

## Bridge/Router Interface (BRIM) User's Guide

CABLETRON  
SYSTEMS



## Notice

Cabletron Systems reserves the right to make changes in specifications and other information contained in this document without prior notice. The reader should in all cases consult Cabletron Systems to determine whether any such changes have been made.

The hardware, firmware, or software described in this manual is subject to change without notice.

IN NO EVENT SHALL CABLETRON SYSTEMS BE LIABLE FOR ANY INCIDENTAL, INDIRECT, SPECIAL, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING BUT NOT LIMITED TO LOST PROFITS) ARISING OUT OF OR RELATED TO THIS MANUAL OR THE INFORMATION CONTAINED IN IT, EVEN IF CABLETRON SYSTEMS HAS BEEN ADVISED OF, KNOWN, OR SHOULD HAVE KNOWN, THE POSSIBILITY OF SUCH DAMAGES.

## Virus Disclaimer

Cabletron has tested its software with current virus checking technologies. However, because no anti-virus system is 100% reliable, we strongly caution you to write protect and then verify that the Licensed Software, prior to installing it, is virus-free with an anti-virus system in which you have confidence.

Cabletron Systems makes no representations or warranties to the effect that the Licensed Software is virus-free.

Copyright © 1999 by Cabletron Systems, Inc. All rights reserved.

Printed in the United States of America.

Order Number: 9031617-04 April 1999

Cabletron Systems, Inc.  
P.O. Box 5005  
Rochester, NH 03866-5005

**SPECTRUM**, **MiniMMAC**, **FNB**, **Multi Media Access Center**, and **DNI** are registered trademarks, and **Portable Management Application**, **IRM**, **IRM2**, **IRM3**, **IRBM**, **ESXMIM**, **ETSMIM**, **EMME**, **EMM-E6**, **ETWMIM**, **FDMMIM**, **FDCMIM**, **MicroMMAC**, **MRXI**, **MRXI-24**, **NB20E**, **NB25E**, **NB30**, **NB35E**, **NBR**, **SEHI**, **STHI**, **TRBMIM**, **TRMM**, **TRMM-2**, **TRMM-4**, **TRMMIM**, **TRXI**, **Media Interface Module**, **MIM**, and **Flexible Network Bus** are trademarks of Cabletron Systems, Inc.

**UNIX** and **OPENLOOK** are trademarks of Unix System Laboratories, Inc. **OSF/Motif** and **Motif** are trademarks of the Open Software Foundation, Inc. **X Window System** is a trademark of X Consortium, Inc. **Ethernet** and **XNS** are trademarks of Xerox Corporation. **Apple** and **AppleTalk** are registered trademarks of Apple Computer, Inc. **Banyan** is a registered trademark of Banyan Systems, Inc. **DECnet** is a registered trademark of Digital Equipment Corporation. **Novell** is a registered trademark of Novell, Inc. **CompuServe** is a registered trademark of CompuServe. **Sun Microsystems** is a registered trademark, and **Sun**, **SunNet**, and **OpenWindows** are trademarks of Sun Microsystems, Inc.

# Restricted Rights Notice

(Applicable to licenses to the United States Government only.)

1. Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Cabletron Systems, Inc., 35 Industrial Way, Rochester, New Hampshire 03867-0505.

2. (a) This computer software is submitted with restricted rights. It may not be used, reproduced, or disclosed by the Government except as provided in paragraph (b) of this Notice or as otherwise expressly stated in the contract.  
  
(b) This computer software may be:
  - (1) Used or copied for use in or with the computer or computers for which it was acquired, including use at any Government installation to which such computer or computers may be transferred;
  - (2) Used or copied for use in a backup computer if any computer for which it was acquired is inoperative;
  - (3) Reproduced for safekeeping (archives) or backup purposes;
  - (4) Modified, adapted, or combined with other computer software, provided that the modified, combined, or adapted portions of the derivative software incorporating restricted computer software are made subject to the same restricted rights;
  - (5) Disclosed to and reproduced for use by support service contractors in accordance with subparagraphs (b) (1) through (4) of this clause, provided the Government makes such disclosure or reproduction subject to these restricted rights; and
  - (6) Used or copied for use in or transferred to a replacement computer.
- (c) Notwithstanding the foregoing, if this computer software is published copyrighted computer software, it is licensed to the Government, without disclosure prohibitions, with the minimum rights set forth in paragraph (b) of this clause.
- (d) Any other rights or limitations regarding the use, duplication, or disclosure of this computer software are to be expressly stated in, or incorporated in, the contract.
- (e) This Notice shall be marked on any reproduction of this computer software, in whole or in part.

## Chapter 1 Introduction

Using the BRIM User's Guide.....	1-2
Related Manuals.....	1-3
BRIM Management Applications .....	1-3
Software Conventions .....	1-4
Using the Mouse .....	1-4
Common Device Window Fields.....	1-6
Using Window Buttons.....	1-7
Getting Help .....	1-8
Using On-line Help.....	1-8
Accessing On-line Documentation.....	1-8
Getting Help from the Cabletron Systems Global Call Center .....	1-9

## Chapter 2 Bridging

Bridging Basics .....	2-1
About Transparent Bridging .....	2-1
About Source Route Bridging .....	2-2
About Source Route-Transparent Bridges.....	2-3
About Source Route-Translational Bridges.....	2-4
Viewing and Managing Bridging Interfaces .....	2-5
The Bridge Status Window .....	2-7
Accessing Bridge Status Window Options.....	2-9
Enabling and Disabling Bridging .....	2-14
Enabling and Disabling Individual Interfaces.....	2-14
Enabling and Disabling All Installed Interfaces.....	2-14
Bridge Statistics .....	2-15
Performance Graphs.....	2-15
Configuring Performance Graphs.....	2-17
Bridge Detail Breakdown .....	2-18
Bridge Port Detail Breakdown.....	2-19
Interface Statistics .....	2-20
CSMACD Statistics .....	2-23
PPP Link Statistics .....	2-25
Dot5 Errors.....	2-28
Source Route Statistics .....	2-31
Spanning Tree .....	2-34
Bridge Level.....	2-35
Bridge Port Level .....	2-37
Configuring Spanning Tree .....	2-39

Changing Bridge Priority .....	2-39
Changing the Spanning Tree Algorithm Protocol Type .....	2-39
Changing Hello Time .....	2-40
Changing Max Age Time .....	2-40
Changing Forwarding Delay Time.....	2-40
Changing Port Priority.....	2-41
Changing Path Cost.....	2-41
Filtering Database .....	2-41
Configuring the Filtering Database .....	2-45
Special Filter Databases.....	2-47
Ethernet and Token Ring Special Filter Database Windows .....	2-48
Special Filter Database Window .....	2-49
Defining and Editing Filters in the Special Database .....	2-51
Changing the Receive Ports.....	2-52
Changing the Port Filtering Action .....	2-52
Setting the Port Filtering Action .....	2-52
Clearing the Port Filtering Action .....	2-53
Enabling and Disabling a Filter .....	2-53
Saving a Set of Filters to a File.....	2-53
Interface Configuration .....	2-54
Bridge Methods .....	2-55
Setting the Bridge Method .....	2-56
Protocol Transmission Methods.....	2-56
Source Route Configuration .....	2-57
Source Routing Information .....	2-58
Source Route Configuration .....	2-59
Making and Setting Changes .....	2-62
Using the Find Source Address Feature.....	2-63
Using the Port Source Addresses Window.....	2-64
Setting the Aging Time.....	2-65
Using the Token Ring Bridge and Port Configuration Windows.....	2-65
Duplex Modes.....	2-70
Setting the Duplex Mode .....	2-72
Ethernet Port Configuration.....	2-72
Fast Ethernet Port Configuration.....	2-73
Setting the Operational Mode for the FE-100TX .....	2-76
Setting the Operational Mode for the FE-100FX .....	2-77
SONET Port Configuration.....	2-77
SONET/SDH Configuration .....	2-77
SONET/SDH Statistics.....	2-79
Configuring SmartTrunking .....	2-85
Configuring Broadcast Suppression .....	2-88
Token Ring Bridge Mode .....	2-89
Setting Token Ring Bridge Mode.....	2-90
Setting Bridge Translation.....	2-91
Enabling and Disabling Auto and Dual Translate Modes .....	2-92
Configuring Token Ring Packet Translation .....	2-93
Configuring Novell Token Ring Packet Translation .....	2-93
The IBM Translation Table .....	2-94
The SNAP Translation Table.....	2-96

Configuring SNAP Translation.....	2-96
Using the Novell Translation Window .....	2-97
Using the Physical View Windows for the ETWMIM .....	2-99
Ethernet Port Physical View .....	2-99
Token Ring Port Physical View .....	2-101

### **Chapter 3 FDDI Applications**

Accessing the FDDI Menu .....	3-2
The Configuration Window.....	3-2
Connection Policy Window .....	3-5
Station List Window .....	3-8
FDDI Performance Window .....	3-10
Setting the Time Interval.....	3-11

### **Chapter 4 WAN Applications**

Accessing the WAN Status Windows .....	4-1
Viewing WAN Interface Status .....	4-2
Configuring the Synchronous Connection.....	4-3
Configuring T-1 Ports.....	4-5
Using the T1 FracTable Configuration Window .....	4-7
Configuring the Fractional Table.....	4-8
Restoring a Fractional Table.....	4-9
Changing the Interface Line Coding .....	4-9
Displaying the WAN Logical View .....	4-10
Changing WAN Logical Settings.....	4-11
Viewing the WAN Port Admin/Status .....	4-12
Synchronous Admin/Status.....	4-12
T1 Admin/Status .....	4-13
Enabling and Disabling WAN T1 Interfaces.....	4-14
Displaying Synchronous Port Statistics .....	4-14

### **Chapter 5 ATM Configuration**

Accessing the ATM Connections Window .....	5-1
Configuring Connections.....	5-3
Adding a New Connection.....	5-4
Deleting a Connection.....	5-5

### **Index**





# Introduction

*Using this guide; related manuals; management applications available for BRIMs; software conventions; getting help; contacting Cabletron Systems Global Call Center*

---


Welcome to the Cabletron Systems' *SPECTRUM® Element Manager Bridge/Router Interface (BRIM) User's Guide*. We have designed this guide to serve as a reference for using SPECTRUM Element Manager for all types of BRIMs.

Cabletron Systems' Bridge/Router Interface Modules (BRIMs) provide flexible, integrated bridging functionality (including traffic filtering by destination, source, type field, and 64-byte data offset, and support for the Spanning Tree Algorithm) or routing functionality to the network hub. By accommodating several media types and topologies, these network interfaces connect to any standard Local Area Network (LAN) or Wide Area Network (WAN).

Since BRIMs plug into Cabletron's Intel i960-based products (e.g., EMM-E6, ESXMIM, MicroMMAC, and MicroMMAC-T), SPECTRUM Element Manager views the hub and bridge/router as a single entity. The following BRIMs are supported by SPECTRUM Element Manager:

- |                              |  |
|------------------------------|--|
| <b>BRIM-A6<br/>BRIM-A6DP</b> | The ATM Bridge/Router Interface Modules feature high speed ATM connectivity (at rates up to 155 Mbps). They are fully compatible with the 9A000 ATM Switch Module for the SmartSwitch® 9000, and Fore Systems' ASX-200. They support the IETF AToM MIB, Multi-Protocol Encapsulation over AAL5, and many other protocols, and meet ATM Forum UNI specifications. The BRIM-A6 supports a single ATM interface, and the BRIM-A6DP supports two redundant ATM interfaces. |
| <b>BRIM-E6</b>               | The Ethernet BRIM has a user-configurable Ethernet Port Interface Module (EPIM) port that provides media flexibility for AUI, UTP, STP, fiber or coax cable.   |

<b>BRIM-E100</b>	The Fast Ethernet BRIM has a user-configurable Fast Ethernet Port Interface Module (FEPIM) port that provides either 100BASE-TX connectivity (via an RJ-45 interface) or 100BASE-FX connectivity (via an MMF interface with an SC connector).
<b>BRIM-F6</b>	The FDDI Bridge/Router Interface Module is a translational FDDI to Ethernet bridge, via media configurable Fiber Port Interface Module (FPIM) connectors using MMF or SMF fiber, or UTP or STP copper. The BRIM-F6 supports the IETF FDDI MIB and dual homing. Its DAS attachment is ANSI-compliant.
<b>BRIM-W6</b>	The WAN Bridge/Router Interface Module supports Synchronous, T1 and E1 connections. It carries PPP, Frame Relay and X.25 protocols through its WAN Port Interface Modules.



*Cabletron Systems has discontinued manufacturing several BRIMs, including the BRIM-WT1, the Cisco BRIM-W/E, the BRIM-A100, the BRIM-T6, the BRIM-FO, the BRIM-FD1, the BRIM-FD2, and the BRIM-F5.*

## Using the BRIM User's Guide

Each chapter in this guide describes one major functionality or a collection of several smaller functionalities that the BRIM adds to the device. This guide contains information about software functions which are accessed directly from the device icon; for information about management functions which are accessed via the SPECTRUM Element Manager platform, consult the *SPECTRUM Element Manager User's Guide*, and *SPECTRUM Element Manager Tools Guide*, and the *Remote Administration Tools User's Guide*.

Chapter 1, **Introduction**, provides a list of related documentation, describes certain software conventions, and shows you how to contact Cabletron Systems' Global Call Center.

Chapter 2, **Bridging**, provides a comprehensive look at all management options associated with the bridge portion of the device, including Bridge Performance Graphs, Spanning Tree, and the Filtering and Special Filtering Databases.

Chapter 3, **FDDI Applications**, describes the FDDI management windows, including Configuration, Connection Policy, Station List, and Performance.

Chapter 4, **WAN Applications**, describes the windows available for WAN Port configuration, and the Synchronous Port Statistics window.

Chapter 5, **ATM Configuration**, describes how to configure Permanent Virtual Circuits (PVCs) for the ATM interface(s) available on the ATM BRIM.

We assume that you have a general working knowledge of Ethernet IEEE 802.3, FDDI, WAN, and ATM type data communications networks and their physical layer components, and that you are familiar with general bridging concepts.

## Related Manuals

The *BRIM User's Guide* is only part of a complete document set designed to provide comprehensive information about the features available to you through SPECTRUM Element Manager. Other guides which include important information related to managing the BRIM include:

Cabletron Systems' *SPECTRUM Element Manager User's Guide*

Cabletron Systems' *SPECTRUM Element Manager Tools Guide*

Cabletron Systems' *SPECTRUM Element Manager Remote Administration Tools User's Guide*

Cabletron Systems' *SPECTRUM Element Manager Remote Monitoring (RMON) User's Guide*

Cabletron Systems' *SPECTRUM Element Manager Alarm and Event Handling User's Guide*

Cabletron Systems' *Network Troubleshooting Guide*

Microsoft Corporation's *Microsoft Windows User's Guide*

For more information about the capabilities of the BRIM and the host device in which it is installed, consult the appropriate hardware documentation.


For more information about the capabilities of SPECTRUM Element Manager for the host device, consult its User's Guide.

## BRIM Management Applications

When a device (e.g., MicroMMAC, NBR-620, EMM-E6, or ESXMIM) has a BRIM installed and enabled, additional applications will be available from the Chassis View window. These applications will depend on the type of BRIM installed in your device.

- **Bridge Status** will be available from the **Device** menu for all devices with BRIMs installed. The Bridging options that are available from the Bridge Status window will vary depending on the device and the type of BRIM you have installed; see **Chapter 2, Bridging**, for details on the bridge applications.
- **Router Utilities (Basic Router Configuration and Advanced Router Configuration)** will be available from the **Tools** menu when a BRIM is installed and enabled on your device. For more information, see the *Routing Services Configuration Guide*.


- **FDDI** menu will be available on devices that have an FDDI BRIM installed and enabled; see **Chapter 3, FDDI Applications**, for details.
- **WAN Status** will be available from the **Device** menu for devices that have a WAN BRIM installed and enabled; see **Chapter 4, WAN Applications**, for details.
- **ATM Connections** will be available from the **Device** menu for devices that have an ATM BRIM installed and enabled; see **Chapter 5, ATM Configuration**, for details.



*If you cannot determine if you have a BRIM installed in your device or are unsure of the type of installed BRIM, the I/F Summary window can help you find the answer. To access this window, select **Device**→**I/F Summary** from your device's Chassis View window. This window's **Description** field lists the interface descriptions for the device, including the type of BRIM installed in the device (e.g., Ctron FDDI BRIM port). The BRIM port(s) will always be listed at the end of the interface list, after all the non-BRIM interfaces.*

## Software Conventions

SPECTRUM Element Manager's device user interface contains a number of elements which are common to most windows and which operate the same regardless of which window they appear in. A brief description of some of the most common elements appears below; note that the information provided here is not repeated in the descriptions of specific windows and/or functions.



*In accordance with Year 2000 compliance requirements, SPECTRUM Element Manager displays and allows you to set all dates with four-digit year values.*

## Using the Mouse

This document assumes you are using a Windows-compatible mouse with two buttons; if you are using a three button mouse, you should ignore the operation of the middle button when following procedures in this document. Procedures within the SPECTRUM Element Manager document set refer to these buttons as follows:




Figure 1-1. Mouse Buttons

For many mouse operations, this document assumes that the left (primary) mouse button is to be used, and references to activating a menu or button will not include instructions about which mouse button to use.

However, in instances in which right (secondary) mouse button functionality is available, instructions will explicitly refer to **right** mouse button usage. Also, in situations where you may be switching between mouse buttons in the same area or window, instructions may also explicitly refer to both **left** and **right** mouse buttons.

Instructions to perform a mouse operation include the following terms:

- **Pointing** means to position the mouse cursor over an area without pressing either mouse button.
- **Clicking** means to position the mouse pointer over the indicated target, then press and release the appropriate mouse button. This is most commonly used to select or activate objects, such as menus or buttons.
- **Double-clicking** means to position the mouse pointer over the indicated target, then press and release the mouse button two times in rapid succession. This is commonly used to activate an object's default operation, such as opening a window from an icon. Note that there is a distinction made between "click twice" and "double-click," since "click twice" implies a slower motion.
- **Pressing** means to position the mouse pointer over the indicated target, then press and hold the mouse button until the described action is completed. It is often a pre-cursor to Drag operations.
- **Dragging** means to move the mouse pointer across the screen while holding the mouse button down. It is often used for drag-and-drop operations to copy information from one window of the screen into another, and to highlight editable text.

## Common Device Window Fields

Similar descriptive information is displayed in boxes at the top of most device-specific windows in SPECTRUM Element Manager, as shown in Figure 1-2.

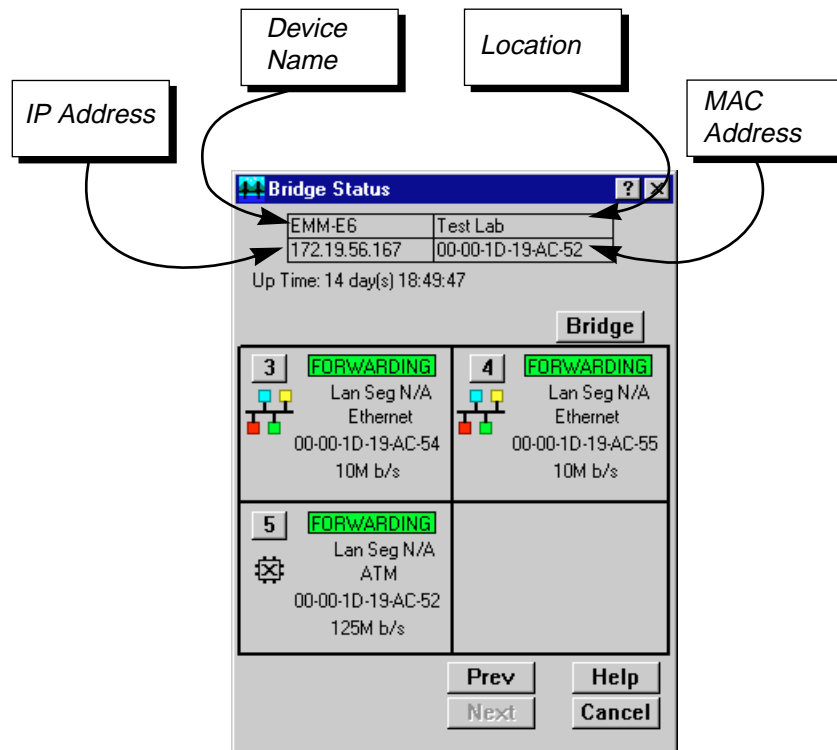


Figure 1-2. Sample Window Showing Group Boxes

### Device Name

Displays the user-defined name of the device. The device name can be changed via the System Group window; see the *Generic SNMP User's Guide* for details.

### IP Address

Displays the device's IP (Internet Protocol) address; this will be the IP address used to define the device icon. IP addresses are assigned via Local Management for the device; they cannot be changed via SPECTRUM Element Manager.

### Location

Displays the user-defined location of the device. The location is entered through the System Group window; see the *Generic SNMP User's Guide* for details.

**MAC Address**

The physical layer address assigned to the interface associated with the IP address used to define the device icon when it was added to SPECTRUM Element Manager. MAC addresses are hard-coded in the device, and are not configurable.

Informational fields describing the boards and/or ports being modeled are also displayed in most windows:

**Board Number**

Displays the number indicating the position of the monitored board in the chassis.

**Port Number**

Displays the number of the monitored port.

**Active Users**

Indicates the number of users processing information through the device's repeater channel, board, or port, as determined by MAC addresses.

**Uptime**

Displays the amount of time, in a days hh:mm:ss format, that the device has been running since the last start-up.

## Using Window Buttons

The **Cancel** button that appears at the bottom of most windows allows you to exit a window and terminate any unsaved changes you have made. You may also have to use this button to close a window after you have made any necessary changes and set them by clicking on the **OK**, **Set**, or **Apply** button.

An **OK**, **Set**, or **Apply** button appears in windows that have configurable values; it allows you to confirm and SET changes you have made to those values. In some windows, you may have to use this button to confirm each individual set; in other windows, you can set several values at once and confirm the sets with one click on the button.

The **Help** button brings up a Help text box with information specific to the current window. For more information on the **Help** button, see [Using On-line Help, page 1-8](#).

The command buttons, for example **Bridge**, call up a menu listing the windows, screens, or commands available for that topic.

Any menu topic followed by ... (three dots) — for example **Statistics...** — calls up a window or screen associated with that topic.


## Getting Help

This section describes different methods of getting help for questions or concerns you may have while using SPECTRUM Element Manager.

### Using On-line Help

You can use the BRIM window **Help** buttons to obtain information specific to the device. When you click on a **Help** button, a window will appear which contains context-sensitive on-screen documentation that will assist you in the use of the window and its associated command and menu options. Note that if a **Help** button is grayed out, on-line help has not yet been implemented for the associated window.


From the **Help** menu accessed from the host device's Chassis View window menu bar, you can access on-line Help specific to the Chassis View window, as well as bring up the Chassis Manager window for reference. Refer to the device's user's guide for information on the Chassis View and Chassis Manager windows.



*All of the online help windows use the standard Microsoft Windows help facility. If you are unfamiliar with this feature of Windows, you can select **Help** from the Windows **Start** menu, or **Help** —>**How to Use Help** from the primary SPECTRUM Element Manager window, or consult your Microsoft Windows product **User's Guide**.*

### Accessing On-line Documentation

The complete suite of documents available for SPECTRUM Element Manager can be accessed via a menu option from the primary window menu bar: **Help** —> **Online Documents**. If you chose to install the documentation when you installed SPECTRUM Element Manager, selecting this option will launch Adobe's Acrobat Reader and a menu file which provides links to all other available documents.



*If you have not yet installed the documentation, the **Online Documents** option will not be able to access the menu file. In order to activate this option, you must run the setup.exe again to install the documentation component. See your **Installation Guide** for details.*



## Getting Help from the Cabletron Systems Global Call Center

If you need technical support related to SPECTRUM Element Manager, or if you have any questions, comments, or suggestions related to this manual or any of our products, please feel free to contact the Cabletron Systems Global Call Center via one of the following methods:

By phone: (603) 332-9400  
*24 hours a day, 365 days a year*

By mail: Cabletron Systems, Inc.  
PO Box 5005  
Rochester, NH 03866-5005

By Internet mail: support@ctron.com

FTP: ftp.ctron.com (134.141.197.25)

*Login* anonymous  
*Password* your email address

By BBS: (603) 335-3358

Modem Setting 8N1: 8 data bits, 1 stop bit, No parity

For additional information about Cabletron Systems products, visit our World Wide Web site: <http://www.cabletron.com/>. For technical support, select **Service and Support**.



# Bridging

*About bridging methods; viewing and managing bridging interfaces; using the Bridge Status window; viewing bridge statistics; using Spanning Tree; using the Filtering Database; configuring duplex modes; using the Interface Configuration window; and setting Bridge Translation*

---

## Bridging Basics

Bridges are used in local area networks to connect two or more network segments and to control the flow of packets between the segments. Ideally, bridges forward packets to another network segment only when necessary.

Bridges are also used to increase the fault tolerance in a local area network by creating redundant bridge paths between network segments. In the event of a bridge or bridge segment failure, an alternate bridge path will be available to network traffic, without significant interruption to its flow.

The method a bridge uses to forward packets, choose a bridge path, and ensure that a sending station's messages take only one bridge path depends on the bridge's type: Transparent (generally used in Ethernet or FDDI environments) or Source Routing (generally used in Token Ring environments), source routing-transparent, or source route-transparent—the two latter being combinations that are found in a mixed network environment.

## About Transparent Bridging

Transparent bridges are most common in Ethernet networks. Individual Transparent bridges monitor packet traffic on attached network segments to learn where end stations reside in relation to each segment by mapping the Source Address of each received frame to the port (and segment) it was detected on. This information gets stored in the bridge's Filtering Database.

When in the Forwarding state, the bridge compares a packet's destination address to the information in the Filtering Database to determine if the packet should be forwarded to another network segment or filtered (i.e., not forwarded). A bridge

filters a packet if it determines that the packet's destination address exists on the same side of the bridge as the source address.

If two or more bridges are connected to the same Ethernet LAN segment—placed in parallel—only a single bridge must be allowed to forward data frames onto that segment. If two or more bridges were forwarding data frames onto the same Ethernet segment, the network would soon be flooded.

With a data loop in the topology, bridges would erroneously associate a single source address with multiple bridge ports, and keep proliferating data by forwarding packets in response to the ever-changing (but incorrect) information stored in their Filtering Database.

To avoid such data storms, Transparent bridges communicate with one another on the network by exchanging Bridge Protocol Data Units (BPDUs) to determine the network topology and collectively implement a Spanning Tree Algorithm (STA) that selects a controlling bridge for each LAN segment. This ensures that only a single data route exists between any two end stations and that topology information remains current.

## About Source Route Bridging

Source Routing is typically used to connect two or more Token Ring network segments. Source Route bridges differ from Transparent bridges in that they do not build and then use a physical address database to make forwarding decisions. Instead, the source end station transmits packets with a header that contains routing information (added by bridges in the network topology during a route discovery process between end stations); once a route has been determined, a Source Route bridge simply reads the header of a source routed packet to determine whether it is a participant in routing the packet.

In Source Routing, sending and receiving devices employ broadcast packets—known as explorer packets—to determine the most efficient route for a message to travel. Generally, before a station sends a message, it will first send a test packet to all stations on the same ring; if the sending station receives a response to this packet, it assumes that the destination station is on the same ring and therefore it will not include routing information in frames sent to that station in the future. Any further packets issued between stations will appear to be transparent-style frames without embedded routing information.

If the sending station does not receive a response to the test packet, it will send explorer packets to the destination; the explorer packets will be propagated by the network's bridges as either All Paths Explorer (APE) packets or as Spanning Tree Explorer (STE) packets. The task of both packet types is to get the destination station to return specific route information to the sending station. They achieve this by including an identifier for each ring the explorer packet traversed and for each bridge between any rings).

Since the data flow on a Source Routed network is determined by end stations (unlike a Transparently bridged network), a looped bridge topology is not an issue for data flow. APE packets are sent from the source station over every

possible bridge path to the end station. The original APE frame contains no routing information (e.g., bridge numbers and ring numbers). As the frame is propagated along all available paths to the destination station, each bridge along the way adds its own bridge and ring numbers to the packet's RIF before forwarding it, thereby providing route information.

In response to each received APE packet, the destination station directs a reply to the sending station. On receiving the replies, the sending station ideally assumes that the first returned reply contains the most efficient route. The sending station then stores the route information and uses it to send subsequent transmissions to the same station.

Because APE frames do increase network traffic, some sites may use STE explorer frames as an alternate method of route discovery. With STE exploration, a Spanning Tree Algorithm (either configured automatically via BPDUs or manually via management) is maintained for the sole purpose of determining how to direct an explorer frame during route discovery.

During the discovery process, a source station will send out STE explorer frames into a bridged topology. If a bridge is in a forwarding state according to Spanning Tree, it will forward an explorer frame onto its attached LAN segment (appending the Bridge and LAN Segment Identifiers in the appropriate area of the RIF); if the bridge is filtering, it will discard the explorer frames. In this fashion, only a single explorer frame will reach each individual LAN segment.

Ultimately, the destination station will receive only a single STE packet, and will respond with APE packets (that return to the sending station on all possible bridge paths) or an STE packet (that returns to the sending station via in the reverse route of the STE explorer packet).

Although the Spanning Tree Algorithm determines the bridge path an STE takes to the destination station, during future communication between the stations, bridges along the route will use Source Routing to forward the packet (i.e., the bridges will read the Routing Information Field in the header of specifically routed frames to decide whether to forward them).


## About Source Route-Transparent Bridges

Because network topologies have developed in which bridges must be able to handle network traffic from end stations which support source routing and others which do not, a hybrid type of bridge—Source Route-Transparent (SRT) — combines elements of both bridging methods.

An end station's network drivers can be configured in software to use a bit setting in the source address portion of a data frame to indicate whether the station is to operate in a Source Route or Transparently bridged network environment. The Routing Information Indicator (RII) bit of the source address is set to 1 if the station is to use Source Routing; if the station is to operate in a Transparently bridged environment, the RII bit is left unchanged (i.e., at 0).

Not all end stations in a Token Ring environment have network drivers which support Source Routing—whether the drivers are improperly configured via management or they simply are not source-route capable.

In a network with a mix of Source Route and Transparent end stations, data frames from both station types must be bridged correctly. An SRT bridge inspects the RII bit setting of incoming frames to determine whether they should be Transparently bridged (if the RII bit was at 0) or Source Routed (if the RII bit was set to 1) to their destination and will use the appropriate bridge method to forward the frame.



*Cabletron has extended the functionality of Ethernet ports on translational bridges, so the ports can be set to Source Route mode. When an Ethernet port is in Source Route mode and receives an SR packet from a Token Ring port, it will save the Source Routing information and send out the packet transparently. When the response comes back, the source routing information will be restored and then sent to the Token Ring port.*

## About Source Route-Translational Bridges

Because SmartSwitch 2000, 6000, and 9000 modules have the ability to combine mixed network topologies, yet another hybrid bridge method—called a Source Route Translational bridge (SR-TB)—is used by a number of these SmartSwitch modules.

An SR-TB bridge supports both Source Routing and Transparent bridging capabilities, with the added requirement of maintaining Source Route information across an FDDI interface—either the SmartSwitch 9000 FNB backplane, or an installed FDDI High Speed Interface Module (HSIM).

An SR-TB bridge does this by “translating” the Token Ring physical frame format (by stripping out routing information, if necessary) so that the frame’s source address can be recognized on an FDDI, Ethernet, or ATM segment; then, when data is returned to the source, the bridge restores the necessary route information to forward it along a bridged Token Ring environment.

For data that is restricted to the Token Ring networks available from the SR-TB bridge’s front panel, the bridging method is user-configurable via local management to be Source Route-only (bridged packets must include RIF information and will be source routed; no transparent bridging is enabled), Source Route-Transparent (bridging method will be determined by whether the RII bit is set), or Transparent only (no source routed packets will be bridged). Remote management of these interfaces is based upon their current mode (as set through local management).

For data that will ultimately be sent across an FDDI interface to an ATM, Ethernet, FDDI, or another Token Ring segment, the Routing Information Field will be stripped from the packet so the packet can be transparently bridged onto Ethernet or FDDI media; however, the RIF information as well as the source address of the packet is stored in a RIF cache of the SR-TB bridge. When data is returned to that source address, the SR-TB bridge can look up the address information in its RIF

cache, append the proper Routing Information onto the packet, and then forward the data to the Token Ring segment.

The RIF cache is a software table that can store up to 8192 entries. An SR-TB bridge updates its RIF cache much like a Transparent bridge dynamically updates its Filtering Database: it learns new address information by listening to incoming packets on each port, saves that information to an Address Database, and—if the address was learned to be Source-Route capable—updates routing information for that source address in the RIF cache. Every time a packet arrives from an FDDI interface for a MAC address that is communicating through the SR-TB bridge's front panel, the RIF cache table is searched for an address/RIF match.

There are configuration issues when a Token Ring module receives a packet from an FDDI interface for a destination address that is unknown, and not in its Address Database or RIF cache. You must configure your SR-TB bridge to treat incoming packets with an unknown destination address as either a Source Route or Transparently bridged packet (since Token Ring end stations attached to the module may or may not support Source Routing).

If the bridge is configured to treat an incoming packet with unknown addresses as a Source Routed frame, it will forward it using either STE or ARE frames. If the bridge is configured to treat an incoming packet with an unknown destination as a Transparently bridged frame, it simply forwards the frame.

After a packet with a previously unknown destination has been bridged successfully, and communication begins between the two end nodes, the RIF cache will be updated and packets will be translated as described previously.

## Viewing and Managing Bridging Interfaces

With SPECTRUM Element Manager, you can view and manage each bridging interface supported by your device, including any installed interface modules, such as BRIMs (Bridge/Router Interface Modules) and HSIMs (High Speed Interface Modules).


You can manage your bridge by using the following windows:

- The **Bridge Status** window provides you with basic information about the current status of the device's bridging interfaces, and allow you to enable or disable bridging at each of those interfaces. The Bridge Status window also lets you access further windows to configure bridging at the device. See [The Bridge Status Window, page 2-7](#), for details.
- Bridge statistics—including the **Performance Graph**, **Interface Statistics**, **CSMACD Statistics**, **PPP Link Statistics**, **Dot5 Error Statistics**, and **Source Route Statistics** windows—graphically display the traffic passing between your bridged networks, and let you compare and contrast traffic and errors processed by each interface. See [Bridge Statistics, page 2-15](#) for details.

- The **Spanning Tree** window shows bridge port information and protocol parameters relating to the Spanning Tree Algorithm—the method of determining the controlling bridge when a series of bridges are placed in parallel; see [Spanning Tree](#), page 2-34, for details.
- With the **Filtering Database** window, you can see the contents of the Static and Learned databases—the two address databases which construct the IEEE 802.1 Source Address Table. The bridge uses the contents of these databases to make its packet filtering and forwarding decisions. You can configure entries in these databases to increase bridging efficiency across your network. See [Filtering Database](#), page 2-41, for details.
- The **Ethernet Special Filter Database** and **Token Ring Special Filter Database** windows let you configure special filtering schemes. With these schemes, you can enter filter parameters for a frame based on the contents of its source or destination address field, type field, or data field (with offset)—then specify the bridging action to take place at each port when a frame matching your specifications is encountered (see [Special Filter Databases](#), page 2-47).
- The **I/F Configuration** port-level menu option invokes the Interface Configuration window, which allows you to select a bridging method for a Token Ring bridging interface. This window also allows you to select one of three transmission methods that should be used when unknown addresses are received from end stations attached to the selected bridge port. See [Using the Novell Translation Window](#), page 2-97, for details.
- The **Source Route Configuration** option enables you to configure source routed traffic passing between bridge ports; see [Source Route Configuration](#), page 2-57, for details.
- The **Bridge Configuration** option allows you to set address and routing information for all interfaces on a Token Ring bridging device, including the Bridge Number and the Virtual Ring Number. It also allows you to set source route bridging parameters at the device level. See [Using the Token Ring Bridge and Port Configuration Windows](#), page 2-65, for details.
- The **Port Configuration** option allows you to view the address and routing information for an individual Token Ring bridging interface. This window displays information that is set at the device level via the Bridge Configuration window, such as the Bridge Number and the Virtual Ring Number. It also allows you to set source route bridging parameters for that port. See [Using the Token Ring Bridge and Port Configuration Windows](#), page 2-65, for details.
- The **Duplex Modes** window lists each interface on your device and whether or not it is operating in Full Duplex mode. The window allows you to switch full duplex mode on and off for each interface on the device. Full Duplex Switched Ethernet (FDSE) mode allows the interface to transmit and receive information simultaneously, effectively doubling the available bandwidth. See [Using the Token Ring Bridge and Port Configuration Windows](#), page 2-65, for details.



- The **SmartTrunk** option invokes the SmartTrunk Configuration and Status window, which enables you to group interfaces logically to achieve greater bandwidth between devices that support this feature. There is no limit to the number of ports that can be included in a single trunk, nor is there a limit to the number of trunked instances that can be supported. See [Configuring SmartTrunking, page 2-85](#), for details.
- The **Broadcast Suppression** window enables you to monitor the number of broadcast packets received by each interface of a selected device, and configure the maximum number of broadcast packets that will be forwarded to other interfaces; see [Configuring SmartTrunking, page 2-85](#), for details.
- The **Token Ring Bridge Mode** window lets you select which type of bridging will be used by the Token Ring bridging device—Transparent, Source Routing, or Source Route Transparent; see [Token Ring Bridge Mode, page 2-89](#), for details.
- The **Bridge Translation** window allows you to control the necessary translation of frames that must occur for Token Ring frames to be bridged to Ethernet. This window offers Auto and Dual modes for translation and includes the IBM and SNAP Translation tables. See [Setting Bridge Translation, page 2-91](#), for details.
- The **Novell Translation** window enables you to configure each module port for translation of Novell packets that are received and transmitted across a Token Ring bridge; see [Using the Novell Translation Window, page 2-97](#), for details.



*The menu options that are available will vary depending on the type of device you are monitoring, and on the type of bridge interfaces supported by the device.*

The following sections detail how to use each of the bridge management windows.

## The Bridge Status Window

The Bridge Status window provides you with basic information about the current status of bridging across your device. Color-coding of each port display allows you to quickly ascertain the status of each interface. The Bridge Status window also lets you access further windows to control bridging at your device.

To access the Bridge Status window from the Chassis View window:

1. Click on the **Device** selection in the menu bar. A menu will appear.
2. Click on **Bridge Status**. The Bridge Status window, [Figure 2-1](#), will appear.

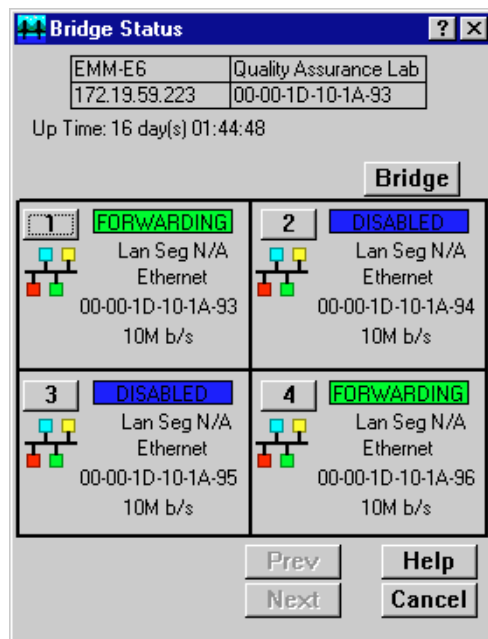



Figure 2-1. The Bridge Status Window



When you first open the Bridge Status window the **Prev** and **Next** buttons will be grayed out, and a message will appear stating that the application is initializing and processing each interface. You will not be able to scroll the display until after **all** the bridging interfaces have been processed. The **Prev** and **Next** buttons are activated when a device supports over four bridge interfaces, so that you can scroll the display to show all interfaces.

The following information is provided by the Bridge Status window for the monitored device as a whole and for each individual bridging interface.

#### Up Time

At the top of the Bridge Status window, you can see the time period (in a days, hours, minutes, seconds format) that has elapsed since the device was last reset or initialized.

### Bridge State on Interface

Indicates the state of bridging over the port interface. Possible bridge states and their corresponding colors are:

- **Forwarding** (green)—The port is on-line and forwarding packets across the bridge from one network segment to another.
- **Disabled** (blue)—Bridging at the port has been disabled by management; no traffic can be received or forwarded on this port, including configuration information for the bridged topology.
- **Listening** (magenta)—The port is not adding information to the filtering database. It is monitoring Bridge Protocol Data Unit (BPDU) traffic while preparing to move to the forwarding state.
- **Learning** (magenta)—The Forwarding database is being created, or the Spanning Tree Algorithm is being executed because of a network topology change. The port is monitoring network traffic, and learning network addresses.
- **Blocking** (orange)—The port is on-line, but filtering traffic from going across the bridge from one network segment to another. Bridge topology information will be forwarded by the port.

### Interface Type

Indicates the interface type which applies to each device bridging port interface (e.g., ethernet). The interface type (*ifType*) is a mandatory object type from the SNMP MIB II Interface (if) Group.

### Bridge Address

Indicates the physical address of the bridge interface.

### Speed

Indicates the speed of the interface in Mbps or Gbps.

## Accessing Bridge Status Window Options

At the top of the Bridge Status window, you can click **Bridge** to access a menu that provides other bridge management options. Depending on which device you are monitoring via SPECTRUM Element Manager, the following bridge management options will be available:

- The **Module Type** window displays a description of the device that is currently being monitored.
- The **Find Source Address** window allows you to discover the bridge interface through which a particular MAC address is communicating; see [Using the Find Source Address Feature](#), page 2-62, for details.
- The **Performance Graph** window displays statistics for traffic across the entire bridge; see [Performance Graphs](#), page 2-15, for details.

- The **Spanning Tree** window allows you to set the Spanning Tree Algorithm parameters for bridging on your device; see [Spanning Tree, page 2-34](#), for details.
- The **SmartTrunk** option invokes the SmartTrunk Configuration and Status window, which enables you to group interfaces logically to achieve greater bandwidth between devices, if both devices support the SmartTrunk feature. There is no limit to the number of ports that can be included in a single “trunk,” nor is there a limit to the number of trunked “instances” that can be supported. See [Configuring SmartTrunking, page 2-85](#), for details.
- The **Filtering Database** window lets you see the contents of the Static and Learned databases—the two address databases which construct the IEEE 802.1 Source Address Table. The bridge uses the contents of these databases to make its packet filtering and forwarding decisions. You can configure the bridge’s acquired and permanent filtering databases to filter or forward traffic across the device. See [Filtering Database, page 2-41](#), for details.
- The **Ethernet Special Filter Database** window lets you configure a special filtering scheme at your bridge. With this scheme, you can enter filter parameters for a frame based on the contents of its source or destination address field, type field, or data field (with offset); you can then specify the bridging action to take place at each port when a frame matching your specifications is encountered. See [Special Filter Databases, page 2-47](#), for details.
- The **Token Ring Special Filter Database** window enables you to define complex filters for transparently-bridged Token Ring frames based upon receive port, source or destination MAC address, Token Ring data type, or data field information (up to 64 bytes); see [Special Filter Databases, page 2-47](#), for details.
- The **Token Ring Bridge Mode** window lets you select which type of bridging will be used by the Token Ring bridging device—Transparent, Source Routing, or Source Route Transparent; see [Token Ring Bridge Mode, page 2-89](#), for details.
- The **Bridge Translation** window allows you to control the necessary translation of frames that must occur for Token Ring frames to be bridged to Ethernet. This window offers Auto and Dual modes for translation and includes the IBM and SNAP Translation tables. See [Setting Bridge Translation, page 2-91](#), for details.
- The **Novell Translation** window enables you to configure each module port for translation of Novell packets that are received and transmitted across a Token Ring bridge; see [Using the Novell Translation Window, page 2-97](#), for details.
- The **Duplex Modes** window allows you to configure duplex mode (on or off) for supporting interfaces on the device; see [Using the Token Ring Bridge and Port Configuration Windows, page 2-65](#), for details.

- **Enable Bridge** and **Disable Bridge** options allow you to administratively activate or deactivate bridging at the device level; see [\(Enabling and Disabling Bridging, page 2-14\)](#), for details.
- The **Bridge Configuration** option opens a window that allows you to set address and routing information for all interfaces on a Token Ring bridging device, including the Bridge Number and the Virtual Ring Number. It also allows you to set source route bridging parameters at the device level. See [Using the Token Ring Bridge and Port Configuration Windows, page 2-65](#), for details.

The individual bridge port index (**1**) menu that you can access from the Bridge Status window will provide the following options, depending on which device you are monitoring through SPECTRUM Element Manager:

- The **Connection Type** window displays a text description of the connection type of the selected bridge interface.
- The **Description** option displays a text description of a bridge interface from the *ifDescr* value of the *ifIndex* related to the selected port.
- The **Performance Graph** window graphically displays the traffic passing between your bridged networks, and lets you compare and contrast traffic processed by each interface; see [Performance Graphs, page 2-15](#), for details.
- The **Source Addressing** window displays the contents of the device's Filtering Database with respect to a selected port. This will display the source MAC addresses that have been detected by the port as it forwards data across the network. The window also lets you set the aging timer that controls how long an inactive MAC address will continue to be stored in the Source Address Database before being aged out. See [Source Route Configuration, page 2-62](#), for details.
- The **PPP Link Status** option invokes the PPP Link Statistics Window, which enables you to view color-coded statistics related to the PPP (Point-to-Point Protocol) link at the selected interface; see [PPP Link Statistics, page 2-25](#), for details.
- The **Source Route Statistics** option opens a window that allows you to view statistics for source routed traffic passing between bridging ports. The window enables you to view the frames that were received, transmitted, and discarded by the bridge. See [Source Route Statistics, page 2-31](#), for details.
- The **I/F Configuration** port-level menu option invokes the Interface Configuration window, which allows you to select a bridging method for a Token Ring bridging interface. This window also allows you to select one of three transmission methods that should be used when unknown addresses are received from end stations attached to the selected bridge port. See [Using the Novell Translation Window, page 2-97](#), for details.
- The **Source Route Configuration** option opens a window that enables you to configure source routed traffic passing between bridging ports; see [Source Route Configuration, page 2-57](#), for details.

- The **Port Configuration** option opens a window that allows you to view the address and routing information for an individual Token Ring bridging interface. This window displays information that is set at the device level via the Bridge Configuration window, such as the Bridge Number and the Virtual Ring Number. It also allows you to set source route bridging parameters for that port. See [Using the Token Ring Bridge and Port Configuration Windows](#), page 2-65, for details.
- The **Dot5 Errors** invokes a window that allows you to view 802.5 statistics for the selected bridging interface on a Token Ring bridging device; see [Dot5 Errors](#), page 2-28, for details.
- The **RMON MAC Layer** option opens the Token Ring Statistics window for Token Ring devices that support RMON, which enables you to view a statistical breakdown of traffic on the monitored Token Ring interface (network segment). Note that if the RMON Default MIB component is disabled, the RMON MAC Layer menu option will launch the Interface Statistics window. Refer to the *RMON User's Guide* for more information on how to enable and disable RMON MIB components. For more information about this menu option, refer to the **Statistics** chapter in the *Remote Monitoring (RMON) User's Guide*, and/or the appropriate device-specific *User's Guide*.
- The **RMON Promiscuous Stats** option opens the Token Ring Promiscuous Statistics window, which allows you to view statistical information about those packets that carry the normal data flow across a bridging interface (network segment). Note that if the RMON Default MIB component is disabled, the RMON Promiscuous Stats menu option will launch the Interface Statistics window. Refer to the *Remote Monitoring (RMON) User's Guide* for more information on how to enable and disable RMON MIB components. For more information about this menu option, refer to the **Statistics** chapter in the *RMON User's Guide*, and/or the appropriate device-specific *User's Guide*.
- The **I/F Statistics** option activates the Interface Statistics Port window, which allows you to view color-coded statistical information about each individual bridge port on the currently monitored device; see [Interface Statistics](#), page 2-20, for details.
- The **Configuration** option opens a window that enables you to configure the selected bridge interface for either full duplex or standard mode; see [Ethernet Port Configuration](#), page 2-72, for details.
- The **Alarm Configuration** option appears as a menu choice for Ethernet devices which support RMON, and invokes the RMON Basic Alarm Configuration window that enables you to create alarms or actions at a specific bridge interface based on rising and falling thresholds for Broadcast/Multicast packets, Kilobits, or Total Errors. Note that if the RMON Default MIB component is disabled, the Alarm Configuration menu option will still appear and the window will still display; however, you will not have the ability to set anything. Refer to the *Remote Monitoring (RMON) User's Guide* for more information on how to enable and disable RMON MIB components. For more

information about this window, refer to the **RMON Alarms and Events** chapter in the *RMON User's Guide*, and/or the appropriate device-specific *User's Guide*.


- The **Statistics** option appears as a menu choice for Ethernet devices which support RMON, and it opens the Ethernet Statistics window, which enables you to view a statistical breakdown of traffic at the monitored Ethernet network segment. Note that if the RMON Default MIB component is disabled, the Statistics menu option will launch the Interface Statistics window. Refer to the *Remote Monitoring (RMON) User's Guide* for more information on how to enable and disable RMON MIB components. For more information about this menu option, refer to the **Statistics** chapter in the *RMON User's Guide*, and/or the appropriate device-specific *User's Guide*.
- The **Sonet/SDH Configuration** window enables you to determine whether any installed FE-100Sx Fast Ethernet Port Interface Modules or APIM-2x ATM Port Interface Modules, both of which provide direct access to SONET (Synchronous Optical Network) networks, will operate according to SONET or SDH (Synchronous Digital Hierarchy) standards; see **SONET/SDH Configuration**, page 2-77, for details.
- The **Sonet Statistics** option opens a window that will let you view some of the statistical information related to any installed FE100-Sx Fast Ethernet Port Interface Modules or APIM-2x ATM Port Interface Modules; see **SONET/SDH Statistics**, page 2-79, for details.
- The **Physical View** option allows you to view the physical state of the Ethernet bridge port through the ETW EtherPhysStatus window and the Token Ring bridge port through the Token Ring Phys Status window when you are monitoring an ETWMIM via SPECTRUM Element Manager; see **Using the Physical View Windows for the ETWMIM**, page 2-99, for details.
- The **CSMACD Stats** option opens a window that enables you to view color-coded statistical information for some Ethernet bridging interfaces, including receive errors, transmission errors, and collision errors. See **CSMACD Statistics**, page 2-23, for details.
- **Enable** and **Disable** options allow you to administratively enable or disable bridging at the selected interface; see **Enabling and Disabling Bridging**, page 2-14, for details.

## Enabling and Disabling Bridging

When you disable a bridge port, you disconnect that port's network from the bridge entirely. The port does not forward any packets, nor does it participate in Spanning Tree operations. Nodes connected to the network can still communicate with each other, but they can't communicate with the bridge and other networks connected to the bridge. When you enable a port, the port moves from the Disabled state through the Learning and Listening states to the Forwarding or Blocking state (as determined by Spanning Tree).

## Enabling and Disabling Individual Interfaces

To disable an individual port interface from the Bridge Status window:

1. Click on the desired **Port** button (  ) to display the port menu.
2. Click on **Enable** to restart bridging on the selected interface, or **Disable** to halt bridging across the selected interface.

To disable an individual port interface from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **Enable** to restart bridging on the selected interface, or **Disable** to halt bridging across the selected interface.

## Enabling and Disabling All Installed Interfaces

To disable bridging across all interfaces installed in a device from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Enable Bridge** to enable bridging across all installed interfaces, or to **Disable Bridge** to disable bridging across all installed interfaces.

To disable bridging across all interfaces installed in a device from the Chassis View window:

1. Click on the **Board Index** of the device of interest; the Board menu will appear.
2. Click on **Enable Bridge** to enable bridging across all installed interfaces, or to **Disable Bridge** to disable bridging across all installed interfaces.



## Bridge Statistics

The following sections describe Statistics windows that are available for the bridge that is being monitored via SPECTRUM Element Manager, both at the device and port levels.

### Performance Graphs

You use Bridge Performance Graphs to view a color-coded strip chart that shows you the traffic being bridged through all networks or an individual network supported by your device. You can configure the display to show frames filtered, forwarded, or transmitted across the device or its individual bridging interfaces, as well as the number of errors experienced at both levels. The graph has an X axis that indicates the 60-second interval over which charting occurs continuously, while its Y axis measures the number of packets or errors that are processed by the device or its bridging interfaces.

You can select the type of errors you wish to monitor by using the available menu buttons. When you click on the error type you wish to view, the name of that error will appear in the button, and the Performance Graph will refresh. The graph will now generate a strip chart based on the newly defined parameters.

At the device level, a **Detail** button on the window allows you to compare the packets forwarded, filtered, or transmitted on all networks supported by the device, as well as errors on all networks.

For a selected bridged network, the **Detail** button allows you to view the number of packets forwarded to, or received from, each other network supported by the device.

To access the device-level Bridge Performance Graph window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Performance Graph**. The Bridge Performance Graph window, [Figure 2-2](#), will appear. (The individual port Bridge Performance Graph windows are similar, except that they display a graph applicable to the selected interface.)

To access the device-level Bridge Performance Graph window from the Chassis View window:

1. Click on the **Board Index** of the device of interest; the Board menu will appear.
2. Click on **Performance Graph**. The Bridge Performance Graph window, [Figure 2-2](#), will appear.

To access the port-level Bridge Performance Graph window from the Bridge Status window:

1. Click on the desired **Port** button (1) to display the Port menu.
2. Click on **Performance Graph**. The Bridge Performance Graph window will appear.

To access the port-level Bridge Performance Graph window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **Performance Graph**. The Bridge Performance Graph window will appear.



The graphic displayed in Figure 2-2 is a device-level window; the window that is displayed at the port level is virtually identical to the one at the device level.

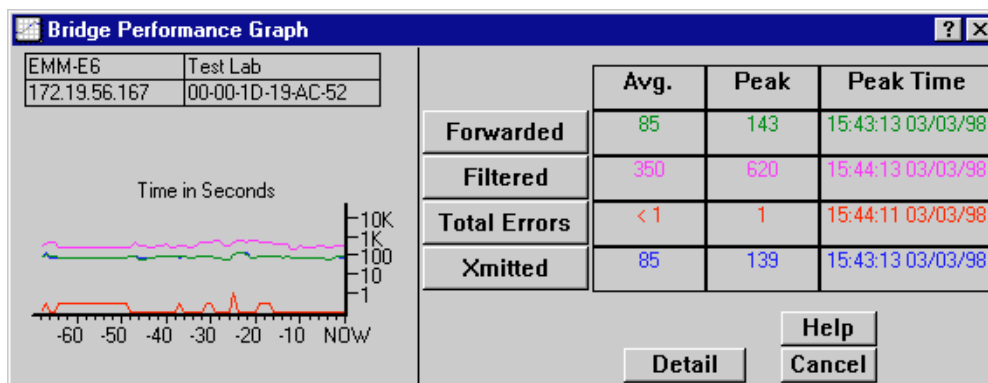


Figure 2-2. Bridge Performance Graph

You can select the following statistics to display in the Bridge Performance Graph or Bridge Port Performance Graph. Statistics are provided numerically (as an average or peak value) and graphically. The device is polled for the graphed information every 2 seconds, and numeric values are updated based on this poll.

The graph updates at the fixed two-second interval. For the first 60 seconds of graphing, you will note the graph lines extending as each interval's data is added to the graph. Once the first 60 seconds has passed, the newest data is added at the right edge of the graph, and the oldest data is scrolled off to the left.

Peak statistics are based on the peak level of activity returned from a single poll since the Performance Graph window was invoked. A date and time is provided for peak levels.

The Average statistics are updated every two seconds as averaged over the previous four poll intervals (i.e., averaged over a sliding eight-second time window).

#### **Frames Forwarded (Green)**

**Forwarded** The number of frames forwarded by an individual port or the device as a whole.

**Nothing** The Frames Forwarded function is currently not measuring any statistics.

#### **Filtered (Magenta)**

**Filtered** The total number of frames filtered by an individual port or the device as a whole.

**Nothing** The Filtered scale is not currently measuring the number of packets filtered by the bridge at the device or port level.

#### **Errors (Red)**

**Total Errors** The total number of errors detected at a single port or on the device as a whole.

**Nothing** The Errors scale is currently not measuring any type of error packets coming through the device or port.

#### **Xmitted (Blue)**

**Xmitted** The total number of frames transmitted by the selected bridge interface, or all bridge interfaces.

**Nothing** The Xmitted scale is not currently measuring the number of packets filtered by the bridge or the individual interface.

### **Configuring Performance Graphs**

To configure the Bridge Performance Graph:


1. Using the mouse, click on **Forwarded** (with green statistics to the right). The Forwarded menu will appear. Click on the desired mode.
2. Click on **Filtered** (with magenta statistics to the right). The Filtered menu will appear. Click on the desired mode.
3. Click on **Total Errors** (with red statistics to the right). The Errors menu will appear. Click on the desired mode.

- Click on **Xmitted** (with blue statistics to the right). The Xmitted menu will appear. Click on the desired mode.

Once you have selected a new mode, it will appear in its respective button, and after the next poll the Performance Graph will refresh and begin to measure using the new mode.

### Bridge Detail Breakdown

The Bridge Detail Breakdown window allows you to compare the number of frames forwarded, filtered, and transmitted on the network segments connected to each interface of your device bridge, as well as the number of errors experienced on each interface.



*The Bridge Detail Breakdown window will not be available if your device has more than 13 bridge ports.*

To access this window from the Bridge performance graph, click on **Detail**. The Bridge Detail Breakdown window, [Figure 2-3](#), will appear.

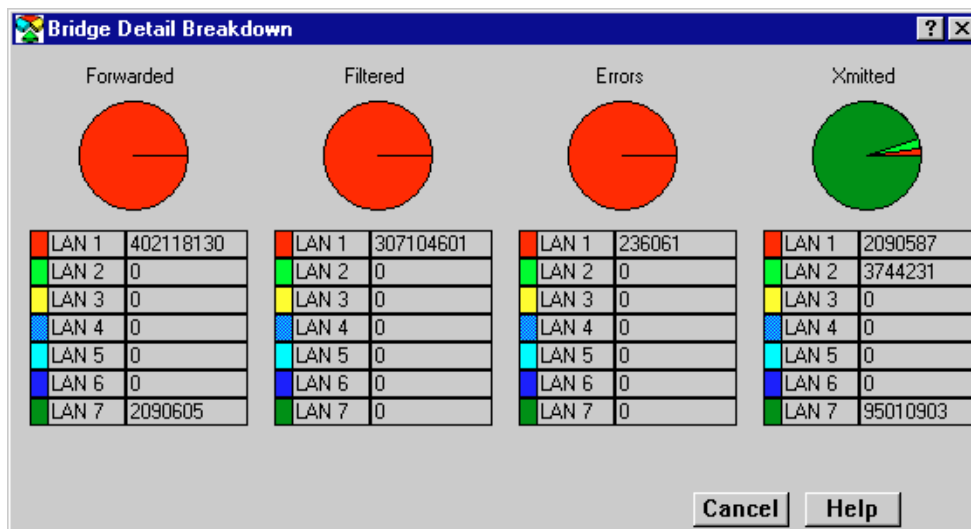


Figure 2-3. The Bridge Detail Breakdown Window

The following information is available for the network segments connected to each of the bridge ports on the device, and any installed BRIM or HSIM port. The information is expressed both numerically and in pie charts. Each port's

network segment has a corresponding color for its statistics or pie chart segments. Depending on your particular bridge and its configuration, the segments are color-coded as follows:

- LAN 1 = light red
- LAN 2 = light green
- LAN 3 = yellow
- LAN 4 = light gray
- LAN 5 = light cyan
- LAN 6 = light blue
- LAN 7 = green
- LAN 8 = red
- LAN 9 = hot pink =
- LAN 10 = light magenta
- LAN 11 = blue
- LAN 12 = cyan
- LAN 13 = black

The values given in these fields are cumulative totals.

#### **Frames Forwarded**

The total number of frames forwarded on each port's network segment, as read from the device after each poll interval.

#### **Filtered**

The total number of frames filtered on each port's network segment, as read from the device after each poll interval.

#### **Errors**

The total number of frames (either inbound or outbound) containing errors which prevented them from being processed by each bridge interface, as reported from the device during the last poll interval.

#### **Xmitted**

The total number of frames transmitted over each port's network segment, as read from the device after each poll interval.

### **Bridge Port Detail Breakdown**

For the selected bridge interface, the Bridge Port Detail Breakdown window allows you to view the number of packets forwarded to or received from each of the other interfaces on your device.

To access the Bridge Port Detail Breakdown window from the port Bridge performance graph, click **Detail**. The Bridge Port Detail Breakdown window, [Figure 2-4](#), will appear.

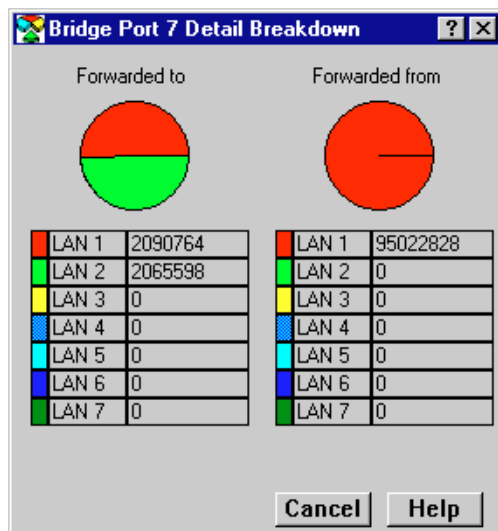


Figure 2-4. The Bridge Port Detail Breakdown Window

The following information is available for each bridge interface on the device. The information is expressed both numerically and in pie charts. The colors corresponding to the forwarding interfaces will vary, depending on which interface is selected.

#### Forwarded to

The number of frames forwarded by the selected bridge interface to each other interface on the bridge, as read from the device after each poll interval.

#### Forwarded from

The total number of frames received by the selected bridge interface from each of the other bridge interfaces, as read from the device after each poll interval.

## Interface Statistics

You can use the interface Statistics window to view color-coded statistical information for each individual bridge port on your device. Statistics are provided for both transmit and receive packets at each port, as well as error and buffering information.

Color-coded pie charts in the middle of the window lets you graphically view statistics for Unicast, Non-Unicast, Discarded and Error packets.

To access the Statistics window from the Bridge Status window:

1. Click on the desired **Port** button ( **1** ) to display the Port menu.
2. Click on **I/F Statistics**. The port I/F Statistics window will appear.

To access the Statistics window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **I/F Statistics**. The port I/F Statistics window will appear.

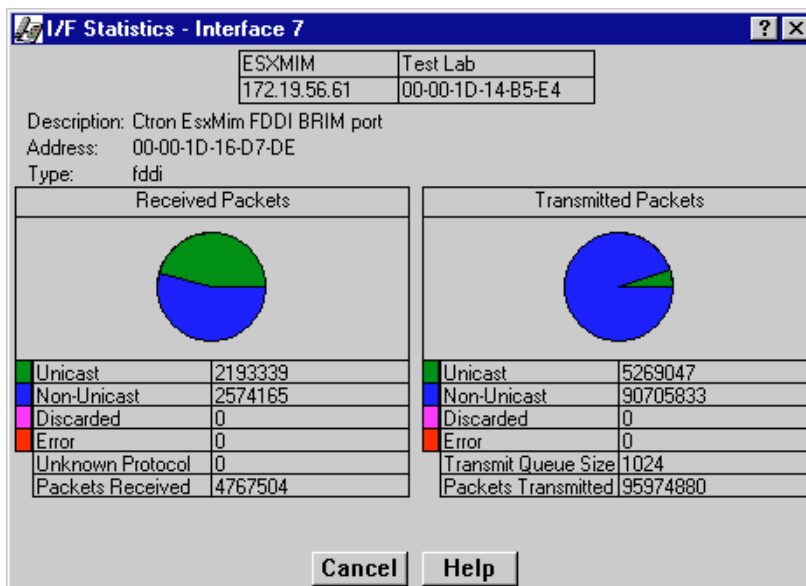


Figure 2-5. I/F Statistics Window

The following three informational fields appear in the upper portion of the window:

#### Description

Displays the interface description for the currently selected port.

#### Address

Displays the MAC (physical) address of the selected port.

#### Type

Displays the interface type of the selected port.



The polling interval is set using the **Device Management** page of the options window accessed via the **Tools**—>**Options** selection from the SPECTRUM Element Manager primary window menu bar. Refer to the **User's Guide** for information on setting device polling intervals.

The following transmit and receive statistics fields are displayed in the lower portion of the window. The first four statistics are also graphically displayed in a pie chart. The statistics are read directly from the device, and are updated with each poll from SPECTRUM Element Manager to the device.

**Unicast**

Displays the number of packets transmitted to, or received from, this interface that had a single, unique source or destination address. These statistics are displayed in the pie chart (color-coded green).

**Non-Unicast**

Displays the number of packets transmitted to, or received from, this interface that had a source or destination address that is recognized by more than one device on the network segment. The non-unicast field includes a count of broadcast packets—those that are recognized by *all* devices on a segment. These statistics are displayed in the pie chart (color-coded dark blue).

**Discarded**

Displays the number of packets which were discarded even though no errors were detected to prevent transmission. One possible reason for discarding such a packet could be to free up buffer space.

The discarding of good packets indicates a very busy network. If a device routinely discards packets, it usually means that network traffic is overwhelming the device. A change in network configuration (such as the addition of a bridge or switch) may help reduce network congestion.

These statistics are displayed in the pie chart (color-coded hot pink).

**Error**

Displays the number of packets received or transmitted that contained errors. These statistics are displayed in the pie chart (color-coded red).

**Unknown Protocol**

Displays the number of packets received which were discarded because of an unknown or unsupported protocol. The device bridge interface will discard the packet and increment this counter if it can't recognize the packet.

**Packets Received**

Displays the number of packets received by this interface.

**Transmit Queue Size**

The number of packets currently queued by the device for transmission from this interface. The amount of device memory devoted to buffer space, and the traffic level on the target network, determine how large the output packet queue can grow before the device begins to discard packets.

**Packets Transmitted**

Displays the number of packets transmitted by this interface.



## CSMACD Statistics

The CSMACD Statistics menu option is available for some Ethernet bridging interfaces. Receive errors, transmission errors, and collision errors are the statistics displayed in this window. Three color-coded pie charts allow you to graphically view the breakdowns of each statistics group.

To access the CSMACD Statistics window from the Bridge Status window:

1. Click on the desired **Port** button ( **1** ) to display the Port menu.
2. Click on **CSMACD Stats**. The device CSMACD Statistics window, [Figure 2-6](#), will appear.

To access the CSMACD Statistics window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **CSMACD Stats**. The device CSMACD Statistics window, [Figure 2-6](#), will appear.

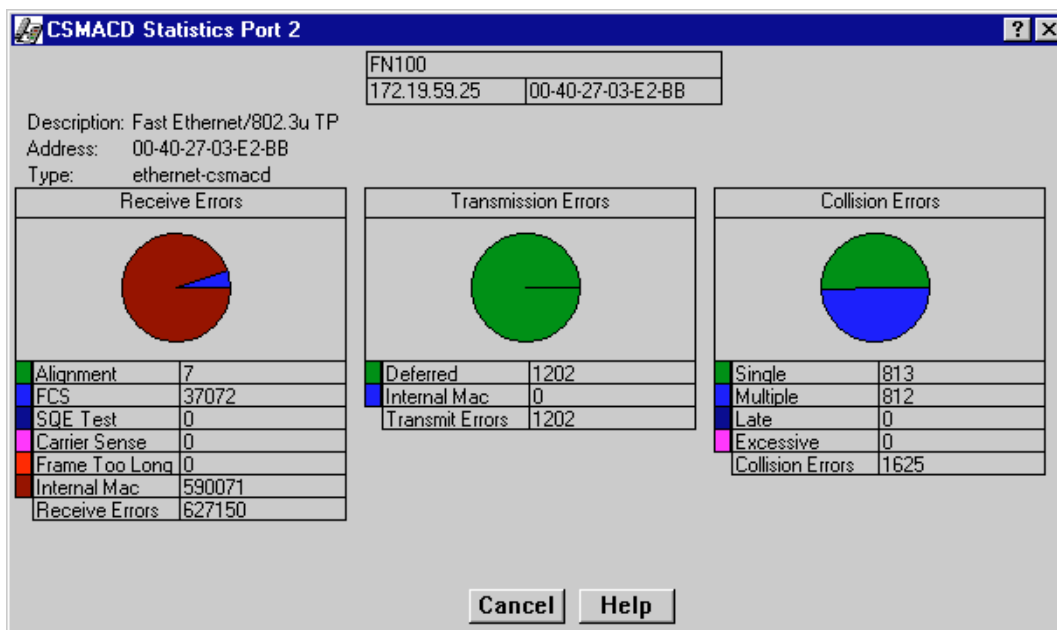


Figure 2-6. CSMACD Statistics Window

Each of the receive, transmission, and collision errors are described in detail below.

### Receive Errors

Indicates the errors detected while the selected interface was receiving a transmission. Possible receive errors are:

- **Alignment**—The number of frames received on a particular interface that contain a nonintegral number of bytes (color-coded green). Misaligned packets can result from a MAC layer packet formation problem, or from a cabling problem that is corrupting or losing data.
- **FCS**—The number of frames received on a particular interface that are an integral number of bytes in length, but do not pass the FCS (Frame Check Sequence) check.

FCS, or Frame Check Sequence, errors occur when packets are somehow damaged in transit. When each packet is transmitted, the transmitting interface computes a frame check sequence (FCS) value based on the contents of the packet, and appends that value to the packet. The receiving interface performs the same computation; if the FCS values differ, the packet is assumed to have been corrupted and is counted as an FCS error.

- **SQE Test**—Displays the number of times that the SQE Test Error message is generated by the PLS sublayer on the selected interface.

The SQE (Signal Quality Error) Test tests the collision detect circuitry after each transmission. If the SQE Test fails, a SQE Test Error is sent to the interface to indicate that the collision detect circuitry is malfunctioning.

- **Carrier Sense**—Displays the number of times that the carrier sense condition was lost or never asserted when attempting to transmit a frame on a particular interface.

Carrier sense describes the action an interface desiring to transmit will take to listen to the communication channel to see if any other interface is transmitting. If a “carrier is sensed,” the sensing interface will wait a random length of time, and then attempt to transmit.

- **Frame Too Long**—Displays the number of frames received on this interface that exceed the maximum permitted frame size.
- **Internal MAC**—The number of frames that could not be received by the interface due to an internal MAC sublayer receive error. These errors are only counted if a Frame Too Long, Alignment, or FCS Error did not occur along with the internal MAC error.
- **Receive Errors**—Displays the total number of receive errors of all types that were detected by the selected interface while it was receiving a transmission.

#### **Transmission Errors**

Indicates the errors that occurred while the selected interface was attempting to transmit frames. Possible transmission errors are:

- **Deferred**— Displays the number of frames for which the first transmission attempt on this interface is delayed because the medium is busy.

- **Internal MAC**—The number of frames for which transmission fails due to an internal MAC sublayer transmit error. This error is only counted in this window if there have not been corresponding Late Collisions, Excessive Collisions, or Carrier Sense Errors.
- **Transmit Errors**—The total number of transmission errors of all types that occurred while the selected interface was attempting to transmit frames.

### Collision Errors

Indicates the collision errors that occurred during transmission from this interface. Possible collision errors are:


- **Single**—Displays the number of transmitted frames on the selected interface for which transmission was prevented by **one** collision.
- **Multiple**—Displays the number of transmitted frames on the selected interface for which transmission was prevented by **more than one** collision.
- **Late**—Displays the number of times that a collision has been detected on this interface later than 51.2 microseconds into the transmission of the packet on a 10 Mbit/s system or later than 5.12 microseconds on a 100 Mbit/s system.
- **Excessive**—Displays the number of transmitted frames on the selected interface for which transmission was prevented by excessive collisions.
- **Collision Errors**—Displays the total number of collision errors of all types that occurred during transmission from this interface.

## PPP Link Statistics

The PPP Link Status option opens the PPP Link Statistics window, which enables you to view color-coded statistics related to the PPP (Point-to-Point Protocol) link at the selected interface.

The Point-to-Point Protocol is a standard method of transporting multiprotocol datagrams over point-to-point links. A PPP Link provides full-duplex communication between the endpoints, allowing a simultaneous bidirectional operation that should maintain the order in which data packets are transmitted.

To access the PPP Link Statistics window from the Bridge Status window:

1. Click on the desired **Port** button (  ) to display the port menu.
2. Click on **PPP Link Status**. The PPP Link Statistics window, [Figure 2-7](#), will appear.

To access the PPP Link Statistics window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **PPP Link Status**. The PPP Link Statistics window, [Figure 2-7](#), will appear.

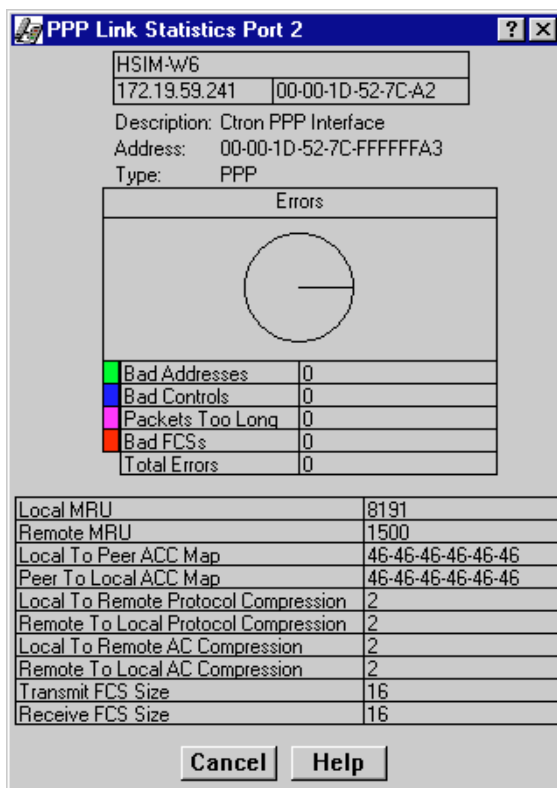


Figure 2-7. PPP Link Statistics Window

Each of the errors and statistics related to the PPP Link at the selected bridging interface is described in detail below.

### Errors

Indicates the errors that occurred which relate to the PPP Link at the selected bridging interface. Possible error types are:

- **Bad Addresses**—The Bad Addresses field displays the number of packets received with an incorrect Address field.
- **Bad Controls**—The Bad Controls field displays the number of packets received on the selected interface that have an incorrect Control field.
- **Packets Too Long**—The Packets Too Long field displays the number of received packets that were discarded because their length exceeded the MRU (Maximum Receive Unit). Note that packets that are longer than the MRU and that are successfully received and processed are not included in the count.
- **Bad FCSs**—The Bad FCSs field displays the number of received packets that were discarded due to having an incorrect FCS (Frame Check Sequence) value.

- **Total Errors**—The Total Errors field displays the total number of errors of all types: Bad Addresses, Bad Controls, Packets Too Long, and Bad FCSs.

### Statistics

Lists statistics fields which are related to the PPP Link at the selected bridging interface. Possible statistics fields are:

- **Local MRU**—The Local MRU field displays the current value of the MRU (Maximum Receive Unit) for the local PPP entity. This value is the MRU that the remote entity uses when sending packets to the local PPP entity. The MRU is the maximum length of data information (included “padded” data octets, but excluding the Protocol field which identifies the datagram’s protocol type) that can be received by this interface. The default MRU size is 1500 octets. The auto-negotiation process may establish another value for MRU if consent is given at both ends of the PPP link (if either the local or remote PPP entity informs the other that larger packets can be sent, or requests that smaller packets be sent).
- **Remote MRU**—The Remote MRU field displays the current value of the MRU (Maximum Receive Unit) established for the remote interface at the other end of the PPP Link. This value is the MRU that the local entity uses when sending packets to the remote PPP entity.
- **Local to Peer ACC Map**—The Local to Peer ACC Map field displays the current value of the Asynchronous Control Character (ACC) Map used for sending packets from the local PPP entity to the remote PPP entity. In effect, this is the ACC Map that is required to ensure that all characters can be successfully transmitted through the local modem. The actual ACC Map used on the transmit side of the link will be a combination of the local node’s *pppLinkConfigTransmitACCMAP* and the remote node’s *pppLinkConfigReceiveACCMAP*.
- **Peer to Local ACC Map**—The Peer to Local ACC Map field displays the Asynchronous Control Character (ACC) Map used by the remote PPP entity when transmitting packets to the local PPP entity. In effect, this is the ACC Map that is required to ensure that the local modem will successfully receive all characters. The actual ACC Map used on the receive side of the link will be a combination of the local node’s *pppLinkConfigReceiveACCMAP* and the remote node’s *pppLinkConfigTransmitACCMAP*.
- **Local to Remote Protocol Compression**—The Local to Remote Protocol Compression field determines whether or not the local PPP entity uses Protocol Compression when transmitting packets to the remote PPP entity.
- **Remote to Local Protocol Compression**—The Remote to Local Protocol Compression field determines whether or not the remote PPP entity uses Protocol Compression when transmitting packets to the local PPP entity.
- **Local to Remote AC Compression**—The Local to Remote AC Compression field determines whether or not the local PPP entity uses Address and Control (AC) Compression when transmitting packets to the remote PPP entity.

- **Remote to Local AC Compression**—The Remote to Local AC Compression field determines whether or not the remote PPP entity uses Address and Control (AC) Compression when transmitting packets to the local PPP entity.
- **Transmit FCS Size**—The Transmit FCS Size field displays the size of the Frame Check Sequence (FCS), in bits, that the local node generates when sending packets to the remote node.
- **Receive FCS Size**—The Receive FCS Size field displays the size of the Frame Check Sequence (FCS), in bits, that the remote node generates when sending packets to the local node.

## Dot5 Errors

The **Dot5 Errors** menu option invokes the Station Statistics window, which enables you to view IEEE 802.5 error statistics reported for a Token Ring bridge interface.

To access the Station Statistics window from the Bridge Status window:

1. Click on the desired **Port** button ( **1** ) to display the port menu.
2. Click on **Dot5 Errors**. The Station Statistics window, [Figure 2-8](#), will appear.

To access the Station Statistics window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **Dot5 Errors**. The Station Statistics window, [Figure 2-8](#), will appear.

IETF dot5 Errors	
Line Errors	0
Burst Errors	0
A.C. Errors	0
Abort Sequences	0
Internal Errors	0
Lost Frames	0
Congestion Errors	0
F.C. Errors	0
Token Errors	0
Soft Errors	0
Hard Errors	0
Signal Loss	195285
Transmit Beacons	0
Recoveries	0
Lobe Wires	0
Removes	0
Singles	0
Frequency Errors	0

Figure 2-8. Dot5 Errors Statistics Window

Each type of IEEE 802.5 error detected by the selected station port is described in detail below.

#### Line Errors

The Line Errors field displays the number of the line errors detected by the selected port. This error indicates the presence of a non-data bit between the starting and ending delimiters of data or a frame check sequence (FCS) error.

#### Burst Errors

The Burst Errors field displays the number of burst errors detected by the selected port. This error indicates a bit information encoding error when there are no transitions between 0 and 1 over five half-bit times.

#### A. C. Errors

The A. C. Errors field displays the number of A. C. errors detected by the selected port. These errors count protocol data units (PDUs) that contain errors in the A or C bits.

### **Abort Sequences**

The Abort Sequences field displays the number of abort sequences transmitted by the selected port. These occur when an adapter has frames to transmit and receives a token, but does not detect an ending delimiter on the token after its access control field. This indicates that the token is corrupted. The station transmits abort delimiters to halt frame transmission before its expected end-frame sequence, re-queues the frame(s) for transmission, but does not release the corrupt token.

### **Internal Errors**

The Internal Errors field displays the number of recoverable internal errors detected by the selected port. These are recorded when a station recognizes a recoverable internal error in its adapter, and removes itself from the ring. This is considered a solid hard error, since the ring automatically reconfigures.

### **Lost Frames**

The Lost Frames field displays the number of non-returned frames detected by the selected port. These occur when a transmitting station's TRR (Timer, Return to Repeat) timer expires (after 4.1 milliseconds) before the end of its frame returns. This timer ensures that the station returns to the data repeat state (i.e., retrieves the token, strips it of data, and issues a new token to the ring). Lost frames are usually caused by a station entering or leaving the ring as the frame is circulating.

### **Congestion Errors**

The Congestion Errors field displays the number of times the selected port has not been able to copy a protocol data unit (PDU) addressed to it because of a lack of internal buffering.

### **F. C. Errors**

The F. C. Errors field displays the number of protocol data units (PDUs) addressed to the selected station with the A bits already set to 1. This error indicates that a possible electrical line disturbance or a duplicate address has occurred on the ring.

### **Token Errors**

The Token Errors field displays the number of times that the selected station, acting as the active monitor, detected an error condition that needed a token transmitted.

### **Soft Errors**

The Soft Errors field displays the number of soft errors detected by the selected port. Although soft errors do not cause ring failure, they degrade the performance of the ring network.

### **Hard Errors**

The Hard Errors field displays the number of immediately recoverable fatal errors detected by the selected port. These are errors which halt normal operation of the ring, and are usually caused by faults in the ring hardware, equipment, or wiring.



**Signal Loss**

The Signal Loss field displays the number of times that the selected port has detected the loss of a signal condition from the ring.

**Transmit Beacons**

The Transmit Beacons field displays the number of beacon frames transmitted by the selected station.

**Recoveries**

The Recoveries field displays the number of claim token frames the monitored station has received or transmitted after a ring purge frame.

**Lobe Wires**

The Lobe Wires field displays the number of open or short circuits detected in the lobe data path.

**Removes**

The Removes field displays the number of Remove Ring Station MAC frame requests detected by the selected port.

**Singles**

The Singles field displays the number of times the selected station has detected that it is the only station on the ring. This error may indicate that the station is the first on the ring or that there is a hardware problem.


**Frequency Errors**

The Frequency Errors field displays the number of times that the selected station detected a larger-than-allowed difference between the incoming frequency and the expected frequency.

## Source Route Statistics

The **Source Route Statistics** menu option invokes the Bridge Source Routing window, which allows you to compare the statistics on frames received, transmitted, and discarded at the Token Ring interfaces of devices that are bridging from a source routing network to a transparent network.

To access the Bridge Source Routing window from the Bridge Status window:

1. Click on the desired **Port** button (  ) to display the port menu.
2. Click on **Source Route Statistics**. The Bridge Source Routing window, [Figure 2-9](#), will appear.

To access the Bridge Source Routing window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **Source Route Statistics**. The Bridge Source Routing window, [Figure 2-9](#), will appear.

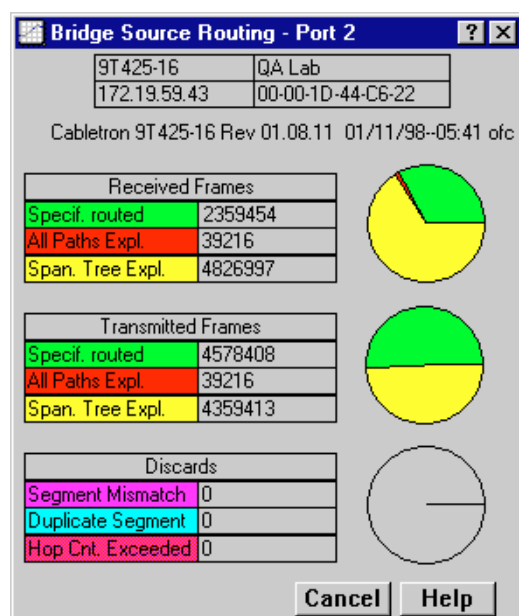


Figure 2-9. The Bridge Source Routing Window

The Bridge Source Routing window provides basic statistics for source routed traffic passing between the bridging ports. Pie charts graphically break down the statistical information. All statistics are calculated since the device was last reset or powered up. The following window fields are listed in the Bridge Source Routing window:

### Received Frames

Frame types received by the bridge ports

- **Specif. Routed**—Displays the total number of Specifically Routed Explorer frames received by the indicated port from its attached segment.

These frames have data and routing information and are following a known route from source to destination.

- **All Paths Expl.**—Displays the total number of All Path Explorer frames received by the indicated port from its attached segment.

When a sending station needs to determine the best route to an intended destination, it transmits an All Paths Explorer (APE) frame. The APE frame contains no routing information; it is propagated along all available paths to

the destination station, which then directs a reply back to the source. The first reply received by the original sending station is considered the most efficient route and is used in subsequent transmissions.

- **Span. Tree Expl.**—Displays the total number of Spanning Tree Explorer (STE) frames received by the indicated port from its attached segment. STE frames, also known as Single Route Broadcast frames, follow the topology established by the Spanning Tree Algorithm.

### Transmitted Frames

List of Frame types transmitted by the bridge ports

- **Specif. Routed**—Displays the total number of Specifically Routed Frames transmitted by the indicated port onto its attached segment.
- **All Paths Expl.**—Displays the total number of All Path Explorer frames transmitted by the indicated port onto its attached segment.
- **Span Tree Expl.**—Displays the total number of Spanning Tree Explorer (STE) frames transmitted by the indicated port onto its attached segment.

### Discards

List of Frames discarded by the bridge ports.

- **Segment Mismatch**—Displays the number of explorer frames discarded because their routing descriptor field contained an invalid value for a segment attached to the port.

The routing information field of a Specifically Routed frame contains LAN Segment In (Ring In)–Bridge Number–LAN Segment Out (Ring Out) information. If the bridge's LAN Segment Out value does not match the LAN Segment Out specified in the frame's Routing Information Field, the bridge logs a Segment Mismatch and discards the frame.

- **Duplicate Segment**—Displays the number of frames discarded because the frame's Routing Information Field identifies a particular segment more than once.
- **Hop Cnt. Exceeded**—Displays the number of All Paths Explorer frames discarded at the specified port because they exceeded the number of routing descriptors (bridge hops) specified by the Hop Count Limit.

## Spanning Tree

The Bridge Spanning Tree window allows you to display and modify the device's bridge port information and protocol parameters relating to the Spanning Tree Algorithm.

In a network design with multiple bridges placed in parallel (i.e, attached to the same LAN), data loops must be prevented. The Spanning Tree Algorithm (STA) is the method that bridges use to communicate with each other to ensure that only a single data route exists between any two end stations.

On a LAN interconnected by multiple bridges, Spanning Tree selects a controlling Root Bridge and Port for the entire bridged LAN, and a Designated Bridge and Port for each individual LAN segment. A Designated Port/Bridge for a LAN segment forwards frames from that LAN towards the Root Bridge, or from the Root Bridge onto the LAN. All other bridge ports attached to that LAN are configured to filter (block) frames.

When data passes from one end station to another across a bridged LAN, it is forwarded through the Designated Bridge/Port for each LAN segment towards the Root Bridge, which in turn forwards frames towards Designated Bridges/Ports on its opposite side.

During the Root Bridge Selection process, all bridges on the network communicate STA information via Bridge Protocol Data Units (BPDUs). With BPDUs, all network bridges collectively determine the current network topology and communicate with each other to ensure that the topology information is kept current.

To access the Bridge Spanning Tree window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Spanning Tree**. The Bridge Spanning Tree window, [Figure 2-10](#), will appear.

To access the Bridge Spanning Tree window from the Chassis View window:

1. Click on the **Board Index** of the device of interest; the Board menu will appear.
2. Click on **Spanning Tree**. The Bridge Spanning Tree window, [Figure 2-10](#), will appear.

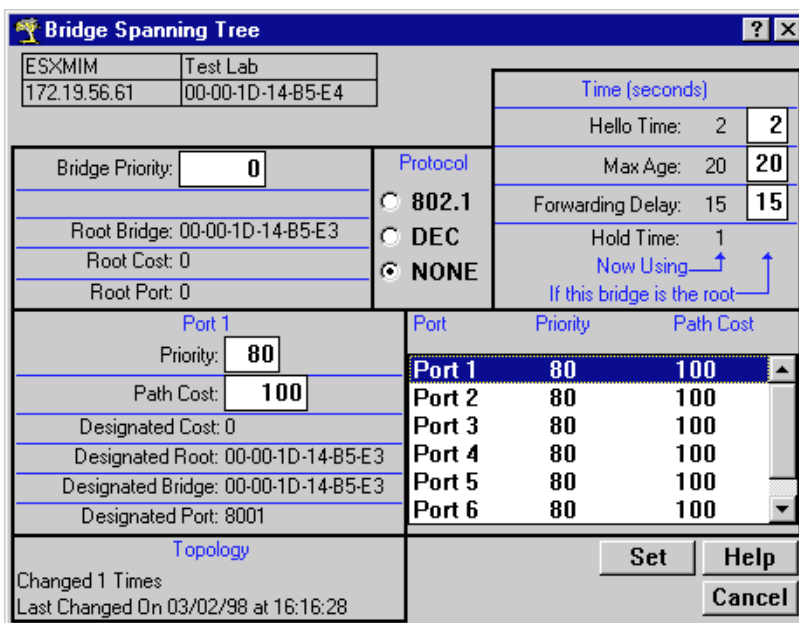


Figure 2-10. Bridge Spanning Tree Window

The Bridge Spanning Tree window displays STA parameters and allows you to alter parameters for the device bridge as a whole, and for each individual bridging interface.


The values displayed apply to the currently-selected bridging interface, highlighted in the lower right quadrant of the window. To view or alter the parameters of another interface, click on the appropriate **Port X** name listed in the quadrant.

## Bridge Level

### Bridge Priority

This field displays the “priority” component of the device’s unique bridge identifier. The Spanning Tree Algorithm assigns each bridge a unique identifier, which is derived from the bridge’s MAC address and the Priority. The bridge with the lowest value of bridge identifier is selected as the Root. A lower priority number indicates a higher priority; a higher priority enhances a bridge’s chance of being selected as the Root.

You can edit this text box to change network topology, if needed. The default value is 8000; the range is 0—FFFF hexadecimal.



*Part of a bridge's Identifier is based on its MAC address. In most network installations, performance differences between bridges may be negligible. You may, however, find your data bottle-necked in installations where both a low-performance bridge and a high-performance bridge are attached to the same LAN segment and the two (or more) bridges have the same Priority component set (e.g., at the default 8000 Hex). In such a scenario you may want to alter the Priority component of the higher performance bridge to ensure that it becomes root for the segment (or overall root). Remember, if Priority components are equal, the bridge on the segment with the lowest MAC address would have a better chance of being selected as the root bridge—as it would have a lower Bridge Identifier. If your bridges come from multiple vendors, they will have different MAC address values (e.g., Cabletron devices have a lower MAC address than 3Com devices); if they come from the same vendor, the bridge with the earlier manufacture date will have the lower MAC address value.*

**Root Bridge**

Displays the MAC address of the bridge that is currently functioning as the Root Bridge.

**Root Cost**

Indicates the cost of the data path from this bridge to the Root Bridge. Each port on each bridge adds a “cost” to a particular path that a frame must travel. For example, if each port in a particular path has a Path Cost of 1, the Root Cost would be a count of the number of bridges along the path. (You can edit the Path Cost of bridge ports as described later.) The Root Bridge's Root Cost is 0.

**Root Port**

This field displays the identifier (the physical index number) of the device bridge port that has the lowest cost path to the Root Bridge on the network. If the device is currently the Root Bridge, this field will read 0.

**Protocol**

Displays the Spanning Tree Algorithm Protocol type the device is currently using. The choices are:

- 802.1
- DEC (DEC Lanbridge 100)
- None

The following four fields display values used for various Spanning Tree timers that are set at the Root Bridge and this bridge. In Spanning Tree operations, the value used for the tree is the one set at the Root Bridge (with the exception of Hold Time, which is a fixed value), but you can change the value for each bridge on your network in the event that it becomes Root.

**Hello Time**

This parameter indicates, in seconds, the length of time the Root Bridge (or bridge attempting to become the Root) waits before resending Configuration BPDUs. The range for this field is 1 to 10 seconds, with a default value of 2 seconds. The Root Bridge sets the Hello Time.

**Max Age**

This parameter displays the bridge's BPDU aging timer. This controls the maximum time a BPDU can be retained by the bridge before it is discarded. During normal operation, each bridge in the network receives a new Configuration BPDU before the timer expires. If the timer expires before a Configuration BPDU is received, it indicates that the former Root is no longer active. The remaining bridges begin Spanning Tree operation to select a new Root. The current Root Bridge on the network sets the Max Age time. The range for this field is 6 to 40 seconds, with a default value of 20 seconds.

**Forwarding Delay**

This parameter displays the time period which elapses between states while the bridge is moving to the Forwarding state. For example, while moving from a Blocking to a Forwarding state, the port first moves from Blocking to Listening to BPDU activity on the network, remains there for the Forward Delay period, then moves to the Learning State (and remains in it for the Forward Delay period), and finally moves into a Forwarding state. This timer is set by the Root Bridge. During a topology change, the Forward Delay is also used as the Filtering Database Aging Time, which ensures that the Filtering Database maintains current topology information.

**Hold Time**

This parameter displays, in seconds, the minimum time that can elapse between the transmission of Configuration BPDUs through a bridge port. The Hold Time ensures that Configuration BPDUs are not transmitted too frequently through any bridge port. Receiving a BPDU starts the Hold Timer. After the Hold Timer expires, the port transmits its Configuration BPDU to send configuration information to the Root. The Hold Time is a fixed value, as specified by the IEEE 802.1d specification.

## Bridge Port Level

The following fields are applicable to each bridge port on the device.

**Priority**

If two or more ports on the same bridge are connected to the same LAN segment, they will receive the same Root ID/Root Cost/Bridge ID information in Configuration BPDUs received at each port. In this case, the BPDU's Port ID information—the transmitting port's identifier and its manageable Priority component—is used to determine which is the Designated Port for that segment.

A lower assigned value gives the port a higher Priority when BPDUs are compared. The allowable range is 0—FF hexadecimal (0—255 decimal); the default is 80 hexadecimal.

**Path Cost**

Displays the cost that this port will contribute to the calculation of the overall Root path cost in a Configuration BPDU transmitted by this bridge port. You can lower a port's Path Cost to make the port more competitive in the selection of the Designated Port—for example, you may want to assign a lower path cost to a port on a higher performance bridge. The allowable range is 1 to 65,535.

**Designated Cost**

Displays the cost of the path to the Root Bridge of the Designated Port on the LAN to which this port is attached. This cost is added to the Path Cost to test the value of the Root Path Cost parameter received in Configuration BPDUs.

**Designated Root**

Displays the unique bridge identifier of the bridge that is assumed to be the Root Bridge.

**Designated Bridge**

Displays the network address portion of the Bridge ID (MAC address/priority component) for the bridge that is believed to be the Designated Bridge for the LAN associated with this port.

The Designated Bridge ID, along with the Designated Port and Port Identifier parameters for the port, is used to determine whether this port should be the Designated Port for the LAN to which it is attached. The Designated Bridge ID is also used to test the value of the Bridge Identifier parameter in received BPDUs.

**Designated Port**

Displays the network address portion of the Port ID (which includes a manageable priority component) of the port believed to be the Designated Port for the LAN associated with this port.

The Designated Port ID, along with the Designated Bridge and Port Identifier parameters for the port, is used to determine whether this port should be the Designated Port for the LAN to which it is attached. Management also uses it to determine the Bridged LAN topology.


**Topology**

This indicates how many times the bridge's Topology Change flag has been changed since the device was last powered up or initialized. It also indicates the time elapsed since the topology last changed. The Topology Change flag increments each time a bridge enters or leaves the network, or when the Root Bridge ID changes.



## Configuring Spanning Tree

The Bridge Spanning Tree window allows you to update the following parameters for your device bridge. When you have finished making changes to the following individual parameters, you must click on **Set** at the bottom of the Spanning Tree window to write the changes to the device.



*Any values you set at the bridge will cause a Topology Change flag to be issued in the next Configuration BPDUs it transmits. This will cause the bridged network to immediately recalculate Spanning Tree and change topology accordingly.*

### Changing Bridge Priority

To change the part of the bridge address that contains the identifier used in the Spanning Tree Algorithm for priority comparisons:

1. Highlight the **Bridge Priority** field.
2. Enter the new identifier, in hexadecimal format; the allowed range is 0-FFFF hexadecimal.
3. Click on **Set**.

The selected Bridge Priority will be applied to the bridge (a lower number indicates a higher priority in the root selection process).

### Changing the Spanning Tree Algorithm Protocol Type

To change the type of protocol used in Spanning Tree:

1. Click the mouse on the appropriate option button: **802.1**, **DEC**, or **None**.
2. Click on **Set**.

The selected Spanning Tree Algorithm protocol type will be applied to the bridge. If you selected None, the Spanning Tree Algorithm will be disabled (if it already was enabled). If STA Protocol Type was changed from None to IEEE 802.1 or DEC, you must restart the bridge for the newly selected STA protocol to be applied.



*All bridges in a network must use the same Spanning Tree version. Mixing Spanning Tree Algorithm protocols will cause an unstable network.*

### Changing Hello Time

If the bridge is the Root Bridge, or is attempting to become the Root, and you want to change the length of time the bridge waits between sending configuration BPDUs:

1. Highlight the **Hello Time** field, and type in a new value.
2. Click on **Set**.

The IEEE 802.1d specification recommends that Hello Time = 2 seconds, with an allowable range of 1 to 10 seconds.

### Changing Max Age Time

If the device is the Root Bridge or attempting to become the Root, and you want to change the maximum time that bridge protocol information will be kept before it is discarded:

1. Highlight the **Max Age** field, and type in a new value.
2. Click on **Set**.


The IEEE 802.1d specification recommends that Max Age = 20 seconds, with an allowable range of 6 to 40 seconds.

### Changing Forwarding Delay Time

If the device is the Root Bridge or attempting to become the Root, and you want to change the time period the bridge will spend in the Listening state (e.g. either listening to BPDU activity on the network while moving from the Blocking to the Learning state or in the Learning state while the bridge is moving from the Listening to the Forwarding state):

1. Highlight the **Forwarding Delay** field, and type in a new value.
2. Click on **Set**.

The IEEE 802.1d specification recommends that Forward Delay = 15 seconds, with an allowable range of 4 to 30 seconds.



*To ensure proper operation of the Spanning Tree Algorithm, the IEEE 802.1d specification recommends that you always observe the following relationship between Forwarding Delay, Max Age, and Hello Time:*

$$2 \times (\text{Forwarding Delay} - 1.0) \geq \text{Max Age} \geq 2 \times (\text{Hello Time} + 1.0)$$

## Changing Port Priority

To change the part of the Port Priority used in priority comparisons:

1. If necessary, select the desired port by clicking the mouse to highlight the port in the lower right quadrant of the window. The lower left quadrant of the window will now allow you to edit parameters for the selected port.
2. Highlight the port **Priority** field, and enter the new priority identifier. Only valid hexadecimal numbers (0 to FF) are allowed in this field. The default is 80 hexadecimal.
3. Click on **Set**. The new port priority will be saved.

## Changing Path Cost

To change the Path Cost:

1. If necessary, select the desired port by clicking the mouse to highlight the port in the lower right quadrant of the window. The lower left quadrant of the window will now allow you to edit parameters for the selected port.
2. Highlight the **Path Cost** field, and type in a new value from 1 to 65535 decimal (default is 100 decimal).
3. Click on **Set**.

The new path cost will be applied to the port.

# Filtering Database

The Filtering Database, which makes up the IEEE 802.1 Source Address Table, is used to determine which frames will be forwarded or filtered across the device's bridging ports.

During initialization, the bridge copies the contents of its Permanent Database to the Filtering Database. Next, the bridge learns network addresses by entering the source address and port association of each received packet into the Filtering Database. When in the Forwarding state, the bridge examines each received packet, checks it against the Special Database (refer to **Special Filter Databases**, page 2-47), and then (if no special filtering applies) compares the destination address to the contents of the Filtering Database.

If the destination address is located on the network from which the packet was received, the bridge filters (does not forward) the packet. If the destination address is located on a different network, the bridge forwards the packet to the appropriate network. If the destination address is not found in the Filtering Database, the bridge forwards the packet to all networks. To keep Filtering Database entries current, older entries are purged after a period of time, which is called the Dynamic Aging Time.

The Filtering Database consists of two separate databases: the Static and the Learned Databases.

The **Static Database** contains addresses that are entered by a network administrator. You add these addresses directly to the database while the bridge is powered up, or to the device's battery-backed RAM so that they are stored on shutdown until the next power-up.


The **Learned Database** consists of addresses that accumulate as part of the bridge's learning process as it is up and running. These do not remain in the Source Address Table when the system is shut down. The Learned Database also contains the addresses that are in the Static Database upon start-up of the bridge.

Entries to the Source Address Table are one of four types: **Permanent**, **Static**, **Dynamic**, or **Learned**.

- **Permanent** entries are addresses that you add to the Static Database (via the Filtering Database window) that are stored in the device's battery-backed RAM. Since they remain in the device on shutdown or restart, they are considered "Permanent."
- **Static** entries are addresses that you add to the Static Database (via the Filtering Database window). These entries remain in the device until it is shut down.
- **Dynamic** entries are addresses that you add to the Static Database (via the Filtering Database window). With the Aging Time feature, you set the time period that these addresses are saved in the Source Address Table. Addresses that have not transmitted a packet during one complete cycle of the aging timer are deleted from the database.
- **Learned** entries are addresses that are added to the Learned Database through the bridge's learning process. With the Aging Time feature, you set the time period that these addresses are saved in the Source Address Table. Addresses which are inactive within a cycle of the aging timer are dropped from the database.

Learned address entries are divided into two types, **Learned** and **Self**. Address entries classified as **Learned** have transmitted frames destined for a device attached to a device port's connected segment. Address entries classified as **Self** are those that have sent a frame with a destination address of one of the device's bridging ports.

At the Filtering Database window (Figure 2-11, page 2-43), you can view the number of entries of each type: Permanent, Static, Dynamic, or Learned.



*Even though new entries into the Filtering Database are added as Static entries by default, note that some devices do not support Static entries. For these devices, once you add an entry into the Filtering Database, it must be changed to a Permanent type before clicking on **OK** to apply the change. If the entry is not changed to a Permanent type before clicking on **OK**, you will receive a Set Failed message.*

A scrollable Address Entry panel allows you to:

- View the address entries in the Filtering Database.
- Alter an entry's type (e.g., from Learned to Permanent, Dynamic, or Static).
- View and configure the bridging action taking place on the packets entering each of the bridging ports.

In addition, you can use buttons to add individual addresses to, or delete them from, these databases, or clear all Permanent, Static, or Dynamic entries in the database.

To access the Filtering Database window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Filtering Database**. The Filtering Database window, [Figure 2-11](#), will appear.

To access the Filtering Database window from the Chassis View window:

1. Click on the **Board Index** of the device of interest; the Board menu will appear.
2. Click on **Filtering Database**. The Filtering Database window, [Figure 2-11](#), will appear.

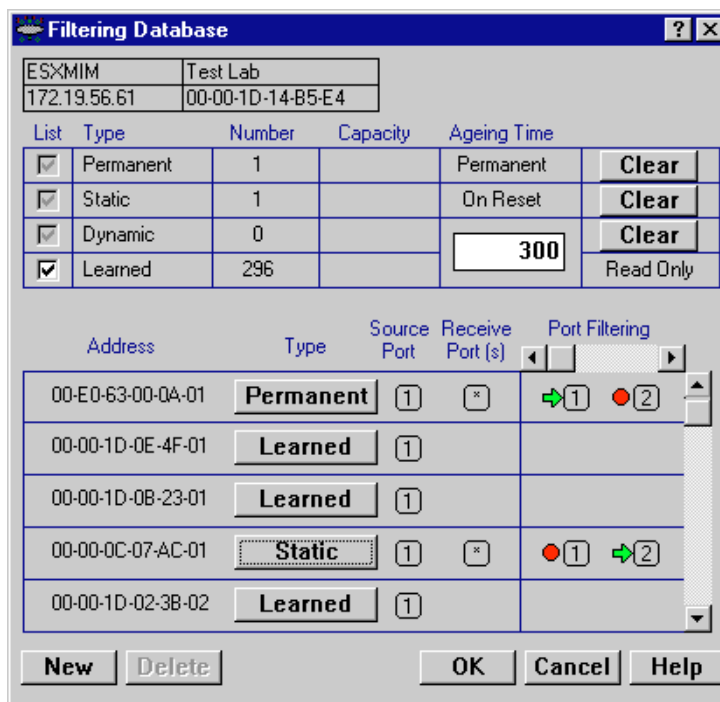


Figure 2-11. The Filtering Database Window

The following fields are listed in the top portion of the Filtering Database window.

**List**

The List checkboxes indicate whether the associated entry type (Permanent, Static, Dynamic, or Learned) will be displayed in the scrollable table of address entries. A check next to the entry type indicates that it will be displayed.

**Type**

Indicates the type of entry in the database.

**Number**

Displays the current number of Permanent, Static, Dynamic, and Learned Address entries.

**Capacity**

Indicates the total capacity of each entry type in the Static and Learned databases.

**Aging Time**

Indicates the length of time, in seconds, that Dynamic and Learned Addresses in the Source Address Table are allowed to remain inactive before they are dropped from the database. The allowable time range for these entries is 10 to 1,000,000 seconds. Aging time is not applicable to Static or Permanent entries. You can configure this field, as described in the next section.

The following fields are applicable to the scrollable Address Entry panel of Filtering Database entries.

**Address**

Lists the addresses for which the bridge's Filtering Database has forwarding and/or filtering information.

**Type**

Indicates the type of an entry in the database. The possible types are Static, Dynamic, Learned, Self, or Permanent. You can alter the entry type, as described in the next section.

**Source Port**

Indicates the port number on which the address entry was first detected. A question mark (?) indicates that the address entry was not a learned entry, but Port Filtering information applies to it (i.e., the entry is a created Permanent, Dynamic, or Static entry and has corresponding filtering information).

**Receive Port**

Indicates the number of the port on which a frame must be received in order for the entry's Port Filtering information to apply. An asterisk (\*) indicates that the receive port is promiscuous, and applies to all ports of the bridge (assuming no conflicting entry applies). You can change the receive port, as described in the following section.

### Port Filtering


Indicates the action that will take place at each bridge port when it receives frames from the selected address entry. A green arrow indicates that the frames received from the address will be forwarded to the port's associated segment (➔<sup>1</sup>). A red circle indicates that frames will be filtered (blocked) from the port's associated segment (●<sup>2</sup>). You can change the Port Filtering action, as described in the next section. (Note that port filtering is scrollable among all the potential ports; however, only two consecutive ports can be viewed simultaneously.)

## Configuring the Filtering Database

You can configure the Filtering Database by:

- Altering the Aging Time for Dynamic and Learned entries.
- Changing the type of entry with the Type buttons.
- Changing the Receive port for the filter.
- Changing the Port Filtering action at each bridge port.
- Adding or deleting individual Filtering Database entries.
- Clearing all Permanent, Static, or Dynamic entries from the Filtering Database.

Note that although configuration changes will appear in the window, no action actually takes place in the bridge's Filtering Database until you click on the **OK** button in the bottom right of the window. This saves the new configuration.




*When you reconfigure the Filtering Database and click **OK**, the screen will clear temporarily and a message will appear to indicate that the information is being updated. When the changes have been successfully set and the Filtering Database has updated, the screen information will be refreshed.*

If you change the window without clicking on **OK**, then attempt to exit the window by clicking on **Cancel**, a text box will appear stating "Changes have been made. Cancel them?". Click on **Yes** to exit the window without changing the Filtering Database, or select **No** to return to the window.

### Altering the Aging Time

To alter the Aging Time for Dynamic and Learned entries:

1. Highlight the **Aging Time** field with the cursor.
2. Type in the new Aging Time (allowable range is 10 to 1,000,000 seconds).



Note that the Filtering Database Aging Time is the same as the Aging Time displayed (and configured) via the Source Addresses window. Setting the Aging Time in the Filtering Database window also changes the time in the Source Addresses window, and vice versa.

### Changing the Type of Entry

You can change any entry type from its current type (Learned, Self, Permanent, Static, or Dynamic) to either a Permanent, Static, or Dynamic entry. To do so:

1. Click on the shadowed **Type** button. A menu will appear with the three types to which the entry can be changed.
2. Highlight the desired type.

### Changing the Receive Port

You can change the Receive port of an address entry in the scrollable panel, so that a frame must be received at the specified port for the filtering action to apply. To do so, click on the **Receive** port in the panel. With each click, the Receive port will cycle to the next port (e.g., from \* (promiscuous), to 1, to 2, to 3, to 4, to 5, etc.).

### Changing the Port Filtering Action

You can change the Port Filtering action at each bridge port from its current action to the opposing action.

1. Maneuver the scroll bar until the desired port is in the Port Filtering panel.
2. Click on the port to alter its filtering action from forwarding frames from the associated address (→**1**), to filtering frames (●**2**) (or vice versa).

### Adding or Deleting Individual Entries

You can add or delete entries individually from the Filtering Database.

To add an address:

1. Click on the **New** button. A window (Figure 2-12) will appear.




Figure 2-12. Filter Database—New Filter Window

2. In the **Filter Address** field, type in the address (Hex format) for which you desire bridging. Be sure to add “-” as a separator between each byte in the address.



3. In the **Receive Port** field, type in the port at which the address must be detected for bridging to take place. If you enter a value of 0 in this field, the Receive Port is considered promiscuous (i.e., any port), and will be designated by an "\*" in the Address Entry panel.
4. Click on **OK**.
5. Specify the **Port Filtering** action on the address entry as described in the previous section.

To delete an address:

1. Click to highlight the address entry in the Address Entry panel that you wish to delete from the filtering database.
2. Click on **Delete**.

#### **Clearing All Permanent, Static, or Dynamic Entries**

To erase all Permanent, Static, or Dynamic entries from the Filtering Database, click on the associated **Clear** button in the upper portion of the window.

## Special Filter Databases

While the Filtering Database defines filters for all packets from a particular source address, the Ethernet Special Filter Database and the Token Ring Special Filtering Databases allow you to filter packets through an Ethernet or a Token Ring bridge, respectively, using a special filtering scheme.

When a packet is received at an Ethernet bridging interface, it is first checked against the Ethernet Special Filter Database to see if any filtering action applies to it. Because of this, an entry in the Ethernet Special Filter Database takes precedence over a filter entry in the Filtering Database that would otherwise apply to the packet.

The Ethernet Special Filter Database allows you to:

- Define and save a filter based on a combination of Source Address, Destination Address, Ethernet Data Type, and Data (including the offset).
- Specify the receive ports at which the filter will take effect.
- Specify the forwarding/filtering action at each bridging port of the device.

When checking for Transparent filtering information, the bridge first checks the Token Ring Special Filter Database to see if any filtering action applies to it. Because of this, a filter entry in the Token Ring Special Filter Database takes precedence over a filter entry in the Filtering Database that would otherwise apply to the packet.

Looking at each enabled filter, starting with the lowest numbered filter, the bridge compares the following fields to the corresponding fields in the received packet:

- Destination address
- Source address
- Ethernet or Token Ring data type
- Up to 16 hex integers (64 bytes) of the data field

In addition, a filter can also specify at which port or ports the packet must be received for the filter to be applicable. If a received packet matches *all* the contents of an enabled filter, the bridge forwards the packet to the defined set of ports.

Filters provide broad configuration flexibility. For example, you can define multiple scenarios for a single filter by specifying different combinations of receive port/destination port. You can use wildcard characters in filter fields to force a match with particular bits of the received packet's destination address, source address, type, or data. You can specify an offset for the data field, to specify the starting point in the data where the bridge looks for the match. For entries that don't match any of the enabled filters, you can configure the bridge to filter or forward the entry or pass the filter / forward decision to the Filtering Database.

## Ethernet and Token Ring Special Filter Database Windows

At the Ethernet and Token ring Special Filter Database windows, [Figure 2-13](#), you can view a list of the special filters for the selected bridge. There are 19 available filters in the Special Filter Database. You can not add any additional filters. You can view five of these filters at a time in the Special Filter Database window. Use the scroll bars to view the other fourteen filters.

When you first open the window, all filters will be undefined. For each field, bytes will be initialized with "match-any" characters (xx) for each digit. Any hexadecimal byte will be accepted as valid for the corresponding wildcard (xx) characters. For example, a Source Address filter defined as "xx-xx-xx-xx-bf-co" will pass the first four bytes of a frame's source address unconditionally, but the last two bytes must match the "bf-co" filter.

To access the Ethernet or Token Ring Special Filter Database window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Ethernet or Token Ring Special Filter Database**. The selected Special Filter Database window, [Figure 2-13](#), will appear.

To access the Ethernet or Token Ring Special Filter Database window from the Chassis View window:

1. Click on the **Board Index** of the device of interest; the Board menu will appear.
2. Click on **Ethernet or Token Ring Special Filter Database**. The selected Special Filter Database window, [Figure 2-13](#), will appear.

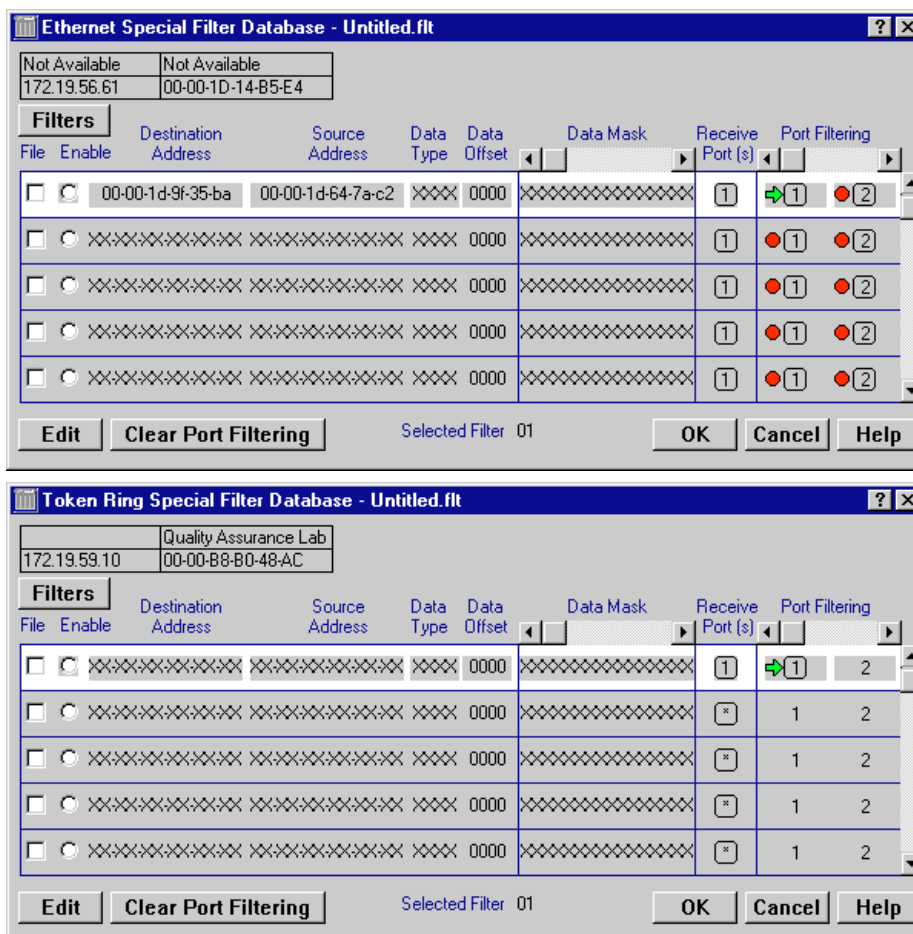


Figure 2-13. Ethernet and Token Ring Special Filter Database Windows

## Special Filter Database Window

The following fields are available in the Special Filter Database window:

### File

An X in this checkbox indicates that the filter is associated with the file name shown in the title bar of the window. If a file has not yet been saved, the title bar will display the filter name “untitled. flt”. A saved file name is only displayed in the title bar after you have opened a saved filter file or saved your current filters.

### Enable

A filled-in circle indicates the filter is enabled.

### Destination Address

Displays a six-byte hexadecimal field for the filter which can be used to filter on Destination Addresses, in whole or in part.

### Source Address

Displays a six-byte hexadecimal field for the filter which can be used to filter on Source Addresses, in whole or in part.

### Data Type

Displays the hexadecimal two-byte field for the filter which can be used to mask out a specified protocol type field. Examples of protocol type are:

- 0800 = IP
- 8137 = Novell
- 0bad = Banyan
- 80f3 = AppletalkARP

### Data Offset

Indicates the offset (in bytes, from the beginning of the data in the packet) where the Data Mask will be applied. The default for this field is 0000 (no data offset). An example of a valid offset to enter into this field is 0016 (16 bytes).

### Data Mask

Displays the 64-byte overlay used to filter on packets. The Data Mask is applied to the packet after the fixed part of the packet, data includes Source Address, Destination Address, and Type fields. The filter applies the mask directly at the start of the data portion of the packet unless there is a Data Offset. If a Data Offset has been defined, the mask will apply to the data that comes after the specified offset in the packet.


### Receive Port(s)

Indicates the ports at which the packet must be received for filtering information to be applied. Note that you can only immediately see one receive port per filter, even though you can set more than one receive port for the filtering action to apply. The receive port field can display each individual device bridge port, or "\*" The "\*" indicates that a packet can be received at any port for the filter to apply (i.e., the port is promiscuous).

### Port Filtering

 **1 forwarding**

Indicates the forwarding/blocking information for the filter at each port on the device. Note that you can only view two ports at a time.

 **2 blocking**

Use the scroll bar at the top of the column to view the hidden ports.

### Selected Filter

This field, visible at the bottom of the window, displays the number of the filter that is currently highlighted. The possible range is from 01-19.

## Defining and Editing Filters in the Special Database

You can edit an existing filter or define a new filter using the following steps:

1. Click to select the filter you wish to edit. The filter is selected when it is highlighted. When the bridge uses the Special Database, it starts with the lowest numbered enabled filter.
2. Click **Edit**. The Special Database Filter window, [Figure 2-14](#), will appear with the following fields:

Destination Address (six-byte hexadecimal field)

Source Address (six-byte hexadecimal field)

Type (two-byte hexadecimal field)

Data Offset (decimal field)

Data Mask (64-byte hexadecimal data mask)


	Destination Address	Source Address	Type	Data Offset
	XXXXXXXXXX	XXXXXXXXXX	XX	0000
<b>64 Byte Data Mask</b>				
1-16	XXXXXXXXXXXXXXXXXX	XXXXXXXXXXXXXXXXXX		
17-32	XXXXXXXXXXXXXXXXXX	XXXXXXXXXXXXXXXXXX		
33-48	XXXXXXXXXXXXXXXXXX	XXXXXXXXXXXXXXXXXX		
49-64	XXXXXXXXXXXXXXXXXX	XXXXXXXXXXXXXXXXXX		

Figure 2-14. The Special Database Filter Window

3. If you are editing an existing filter, the fields will reflect the current configuration. A filter that has not yet been defined will have wildcards (xx) in every field.

If you want to completely reconfigure an existing filter, click on **Clear**. This will revert all the fields to all xx's.

4. Highlight the field which you want to define, and enter the appropriate information.
5. When you have finished defining the filter, click on **OK**. This will save the filter you created and return you to the Special Filter Database window, where the configured filter will be displayed.



If you do not wish to save what you have entered in the Special Database Filter Window, click on the **Cancel** button. This will cancel what you have entered into this window and return you to the Special Filter Database window.

6. Click on **OK** to save the changes you have made and exit the Special Database Filter window.

## Changing the Receive Ports

You can set the receive ports in the Special Filter Database window either before or after you define a filter. These are the ports at which the frame must be received for the filtering parameters to apply. The default selection is Port 1.

To designate a receive port, click on the receive port icon (1) for the filter. As you click on the icon, it will cycle through the ports (e.g., 1, 2, 3, 4, etc. until the end of the interface table, and “\*”). When you have selected a port, you can set the port filtering action that will apply when the packet is received at that particular port (refer to the following section for further information).

In this fashion, you can specify all receive ports at which the packet must be received and the designated filtering action which will apply when the packet is received at each port. Selecting “\*” (promiscuous or any port) will apply the filter and its specified filtering action to all ports on the device.

Remember that you can only view a single receive port and its filtering action. To check all receive ports for a single filter, you must click on the receive port icon to cycle through the series of ports.

## Changing the Port Filtering Action

Use the port icons under the Port Filtering section of the Special Filter Database window to determine the port filtering action associated with the filter when it is received at a specified receive port. You can select the port filtering action either before or after defining the filter. By default, the filtering action is initially not set at any port. You must click on a port to invoke the filtering action symbols. After the first port is set (either to filtering or blocking), the remaining ports in the filter are set to blocking until you specify otherwise.

### Setting the Port Filtering Action

When you set the port filtering action for a filter, you determine whether the port will block or forward packets which match the filter’s specifications. To set port filtering action, click on the desired port icon (e.g., 1, 2, 3, 4, 5, 6, up to 32) to toggle from blocking (●(2)) to forwarding (➔(1)) or vice versa.

You can set the port filtering action for the bridging port on each port of the device, as well any BRIM ports.

### Clearing the Port Filtering Action

When you clear the port filtering action of a filter, all ports that were configured to forwarding or blocking will be reset to no action. Note that when you clear port filtering for a filter, the filtering or blocking action will be simultaneously cleared at all of its receive ports.

In order to clear the port filtering action, use the following steps.

1. Click to select the filter whose port filtering action you would like to disable.
2. Click on **Clear Port Filtering**. This will clear the port filtering action for the selected filter at all of its receive ports. The port filtering symbols will appear in cleared mode.

### Enabling and Disabling a Filter

To determine if a filter is enabled, check the **Enable** option button.

To enable a filter:

1. Click on the empty **Enable** button. When the button is filled () , the filter is enabled.

To disable a filter:

1. Click on the filled **Enable** button. When the button is empty () , the filter is disabled.

### Saving a Set of Filters to a File

When you have defined a set of filters, you can save that set to a file. This allows you to conveniently recall a series of filters when the need arises.

To save a set of filters:

1. Make sure that all filters that you want contained in the set have the File checkbox checked.
2. Click on **Filters**. A menu will appear.
3. Click on **Save As**. A standard Microsoft Windows Save File window will appear.
4. In the **File name** field, specify the file name and file path in which you want to save the filter series.
5. Click on **OK**. The file will be saved as indicated.

To update the file while it is still open, click on the **Save** selection from the Filters pull-down menu.

To open an existing file containing a filter set:

1. Click on **Filters**. A menu will appear.
2. Click on **Open**. A standard Microsoft Windows Open File window will appear.
3. To specify the file:

In the **File name** field, specify the file to open by path and name

**or**

Use the **Look in** drop-down list box and associated file list to select the desired file, and click to highlight it.

4. Click on **Open**.

The filters will appear in the Special Filter Database window, with all parameters (File, Enable, Source and Destination Address, Data Type and Offset, Data Mask, Receive Port, and Port Filtering Action) displayed as they were configured at the last file save.

## Interface Configuration

The I/F Configuration port-level menu option invokes the Interface Configuration window, which allows you to select a bridging method for a Token Ring bridging interface. You can also make this selection via the Token Ring Bridge Mode window; see [Token Ring Bridge Mode, page 2-89](#), for details.

This window also allows you to select one of three transmission methods that should be used when unknown addresses are received from end stations attached to the selected bridge port.

To access the Interface Configuration window from the Bridge Status window:

1. Click on the desired **Port** button (**1**) to display the Port menu.
2. Click on **I/F Configuration**. The Interface Configuration window, [Figure 2-15](#), will appear.

To access the Interface Configuration window from the Chassis View window:

1. Click on the appropriate bridge port index to access the Port menu.
2. Click on **I/F Configuration**. The Interface Configuration window, [Figure 2-15](#), will appear.



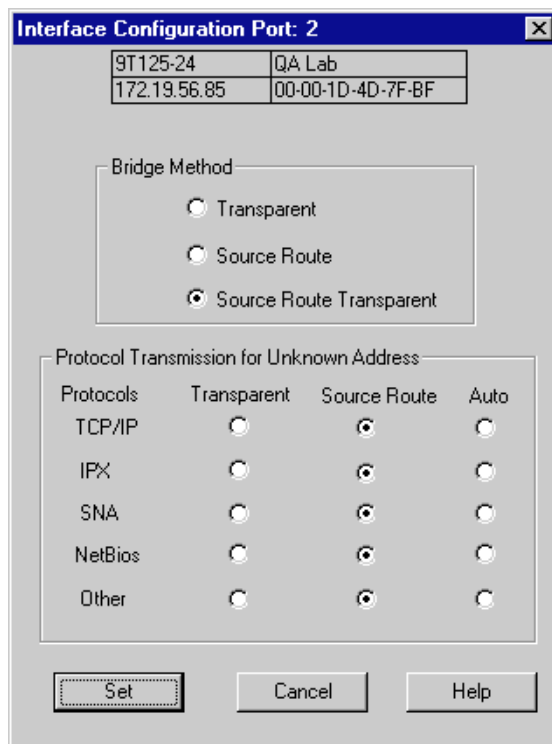


Figure 2-15. Interface Configuration Window

## Bridge Methods

The options available in the Bridge Method field are as follows:

### Transparent

When the bridge is set to Transparent mode, the bridge will only transmit transparent frames from the Token Ring connection. If a source route frame is received by the bridge, the Source Route information in the frame will be dropped from the packet. (A transparent frame is the same as a source route frame without a RIF—Routing Information Field.)

### Source Routing

When the bridge is set to Source Routing mode, the bridge will only transmit source route frames from the Token Ring connection. You should set the bridging mode to Source Route when you are bridging from Ethernet to Token Ring. The source route information (as configured via the Ethernet port's **Source Route Configuration** window, [page 2-57](#)) will be appended to the RIF for frames transmitted on the Token Ring.

### Source Route Transparent

When the bridge is set to Source Route Transparent, the bridge will transmit both transparent and source route frames. The frames received which have source route information will be transmitted as source route, while frames received that are transparent will be transmitted as transparent.

## Setting the Bridge Method

1. Click on the option button next to the bridging mode you would like your Token Ring bridge port to use: **Transparent Bridge**, **Source Routing**, or **Source Route Transparent**.
2. Click on **Set** to apply the desired mode.

## Protocol Transmission Methods

The options in the Protocol Transmission for Unknown Address field are as follows:

### TCP/IP

Determines whether IP frames received at the interface should be forwarded as transparent frames, source route frames, or both.

### IPX

Determines whether IPX frames received at the interface should be forwarded as transparent frames, source route frames, or both.

### NetBIOS

Determines whether NetBIOS frames received at the interface should be forwarded as transparent frames, source route frames, or both.

### SNA

Determines whether SNA frames received at the interface should be forwarded as transparent frames, source route frames, or both.

### Other

Determines whether frames of all other protocols not mentioned above (IP, IPX, NetBIOS, and SNA) that are received at the interface should be forwarded as transparent frames, source route frames, or both.

If **Transparent** is selected, the frame is forwarded out of the bridge interface as a transparent frame. If **Source Route** is selected, the frame is forwarded out of the bridge interface as a source route frame. If **Auto** is selected, the frame is forwarded out of the bridge interface as both a transparent frame and as a source route frame.

To select the transmission method for TCP/IP, IPX, SNA, NetBIOS or Other protocols:

1. Click on the option button next to the transmission method you would like your Token Ring bridge port to use: **Transparent**, **Source Route**, or **Auto**.
2. Click on **Set** to apply the desired mode.

## Source Route Configuration

With the Source Route Configuration window, you can view address and routing information, and set source route bridging parameters for bridging interfaces.

To access the Source Route Configuration window from the Bridge Status window:

1. Click on the desired **Port** button ( **1** ) to display the port menu.
2. Click on **Source Route Configuration**. The Source Route Configuration window, [Figure 2-16](#), will appear.

To access the Source Route Configuration window from the Chassis View window:

1. Click on the appropriate bridge port index to access the Port menu.
2. Click on **Source Route Configuration**. The Source Route Configuration window, [Figure 2-16](#), will appear.

9T125-24	QA Lab
172.19.56.85	00-00-1D-4D-7F-BF

Spanning Tree Mode: Manual    Bridge Number: **1**  
 Bridge Type: SRT

IP Address	172.19.56.85
Subnet Mask	255.255.0.0
MAC Address	00-00-B8-B2-FE-03
Local Segment	<b>0003</b>
Target Segment	<b>FFFF</b>
Hop Count Limit	<b>6</b>
Spanning Tree Expl.	Enabled

Spanning Tree Mode  
 Auto     Manual  
 Enable  
 Disable

Buttons: Set, Cancel, Help

Figure 2-16. Source Route Configuration Window

## Source Routing Information

Source Routing is a bridging technique developed by IBM and the 802.5 standards committee in which a bridge routes frames based on the contents of their media access control frame header, rather than by maintaining a filtering database to determine whether a packet should be forwarded or filtered. Source Routing functions as follows:

- An end point station transmits discovery (explorer) frames to a particular destination address in order to seek the best route through a bridged topology to that node. These frames are broadcast over the entire network.


In a network topology with parallel bridges, multiple paths may be available to the same destination. In this case, the explorer frame may be further defined as:

- **All Routes Explorer**, so that all possible routes to the destination are recorded, and multiple explorer frames can reach the same segment.
- **Spanning Tree Explorer** (also known as **Single Route Broadcast**), so that only one path is possible to a segment (i.e., through a designated bridge in a Spanning Tree topology), and only one explorer frame will be forwarded onto each segment. The Spanning Tree can be configured either automatically (i.e., by algorithm) or manually.
- When a source routing bridge processes the explorer frame, it adds a unique identifier to the frame in a reserved portion of the frame. This identifies the segment the frame was received from, followed by the specific bridge, and finally the segment it was forwarded onto.
- When the discovery frame (or frames if more than one route is possible) reaches its destination, it contains a complete record of bridge hops on its route.
- The destination address then returns All Routes Explorer (using specifically routed frames) and Single Route (Spanning Tree) Explorer frames (using All Routes Broadcast frames), to the source address.
- The source station selects one path from the returned explorer frames, and includes that path specification (with bridge and segment identifiers) in subsequent transmissions to that particular destination.

All bridges in the topology then examine the routing information field of a specifically source routed frame and either forward it if there is a match in the routing information—or if it is an All Routes explorer frame—or discard it.

## Source Route Configuration

The Source Route Configuration window allows you to view IP address and routing information, and to view and set source route bridging parameters for any bridging device which supports this menu option.



*It is recommended that the device be restarted when changes are made that affect source route bridging in order to clear the buffers, but you do not need to restart for the changes to take effect.*

The following fields are available in the Source Route Configuration Window:

#### **IP Address**

This field displays the Internet Protocol (IP) address, which acts as a logical identifier on the network, currently assigned to each port on the device. This is needed for SNMP network management capability. The IP address is expressed in dotted decimal notation (four decimal values between 0 and 255, separated by a period, e.g., 255.255.255.255).

This field can only be edited (with the correct security access) via Local Management for the device. Refer to the appropriate device-specific *User's Guide* for more information.

#### **Subnet Mask**

A subnet mask is used by a device to determine whether a destination address exists within its own subnetwork (logical division of the network by router or gateway) and can be reached directly, or whether it is unknown and therefore must be delivered to a router (as specified by the device's IP routing table or default gateway address).

A subnet mask should be set at the device if it will issue SNMP traps in a routed environment, so that the trap messages it generates will be routed correctly.

A subnet mask acts as a filter for destination IP addresses. It is a 32-bit quantity in which all bits that correspond to the network portion (both site and subnet identifying bits) of the device's IP address are set to 1, and all bits that correspond to the host portion are set to 0.

The device will logically AND a destination trap IP address with the subnet mask to determine which portion of the address identifies the network/ subnetwork. The device then compares the result on a bit-to-bit basis with the network-identifying bits in its own IP address. If the network portions match, the bridging device transmits the trap onto its subnetwork. If they do not match, the device transmits the trap through a router or gateway.

This field can only be edited (with the correct security access) via Local Management for the device (or the MIBTree utility). Refer to the appropriate device-specific *User's Guide* for more information.

#### **MAC Address**

This field displays the Media Access Control (MAC) layer address which identifies the ports/interfaces of the bridging device on a network. This six-byte address is set at the factory and is unique to each interface. Each byte is identified in bit order starting with the most significant bit. You cannot configure this field.

The following fields apply to the Source Route Configuration window:

**Local Segment**

This field displays the unique segment number that identifies the segment attached to the selected interface (either of the Token Ring or FDDI interfaces). The bridge adds the Local Segment number to the routing information field of source route discovery frames. Valid values range from 0 to 4095.

**Target Segment**

This field displays the unique segment number of the target segment that the source routed frame will be forwarded to. Valid values range from 0 to 4095.

**Hop Count Limit**

The maximum number of routing descriptors (i.e., bridge hops) allowed for an All Routes Explorer or a Spanning Tree Explorer frame received by the device. This will reduce the unnecessary propagation of explorer frames through the network.

You can use the **Set** button at the bottom of the window to change the Hop Count for the port, as explained in [Making and Setting Changes, page 2-62](#). The allowable range of values for this field is 0 to 28.

**Spanning Tree Expl.**

This read-only field displays the action currently being applied to Spanning Tree Explorer frames received by the indicated port. This field will appear in one of two ways:

- If the Spanning Tree Mode for the bridge is set to **Auto** (as explained in the following section), this field will display the Spanning Tree Port State for the indicated port.

If set to **Auto**, the device is subject to the Spanning Tree Algorithm. Each port will treat incoming frames according to its current Spanning Tree bridging state (i.e., Forwarding, Disabled, Listening, Learning, Blocking, or Broken).

- If the Spanning Tree Mode is set to **Manual** (as explained in the following section) this field will display either **Enabled** or **Disabled** as the Spanning Tree Port Enable State for the indicated port.

### Bridge Number

The Bridge Number uniquely identifies a bridge port when more than one bridge is used to span the same two segments. The Bridge Number should be in the range of 0 to 15.

You can use the **Set** button at the bottom of the window to change the bridge number of the port, as explained in [Making and Setting Changes, page 2-62](#). Current source routing protocols allow a range of 0 to 15 (0–F hexadecimal) for the bridge number identifier. If no bridge number is assigned to the device, a default value of 1 will appear in this field.

### Spanning Tree Mode

Indicates how a port on the device will behave with an incoming single-route broadcast (Spanning Tree Explorer—STE) frame. You can configure this field with the option buttons and checkboxes, or via the MIBTools utility or local management.

This field allows you to configure a Spanning Tree for your network. You can set the Spanning Tree Mode to Auto or Manual using the option buttons. We recommend that all bridges in your network topology have the same setting for Spanning Tree Mode (i.e., all set to Auto or all set to Manual).

**Auto** If the Spanning Tree Mode is set to **Auto**, a port that implements the Spanning Tree Algorithm (STA) and is enabled and in the forwarding state will accept and relay STE frames onto its attached segment.

Using STA, a bridge port will only forward frames if it is the designated port for its attached segment. A port is “designated” for its segment if it has the lowest Root Path Cost of all bridge ports attached to that segment. The Root Path Cost is the lowest total path cost calculated by adding the costs of each port along the path of a frame that traverses the bridge topology from the root to that port (including its own path cost).

If two ports on a segment have equal Root Path Costs, the port on the bridge with the highest priority bridge identifier (for convenience sake, that have the lowest numerical value) will be chosen as the root port.

You can affect Spanning Tree topology by changing the device’s bridge priority (Bridge Label) and path cost for its port pair (path cost increment) via the [Spanning Tree](#) window, [page 2-34](#).

**Manual** If the Spanning Tree Mode is set to **Manual**, you can manually configure the bridge to forward STE frames (i.e., manually establish a Spanning Tree for STE frames by determining which bridge in a parallel series of bridges will forward these frames).

If you set the Spanning Tree Mode to **Manual**, you can use the Enable or Disable checkboxes to set a port’s Spanning Tree Enable State to:

- **Enabled** (participating in frame relay).

- **Disabled** (not participating in the bridging process or in operation of the Spanning Tree Algorithm and protocol). If the Spanning Tree Mode is set to Disabled, the bridge port will not send or accept any STE frames. Any STE frame received will be discarded. The **Spanning Tree Expl** field at the Configuration window, and the **STE Frames** field at the Status window will both read “Disabled.”

## Making and Setting Changes

The Source Route Configuration window allows you to affect changes for the following Source Route Bridging parameters: Bridge Number, Local Segment, Target Segment, Hop Count Limit, and the device’s Spanning Tree Mode.

To make a change to Bridge Number, Local Segment, Target Segment, or Hop Count Limit, use the mouse to highlight the existing value in the desired field, and type in a new value.

To set the Spanning Tree Mode to **Auto** or **Manual**, click on the option button next to the appropriate selection. If set to **Auto**, a Spanning Tree Algorithm will calculate the device’s priority in a series of parallel bridges to determine a root bridge on the network. If set to **Manual**, you configure a Spanning Tree by administratively enabling or disabling each bridging port on the network.

When the device’s Spanning Tree Mode is set to **Manual**, you can change how a bridge port will treat a Spanning Tree Explorer frame. Use the **Enable** checkbox to allow STE frame forwarding at the port, or use the **Disabled** checkbox to prevent STE frame forwarding at the port. Click on the **Enabled** or **Disabled** checkbox to make your selection.

When you make changes in the Source Route Configuration window, they are not implemented at the device until you click on the **Set** button. This will cause the device to reboot. Since rebooting the device will bring it down for several minutes, a “Reset with new parameters?” pop-up dialog box will appear to ensure that you are ready. Click on **OK** to set the changes, or **Cancel** to return to the Source Route Configuration window.

## Using the Find Source Address Feature

You can select the Find Source Address option to discover which bridging interface a specified source MAC address is communicating through. When you select the Find Source Address option, a search is made of the 802.1d Bridge Filtering Database to discover the bridge interface associated with the address that you specify. If the search is successful, the corresponding interface will flash in the Chassis View window. See [Filtering Database, page 2-41](#), for details.

Use the Find Source Address feature as follows:

1. Click to display the **Device** menu.
2. Click again on **Find Source Address**. The following window will appear.



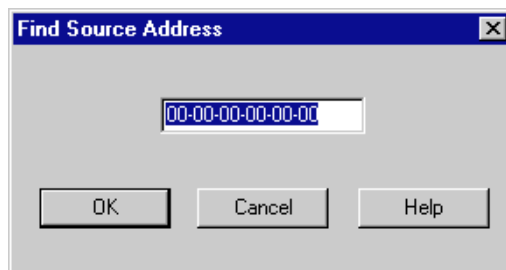


Figure 2-17. Find Source Address Window

3. In the text field in the middle of the window, enter a valid MAC address in hexadecimal format and then click **OK**.

If the address is found in the 802.1d Bridge Filtering Database, the port through which the address is communicating will flash in the front panel Chassis View display.

If the address is not found in the Filtering Database, a separate window will appear with a “Can’t Find Source Address” message.

## Using the Port Source Addresses Window

You can use the port-level Source Addresses window to view all the MAC addresses that are communicating through a selected bridge interface.

To open the Source Addresses window from the Bridge Status window:

1. Click on the desired **Port** button ( **1** ) to display the Port menu.
2. Click on **Source Addressing**. The Port Source Addresses window, [Figure 2-18](#), will appear.

To open the Source Addresses window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **Source Addressing**. The Port Source Addresses window, [Figure 2-18](#), will appear.

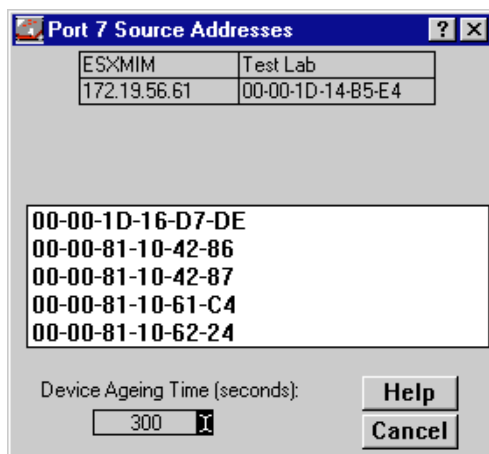



Figure 2-18. Port Source Addresses Window

The Port Source Addresses window displays the MAC addresses of all devices that have transmitted packets that have been forwarded through the selected bridging interface during the last cycle of the Filtering Database's defined aging timer (learned addresses that have not transmitted a packet during one complete cycle of the aging timer are purged from the Source Address Table). For more information, see [Filtering Database, page 2-41](#).



*The aging time displayed in the Port Source Addresses window is the same as the aging time displayed in the Filtering Database window. The aging time can be set from either window, and any changes to its value will be reflected in both locations.*

## Setting the Aging Time

The Filtering Database Aging Time is user-configurable through the Device Aging Time window.

To alter the Aging Time for Dynamic and Learned entries:

1. Click the **I-bar cursor** (I) next to the **Device Aging Time** field. The Device Aging Time window, [Figure 2-19](#), will appear.



Figure 2-19. Device Aging Time Window

2. Type in the new Aging Time, in seconds, then click on **OK**. The allowable range is 10 to 1000000 seconds; the default is 300 seconds.

## Using the Token Ring Bridge and Port Configuration Windows

The Bridge Configuration and the Port Configuration windows available for Token Ring devices look similar and are used for similar purposes, with the only exception being that the former window contains values that apply to the device as a whole, while the latter contains values that apply to the selected port.

The Bridge Configuration window provides a global capability to configure all of the Token Ring bridging interfaces on a device simultaneously as well as set the bridge number and virtual ring number (target ring).

The Port Configuration window provides the capability to configure individual Token Ring bridging interfaces on a device. This window displays the information that is set at the device level via the Bridge Configuration window, such as the Bridge Number and the Virtual Ring Number—both of which are read-only fields in the Port Configuration window.

The Ring Number field is the only field that is not common to both windows, because this value cannot be set globally on a device. It appears in the Port Configuration window only, since the value assigned to this field must be unique to each interface.

To access the Bridge Configuration window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Bridge Configuration**. The Bridge Configuration window, [Figure 2-20](#), will appear.

To access the Bridge Configuration window from the **Chassis View window**:

1. Click on the **Board Index** of the device of interest; the Board menu will appear.
2. Click on **Bridge Configuration**. The Bridge Configuration window, [Figure 2-20](#), will appear.

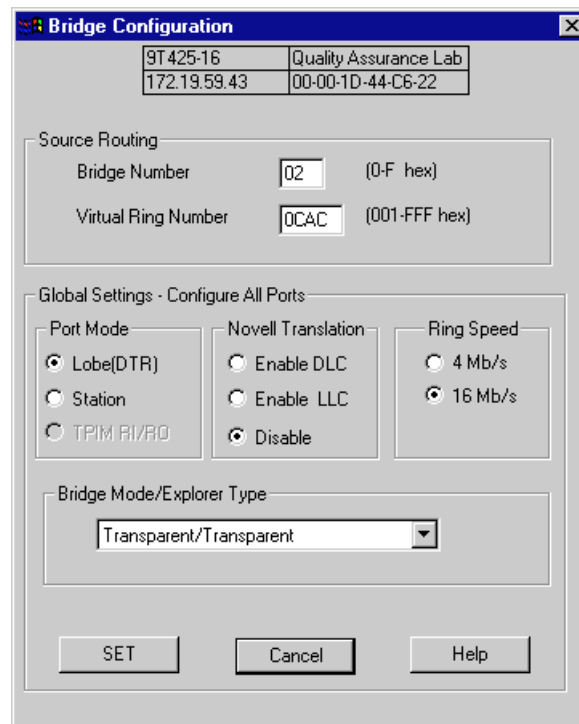


Figure 2-20. Bridge Configuration Window

To access the Port Configuration window from the Bridge Status window:

1. Click on the desired **Port** button (**1**) to display the port menu.
2. Click on **Port Configuration**. The Port Configuration window, [Figure 2-21](#), will appear.

To access the Port Configuration window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **Port Configuration**. The Port Configuration window, [Figure 2-21](#), will appear.

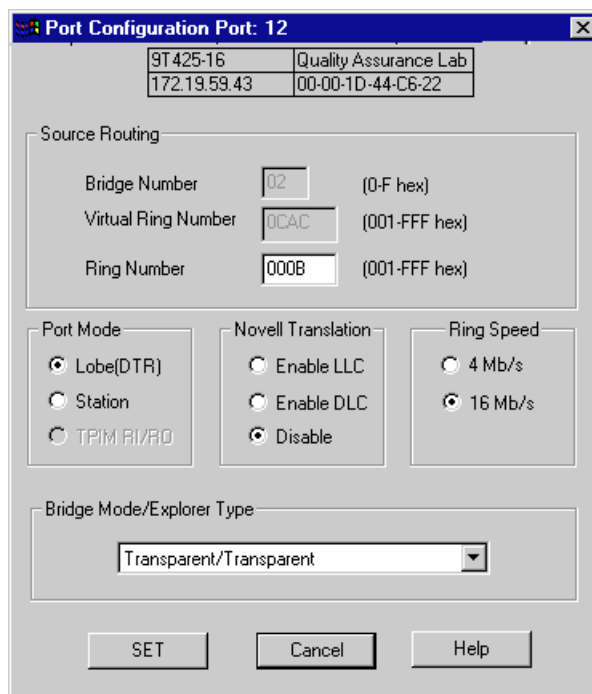


Figure 2-21. Port Configuration Window

The Bridge Configuration and Port Configuration window fields are defined as follows:

#### **Bridge Number**

Displays and allows you to set the number assigned to Token Ring bridge. This value is limited to the hexadecimal range of 0 through F. This field is settable in the Bridge Configuration window and read-only in the Port Configuration window.

#### **Virtual Ring Number**

Displays and allows you to set the number of the target segment connected to the selected bridge. This value is limited to the hexadecimal range of 001 through FFF. This field is settable in the Bridge Configuration window and read-only in the Port Configuration window.

#### **Ring Number**

Displays and allows you to set the segment number that uniquely identifies the segment to which this port is connected. This value is limited to the hexadecimal range of 001 through FFF. This field appears in the Port Configuration window only.

### Port Mode

Displays the three port mode options that are defined as follows:

- **Lobe**—allows direct-attach station connections (acting as a concentrator port).
- **Station**—provides station emulation.
- **TPIM RI/RO**—allows attachment to fiber TPIMs on standard workgroup hubs or direct-attachment to devices, such as servers, via fiber adapter cards. If the device does not support this mode of operation, this option will be grayed out.

### Novell Translation

Displays the three bit-order options that are available for translation—Enable LLC (Logical Link Control Translation), Enable DLC (Data Link Layer Translation), and Disable (No translation will take place). See [Using the Novell Translation Window, page 2-97](#), for more information.

### Ring Speed

Displays the selected ring speed, 4 Mbps or 16 Mbps.

### Bridge Mode/Explorer Type

Displays the available bridging mode and explorer frame type combinations:

<b>Source Route Transparent/Transparent</b>	Indicates the bridge forwards packets so they know the route and the devices they pass through to reach their destination. Explorer frames have no embedded routing information
<b>Source Route Transparent/Are</b>	Indicates that the bridge forwards packets so they know the route and the devices they pass through to reach their destination. ARE (All Routes Explorer) packets are sent to the destination station if the sending station does not receive a response to a test packet.
<b>Source Route Transparent/STE</b>	Indicates that the bridge forwards packets so they know the route and the devices they pass through to reach their destination. STE (Spanning Tree Explorer) packets are sent to the destination station if the sending station does not receive a response to a test packet.
<b>Source Route/ARE</b>	Indicates that the bridge forwards packets so they do not know the route or the devices they pass through to reach their destination. ARE (All Routes Explorer)

packets are sent to the destination station if the sending station does not receive a response to a test packet.

**Source Route/STE**

Indicates that the bridge forwards packets so they do not know the route or the devices they pass through to reach their destination. STE (Spanning Tree Explorer) packets are sent to the destination station if the sending station does not receive a response to a test packet.

**Transparent/Transparent**

Indicates that the two network segments are connected so that a single data route exists between any two end stations. Explorer frames have no embedded routing information



*If the ports on the selected bridge have different configurations, none of these options are selected. Do not select any of these options unless you want to set all ports on the selected bridge to a single mode.*

To set the Bridge number or the Virtual Ring Number in the Bridge Configuration window:

1. Click in the Bridge Number or the Virtual Ring Number field in the upper portion of the Bridge Configuration window. Enter a hexadecimal value between 0 and F in the Bridge Number field, or a hexadecimal value between 001 and FFF in the Virtual Ring Number field.
2. Click on **Set** to apply the change, or click on **Cancel** to exit the window without applying the change.

To set the Ring Number in the Port Configuration window:

1. Click in the Ring Number field in the upper portion of the Port Configuration window. Enter a hexadecimal value between 001 and FFF.
2. Click on **Set** to apply the change, or click on **Cancel** to exit the window without applying the change.

To set the Port Mode, Novell Translation, or Ring Speed globally for all bridge interfaces on a device in the Bridge Configuration window or for an individual interface in the Port Configuration window:

1. Click on the empty option button adjacent to one of the choices in the selected field. When the option button is filled (●), the selected choice will be enabled.
2. Click on **Set** to apply the change, or click on **Cancel** to exit the window without applying the change.

To select the Bridge Mode/Explorer Type globally for all bridge interfaces on a device in the Bridge Configuration window or for an individual interface in the Port Configuration window:

1. Click in the Bridge Mode/Explorer Type pull-down list box. Select one of the available choices: **Source Route Transparent/Transparent**, **Source Route Transparent/ARE**, **Source Route Transparent/STE**, **Source Route/ARE**, **Source Route/STE**, and **Transparent/Transparent**. The default selection is Transparent/Transparent.
2. Click on **Set** to apply the change, or click on **Cancel** to exit the window without applying the change.

## Duplex Modes

Some of the bridge interfaces on a device will support Full Duplex Switched Ethernet (FDSE) mode. Enabling full duplex mode on an interface allows the interface to receive and transmit packets at the same time, effectively doubling the available bandwidth.

On an Ethernet connection that is not using full duplex mode, the interface can either transmit or receive packets. The interface has to wait for one activity to be completed before switching to the next activity (receive or transmit).

Using the full duplex mode allows for faster transmission of packets over Ethernet connections because the bridging interface can transmit and receive packets; the interface does not have to wait for one activity to be completed before switching to the next one.



*Full Duplex should **only** be enabled on an interface that has a connection to a single destination address at the other end of the connection (i.e., it is not a segment with an attached repeater cascading the connection to multiple destination addresses).*

*Full Duplex mode disables the collision detection circuitry at the interface, so that both Transmit and Receive wires can be used simultaneously. With a single destination address at the other end of the connection (for example, if the connection was to a full duplex interface on another switching module, or if a single file server was connected to the full duplex switch port), this essentially doubles the available bandwidth from 10 Mbit/sec to 20 Mbit/sec. Note that the interface at the other end of the connection must also have Full Duplex enabled at the attached interface.*

*Full Duplex mode **must** be disabled if the interface is communicating with multiple destinations simultaneously (i.e., if a repeater is cascaded from the interface), since Ethernet relies on Collision Sense for proper operation.*



### The Duplex Modes Window

The bridge-level Duplex Modes window allows you to enable and disable full duplex mode capability for each bridging interface on your device. The window lists each interface on the device and whether full duplex is “ON” or “OFF” for each interface.

To access the Duplex Modes window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Duplex Modes**. The Duplex Modes window, [Figure 2-22](#), will appear.

To access the Duplex Modes window from the Chassis View window:

1. Click on the **Board Index** of the device of interest; the Board menu will appear.
2. Click on **Duplex Modes**. The Duplex Modes window, [Figure 2-22](#), will appear.




Figure 2-22. Duplex Modes Window

The following information is displayed in the Duplex Modes window:

#### Interface

Lists the bridging interfaces available on the device (Interface 1, Interface 2, and so on).

#### Full Duplex

Displays the current state of full duplex on each interface. Possible values for this field are as follows:

#### Connect A

Indicates that the interface is connected to MMAC Channel A and does not support full duplex mode (Interface 1 only). You will not be able to change the value of this field from this window.

<b>ON</b>	Indicates that full duplex mode is being used on this interface.
<b>OFF</b>	Indicates that full duplex mode is not being used on this interface.
<b>N/A</b>	Indicates that full duplex mode is not available on this interface.

### Setting the Duplex Mode

You set an interface to use or not use Full Duplex Switched Ethernet by turning the full duplex capability ON or OFF from this window.


To turn the full duplex mode ON or OFF:

1. In the Duplex Modes window, highlight the interface you want to change.
2. Double-click on the highlighted interface. The interface list will be briefly grayed-out as the set is being made to the device.

If the set is successful, the interface list will reactivate and the **Full Duplex:** indicator will switch from **ON** to **OFF** or **OFF** to **ON**.

If you attempt to set an interface to full duplex mode that does not support this feature, you will receive a “Set Failed” error message.

3. Click on **Cancel** to close the window.




*Because full duplex configuration takes place as you set each change individually, any changes that have been completed up to the point of clicking on **Cancel** will have been set at the device. Make sure that you have undone any unwanted changes before exiting the window.*

## Ethernet Port Configuration

You can also configure duplex modes from the Port Configuration window.

To access the Port Configuration window from the Bridge Status window:

1. Click on the desired **Port** button (  ) to display the Port menu.
2. Click on **Configuration**. The Port Configuration window for the selected Ethernet interface, [Figure 2-23](#), will appear.

To access the Port Configuration window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **Configuration**. The Port Configuration window for the selected Ethernet interface, [Figure 2-23](#), will appear.




Figure 2-23. Port Configuration Window

This window will indicate which mode is being used on the interface, standard mode or full duplex mode.

#### Standard Mode

Standard mode is being used on this interface. In standard mode, the interface can transmit *or* receive packets. The interface has to wait for one activity to be completed before switching to the next activity (receive or transmit).

#### Full Duplex

Full duplex mode is being used on this interface. In full duplex mode, the interface receives and transmits packets at the same time.

You set an interface to use standard or full duplex by selecting the appropriate mode from this window. When you open the Port Configuration window the currently used mode appears selected.

To change the mode from standard to full duplex mode or from full duplex to standard mode, click in the option button of the appropriate option and then click on **Apply**. To cancel the action without applying any changes, click on **Cancel**.

## Fast Ethernet Port Configuration

You can use the port-level Fast Ethernet Configuration window to manually configure 100Base-TX Fast Ethernet ports and FE100-TX Fast Ethernet Interface Modules (FEPIMs) for 10Base-T and 100Base-TX full or half duplex operation. You can also configure them to auto-negotiate with the device at the other end of the connection, based upon each device's Advertised and Remote Capabilities.

If you are monitoring a device with 100Base-FX Fast Ethernet ports, you can use the Fast Ethernet Configuration window to manually configure them to full or half duplex operation. No auto-negotiation is available for the 100Base-FX ports, and by extension, no Advertised or Remote capabilities.

From this window you can manually set the operational mode of the port, determining the speed of the port (10 Mbps or 100 Mbps), and whether it uses full duplex or standard mode bridging.

You can also set a 100Base-TX port to auto-negotiation so that the appropriate operational mode can be determined automatically (using the Advertised Abilities of the local interface that you determine, and the Remote Capabilities of the Remote Link). The mode you set will determine the speed of the port and whether it uses full duplex or standard mode bridging.

To access the Fast Ethernet Configuration window from the Bridge Status window:

1. Click on the desired **Port** button (1) to display the Port menu.
2. Click on **Configuration**. The Fast Ethernet Configuration Port X window (where X represents the port number of the selected interface), Figure 2-24, will appear.

To access the Fast Ethernet Configuration window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **Configuration**. The Fast Ethernet Configuration Port X window (where X represents the port number of the selected interface), Figure 2-24, will appear.

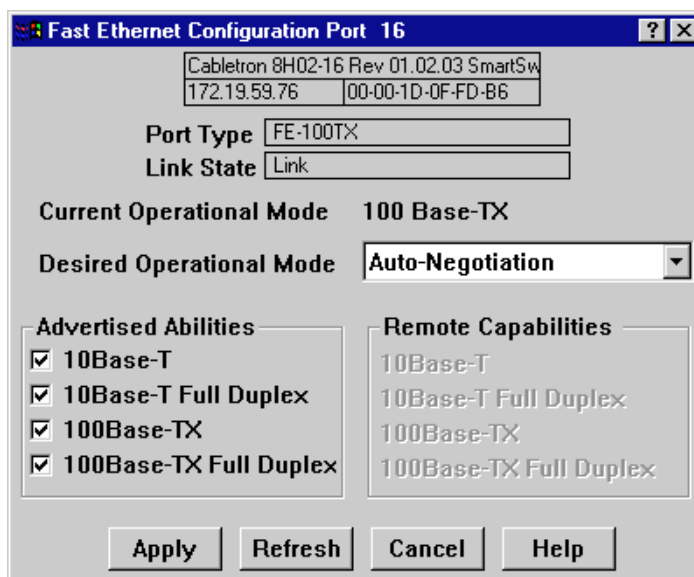


Figure 2-24. Fast Ethernet Configuration Port X Window

From this window you can manually set the operational mode of the port, or—for 100Base-TX interfaces—set the port to auto-negotiation so that the appropriate operational mode can be determined automatically. The mode you set will

determine the speed of the port and whether it uses full duplex or standard mode bridging.

The following information about the selected Fast Ethernet port is displayed:

**Port Type**

Displays the type of Fast Ethernet port for example, FE-100TX or FE-100FX.

**Link State**

Displays the connection status of the selected port: Link or No Link.

**Current Operational Mode**

Displays the mode that the port is operating in at the present time. Possible operational modes include 10Base-T, 10Base-T Full Duplex, 100Base-TX, 100Base-TX Full Duplex, 100Base-FX or 100Base-FX Full Duplex.

If no current operational mode is returned, it indicates the port is operating under auto-negotiation.

**Desired Operational Mode**

Displays the operational mode that you want to configure for this port. The following operational modes are available for each port:

**FE-100TX**      Auto-Negotiation, 10Base-T, 10Base-T Full Duplex, 100Base-TX, and 100Base-TX Full Duplex.

**FE-100FX**      100Base-FX and 100Base-FX Full Duplex

See [Setting the Operational Mode for the FE-100TX, page 2-76](#), and [Setting the Operational Mode for the FE-100FX, page 2-77](#), for details.

**Advertised Abilities**


This field works in conjunction with auto-negotiation on FE-100TX ports. During auto-negotiation, the local hardware will advertise all selected modes in descending bandwidth order: 100Base-TX Full Duplex, 100Base-TX, 10Base-T Full Duplex, and 10Base-T.

Of the selected abilities, the highest mode available on the port on the other side of the connection will automatically be used. The Advertised Abilities will only be used when auto-negotiation is enabled.

**Remote Capabilities**

This field displays the advertised abilities of the remote hardware at the other end of the link from the FE-100TX port. Again, possible advertised abilities by the remote partner include 10Base-T, 10Base-T Full Duplex, 100Base-TX, or 100Base-TX Full Duplex.

If auto-negotiation is not enabled or supported at either the local or remote interface, or if there is no active link, all entries in this field will be grayed out.



*Auto-negotiation is not available on the FE-100FX; therefore, the Advertised Abilities and Remote Capabilities section of the Fast Ethernet Configuration window will be grayed out when you are viewing the port configuration of an FE-100FX.*



*If you choose to select a specific mode of operation (rather than auto-negotiation), you should be sure that the link partner supports the same mode. Otherwise, no link will be achieved.*

*If you select a full duplex mode and the link partner supports the same wire speed but not full duplex, a link will be achieved, but it will be unstable and will behave erratically.*

*If you select auto-negotiation, the local node will try to match the mode of the link partner, even if the link partner is not set to auto-negotiate, and even if the local node must use a mode which it is not currently advertising.*

### Setting the Operational Mode for the FE-100TX

You can manually set the FE-100TX to use any one of four operational modes. You can also set the port to auto-negotiation, which allows the port to determine for itself the best operational mode using the Advertised Abilities and Remote Capabilities of the local and remote interface, respectively.

If you want to manually configure the mode:

1. Click on the **Desired Operational Mode** list-box, and select one of the following modes:

**10Base-T**—10 Mbps connection, Standard Mode

**10Base-T Full Duplex**—10 Mbps connection, Duplex Mode

**100Base-TX**—100 Mbps connection, Standard Mode

**100Base-TX Full Duplex**—100 Mbps connection, Duplex Mode

2. Click on **Apply**. The mode that you have chosen will be set at the port.

If you want the port to use auto-negotiation:

1. Click on the **Desired Operational Mode** list-box and select **Auto Negotiation**.
2. Click in the Advertised Abilities check boxes to select either **10Base-T**, **10Base-T Full Duplex**, **100Base-TX**, or **100Base-TX Full Duplex**.
3. Click on **Apply**.

When an active link is established, the operational mode will be dynamically set based on the modes selected in the Advertised Abilities field and the speeds and modes supported by the attached device; see the definition for [Advertised Abilities](#) on page 2-75.

### Setting the Operational Mode for the FE-100FX

You can manually set the FE-100FX to use either of two operational modes:

1. Click on the **Desired Operational Mode** list-box, and select one of the following modes:
  - **100Base-FX**—100 Mbps connection, Standard Mode
  - **100Base-FX Full Duplex**—100 Mbps connection, Duplex Mode
2. Click on **Apply**. The mode that you have chosen will be set at the port.

## SONET Port Configuration

The FE100-Sx series of Fast Ethernet Port Interface Modules, and the APIM-2x series of ATM Port Interface Modules provide SONET (Synchronous Optical Network) access for some of Cabletron's devices. SONET interfaces link high-speed local or metropolitan area networks by using an OC-3 connection (leased from your local telco or Internet service provider) to a SONET ring.

If your device is equipped with an FE100-Sx or an APIM-2x port interface module, you can use the SONET/SDH Configuration window to set its operating parameters, and the SONET/SDH Statistics window to view performance information for the interface (which can tell you if your telco/service provider is meeting any guarantees regarding network reliability).

### SONET/SDH Configuration

The SONET/SDH Configuration window lets you determine whether your FE-100Sx or APIM-2x port interface module will operate according to SONET or SDH (Synchronous Digital Hierarchy) standards.

SONET is the ANSI (American National Standards Institute) standard for the optical transport of data according to the transmission standards in effect in North America (United States/Canada), Korea, Taiwan, and Hong Kong. ANSI sets industry standards in the U.S. for the telecommunications industry, among other industries.

The basic SONET building block signal (transmitted at 51.84 Mbps) is referred to as STS-1 (Synchronous Transport Signal Level 1). SONET can multiplex (or combine) STS-1 signals into STS-N signals, where N is some integer multiple of STS-1 signals.

The ITU, or International Telecommunications Union (formerly known as the CCITT—the Consultative Committee on International Telegraph and Telephone) incorporated the SONET standard into its Synchronous Digital Hierarchy (SDH) recommendations, which address differences between the European and North American transmission standards. The ITU sets standards for international communications (except for nations adhering to ANSI standards). SDH is a world standard, and as such, the SONET standard is considered a subset within it.


The SDH transmission hierarchy uses the STM-1 (Synchronous Transfer Module Level 1) as its basic building block signal (transmitted at 155.52 Mbps). Again, there are STM-N signals, which are STM-1 signals that have been multiplexed into a higher signaling rate.

Table 2-1. SONET/SDH Transmission Hierarchies

SONET	Bit Rate	SDH
STS-1/OC-1	51.84 Mbps	—
<b>STS-3/OC-3</b> (supports FE-100Sx and APIM-2x in SONET operational mode)	155.52 Mbps	STM-1 (supports FE-100Sx and APIM-2x in SDH operational mode)
STS-12/OC-12	622.08 Mbps	STM-4
STS-24/OC-24	1244.16 Mbps	—
STS-48/OC-48	2588.32 Mbps	STM-16
STS-192/OC-192	9953.28 Mbps	STM-64

You should be sure that the operational mode for both the local and remote ends of the SONET connection is set appropriately for your region. Setting the wrong operational mode may cause errors to be generated during transmission, since there are slight differences in framing SONET and SDH signals.

To access the SONET/SDH Configuration window from the Bridge Status window:

1. Click on the desired **Port** button (  ) to display the Port menu.
2. Click on **SONET/SDH Configuration**. The SONET/SDH Configuration: Port X window, [Figure 2-25](#), will appear.

To access the SONET/SDH Configuration window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **SONET/SDH Configuration**. The SONET/SDH Configuration: Port X window, [Figure 2-25](#), will appear.






Figure 2-25. SONET/SDH Configuration Window

To set the operational mode of the SONET port via the SONET/SDH Configuration window:

1. Click on the option button adjacent to the appropriate selection, **SONET** or **SDH**, to choose the data transmission standard to be used by the interface.
2. Click on **Apply** to set your change at the interface, or **Cancel** to exit the SONET/SDH Configuration window without applying any changes.

## SONET/SDH Statistics

SONET/SDH statistics are available for each FE100-Sx, APIM-2x, or other SONET port interface modules installed in your device. The same statistics apply whether you have configured the interface to operate according to SONET or SDH transmission standards.

The FE100-Sx and the APIM-2x port interface modules are SONET path-terminating equipment (PTE). They act as an endpoint of an end-to-end connection between themselves and another similar port interface module. As endpoints, they are capable of generating and receiving the **Path Overhead** information contained within the SPE (Synchronous Payload Envelope) of the base-level SONET or SDH signals. Simply put, overhead is the extra bits in the digital stream that relay information besides traffic signals.

The Path Overhead provides for end-to-end performance monitoring of the link, the signal label (the content of the SPE, including status of mapped payloads), the path's current status, and path trace capabilities.

The SONET/SDH Statistics window enables you to view some of the error information contained within the Path Overhead that your FE100-Sx or APIM-2x is receiving from the remote endpoint.

The window will inform you whether there have been specific defects experienced on the SONET link, and if the network has experienced any significant unavailability time as a result.

With a SONET link, there are three levels of error conditions—**anomalies**, **defects**, and **failures**.

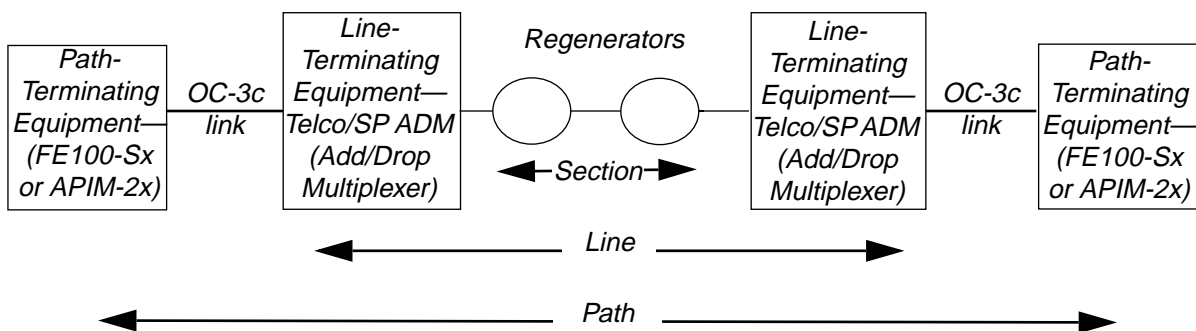
- **Anomalies** are small discrepancies between a desired and actual characteristic of an item, which when occurring singly will not interrupt the ability of the SONET network elements to perform their required functions.
- **Defects** indicate that anomalies have reached a level where the ability of the SONET network elements to perform their required functions has been interrupted. Defects are used in performance monitoring and in determining the fault's cause, and have impact on consequent actions on the network.
- **Failures** indicate that a network element has been unable to perform its required functions beyond a maximum time allocated to a given error condition.

These errors can occur in any of the four optical layers of a SONET network, which are (in order from lowest to highest layer in the hierarchy) the physical Medium, Section, Line, and Path layers.

- The **Medium layer** is the Photonic layer that physically converts electrical signals to optical signals.
- The **Section layer** deals with the transport of frames across the optical medium, including framing and scrambling data for transmission, the error monitoring and maintenance between section-layer elements (such as signal regenerators/repeaters), and orderwire (provisioning channels).
- The **Line layer** is responsible for reliably transporting the higher-level Path layer payload and overhead across the physical medium. It is responsible for synchronizing (clocking) the data transmission, multiplexing signals into a single channel, error monitoring and maintenance between line-layer elements (such as Add/Drop Multiplexers), and switching to secondary data paths should the primary path experience failure.
- The **Path layer** transports services between path-terminating equipment. It maps signals into a format required by the line layer, and reads, interprets, and modifies path overhead for performance monitoring and automatic protection switching.

Error reporting occurs at the Section, Line, and Path layers, and is carried within the corresponding SONET overhead. In terms of the SONET protocol stack, the three layers with overhead are mapped to the SONET link as shown in the following diagram.

The statistics and errors indicators provided in the SONET/SDH statistics window are taken from both the end-to-end Path layer, and from the Section layer between the FE100-Sx or APIM-2x and the Add/Drop Multiplexer to which it is connected. They reflect errors that may be occurring on your customer premises equipment, as well as errors that may be occurring at the Line or Section layers within the SONET MAN/WAN ring itself.



To access the SONET/SDH Statistics window from the Bridge Status window:

1. Click on the desired **Port** button ( **1** ) to display the Port menu.
2. Click on **SONET/SDH Statistics**. The SONET/SDH Statistics window for that interface, [Figure 2-26](#), will appear.

To access the SONET/SDH Statistics window from the Chassis View window:

1. Click on the appropriate port index to access the Port menu.
2. Click on **SONET/SDH Statistics**. The SONET/SDH Statistics window for that interface, [Figure 2-26](#), will appear.

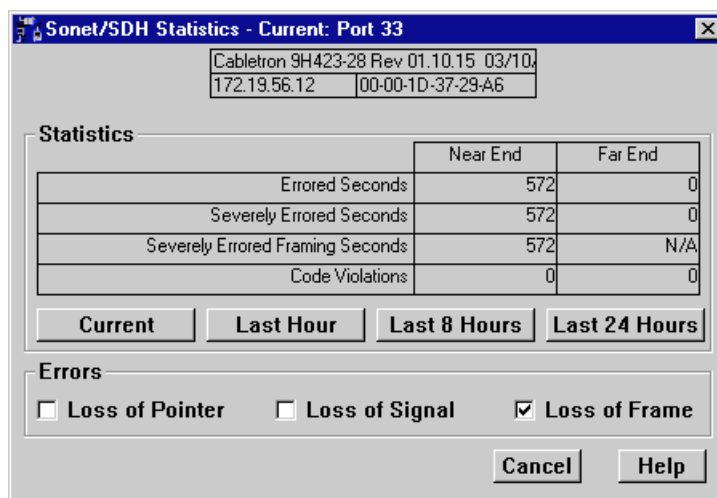


Figure 2-26. The SONET/SDH Statistics Window

## Errors

The Errors indicators at the bottom of the SONET/SDH Statistics window show the status of the SONET link as reported by the monitored interface, by indicating whether the link has experienced **Loss of Pointer**, **Loss of Signal**, or **Loss of Frame** defects or failures during the current 15-minute interval.

Note that Loss of Pointer is detected at the Path level on the SONET link, meaning that the error occurred anywhere on the end-to-end link between the connected FE100-Sx or APIM-2x devices that are customer premises equipment (CPE); Loss of Signal and Loss of Frame are detected at the Section level, meaning that the error occurred on the SONET section between the monitored CPE device and the ADM node (line-terminating equipment—LTE) to which it is connected.

Note also that these indicators simply show which error conditions have been detected during the last 15-minute interval; they do not alter the display of the statistics above.

- **Loss of Pointer**—SONET uses “pointers” to compensate for frequency and phase variations as data is being transmitted across the optical network, so that data is not delayed or lost on the network. Basically, a pointer is a data offset value that indicates where in the frame that the “payload” (user data and path overhead) begins, so that it can be differentiated from the “transport overhead” (the information in the frame used for transporting it across the SONET network).

A Loss of Pointer (LOP) **defect** occurs when either a valid pointer is not detected in eight consecutive SONET STS-N frames, or when eight consecutive frames are detected with the New Data Flag (NDF) set without being validly combined into an STS-N(c)—a concatenated STS-N signal—to carry a larger payload.

An LOP defect is cleared when three consecutive frames are detected with either a valid pointer and a normal NDF, or a valid concatenation indicator. Note that incoming Alarm Indicator Signals (which are alarm messages generated by the line and section layers that are propagated along the path to indicate a loss of signal condition on upstream network elements) cannot contribute to an LOP defect.

A Loss of Pointer **failure** is declared when a defect condition persists for a period of 2 to 3 seconds; the LOP failure is cleared when there is no defect condition detected for 9.5 to 10.5 seconds

- **Loss of Signal**—Incoming SONET signals are monitored for Loss of Signal (LOS) errors, which indicate a loss of physical signal failure (either optical or electrical) at the source (e.g., a laser failure) or in the transmission facility (e.g., a fiber cut). Loss of signal is detected in the data (before scrambling) by an “all zeros” pattern, which indicates that there are no light pulses for OC-N optical interfaces (on the line-terminating equipment or a regenerator), or no voltage transitions for STS-1 or STS-3 electrical interfaces (on path-terminating equipment, such as the FE100-Sx or APIM-2x).

A state of no transitions that lasts 2.3  $\mu$ s (microseconds) or less is insignificant.

A state of no transitions that lasts between 2.3  $\mu$ s and 100  $\mu$ s is declared an LOS *defect*. The LOS defect is cleared after a 125  $\mu$ s interval (the time required to transmit one frame on a SONET network) during which no LOS defect is detected.

If the LOS defect persists for a period of 2 to 3 seconds, an LOS *failure* will be declared, an alarm indicator will be set, and an alarm message will be sent to an Operations Systems application (responsible for overseeing the entire network). The LOS failure is cleared when the LOS defect is absent for a period of 9.5 to 10.5 seconds.

A Loss of Signal may also be detected if the received signal level (e.g., the incoming optical power) falls below a Bit Error Rate (BER) threshold of 1 in  $10^3$ . A BER is the number of coding violations detected in an interval of time (usually one second). A predicted BER of 1 in  $10^3$  means that during each second, there is an error ratio of 1 errored bit per 1,000 bits sent. This state clears when two consecutive framing patterns are received, and no "all zeros" LOS conditions are detected in the intervening time (one frame).

Note that for path- or line-terminating SONET network elements, LOS failure detection is also linked to the declaration or clearing of Loss of Frame (LOF) failures (described below). If there was a previously existing LOF failure at the time an LOS failure is declared, the LOF failure will be cleared; if an existing LOS failure is cleared, but LOF failure conditions still exist, an LOF failure will be immediately declared on clearing the LOS failure.

- **Loss of Frame**—SONET frames uses A1 and A2 framing bytes in the section overhead to indicate the beginning of the frame. An Out of Frame (OOF) alignment defect (also known as a Severely Errored Frame—SEF—defect) occurs when four consecutive SONET frames are received with invalid patterns in these framing bytes. This defect is cleared when two consecutive SONET frames are received with valid framing patterns.

A Loss of Frame (LOF) **defect** occurs when this OOF/SEF defect persists for a period of 3 milliseconds. This defect is cleared when the incoming signal remains continuously in-frame for a period of 1 to 3 milliseconds.

An LOF **failure** is declared when an LOF defect persists for a period of 2 to 3 seconds (except when a Loss of Signal defect or failure is present, as described above). An LOF failure is cleared if an LOS failure is declared, or when the LOF defect is absent for 9.5 to 10.5 seconds.

### Statistics

Statistics are given for both the Near-End and Far-End of the SONET/SDH path. Far-end statistics are taken from the far-end block error code (FEBC)—used to indicate that the remote entity at the far-end of the path has detected errored data—within the Path Overhead of SONET frames.

You can view statistics for the current 15-minute interval, or accumulated over the last one-, eight-, or 24-hour period by clicking on the appropriate selection button.

- **Errored Seconds**—The counter associated with the number of Errored Seconds, or Far-End Errored Seconds, encountered by a SONET/SDH Path in the specified interval.

An Errored Second (ES) is a second with one or more coding violations (bit parity errors) at the associated layer reported at the Section, Line, or Path layer of the SONET link, **or** a second during which at least one or more incoming defects (e.g., Loss of Signal, Loss of Pointer, or Loss of Frame) has occurred at that layer. Coding Violations are Bit Interleaved Parity (BIP) errors that are detected in the incoming signal (as described below).

- **Severely Errored Seconds**—The number of Severely Errored Seconds, or Far-End Severely Errored Seconds, encountered by a SONET/SDH Path in the specified interval.

A Severely Errored Second (SES) is a second with X or more coding violations (bit parity errors) reported at the Section, Line, or Path layer of the SONET link, **or** a second during which at least one or more incoming defects (e.g., Loss of Signal, Loss of Pointer, or Loss of Frame) has occurred at that layer. The statistic provided in this field is provided by the STS-Path level of the link.

Values of X at each layer depend on the link's line rate and the Bit Error Rate. For the STS-Path layer, with a line rate of 51.84 Mbps (STS-1) and a BER of  $1.5 \times 10^{-7}$ , X is **9**; with a line rate of 155.52 Mbps (STS-3) and a BER of  $1 \times 10^{-7}$ , X is **16**.

If the FE100-Sx or APIM-2x is experiencing consecutive Severely Errored Seconds, it may indicate an impending period of network unavailability (which begins at the onset of 10 consecutive SESs). Periods of unavailability can severely impact service (e.g., the disconnection of switched services). Availability is restored at the onset of 10 consecutive error-free seconds.

- **Severely Errored Framing Seconds**—The counter associated with the number of Severely Errored Framing Seconds encountered by a SONET/SDH Section in the specified interval. A Severely Errored Framing Second (SEFS) is a second containing one or more SEF events. This counter is only counted at the Section Layer, and is not available as a Far-End counter.
- **Code Violations**—The number of Coding Violations (CVs) encountered by a SONET/SDH Path interface, or the number of Far-End Coding Violations reported via the far-end block error count to the monitored SONET/SDH Path interface, in the specified interval.

Coding Violations are Bit Interleaved Parity (BIP) transmission errors that are detected in the incoming signal. Bit Interleaved Parity is a check at the receiving interface that groups all bits in a block into a unit (e.g., a byte), then verifies the block for parity for each bit position in the group by making sure that the number of bits set to the value '1' is either even or odd, as reported by the transmitting entity.

## Configuring SmartTrunking

The SmartTrunk menu option invokes the SmartTrunk Configuration and Status window, which allows you to group interfaces logically to achieve greater bandwidth between devices when both devices support this feature. There is no limit to the number of ports that can be included in a single “trunk.”



*SmartTrunking is designed to work in the traditional bridging mode only, and is not available if a switch is in the SecureFast VLAN mode. The SecureFast VLAN architecture supports a fully-meshed topology, which has benefits similar to SmartTrunking.*

To access the SmartTrunk Configuration and Status window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **SmartTrunk**. The SmartTrunk Configuration and Status window, [Figure 2-27](#), will appear.

To access the SmartTrunk Configuration and Status window from the Chassis View window:

1. Click on the **Board Index** of the device of interest; the Board menu will appear.
2. Click on **SmartTrunk**. The SmartTrunk Configuration and Status window, [Figure 2-27](#), will appear.

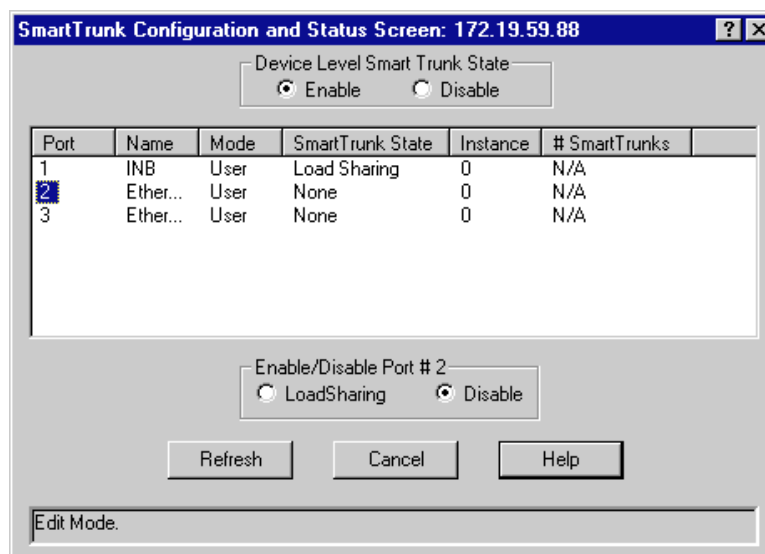


Figure 2-27. The SmartTrunk Configuration and Status Window

The SmartTrunk Configuration and Status window displays all of the ports on the selected device. The following information is given for each port:

**Port**

Displays each port on the selected module. Use the scroll bar to the right of the list box to view information for all available ports.

**Name**

Displays the interface description of the selected port.

**Mode**

Displays the connection type for each port, either **User** or **Network**. **User** connections do not participate in SmartTrunking; **Network** connections do. At least two ports (from two separate chassis) must be designated as **Network** connections to participate in SmartTrunking. All FNB interfaces must be designated as **User** connections.

**SmartTrunk State**

Displays the current operating state of each listed port. The possible states include:

- **None**—The port is operating as a normal switch port.
- **Blocking**—The port is load sharing, but in the blocked mode. While the module performs the function of determining if there is a network loop, data is temporarily blocked on new SmartTrunk ports and on any port that becomes newly linked.
- **Load Sharing**—The port is actively load-sharing with other ports.

**Instance**

Displays the number of ports associated with each loop.

**# SmartTrunks**

Displays the total number of load-sharing ports in the loop.

The only configurable fields in the SmartTrunk Configuration and Status window are the two fields with option buttons, each with two possible settings:

**SmartTrunk** (with the options of **Enable** and **Disable**) and **Enable and Disable Port # X** (with the options of **LoadSharing** and **Disable**).



*When you first open the SmartTrunk Configuration and Status Screen, the **Enable and Disable Port # X** field will be labeled **SmartTrunk State Port #**. After you click on a port number in the list box, the field title will change to **Enable and Disable Port # X**.*



To enable or disable SmartTrunking at the device level:

1. Click on the option button adjacent to the value you wish to set in the Device Level SmartTrunk field: **Enable** or **Disable**.

When the option button is filled, the following confirmation window (Figure 2-28) will appear:




Figure 2-28. Device Level SmartTrunking Confirmation Window

2. Click on **Yes** to apply your selection, or **No** to exit the confirmation window without applying the change.
3. Click on **Refresh** to ensure that changes are applied.

To enable or disable SmartTrunking on an individual interface:

1. Click to select the interface whose Load Sharing status you wish to change. The interface number will then be listed as "X" in the Enable/Disable Port # X field.
2. Click on the option button adjacent to the value you wish to set: **LoadSharing** or **Disable**. When the option button is filled, the following confirmation window (Figure 2-29) will appear:





Figure 2-29. Example LoadSharing Confirmation Window

3. Click on **Yes** to apply your selection, or **No** to exit the confirmation window without applying the change.
4. Click on **Refresh** to ensure that changes are applied.



When you first open the SmartTrunk Configuration and Status Screen, the **Enable and Disable Port # X** field will be labeled **SmartTrunk State Port #**. After you click on a port number in the list box, the field title will change to **Enable and Disable Port # X**.

## Configuring Broadcast Suppression

Excessive broadcasts to all ports, or broadcast storms, can result in severe network performance problems, and possibly cause the network to crash. Devices which support the broadcast suppression feature provide automatic protection against broadcast and multicast storms.

In many ways, broadcast suppression is similar to filtering. To protect against storms, an acceptable rate for broadcast traffic across a port is defined. Once the user-defined threshold has been reached on an interface, broadcast frames will be dropped and an SNMP trap message will be sent to the network management station.

To access the Broadcast Suppression window:

1. Click on the **Device** menu from the Chassis View window.
2. Click on **Broadcast Suppression**. The Broadcast Suppression window, [Figure 2-30](#), will appear.

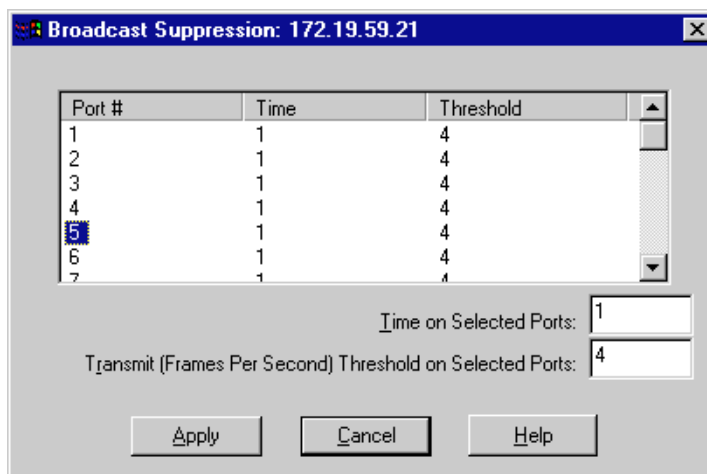


Figure 2-30. The Broadcast Suppression Window

In the Broadcast Suppression Window, each interface of the device that is being monitored can be individually configured for automatic broadcast and multicast storm protection.

You can also define what level of broadcasts the device will recognize as a broadcast storm by specifying the number of broadcast packets that can be transmitted within a given time period.

To configure a port for broadcast storm protection:

1. Click to highlight the entry for the port you wish to configure for automatic broadcast storm protection.
2. In the **Time on Selected Ports** field, enter the desired time period in seconds. Note that a value of 0 will disable the threshold alarm.
3. In the **Transmit (Frames Per Second) Threshold on Selected Ports** field, enter the number of broadcast packets that will be the threshold for the time period set in Step 2.
4. Click **Apply** and your settings will be added to the window. Click **Cancel** to close the window.

## Token Ring Bridge Mode

The Token Ring Bridge Mode window allows you to choose one of the three different modes of bridging on the device's Token Ring bridge port: Source Route Transparent, Source Routing, or Transparent. The default setting is Source Route Transparent.

To access the Token Ring Bridge Mode window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Token Ring Bridge Mode**. The Token Ring Bridge Mode window, [Figure 2-31](#), will appear.

To access the Token Ring Bridge Mode window from the Chassis View window:

1. Click on the **Board Index** of the device of interest; the Board menu will appear.
2. Click on **Token Ring Bridge Mode**. The Token Ring Bridge Mode window, [Figure 2-31](#), will appear.

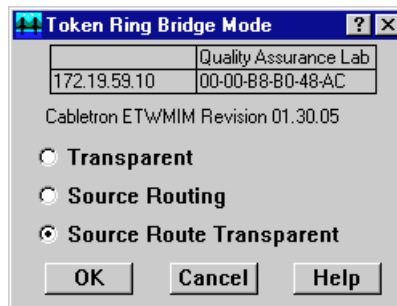


Figure 2-31. Token Ring Bridge Mode Window

The following options are available in the Token Ring Bridge Mode Window:

### Transparent

When the bridge is set to Transparent mode, the bridge will only transmit transparent frames from the Token Ring connection. If a source route frame is received by the bridge, the Source Route information in the frame will be dropped from the packet. (A transparent frame is the same as a source route frame without a RIF—Routing Information Field.)

### Source Routing

When the bridge is set to Source Routing mode, the bridge will only transmit source route frames from the Token Ring connection. You should set the bridging mode to Source Route when you are bridging from Ethernet to Token Ring. The source route information (as configured via the Ethernet port's [Source Route Configuration](#) window, [page 2-57](#)) will be appended to the RIF for frames transmitted on the Token Ring.

### Source Route Transparent

When the bridge is set to Source Route Transparent, the bridge will transmit both transparent and source route frames. The frames received which have source route information will be transmitted as source route, while frames received that are transparent will be transmitted as transparent.

## Setting Token Ring Bridge Mode

1. Click on the option button next to the bridging mode you would like your Token Ring bridge port to use: **Transparent Bridge**, **Source Routing**, or **Source Route Transparent**.
2. Click on **OK** to close the window and set the bridge to the desired mode.

## Setting Bridge Translation

When bridging between Ethernet and Token Ring networks, it is necessary to alter the MAC (Media Access Control) layer information. The Bridge Translation window controls the default frame translation that will occur for Token Ring frames that are bridged to Ethernet.

If the device that is being monitored via SPECTRUM Element Manager is set to operate in Auto Translation Mode, it has the ability to learn the frame type that is used for a given Source Address from its Ethernet MAC address.

However, when the Token Ring bridge port receives a frame from an unlearned Ethernet MAC address, it is necessary to configure the default Token Ring-to-Ethernet frame translation. When the bridge is operating in Auto Translation Mode, the learned frame type for a given Ethernet MAC address will override the default setting which may be configured in this window.

To access the Bridge Translation window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Bridge Translation**. The Bridge Translation window, [Figure 2-32](#), will appear.

To access the Bridge Translation window from the Chassis View window:

1. Click on the **Board Index** of the bridging device of interest; the Board menu will appear.
2. Click on **Bridge Translation**. The Bridge Translation window, [Figure 2-32](#), will appear.

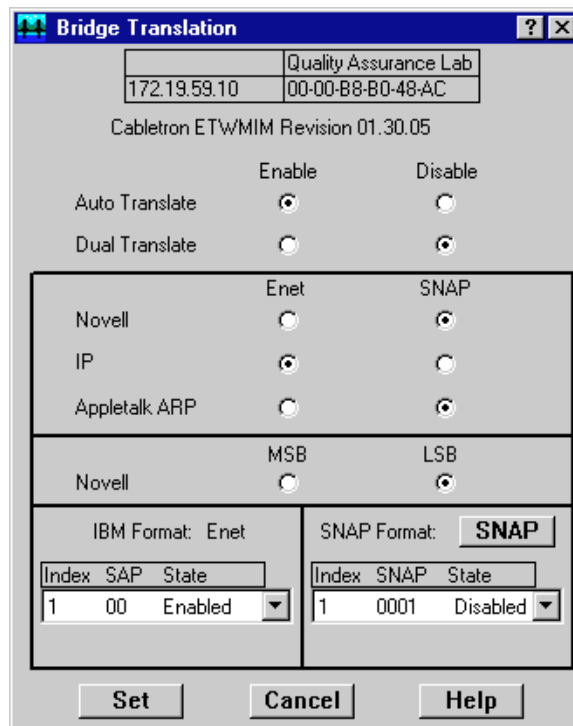


Figure 2-32. Bridge Translation Window


## Enabling and Disabling Auto and Dual Translate Modes

When Auto Translate mode is enabled, the bridge will learn, from its Ethernet ports, which frame format each source address is using. When the bridge forwards a Token Ring packet to one of these learned source addresses, it will automatically translate the packet to the correct frame format.

When Dual Translate mode is enabled, the bridge will translate a Token Ring broadcast or multicast packet to both Ethernet frame formats (i.e., Ethernet Type-II and Novell 802.3) when the format has not been learned previously.

To enable or disable translation:

1. Click on the empty **Enable** or **Disable** option button. When the option button is filled (●), the selected translation mode is enabled or disabled.



*When Auto Translate mode is enabled, it will override any other settings you configure in this window.*

## Configuring Token Ring Packet Translation

The middle section of the Bridge Translation window allows you to select the Ethernet frame format the bridge will use to translate various packet types that have been transported across Token Ring media.




Figure 2-33. Token Ring Packet Translation


The three Token Ring packet types you can configure for translation are:

- Novell packets
- IP packets (including ARP packets)
- AppleTalk ARP packets

You can set each type of packet to either **Enet** (Ethernet) or **SNAP** (Token Ring 802.5 with SNAP) packet formats.

To set a selected packet type to Enet or SNAP:

1. Click on the empty **Enet** or **SNAP** option button, located to the right of the selected packet type. When the option button is filled () , the selected packet format is enabled.



*For the Novell Token Ring packet the SNAP format is invalid. When the SNAP format is selected the packet will actually be translated into Novell 802.3 format.*

*If you are using Appletalk Phase-II, the mode of Appletalk ARP should be set to SNAP.*

## Configuring Novell Token Ring Packet Translation

The next section of the window allows you to select the bit ordering of the hardware addresses located in the data field of a Novell Token Ring packet. The data field contains all bytes of the packet, with the exception of the MAC header and CRC byte. For more information about bit ordering and translation, see [Using the Novell Translation Window, page 2-97](#).




Figure 2-34. Configuring Bit Order Translation

There are two bit-order options:

- **LSB**—Least Significant Bit
- **MSB**—Most Significant Bit

To set the bit order:

1. Click on the empty **LSB** or **MSB** option button. When the option button is filled (☉), the selected bit order is enabled.



*Novell recommends setting the bit order to LSB (Least Significant Bit) when bridging between Token Ring and Ethernet. When the bridge is set to operate in LSB mode, all the NIC cards on the Token Ring network should also be configured to operate in LSB mode.*

## The IBM Translation Table

The IBM Translation Table allows you to configure the mode of frame translation for IBM SNA frames bridged from a Token Ring network to an Ethernet network. The device will translate IBM SNA frames using any of the five default SAP (Service Access Point) addresses to Ethernet Type-II frames with a type code of 80D5; these default SAP addresses are shown in the IBM Format Table as the first five entries.

Index	SAP	State
1	00	Enabled

Figure 2-35. The IBM Format Table

If the state on the entry is enabled, then the translation for the Token Ring frame with the given SAP address will be translated to Ethernet Type-II frame format.

If the state of the entry is disabled, the frame will be bridged as a 802.2 frame maintaining the SAP address information.

From the IBM Translation window you can change the first five default SAP address entries, or add up to 15 additional entries to the table for the purpose of translating IBM SNA frames.

To view all the entries in the table you must click on the scroll box to the right of the displayed entry. A small scrollable window will appear. You can scroll through the window to view all the entries.



The table contains the following information fields:

**Index**

Displays the index of an entry in the Token Ring-to-Ethernet Translation table for IBM protocols. This number can be from 1 to 20.

**SAP**

Displays the SAP (Service Access Point) value of an IBM protocol. The bridge will translate all Token Ring 802.5 packets without a SNAP header and with this SAP value to Ethernet packets with Ethernet type set to 80D5. The default values included in the table are 0x00, 0x04, 0x08, 0xF0, and 0xFC; you can change any of these defaults, and/or add up to 15 additional SAP values.

**State**

Displays whether or not this entry will be used in the packet translation process. If the state is **Enabled** the entry is used; if the state is **Disabled** the entry is not used.

To configure IBM Translation:

1. Click on the scroll box next to the entry and the scrollable window of entries will appear.
2. Click on the entry you want to change. The IBM Translation Table Entry Window, [Figure 2-36](#), will appear.

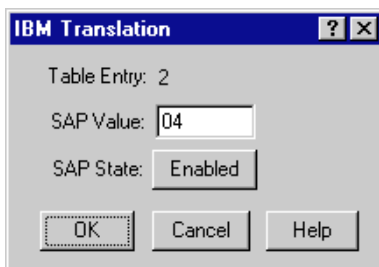


Figure 2-36. IBM Translation Table Entry Window

The Table Entry field will display the index of the entry you are viewing and configuring.

1. To change the **SAP Value** of this entry, highlight the current value and type a new one. The SAP Value must be one hexadecimal octet.
2. To set whether this entry will be used in the translation process, click on the **SAP State** button to toggle between Enabled and Disabled.
3. Click on **OK** to save your changes to the IBM Translation table. The new values will now appear in the window.

## The SNAP Translation Table

The SNAP format table contains Ethernet type field entries that the bridge uses to determine how to translate a Token Ring (802.5 with SNAP) packet to an Ethernet packet. If a Token Ring packet is received that matches the Ethernet type field of an enabled entry, the bridge will translate the Token Ring packet to the frame format specified by the button above the table. Otherwise, it will translate the packet to the opposite frame format.




Figure 2-37. SNAP Format Table

The SNAP Format Table contains the following fields:

### Index

Displays the index of an entry in the Token Ring-to-Ethernet Translation table for Ethernet type fields. The index can range between 1 and 20.

### SNAP

Displays the Ethernet type field value which the bridge will use to translate all Token Ring 802.5 packets with a SNAP header and this type field to the format specified by the button. The Ethernet type field value must be two octets.

### State

Displays if this entry will be used in the packet translation process. If the state is **Enabled** the entry is used; if the state is **Disabled** the entry is not used.

## Configuring SNAP Translation

When you choose the SNAP format, you select the Ethernet frame format to which the bridge translates all Token Ring packets with a SNAP header whose Ethernet type resides in the table and whose entry is enabled.

To configure the SNAP Translation Table:

1. Click on the **SNAP Format** button on the Bridge Translation window. It will toggle between the **Enet** and **SNAP** formats. The format that appears in the button is the one that is currently being used.
2. Click on the scroll box next to the entry and the scrollable window of entries will appear.

- Click on the entry you want to change. The SNAP Translation table, [Figure 2-38](#), will appear.

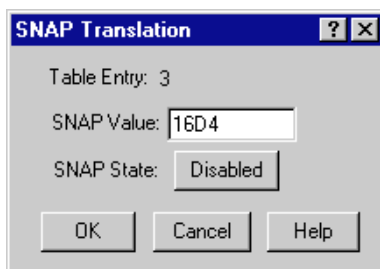


Figure 2-38. SNAP Translation Table

The Table Entry field will display the index of the entry you are viewing and configuring.

- To change the **SNAP Value** (the type field) of this entry, highlight the current value and type a new one. The SNAP value must be two octets.
- To set whether this entry will be used in the translation process, click on the **SNAP State** button to toggle between enabled and disabled.
- Click on **OK** to save your changes to the SNAP Translation table. The new values will now appear in the table.

## Using the Novell Translation Window

The Token Ring-specific Novell Translation window allows you to configure each bridge port for translation of Novell packets that are received and transmitted across a Token Ring bridge.

The need for translation arises from the way that stations on Token Ring, FDDI, and Ethernet media read and write bits. Ethernet transmits and receives each address byte in Least Significant Bit first (LSB) order; Token Ring transmits and receives each address byte in Most Significant Bit first (MSB) order; and FDDI transmits each address byte using MSB, and receives using LSB. Translational bridges must use the appropriate form of LSB/MSB address bit conversion to allow the destination station to properly recognize its address when a packet arrives. This conversion is necessary for both IP and IPX (Novell protocol) packets.

When you select a translation mode, the MAC addresses in the Novell header or in the MAC header of Novell packets received will be converted from most significant bit format to least significant bit format. LLC layer translation converts the MAC address in the Novell header; Data Link Layer Translation converts the MAC address in the MAC header.

To access the Novell Translation window from the Bridge Status window:

1. Click on **Bridge** to display the Bridge menu.
2. Click on **Novell Translation**. The Novell Translation window, [Figure 2-39](#), will appear.

To access the Novell Translation window from the Chassis View window:

1. Click on the **Board Index** of the bridge module of interest; the Board menu will appear.
2. Click on select **Novell Translation**. The Novell Translation window, [Figure 2-39](#), will appear.

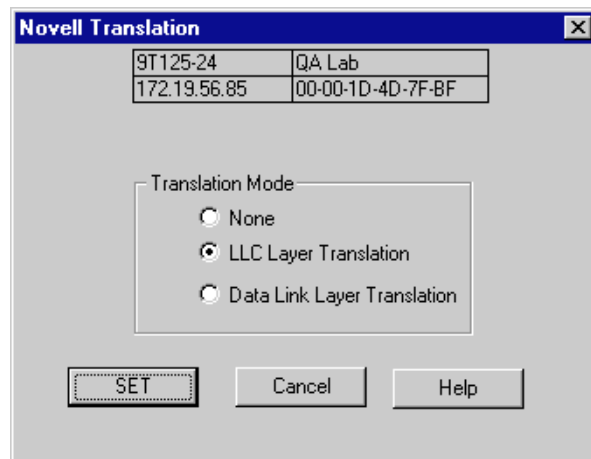


Figure 2-39. Novell Translation Window

Translation modes available are:

- Data Link Layer Translation
- LLC Layer (Logical Link Control) Translation
- None

If **Data Link Layer Translation** is selected, translation is performed within the data link layer for Novell packets. This translation provides most significant bit to least significant bit conversion on the source/destination MAC addresses located in the MAC header.

Devices that support Data Link Layer Translation include:

9T125-08  
9T125-24  
9T425-16  
3T0x-04

If **LLC Layer Translation** is selected, translation is performed within the logical link control layer for Novell packets. This translation provides most significant bit to least significant bit conversion on the source/destination MAC addresses located in the IPX header. The default choice is LLC Layer Translation.

Devices that support LLC translation include:

9T122-08  
9T122-24  
9T425-16  
3T0x-04

If **None** is selected, no translation is performed. This option should be selected when IPX packets are not traveling from Token Ring to either FDDI or Ethernet media.

To select a translation mode:

1. Click on the empty option button adjacent to the desired translation mode.
2. Click **Set** to implement the change.

## Using the Physical View Windows for the ETWMIM

When you are viewing and configuring the bridging capabilities of an ETWMIM, there are a couple special views available. These are described below.

### Ethernet Port Physical View

The Physical View allows you to view the physical state of the Ethernet port when you are monitoring an ETWMIM via SPECTRUM Element Manager.

To use the Physical View option from the Bridge Status window:

1. Click on the Ethernet bridge port (Port 1). The Ethernet bridge port menu will appear.
2. Click on **Physical View**. The ETWMIM EtherPhysStatus (Ethernet Physical Status) window, [Figure 2-40](#), will appear.

To use the Physical View option from the Chassis View window:

1. Click on the Ethernet bridge interface (Port 1). The Ethernet bridge port menu will appear.
2. Click on **Physical View**. The ETWMIM EtherPhysStatus (Ethernet Physical Status) window, [Figure 2-40](#), will appear.

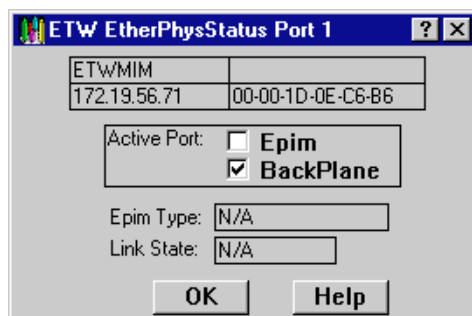


Figure 2-40. Ethernet Port Physical View

The following status fields are available in the Ethernet Port Physical View window:

#### Active Port

This field will have an enabled check box next to the active port configuration option you have selected for your ETWMIM Ethernet port.

- If you have configured the ETWMIM for use with the Ethernet backplane connection, the X will appear in the BackPlane checkbox.
- If you have opted to use a front panel EPIM for your Ethernet connection, the X will appear in the EPIM checkbox.

You cannot change your active port configuration from this window. It must be changed physically on the ETWMIM itself.

#### Epim Type

This field will show the type of EPIM you have installed via the front panel of your ETWMIM, if applicable. The types of EPIMs are listed below, along with the type of segment each will be connected to.

- **EPIM-T**—10BASE-T Twisted Pair Segment
- **EPIM-F1/F2**—Fiber Optic Link Segment
- **EPIM-F3**—Single Mode Fiber Optic Link Segment
- **EPIM-C**—Thin-net segment
- **EPIM-A**—AUI cable segment
- **EPIM-X**—AUI cable segment
- **EPIM Unknown**—The EPIM type cannot be determined
- **N/A**—The backplane connection is being used.

#### Link State

This field will display the link state of the EPIM Ethernet port. The possible states are:

- **Linked**—indicates a link has been established on the EPIM.
- **Unlinked**—indicates a link has not been established on the EPIM.

- **Unknown**—indicates the status of the EPIM link is unknown, or not valid for the type of EPIM installed.
- **N/A**—indicates that the backplane connection is being used.

## Token Ring Port Physical View

The Physical View option allows you to view and configure the physical set up of the Token Ring port when you are monitoring an ETWMIM via SPECTRUM Element Manager.

To use the Physical View option from the Bridge Status window:

1. Click on the Token Ring bridge port (Port 2). The Token Ring bridge port menu will appear.
2. Click on **Physical View**. The ETWMIM Token Ring Phys(ical) Status window, [Figure 2-41](#), will appear.

To use the Physical View option from the Chassis View window:

1. Click on the Token Ring bridge port (Port 2). The Token Ring bridge port menu will appear.
2. Click on **Physical View**. The ETWMIM Token Ring Phys(ical) Status window, [Figure 2-41](#), will appear.

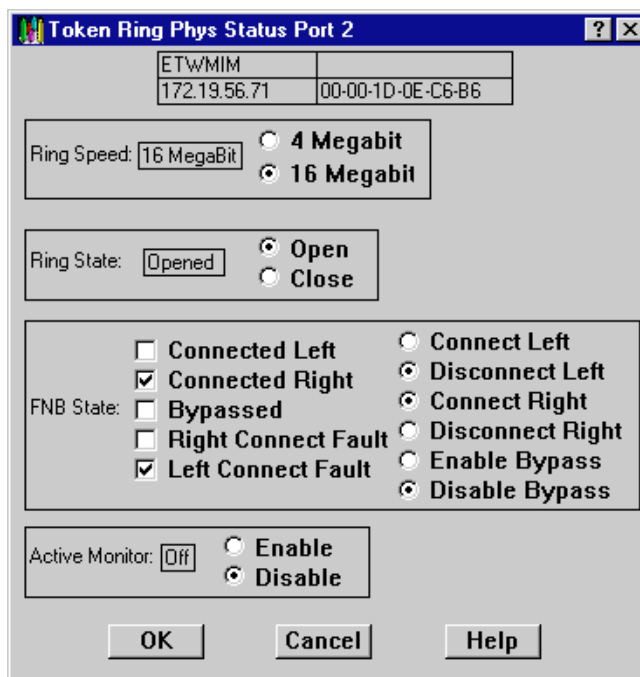


Figure 2-41. Token Ring Port Physical View

The following Status Fields are available in the Token Ring Port Physical View window:

### Ring Speed

Displays the current ring speed configured for your Token Ring port. You can change the ring speed from this window by clicking on the option button next to the desired ring speed: **4 Megabits/second** or **16 Megabits/second**. When you reconfigure the ring speed, the new speed will appear in the text box in this field.

### Ring State

Displays the state of the ETWMIM's Token Ring MAU with respect to the ring. When the ring is "open," the Token Ring MAU is participating in the ring poll process and is receiving and transmitting data onto the ring. When the ring is "closed," the MAU is removed from the ring, and data is not being transmitted or received on the ring. You can change the ring state from this window by clicking on the option button next to the desired option: **Open** or **Close**. If you successfully reconfigure the ring state, the new state will appear in the text box in this field.

### FNB State

The FNB State section displays, and lets you configure, the state of the backplane FNB connectors on the ETWMIM.

The right-hand side of the window displays the current connection configuration for the FNB connectors on the ETWMIM, and lets you alter those options by using the appropriate option button selections:

- **Connect Left** indicates that the ETWMIM is/will be connected on the FNB to the first board to its left in the MMAC chassis with a valid right FNB connection.
- **Disconnect Left** indicates that the ETWMIM is/will be disconnected on the FNB from any boards to its left in the MMAC chassis.
- **Connect Right** indicates that the ETWMIM is/will be connected on the FNB to the first board to its right in the MMAC chassis with a valid left FNB connection.
- **Disconnect Right** indicates that the ETWMIM is/will be disconnected on the FNB from any boards to its right in the MMAC chassis.
- **Enable Bypass** indicates that the ETWMIM is/will be in bypass state. It will not be connected to any boards on its left or right. In a shunting chassis, the FNB will bypass the board to maintain the integrity of the ring across the chassis.
- **Disable Bypass** indicates that the ETWMIM is/will be inserted into the FNB, according to the established FNB connection options above.



The left-hand side of the window indicates the results of the current FNB configuration, with an X next to the appropriate state of the FNB connection: **Connected Left**, **Connected Right**, **Bypassed**, **Right Connection Fault**, or **Left Connection Fault**. For example, if you choose Connect Right and Disconnect Left, then click **OK**, the Connected Right and Left Connect Fault fields will appear with an X next to them.

#### **Active Monitor**

This field allows you to configure whether or not the ETWMIM's onboard management station will engage in the active monitor contention process, which occurs as part of the recovery procedures initiated after certain ring error situations.

If you select **Enable**, the station will contend in the process used to establish a ring station as an Active Monitor.

If you select **Disable**, the station will not contend, even if the contention process is activated for the ring. Note that if the ETWMIM is currently serving as the active monitor, it will continue in that role until the next contention.

The box to the left of the choices will reflect your actions by displaying **On** when the Active Monitor has been enabled, and **Off** when the Active Monitor has been disabled.



# FDDI Applications

*Concentrator Configuration window; Connection Policy window; Station List window; Concentrator Performance window*

---

The FDDI menu, available on devices with an installed and enabled FDDI BRIM, lets you access windows to view information about the Station Management (SMT) entity supported by your FDDI BRIM. The SMT entity provides the system management services for the FDDI protocols, including connection management, node configuration, management statistics, and ring recovery. SMT is composed of various subcomponent functions, including Connection Management (CMT) and Ring Management (RMT).

The submenus that provide FDDI management are:

- **Configuration**, which displays the current configuration of the FDDI BRIM, and the status of the ring; see [The Configuration Window, page 3-2](#), for details.
- **Connection Policy**, which displays the types of connections between the four FDDI PHY (port) types – A, B, M, and S – that will be allowed by the FDDI BRIM; see the [Connection Policy Window, page 3-5](#) for details.
- **Station List**, which displays the configuration of the FDDI BRIM managed ring, including the number of nodes, node addresses (both Canonical and MAC), node class, and current ring topology; see [Station List Window, page 3-8](#) for details.
- **Performance** displays the number of frames transmitted and received on the ring, error and lost frames detected on the ring, and the number of ring initializations; see [FDDI Performance Window, page 3-10](#).

## Accessing the FDDI Menu

1. Click on the **FDDI** menu in your device Chassis View window and click again to select the appropriate Station Management (**SMT**) entity to reveal the following FDDI menu (Figure 3-1).





Figure 3-1. The FDDI Menu

2. Click on the desired selection. When you select one of these options, the associated FDDI window will appear.



*The title bar of the selected FDDI submenu windows will display the index number of the SMT entity for which information is being displayed.*

## The Configuration Window

The Configuration window, Figure 3-2, displays the FDDI BRIM configuration, the operating state of the FDDI ring, and the parameters relating to ring initialization.

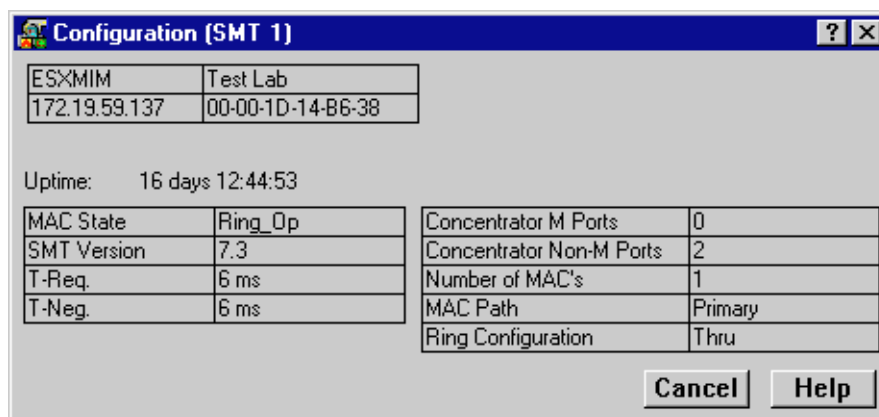


Figure 3-2. The FDDI Configuration Window

The FDDI Configuration window displays the following fields:

### MAC State

This field indicates the current state of the MAC on the FDDI ring associated with the selected SMT entity. The RMT component of SMT monitors MAC operation and takes actions necessary to aid in achieving an operational ring. RMT occurs on a per-MAC basis and aids in the detection and resolution of failures, such as stuck beaconing and the presence of duplicate addresses.

- **Not Available**—there is no MAC on the FDDI ring associated with this SMT entity, or the selected SMT entity is not attached to the main ring through the backplane FNB A and B ports.
- **Ring-Op**—the ring is functioning normally. While in this state, the MAC being managed is part of an operational FDDI ring.
- **Isolated**—the SMT has just initialized RMT or RMT has entered this state during a path test (trace) after ring beaconing; RMT is not aware of the ring path or state.
- **Non-Op**—the MAC being managed by the selected SMT is participating in ring recovery; the ring is not operational.
- **Detect**—the claim (beacon) process of the FDDI ring protocol has exceeded one second. In this state, the ring is still alive, but no data is being transmitted. This may indicate a problem on the ring, including the detection of duplicate address conditions.
- **Non-Op-Dup**—the ring is not operational; the address of the MAC under the control of the SMT entity is a duplicate of another MAC on the ring. The duplicate address condition prevented ring recovery and initialization after a claim and beacon process. This state will not occur unless you are using locally-administered addresses, as factory-set MAC addresses are unique.
- **Ring-Op-Dup**—the ring is operational; however, the address of the MAC under control of the SMT entity has been found to be a duplicate of another MAC on the ring. Corrective actions will be attempted before the duplicate address condition causes ring initialization to fail after the claim and beacon recovery process. Like Non-Op-Dup, this state will not occur unless you are using locally-administered addresses.
- **Directed**—the beacon process did not complete within 7 seconds. The selected SMT has directed the controlled MAC to send beacon frames to notify the other stations that a serious problem exists on the ring, and a Trace state is soon to follow.
- **Trace**—a problem exists on the ring which could not be corrected during the beaconing process, and a Trace has been initiated. During a Trace (or Path Test), the SMT sends a signal that forces its nearest upstream neighbor to remove from the ring and conduct a self-test. If the ring does not recover, each subsequent upstream station will be forced to remove from the ring and conduct self-tests until the problem has been corrected. While the test is being conducted, ring management re-enters the isolated state.

**SMT Version**

Displays the version number of the Station Management (SMT) entity. SMT frames have a version ID field that identifies the structure of the SMT frame Info field. The version number is included in the SMT frame so that a receiving station can determine whether or not its SMT version is able to communicate with the SMT version of another station. Knowing the version number allows the station to handle version mismatches. Each station supports a range of SMT versions. The supported version range is identified within the ietf-fddi MIB by two smfTable attributes: *snmpFddiSMTLoVersionId* and *snmpFddiSMTHiVersionId*. If a received frame is not within the supported version range, the frame is discarded. SMT provides the system management services for the FDDI protocols, including connection management, node configuration, error recovery, and management frame encoding.

**T-Req. (Requested Target Token Rotation Time)**

The token rotation time bid made by the selected SMT entity during ring initialization. Each station detecting that the ring must be initialized begins a claim token process and issues a stream of Claim Frames, which negotiate the value assigned to the Target Token Rotation Time (TTRT). The information field of these frames contains the issuing station's bid for the value of TTRT. Each claiming station inspects incoming Claim frames (from other issuing stations) and either continues its own bid (and removes the competing Claim Frame from the ring) or defers (halts transmission of its own bid and repeats the competing bid) according to the following hierarchy of arbitration:

- A Claim Frame with the lowest TTRT bid has precedence.
- If the values of TTRT are equal, the frame with the longest source address (48 vs. 16 bits) has precedence.
- If both TTRT value and source address length are equal, the frame with the highest address has precedence.

The FDDI BRIM is shipped with a T-Req = 83 msec (earlier versions of firmware) or 6 msec (later firmware versions). T-Req is stored within the MIB in units of nanoseconds (one billionth of a second) rather than milliseconds (one thousandth of a second); SPECTRUM Element Manager converts nanoseconds to milliseconds for display purposes. You can use any SNMP Set Request tool to edit the T-Req value; just remember that you must enter your value in nanoseconds, rather than milliseconds.

**T-Neg. (Negotiated)**

The winning time negotiated in the ring initialization sequence.

**Concentrator M Ports**

This field displays the number of Master (M) ports on the modular concentrator controlled by the FDDI BRIM. A Master port is a port that provides a connection for Single Attachment Station (SAS) devices to the FDDI network.

**Concentrator Non-M Ports**

This field displays the number of non-Master ports (A, B, or S ports) on the modular FDDI BRIM concentrator.

**Number of MACs**

The number of Media Access Control entities present in the FDDI BRIM, indicating the number of ring port pairs. For the FDDI BRIM, this number will be 1.

**MAC Path**

This field displays the configuration of the MAC with respect to the logical ring, as determined by the Connection Management (CMT) portion of SMT. CMT controls the establishment of a media attachment to the FDDI network, the connections with other nodes in the ring, and the internal configuration of the various entities within a node. CMT provides the link confidence test and specifies a Link Error Monitor (LEM) which monitors active links on a per-link basis to ensure that failing links are detected and, if required, removed from the network. Possible values are:

- **Primary 1**—the Primary 1 FDDI ring is being used.
- **Secondary 1**—the Secondary 1 FDDI ring is being used.
- **Primary 2**—the Primary 2 FDDI ring is being used.
- **Secondary 2**—the Secondary 2 FDDI ring is being used.
- **Local**—the MAC is not inserted into a primary or secondary path of a dual ring, but may be connected to one or more other nodes.
- **Isolated**—the MAC has no connection to the ring or other concentrator ports.
- **Unknown or ?**—SPECTRUM Element Manager cannot determine the MAC path for the FDDI BRIM.

**Ring Configuration**

The current configuration of the MAC and physical layers of the A and B ports.

## Connection Policy Window

The SMT Connection Policy determines which types of connections are allowed among the four FDDI port types: A, B, M (Master), and S (Slave). FDDI protocol forbids Master->Master connections; all other connection types are legal, although some are considered to be undesirable.

The following table summarizes the FDDI connection rules:

Table 3-1. FDDI Connection Rules

	A	B	S	M
A	V, U	V	V, U	V, P
B	V	V, U	V, U	V, P
S	V, U	V, U	V	V
M	V	V	V	X

- V – valid connection
- X – illegal connection
- U – undesirable (but legal) connection
- P – valid, but when both A and B are connected to M ports (a dual-homing configuration), only the B connection is used.



*Though technically legal under FDDI connection rules, undesirable connections will cause a twisted or wrapped ring.*

The Connection Policy window, [Figure 3-3](#), lists potential connection types in a “Reject X-Y” format, where X represents a port on the FDDI BRIM, and Y represents the attaching node. A check in the checkbox next to a Connection Policy indicates that it is an illegal connection.

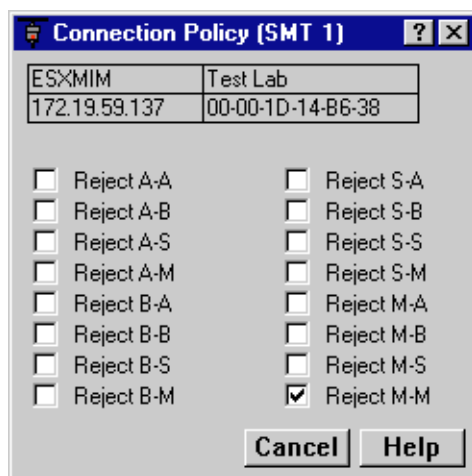



Figure 3-3. The Connection Policy Window





You can use any SNMP Set Request or MIB tool to edit the Connection Policy for your device by setting the `fdimibSMTConnectionPolicy` MIB OID (part of the MIBII FDDI Transmission MIB (RFC1512)). `fdimibSMTConnectionPolicy` is simply a 16-bit integer value (ranging from 32768 to 65535) that corresponds to the connection policy (in the “Reject X-Y” format, where X represents a port on the FDDI Switch Module, and Y represents the attaching node).

To set the connection policy for the device, total the bit values corresponding to the desired connection policy according to the table below, and then use your SNMP Set Request or Mib tool to set the value for the appropriate SMT index. For example, to set a connection policy that disallowed the undesirable A-A or B-B connections you would set the `fdimibSMTConnectionPolicy` MIB OID to 32,801: 32,768 (reject M-M, required) + 32 (reject B-B) + 1 (reject A-A).

Policy	Power
reject A-A	$2^0$ (1)
reject A-B	$2^1$ (2)
reject A-S	$2^2$ (4)
reject A-M	$2^3$ (8)
reject B-A	$2^4$ (16)
reject B-B	$2^5$ (32)
reject B-S	$2^6$ (64)
reject B-M	$2^7$ (128)
reject S-A	$2^8$ (256)
reject S-B	$2^9$ (512)
reject S-S	$2^{10}$ (1,024)
reject S-M	$2^{11}$ (2,048)
reject M-A	$2^{12}$ (4,096)
reject M-B	$2^{13}$ (8,192)
reject M-S	$2^{14}$ (16,384)
reject M-M	$2^{15}$ (32,768 – a permanently set value for this bit)

Each device has its own connection policy; however, when two devices attempt to connect, their combined established connection policies dictate the connections that will be allowed. In an attempted connection between two nodes, the most lenient policy will determine whether the connection (as long as it is legal) can be made. For example, if two FDDI nodes attempt an A→A connection, and this connection is not allowed at one FDDI node but allowed at the other, the connection would be accepted. If the connection policy at both nodes disallows the connection, the connection will be rejected.

This is a read-only window; you currently cannot edit the FDDI BRIM’s connection policy directly from this window.

## Station List Window

The Station List displays the configuration of the FDDI BRIM managed ring, including the number of nodes on the ring, node addresses (both Canonical and MAC), node class, and ring topology.

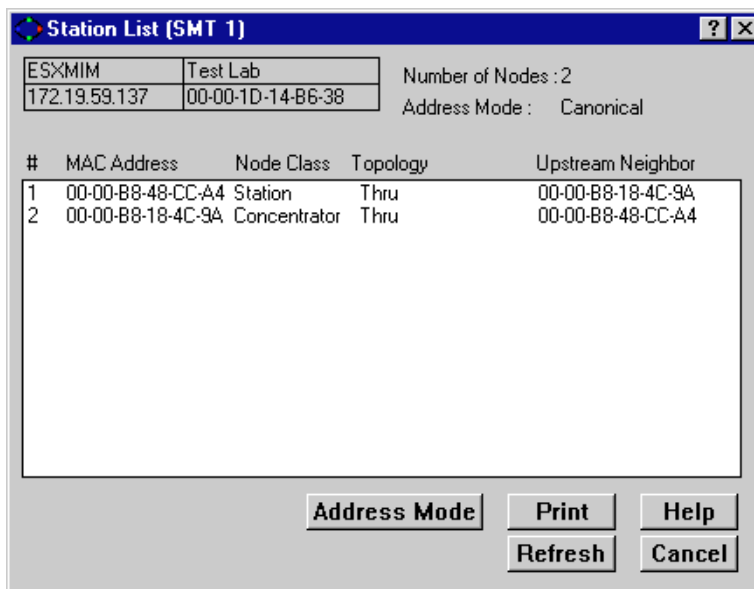


Figure 3-4. The Station List Window

The upper section of the Station List window displays information about the FDDI BRIM managed ring.

### Number of Nodes

The number of stations inserted into the FDDI ring to which the FDDI BRIM MAC is connected.

### Address Mode

Displays the current mode being used to display the addresses of the devices in the Station List. The two possible modes are Canonical (FDDI) or MAC (Ethernet). To change the current Address Mode, click on the **Address Mode** button at the bottom of the window. The current address mode will change in the Address Mode field and the Stations panel.

The **Stations Panel** section displays a list of the stations on the ring to which the selected SMT is connected, in ring sequence beginning with the MAC, along with each station's node class and current topology.

**#**

An index number assigned to each station that indicates its position on the ring in relation to the FDDI BRIM. The monitored FDDI BRIM is always 1.

**MAC Address**

Displays the manufacturer-set MAC address of the node inserted into the ring. MAC addresses are hard-coded into the device and are not configurable.

**Node Class**

Displays the type of ring device. Possible values are:

Station	Indicates an FDDI node capable of transmitting, receiving, and repeating data.
Concentrator	Indicates an FDDI node that provides attachment points to the ring for stations that are not directly connected to the dual ring.

**Topology**

Indicates the node's MAC configuration topology.

**Upstream Neighbor**

Displays the hardware address (in Canonical or MAC format, as currently selected) of each node's upstream neighbor.



*Note that the information displayed in the Station List is static once the window is opened; for updated information, click on the **Refresh** button. If the number of nodes exceeds the panel size, scroll bars will appear in the list box that will allow you to scroll through the station list to view the node of interest.*

## FDDI Performance Window

The FDDI Performance window, [Figure 3-5](#), provides graphical and numeric performance statistics for the FDDI BRIM, including transmit frames, receive frames, frame errors, lost frames, and ring ops.

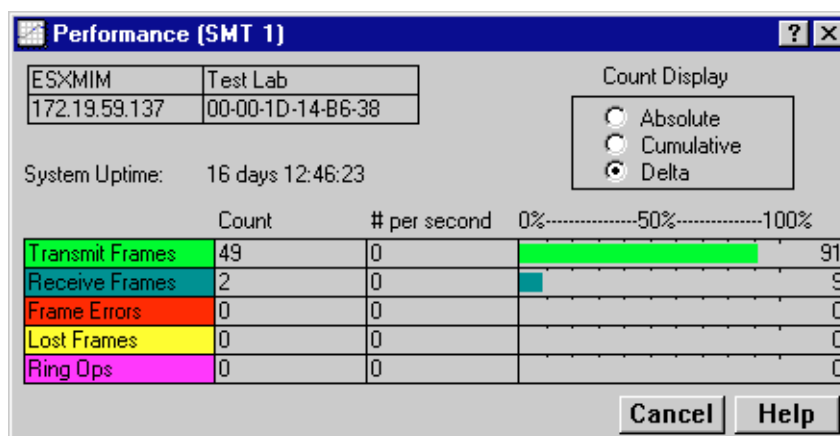


Figure 3-5. The FDDI Performance Window

The FDDI Performance window displays performance data in three formats:

- **Count** the number detected of each performance type for the selected interval.
- **Rate**—# per second the number of each performance type per second, as averaged over the selected interval.
- **Graphically**—0% - 50% - 100% the percentage of each performance type with respect to the total network load processed by the FDDI BRIM during the last interval (e.g., a transmit frames rate of 75% during a delta interval indicates that of all frames *processed* by the selected interface, 75% were *transmitted* by that interface).

The FDDI Performance window displays performance data in the following fields:

### Transmit Frames

The number of frames transmitted by the FDDI BRIM's MAC during the selected interval.

### Receive Frames

The number of frames received by the FDDI BRIM's MAC during the selected interval.

**Frame Errors**

The number of error frames detected by the FDDI BRIM's MAC during the selected interval that had not been detected previously by other stations. Error frames may include frames with an invalid Frame Check Sequence (FCS), with data length errors, or with internal errors that prevent the MAC from transferring the frame to the Logical Link Control (LLC) layer.

**Lost Frames**

The number of frames detected by the FDDI BRIM's MAC during the selected interval that have an unknown error, so their validity is in doubt. When the FDDI BRIM's MAC encounters a frame of this type, it increments the Lost Frame counter and strips the remainder of the frame from the ring, replacing it with idle symbols.

**Ring Ops**

The number of times the ring has entered the "Ring Operational" state from the "Ring Not Operational" state during the selected interval. This counter updates when the FDDI BRIM's MAC informs Station Management (SMT) of a change in Ring Operation status.

## Setting the Time Interval

You can change the time interval during which the performance statistics are collected. To do so, click to select the desired option button in the Count Display panel (in the top right hand corner of the window):

- Absolute—statistical counts recorded since the FDDI BRIM was last started.
- Cumulative—statistical counts recorded since the Performance window was opened.
- Delta—statistical counts recorded during a single polling interval that is set for SPECTRUM Element Manager (refer to the *SPECTRUM Element Manager User's Guide* for more information).



# WAN Applications

*Viewing WAN Interface Status; configuring the synchronous and T1 connection; displaying the WAN logical settings; Viewing the Wan Port AdminStatus; enabling and disabling WAN interfaces; displaying Synchronous port statistics*

---

This chapter describes the options available from the WAN Status menu when a WAN BRIM is installed in a device. This option allows you to set up and view the connections for a WAN BRIM, and view protocol and synchronous port statistics.

The following windows are available from the WAN Status menu:

- The **WAN Interface Status** window displays the configuration settings of the two available WAN ports (Synchronous and T1); see [Viewing WAN Interface Status, page 4-2](#), for details.
- The **WAN Logical View** window displays status information about the logical interface(s) that comprise your physical WAN port; see [Displaying the WAN Logical View, page 4-10](#), for details.
- The **Admin/Status** window displays status information for your WAN port; see [Viewing the WAN Port Admin/Status, page 4-12](#), for more information. Depending on the type of port that is active, the Admin/Status window for the Synchronous port ([page 4-12](#)) or the T1 interfaces ([page 4-13](#)) will display.
- The **Synchronous Port Statistics** window displays the traffic going through the Synchronous port on your WAN BRIM. This menu option will only be available when the Synchronous Port is configured as the active port. See [Displaying Synchronous Port Statistics, page 4-14](#), for details.

## Accessing the WAN Status Windows

1. Click on **Device** in the Chassis View menu bar to access the Device menu.
2. Click on **WAN Status** and then right to the appropriate selection. When you select one of these options, the associated WAN Status window will appear.

## Viewing WAN Interface Status

The WAN Interface Status window displays the configuration of the synchronous and/or T1 ports on a BRIM-W6 or BRIM-WT1. This window has two port configuration sections – one for each WAN port on the device, whether it is synchronous or T1. You can use this window to determine which WAN port is the active port and set a port to be the primary port. You can also use the **Configure** button to access the T-1 Port and Synchronous Port Configuration windows.

To open the WAN Interface Status window:

1. Click on **Device** on the Chassis View menu bar to access the Device menu.
2. Click on **WAN Status** and then on **Physical View**. The WAN Interface Status window, [Figure 4-1](#), will appear.




Figure 4-1. WAN Interface Status Window

The WAN Interface Status window displays the following fields:

### Type

Displays the type of WAN port: **Synchronous** or **T1**.

### Current Active

Displays which WAN port is currently being used on the device. A check will appear in the box under Port 1 or Port 2.


### Current Primary

Displays which port has been configured to be the Primary port.



**Desired Primary**

Allows you to assign a port to be your primary port by clicking on the option button under the desired port.



*In future releases, you will be able to configure redundancy for the device. At this time, by choosing the Desired Primary, you are actually choosing the Current Active.*

## Configuring the Synchronous Connection

The Sync Port Configuration window displays the current configuration of your synchronous WAN connection. The drop-down menus allow you to change the configuration — Type, Inspeed, and Outspeed — of the synchronous port.

To access this window from the WAN Interface Status window:

1. Click on the **Configure** button that appears under the Port 1 - Synchronous column. The Sync Port Configuration window, [Figure 4-2](#), will appear.

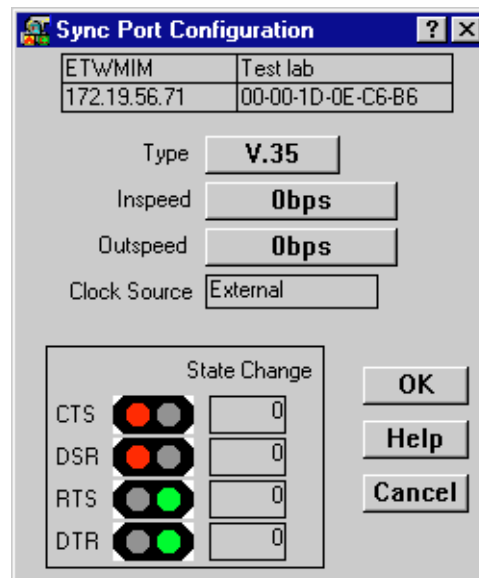


Figure 4-2. Sync Port Configuration Window

The Sync Port Configuration window contains the following fields:

**Type**


Displays the type of connection used at the port. Click the associated button to select **RS422**, **RS232**, **V.35**, and **Other**. You can change the type of connection displayed in this window to match the configuration of your physical synchronous port.

**Inspeed**

Displays the speed of input communications of the synchronous port. Click the associated button to select **2.048 Mbps**, **1.54 Mbps**, **256 Kbps**, **128 Kbps**, **64 Kbps**, **56 Kbps**, **38.4 Kbps**, **19.2 Kbps**, **14.4 Kbps**, and **9.6 Kbps**. The selected inspeed will appear in the button in bits per second (bps) format.

**Outspeed**

Displays the speed of output communications of the synchronous port. Click the associated button to select: **2.048 Mbps**, **1.54 Mbps**, **256 Kbps**, **128 Kbps**, **64 Kbps**, **56 Kbps**, **38.4 Kbps**, **19.2 Kbps**, **14.4 Kbps**, and **9.6 Kbps**. The selected outspeed will appear in the button in bits per second (bps) format.



*The Inspeed and Outspeed are determined by the CSU/DSU attached to the synchronous port. If you change the Inspeed or Outspeed displayed in this window, it will not have an effect on the actual transmit or receive speed of the connection. If the communication speeds do change, you should change the values in this window to reflect these changes.*

**Clock Source**

Displays the source of the port's bit rate clock. The clock source will always be **External** on a synchronous port.

**State Change**

The bottom half of this window displays the traffic for individual signals on your synchronous port. The traffic signals are: **CTS** (Clear to Send), **DSR** (Data Set Ready), **RTS** (Request to Send), and **DTR** (Data Terminal Ready). The red and green indicators reflect whether the indicated signal is disabled (off) or enabled (on), respectively. The **State Change** text box to the right of each of these fields will display the number of times the signal has changed state from enabled to disabled, or vice versa.

To modify the synchronous port parameters:

1. Click on the **Type**, **Inspeed**, or **Outspeed** command buttons.
2. From the Type, Inspeed, or Outspeed drop-down menus, select the new parameter.

## Configuring T-1 Ports

The T-1 Port Configuration window displays the configuration of your WAN T1 connection. The drop-down menus in this window allow you to change the configuration of the T1 connection without accessing Local Management. You can use the **FracTable** button to access the WAN FracTable Configuration window, which allows you to configure your timeslots.

To access the T-1 Port Configuration window from the WAN Interface Status window:

1. Click on the **Configure** button that appears under the Port 2 - T1 column. The T-1 Port Configuration window, [Figure 4-3](#), will appear.

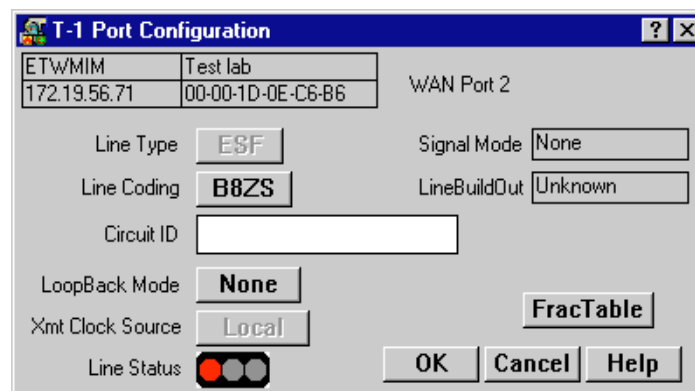


Figure 4-3. T-1 Port Configuration Window

The T-1 Port Configuration window contains the following fields:

### WAN Port Name

Displays the unique value for the T1 port on the BRIM. In this example the value is **WAN Port 2**.

### Line Type

Displays the type of service you are using over your T1 line. Click on the associated button to set the Line Type per your WAN service provider's instructions: **ESF** (Extended Super Frame DS1) or **D4** (AT&T D4 format DS1). In [Figure 4-3](#), the Line Type field is grayed out because **D4** is not supported on ETWMIM devices.

### Line Coding

Displays the line coding for the Full T1 line. Click on the associated button to set the Line Coding per your WAN service provider's instructions: **JBZS**, **B8ZS**, **AMI**, and **None**. The default for this field is **B8ZS**.



If AMI is chosen, the line code selection on the individual interfaces should **not** be set to **None**; for more information on setting the interface line coding refer to [Configuring the Fractional Table](#), page 4-8.

#### **Circuit ID**

Allows you to enter a character string specified by the circuit vendor as a circuit identifier. This is used for a reference during troubleshooting processes.

#### **Loopback Mode**

Displays the loopback configuration of the DS1 interface. Click on the associated button to select **No Loop**, **Payload**, or **Line Loop**.

#### **Xmt Clock Source**

Displays the T1 Transmit Clock Source. Click on the associated button to select **Loop Timing**, **Local Timing**, or **Through Timing**. The default setting is Loop Timing, which indicates that the recovered receive clock is used as the transmit clock; Local Timing indicates that an internal clock source is used. Through Timing is not supported by the WAN BRIM at this time.

#### **Line Status**

Indicates the status of the T1 circuit. Green indicates that there are no alarms present. Red indicates that data communications have been disrupted at the monitored device, due to a loss of frame synchronization or a loss of signal which has occurred for more than two to three seconds. A yellow condition indicates that the far end equipment has experienced a red alarm condition, and subsequently reported it to the near end.

#### **Signal Mode**

Displays the type of signaling that is in use on the T1 line. The possible signal modes are **None**, **robbedBit**, or **messageOriented**. None indicates that no bits are reserved for signaling on this channel. When T1 Robbed Bit Signaling is in use, this field will display robbedBit. When Common Channel Signaling is in use on channel 24 of a T1, this field will display messageOriented.

#### **LineBuildOut**

Displays the value of the Line Buildout setting. This setting controls the amount of attenuation of the T1 signal. You can set this value via Local Management. The possible settings are 0 db, -7.5db, and -15 db; the default is zero. Consult you local telephone carrier, before you change this value.

## Using the T1 FracTable Configuration Window

The FracTable Configuration window allows you to assign your interfaces to the 24 channels on the T1 line according to the mapping provided by your WAN carrier.

To open the T1 FracTable Configuration window from the T1 Port Configuration window:

1. Click on the **FracTable** button. The T1 FracTable Configuration window, [Figure 4-4](#), will appear.

This figure illustrates an example configuration of a T1 connection that was set according to the mapping provided by the wide area carrier.

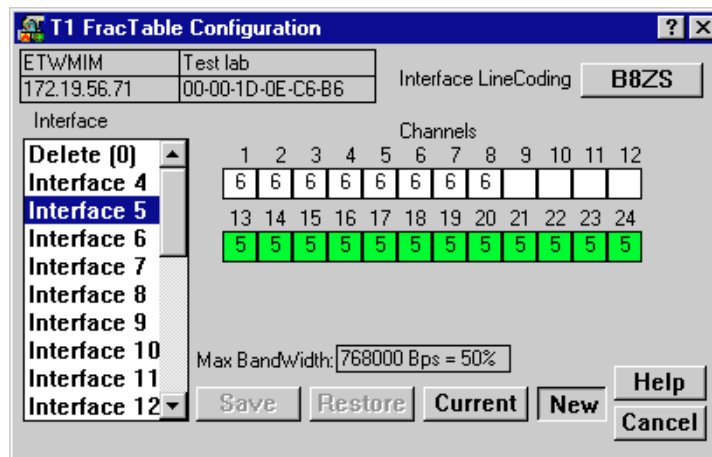


Figure 4-4. WAN T1 FracTable Configuration Window

The T1 FracTable Configuration window displays the following information:

### Interface LineCoding

Displays the Line Coding set for this interface. Click on the associated button per your WAN service provider's instructions to select: **JBZS**, **InvHDLC**, and **None**.



*None* (the default value) should **not** be selected when you are using AMI line coding on the T1 line; for more information on setting the T1 line coding refer to [WAN Port Name](#), page 4-5.

### Interface

Displays the *ifIndex* interface index value for each interface, which is a unique identifier for each physical and logical interface on the device. You use the *ifIndex* to map the interface on the fractional table of a T1 port.

The range of interface indices for a BRIM's interfaces will start after the indexing of interfaces on the host device, and end at the last interface supported by the WAN port. For example, for an ETWMIM that has a BRIM-W6 with a single 24-channel WPIM-T1 installed in the first BRIM port, the BRIM interfaces will be indexed from 3 (since *ifIndex* 1 and 2 are the ETWMIM Ethernet and Token Ring interfaces) to 26 (since the WPIMs supports 24 interfaces). If there was a WPIM-SY Synchronous port installed in the first BRIM port, the T1 port interfaces would be indexed from 4 to 27 (since *ifIndex* 1 and 2 are the Ethernet and Token Ring interfaces, and *ifIndex* 3 is the Synchronous port interface).

### Channels

Displays the 24 channels or timeslots available for interface assignments.


### Max Bandwidth

Displays the bandwidth you have assigned to the current interface. This field also contains a percentage of the total T1 line bandwidth that the interface is using. The bandwidth will not reflect your configuration until it has been saved.

## Configuring the Fractional Table

When you open the *FracTable* Configuration window, the current configuration of the fractional table will be displayed. The channels that are currently assigned to the selected interface will be highlighted green.


1. Click on the **New** button and use a "scratchpad" copy of the current configuration to make changes on. When you use the "scratchpad" you can switch back and forth to compare the current configuration that is being used by the device and the configuration that you have set up on the "scratchpad". You switch back and forth between the **New** and **Current** templates by clicking on the **New** and **Current** buttons at the bottom of the window.
2. Click to select the interface that you wish to assign to the channel(s). The interface will appear highlighted in blue. In the preceding illustration, Interfaces 5 and 6 have been assigned to the device's channels.



*We recommend that you use consecutive interfaces starting at the first available interface.*

3. Click on the boxes under the numbers of the channels that your wide area carrier has assigned. The interface number will then appear in the box, and the box will be colored in turquoise. In the preceding illustration, the wide area carrier assigned channels (timeslots) 1-8 and 13-24 for use.
4. Repeat steps 2 and 3 for any other interfaces you want to assign to channels.

5. If a channel is not being used by an interface you should disable it by selecting **Delete (0)** in the interface scroll box and then clicking on the selected channel. In the preceding window, channels 9-12 are not assigned and are disabled.
6. When you are done making your changes click on the **Save** button. The “scratchpad” will then be copied to the current fractional table being used by the device. A confirmation window will appear asking you to confirm the choice to write over the current configuration.



*After you have saved a configuration or when you first open this window, the channels that are currently assigned to the selected interface will be highlighted in green.*

### Restoring a Fractional Table

If you have made changes to the “scratchpad” table, but have not saved these changes, you can revert to the current fractional table and erase the changes you have made to the New or “scratchpad” table.

To revert to the current fractional table:

1. Click on the **Restore** button. The fractional table currently being used by the device will appear, and is copied to the New or “scratchpad” table.

### Changing the Interface Line Coding

You can change the line coding used by each individual interface on your device.

1. Select the interface for which you want to change the line coding.
2. Click on the **Interface Line Coding** button. A drop-down menu will appear.
3. Select the line coding you want to use on that interface.



***None** (the default value) should **not** be selected when you are using AMI line coding on the T1 line; for details on setting the T1 line coding, see [WAN Port Name](#), page 4-5.*

## Displaying the WAN Logical View

The WAN Logical View window displays information about the logical interfaces that comprise your physical WAN port. Although the WAN Logical View window has identical fields for the T1 and the Synchronous ports, the number of interface entries depends on the type of port. The Synchronous port will have a single interface entry; the T1 port will have 24 interface entries.

To open the WAN Logical View window:

1. Click on **Device** in the Chassis View menu bar to access the Device menu.
2. Click on **WAN Status** and **Logical View**. The WAN Logical View window, [Figure 4-5](#), will appear.

IF	Protocol	Compression	MTU	Line Coding	CRC Length
4	None	Off	0 bytes	None	16 bits
5	None	Off	0 bytes	None	16 bits
6	None	Off	0 bytes	None	16 bits
7	None	Off	0 bytes	None	16 bits
8	None	Off	0 bytes	None	16 bits
9	None	Off	0 bytes	None	16 bits
10	None	Off	0 bytes	None	16 bits
11	None	Off	0 bytes	None	16 bits
12	None	Off	0 bytes	None	16 bits
13	None	Off	0 bytes	None	16 bits
14	None	Off	0 bytes	None	16 bits
15	None	Off	0 bytes	None	16 bits
16	None	Off	0 bytes	None	16 bits
17	None	Off	0 bytes	None	16 bits

Figure 4-5. The WAN Logical View Window (T1)

The WAN Logical View window displays the following fields:

### IF

Displays the interface index; a unique value for each logical interface supported by this device.

### Protocol

Displays the active Link Layer protocol: **PPP** (Point to Point), **Frame Relay**, **LEX**, or **None**.

### Compression

Display if data compression is **On** or **Off**. Data compression is not supported by the WAN BRIM at this time; compression will always be de-activated or **Off**.



**MTU**


Displays the MTU (Maximum Transfer Unit) for this interface. The MTU is the largest packet size that can be transmitted on the selected interface.

**Line Coding**

Displays the line coding set for this interface: **INV-HDLC**, **JBZS**, or **None**. None (the default value) is displayed when the line coding being used on the interface is B8ZS.

**CRC Length**

Displays the length of the CRC (Cyclical Redundancy Check) for this interface.



The information in this window is static; use the **Refresh** button to view updated information.

## Changing WAN Logical Settings

You can change the protocol setting from the WAN Logical View window.

1. In the list box, click on the interface line of interest. The selected interface line will be highlighted blue and the WAN Logical Settings window, as shown in [Figure 4-6](#), will appear.




Figure 4-6. The WAN Logical Settings Window

2. Click on the Protocol button to select **PPP**, **Frame Relay**, **LEX** or **None**.
3. Click on **OK** to save changes and exit the window, or **Cancel** to exit the window without saving changes.

After exiting the Logical Settings window, the WAN Logical View window will update with the changes you made.

## Viewing the WAN Port Admin/Status

The Admin/Status window displays status information for your WAN port. Depending on the type of port that is active, the Admin/Status window for the Synchronous Port (Figure 4-7) or the T1 interfaces will display (Figure 4-8).

To access the Admin/Status window:

1. Click on **Device** on the Chassis View menu bar to access the Device menu.
2. Click on **WAN Status** and **Admin/Status**. The AdminStatus window for the Synchronous Port (Figure 4-7) or T1 Interfaces (Figure 4-8) will appear.

### Synchronous Admin/Status

The Synchronous Admin/Status window displays the operational state of your Synchronous connection, and allows you to enable or disable the port.

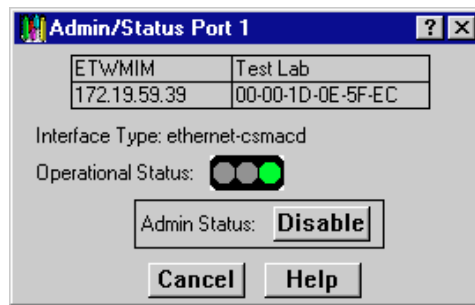


Figure 4-7. AdminStatus (Synchronous) Window

The Admin/Status Synchronous Port window displays the following fields:

#### Interface Type

Displays the interface type of the port (e.g., **ethernet-csmacd**).

#### Operational Status

Displays the operational status of the port, as indicated by a traffic light display:



Green indicates the interface's operational status is Up.



Red indicates that the interface's operational status is Down. No data can be received or transmitted by the interface. Note that an interface can be operationally down, even though it is administratively enabled (i.e., the interface has experienced an operational failure, regardless of its administrative state).



Yellow indicates that the port is in a test mode, or in some transitional state between the disabled and enabled states.

### Admin Status

The Admin Status toggle button lets you administratively **Enable** or **Disable** the port. The Operational Status indicator will reflect the results of the administrative action.

## T1 Admin/Status

The T1 Admin/Status window (Figure 4-8) displays the administrative status or operational status of your T1 interfaces, and also allows you to enable or disable any of the 24 possible T1 interfaces.

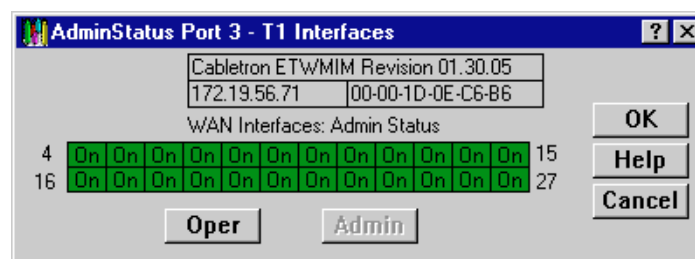


Figure 4-8. AdminStatus T1 Interfaces Window

The AdminStatus T1 Interface window displays the following fields:

### WAN Interfaces

Indicates whether the administrative (Admin Status) or operational (Oper Status) state is in effect for the interface display. The corresponding command button will be grayed-out.

### Interface Display

Graphically displays the status of each of the 24 logical interfaces on your T1 connection as selected by the command buttons, either **Oper** or **Admin**. If Oper is selected, each interface will display its actual, operational status: On (green) if it is up, or Off (red) if it is down. If Admin is selected, each interface will display its administrative status as set by management: On (green) if it has been administratively enabled, or Off (red) if it has been administratively disabled.

You can also access a menu from each interface to administratively enable or disable the interface.

### Oper and Admin

The **Oper** and **Admin** buttons let you change the interface display between Oper Status and Admin Status. The command button of the active status will be grayed-out and inactive.

## Enabling and Disabling WAN T1 Interfaces

From the Interface drop-down menus, you can administratively enable or disable any of the 24 possible T1 interfaces.

1. From the WAN AdminStatus T1 Interfaces window, click on the desired **Interface** button.
2. Click on **Enable** to enable the interface, or **Disable** to disable the interface. Note that the interface display may take a few moments to update your selections.

## Displaying Synchronous Port Statistics

The Sync Port Statistics window displays statistics for the traffic going through your synchronous WAN connection. Statistics are displayed in numeric form and in graphic form as a pie chart. The Statistics option is only available when the Synchronous Port is active.

To access the Sync Port Statistics window for the synchronous port:

1. Click on **Device** in the Chassis View menu bar to access the Device menu.
2. Click on **WAN Status** and **Statistics**. The Sync Port Statistics window, as shown in [Figure 4-9](#), will appear.

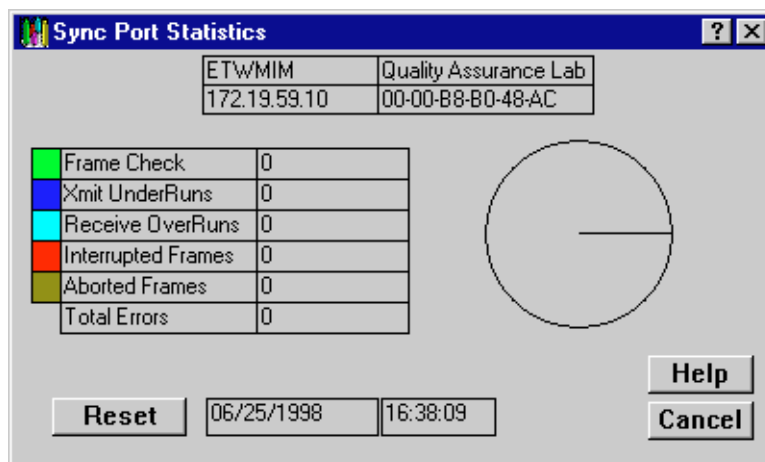


Figure 4-9. Sync Port Statistics

The Sync Port Statistics window displays the following fields:

### Frame Check

The number of frames with an invalid frame check sequence input from the port. Frame Check statistics are color-coded Green.

**Xmit UnderRuns**

The number of frames that failed to be transmitted on the port because data was not available to the transmitter in time. Xmit UnderRuns are color-coded Blue.

**Receive OverRuns**

The number of frames that failed to be received on the port because the receiver did not accept the data in time. Receive OverRuns are color-coded Cyan.

**Interrupted Frames**

The number of frames that failed to be received or transmitted on the port due to loss of modem signals since port state was "up". Interrupted Frames are color-coded Red.

**Aborted Frames**

The number of frames aborted on the port due to receiving an abort sequence. Aborted Frames are color-coded Brown.

**Total Errors**

The total number of frames with errors that have passed through the port.



*The **Reset** button will reset all the Statistics counters in this window to zero and restart counting statistics from the time of reset.*



# ATM Configuration

*Viewing connection data; configuring Permanent Virtual Circuits (PVCs); adding and deleting connection entries*

---

The ATM Connections option will be available when you have an ATM BRIM installed and enabled in your device. The ATM interfaces provided by an ATM BRIM provide the connectivity that allows you to merge ATM network segments with traditional LAN technologies.

An ATM network uses two types of virtual channels, or circuits: Switched Virtual Circuits, or SVCs, and Permanent Virtual Circuits, or PVCs. SVCs are created and dismantled dynamically on an as-needed basis, and require no management definition. PVCs, however, must be manually configured. The Current ATM Connections window provides the means for accomplishing these configurations.

## Accessing the ATM Connections Window

To open the ATM Connections window:

1. Click on **Device** on the Chassis View menu bar to access the Device menu.
2. Click on **ATM Connections**. The Current ATM Connections window, [Figure 5-1](#), will appear.

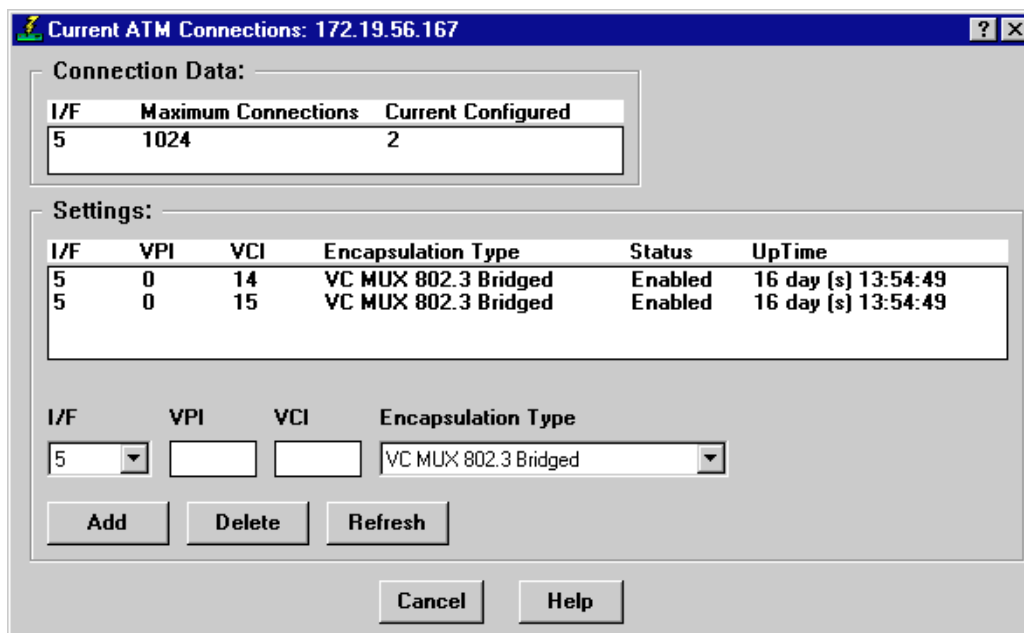


Figure 5-1. The Current ATM Connections Window

The Current ATM Connections window provides the following information:

**Connection Data**

The Connection Data fields provide the following information about each ATM interface available on the device:

- |                     |   |
|---------------------|---|
| I/F                 | Displays the index number assigned to each ATM interface present on the selected BRIM device. |
| Maximum Connections | Displays the maximum number of connections allowed by current device firmware.                |
| Current Configured  | Displays the number of Permanent Virtual Circuits, or PVCs, currently configured.             |

**Settings**

The Settings portion of the window contains a list box which displays information about each of the currently configured PVCs, as well as the fields used to configure new connections:

- |     |  |
|-----|--|
| I/F | Displays the device interface on which the PVC was configured. |
|-----|--|



VPI	Displays the Virtual Path Identifier assigned to the connection. Virtual Path Identifiers are used to group virtual connections, allowing for channel trunking between ATM switches. Each VPI can be configured to carry many different channels (designated by VCIs) between two points.
VCI	Displays the Virtual Channel Identifier assigned to the connection; allowable values are 0 - 1023 for each VPI. Remember, it is the combined VPI and VCI designations assigned to a channel that creates the grouping of virtual connections.
Encapsulation Type	Displays the method used to encapsulate LAN packets on the selected circuit. Current versions of ATM BRIM firmware use 802.3 VC-based multiplexing for bridging protocols (designated VC MUX 802.3 Bridged); future versions will add support for ATM Forum LAN Emulation and Cabletron's SecureFast Switching. You may also see some connections assigned a type of "other"; these are default connections that cannot be modified or deleted.
Status	Displays the current administrative status of the connection: <b>Enabled</b> or <b>Disabled</b> . For current versions of firmware, all connections are enabled by default, and cannot be disabled.
UpTime	Displays the length of time the connection has been enabled.

## Configuring Connections

You can use the command buttons (**Add**, **Delete** or **Refresh**) at the bottom of the Current ATM Connections window to add, modify or delete a Permanent Virtual Circuit (PVC) or refresh the window.


Add	Adds a new connection or modifies an existing one, using the parameters entered in the fields below the list box. A confirmation window will appear for additions and modifications.
-----	--

Delete	Deletes the selected connection; a confirmation window requires that you confirm the deletion.
Refresh	Refreshes the connection information displayed in the window.

## Adding a New Connection


To configure new Permanent Virtual Circuits (PVCs), enter the following information in the text fields which appear just below the connections list box:

1. In the **I/F** menu, select the interface for which you wish to configure a connection. All available ATM interfaces will be listed in this menu. Note that depending on the type of your host device, you will have a fixed number of interface selections:
  - EMM-E6: 5-20
  - ESXMIM: 7-22
  - ESX-1320/ESX-1380: 13-28
  - MicroMMAC: 2-14
  - NBR-420: 4-19
  - NBR-620: 6-21
2. In the **VPI** text box, enter the Virtual Path Identifier you wish to assign to this connection. Currently, the only allowable value is 0; remember, the VPI you assign will be used to group virtual connections, allowing for channel trunking between ATM switches.
3. In the **VCI** text box, enter the Virtual Channel Identifier you wish to assign to this connection. Allowable values are 0 to 1023 for each VPI. Again, remember that it is the combination of VPI and VCI that will be used to direct cells through the intermediate switches between the source and destination.




Currently, versions of Cabletron's ATM device firmware restrict the range of VPI and VCI values that you can set. VPI values for the BRIM-A6 and BRIM-A6DP are restricted to a value of 0; VCI values are restricted to a range of 0-1023.

4. In the **Encapsulation Type** menu, select the desired encapsulation type. Current versions of ATM BRIM firmware use 802.3 VC-based multiplexing for bridging protocols (designated VC MUX 802.3 Bridged); future versions will add support for additional encapsulation methods.



Selecting any of the other encapsulation types listed in the field's menu will cause a "Set Failed" error when you attempt to add the new connection.

5. Click the **Add** button to add the new permanent circuit to the ATM interface. The circuit is automatically enabled, and will remain in place until it is manually removed.



*Note that for the BRIM-A6DP which supports a redundant APIM, any PVCs that you configure will automatically be created on both APIMs.*

## Deleting a Connection

To delete an existing PVC:

1. In the connections list box, click to select the connection you wish to delete.
2. Click on the **Delete** button. A confirmation window will appear, listing the parameters assigned to the connection and asking you to verify that you wish to delete it. Click on **OK** to proceed with the deletion, or on **Cancel** to terminate the deletion.



**A**

- Aborted Frames 4-15
- Absolute 3-11
- Accessing Other Management Options
  - Duplex Modes window 2-10
  - Ethernet Special Filter Database 2-10
  - Filtering Database window 2-10
  - Module Type window 2-9
  - Performance Graph 2-9
  - Spanning Tree window 2-10
- Active Monitor 2-103
- Active Port 2-100
- Active Users 1-7
- Address 2-21, 2-44
- Address Mode 3-8
- Address Mode button 3-8
- AdminStatus 4-13
- AdminStatus button 4-13
- AdminStatus window 4-12
  - Synchronous 4-12
  - T1 4-13
- Aging Time 2-44
- Alarm Configuration window 2-12
- All Paths Explorer (APE) packet 2-2
- ATM 5-1
- ATM Connections
  - Configuring 5-3
- Auto Translate mode 2-92

**B**

- Bit Interleaved Parity 2-84
- Board Number 1-7
- Bridge Address 2-9
- Bridge Configuration window 2-6, 2-11, 2-65
- Bridge Detail Breakdown window 2-18
- Bridge Performance Graph window 2-15
  - fields 2-16
- Bridge Port Detail Breakdown window 2-19
- Bridge Port Level Fields 2-37
- Bridge Priority 2-35
  - Changing 2-39
- Bridge Protocol Data Units (BPDUs) 2-2

- Bridge Spanning Tree window 2-34, 2-35
  - Changing parameters 2-39
- Bridge State on Interface 2-9
- Bridge Status 1-3
- Bridge Status window 2-5, 2-7
  - Accessing Other Management Options 2-9
  - Information Fields 2-8
- Bridge Translation
  - Setting 2-65
- Bridge Translation window 2-91, 2-92
- Bridge/Router Interface Modules (BRIMs) 1-1
  - BRIM
    - FDDI 3-6, 3-10
    - supported 1-1
  - BRIM Management Applications 1-3
  - BRIM-A6 1-1
  - BRIM-A6DP 1-1
  - BRIM-E100 1-2
  - BRIM-E6 1-1
  - BRIM-F6 1-2
  - BRIMs 1-1
  - BRIM-W6 1-2
- Broadcast packet 2-2
- Broadcast Suppression
  - window 2-88
- Broadcast suppression 2-88
- Broadcast Suppression window 2-7

**C**

- Cancel button 1-7
- Capacity 2-44
- Channel trunking 5-2
- Channels 4-8
- Circuit ID 4-6
- Claim token process 3-4
- Clock Source 4-4
- CMT 3-5
- Coding Violations 2-84
- Command buttons 1-7
- Common Window Fields 1-6
- Compression 4-10
- Concentrator 3-9

- Concentrator M Ports 3-4
- Concentrator Non-M Ports 3-5
- Configuration 3-1
- Configuration window 2-12
- Configure button 4-3, 4-5
- Connect A 2-71
- Connection Management 3-5
- Connection Policy 3-1
  - window 3-6
- Connection rules 3-6
- Connection Type window 2-11
- CRC Length 4-11
- CSMACD Statistics window 2-5
- CSMACD Stats window 2-13
- Cumulative 3-11
- Current Active 4-2
- Current button 4-8
- Current Primary 4-2

## D

- Data Mask 2-50, 2-51
- Data Offset 2-50, 2-51
- Data Type 2-50, 2-51
- Delta 3-11
- Description 2-21
- Description window 2-11
- Designated Bridge 2-38
- Designated Cost 2-38
- Designated Port 2-38
- Designated Root 2-38
- Desired Primary 4-3
- Destination Address 2-50, 2-51
- Detect 3-3
- Device Name 1-6
- Directed 3-3
- Disable bridge port 2-14
- Discarded 2-22
- Dot5 Error Statistics window 2-5
- Dot5 Errors window 2-12
- Dual Translate mode 2-92
- Dual-homing 3-6
- Duplex Modes 2-65
  - Setting 2-72
- Duplex Modes window 2-6, 2-71
- Dynamic entries 2-42

## E

- Enable 2-49
- Enable Auto Translation Mode 2-92
- Encapsulation Type 5-3
- Entries
  - Clearing All 2-47
- Epim Type 2-100
- Error 2-22
- Errored Seconds 2-84
- Errors 2-17, 2-19, 4-15
- Ethernet Port Physical View 2-100
- Ethernet Special Filter Database window 2-6
- Explorer packet 2-2

## F

- far-end block error 2-83
- FDDI
  - BRIM 3-10
  - FDDI Configuration window 3-2
  - FDDI connection rules 3-6
  - FDDI Menu 1-4, 3-2
  - FDDI Performance window 3-10
    - Statistics 3-10
  - FDDI protocol 3-5
  - FE100-Sx 2-79
  - File 2-49
  - Filter
    - Enabling and Disabling 2-53
  - Filter Address 2-46
  - Filter Database
    - New Filter Window 2-46
  - Filtered 2-17, 2-19
  - Filtering Database 2-1, 2-41
    - Configuring 2-45
    - window 2-43
  - Filtering Database window 2-6
  - Filters
    - Saving a Set 2-53
  - Filters button 2-53
  - FNB State 2-102
  - Forwarded from 2-20
  - Forwarded to 2-20
  - Forwarding Delay 2-37
  - Forwarding Delay Time
    - Changing 2-40
  - FracTable button 4-7
  - Fractional Table 4-7
    - Configuring 4-8
    - Restoring 4-9

Frame Check 4-14  
Frame Errors 3-11  
Frame Relay 4-11  
Frames Forwarded 2-17, 2-19  
Full Duplex 2-71

## G

Getting Help 1-8  
Global Call Center 1-9  
Grouping of virtual connections 5-3

## H

Hello Time 2-37  
    Changing 2-40  
Help button 1-7, 1-8  
Hold Time 2-37

## I

I/F Configuration window 2-6, 2-11  
I/F Statistics window 2-12  
IBM Translation Table 2-94  
IF 4-10  
Index 2-95, 2-96  
Individual Entries  
    Adding or Deleting 2-46  
Inspeed 4-4  
Instance 2-86  
Interface 2-54, 2-71, 4-7  
Interface button 4-14  
Interface Line Coding 4-7  
    Changing 4-9  
Interface Line Coding button 4-9  
Interface Statistics window 2-5, 2-20  
Interface Type 2-9, 4-12  
Interrupted Frames 4-15  
IP Address 1-6  
Isolated 3-3, 3-5

## L

Learned Database 2-42  
Learned entries 2-42  
LEX 4-11  
Line Coding 4-5, 4-11  
Line Status 4-6  
Line Type 4-5  
LineBuildOut 4-6  
Link State 2-100  
List 2-44

Local 3-5  
Location 1-6  
Logical View 4-10  
Loopback Mode 4-6  
Loss of Frame 2-83  
Loss of Pointer 2-82  
Loss of Signal 2-82  
Lost Frames 3-11  
LSB 2-94

## M

MAC Address 1-7, 3-9  
MAC Path 3-5  
MAC State 3-3  
Master 3-5  
Max Age 2-37  
Max Age Time  
    Changing 2-40  
Max Bandwidth 4-8  
Mode 2-86  
Mouse usage 1-4  
MSB 2-94  
MTU 4-11

## N

N/A 2-72  
Name 2-86  
Network design 2-34  
New button 2-46, 4-8  
New Filter Window 2-46  
Node Class 3-9  
Non-Op 3-3  
Non-Op-Dup 3-3  
Non-Unicast 2-22  
Not Available 3-3  
Number 2-44, 3-9  
Number of MACs 3-5  
Number of Nodes 3-8

## O

OFF 2-72  
OK button 1-7  
ON 2-72  
Operational Status 4-12  
Outspeed 4-4

## P

- Packets Received 2-22
- Packets Transmitted 2-22
- Path Cost 2-38
  - Changing 2-41
- Performance 3-1
- Performance Graph window 2-5, 2-11
- Permanent entries 2-42
- Permanent Virtual Circuits (PVCs) 5-1
- Physical View 2-99, 2-101, 4-2
- Physical View window 2-13
- Port # 2-86
- Port Configuration window 2-6, 2-12, 2-65
- Port Filtering 2-45, 2-50
- Port Filtering Action
  - Changing 2-46, 2-52
  - Clearing 2-53
  - Setting 2-52
- Port Number 1-7
- Port Priority
  - Changing 2-41
- PPP 4-11
- PPP Link Statistics window 2-5
- PPP Link Status window 2-11
- Primary 1 3-5
- Primary 2 3-5
- Priority 2-37
- Protocol 2-36, 4-10

## R

- Receive Frames 3-10
- Receive OverRuns 4-15
- Receive Port 2-44, 2-47
  - Changing 2-46
- Receive Port Icon 2-52
- Receive Ports 2-50
  - Changing 2-52
- Refresh button 3-9
- Related Manuals 1-3
- Remote Capabilities 2-75
- Requested Target Token Rotation Time 3-4
- Restore button 4-9
- Ring Configuration 3-5
- Ring Ops 3-11
- Ring Speed 2-102
- Ring State 2-102
- Ring-Op 3-3
- Ring-Op-Dup 3-3
- RMON MAC Layer window 2-12

- RMON Promiscuous Stats window 2-12
- Root Bridge 2-36
- Root Bridge Selection process 2-34
- Root Cost 2-36
- Root Port 2-36
- Router Config 1-3

## S

- SAP 2-95
- SDH 2-78
- Secondary 1 3-5
- Secondary 2 3-5
- Selected Filter 2-50
- Set button 1-7
- Setting full duplex mode 2-72
- Severely Errored Framing Second 2-84
- Severely Errored Seconds 2-84
- Signal Mode 4-6
- Slave 3-5
- SmartTrunk Configuration and Status window 2-7, 2-85
- SmartTrunk State 2-86
- SmartTrunks 2-86
- SMT 3-2
- SMT Version 3-4
- SNAP 2-96
- SNAP Format
  - selecting 2-96
- SNAP format table 2-96
- SNAP Translation Table
  - configuring the table 2-96
- SONET 2-77
- Sonet Statistics window 2-13
- SONET/SDH 2-77
  - Coding Violations 2-84
  - configuration 2-77
  - Errored Second 2-84
  - errors 2-80
  - Errors indicators 2-82
  - Loss of Frame 2-83
  - Loss of Pointer 2-82
  - Loss of Signal 2-82
  - optical layers 2-80
  - Severely Errored Framing Second 2-84
  - Severely Errored Second 2-84
  - Statistics 2-79, 2-83
  - Statistics window 2-81
- SONET/SDH configuration 2-77
- Sonet/SDH Configuration window 2-13



SONET/SDH transmission hierarchy 2-78  
 Source 2-51  
 Source Address 2-50, 2-51  
 Source Address Table 2-6, 2-41  
 Source Addressing window 2-11  
 Source Port 2-44  
 Source Route Configuration window 2-6, 2-11  
 Source Route Statistics window 2-5, 2-11  
 Source Route Transparent mode 2-56, 2-90  
 Source Routing 2-2  
 Source Routing mode 2-56, 2-90  
 Spanning Tree Algorithm (STA) 2-2  
 Spanning Tree Algorithm Protocol Type  
   Changing 2-39  
 Spanning Tree Explorer (STE) packet 2-2  
 Spanning Tree window 2-6  
 Special Filter Database  
   Defining and Editing Filters 2-51  
   window 2-48  
   Window Fields 2-49  
 State 2-95, 2-96  
 State Change 4-4  
 Static Database 2-42  
 Static entries 2-42  
 Station 3-9  
 Station List 3-1  
 Station Management 3-2  
 Statistics window 2-13  
 Subnet Mask 2-60  
 Switched Virtual Circuits (SVCs) 5-1  
 Sync Port Configuration window 4-3, 4-4  
 Synchronous Port Statistics window 4-1  
   fields 4-14

## T

T1 AdminStatus 4-13  
 T1 connection 4-7  
 T1 interfaces 4-14  
 T-1 Port Configuration window 4-5  
   fields 4-5  
 T1 Port Configuration window 4-7  
 Technical support 1-9  
 T-Neg. 3-4  
 Token Ring Bridge Mode 2-89  
   window 2-7, 2-90  
 Token Ring Bridge Mode window 2-10  
 Token Ring Port Physical View 2-101  
 Token Ring Special Filter Database  
   window 2-6, 2-10

Topology 2-38, 3-9  
 Total Bridge Detail Breakdown window  
   Color-code 2-19  
 Trace 3-3  
 Transmit Frames 3-10  
 Transmit Queue Size 2-22  
 Transparent mode 2-55, 2-90  
 T-Req. 3-4  
 Twisted ring 3-6  
 Type 2-21, 2-44, 4-4  
 Type of Entry  
   Changing 2-46

## U

Unicast 2-22  
 Unknown 3-5  
 Unknown Protocol 2-22  
 Up Time 2-8  
 Upstream Neighbor 3-9  
 Uptime 1-7  
 Using the BRIM User's Guide 1-2

## V

VC MUX 802.3 Bridging 5-3, 5-4  
 Virtual Channel Identifier (VCI) 5-3  
 Virtual Path Identifier (VPI) 5-2

## W

WAN Interface Status window 4-1, 4-2, 4-5  
 WAN Interfaces 4-13  
 WAN Logical Settings  
   Changing 4-11  
 WAN Logical Settings window 4-11  
 WAN Logical View window 4-1, 4-10, 4-11  
 WAN Port Name 4-5  
 WAN Status 1-4, 4-1, 4-10, 4-12  
 WAN T1 Interfaces  
   Enabling and Disabling 4-14  
 Wrapped ring 3-6

## X

Xmit Clock Source 4-6  
 Xmit UnderRuns 4-15  
 Xmitted 2-17, 2-19

