implementing SNTP

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):

- **A 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
- **A multicast server** — Periodically sends a unsolicited message to a designated IP local broadcast address or multicast group address and expects no requests from clients.
- **An 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.

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 references for more information on these protocols:

- **RFC 959** — File Transfer Protocol Specification
- **RFC 1350** — Trivial File Transfer Protocol Specification
- **RFC 2030** — Simple Network Time Protocol, v4.0, Specification
- **RFC 1305** — Network Time Protocol, v3.0, Specification
- **RFC 868** — Time Protocol Specification

4

Physical Port Numbering

The CoreBuilder® 3500 follows a specific set of rules for assigning physical port numbers. This chapter describes the physical port numbering on the system. It covers the following information:

- Port Numbering Overview
- Key Guidelines for Implementation
- Examples of Port Numbering
- Effects of Removing a Module
- Effects of Replacing Modules

Port Numbering Overview

Before you configure your system, read this chapter to become familiar with the physical port numbering scheme on the system. Understanding the port numbering scheme enables you to:

- Manage your bridge ports, especially if you use trunking, as described in Chapter 8.
- Accurately define your virtual LANs (VLANs), as described in Chapter 9.

Numbering Rules

Your system supports up to 24 ports, numbered consecutively:

- Port 1 represents the first port associated with a module in Slot 1 (the top left slot on the system) and continues for the rest of the ports associated with Slot 1.
- Numbering continues for the ports associated with a module in Slot 2 (top right).
- Numbering continues for the ports associated with a module in Slot 3 (bottom left).
- Numbering continues for the ports associated with a module in Slot 4 (bottom right).

See Figure 6 later in this chapter for an example.

Additional rules:

- Port numbering is consecutive, regardless of module type.
- Numbering skips over an empty slot and continues with the ports associated with the next occupied slot.
- Numbering includes unused ports.

For several examples of port numbering, see "Examples of Port Numbering" later in this chapter.

Supported Module Types

The port numbering range depends on the type of modules that you have configured into your system. For example, at Release 2.0.0, the system supported the following modules:

- Up to four 10/100BASE-TX Ethernet modules, each with 6 ports that have RJ-45 connectors
- Up to four 100BASE-FX Ethernet modules, each with 6 ports that have SC connectors
- Up to four 1000BASE-SX or 1000BASE Gigabit Interface Converter (GBIC) Ethernet modules, each with 1 port (up to four Gigabit Ethernet ports per system). The 1000BASE GBIC module requires CoreBuilder 3500 system software at release 1.2.0 or higher. Each Gigabit Ethernet module uses a trunk resource, so keep track of your trunk resources (maximum of 4) when you add a Gigabit Ethernet module. See Chapter 8 for information on trunking and trunking resources.
- Up to four FDDI modules, each with 6 ports
- Any combination of these modules. For example, you can have one Gigabit Ethernet module in Slot 1 (port 1), one FDDI module in Slot 2 (ports 2–7), one 10/100BASE-TX Ethernet module in Slot 3 (ports 8–13), and one 100BASE-FX Ethernet module in Slot 4 (ports 14–19).

Key Guidelines for Implementation

To ensure that you understand the port numbering that the system reports for certain aspects of your configuration (bridging information, trunks, FDDI ports, and VLANs), observe these guidelines when you configure your system:

- Determine your physical port configuration before you attempt to configure any bridging parameters.

Trunking

- If you use *trunking* to group ports, configure your trunks *before* you attempt to configure any VLANs. Be sure that you understand how trunking associates a group of ports with a trunk. (See Chapter 8.) These associations affect the following situations:
  - When you perform an operation for which you must specify bridge ports (for example, when you define VLANs), you must use the lowest-numbered port in each trunk to represent the trunk. The operation that you perform then applies to all ports in the trunk.
  - When you view information that applies to more than one port (for example, bridging displays for trunks), the port number field identifies all ports in the trunk. A VLAN summary display lists all physical ports to indicate which physical system connectors can receive or transmit frames within a VLAN. (You must use the VLAN detail display to see trunk port groups.)

FDDI DAS pairs

- By default, FDDI ports are single-attached station (SAS) M-ports, where each port is selectable as a bridge port. If you configure FDDI ports as *dual-attach station (DAS) pairs*, you associate two FDDI ports with each DAS pair and only the lowest-numbered port in the pair is selectable as a bridge port. Configure the appropriate number of DAS pairs *before* you configure any VLANs. Be sure that you understand how a DAS configuration associates the two FDDI ports in the pair. (See Chapter 6.) These associations affect the following situations:
  - When you perform an operation for which you must specify bridge ports (for example, when you define VLANs), you must use the lowest-numbered port in each DAS pair to represent the DAS pair. The operation that you perform then applies to both ports in the DAS pair.
  - When you view information that applies to more than one port (for example, bridging displays for DAS pairs), the port number field identifies both ports in the DAS pair. A VLAN summary display lists all physical ports to indicate which physical system connectors can receive or transmit frames within a VLAN. (You must use the VLAN detail display to see the DAS pairs.)

The configuration of trunks or DAS pairs does not change the port numbering scheme shown in displays such as Ethernet statistics displays or bridge port displays. If you have created trunks or FDDI DAS pairs, however, be aware that a group of ports is associated with each trunk or DAS pair. Therefore, a display such as a bridge port display groups the ports associated with each trunk or DAS pair. See the Command Reference Guide *for examples of the bridge port display commands.*

Examples of Port Numbering

This section provides sample configurations that illustrate port numbering on the system.

Example 1: Fully Loaded System

For a fully loaded system (4 occupied slots) with Fast Ethernet ports, the ports are numbered 1 through 24, starting top left to top right, and then continuing bottom left to bottom right, as shown in Figure 6. (The figure shows the 10/100BASE-TX module.)

Figure 6 Port Numbering for a System with Four Fast Ethernet Modules

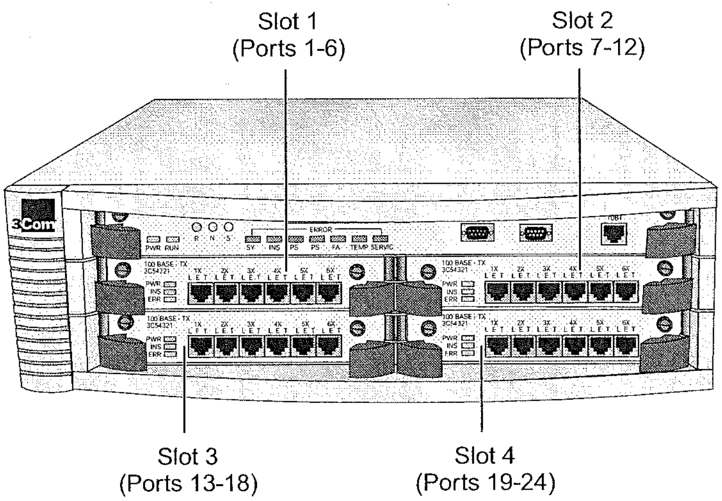

Example 2: Empty Slot in the System

When you have an empty slot, the port numbering includes no ports for that slot. With three Fast Ethernet modules, for example, you have 18 ports, which are numbered according to their position in the system.

For example, if the top-right slot is empty (slot 2), the ports are numbered as shown in Figure 7. (The figure shows the 10/100BASE-TX module.)

Figure 7 Port Numbering for a System with an Empty Slot

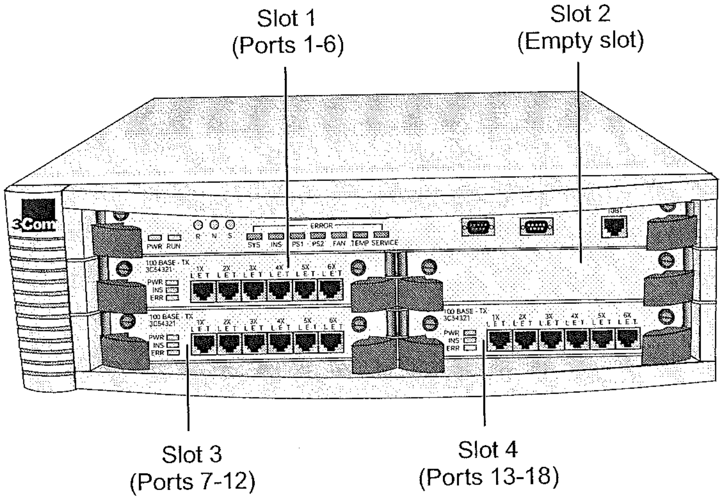

Example 3: Gigabit Ethernet Module with Other Modules

When you have a system with one Gigabit Ethernet module and three Fast Ethernet modules, port numbering accounts for the single port on the Gigabit Ethernet module, as shown in Figure 8.

Figure 8 Port Numbering for a System with a Gigabit Ethernet Module

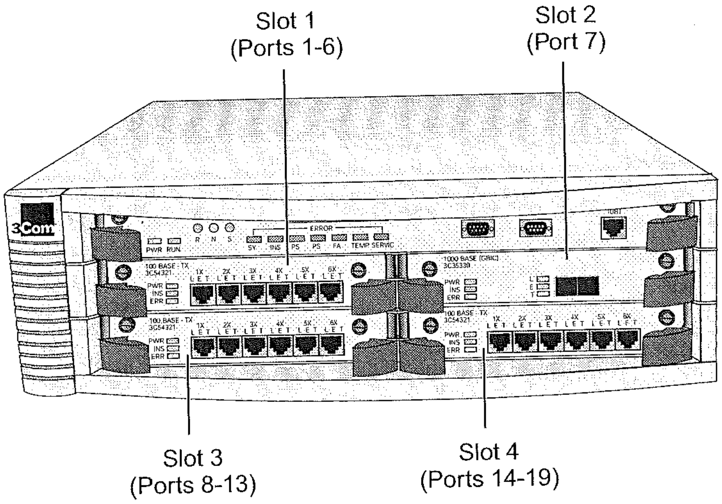

Example 4: FDDI Module with Other Modules

An FDDI module has six FDDI ports (two rows of three ports). Figure 9 shows an FDDI module in slot 1. The top row's ports are numbered 1 through 3 and the bottom row's ports are numbered 4 through 6. Slots 2 and 3 have 10/100 Fast Ethernet modules, and Slot 4 has a Gigabit Ethernet module.

When two FDDI ports are configured as a dual-attach station (DAS) pair, there is one bridge port using two physical (fiber) connectors. The anchor port is the A-port of the DAS port pair. If you configure two FDDI ports as a DAS pair, you must specify the lowest-numbered (anchor) port in the DAS pair and the other port in the pair becomes unselectable.

For example, for the FDDI module shown in slot 1, the three configurable DAS pairs have ports 1 and 4, ports 2 and 5, and ports 3 and 6. When specifying bridge ports (for example, for VLANs), you specify port 1 to represent the first DAS pair, port 2 to represent the second DAS pair, and port 3 to represent the third DAS pair. For more information about FDDI configurations, see Chapter 6.

Figure 9 Port Numbering for a System with an FDDI Module

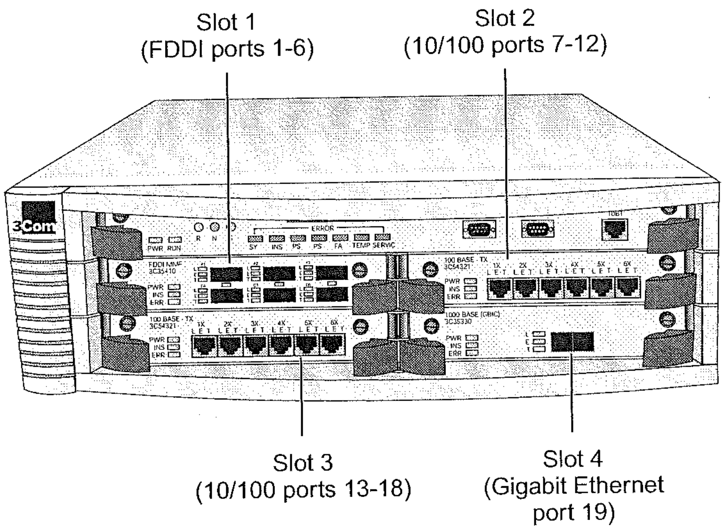

Effects of Removing a Module

When you remove a module and leave the slot empty, a number of changes occur.

Port-Numbering Changes

The ports are sequentially renumbered when you remove a module from slot 1, 2, or 3. Removing a module in slot 4 does not cause renumbering, only a loss of those ports.

Example If you have a fully loaded system with four 10/100BASE-TX modules and you remove the module in slot 3 (ports 13-18), the ports associated with the module in slot 4 (formerly numbered 19-24) are renumbered to 13-18. (See Figure 6.)

VLAN Changes

When you remove a module, VLAN changes occur as follows:

- If you have a VLAN that contains ports that have been renumbered, the renumbered ports now appear in the VLAN summary display.

  Example If a VLAN contained ports 20 through 22 before you removed the module in slot 3, these ports show up as ports 14 through 16 in the VLAN summary after you remove the module.

- If you have a VLAN that includes ports associated with the removed module, those ports are removed from the VLAN and the VLAN summary display no longer shows those ports. (This change includes trunk ports.)

  Example If a VLAN contained ports 17 through 24 before you removed the module in slot 3, the removal of the module in slot 3 causes the removal of previous ports 17 and 18 from the VLAN. (See Figure 6.) The VLAN then contains the renumbered ports 13 through 18 (previously ports 19-24).

- If there are no remaining ports in the VLAN once you remove the module, the VLAN summary display shows the VLAN without any ports.

See Chapter 9 for more information about VLANs.

Trunk Changes

When you remove a module, trunk changes occur as follows:

- If you have a trunk that includes ports associated with the removed module, the trunk display shows that the trunk has `Missing` ports.

  Example If you had a trunk on ports 17 through 20 before you removed the module in slot 3, the removal of that module causes the trunk to have two missing ports (17 and 18). (See Figure 6.) It now has renumbered ports 13 and 14 (previously ports 19 and 20).

- If there are no remaining ports in the trunk after the module is removed, the trunk summary display shows the trunk without any ports.

  Example If you had a trunk with ports 13 through 16 before you removed the module in slot 3, the trunk summary now shows an empty port list.

See Chapter 8 for more information on trunking.

Effects of Replacing Modules

When you remove a module, a number of changes occur, depending on the replacement module.

Replacing Modules of the Same Type or Same Number of Ports

If you remove a module that does not have any trunks or DAS ports and replace it with another module that has the same number of ports, the following changes occur:

- The port numbering is not affected — 10/100 Ethernet and FDDI modules can be exchanged without affecting the port numbers.

One Gigabit Ethernet module in any slot other than slot 4 must be replaced by another Gigabit Ethernet module to prevent port renumbering.

- The system remembers ports that were members of a VLAN. When another module is inserted into the empty slot, the ports are added back into the VLAN.

Replacing Modules of Different Types

More complicated changes occur when you swap six-port modules and one-port Gigabit Ethernet modules, replace FDDI modules that have DAS port pairs (because a DAS pair uses one bridge port to represent two physical ports), or replace modules on which you have trunks defined (because only the anchor port is used to define a trunk in a VLAN).

Port-Numbering Changes

Swapping six-port 10/100 Ethernet modules and FDDI modules does not cause the system to renumber ports, but swapping six-port modules and Gigabit Ethernet modules *does* cause the system to renumber ports.

Example

If you have four 10/100 Ethernet modules, and you replace a 10/100 Ethernet module in slot 1 with a Gigabit Ethernet module, the ports are renumbered as follows: ports 1-6 become port 1, ports 7-12 become ports 2-7, ports 13-18 become ports 8-13, and ports 18-24 become ports 14-19.

VLAN Changes

- If you replace a six-port module with a Gigabit Ethernet module, the ports are renumbered, and any preexisting VLANs now include the Gigabit Ethernet port *only* if the VLANs previously included the first port of the six-port module.

  Example If a VLAN contained ports 1 through 12 before you replaced the 10/100 Ethernet module in slot 1 with a Gigabit Ethernet module, the VLAN contains ports 1 through 7 after the change. (Port 1 is the Gigabit Ethernet port.)

- If a VLAN is defined over a Gigabit Ethernet module and you replace the module with a six-port module, only the *first* port of the new six port module is included in the VLAN after the change.

  Example If a VLAN is defined over a Gigabit Ethernet module in slot 1 and six ports in slot 2 (that is, the VLAN has ports 1 through 7 configured) and you replace the Gigabit Ethernet module with an FDDI module with all SAS ports, the VLAN contains ports 1,7 through 12 after the change.

- If a VLAN has a DAS pair on an FDDI module, and the stationMode for that DAS port pair is changed to SAS (or the FDDI module is replaced by a six-port module), only the first SAS port of the previous DAS port pair is included in the VLAN after the change.

  Example If a VLAN is defined over three DAS ports of an FDDI module in slot 1 and six Ethernet ports in slot 2 (that is, the VLAN has ports 1-12), and you change the FDDI ports from DAS to SAS, the VLAN contains ports 1 through 3 and 7 through 12 after the change.

Trunk Changes

- If you remove a module of a specific type that has trunks and replace it with a module of another type, the new ports do not become part of the trunk. When you define a trunk, the trunk is associated with a specific media type (100 Mb, Gigabit, or FDDI).

Example If you replace a 10/100 Ethernet module in slot 1 (that has a trunk on ports 5 and 6) with an FDDI module, the new FDDI ports 5 and 6 do *not* become part of the trunk. In this case, removing the 10/100 Ethernet module causes the 100 Mb trunk to lose all of its ports, although the trunk itself remains configured on the system.

Special Case: If you have four trunks and you replace a module of a given type with a Gigabit Ethernet module, the system cannot recognize the new Gigabit Ethernet module, because this module type uses a trunk resource. In this case, you must remove one of your trunks before you add the Gigabit Ethernet module (for example, a trunk associated with the removed module).

- If you replace a module that has a trunk spanning another module, after the change, the trunk is missing the ports associated with the removed module.

Example If a trunk spans two 10/100 Ethernet modules (ports 5-8) in slots 1 and 2, and the module in slot 1 is replaced by an FDDI module, after the change, the trunk display shows ports 5 and 6 as `Missing`, but the trunk still has the ports from the other Ethernet module (ports 7 and 8).

For more information on trunking, see Chapter 8. For information on VLANs, see Chapter 9.

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
- Flow Control
- PACE Interactive Access
- 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
- **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 on the CoreBuilder® 3500:

- **Port state** — Whether a port is enabled and connected to a cable (on-line) or disabled (off-line)
- **Port label** — An alphanumeric port identifier
- **Port mode** — Port speed (10 Mbps, 100 Mbps, 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** — 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:

- **Trunking** — Increases bandwidth between switches and servers
- **Trunk Control Message Protocol (TCMP)** — Increases the availability of trunked links by handling physical configuration errors
- **Gigabit Interface Converter (GBIC)** — A Gigabit Ethernet port media type that allows you to hot-swap one media connector without affecting the other connectors

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 Open Systems Interconnection (OSI) Layer 2. 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 trunking. A trunk enhances network availability because its underlying TCMP technology detects and handles physical configuration errors in point-to-point configurations. 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
- Increased support for workstation moves, adds, changes, and upgrades
- Low-cost expansion of switch-to-switch or switch-to-server bandwidths without having to change device modules or cabling
- With PACE Interactive Access, reduction of network jitter, improved network timing, and optimization of LAN bandwidth use

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), 1000 Mbps (Gigabit Ethernet)
- **Port state** — Whether a port is enabled and connected to a cable (on-line) or disabled (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.
- **Trunking** — A technology that combines multiple Fast Ethernet or Gigabit Ethernet ports into a single high-speed channel, thereby increasing bandwidth between switches and between servers and switches
- **Trunk Control Message Protocol (TCMP)** — A protocol that detects and handles physical configuration errors in a point-to-point configuration, thereby increasing availability of trunked links
- **Gigabit Interface Converter (GBIC)** — A Gigabit Ethernet port media type that allows you to hot-swap one media connector without affecting the other connectors

- **PACE® Interactive Access** — 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 Frame Processing

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

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

Figure 10 shows the order in which frame discard tests are made.

Figure 10 How Frame Processing Affects Ethernet Receive Frame Statistics

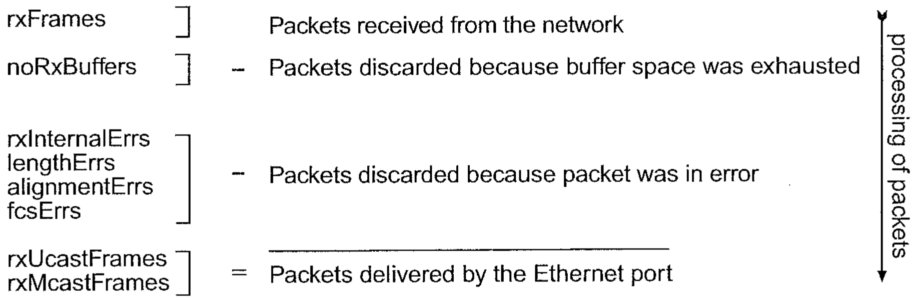
rxFrames
Packets received from the network
noRxBuffers
–
Packets discarded because buffer space was exhausted
rxInternalErrs
lengthErrs
alignmentErrs
fcsErrs
–
Packets discarded because packet was in error
rxUcastFrames
rxMcastFrames
=
Packets delivered by the Ethernet port
processing of packets

Frames also may be delivered directly to an Ethernet port by bridge, router, or management applications. A transmitted frame 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 frame transmission.

Figure 11 shows the order in which these discard tests are made.

Figure 11 How Frame Processing Affects Ethernet Transmit Frame Statistics

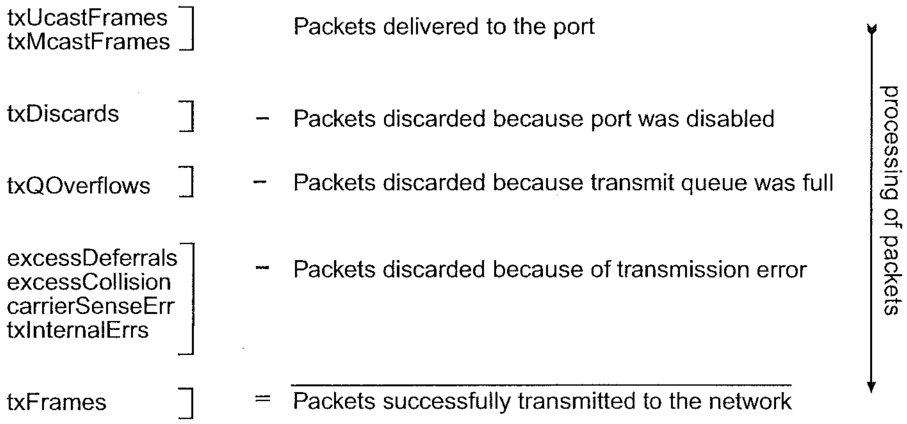
txUcastFrames
txMcastFrames
Packets delivered to the port
txDiscards
– Packets discarded because port was disabled
txQOverflows
– Packets discarded because transmit queue was full
excessDeferrals
excessCollision
carrierSenseErr
txInternalErrs
– Packets discarded because of transmission error
txFrames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

Key Guidelines for Implementation

Consider these important factors when you implement 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 5. In areas that may benefit from 1000 Mbps pipelines, you may be able to substitute trunked Fast Ethernet, subject to the issues raised in Chapter 8.

Table 5 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+

Trunks

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 area
- 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 port-numbering information on the CoreBuilder 3500, see Chapter 4. For more information 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

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- 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 that are connected to a network cable.

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

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- 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 6 lists Ethernet port types on your system, whether they support autonegotiation, and which features they negotiate.

Table 6 Port Types and Autonegotiation Attributes

Port Type	Supports Autonegotiation?	Negotiable Attributes	Default Values for Negotiable Attributes
10/100BASE-TX	Yes	Port speed	10 Mbps
		Duplex mode	Half-duplex
100BASE-FX	No	Not applicable	Not applicable
1000BASE-SX	Yes	Duplex mode	Full-duplex
		Flow control	If autonegotiation is enabled, the system's best effort is On
1000BASE-LX GBIC	Yes	Duplex mode*	Full-duplex*
		Flow control	If autonegotiation is enabled, the system's best effort is On
1000BASE-SX GBIC	Yes	Duplex mode*	Full-duplex*
		Flow control	If autonegotiation is enabled, the system's best effort is On

* LX GBIC, and SX GBIC duplex modes are fixed at full-duplex at this release.

- **10/100BASE-TX ports** — Enabling autonegotiation causes both the port speed and duplex mode attributes to be autonegotiated.
- **100BASE-FX ports** — No autonegotiation of duplex mode occurs. The port speed is fixed at 100 Mbps. The default duplex mode is half-duplex.
- **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.
- 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.
- Autonegotiation is enabled by default on the ports that support it.

Port Mode

You can change the port speed and duplex mode for the 10/100BASE-TX ports and the duplex mode for 100BASE-FX ports. You cannot change the port speed or duplex mode for Gigabit Ethernet ports.

Important Considerations

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- The device that is connected to each port must be configured for the same port mode. If the port speeds differ, the link does not come up. If the duplex modes differ, link errors occur.
- Gigabit Ethernet ports do not support mode options. The value `all` refers only to ports that support port mode options.
- If you change to full-duplex mode on the port, a message indicates that collision detection will be disabled unless you configure the connected device to the same duplex mode.
- Disable autonegotiation on any port on which you are setting a specific port mode.
- Table 7 lists the duplex port mode options available for each port type.

Table 7 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	
100BASE-FX	100full	100 Mbps, full-duplex	100half
	100half	100 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

Table 8 lists the effects of flow control options.

Table 8 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
off	Port ignores flow control packets and does not generate flow control packets.	Gigabit Ethernet Fast Ethernet
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

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- The default setting for flow control is off.
- The system does not count flow control packets in receive or transmit statistics.

PACE Interactive Access

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

Important Considerations

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- 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.

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

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 sending stations stop transmitting, send a collision alert to other stations, and wait a random amount of time before trying again.

Media Specifications

Table 9 summarizes the system's Ethernet media options.

Table 9 Ethernet Media Specifications

Type	Speed	Media	Connector	Recommended Distance (max)
10/100BASE-TX	10/100 Mbps	Category 5 UTP	RJ-45	100 m
100BASE-FX	100 Mbps	single-mode fiber	SC	20 km
		multimode fiber	SC	412 m (half-duplex) 2 km (full-duplex)
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 GBIC	1000 Mbps	single-mode fiber	GBIC	5 km (9 micron) (qualified for up to 10 km)
		multimode fiber	GBIC, with duplex SC conditioned launch cable	550 m (62.5 and 50 micron @ all modal bandwidths)
1000BASE-SX GBIC	1000 Mbps	multimode fiber	GBIC	550 m (62.5 and 50 micron @ all modal bandwidths)

Gigabit Ethernet Interface Converter (GBIC) ports are hot-swappable, that is, you can replace one GBIC connector while the other connectors continue to carry traffic.

To ensure optimal compatibility, performance, and regulatory compliance, use only GBIC transceivers and conditioned launch cables that 3Com supports. For information about currently supported GBIC specifications and conditioned launch cables, see the 3Com Web site:

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

Related Reading

For information about Ethernet media options, see the *CoreBuilder 3500 Getting Started Guide*.

6

Fiber Distributed Data Interface (FDDI)

This chapter provides an overview, key concepts, guidelines, and other key information about how to configure Fiber Distributed Data Interface (FDDI) in your system. This chapter covers these topics:

- FDDI Overview
- Key Concepts
- Key Guidelines for Implementation
- FDDI Stations
- FDDI Paths
- FDDI MACs
- FDDI Ports
- Station Mode (DAS and SAS)
- Sample FDDI Configurations

You can manage FDDI in either of these ways:

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

FDDI Overview

Fiber Distributed Data Interface (FDDI) is a standards-based solution that provides fast and reliable data transfer on a local area network (LAN). FDDI technology, which supports data transfer of 100 million bits per second (100 Mbps), was developed by the American National Standards Institute (ANSI).

Features

FDDI technology:

- Uses optical fiber as its transmission medium, providing security, low signal loss, and high bandwidth data communication.
- Supports simultaneous connection of over 500 nodes on a ring, with up to 2 kilometers (1.2 miles) between adjacent nodes, and up to 200 kilometers (124 miles) of total fiber length.
- Uses a token-passing protocol for access to the network.
- Uses a dual-ring approach: a combination of two independent counter-rotating rings, each running at a data rate of 100 Mbps.
- Is the first LAN technology to provide an embedded network management capability.

Benefits

FDDI offers numerous benefits, many of which originate from the use of fiber-optic cable instead of copper cable.

- The FDDI standard specifies a data rate of 100 Mbps, which allows more data to be sent over optical fiber.
- The distance between nodes using multimode fiber is up to 2 km, which allows for a larger group of network users.
- Radio frequency interference (RFI) or electromagnetic interference (EMI) do not affect fiber-optic cable.
- Fiber-optic cable uses a dual ring topology that:
  - Provides fault tolerance and isolation.
  - Allows for ring wrapping in the event of a fault.
- FDDI uses a token access method that:
  - Supports larger networks.
  - Exploits the cable bandwidth more fully.
  - Eliminates collisions, similar to Carrier Sense Multiple Access/Collision Detect (CSMA/CD).

Key Concepts

Before you implement FDDI in your system, review the following FDDI standards, key concepts, and key terms.

Related Standards

The industry guideline for FDDI technology is divided into four major standards:

- **Physical Medium Dependent (PMD)** — Specifies the characteristics of the fiber-optic medium, the connectors that attach stations to the fiber-optic medium, the transmission wavelength, the power requirements for transmitters, and the methods for optically bypassing inactive stations.
- **Physical (PHY)** — Specifies data encoding and decoding, clock speed and clocking scheme, data framing, and the control symbols used in the network.
- **Media Access Control (MAC)** — Specifies access to the medium, token passing, addressing, data checking, frame generation and reception, error detection and recovery, and the bandwidth allocation among the stations.
- **Station Management (SMT)** — Specifies the FDDI station and ring configurations, initialization and maintenance of station-to-station connections, and the control required for the proper operation of stations in an FDDI ring.

These four standards are always described in relation to the Open Systems Interconnection (OSI) Reference Model. This model was established by the International Standards Organization (ISO) to standardize digital data communications. Each FDDI station is made up of logical entities that conform to the four standards. These entities represent the active services or management elements within OSI.

Figure 12 illustrates the relationship of FDDI entities to the OSI Reference Model. Network attachments communicate with each other using predetermined protocols. The model divides these communication protocols into seven layers, which are defined so that each layer only requires services from the layer below it.

Figure 12 FDDI Relationship to OSI Reference Model

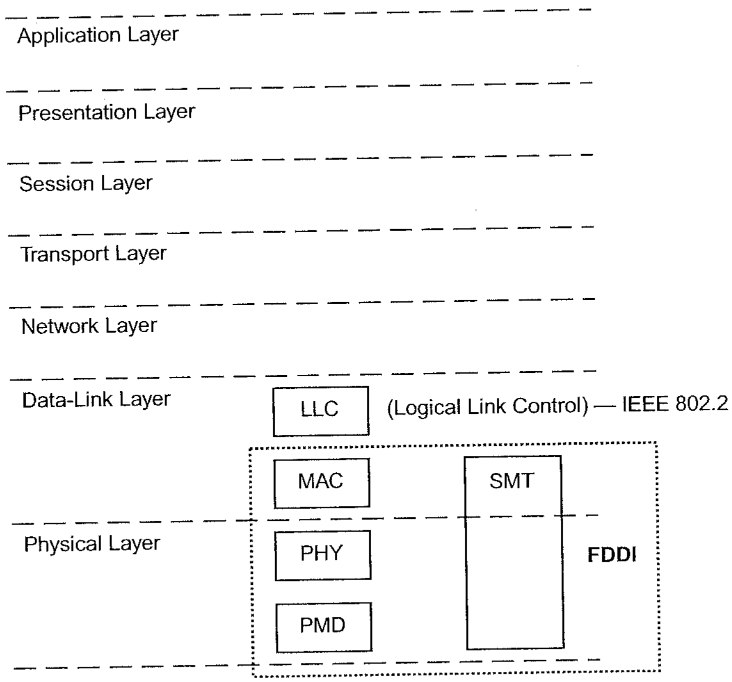

FDDI Network Topologies

The term *network topology* refers to the ways that stations are interconnected within a network. An FDDI network topology may be viewed at two distinct levels:

- **Physical topology** — A network's physical topology is defined by the arrangement and interconnection of its nodes. The FDDI physical topology is a *ring of trees*. See Figure 13.

Figure 13 Physical Topology

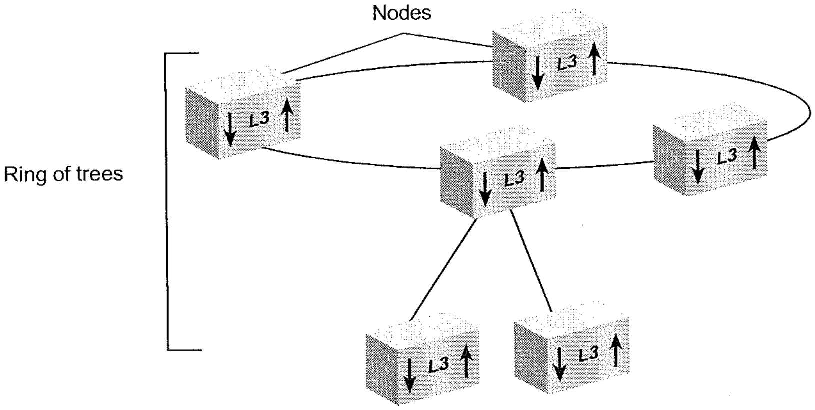

- **Logical topology** — A network's logical topology is defined by the paths through which tokens and data flow in the network. The FDDI logical topology is a *dual ring*. See Figure 14.

Figure 14 Logical Topology

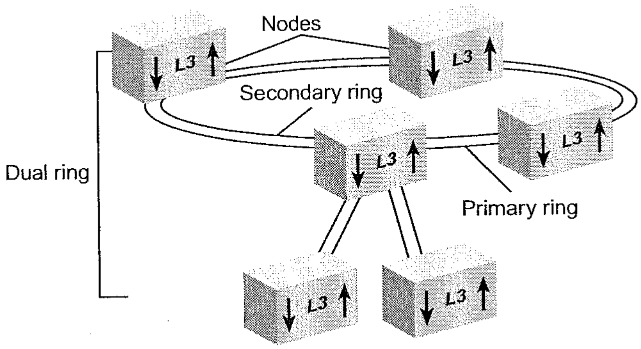

Physical Topology: A Ring of Trees

The FDDI ring consists of dual-attach stations (DASs) and dual-attach connectors (DACs). The DACs on the ring allow you to attach *trees*. The trees consist of *branches* of single-attach stations (SASs) and DASs that are star-wired off of the concentrators. This kind of network is highly reliable, provides a single, fault-tolerant ring, offers fault isolation, and allows centralized management. See Figure 15.

Figure 15 Ring of Trees

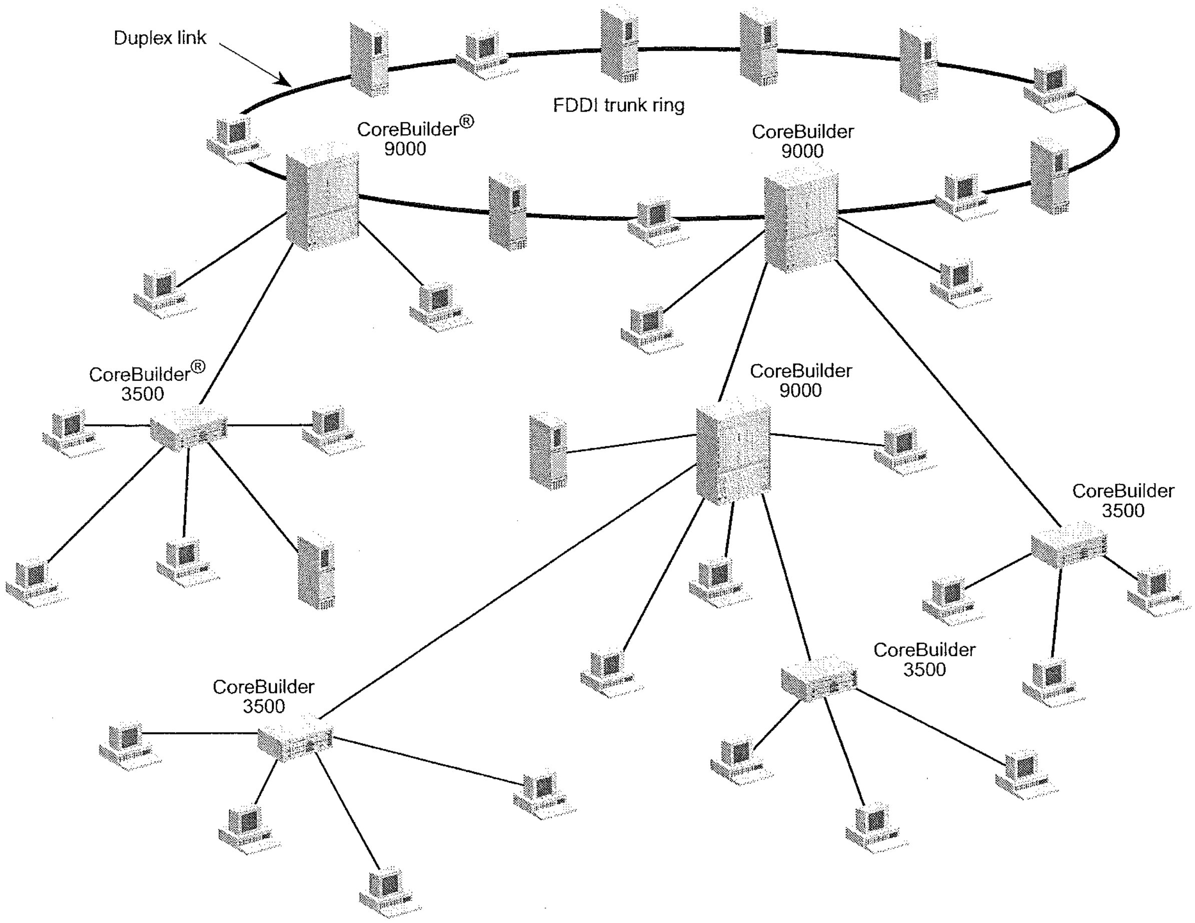

All physical connections in an FDDI topology are *duplex links* (a pair of insulated fiber-optic conductors). Both the FDDI ring and the ring of trees that are created through concentrators are made up of duplex links. Interconnect the nodes in an FDDI network to form at *most* one ring.

If a topology is legal, when physical connections and nodes fail or are removed from the network, one or more legal FDDI topologies are formed. So subsets of legal topologies are also legal. Examples of legal FDDI topologies include the dual ring with trees, the dual ring without trees, and the single tree. For information about legal topologies, see "Setting the Connection Policies" later in this chapter.

Logical Topology: The Dual Ring

A legal FDDI topology consists of at most two separate logical rings: the primary ring and the secondary ring. These logical rings are formed from the physical links that make up the Physical Layer connections. For example, a set of DASs that are connected into a closed loop form an FDDI dual ring (that is, A to B; B to A). Each ring is a logical ring, that is, a separate data path with its own token.

Functionally, the dual ring provides a high degree of reliability to a LAN. When an FDDI network is in normal operation, only the primary ring transmits and receives data. The secondary ring may also carry data, but it is typically used as a backup in case there is a connectivity problem in the primary ring or in one of the nodes on the ring.

When a single fault takes place on an FDDI dual ring, recovery can be made by joining the two rings between the two nodes that are adjacent to the fault. Doing this creates a single logical ring, which results in a wrapped configuration. A wrapped ring is a legal FDDI topology. In the same way, when many faults take place, several disjointed logical rings are created, producing multiple FDDI topologies.

Nodes and Attachments

An FDDI network is made up of stations, concentrators, and switches that contain active services or management elements that conform to the ANSI FDDI standards. These stations and concentrators are connected to optical fiber medium and are attached in the prescribed manner set forth in the FDDI standards to allow reliable data transmission. Connections are made through FDDI ports and are managed by FDDI MACs.

Nodes

An FDDI network is made up of logically connected *nodes*. This generic term is used to refer to any active *station* or *concentrator* in an FDDI network.

- **Station** — Any addressable node on an FDDI network that can transmit, repeat, and receive information. A station contains only one SMT, and *at least one* MAC, one PHY, and one PMD.
- **Concentrator** — An FDDI station with additional PHY/PMD entities, beyond those required for its own connection to an FDDI network. These additional PHY/PMD entities (M ports) connect other FDDI stations, including other concentrators, in a tree topology.

Attachments

Attachments refer to how a node, station, or concentrator is connected to an FDDI network. They are classified as *single attachment* and *dual attachment*. Concentrators can be classified as *null attachment* when the A and B ports are either not present or not used.

- **SAS** — Single Attachment Station. A station or concentrator that has only one physical connection to an FDDI network. The single attachment cannot accommodate a dual (counter-rotating) ring. A single attachment station or concentrator has an S port that attaches to an M port within a concentrator tree.
- **DAS** — Dual Attachment Station. Any station or concentrator that has two physical connections to an FDDI network. This type of attachment can accommodate a dual (counter-rotating) ring. A dual attachment station has one A-B port pair; a dual attachment concentrator has an A-B port pair and at least one M port.

Node Types

Six station and concentrator types are used to describe station configurations and topologies. Table 10 lists these node types and their abbreviations.

Table 10 Node Types and Abbreviations

Node Type	**Abbreviation**
Single MAC-Dual Attachment Station	SM-DAS
Dual MAC-Dual Attachment Station	DM-DAS
Single Attachment Station	SAS
Dual Attachment Concentrator	DAC
Single Attachment Concentrator	SAC
Null Attachment Concentrator	NAC

Figure 16 shows how these six node types may connect to an FDDI dual ring.

Figure 16 Examples of FDDI Node Types

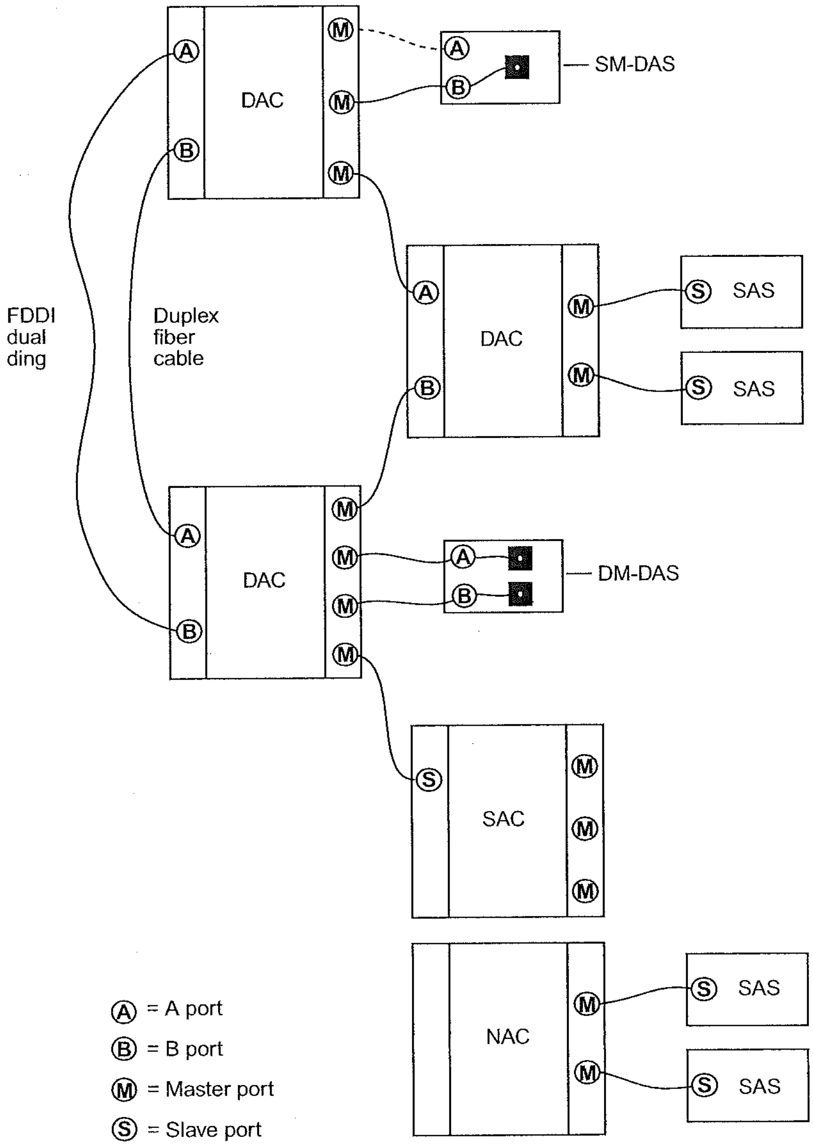

Dual Homing

When the operation of a dual attachment node is crucial to your network, a configuration called *dual homing* can provide added reliability. Using dual homing you can determine a station's operation by setting the appropriate configuration policy. You can configure the dual-homed station with both links active or with one link active and one connection withheld as a backup. The backup connection becomes active only if the primary link fails. See Figure 17.

Figure 17 Dual Homing

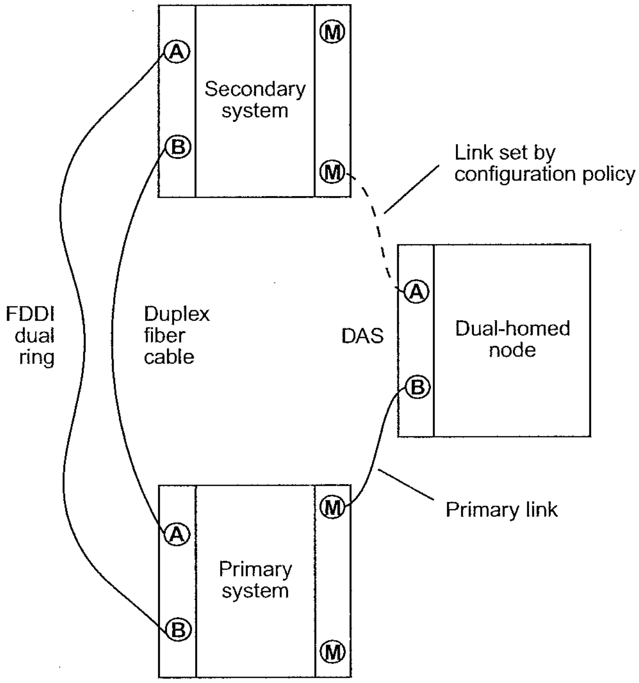

FDDI Stations

Each FDDI station has one Station Management (SMT) entity to provide connection management, ring management, and operational management to the FDDI network. SMT specifies a set of services and signaling mechanisms that are dedicated to FDDI network management. It manages those services of each station on the FDDI network that are specific to the Physical Layer and the MAC portion of the Data Link Layer.

The goal of SMT is to completely define shared medium-management services to guarantee the interoperability of FDDI network equipment from multiple vendors.

SMT Operation

The operation of SMT falls into three broad categories:

- **Physical Connection Management (PCM)** — Establishes and maintains point-to-point physical links between neighboring ports. It provides all the signaling necessary to initialize connections, withhold marginal connections, and support maintenance.
- **Configuration Management (CFM)** — Interconnects PHYs and MACs on paths to achieve proper station configuration and network topology.
- **Ring Management (RMT)** — Manages a MAC's operation in an FDDI ring. RMT detects stations that are *stuck* in the beacon process and initiates the trace function. RMT locates duplicate addresses that might prevent the ring from operating.

FDDI MIB

The FDDI Management Information Base (MIB) defines the collection of information that is available to network management about an FDDI station. The MIB uses an object-oriented approach similar to that used in OSI management standards.

FDDI-managed objects include SMT (that is, the SMT of the station), MACs, paths, and ports. Each of these objects has a collection of attributes such as statistics, error counters, configuration information, event notifications, and actions.

You can access a station's MIB locally through a local management interface or remotely through a management protocol such as Parameter Management Frame (PMF) or Simple Network Management Protocol (SNMP). The SMT standard specifies the meaning and encoding of each MIB attribute.

Frame-based Protocols

SMT provides a number of frame-based services that higher level management functions use to manage stations on the network and to gather information about them. Frame-based protocols:

- Gather network statistics
- Detect, isolate, and resolve faults in the network
- Tune FDDI configuration and operational parameters to meet application and connectivity requirements

SMT has six key frame-based protocols:

- **Neighbor Notification** — Allows SMT to learn the addresses of the logical neighbors of each MAC in a station. This information is useful in detecting and isolating network faults.
- **Parameter Management** — Performs the remote management of station attributes. It operates on all SMT MIB attributes, attribute groups, and actions.
- **Status Reporting** — Allows a station to notify network managers about events such as station configuration changes and network errors.
- **Status Polling** — Provides a mechanism to obtain station status remotely through a request/response protocol.
- **Echo** — Performs loopback testing on the FDDI dual ring.
- **Synchronous Bandwidth Allocation** — Allocates synchronous bandwidth and monitors both synchronous and total bandwidth.

Primary and Secondary Paths

FDDI's dual, counter-rotating ring is made up of a primary and secondary ring. You can be connect FDDI stations to either ring or to both rings simultaneously. Data flows downstream on the primary ring in one direction from one station to its neighboring station. The secondary ring serves as a redundant path and flows in the opposite direction. When a link or station failure occurs, the ring *wraps* around the location of the failure, creating a single logical ring.

Paths represent the segments of a logical ring that pass through a station. An FDDI station can contain two paths:

- **Primary path** — The segment or segments of the primary ring that pass through a station. Conditions may exist in parts of the network that cause the path to be in a different ring. The primary path must be present in all nodes on the network.
- **Secondary path** — The segment or segments of the secondary ring that pass through a station. Conditions may exist in parts of the network that may cause the path to be in a different ring.

Media Access Control

The Media Access Control (MAC) uses a token-passing protocol to determine which station has control of the physical medium (the ring). The MAC delivers frames to their destinations by scheduling and performing all data transfers.

MAC Services

Some of the services that the MAC performs include:

- Frame repetition and reception
- Frame removal
- Frame validity criteria checking
- Token capture
- Token rotation
- Ring initialization
- Beacon process

MAC services are provided by all conforming stations that are attached to the FDDI network.

MAC Operation

The MAC controls access to the physical medium by passing a token around the ring. When a station receives the token, the station may transmit a frame or a sequence of frames. When a station wants to transmit, it removes the token from the ring and transmits the queued frames. After transmission, the station issues a new token, which the downstream station uses.

Stations that are not transmitting only repeat the incoming symbol stream. When repeating, the station determines whether the information was destined for it by comparing the destination address to its own address. If it sees a match, the MAC processes subsequent received symbols or sends them to the Logical Link Control (LLC) in the data-link layer for translation.

Ports

As parts of the Physical Layer, the PHY and PMD entities work together to support each link between FDDI stations. These entities provide the protocols that support the transmission and reception of signals between stations, as well as the optical fiber hardware components that link FDDI stations together. Within an FDDI station, the PHY and PMD entities make up a *port*. Together, they create a PHY/PMD pair that connects to the fiber-optic media and that provides one end of a physical connection with another station.

Ports at both ends of a physical connection determine the characteristics of that physical connection. The protocols that are executed at each port determine whether the connection is accepted or rejected. A connection is accepted if at least one station's policy allows such a connection. A connection is rejected if each station has a policy that disallows the connection.

Each port is one of four types: A, B, M, and S.

- **A port** — Connects to the primary ring on the incoming fiber and the secondary ring on the outgoing fiber. A properly formed FDDI dual ring is composed of a set of stations with the A port of one station connected to the B port of the neighboring station.
- **B port** — Connects to the incoming fiber of the secondary ring and the outgoing fiber of the primary ring.
- **M port** — Used by a concentrator station to provide connections within a concentrator tree. Also referred to as *Master port*.
- **S port** — Used by a single attachment station to provide attachment to an M port within a concentrator tree. Also referred to as *Slave port*.

Key Guidelines for Implementation

Consider the following guidelines when you configure and implement FDDI in your system:

- A high frame error rate often indicates a faulty station on the FDDI ring or a dirty FDDI connector or cable.
- If there is something wrong on your network, you may want to turn off data (user) traffic for a MAC by disabling LLC service. Although you have disabled data traffic from the MAC, the MAC still participates in neighbor notification and is visible to network management.
- The FDDI MAC path selections depend on the stationMode configuration (DAS or SAS).
- FDDI ports can be type A or type B for DAS ports.
- FDDI ports can be type S or type M for SAS ports.

- Before the new FDDI stationMode takes effect, you must reboot your system.
- You cannot modify fddi-stationMode port-pairs when any of the ports in the pair are members of a trunk.
- In a DAS configuration, the Activity LED from the secondary port (channel B) is not used by the port. If network traffic is passing through Channel B, the Activity LED on Channel A indicates port activity.

FDDI Stations

You can set the following FDDI station parameters:

- Connection policies
- Neighbor notification timer
- Status reporting

Setting the Connection Policies

The connectPolicy attribute is a bit string that represents the connection policies that are in effect on a station. A connection's *type* is defined by the types of the two ports involved (A, B, M, or S) in the connection. You can set the corresponding bit for each of the connection types that you want a particular station to reject.

The system's FDDI ports can be of type A or type B. By default, all connections to the systems FDDI ports are valid. Table 11 lists the possible connections to reject and their corresponding bits.

Effects and Consequences

When you set the connection policies, consider the following:

- By default all connections are valid on the system. An M-to-M connection is accepted so that a system port can be connected to another system port.
- Although an M-to-M connection is illegal within the FDDI standard, the CoreBuilder 3500 system allows this connection.

Table 11 Bit to Set for Rejecting a Station Connection

This Connection Is Rejected **(System port - Remote port)**	**If This Bit Is Set**	**Connection Rules**
A-A	0	Undesirable peer connection that creates twisted primary and secondary rings; notify station management (SMT).
A-B	1	Normal trunk ring peer connection.
A-S	2	Undesirable peer connection that creates a wrapped ring; notify SMT.
A-M	3	Tree connection with possible redundancy. The node may not go to Thru state in Configuration Management (CFM). In a single MAC node, Port B has precedence (with defaults) for connecting to a Port M.
B-A	4	Normal trunk ring peer connection.
B-B	5	Undesirable peer connection that creates twisted primary and secondary rings; notify SMT.
B-S	6	Undesirable peer connection that creates a wrapped ring; notify SMT.
B-M	7	Tree connection with possible redundancy. The node may not go to Thru state in CFM. In a single MAC node, Port B has precedence (with defaults) for connecting to a Port M.
S-A	8	Undesirable peer connection that creates a wrapped ring; notify SMT.
S-B	9	Undesirable peer connection that creates a wrapped ring; notify SMT.
S-S	10	Connection that creates a single ring of two slave stations.
S-M	11	Normal tree connection.
M-A	12	Tree connection with possible redundancy.
M-B	13	Tree connection with possible redundancy.
M-S	14	Normal tree connection.
M-M	15	Illegal connection that creates a tree of rings topology.

Setting Neighbor Notification Timer

The T-notify attribute is a timer that the Neighbor Notification protocol uses to indicate the interval of time between the generation of Neighbor Information Frames (NIF). NIF frames allow stations to discover their upstream and downstream neighbors. The T-notify value has a range of 2 through 30 seconds, with a default value of 30 seconds.

Effects and Consequences

When you set the neighbor notification timer, consider the following:

- By setting the T-notify value low, your network reacts quickly to station changes, but more bandwidth is used.
- By setting the T-notify value high, less bandwidth is used, but your network does not react to station changes as quickly.

Enabling and Disabling Status Reporting

The statusReporting attribute controls whether a station generates Status Report Frames (SRFs) to report events and conditions to network management stations. By default, status reporting is `enabled`. If you do not have an SMT management station that listens to these event reports or if you use SNMP to monitor FDDI events on all FDDI end-stations, you can set this attribute to `disabled` so that the station does not generate SRFs.

FDDI Paths

You can display FDDI path information and set the time values of the following attributes:

- tvxLowerBound
- tmaxLowerBound
- maxTreq

Setting tvxLowerBound

The tvxLowerBound attribute specifies the minimum time value of fddiMAC TvxValue that any MAC that is configured on this path uses. A MAC uses its valid transmission timer (TVX) to detect and recover from certain ring errors. If a valid frame has not passed through a MAC during the time indicated by fddiMACTvxValue, the MAC reinitializes the ring.

Effects and Consequences

When you set the tvxLowerBound attribute, consider the following:

- By adjusting the tvxLowerBound value, you specify how quickly the ring recovers from an error. The lower that you set this value, the faster the network reacts to problems, but the ring may reinitialize when there is no problem. The recommended value for tvxLowerBound is `2500` microseconds.
- The higher that you set the tvxLowerBound value, the less chance of frequent reinitializations, but the network takes longer to recover from errors. The recommended value for tvxLowerBound is `2500` microseconds.

Setting tmaxLowerBound

The tmaxLowerBound attribute specifies the minimum time value of fddiMAC T-Max that any MAC that is configured on this path uses. This value specifies the boundary for how high T-Req (the requested token rotation time) can be set.

Setting maxT-Req

The maxT-Req attribute specifies the maximum time value of fddiMACT-Req that any MAC that is configured onto this path uses. T-Req is the value that a MAC bids during the claim process to determine a ring's operational token rotation time, T_Opr. The lowest T-Req bid on the ring becomes T_Opr.

Effects and Consequences

When you set the maxT-req, consider the following:

- When T_Opr is a low value, the token rotates more quickly, so token latency is reduced. However, more of the ring's available bandwidth is used to circulate the token.
- Higher values of T_Opr use less bandwidth to circulate the token, but they increase token latency when the ring is saturated.

FDDI MACs

You can display MAC statistics and configure the following parameters:

- MAC FrameErrorThreshold
- NotCopiedThreshold
- Logical Link Control (LLC) service

Setting the Frame Error Threshold

The FrameErrorThreshold attribute determines when the system generates a MAC condition report because too many frame errors have occurred. A frame error occurs when a frame becomes corrupted.

Station Management (SMT) monitors the ratio of frame errors to all frames that are transmitted within a certain period of time. The FrameErrorThreshold setting determines at what percentage the frame errors are significant enough to report to network management. The threshold value is expressed in a percentage based on 65536 (which is 100 percent). For example, to set the threshold at 1 percent, the value is 655 (the system default). The lower that you set the percentage, the more likely it is for SMT to report a problem.

Effects and Consequences

When you set the frame error threshold, consider the following:

- A high error rate often indicates a faulty station on the FDDI ring or a dirty FDDI connector.

Setting the Not Copied Threshold

The NotCopiedThreshold attribute determines when the system generates a MAC condition report because too many frames could not be copied. Not-copied frames occur when there is no buffer space available in the station (which in turn indicates congestion in the station).

SMT monitors the ratio of frames that are not copied to all frames that are transmitted within a certain period of time. The NotCopiedThreshold setting determines at what percentage the number of frames that are not copied is significant enough to report to network management. The threshold value is expressed in a percentage based on 65536 (which is 100 percent). For example, to set the threshold at 1 percent, the value is 655 (the system default). The lower that you set the percentage, the more likely it is for SMT to report a problem.

Enabling and Disabling LLC Service

The Logical Link Control (LLC) service allows LLC frames to be sent and received on the MAC. LLC frames are all data frames that are transmitted on the network. If there is something wrong on your network, turn off data (user) traffic for a MAC by disabling LLC service. Although you have disabled data traffic from the MAC, the MAC still participates in neighbor notification and is visible to network management.

FDDI Ports

You can display port statistics and configure the following port parameters:

- lerAlarm
- lerCutoff
- port labels

Setting lerAlarm

The lerAlarm attribute is the link error rate (LER) value at which a link connection generates an alarm. If the LER value is greater than the alarm setting, then SMT sends a Status Report Frame (SRF) to the network manager software indicating a problem with a port.

Effects and Consequences

When you set the lerAlarm attribute, consider the following:

- The lerAlarm value is expressed as the absolute value of the exponent (such as 1 x 10^{-10}).
- A healthy network has an LER exponent between 1 x 10^{-10} and 1 x 10^{-15}.
- Set the lerAlarm value below these values so that you only receive alarms if your network is in poor health. The SMT Standard recommended value is 8.

Setting lerCutoff

The lerCutoff attribute is the link error rate estimate at which a link connection is disabled. When the lerCutoff value is reached, the PHY that detected a problem is disabled.

Effects and Consequences

When you set the lerCutoff attribute, consider the following:

- The lerCutoff value is expressed as an exponent (such as 1 x 10^{-10}).
- A healthy network has an LER exponent between 1 x 10^{-10} and 1 x 10^{-15}.
- Set the lerCutoff below these values so that a port is only removed only as a last resort. The SMT Standard recommended value is 7.
- The lerCutoff value must be lower than the lerAlarm value so that the network manager software is alerted to a problem before the PHY (port) is actually removed from the network.

Setting Port Labels

Port labels serve as useful reference points and as an accurate means of identifying your ports for management. Label your FDDI ports for easy identification of the devices attached to them (for example, workstation, server, FDDI backbone).

Station Mode (DAS and SAS)

You can modify the FDDI station mode that is assigned to a specific port number to either DAS (Dual Attachment Station) or SAS (Single Attachment Station) S port or M port. For the new station mode to take effect, you must reboot your system.

Single Attachment Station (SAS)

If you configure the FDDI ports as single-attached stations, each port is selectable as a bridge port. You can select your SAS ports to be either S or M ports.

Dual Attachment Stations

If you configure the FDDI ports as dual-attached stations, you must specify the lowest-numbered (anchor) port in the DAS pair. The other port becomes unselectable.

Effect and Consequence

When you set the station mode, consider the following:

- When you modify the station mode, any FDDI ports that are associated with a VLAN or a trunk are removed from the VLAN or trunk.

Sample FDDI Configurations

You can install your system into many possible FDDI configurations. Figure 18 shows systems attached to an FDDI dual ring. The connection to the dual ring is made by the A and B ports on the system. DASs, excluding concentrators, may be attached to the dual ring, as shown.

⚠ ***CAUTION:*** *3Com strongly recommends that you connect equipment that can be turned on and off, such as workstations, only through concentrators. Connect intermediate systems that are seldom turned off, such as bridges and routers, to the FDDI dual ring only if they are equipped with an optical bypass switch. These precautions protect the integrity of the dual ring.*

Figure 18 Sample FDDI Configuration

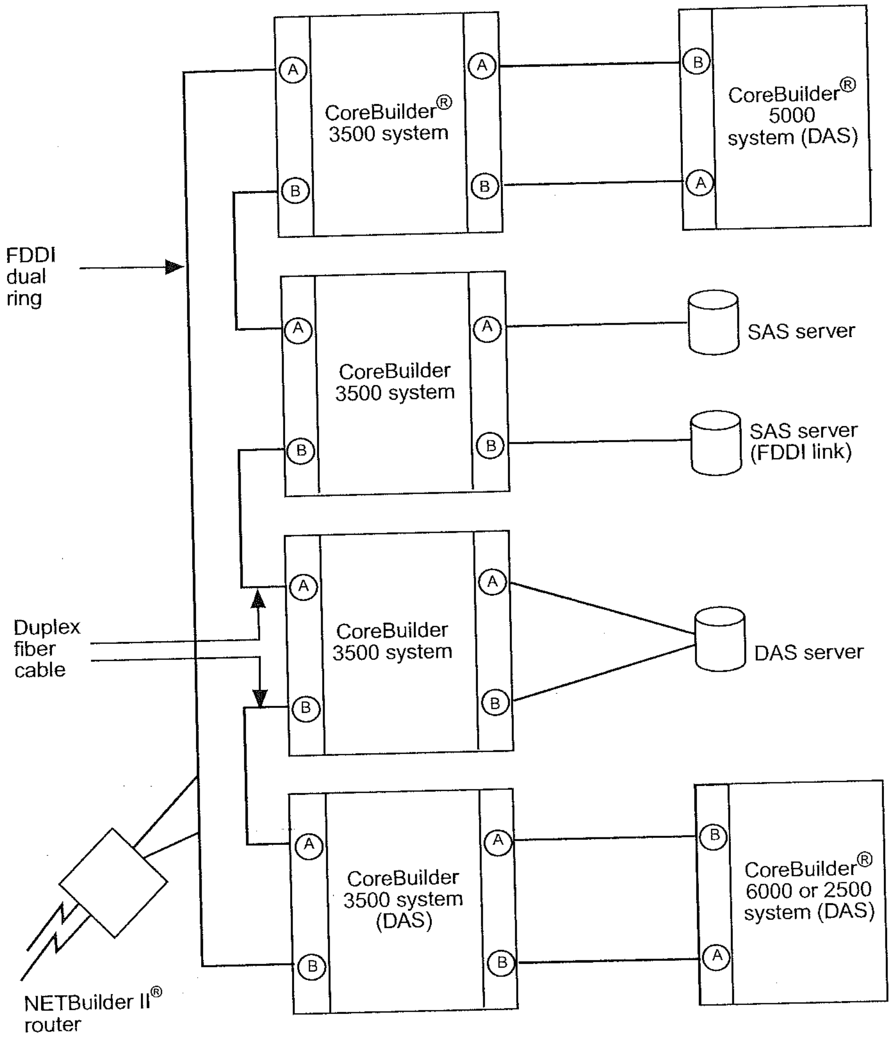

Standards, Protocols, and Related Reading

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

Requests For Comments (RFCs)

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

- **RFC 1130** — A Proposed Standard for the Transmission of IP Datagrams over FDDI Networks
- **RFC 1390** — Transmission of IP and ARP over FDDI Networks
- **RFC 1512** — FDDI Management Information Base

You can obtain copies of RFCs from 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)
- Electronic Industry Association (EIA)
- 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 FDDI, be sure to refer to the following books:

- *Understanding FDDI: A 100Mbps Solution for Today's Corporate LANs*. Andrew Mills, Prentice Hall, 1995.
- *FDDI: A High Speed Network*. Amit Shah, G. Ramakrishnan, Akrishan Ram (Contributor), Prentice Hall, 1993

7

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
- IP Fragmentation
- IPX SNAP Translation
- Broadcast and Multicast Limit for Bridge Ports
- GARP VLAN Registration Protocol (GVRP)
- 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 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.

You 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 (STP) 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.
- **GARP VLAN Registration Protocol (GVRP)** — You can enable your system to transmit and receive VLAN information using GVRP. GVRP is also addressed in Chapter 9.
- **IP fragmentation** — When Fiber Distributed Data Interface (FDDI) stations transmit IP packets that are too large for standard Ethernet to handle, IP fragmentation allows your system to reformat large packets into smaller sizes that can be bridged to Ethernet networks.
- **IPX SNAP translation** — IPX SNAP translation allows any 802.3_RAW IPX packets that are forwarded from Ethernet to FDDI to be translated to FDDI_SNAP (instead of FDDI_RAW), and vice versa.

The following bridging topics are covered in other chapters:

- **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, including discussion of GVRP, see Chapter 9.
- **Trunking** — You can configure your system to aggregate multiple bridge port links into a single point-to-point trunk. Trunking allows you to increase bandwidth and redundancy without replacing cabling. For more information about trunking, see Chapter 8.

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 32 K addresses in its address table.

Aging Addresses

A dynamic address remains in the bridge's 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 dynamic 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 be 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 up the LAN bandwidth and eventually affects the bridge's processing capability.

A backup or redundant path remains 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 the values in 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 19.

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" later in this chapter.

Figure 19 STP Blocks Redundant Links

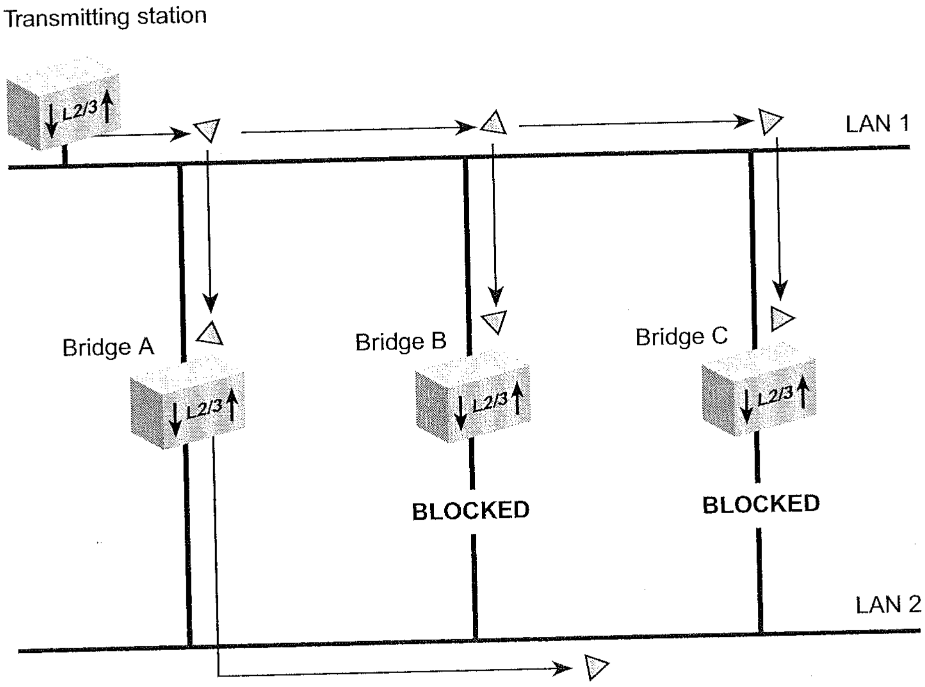

How the Spanning Tree Protocol Works

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 20 shows the hierarchy of the STP bridges and their ports.

Figure 20 Hierarchy of the Root Bridge and the Designated Bridge

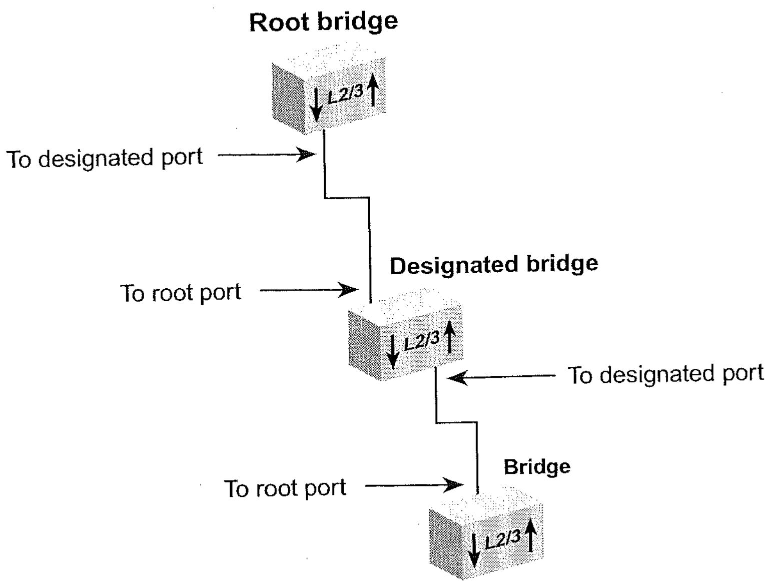

Actions That Result from CBPDU Information

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 21 shows a bridged network with its STP elements.

Figure 21 STP Root and Designated Bridges and Ports

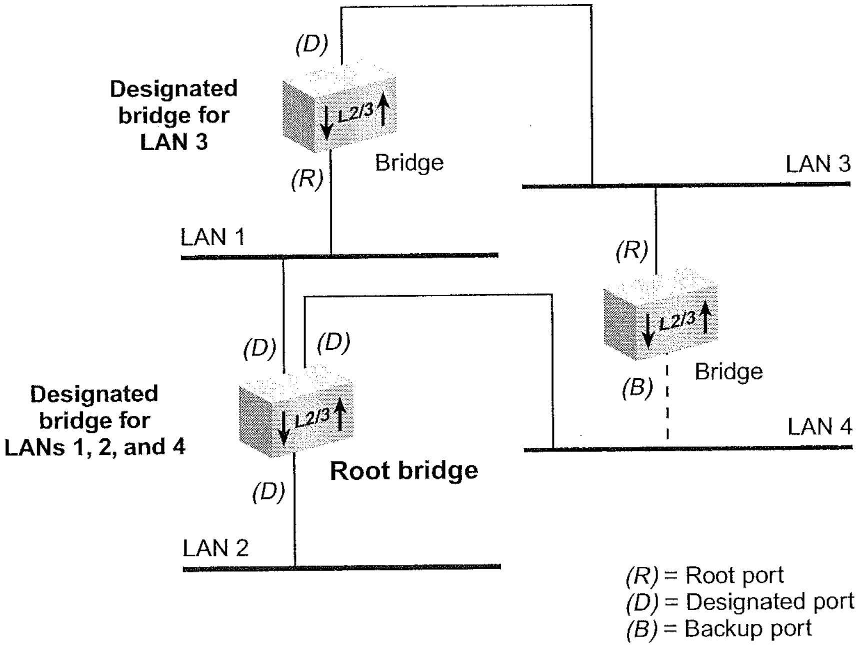

Contents of CBPDUs

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

Figure 22 and Figure 23 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 22 shows the network at the start of the STP topology calculation. Figure 23 shows the network after the STP topology has stabilized.

Figure 22 Starting the Spanning Tree Calculation

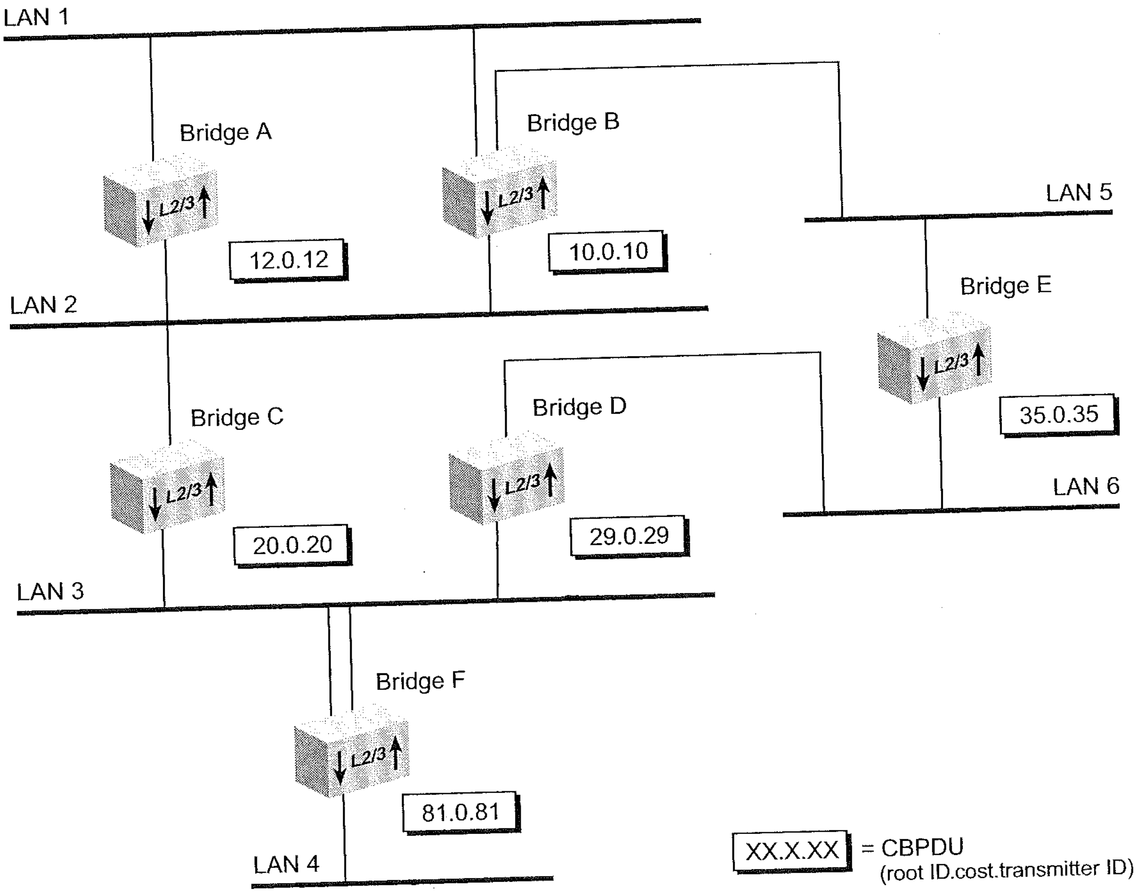

Figure 23 Spanning Tree Topology Calculated

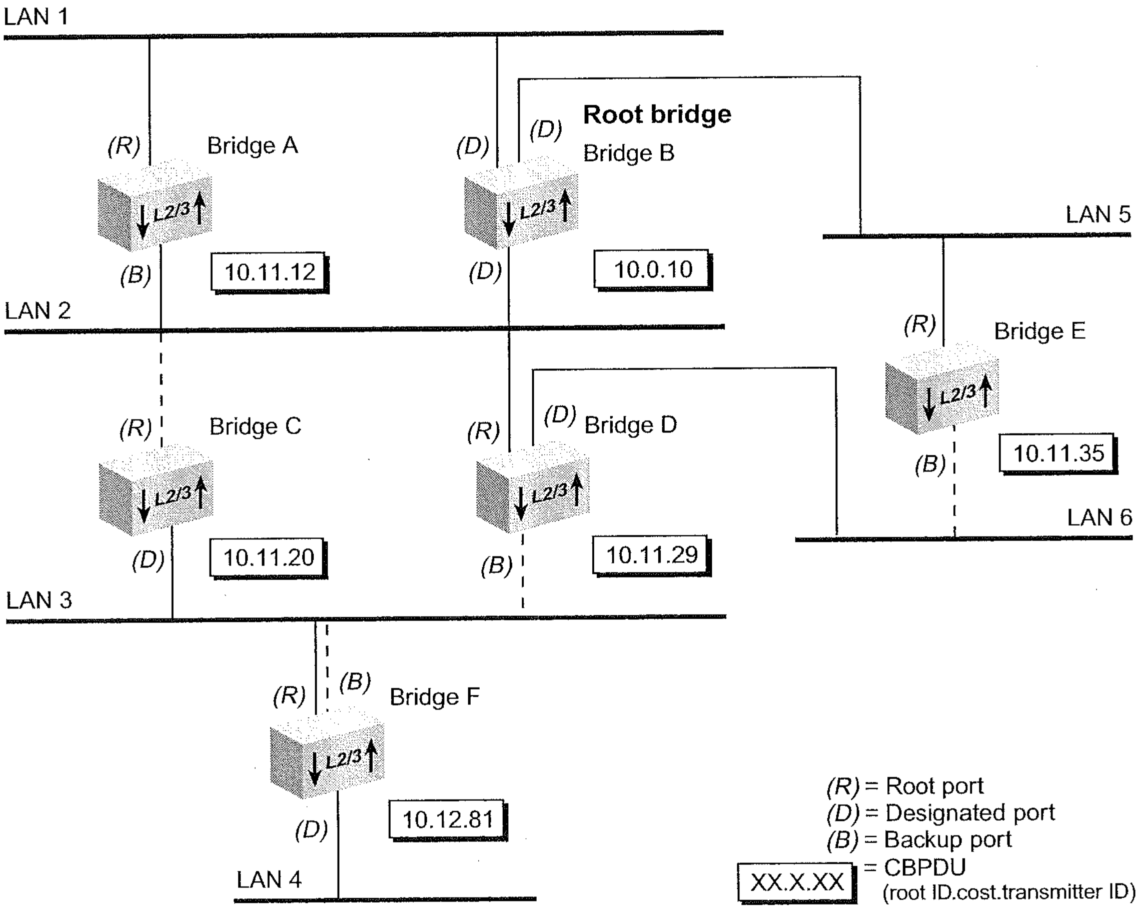

Determining the Root Bridge

The root ID portion of the CBPDU determines which bridge actually becomes the root bridge. In Figure 22, 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 23, 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 23, 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 23, 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 back 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 24 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 24 Factors in Spanning Tree Port State Transitions

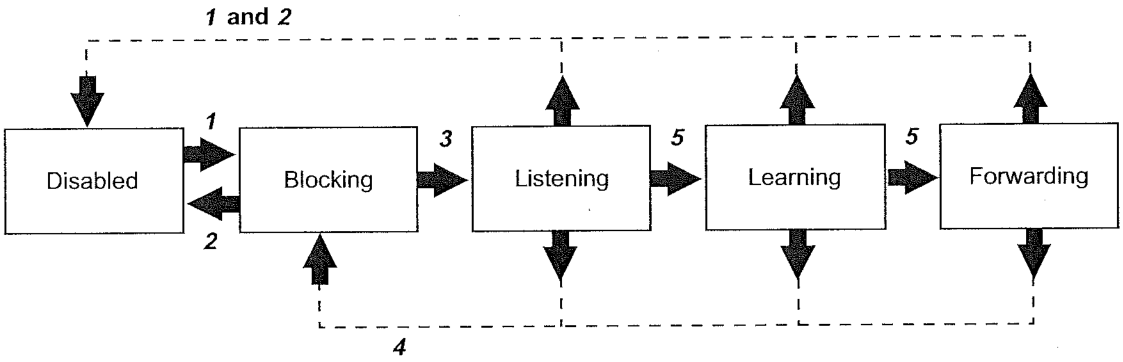

1 Port enabled by either network administrator or initialization

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 24, 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 from its address table. This action ensures that the bridge learns the correct addresses and paths and continues to forward packets to the correct LANs.

Key Guidelines for Implementation

Consider the following guidelines when you configure bridge-wide and bridge port parameters on your system:

- When you disable bridge-wide STP, the bridge cannot participate in the algorithms for loop detection.
- 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 port 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 4 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 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 has a unique address table.
- The system includes an "ignore STP mode" option that affects VLAN configurations. See Chapter 9 for more information or see the *Command Reference Guide.*
- GVRP is useful only when there are other switches or NICs in the network that support GVRP.
- You can define up to 32 bridge ports on the system. One consideration is that if you configure two or more ports of any technology type to form a trunk (a single logical bridge port), the system counts all ports in the trunk toward the bridge port limit.

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 set 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 all devices are configured with the same group address.

Administering STP Parameters on Bridge Ports

You can enable, disable, or remove the Spanning Tree Protocol for one or more ports on the system. This setting affects the operation of a port only if the STP is enabled for the bridge. 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 user-defined packet filter indicated not to receive the frame.

A frame that is forwarded from a physical interface to a bridge port is then transmitted to 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.
- A user-defined packet filter indicated not to forward the frame.

MAC Address Table

The 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 32 K 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.

Address Threshold

The address threshold is the value at which the system reports the total number of addresses that are known. Specifically, when this threshold is reached, the system generates the SNMP trap *addressThresholdEvent*.

The range of values that you can enter for this parameter is between 1 and 1 plus the maximum address table size (32 K). Setting the address threshold to one greater than the address table size prevents the system from generating events because the limit can never be reached.

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.)

IP Fragmentation

Standard FDDI allows larger maximum packet sizes than standard Ethernet. FDDI stations that transmit IP packet sizes larger than approximately 1500 bytes wish cannot communicate with stations on an Ethernet LAN. If the system receives such packets and they are destined for one or more Ethernet LANs, it filters them — except when IP fragmentation is enabled.

When you enable IP fragmentation, the system breaks up large FDDI packets into smaller packets before bridging them to Ethernet.

IPX SNAP Translation

IPX SNAP Translation allows an alternative method of translating IPX packets from Ethernet to FDDI and vice-versa.

- When IPX SNAP translation is enabled, any 802.3_RAW IPX packets that are forwarded from Ethernet to FDDI are translated to FDDI_SNAP. Likewise, SNAP IPX packets that are forwarded from FDDI to Ethernet are translated to 802.3_RAW packets.
- When IPX SNAP translation is disabled, the system uses standard IEEE 802.1H bridging to translate 802.3_RAW packets to FDDI_RAW packets.

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 multicast and broadcast packets that are received in that second of time are dropped. 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 K frames per second (approximately 1000 frames per second). To determine an appropriate limit, measure the normal amount of broadcast or multicast traffic on your network.
- If you have IP multicast application traffic on your network, be sure that any limits that you configure do not constrain these traffic flows.
- 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.
- There are similar options available through the Quality of Service menu. For more information, see Chapter 17.

GARP VLAN Registration Protocol (GVRP)

To activate GVRP on the system, you enable the GARP VLAN Registration Protocol (GVRP) first on the bridge and then on individual bridge ports.

On a port-by-port basis, GVRP allows the system to automatically learn the presence of and updates to 802.1Q VLANs. GVRP simplifies the management of IEEE 802.1Q VLAN configurations in large networks by making aspects of VLAN configuration dynamic.

GVRP maintains a database of VLAN member ports as the bridge learns about them. Specifically, GVRP tracks which ports are added to and removed from each VLAN and communicates this information to other GVRP-aware bridges. The bridges then determine active topologies for the network and for each VLAN using STP to prevent network loops.

GVRP operates only on ports that are in the STP forwarding state. If GVRP is enabled, a port that changes to the STP forwarding state automatically begins to participate in GVRP. A port that changes to an STP state other than forwarding no longer participates in GVRP. For more information about GVRP and VLANs, see Chapter 9.

Important Considerations

To use GVRP, consider the following:

- GVRP updates are not sent out to any blocked STP ports. GVRP operates only on ports that are in the STP forwarding state.
- GVRP is disabled by default on the bridge and on all bridge ports.
- Enabling GVRP determines whether the VLAN origin for a port-based VLAN is dynamic (GVRP enabled) or static (GVRP disabled).
- To maximize the effectiveness of GVRP, it should be supported in as many end stations and network devices as possible.
- Based on updates from GVRP-enabled devices, GVRP allows the system to dynamically create a port-based VLAN (unspecified protocol) with a specific VLAN ID and a specific port.
- On a port-by-port basis, GVRP allows the system to learn about GVRP updates to an existing port-based, protocol-based, or network-based VLAN with that VLAN ID and IEEE 802.1Q tagging.
- VLANs that are created dynamically with GVRP exist only as long as a GVRP-enabled device is sending updates — if the devices no longer send updates, or if GVRP is disabled, or if the system is rebooted, all dynamic VLANs are removed.

- GVRP manages the active topology, not nontopological data such as VLAN protocols. If a local bridge needs to classify and analyze packets by VLAN protocols, you must manually configure protocol-based VLANs and simply rely on GVRP to send VLAN ID updates. But if the local bridge needs to know only how to reach a given VLAN, then GVRP provides all necessary information.
- The VLAN topologies that GVRP learns are treated differently from VLANs that are statically configured. Although static updates are saved in nonvolatile RAM, GVRP's dynamic updates are not. When GVRP is disabled, the system deletes all VLAN interfaces that were learned through GVRP and leaves unchanged all VLANs that were configured through the Administration Console or through the Web management software.

Standards, Protocols, and Related Reading

For more information about bridging, STP, and GVRP consult the following standards:

- **IEEE 802.1D** — This standard specifies the requirements to which your system, as a transparent bridge, complies.
- **IEEE 802.1Q** — This standard defines GVRP, tagging, and the dynamic registration of VLANs.

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

`http://www.ieee.org`

8 TRUNKING

This chapter provides guidelines, limitations, and other important information about how to implement the trunking function for CoreBuilder® 3500 systems. 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 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 and allows you to combine multiple Fast Ethernet, Gigabit Ethernet, or FDDI ports into a single high-speed link between two switches (see Figure 25).

Figure 25 Example of a Trunk

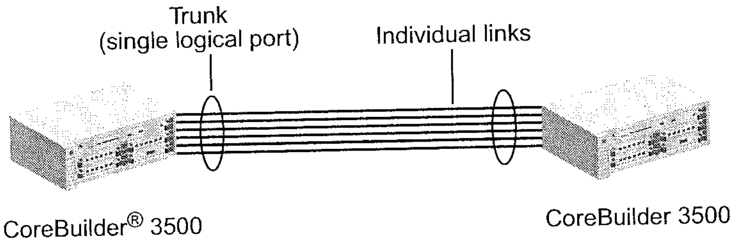

The system treats trunked ports in the same way that it treats individual ports. Also, all higher-level network functions — including Spanning Tree algorithms, VLANs, and Simple Network Management Protocol (SNMP) management — do not distinguish a trunk from any other network port.

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, Gigabit Ethernet, or Fiber Distributed Data Interface (FDDI) ports into a single high-speed channel.

If Gigabit Ethernet is not available, you can use trunked Fast Ethernet 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 26 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 virtual LAN (VLAN) information, but you cannot select bridge ports 3 or 4 since they are part of the trunk.

Figure 26 Bridge Port Numbering After Trunking

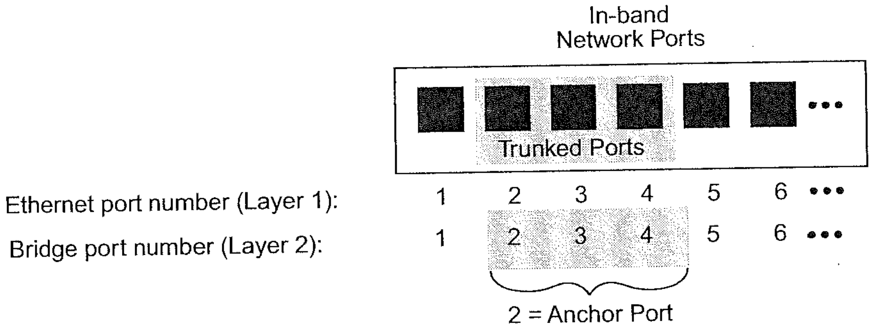

Regardless of whether you define trunking, the physical port numbering on your system remains the same.

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 9), 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.

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

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.

Key Guidelines for Implementation

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

General Guidelines

- Create trunks before you define VLANs.
- The system supports four point-to-point trunks, each built from up to eight 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 FDDI, Fast Ethernet, 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 scenarios:
  - 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 currently a proprietary implementation. No *de facto* standards currently exist.

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 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 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.

- 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 14 shows that given a choice, trunking Fast Ethernet may not be a cost-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 14 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 10 links) (1200 Mbps full duplex)

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, such as the FDDI station mode [dual attach station (DAS) or single attach station (SAS)].
- Trunk names can be no 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, FDDI, Fast Ethernet, 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 or full-duplex).
- Each Gigabit Ethernet module that you install takes up one of the switch's four trunk resources (but does not itself constitute a trunk). If you have two or more Gigabit Ethernet modules, you can trunk them together to free up switch trunk resources. For example, if you install three Gigabit Ethernet modules, the switch allows only one additional trunk. But if you trunk the Gigabit Ethernet modules, the switch supports three additional trunks after you reboot.

  If you add a Gigabit Ethernet module to a switch that has four trunks already defined, the module does not power up, and you receive an error message.

- When you create a 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, simply 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).
- If you upgrade from Version 1.1 and exceed four trunk channels, the Gigabit Ethernet port is not initialized and an error message is posted to the system log. When this situation occurs:
  - The Gigabit Ethernet port MIB returns a `config error` state.
  - The Gigabit Ethernet port is disabled.
  - The console displays `Configuration incompatible please check release notes.`
  - The system error LED lights.
- Doing 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 you that you are modifying.
- To avoid configuration errors, do not modify FDDI station mode port-pairs when any of the ports in the pair are members of a trunk.
- If you have more than one media type on your system (for example, Fast Ethernet 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.
- In an FDDI trunk:
  - You cannot modify FDDI station mode port pairs when any of the ports in the pair are in a trunk.
  - When you modify the station mode, any FDDI ports that are associated with VLANs or a trunk are removed from the VLAN or trunk.
- If you change an FDDI port pair from SAS to DAS, select the pair using only the lower of the two port numbers, as you do with a trunk anchor port.

- You cannot change some port characteristics within a trunk. For example, in an FDDI trunk, you cannot change a trunked DAS port to a SAS port.

  Here is an example of how to change the FDDI station mode of a trunk:

a Remove the desired trunk.

b Reboot and then change the station mode.

c Reboot and redefine the trunk (and any affected VLANs).

d Reboot.

To avoid configuration errors, do not modify FDDI-station mode port-pairs when any of the ports in the pair are members of a trunk.

Removing Trunks

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

Important Considerations

- 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.
- Because each Gigabit Ethernet module uses an internal trunk resource towards the system limit of four, keep in mind how many trunk resources may be used when you remove a trunk. For example, if your system has a trunk with two Gigabit Ethernet ports (which consolidates two trunk resources into one) plus three other trunks, and you then try to untrunk the two Gigabit Ethernet ports, you will exceed the trunk resource limit. The untrunked Gigabit Ethernet ports try to take over two separate trunk resources (for an illegal total of 5), and the system sends a warning message like the following:

```
Unable to remove trunk(s).
Internal trunk resource limit would be exceeded.
```

Standards, Protocols, and Related Reading

The system supports these Ethernet standards:

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

3Com trunking technology interoperates with similar technology from other vendors, including Sun Microsystems and Cisco Systems.

9 VIRTUAL LANs

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

This chapter covers the following topics:

- VLAN Overview
- Key Concepts
- Key Guidelines for Implementation
- VLAN allOpen or allClosed Mode
- Ignore STP Mode
- Port-based VLANs
  - The Default VLAN
  - Static Port-based VLANs
  - Dynamic Port-based VLANs Using GVRP
- Protocol-based VLANs
- Network-based IP VLANs
- Rules of VLAN Operation
- Modifying and Removing VLANs
- Monitoring VLAN Statistics

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 end 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 or a Layer 3 network.

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 subnet 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 is 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). These mechanisms work 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 network administrators 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, thereby 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.
- Help avoid flooding and minimize broadcast and multicast traffic.
- Reduce the need for routing to achieve higher network performance, ease of administration, and reduced costs.
- Control communication among broadcast domains.

Features

Your 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. Using allClosed mode also enables you to use another VLAN feature called Ignore STP mode. The chosen VLAN mode dictates the requirements for the port-based, protocol-based, and network-based VLANs. See "Terminology" for more information about the VLAN modes and Ignore STP Mode.
- **Configurable types of VLANs** — The system allows you to configure different types of VLANs for controlling the flow of traffic through a network:
  - **Port-based VLAN** — Determines VLAN membership using a group of ports. By default, your system provides a special port-based VLAN that contains all ports without tagging. This special VLAN is called the *default VLAN*. It always uses the VLAN ID of 1, the name `Default`, and the protocol type unspecified. See "The Default VLAN" later in this chapter for more information.

    The system also supports both *static* and *dynamic* port-based VLAN configuration if you choose to set it up that way. See "Static Port-based VLANs" and "Dynamic Port-based VLANs Using GVRP" later in this chapter for more information.
  - **Protocol-based VLAN** — Determines VLAN membership using a group of ports that share one or more protocol types. In addition to the user-defined protocol-based VLANs, the system supports a special type of protocol-based VLAN called a *router port IP VLAN*. This type of VLAN, which the system generates when you define an IP interface as a router port IP interface, requires allClosed mode. See "VLANs Created by Router Port IP Interfaces" later in this chapter for more information.
  - **Network-based VLAN** — Determines IP VLAN membership for a group of ports that are configured for IP and a specific network address.

VLAN hierarchy

The VLAN type classification is hierarchical: a protocol-based VLAN is a special type of port-based VLAN, and a network-based VLAN is a special type of an IP protocol-based VLAN. This hierarchy allows you to use a combination of VLAN types to group users and traffic types.

You can either configure network-based IP VLANs (IP VLANs with unique Layer 3 IP addresses) or you can define a single VLAN with the protocol type IP and then define multiple IP routing interfaces for that single IP VLAN. See Chapter 11 for more information about defining VLAN-based routing interfaces.

- **Per-port IEEE 802.1Q tagging** — Selecting IEEE 802.1 tagging on a per-port basis dictates that frames be encapsulated and tagged as specified in the IEEE 802.1Q standard. See "Port-based VLANs","Protocol-based VLANs", and "Network-based IP VLANs" later in this chapter for specific information on tagging for the types of VLANs.

Key Concepts

Before you configure VLANs, review the following key concepts.

Related Standards and Protocols

The following standards and protocols apply to the VLANs that you can configure on your system:

- **IEEE 802.1Q** — A proposed standard for VLANs, it is aimed at:
  - Defining an architecture to logically partition bridged LANs and provide services to defined user groups, independent of physical location.
  - Allowing interoperability among multivendor equipment.

  IEEE 802.1Q defines the bridging rules for VLANs, that is, ingress and egress rules, as defined in "Key Concepts" (and described in detail in "Rules of VLAN Operation" later in this chapter).

  The standard also specifies a tag format that embeds explicit VLAN membership information in a 12-bit VLAN ID (VID) that provides 4094 possible VLANs. (Standard IEEE 802.1p uses this same frame format, but also takes advantage of an additional 3 bits for specifying the priority levels used for class of service differentiation.)
- **Generic Attribute Registration Protocol (GARP)** — This protocol is defined in IEEE 802.1p, which is a supplement to the IEEE 802.1D standard. GARP is a Layer 2 transport mechanism that allows switches and end systems to propagate information across the switching domain.
- **GARP VLAN Registration Protocol (GVRP)** — This protocol, which is defined in IEEE 802.1Q, defines dynamic registration of VLANs that use IEEE 802.1Q tagging (the VLAN ID).

VLAN IDs

Each VLAN is identified by its VLAN ID (VID). For VLANs that you create, the system keeps track of its used VLAN ID numbers to help you select the next available VLAN ID. Outgoing data frames are tagged per IEEE 802.1Q (which specifies the VID) if tagging is enabled on the transmit port for that VLAN. Tagged IEEE 802.1Q data frames that are received on the system are assigned to the VLAN that corresponds to both the VID contained in the tag and the protocol type.

Be aware of these additional guidelines:

- The default VLAN always uses the reserved VID of 1.
- Before you assign a VID, review the information in Table 15.

Table 15 Assigning ID Numbers to VLANs

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

- If you rely on dynamic configuration to create a port-based VLAN based on GVRP updates, the VID is the unique IEEE 802.1Q VID.

When you define a router port IP interface, the system automatically creates a router port IP VLAN and assigns it the next available VID. See Chapter 11 for information on router port IP interfaces.

Terminology

The following terms apply to VLANs:

- **Default VLAN** — The predefined port-based VLAN interface on your system that always uses VID 1, the protocol type unspecified, and the name Default. The default VLAN also initially includes all of the bridge ports without any tagging, but you can modify the bridge ports and tag status of the default VLAN. If you maintain the default VLAN and you install a new module, the system adds all ports that are associated with the new module to the default VLAN. See "The Default VLAN" for more detailed information.
- **VLAN origin** — Whether the VLAN was created in one of the following ways:
  - **Statically** — The VLAN display shows an origin of static if you define the VLAN.
  - **Dynamically** — The VLAN display shows an origin of GVRP if the system learned the VLAN dynamically through GVRP.
  - **Router** — The VLAN display shows an origin of router if you have defined a router port IP interface on a single bridge port. When you define a router port IP interface, you must place the system in allClosed mode. This removes any allOpen VLANs and re-creates the default VLAN. See Chapter 11 for more information on defining router port IP interfaces.
- **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 (although data can still be *routed* between VLANs). The default VLAN mode is allOpen. See "VLAN allOpen or allClosed Mode" for more information.
- **Ignore STP mode** — A per-VLAN mode that determines whether the system ignores the blocking Spanning Tree Protocol (STP) mode for the ports of a designated VLAN. (One instance of STP runs on the system, but you can disable it for each VLAN.) Ignore STP mode is only available in allClosed mode; it is disabled by default. It allows the user to select (for each VLAN) which VLANs ignore STP blocked ports. This mode is typically used for VLANs that have router interfaces that choose to ignore the STP state. It allows routing (or bridging) over a port that is blocked by STP. See "Ignore STP Mode" later in this chapter for more information.

- **Protocol suite** — The protocol family that is associated with a protocol-based VLAN. Protocol-based VLANs can be associated with one or more protocol suites. The protocol suite is unspecified for the default VLAN and all port-based VLANs.
- **Layer 3 address** — The network or subnetwork address that is associated with a network-based IP VLAN.
- **Tagging type** — On a per-port basis, whether there is explicit VLAN membership information (the IEEE 802.1Q header and the VLAN ID or VID) in each frame. You can specify no tagging or IEEE 802.1Q tagging.
- **Port membership** — The bridge ports that you assign to be part of the VLAN.

If you have created trunks, you must specify the anchor port (the lowest-numbered port) port in the trunk when you define the VLAN interface. All bridge ports are initially part of the default VLAN.

- **VLAN name** — The 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.
- **Dynamic VLAN configuration** — Using the GARP VLAN Registration Protocol (GVRP), this configuration enables dynamic VLAN configuration of port-based VLANs and dynamic updates of IEEE 802.1Q tagged port-based VLANs.
- **Ingress and egress rules** — *Ingress rules* determine the VLAN to which an incoming frame belongs. If it 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, see "Rules of VLAN Operation" later in this chapter.

Key Guidelines for Implementation

This section provides a series of guidelines to consider when you use VLANs. The guidelines are organized as follows:

- Network-based VLANs vs. multiple interfaces per VLAN
- VLANs created by router port IP interfaces
- Number of VLANs
- General guidelines

Network-based VLANs vs. Multiple Interfaces per VLAN

You can either configure network-based IP VLANs (IP VLANs with unique Layer 3 IP addresses) or you can define a single VLAN with the protocol type IP and then define multiple IP routing interfaces for that single protocol-based VLAN (an IP VLAN).

If you decide to convert an existing network-based VLAN to a protocol-based VLAN that has multiple interfaces associated with it, use the following procedure:

1. Remove one or more network-based VLANs.
2. Define an IP VLAN or a VLAN that supports IP as one of its protocols.
3. Define multiple IP interfaces (with different IP addresses) to use that IP VLAN.

You can define up to 32 IP interfaces on the system, including IP routing interfaces for static VLANs, router port IP VLANs, or any combination of static VLANs and router port IP VLANs.

If you define multiple interfaces for an IP VLAN, you cannot subsequently modify that IP VLAN to supply Layer 3 address information. If only one routing interface is defined for the IP VLAN, then you can supply Layer 3 address information as long as it matches the Layer 3 information that is specified for the routing interface.

If you use network-based VLANs, you are limited to defining only *one* IP routing interface for that VLAN. When you define an IP routing interface with the interface type *vlan*, the system does not allow you to select a network-based IP VLAN that already has a routing interface defined for it. For more information on IP routing interfaces, see Chapter 11.

VLANs Created by Router Port IP Interfaces

By default, your system uses a routing over bridging model, in which any frame is bridged before it is potentially routed. If you want to define IP routing interfaces that use a routing versus bridging model, however, you can bypass your static VLAN configuration and instead go directly to defining an IP interface on a single router port (a router port IP interface). That process is described in this section.

If you define a router port IP interface, note the following information:

- Defining an IP interface for a router port requires the interface type port. Defining an IP interface for a configured IP VLAN requires you to specify the interface type vlan.
- The IP interface definition procedure for a router port requires that you place the system in allClosed mode. The allClosed mode prevents MAC addresses from being shared between the router port IP VLAN and any other VLANs and enables the router port to ignore Spanning Tree states on the port.
- Once you define the router port IP interface and change the VLAN mode to allClosed, the following events occur:
  - The system deletes all other VLANs and redefines the default VLAN. You must redefine any VLANs that you had configured, keeping in mind that unicast traffic will no longer be forwarded between VLANs. You must define routing interfaces to allow forwarding between VLANs. Also, you cannot specify the bridge port owned by the router port IP interface in any VLAN that you configure or modify.
  - The system creates a special protocol-based VLAN called a router port IP VLAN and assigns to it the next available VID. The VLAN displays identify the origin of a router port IP VLAN as router, as well as the port that is owned by the router port IP interface. You cannot modify or remove a router port IP VLAN, nor can you change its Ignore STP mode (which is always enabled).
- To disable bridging entirely for the router port, remove that port from the default VLAN.

For more information on defining a router port IP interface, see Chapter 11.

Number of VLANs

Your system supports a maximum of 64 VLANs based on a physical limit of 125 VLAN table entries. To determine the number of VLANs of any type that you can have on the system, use the following equation:

Number of VLANs supported =
(125 divided by the number of protocol suites) minus 3

Important Considerations

- When you use the VLAN equation to calculate the number of VLANs that you can have on your system, keep in mind that the equation provides an estimate. Your system may allow additional or fewer VLANs, depending on your configuration, use of protocol suites, and chosen tag style. If, for example, you are using the Release 3.0 VLAN tag style of all ports, this formula generally yields a maximum number of VLANs. If you use the Release 1.2 tag style of taggedVlanPorts, then this formula generally yields a minimum number of VLANs.
- The number of allowable VLANs includes the default VLAN.

Determining the Number of Protocol Suites

To perform the calculation, first determine the total number of protocol suites used on your system. Use the following guidelines:

- IP counts as one protocol suite for IP VLANs.
- AppleTalk counts as one protocol suite for AppleTalk VLANs.
- Generic IPX, which uses all four IPX types, counts as four protocol suites. (Each IPX type alone counts as one.)
- DECnet counts as one protocol suite for DECnet VLANs.
- The unspecified type of protocol suite counts as one for the default VLAN or port-based VLANs. (Even if you have *only* the unspecified protocol suite on the system, the limit is still 64 VLANs.)
- If you are using GVRP (for dynamic port-based VLANs), use the type unspecified in the VLAN formula

Remember to include the unspecified type for the default VLAN, even if you have removed the default VLAN and do not have another VLAN defined with the unspecified protocol type.

In addition to the limit on the number of VLANs, you are limited to 15 different protocols that can be implemented by the protocol suites on the system. See Table 19 later in this chapter for a list of supported protocol suites and the number of protocols within each suite.

VLAN Equation Examples

Example 1 You have 7 protocol suites on the system (IP, AppleTalk, unspecified for the default VLAN, and generic IPX, which counts as 4 protocol suites):

(125 / 7) – 3 = 14

In this configuration, the system supports a minimum of 14 VLANs. Per Table 19, these 7 protocol suites use 10 protocols: 3 IP, 2 AppleTalk, 1 unspecified, and 4 generic IPX.

Example 2 You have 5 protocol suites on the system (IP, unspecified, AppleTalk, IPX 802.2 Sub-Network Access Protocol [SNAP], and IPX 802.3 Raw):

(125 / 5) – 3 = 22

In this configuration, the system supports a minimum of 22 VLANs. Per Table 19, these 5 protocol suites use 7 protocols: 3 IP, 1 unspecified, 2 AppleTalk, 1 IPX 802.2 SNAP, and 0 IPX 802.3 Raw, because it does not use an Ethernet protocol type.

If you are upgrading your system from Release 1.2 and the VLAN resource limit is reached during a power up with a serial port console connection, use the Administration Console option `bridge vlan vlanAwareMode` *to change the VLAN aware mode to taggedVlanPorts. See "VLAN Aware Mode" later in this chapter for more information.*

General Guidelines

- The VLAN mode of allOpen or allClosed applies to *all* VLANs associated with the system (static, dynamic, or router port). Configure the VLAN mode *before* you define any static VLANs. (As part of the configuration procedures for a router port IP interface, you must place the system in allClosed mode; see Chapter 11.)

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

- If you configure the system for allClosed mode, you can enable Ignore STP mode on any VLAN. You can also disable STP on a any port for either allOpen or allClosed mode by using a bridge port option. (Use `bridge port stpState` on the Administration Console.) See Chapter 7 for bridging information. Also see "Ignore STP Mode" later in this chapter.
- To take advantage of GVRP for dynamic configuration or dynamic updates of port-based VLANs, verify that GVRP is enabled as both a bridge-wide and a bridge-port parameter. See Chapter 9 for information about bridging parameters. See "Dynamic Port-based VLANs Using GVRP" for information about GVRP.
- You can configure overlapping VLANs if they have some distinguishing characteristic. For example, a bridge port can be shared by multiple VLANs as long as the shared port has a distinguishing characteristic for the shared port, such as protocol type or tagging type. In allClosed mode, you must tag overlapped ports of any network-based VLANs. See "Network-based IP VLANs" later in this chapter.
- Per-port tagging requirements depend on whether the hosts connected to the port are configured for IEEE 802.1Q tagging. Per-port tagging is also required to differentiate between overlapped ports of the same protocol type and between overlapped IP Layer 3 VLANs in allClosed mode.
- Consider maintaining the system's default VLAN. The default VLAN preserves the flooding of unspecified traffic, since it initially contains all of the system's bridge ports, with unspecified protocol information and no tagging.

- To establish routing between static VLANs and configure a VLAN interface to support one or more routing protocols, configure the VLAN for the protocols *before* you configure a routing interface. For protocols other than IP, the system does not define the routing interface for a protocol if a VLAN for that protocol does not exist.

  If you define an IP interface and specify vlan as the interface type, the system does not define the IP routing interface unless you have an IP VLAN configured. See the appropriate routing chapter for an overview of your routing options and guidelines. See Chapter 11 for information on defining either an IP router interface (for a static IP VLAN) or a router port IP interface.
- If you plan to use trunks, define the appropriate 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 the VLAN and places them in the default VLAN.) See "Trunking and the Default VLAN" for more information. 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.
- 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 the system transmits the frame, it 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, the frame is transmitted untagged.

VLAN allOpen or allClosed Mode

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

3Com's use of the term "allOpen" is equivalent to the IEEE Standard 802.1Q term "Shared VLAN Learning" (SVL). The term "allClosed" is equivalent to the IEEE 802.1Q term "Independent VLAN Learning" (IVL). 3Com imposes the restriction that you must choose 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 (VLANs with an origin of static).
- As part of the configuration procedures for a router port IP interface, you must place the system in allClosed mode. Once you define a router port IP interface (and the system creates the router port VLAN), you cannot change the VLAN mode until you delete the router port IP interface.
- Select the VLAN mode as follows:
  - **allOpen** — Use this less restrictive mode if you have no security issues about the forwarding of data between VLANs. The allOpen mode is the default VLAN mode for all VLANs that you create. It permits data with a unicast MAC address to be forwarded between VLANs. For example, data received on IP VLAN 2 with a destination of IP VLAN 3 is forwarded there.

    The allOpen mode implies that the system uses a single bridge address table for all of the VLANs on the system (the default configuration).
  - **allClosed** — Use this more restrictive mode if you are concerned about security between VLANs. Data cannot be forwarded between VLANs (although data can still be routed between VLANs). The allClosed mode implies that each VLAN that you create has its own address table. Router port IP interfaces require allClosed mode.

- If you are using allClosed mode and STP on the system (with multiple routes to a destination), you can also specify a mode called *Ignore STP mode* to disable STP blocking for a specified static VLAN. (Although each VLAN has its own address table, there can be only one instance of STP on the system.) See "Ignore STP Mode" for information on this mode. To disable STP blocking on any port with allOpen or allClosed VLANs, use the `bridge port stpState` option on the Administration Console. See Chapter 7 for bridging information.
- Your choice of the VLAN mode affects how you manipulate bridge port addresses (via the 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 addresses without specifying a VLAN interface index.

Modifying the VLAN Mode

To change your VLAN mode, perform these procedures:

1. Delete all routing interfaces (including router port IP interfaces) that you have configured on the system. You cannot change the mode if you have router interfaces defined on the system.
2. Using your configuration tool (for example, the Administration Console or the Web Management applications), 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.
3. Reconfigure your VLANs and redefine your routing interfaces.

For the specific commands for these procedures, see the *Command Reference Guide*.

Mode Requirements

Table 16 shows the requirements for defining static VLANs in allOpen and allClosed mode.

Table 16 Mode Requirements for Static VLANs

Type of Static VLAN	Requirements
Port-based	For *nonoverlapped* port-based VLANs: ■ Protocol type: unspecified ■ Separate member ports. That is, each port-based VLAN owns a different set of ports. For *overlapped* port-based VLANs: ■ Protocol type: unspecified ■ IEEE 802.1Q tagging for shared ports. That is, the shared ports can employ a tagging mode of none in only one VLAN; shared ports in all other VLANs must use IEEE 802.1Q tagging.
Protocol-based	For *nonoverlapped* protocol-based VLANs: ■ Either the protocol type or the member ports are unique per VLAN For *overlapped* protocol-based VLANs (multiple VLANs of the same protocol type that share ports): ■ IEEE 802.1Q tagging for shared ports. That is, the shared ports can employ a tagging mode of none in only one of the same protocol type VLANs; shared ports in all other VLANs of the same protocol type must use IEEE 802.1Q tagging.
Network-based (IP VLAN only)	■ A Layer 3 address that is unique per network-based VLAN ■ For *allOpen mode*, no tagging restrictions on the shared ports ■ For *allClosed mode*, IEEE 802.1Q tagging for shared ports. That is, the shared ports can employ a tagging mode of none in only one of the network-based VLANs; shared ports in all other network-based VLANs must use IEEE 802.1Q tagging.

Ignore STP Mode

When you use allClosed VLAN mode on your system, you can enable the system to ignore the Spanning Tree Protocol (STP) mode on a per-VLAN basis, that is, to ignore STP blocked ports for static protocol-based VLANs associated with routing interfaces. (When STP detects multiple paths to a destination, it blocks all but one of the paths.)

If you have configured router port IP interfaces on your system (so that the system generates router port VLANs owned by the router IP interfaces), ignore STP mode is automatically enabled and you cannot disable it.

Important Considerations

- Ignore STP mode is disabled by default for static VLANs.
- You can use this mode *only* when the system is in allClosed mode.
- Ignore STP mode is useful when you have redundant router connections between systems that have STP enabled. In this situation, if you want to create multiple VLANs and use one VLAN for routing, you can configure your system to ignore the STP blocking mode for that VLAN. This setting avoids disruptions to routing connectivity based on the STP state.
- To disable STP blocking on a *per-port* basis with allOpen or allClosed VLANs, use the bridging option (`bridge port stpState` on the Administration Console). See Chapter 7 for bridging information.

Ignore STP mode affects bridging as well as routing. If you have STP enabled on the system and you have redundant bridged paths between systems with different VLANs, STP blocks one of the paths unless you enable Ignore STP mode. See Figure 28 later in this chapter for an example of redundant paths between systems that have different port-based VLANs.

Example of Ignore STP Mode

Figure 27 shows two paths available if a workstation associated with IP VLAN E wants to communicate with a server associated with IP VLAN D. STP blocks the routed as well as bridged traffic for the one path unless you enable Ignore STP Mode for the routed IP VLANs. With the blocking removed for IP routed traffic, the best path is used.

Figure 27 Using Ignore STP Mode

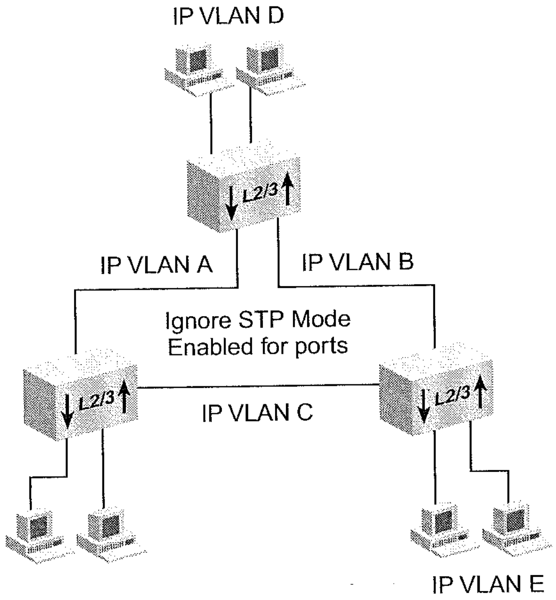

VLAN Aware Mode

VLAN aware mode accommodates the difference in VLAN resource usage as well as tagged-frame ingress rules between Release 1.2 and Release 3.0 of the system software. For more information on ingress rules, see "Rules of VLAN Operation" later in this chapter. (The Release 1.2 ingress rules in allOpen mode mandated that incoming tagged frames assigned to one of the configured VLANs if the VID of the frame matched that of the VLAN *and* if a port in that VLAN were tagged.)

The VLAN aware mode, which you set with the Administration Console option `bridge vlan vlanAwareMode`, reflects the difference in VLAN resource usage and modes of tagging as follows:

- At Release 1.2, all bridge ports were *not* VLAN aware (tagging aware) unless they were assigned to a VLAN that has one or more tagged ports.
- At Release 2.0 and later, all bridge ports become VLAN aware after a software update or after an NV data reset and do not have to be explicitly tagged in order to forward tagged frames.

This difference in resource usage and modes of tagging has the following impact: After you upgrade the system from 1.2 to 3.0, the release uses VLAN resources differently than it did at Release 1.2 and may cause a change in the total number of allowable VLANs.

VLAN aware mode is currently supported only through the Administration Console, not through Web Management or SNMP.

Initial installation of Release 3.0 provides a default VLAN aware mode of allPorts, which is consistent with the Release 3.0 ingress rules and resource allocation.

If you upgrade your system from Release 1.2 to a later release and the VLAN resource limit is reached during a power up with a serial port console connection, the system displays an error message similar to the following one to identify the index of the VLAN that it was unable to create:

```
Could not create VLAN xx - Internal resource threshold
exceeded
```

In this situation, the system removes all bridge ports from the VLAN that it could not restore from nonvolatile (NV) data, although it does maintain the previously stored NV data. To restore your VLANs after you see the resource error message, use the `bridge vlan vlanAwareMode` option and then set the VLAN aware mode to taggedVlanPorts. If VLANs are already defined, the system prompts you to reboot to put the new mode into effect.

If you do not see the VLAN internal resource error message, maintain the default VLAN aware mode of allPorts. In this case, the system can accommodate the number of Release 1.2 VLANs, but it now uses different ingress rules for tagged frames.

The Administration Console options `bridge vlan summary` and `bridge vlan detail` display the current VLAN aware mode after the VLAN mode (allOpen or allClosed).

Port-based VLANs

Port-based VLANs logically group together one or more bridge ports on the system and use the generic protocol type unspecified. Each arbitrary collection of bridge ports is designated as a *VLAN interface.* This 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
- Static port-based VLANs that you create
- Dynamic port-based VLANs created using GVRP

An alternative to port-based VLANs is packet filtering using port groups, as described in Chapter 10.

The Default VLAN

The system predefines a port-based VLAN to initially include all of the system's bridge ports without any tagging. For example, if you have four 10/100 Ethernet modules (24 bridge ports) installed on your system, the default VLAN initially contains all 24 ports.

The default VLAN always uses the VID of 1, the name Default, and the protocol type unspecified. No other VLAN than the default VLAN can use a VID of 1.

The default VLAN is the flood domain in either of these cases:

- 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 receiving the data.

See "Rules of VLAN Operation" later in this chapter.

Modifying the Default VLAN

The default VLAN is always associated with the VID of 1, the unspecified protocol type, and the name `Default`. Initially, the default VLAN is also associated with all ports and no tagging. Keeping the default VLAN intact ensures that the system accommodates the addition of a module by automatically adding the new module's bridge ports to the default VLAN. If necessary, the system also renumbers its ports when you add the module.

If necessary, you can modify (or remove) the default VLAN on the system. For example, you may want to modify the default VLAN to remove certain ports. Such a change does not prevent the system from adding a new module's bridge ports to the default VLAN.

However, the following changes *do* prevent the system from adding a new module's bridge ports to the default VLAN:

- If you modify the default VLAN to remove all ports
- If you remove the default VLAN completely. Even if you subsequently redefine the default VLAN, the system will not add bridge ports to the newly defined default VLAN.
- If you modify the default VLAN to tag a port

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 do not have 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 affect the default VLAN in the following ways:

- If you have only the default VLAN with all ports and you define a trunk (or subsequently remove a trunk), the ports 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 subsequently define a trunk for ports in one of the other VLANs, the system removes those ports from that other 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 ipvlan1: ports 5-11	Define a trunk with ports 7,8	default VLAN: ports 1-4, 7-8 ipvlan1: ports 5-6, 9-11

- If you have the default VLAN as well as other VLANs and you subsequently modify an existing trunk that has ports in one of the VLANs, any port 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 ipvlan1: 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 ipvlan1: ports 6-11 (port 6 is new anchor port)

Trunking with the default VLAN removed

If you remove the default VLAN, the system has nowhere to return ports 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
ipvlan1: ports 1-11	Define trunk with ports 5-8	ipvlan1: ports 1-4, 9-11

- If you have VLANs (but no default VLAN) and you 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
ipvlan1: ports 1-11 (ports 5-8 are trunk ports)	Modify existing trunk to have ports 6-8 (remove port 5, the anchor port)	ipvlan1: ports 1-4, 6-11 (port 6 is new anchor port)

See Chapter 8 for information on using trunks.

Static Port-based VLANs

You can explicitly configure port-based VLAN interfaces instead of relying on GVRP to dynamically create port-based VLAN interfaces.

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 or unless you have DAS ports defined on an FDDI module.
- 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, not the other bridge ports in the trunk, is selectable when you are creating VLANs. For more information, see Chapter 8.

- If you define FDDI DAS ports, select the lowest-numbered port in the DAS pair when you define the ports in the VLAN. The higher-numbered port in the DAS pair is not selectable. See Chapter 6.
- 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 IEEE 802.1Q tagging for each VLAN.
- Port-based VLANs use the protocol type unspecified.
- To define a port-based VLAN interface, specify this information:
  - A VID in the range 2 through 4094, 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. For FDDI DAS ports, specify the lowest-numbered port in the DAS pair.
  - The protocol type unspecified.
  - Tag status (none or IEEE 802.1Q).
  - The unique name of the VLAN interface.

Example 1: Nonoverlapped VLANs

Figure 28 shows two systems that have nonoverlapping port-based VLANs and no port tagging. Ports 1 through 4 on Device1 make up the VLAN called unspecA, while ports 5 through 8 make up unspecB. 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 unspecA are assigned to unspecA, regardless of the data contained in the frames. 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 versus allClosed VLANs. For example, for allClosed VLANs, if a frame is received on a port in unspecA with a destination address that is known in the address table of unspecB, the frame is flooded throughout unspecA because it has an unknown address for unspecA. For allOpen VLANs, there is one address table; therefore; the frame is forwarded to the port that corresponds to the known destination address. However, if the transmit port is not a member port of unspecA, the frame is transmitted untagged, regardless of that port's tag status on unspecB.

In Figure 28, if STP is enabled, STP blocks one of the paths unless you enable Ignore STP mode. See "Ignore STP Mode" earlier in this chapter for more information.

Figure 28 Port-based VLANs Without Overlapped Ports

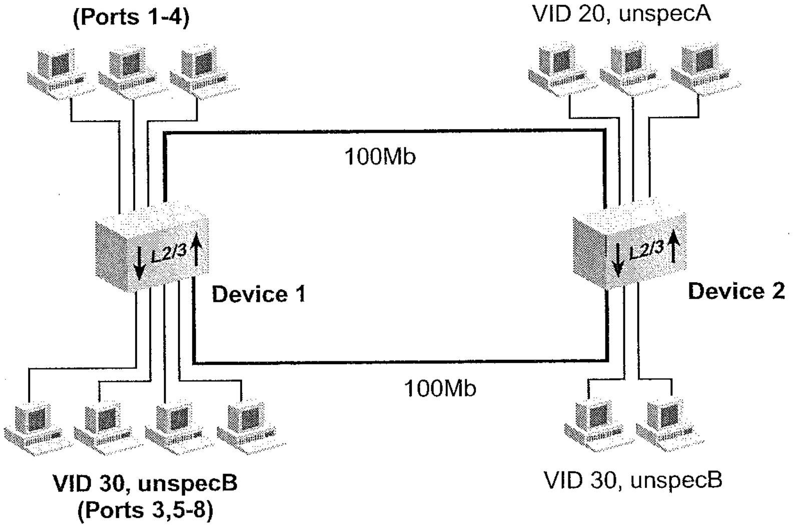

Table 17 shows the information that can be used to configure these VLANs *without* overlapped ports on Device 1 (the device on the left):

Table 17 Port-based VLAN Definitions Without Overlapped Ports for Device 1

unspecA	**unspecB**
VLAN Index 2	VLAN Index 3
VID 10	VID 15
Bridge ports 1-4	Bridge ports *5-8*
Protocol type unspecified	Protocol type unspecified
Per-port tagging: ■ Ports 1-4 — *none*	Per-port tagging: ■ Ports 5-8 — *none*
VLAN name unspecA	VLAN name unspecB

Example 2: Overlapped VLANs

Figure 29 shows port-based VLANs that *overlap* on bridge port *3*.

Figure 29 Port-based VLANs with Overlapped Ports

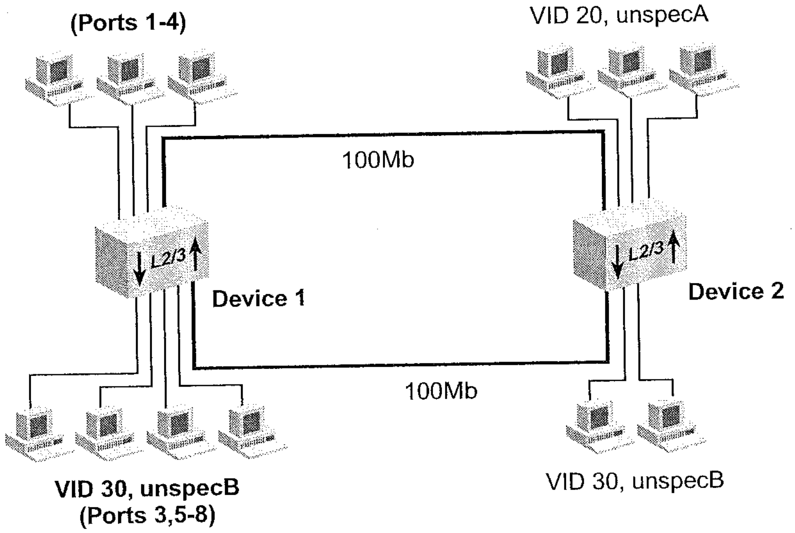

Table 18 shows the information that you use to configure these VLANs *with* overlapped ports on Device 1:

Table 18 Port-based VLAN Definitions with Overlapped Ports for Device 1

unspecA	unspecB
VLAN Index 2	VLAN Index 3
VID 20	VID 30
Bridge ports *1-4*	Bridge ports *3, 5-8*
Protocol type unspecified	Protocol type unspecified
Per-port tagging: Ports 1-4 — *none*	Per-port tagging: ■ Port 3 — *IEEE 802.1Q* ■ Port 5 — *IEEE 802.1Q* ■ Ports 6-8 — *none*
VLAN name unspecA	VLAN name unspecB

If you plan for your VLAN to include trunk ports, specify the anchor port (lowest-numbered port) associated with the trunk. For example, if ports 5 through 8 in unspecB were associated with a trunk, you specify only bridge port 5 to define the VLAN to include all of the physical ports in the trunk (ports 5 through 8). The IEEE 802.1 Q tagging applies to all ports in the trunk.

Dynamic Port-based VLANs Using GVRP

GARP VLAN Registration Protocol (GVRP) can help you simplify the management of VLAN configurations in your larger networks.

GVRP allows the system to:

- Dynamically create a port-based VLAN (unspecified protocol) with a specific VID and a specific port, based on updates from GVRP-enabled devices.
- Learn, on a port-by-port basis, about GVRP updates to an existing port-based VLAN with that VID and IEEE 802.1Q tagging.
- Send dynamic GVRP updates about its existing port-based VLANs.

GVRP enables your system to advertise its manually configured IEEE 802.1Q VLANs to other devices supporting GVRP. Because the VLANs are advertised, GVRP-aware devices in the core of the network need no manual configuration to pass IEEE 802.1Q frames to the proper destination. The method of VLAN advertisement used by all GVRP-capable switches involves protocol data units (PDUs), similar to the method used by STP. GVRP-capable devices send their updates to a well-known multicast address to which all GVRP-capable devices listen for information changes.

Enabling GVRP lets the system dynamically adjust active network topologies in response to configuration changes in one or more VLANs. GVRP then advertises VLAN changes on each bridge to all other GVRP bridges in the network.

Important Considerations

To use GVRP, consider the following:

- To take advantage of dynamic IEEE 802.1Q VLAN configuration, enable GVRP as an entire bridge state and then as an individual bridge port state for the appropriate ports. See Chapter 7. By default, GVRP is disabled as both a bridge state and a bridge port state. If GVRP is enabled, the VLAN origin for a port-based VLAN is dynamic (with GVRP). When GVRP is disabled, the VLAN origin is either static (traditional static VLAN without GVRP) or router (router port).
- In a GVRP environment, devices must be GVRP-enabled (that is, support GVRP). These devices may be end stations with 3Com's Dynamic*Access*® software or other switches that explicitly enable GVRP.
- VLANs created dynamically with GVRP exist only as long as a GVRP-enabled device is sending updates. If the devices no longer send updates, or GVRP is disabled, or the system is rebooted, all dynamic VLANs are removed.
- GVRP updates are not sent out on any blocked STP ports. GVRP operates only on ports that are in the STP forwarding state. If GVRP is enabled, a port that changes to the STP forwarding state begins to participate in GVRP. A port that changes to an STP state other than forwarding no longer participates in GVRP.

- The VLAN topologies that GVRP learns are treated differently from VLANs that are statically configured. GVRP's dynamic updates are not saved in NVRAM, while static updates are saved in NVRAM. When GVRP is disabled, the system deletes *all* VLAN interfaces that were learned through GVRP and leaves unchanged all VLANs that were configured through the Administration Console, SNMP, or the Web Management software.
- GVRP manages the active topology, not nontopological data such as VLAN protocols. If you need to classify and analyze packets by VLAN protocols, you manually configure protocol-based VLANs. But if the system needs to know only how to reach a given VLAN, then GVRP provides all necessary information.
- A GVRP-created VLAN is useful in situations where only Layer 2 switching needs to be performed for that VLAN. (Routing between a GVRP-created VLAN and another VLAN can be performed with an external router.) Because GVRP-created VLANs are assigned the unspecified protocol type, router interfaces cannot be assigned to them. Therefore, all communication within a GVRP-created VLAN is constrained to that VLAN in allClosed mode; in allOpen mode, only unicast frames with a known destination address can be transmitted to another VLAN.

Example: GVRP

Figure 30 shows how a GVRP update (with the VID) sent from one end station is propagated throughout the network.

Figure 30 Sample Configuration Using GVRP

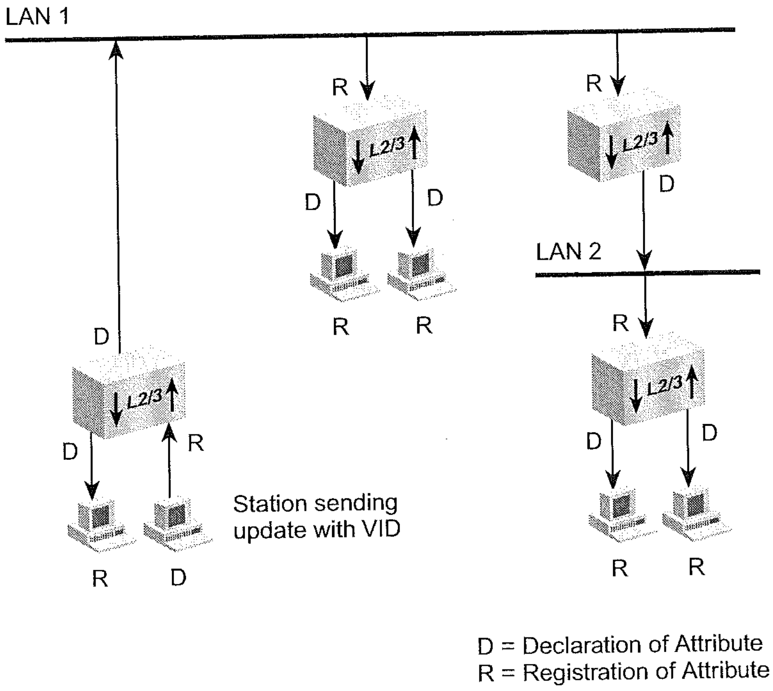
LAN 1
R
L2/3
D
D
R
R
R
L2/3
D
LAN 2
R
L2/3
D
D
R
R
D
L2/3
D
R
Station sending update with VID
R
D
D = Declaration of Attribute
R = Registration of Attribute

Protocol-based VLANs

Protocol-based VLANs enable you to use protocol type and bridge ports as the distinguishing characteristics for your VLANs. When you select a protocol such as IP, you do so based on the guidelines in this section.

Important Considerations

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

- If you plan to use the VLAN for *bridging* purposes, select one or more protocols per VLAN. Select them one protocol at a time.
- If you plan to use the VLAN for *routing*, select one or more protocols per VLAN, one protocol at a time, and subsequently define a routing interface for each routable protocol that is associated with the VLAN.
- The system supports routing for three protocols: IP, IPX, and AppleTalk.
- To define a protocol-based VLAN interface, specify this information:
  - The VID of your choice (except 1 or any VID already assigned), or accept the next available VID.
  - The bridge ports that are part of the VLAN interface. (If you have trunk ports, specify the anchor port for the trunk.)
  - The protocol for the specified ports in the VLAN.
  - Tag status (none or IEEEE 802.1Q). IEEE 802.1Q tagging must be selected for ports that overlap on both port and protocol (for example, if two IPX VLANs overlap on port 3).
  - The name you want to assign to this VLAN interface.
- If you use IP as the protocol and also specify a Layer 3 address, the protocol-based VLAN becomes a *network-based VLAN*.

You can either configure network-based IP VLANs (IP VLANs with unique Layer 3 IP addresses) or you can define a single protocol-based VLAN with the protocol type IP and then define multiple IP routing interfaces for that VLAN. For more information on network-based VLANs, see "Network-based IP VLANs" later in this chapter. For more information about IP interfaces, see in Chapter 11.

Selecting a Protocol Suite

The protocol suite describes which protocol entities can comprise a protocol-based VLAN. For example, the system's VLANs support the IP protocol suite, which has three protocol entities (IP, ARP, and RARP).

Table 19 lists the protocol suites that the system supports, as well as the number of protocols that are associated with each protocol suite.

Table 19 Supported Protocol Suites for VLAN Configuration

Protocol Suite	Protocol Entities	Number of Protocol Suites (PVIDs)	Number of Protocols in Suite
IP	IP, ARP, RARP (Ethernet Version 2, SNAP PID)	1	3
Novell IPX	IPX supports these IPX types:	4	4
	■ IPX - type II (Ethernet Version 2)	1	1
	■ IPX - 802.2 LLC (DSAP/SSAP value 0xE0 hex)	1 1	0* 0*
	■ IPX - 802.3 Raw (DSAP/SSAP value 0xFF hex)	1	1
	■ IPX - 802.2-SNAP (DSAP/SSAP value 0xAA hex)		
AppleTalk	DDP, AARP (Ethernet Version 2, SNAP PID)	1	2
Xerox XNS	XNS IDP, XNS Address Translation, XNS Compatibility (Ethernet Version 2, SNAP PID)	1	3
DECnet	DEC MOP, DEC MOP Remote Console, DEC DecNet Phase IV, DEC LAT, DEC LAVC (Ethernet Version 2, SNAP PID)	1	5
SNA	SNA Services over Ethernet (Ethernet Version 2 and DSAP/SSAP values 0x04 and 0x05 hexadecimal)	2	1
Banyan VINES	Banyan (Ethernet Version 2, DSAP/SSAP value 0xBC hexadecimal, SNAP PID)	1	1
X25	X.25 Layer 3 (Ethernet Version 2)	1	1
NetBIOS	NetBIOS (DSAP/SSAP value 0xF0 hexadecimal)	1	0*
Default	Default (all protocol types)	1	1
(unspecified)	No protocol types		

* This protocol does not use an Ethernet protocol type.

The system imposes two important limits regarding the number of VLANs and the number of protocols:

- **Number of VLANs supported on the system** — To determine the minimum number of VLANs that the system can support, use the equation described in "Number of VLANs" earlier in this chapter. The system supports a maximum of 64 VLANs.
- **Maximum number of protocols** — Use the value 15 as the maximum number of protocols that can be implemented on the system. A protocol suite that is used in more than one VLAN is counted only once toward the maximum number of protocols. For example, the DECnet protocol suite uses 5 of the available 15 protocols, regardless of the number of VLANs that use DECnet.

Example: Protocol-based VLANs for Bridging

Figure 31 is an example of a VLAN bridging configuration that contains three protocol-based VLANs (two IP and one IPX) that overlap on an FDDI link (port 1 in each VLAN). (You can configure the link to be part of a trunk, as described in Chapter 8.) The end stations and servers are on 100Mbps ports, with traffic segregated by protocol. They are aggregated over the FDDI link.

Figure 31 Example of a Bridging Protocol-based VLAN Configuration

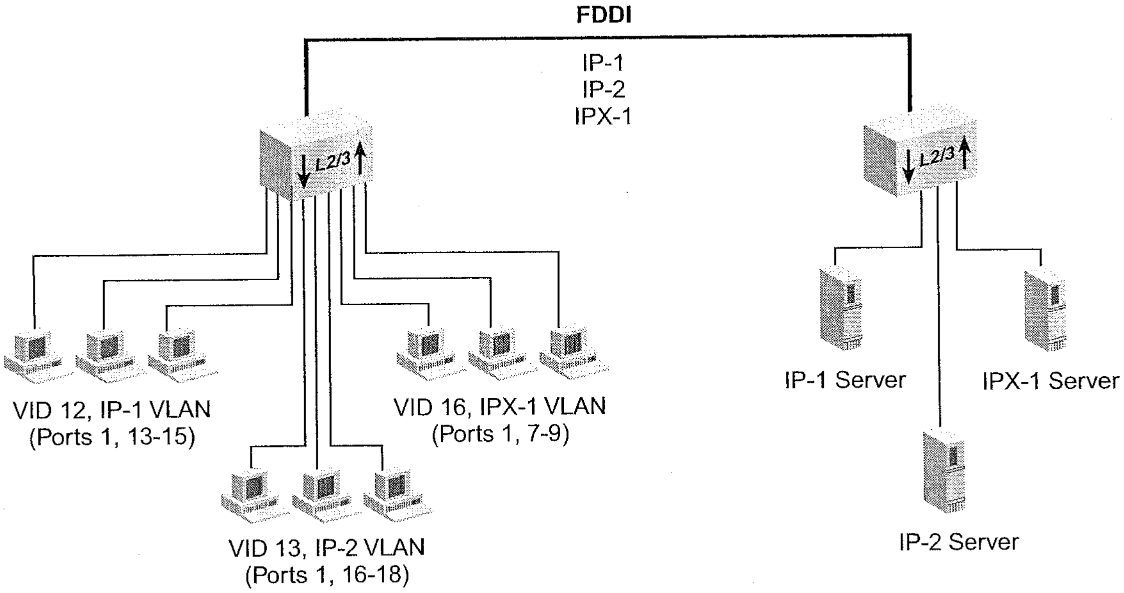

Table 20 shows the information that can be used to configure these VLANs on Device 1 (the device on the left):

Table 20 Sample Protocol-based VLAN Definitions

IP-1 VLAN	**IP-2 VLAN**	**IPX-1 VLAN**
VLAN Index 2	VLAN Index 3	VLAN Index 4
VID 12	VID 13	VID 16
Bridge ports *1, 13-15*	Bridge ports *1, 16-18*	Bridge Ports *1, 7-9*
Protocol type *IP*	Protocol type *IP*	Protocol type *IPX-802.3*
No Layer 3 address	No Layer 3 address	No Layer 3 address
Per-port tagging: ▪ Port 1 — *IEEE 802.1Q* ▪ Ports 13-15 — *none*	Per-port tagging: ▪ Port 1 — *IEEE 802.1Q* ▪ Ports 16-18 — *none*	Per-port tagging: ▪ Port 1 — *none* ▪ Ports 7-9 — *none*
VLAN name *IP-1*	VLAN name *IP-2*	VLAN name *IPX-1*

Establishing Routing Between VLANs

Your system supports routing using IP, IPX, and AppleTalk VLANs. If VLANs are configured for other routable network layer protocols, the VLANs can communicate between those protocols only through an external router.

The system's routing over bridging model allows you to configure routing protocol interfaces based on a static VLAN defined for one or more protocols. You must first define a VLAN to support one or more protocols and then assign a routing interface for each protocol associated with the VLAN. (You can also opt to use a routing versus bridging model by defining a router port IP interface, as defined in Chapter 11.)

Because the system supports router port IP interfaces as well as IP router interfaces for static VLANs, you must specify the interface type vlan when you define an IP interface for a static VLAN.

Important Considerations

To create an IP interface that can route through a static VLAN, you must:

1. Create a protocol-based IP VLAN for a group of bridge ports. If the VLAN overlaps with another VLAN at all, define it in accordance with the requirements of your VLAN mode.

   This IP VLAN does not need to contain Layer 3 information. An IP VLAN with Layer 3 information is a network-based VLAN. See "Network-based IP VLANs" later in this chapter.

2. Configure an IP routing interface with a network address and subnet mask, and specify the interface type vlan.

3. Select the IP VLAN index that you want to "bind" to that IP interface.

   If Layer 3 information is provided in the IP VLAN for which you are configuring an IP routing interface, the subnet portion of both addresses must be compatible. For example:

   - IP VLAN subnet 157.103.54.0 with subnet mask of 255.255.255.0
   - IP host interface address 157.103.54.254 with subnet mask of 255.255.255.0

   Layer 2 (bridging) communication is still possible within an IP VLAN (or router interface) for the group of ports within that IP VLAN:

   - For allClosed VLANs, IP data destined for a different IP subnetwork uses the IP routing interface to reach that different subnetwork even if the destination subnetwork is on a shared port.
   - For allOpen VLANs, using the destination MAC address in the frame causes the frame to be bridged; otherwise, it is routed in the same manner as allClosed VLANs.

4. Enable IP routing.

You perform similar steps to create IPX and AppleTalk routing interfaces. For more information, see the routing chapters in this guide (for routing protocols such as IP, OSPF, IPX, and AppleTalk).

Example: Protocol-based VLANs for Routing

Figure 32 shows a VLAN configuration that contains three IP VLANs without overlapped ports.

Figure 32 Sample VLAN Routing Configuration

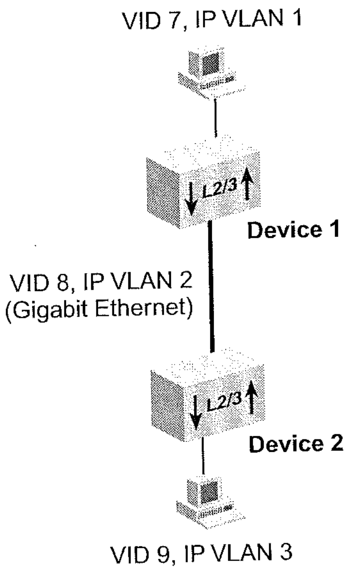

Table 21 shows the information that is used to configure these routing VLANs:

Table 21 Sample Protocol-based VLAN Definitions for Routing

IP VLAN1 (Device 1)	IP VLAN2 (Devices 1 and 2)	IP VLAN3 (Device 2)
VLAN Index 2	VLAN Index 3	VLAN Index 4
VID 7	VID 8	VID 9
Bridge port 6	Bridge port 13 on device 1 Bridge port 1 on device 2	Bridge port 8
Protocol type *IP*	Protocol type *IP*	Protocol type *IP*
No Layer 3 address	No Layer 3 address	No Layer 3 address
Per-port tagging: Port 6 — *none*	Per-port tagging: Ports 1 and 13 — *none*	Per-port tagging: Port 8 — *none*
VLAN name *"IP VLAN 1"*	VLAN name *"IP VLAN 2"*	VLAN name *"IP VLAN 3"*

Network-based IP VLANs

For IP VLANs only, you can configure network-layer subnet addresses. With this additional Layer 3 information, you can create multiple independent IP VLANs with the same bridge ports. Untagged frames are assigned to a network-based VLAN according to both the protocol (IP) and the Layer 3 information in the IP header. Assigning Layer 3 address information to IP VLANs is one way that network administrators can manage their IP routing interfaces by subnetwork.

Because network-based IP VLANs accommodate multiple routing interfaces over the same set of ports without tagging, this option can be useful in allOpen mode. In allClosed mode, overlapped network-based IP VLANs must be IEEE 802.1Q tagged, which means that the system does not use the Layer 3 information.

Important Considerations

When you create a network-based VLAN interface, review these guidelines:

- You can either configure network-based IP VLANs (IP VLANs with unique Layer 3 IP addresses) or you can define a single protocol-based VLAN with the protocol type IP and then define multiple IP routing interfaces for that VLAN.
- The network information is used only when multiple network-based VLANs are defined on a particular port. In situations where there is only one network-based VLAN defined on a port, the VLAN is treated as an ordinary IP protocol-based VLAN, and network-based information is ignored.
- When they are overlapped, network-based VLAN interfaces take precedence over protocol-based and port-based VLAN interfaces.

- You can define only *one* IP routing interface for a network-based VLAN. When you define an IP routing interface with the interface type vlan, the system does not allow you to select a network-based IP VLAN that already has a routing interface defined for it. For more information about IP routing interfaces, see Chapter 11.
- If you define multiple interfaces for an IP VLAN (instead of defining a network-based VLAN), you cannot subsequently modify that IP VLAN to supply Layer 3 address information. If only one routing interface is defined for the IP VLAN, then you can supply Layer 3 address information as long as it matches the Layer 3 information specified for the routing interface.
- In allClosed VLAN mode, you must supply IEEE 802.1Q tagging for any overlapped ports. Therefore, this feature has no added benefit. When IEEE 802.1Q tagging is implemented, implicit VLAN membership information such as the protocol or Layer 3 IP network address is not used; the frame is assigned to the VLAN based solely on the tag VID and the receive port.
- In allOpen mode, you are not required to supply the IEEE 802.1Q tagging. To ensure line-speed throughput for overlapped network-based IP VLANs in allOpen mode, however, you should still supply the IEEE 802.1Q tagging.

Example of Network-based VLANs

Figure 33 shows two network-based IP VLAN interfaces. The IPVLAN2 interface includes trunk ports and defines the protocol type IP, a Layer 3 address, a subnet mask, and IEEE 802.1Q tagging on bridge ports 6 and 7 (the anchor port for the trunk that uses ports 7 and 8). The IPVLAN3 interface defines IP and a different Layer 3 address; it uses exactly the same ports as IP VLAN2, with IEEE 802.1Q tagging on bridge ports 6 and 7.

Figure 33 Network-based VLANs with Overlapped Ports

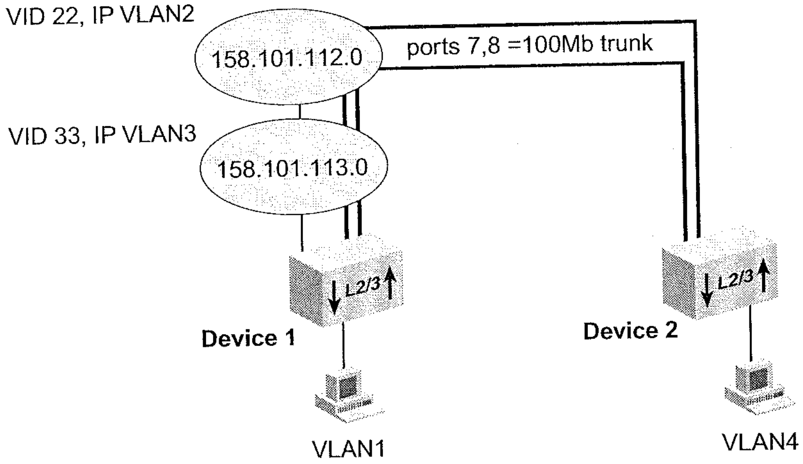

Table 22 shows the information that can be used to configure the two overlapped IP VLANs on Device 1:

Table 22 Network-based IP VLAN Definitions with Overlapped Ports

IP VLAN2	IP VLAN3
VLAN Index 2	VLAN Index 3
VID 22	VID 33
Bridge ports 6, *7* (7 is anchor port for a trunk that uses ports 7 and 8)	Bridge ports 6,*7* (7 is anchor port for a trunk that uses ports 7 and 8)
Protocol type *IP*	Protocol type *IP*
158.101.112.0 Layer 3 address 255.255.255.0 mask	158.101.113.0 Layer 3 address 255.255.255.0 mask
Per-port tagging: ■ Port 6 — *IEEE 802.1Q* ■ Anchor port 7 — *IEEE 802.1Q*	Per-port tagging: ■ Port 6 — *IEEE 802.1Q* ■ Anchor port 7 — *IEEE 802.1Q*
VLAN name *IPVLAN2*	VLAN name *IPVLAN3*

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, port membership, protocol type, and Layer-3 (network) address information, 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 has the most specific match. The system uses this protocol match hierarchy to find the most specific match:

1. IEEE 802.1Q tag VID value, if the frame is tagged
2. A specific protocol match (for example, IP, IPX, or AppleTalk)
3. Either the default VLAN (an untagged, unspecified protocol type VLAN with all ports and a VID of 1) or any VLAN that has the unspecified protocol type
4. The *null VLAN*, a special VLAN that the system uses if the frame cannot be assigned to any VLAN. This VLAN has no ports and has no address table (in allClosed mode).

The CoreBuilder 3500 Release 3.0 ingress rules are classified according to the tag status of the frame and the VLAN mode (allOpen for open VLANs or allClosed for closed VLANs). For the ingress rules, the system considers a priority tagged frame to be an untagged frame.

The flow chart in Figure 34 shows the VLAN ingress rules for the system at Release 3.0.

Figure 34 Flow Chart for 3.0 Ingress Rules

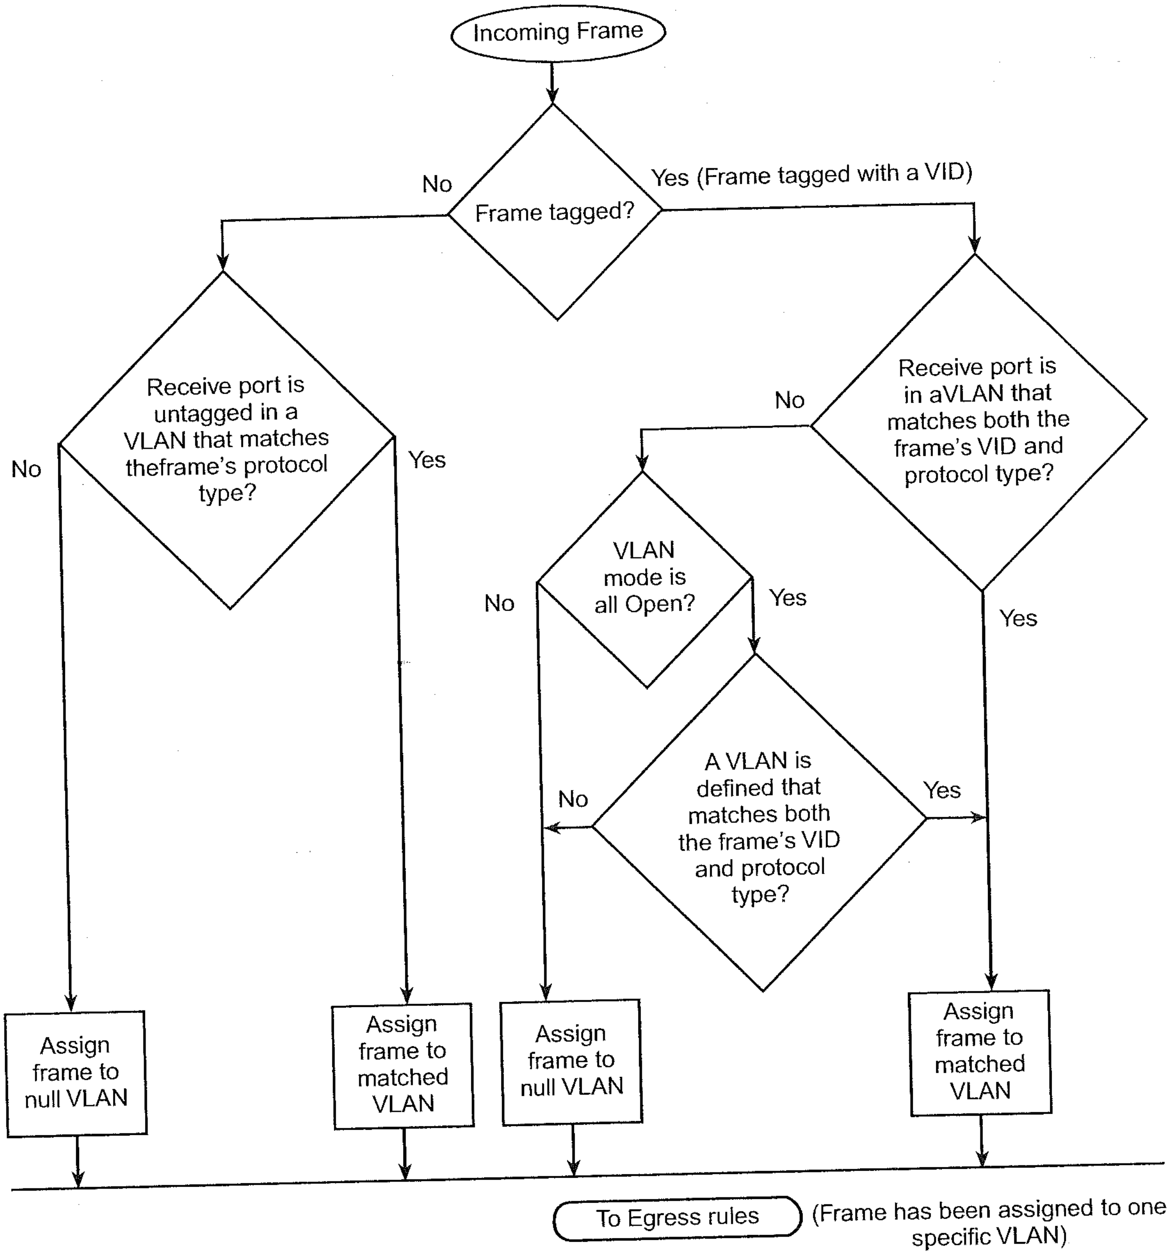
Incoming Frame
Frame tagged?
No
Yes (Frame tagged with a VID)
Receive port is untagged in a VLAN that matches theframe's protocol type?
No
Yes
Receive port is in aVLAN that matches both the frame's VID and protocol type?
No
Yes
VLAN mode is all Open?
No
Yes
A VLAN is defined that matches both the frame's VID and protocol type?
No
Yes
Assign frame to null VLAN
Assign frame to matched VLAN
Assign frame to null VLAN
Assign frame to matched VLAN
To Egress rules
(Frame has been assigned to one specific VLAN)

The ingress rules for tagged frames vary for the various system releases. Table 23 summarizes the differences.

Table 23 Ingress Rules for IEEE 802.1Q Tagged Frames Based on VLAN Mode and Software Release Number

VLAN Mode	Release 1.2	Release 2.0	Release 3.0	Action Without Required Match
allOpen	The tagged frame is assigned to one of the configured VLANs if: ■ The VID of the frame matches that of a VLAN *and* ■ A port in that VLAN is tagged	The tagged frame is assigned to one of the configured VLANs if: ■ The VID of the frame matches that of a VLAN	The tagged frame is assigned to one of the configured VLANs if: ■ The VID of the frame matches that of a VLAN *and* ■ The protocol type of the frame matches that of the same VLAN	The frame is assigned to the null VLAN. It can still be forwarded (untagged) if the destination address of the frame is associated with another port in the bridge address table.
allClosed	The tagged frame is assigned to one of the configured VLANs if: ■ The receive port is in a VLAN with a VID that matches that of the frame *and* ■ A port in that VLAN is tagged	The tagged frame is assigned to one of the configured VLANs if: ■ The receive port is in a VLAN with a VID that matches that of the frame	The tagged frame is assigned to one of the configured VLANs if: ■ The receive port is in a VLAN with a VID matching that of the frame *and* ■ The protocol type of the frame matches that of the same VLAN	The frame is assigned to the null VLAN and dropped.

Egress Rules

These rules determine whether the *outgoing* frame is forwarded, filtered (dropped), or flooded; they also determine the frame's tag status. Although the same standard bridging rules apply to both open and closed VLANs, they result in different behavior depending on the allOpen mode (one address table for the system) versus allClosed mode (one address table for each VLAN).

Standard Bridging Rules for Outgoing Frames

The frame is handled according to these bridging rules:

- If the frame's destination address matches an address that was previously 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 the receive port and transmit port are in the same VLAN or the system is in allOpen mode.
- 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). Those frames assigned to the null VLAN are not flooded to any ports because no ports are associated with the null VLAN. See "Examples of Flooding and Forwarding Decisions" later in this chapter.
- 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 either a management/configuration frame (such as a RIP update, SNMP get/set PDU, Administration Console Telnet packet, or a Web Management Interface http packet), or it is a routed packet. If it is a routed packet, the system performs the routing functions described in the appropriate routing chapter (for example, IP, OSPF, IPX, or AppleTalk).

For example, if a frame is associated with VLAN A and has a destination address associated with VLAN B, the frame is flooded over VLAN A in allClosed mode but forwarded untagged in allOpen mode.

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. Priority tagged frames for QoS use the same frame format as IEEE 802.1Q tagging but with a VID of 0. Priority tagged frames received by the system are transmitted as either untagged frames (that is, no priority tagging) or IEEE 802.1Q tagged frames.

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

If the transmit port is not a member of the assigned VLAN, the frame is transmitted untagged. For VLANs in allOpen mode, this result may occur in either of these situations:

- *If the frame is assigned to the null VLAN. (The frame can still be forwarded if the address was statically entered in the address table or dynamically learned on another VLAN.)*
- *If the frame is assigned to a specific VLAN but the transmit port is* not *part of this VLAN.*

Examples of Flooding and Forwarding Decisions

This section provides several examples of flooding and forwarding decisions.

Example 1: Flooding Decisions for Protocol-based VLANs

Table 24 shows how flooding decisions are made according to three VLANs that are set up by protocol (assuming a 12-port configuration). In this example, ports and frames are untagged and the destination address is unknown, multicast, or broadcast.

Table 24 Protocol-based VLANs and Flooding Decisions

Index	VID	VLAN Name	Ports
1	1	Default	1 – 12
2	2	IP1	1 – 8
3	3	IPX1	9 – 11

Untagged data received on this port	Is flooded on this VLAN	Because
IP - port 1	IP1, VID 2	IP data received matches IP1 on the source (receive) port.
IPX - port 11	IPX1, VID 3	IPX data received matches IPX1 on the source port.
XNS - port 1	Default, VID 1	XNS data received matches no protocol VLAN, so the Default VLAN is used.