# Configuring System Software
# Volume I

Software Version 7.50, Site Manager Version 1.50

# Reading Path

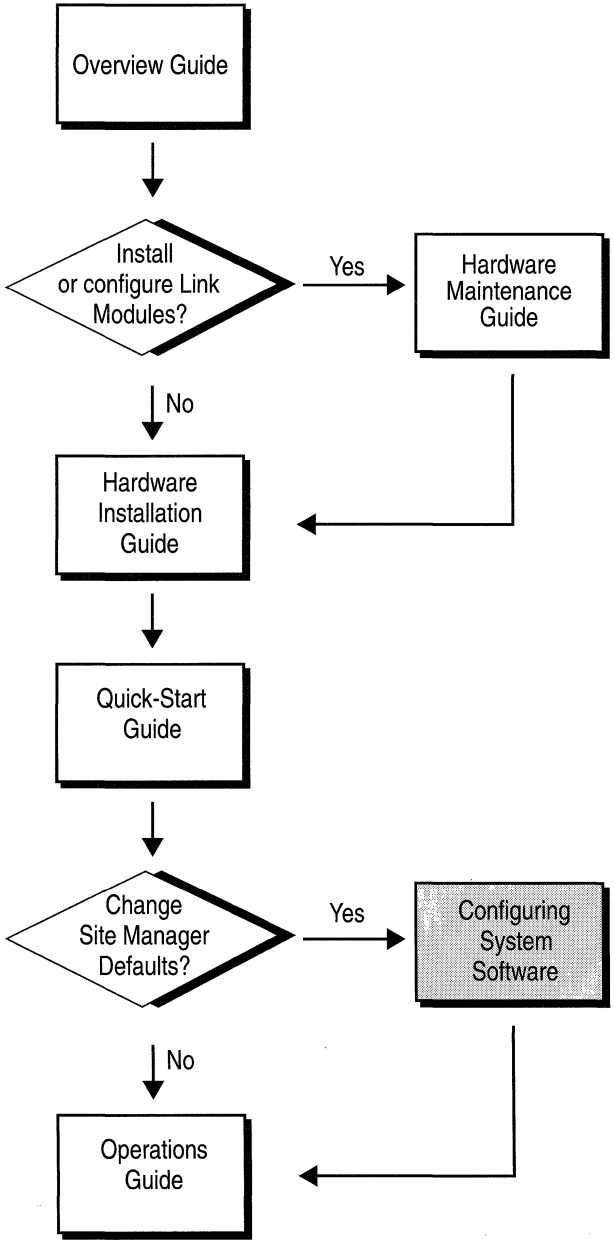

**Part Number: 105543, Revision A**



**Wellfleet Communications, Inc., 15 Crosby Drive, Bedford, MA 01730**

## Software License

This license governs the licensing of all Wellfleet software (Software) provided to licensee for use with Wellfleet equipment (Equipment). Licensee is provided with Software in machine-readable form and related documentation. The Software provided under this license is proprietary to Wellfleet and to third parties from whom Wellfleet has acquired license rights. Wellfleet does not grant any Software license whatsoever, either explicitly or implicitly, except by acceptance of an order for either a Software license or for a Wellfleet product that is packaged with Software. Each such license is subject to the following restrictions:

1. Licensee is granted a license to use the Software when payment for the license fee is made. Upon receipt of payment, licensee is granted a personal, nontransferable, nonexclusive license to use the Software with the specific item of Equipment with which or for which it was originally acquired, including use at any of licensee's facilities to which the Equipment may be transferred, for the useful life of the Equipment unless earlier terminated by default or cancellation. Use of the Software shall be limited to such specific item of Equipment and to such facility. Software which is licensed for use on hardware not offered by Wellfleet (e.g. Site Manager) is not subject to restricted use on any Equipment, however, unless otherwise specified in the Documentation, each licensed copy of such Software may only be installed on one item of hardware at any time.
2. Licensee may use the Software with the backup Equipment only if the Equipment with which or for which it was acquired in inoperative.
3. Licensee may make a single copy of the Software (but not firmware) for safekeeping (archives) or backup purposes.
4. Licensee may modify Software (but not firmware), or combine it with other software, subject to the provision that those portions of the resulting software which incorporate licensed Software are subject to the restrictions of this license. Licensee shall not make the resulting software available for use by any third party.
5. Wellfleet and third parties from whom Wellfleet has acquired license rights shall at all times retain title to and ownership of their respective portions of the Software including new versions, new releases, updates and modifications provided to licensee. Licensee agrees and acknowledges that licensee will obtain only such rights to a license or sublicense for the Software as are specifically provided herein.

## Software License (continued)

6. Licensee shall not provide, or otherwise make available, any Software, in whole or in part, in any form, to any third party. Third parties do not include consultants, subcontractors or agents of licensee who have licensee's permission to use the Software at licensee's facility, and who have agreed in writing to use the Software only in accordance with the restrictions of this license.

7. Third party owners from whom Wellfleet has acquired license rights to software that is incorporated into Wellfleet products shall have the right to enforce the provisions of this license against licensee.

8. Licensee shall not remove or obscure any copyright, patent, trademark, trade secret or similar intellectual property or restricted rights notice within or affixed to any Software and shall reproduce and affix such notice on any backup copy of Software or copies of software resulting from modification or combination performed by licensees as permitted by this license.

9. Notwithstanding any foregoing terms to the contrary, if Customer licenses the Product "Site Manager", Customer may duplicate and install the Site Manager Software as specified in the Documentation. This right is granted solely as necessary for use of the Site Manager Software on hardware installed within Customer's network. [Note: For licensees in the European Community, the Software Directive enacted by the Council of European Communities Directive dated 14 May 1991 shall apply for interoperability purposes. Licensee must notify Wellfleet in writing of any such intended examination of the Software and Wellfleet may provide review and assistance.]

10. Licensee shall not reverse assemble, reverse compile, or in any way reverse engineer the Software.

11. This license will automatically terminate upon improper handling of Software, such as by disclosure, or Wellfleet may terminate this license by written notice to licensee if licensee fails to comply with any of the material provisions of this license and fails to cure such failure within thirty (30) days after the receipt of written notice from Wellfleet. Upon termination of this license, licensee shall discontinue all use of the Software and return the Software and related documentation, including all copies, to Wellfleet.

12. Licensee's obligations under this license shall survive expiration or termination of this license.

## FCC Compliance Notice: Radio Frequency Notice

The following notice regarding compliance with Federal Communications Rules pertain to the Backbone Node.

This equipment generates, uses, and can radiate radio-frequency energy. If you do not install and use this equipment according to the instruction manual, this product may interfere with radio communications. This product has been tested and found to comply with the limits for a Class A computing device, pursuant to Subpart J of Part 15 of FCC Rules; compliance with these limits provides reasonable protection against radio interference when such equipment is operated in a commercial environment. Operating this equipment in a residential area is likely to interfere with radio communications; in which case, the user, as his/her own expense, must correct the interference.

Wellfleet shielded cables must be used with this unit to ensure compliance with the Class A limits.

## Canadian Department of Communications Radio Interference Regulations

This digital apparatus (the Backbone Node) does not exceed the Class A limits for radio-noise emissions from digital apparatus as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Le présent appareil numérique (le Feeder Node, le Link Node, et le Concentrator Node) n'émet pas de bruits radioélectriques dépassant les limites applicables au appareils numériques de Classe A prescrites dans Le Réglement sur Le Brouillage Radioélectrique Édité par Le Ministère des Communications du Canada.

## SITE MANAGER SOFTWARE

SITE MANAGER SOFTWARE IS AVAILABLE FOR INSTALLATION ON EITHER SUN SPARCSTATIONS OR DOS-BASED PERSONAL COMPUTERS (PCs). SITE MANAGER MAY BE INSTALLED ON AN UNLIMITED NUMBER OF CUSTOMER SUN SPARCSTATIONS. HOWEVER, SITE MANAGER FOR DOS PCs INCLUDES DISTINCT CORPORATION'S IP RUNTIME SOFTWARE WHICH CAN BE COPIED AND INSTALLED ON UP TO 15 PCs PER NETWORK IN CONJUNCTION WITH WELLFLEET SITE MANAGER FOR DOS PCs.

# Table of Contents

## Chapter 4

### Configuring Frame Relay



## Chapter 5

### Configuring SMDS



## Chapter 6

### Configuring AppleTalk

## Chapter 7

### Configuring the Bridge



## Chapter 8

### Configuring Source Routing



## Chapter 9

### Configuring DECnet Phase IV

## Chapter 10

### Configuring IP



## Chapter 11

### Configuring OSPF



## Chapter 12

### Configuring IPX

## Chapter 13

### Configuring SNMP



## Chapter 14

### Configuring VINES



## Chapter 15

### Configuring XNS

## Chapter 16

### Configuring Filters



## Chapter 17

### Configuring Protocol Prioritization



## Chapter 18

### Booting the BN with the Config File

## Appendix A

### Site Manager Default Settings



## Appendix B

### IEEE Assigned Codes

## Appendix C

### Converting Existing Traffic Filters

# About this Guide

## Audience and Scope

This guide is intended for experienced network managers who will be configuring the Backbone Node (BN). Such individuals should know the network topology in which the BN will operate.

Network managers should understand all protocols required in their networks.

This guide describes how to use the Site Manager's Configuration Manager application to set BN parameters in one of three modes: local, remote, or dynamic.

## Organization

The Configuring System Software Guide consists of two volumes. Volume I contains Chapters 1 through 11. Volume II contains Chapters 12 through 18 and Appendixes A through C. Both volumes contain inclusive Tables of Contents, and Indexes.

# How to Use this Guide

Refer to the following table for instructions on how to use this guide. The table also indicates in which volume the information can be found.

| For Instructions on: | Refer to: | Vol: |
|---|---|---|
| Using the Site Manager User Interface | *Site Manager User Interface* | I |
| Familiarizing yourself with the Configuration Manager's functionality and performing local, remote, or dynamic configuration | *Configuration Manager Overview* | I |
| Configuring data-link layer connections to networks and adding protocols to circuits | *Configuring Circuits* | I |
| Configuring Frame Relay | *Configuring Frame Relay* | I |
| Configuring SMDS | *Configuring SMDS* | I |
| Configuring AppleTalk | *Configuring AppleTalk* | I |
| Configuring the Bridge | *Configuring the Bridge* | I |
| Configuring Source Routing | *Configuring Source Routing* | I |
| Configuring DECnet Phase IV | *Configuring DECnet* | I |
| Configuring the IP Router | *Configuring the IP Router* | I |
| Configuring OSPF | *Configuring OSPF* | I |
| Configuring IPX | *Configuring IPX* | II |
| Configuring the BN to generate traps, and to restrict SNMP access to the BN | *Configuring the SNMP Agent* | II |

| For Instructions on: | Refer to: | Vol: |
|---|---|---|
| Configuring VINES | *Configuring VINES* | II |
| Configuring XNS | *Configuring XNS* | II |
| Configuring traffic filters | *Configuring Filters* | II |
| Configuring Protocol Prioritization | *Configuring Protocol Prioritization* | II |
| Implementing configuration changes in local and remote mode, and saving dynamically made changes | *Booting the BN with the Config File* | II |

**Note:** Appendix A contains Site Manager default settings. Appendix B contains IEEE Assigned Codes. Appendix C provides guidelines for creating Version 7 traffic filters identical to your Version 5 traffic filters.

# Document Set

The following guides complete this documentation set:

*Overview Guide*

Describes the user interface, called the Site Manager application, the system software, and the router hardware.

*Hardware Installation Guide*

Describes how to physically install the router hardware

*Quick-Start Guide*

Describes how to configure the router's initial IP network interface, install the Site Manager application software, and remotely create a pilot configuration for the Wellfleet router using the Site Manager.

*Hardware Maintenance Guide*

Describes how to access the interior of the Wellfleet router, replace the hardware, and how to read the LEDs.

*Operations Guide*

Describes how to use the Site Manager to perform day-to-day operations and how to use the Technician Interface to perform software maintenance.

If you are missing any guides, contact Wellfleet Customer Support at 1-800-2LANWAN.

# Conventions

This document set uses the following conventions:

| Convention: | Denotes: |
|---|---|
| *filename* | Italics denote file and directory names. |
| **command** | Bold text denotes text the user needs to enter. |
| Protocols/DECnet | The slash character (/) separates menu and option names in instructions; this example identifies the DECnet option in the Protocols menu. |

# Chapter 1

## Site Manager User Interface

## List of Figures

# Site Manager User Interface

## About this Chapter

This chapter describes how to enter and exit the Site Manager application; how to determine which Site Manager version you are running; and describes the Site Manager user interface, an intuitive, window-based graphical user interface (GUI) with point-and-click access.

## Entering and Exiting the Site Manager

To enter the Site Manager application, move to the directory in which you installed the application (see the *Release Notes* for installation instructions), type **wfsm** and press the carriage return. The Wellfleet Site Manager Window appears (see Figure 1-1).

To exit the Site Manager application, select the File/Exit option in the Wellfleet Site Manager Window. A pop-up window appears asking if you would like to terminate this session. Click on the Ok button to exit the application.

## Determining the Site Manager Version

To determine the Site Manager version you are running, select the File/Site Manager Version option in the Wellfleet Site Manager Window (see Figure 1-1). A pop-up window appears containing the version number. Click on the Cancel button to exit the window.

# Window-Based User Interface

The Site Manager user interface is a hierarchical window system. The top window in the hierarchy is the Wellfleet Site Manager Window (see Figure 1-1), which is the first window displayed when you enter the application.

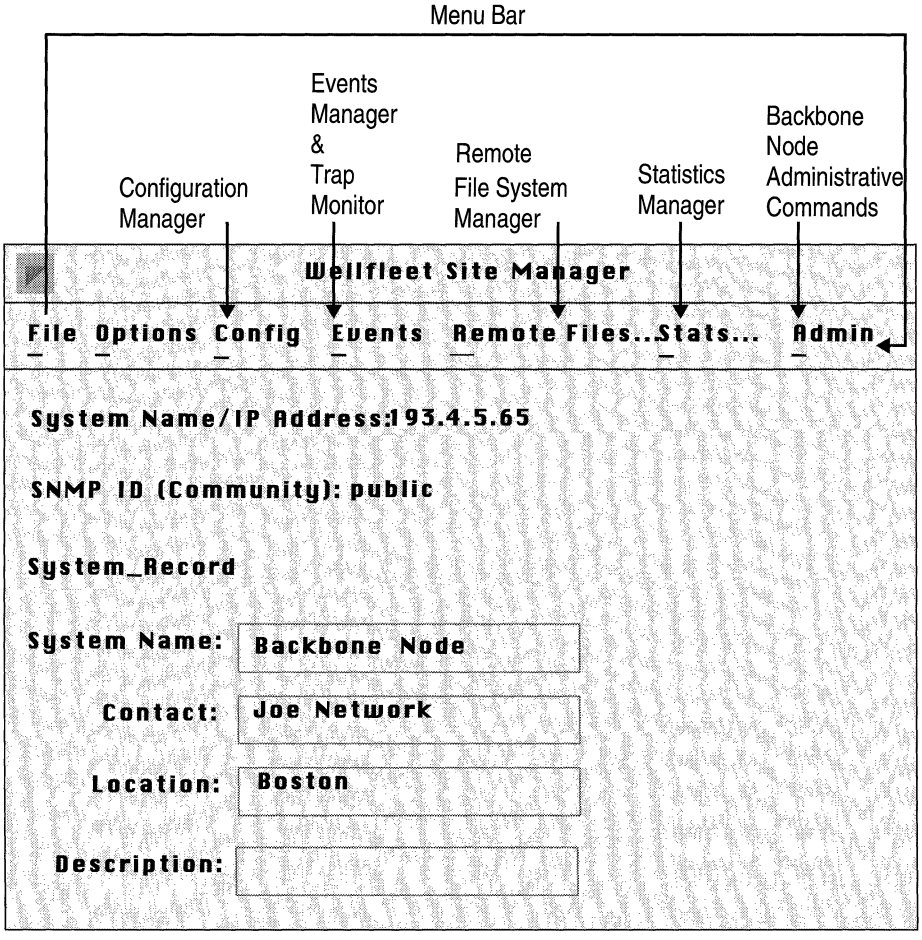


**Figure 1-1. Wellfleet Site Manager Window**

The menu bar in the Wellfleet Site Manager Window provides access to the five management applications (the Configuration Manager, the Events Manager, the Trap Monitor, the Remote File System Manager, and the Statistics Manager). You can run the five management applications simultaneously.

The Wellfleet Site Manager Window also allows you to execute BN administrative commands, by selecting Admin and the appropriate command option.

## Active Windows

An active window allows you to perform functions and/or enter information. Each management application allows you to have multiple active windows simultaneously. However, in some instances, windows displayed on your console will be inactive. If you try to work with an inactive window, the Site Manager highlights the active windows.

## Window Types

Every Site Manager window has either a menu bar (for example, the Wellfleet Site Manager Window in Figure 1-1) or function buttons (for example, the RIP Interface Parameters Window in Figure 1-2). Basically, windows with menu bars allow you to perform many functions; whereas, windows with function buttons allow you to enter and modify data associated with a single function only.

## Identifying Windows

Every Site Manager window has a title that can appear in three locations, as follows:

- Centered at the top of the window (for example, the Wellfleet Site Manager Window in Figure 1-1).
- Displayed below the function buttons (for example, the RIP Interface Parameters Window in Figure 1-2).
- Displayed above the function buttons (for example, the Delete IP Adjacent Host Window in Figure 1-3).

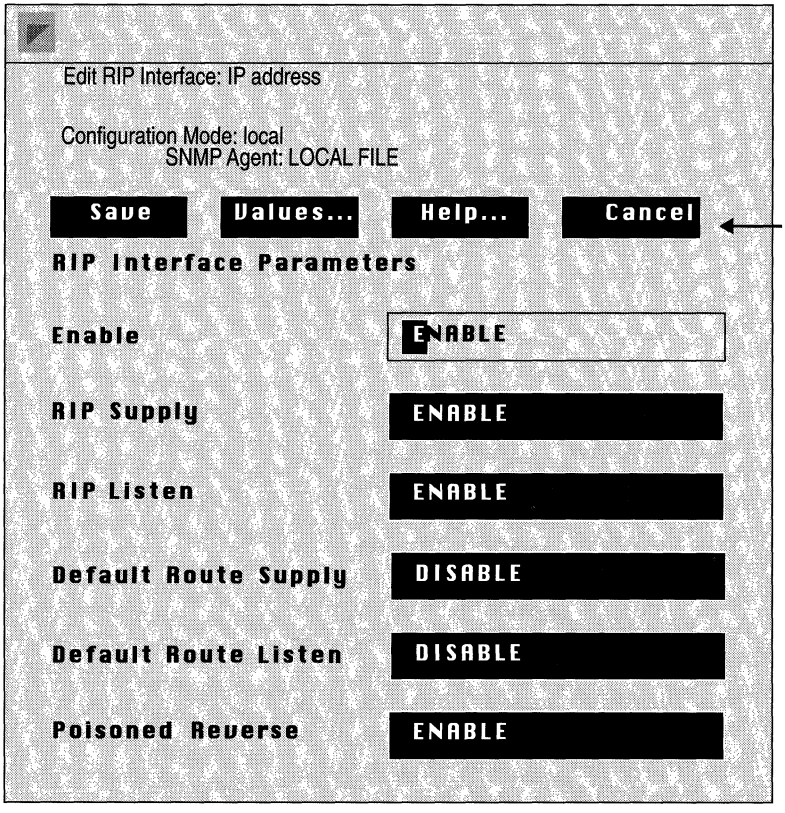


**Figure 1-2. RIP Interface Parameters Window**

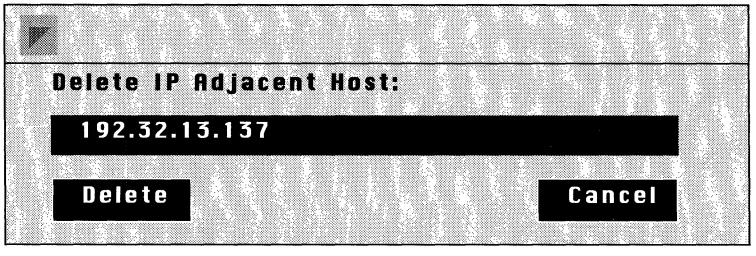


**Figure 1-3. Delete IP Adjacent Host Window**

# Working with Windows

The Site Manager is easy to use. The program's point-and-click access reduces the likelihood of typographical errors, allowing you to configure most parameters and to execute functions simply by pressing a mouse button. The application also allows you to enter data directly from the keyboard.

The following sections describe how to work with Site Manager windows. The first section describes window conventions that you should familiarize yourself with before using the Site Manager. The remaining sections describe how to enter, display, delete, and save data; how to display on-line help; and how to exit windows.

## Window Conventions

Site Manager windows use the following conventions:

- Menu labels, menu options, and button labels that end with three dots (...) generate a new window when selected.

  For example, the Site Manager generates a new window when you select the menu label Remote Files in the Wellfleet Site Manager Window (see Figure 1-1), when you click on the Help button in the RIP Interface Parameters Window (see Figure 1-2), and when you select the Protocols/IP/Global option in the Wellfleet Configuration Manager Window (see Figure 1-4).

- Menu labels with no dots and menu options followed by a shaded arrow ( ▶ ) display a menu or submenu when selected.

  For example, the Protocols menu and the IP submenu in the Wellfleet Configuration Window (see Figure 1-4).

- Button labels with no dots (...) perform a function when you click on them.

  For example, when you click on the Save button in the RIP Interface Parameters Window (see Figure 1-2), the Configuration Manager saves the displayed data and exits the window.

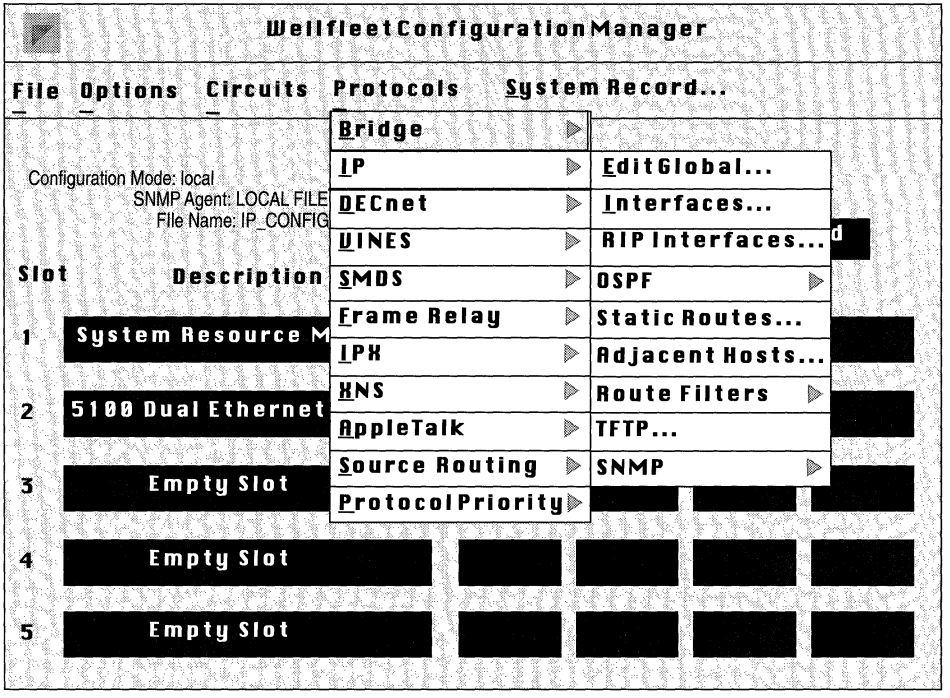


**Figure 1-4. Wellfleet Configuration Manager Window**

- Underlines in menu labels and menu options identify keyboard short cuts.

  For example, to use the keyboard to select the Protocols/IP option in the Wellfleet Configuration Manager Window (see Figure 1-4), you would type the letters **P** and **I** (the Site Manager is not case sensitive).

- PF# in menu labels identify function key shortcuts.

  For example, to select the Trap Monitor menu option in the Events menu displayed in the Wellfleet Site Manager Window (see Figure 1-5), you would press the F6 function key.

- The Site Manager displays in grey menu options that are not active for a particular window.

  For example, in Figure 1-4, the Site Manager displays the Bridge menu options in grey, because that configuration does not support the Bridge protocol. If you were to add Bridge support to that configuration, the Site Manager would display the Bridge menu option in full color and allow you to select it.

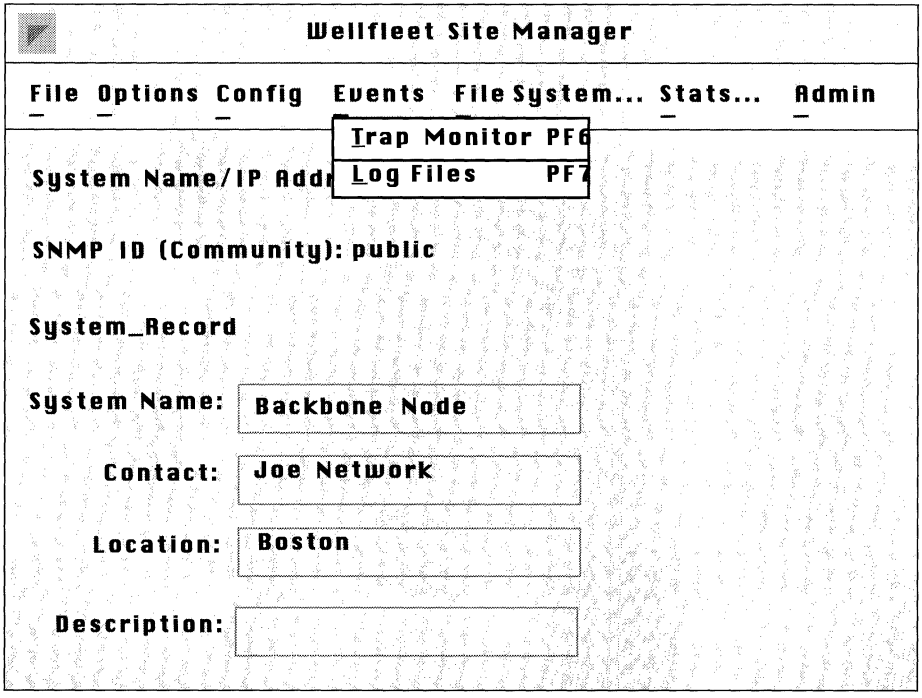


**Figure 1-5. Wellfleet Site Manager Window**

## Entering Data

The Site Manager allows you to use the mouse and the keyboard to enter data. The following sections describe both procedures.

### Using the Mouse

In windows containing a Values button, the Configuration Manager allows you to position the cursor in a particular data field and then click on the Values button to display a pop-up window with the proper range of settings for that data field. Figure 1-6 shows the Values pop-up window for the Default Volume parameter in the TFTP Parameters Window. To use a Values pop-up window, simply select or enter the desired setting in the pop-up window and click on the Save button. The Site Manager automatically enters your selection in the data field.

### Using the Keyboard

To use the keyboard to enter data in a field, do one of the following:

- Use the mouse to select the existing data entry in the field and then use the keyboard to enter your data; the data you enter will replace the selected entry.

  Double clicking on a word selects the word. Triple clicking on a word selects the entire data entry in the field.

- Position the cursor in the field, toggle the insert key to the proper insert mode, and then use the keyboard to enter data.

  The insert key is one of the number keys on the right side of your keyboard (see Figure 1-7). The insert key toggles the insert mode between insert and overwrite. When you enter data in insert mode, the cursor appears as a block (▯), and your entry is added to the existing entry. When you enter data in overwrite mode, the cursor appears as an I (I), and your entry overwrites the existing entry.

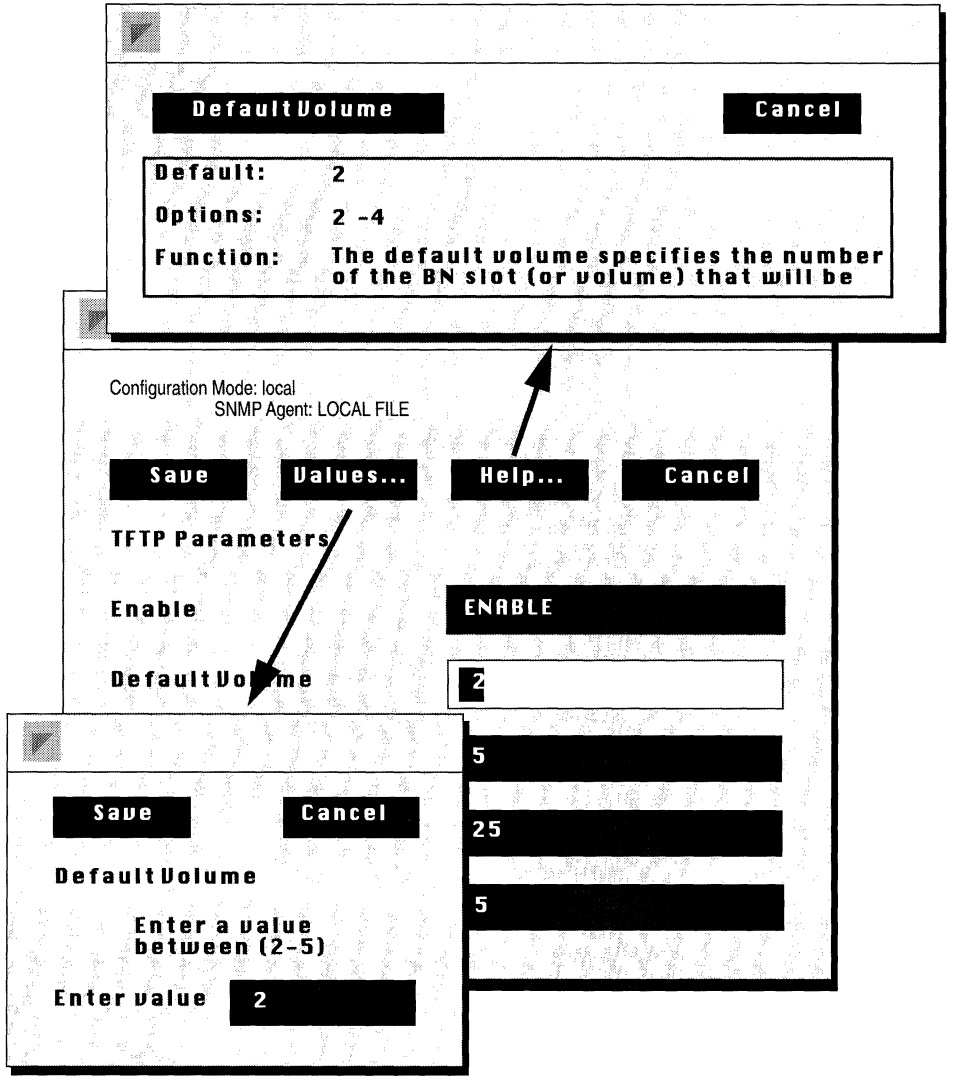


**Figure 1-6. Values and Help Pop-Up Windows for Default Volume Parameter**

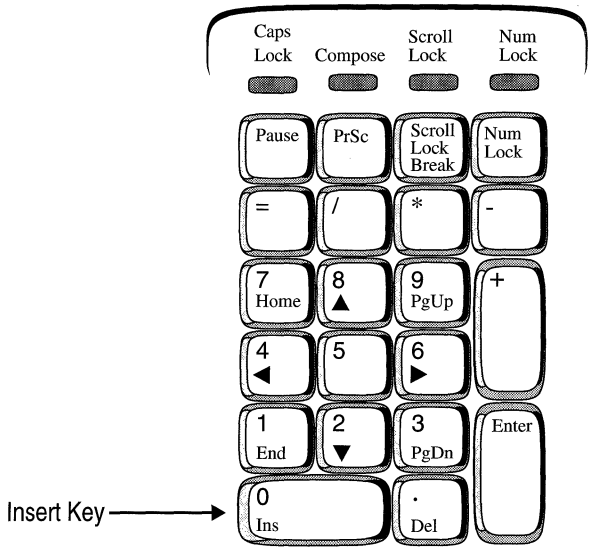


**Figure 1-7. Insert Key**

## Displaying Hidden Data

Some windows may contain hidden data fields. These windows have a jagged line underneath the last visible data field signalling that there are more fields. To see the additional field, click into the last visible data field and use the down arrow key on your number key pad. Or, simply click on the jagged line. If no additional field exists, the Site Manager does nothing.

Some data fields may contain hidden data. To see the hidden data, click into the data field and use the right arrow key on your number key pad.

Windows that display lists of data provide scroll bars that indicate whether hidden data exists and that allow you to access any hidden data. For example, the Module List Window in Figure 1-8 indicates that the window contains hidden data, because the scroll bar does not stretch the entire length of the data display. To reveal the hidden data, use the mouse to drag down the scroll bar.

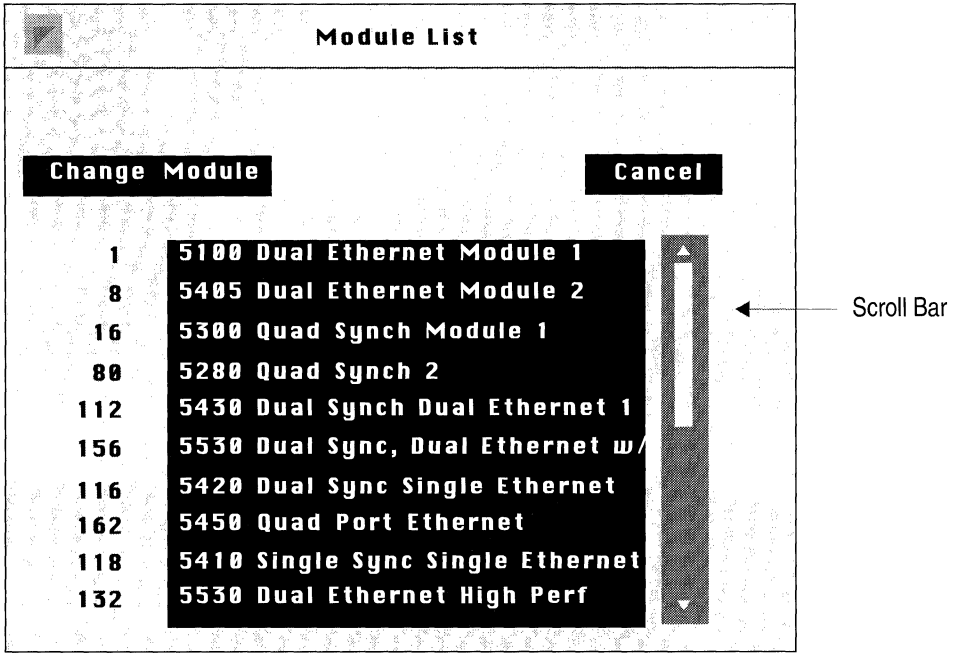


**Figure 1-8. Module List Window**

## Displaying Help Windows

In windows containing a Help button, the Configuration Manager allows you to display pop-up Help windows. Simply place the cursor in the field for which you would like a Help Window displayed and then click on the Help button. Figure 1-6 shows the Help window for the Default Volume parameter in the TFTP Parameters window.

## Exiting Windows

To exit a window containing a menu bar, select the File/Exit or File/Cancel option. To exit a window containing function buttons, click on the Cancel button.

# Chapter 2

## Configuration Manager Overview

## List of Figures

# Configuration Manager Overview

## About this Chapter

This chapter provides:

- An overview of the Configuration Manager's functions
- A description of each of the Configuration Manager's three operating modes (local, remote, and dynamic), and an explanation of how the Configuration Manager operates in each mode
- A description of how to use the Configuration Manager in each operating mode to configure the BN, including specific configuration steps
- A description of how to use the Configuration Manager to configure the BN's connection to the Technician Interface (TI)
- A description of how to use the Configuration Manager to specify BN administrative information

You should read this chapter before you use the Configuration Manager to configure a BN.

# Configuration Functions

The Configuration Manager allows you to perform the following BN configuration functions:

- Configure circuits.

  Circuits are the data-link layer connections between the BN and attached network(s).

- Configure routing/bridging protocols.

  Routing/bridging protocols are added to circuits to form network interfaces.

- Configure filters.

  Filters are used mainly for security reasons; they enable the BN to selectively relay or drop incoming packets.

- Assign priority to certain types of traffic.

  Protocol prioritization allows you to assign priority to certain types of traffic that you choose.

- Reconfigure the Technician Interface (TI) console parameters.

  A command-line interface, the TI is a back up to the Site Manager. The console port on the System Resources Link Module (SRM-L) connects the BN to the TI.

In addition, the Configuration Manager allows you to specify a BN's hardware configuration and certain administrative information about the BN. The following sections provide an overview of each function.

## Configuring Circuits

The Configuration Manager enables you to configure circuits quickly—default parameter settings minimize the parameters you must specify. You simply select the connector for which you wish to configure a circuit from a graphic display in the Add Circuit Window and then click on the Save button. The Configuration Manager provides a default circuit name and automatically configures the requisite physical-layer (or line) parameters appropriate for the selected connector.

After configuring the circuit, the Configuration Manager automatically prompts you to select the network protocols you wish to run over the circuit so that you can configure a network interface. You specify only a few parameters for each added protocol. The application provides default settings for the remaining parameters needed to configure the interface; however, the application allows you to edit certain interface-specific parameters at this point.

After you have added a circuit (and optionally added network protocols), the Configuration Manager allows you to access and edit all parameters associated with the circuit. Specifically, you can:

- Delete or rename the circuit.
- Edit the physical-layer, or line, parameters (Ethernet, FDDI, synchronous, E1, T1, Token Ring, or HSSI) associated with the circuit.
- Add network protocols to the circuit to configure a network interface.
- Edit interface-specific parameters associated with supported network protocols.
- Delete network protocols from the circuit.
- Move a circuit from one network interface to another.
- Add multiple IP addresses to a single circuit that supports IP.

Any changes you make affect only the circuit or network interface (if you added network protocols) you are reconfiguring. Each time you add a new circuit, *the Site Manager again uses its own default settings*, plus the few parameters that you specify, to configure the new circuit.

## Configuring Routing/Bridging Protocols

The Configuration Manager allows you to access and reconfigure all parameters associated with each protocol that you have added to a circuit. You can access these parameters on a system-wide or interface-specific basis.

## Configuring Traffic Filters

The Configuration Manager allows you to configure traffic filters that instruct the router to selectively relay or drop a packet, frame, or datagram based on standard protocol fields or user-defined fields. You can configure traffic filters for the following protocols: the Bridge, IP, DECnet Phase IV, VINES, Source Routing, IPX, and XNS.

## Configuring Protocol Prioritization

The Configuration Manager allows you to assign priority to different types of traffic transmitted on an individual synchronous line interface. Priority assignment is based on fields with in the datalink or IP headers, or by length.

## Configuring the BN Connection to the Technician Interface (TI)

The console port on the System Resources Link Module connects the BN to the Technician Interface (TI). The Site Manager provides default settings for the console configuration (Appendix A lists the default settings); however, you can use the Configuration Manager to access and reconfigure all console parameters.

## Specifying BN Hardware Configuration

During dynamic and remote mode (operating modes are described later in this chapter), the Configuration Manager uses SNMP to retrieve and display the BN's hardware configuration in the hardware display in the Wellfleet Configuration Manager Window (see Figure 2-1). Thus, you do not specify the hardware configuration in dynamic or remote mode.

However, in local mode, the Configuration Manager requires you to specify the hardware configuration whenever you create a new configuration file, and when you edit a local configuration file for which there is no local hardware configuration file. When you save a local configuration, the Configuration Manager automatically creates a local hardware configuration file that is associated with that BN configuration file.

### Specifying Administrative Information

The Configuration Manager allows you to specify a system name, a system contact, and a system location for the BN. Each time you use the Site Manager to access the BN, the Site Manager retrieves the administrative information and displays it in the Wellfleet Site Manager Window (see Figure 2-2), which is the window displayed at Site Manager start-up.

## Operating Modes

The Configuration Manager allows you to perform all configuration functions in one of three modes: local, remote, or dynamic. You specify the operating mode from the Wellfleet Site Manager Window (see Figure 2-2) by selecting Config and the appropriate operating mode option. When you operate the Configuration Manager in remote or dynamic mode, you configure SNMP options that allow the Configuration Manager to access the BN.

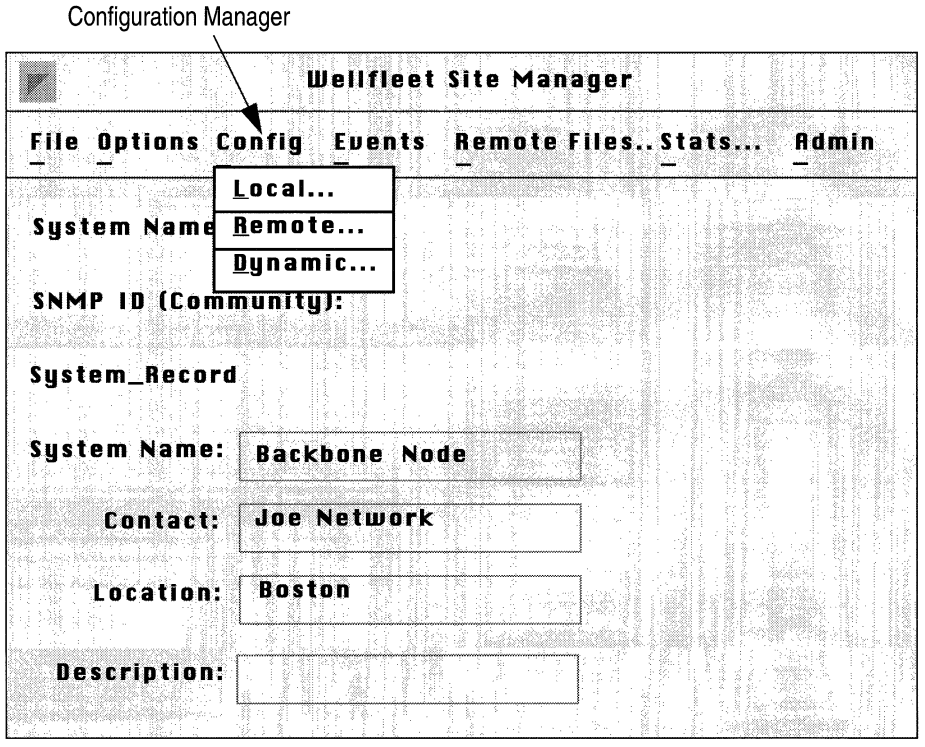


**Figure 2-1. Wellfleet Site Manager Window**

For each function you can perform, the Configuration Manager displays the same windows in the same sequence regardless of the operating mode. For example, the Configuration Manager displays the same sequence of windows when you configure the Routing Information Protocol (RIP) protocol in local mode, as it displays when you configure RIP in dynamic mode.

The Configuration Mode field in the upper-left corner of each window identifies the Configuration Manager's current operating mode. For

example, Figure 2-2 shows the RIP Interface Parameters Window when the application is running in local mode.

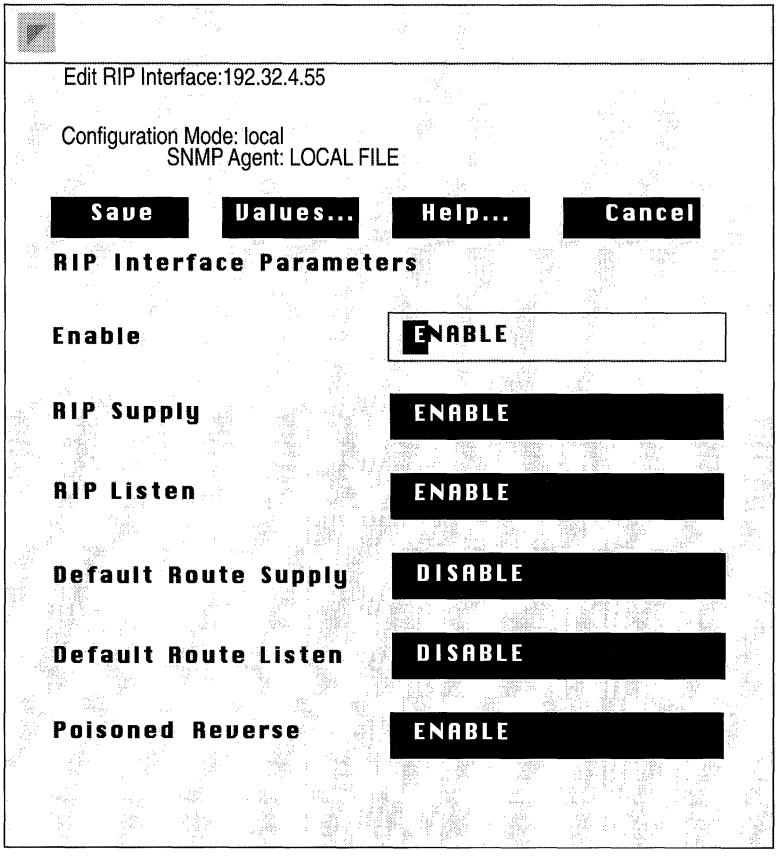


**Figure 2-2. RIP Interface Parameters Window**

The following sections describe each operating mode, and how to run the Configuration Manager in each of the three modes; refer to Figure 2-3 as necessary.

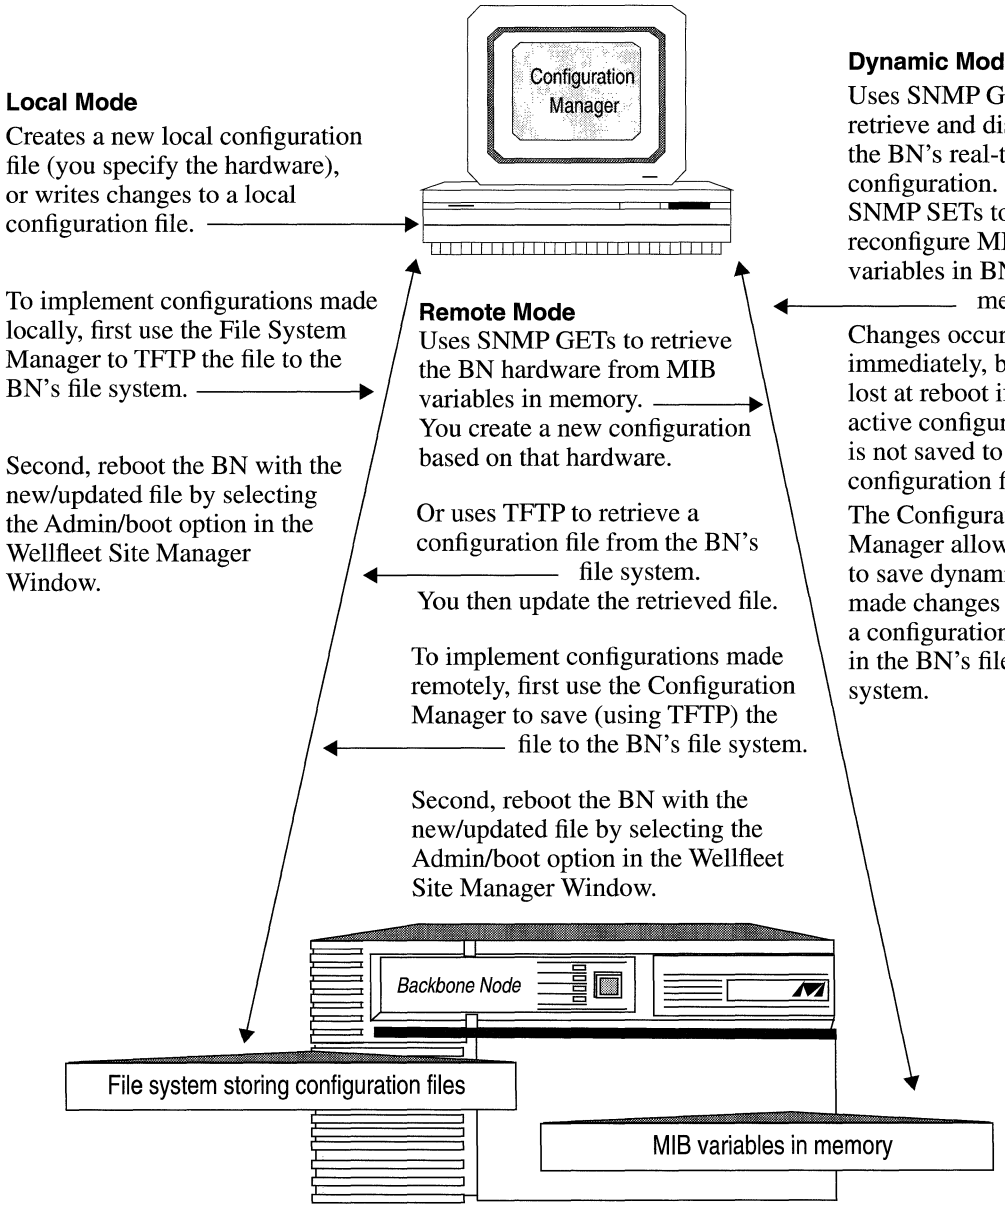


**Figure 2-3. Local, Remote, and Dynamic Operating Modes**

## Local Mode

When you run the Configuration Manager in local mode, you either create or edit a local configuration file on the Site Manager workstation. Unlike remote and dynamic modes, the Configuration Manager does not access a BN; therefore, you may have to specify the BN's hardware configuration (the Configuration Manager automatically retrieves and displays a local hardware configuration file if one exists).

In local mode, you must specify hardware when you:

- *Create* a new configuration file.

  However, when you save a configuration file in local mode, the Configuration Manager automatically creates a hardware configuration file (named *config_file_name.hdw*) that it associates with that configuration file. The application saves both the configuration and hardware file locally on the Site Manager workstation.

- Edit an existing local file for which there is *no local hardware configuration file*.

  When you open a local configuration file, the Configuration Manager automatically opens the associated hardware configuration file. However, in some instances there may be no hardware configuration file. For example, if you use the File System Manager to TFTP a configuration file to the Site Manager workstation, when you use the Configuration Manager to open that file for local editing, there will be no local hardware file associated with it. In such a situation, you must specify the hardware; once you have done so, the Configuration Manager will automatically display the configured circuits in the Wellfleet Configuration Manager Window.

Once you have specified a BN's hardware configuration in local mode, you can go on to configure circuits, routing/bridging protocols, filters and priority assignments in the same manner as if you were running the application in remote or dynamic mode. When you have configured all required parameters, you save the file to the Site Manager workstation. To implement your configuration, you must use the File

System Manager to TFTP the configuration file to the BN, and then use the Site Manager to reboot the BN with the file. *Performing Local Configuration* lists the configuration steps you must follow when operating the Configuration Manager in local mode.

## Remote Mode

You use remote mode if you can access the BN over the network but want to implement the configuration at a later date. In order to run the Configuration Manager in remote mode, you must first configure SNMP options that identify the BN you wish to configure and that provide the Site Manager with an SNMP community that has access to the BN. *Enabling SNMP Access to the Backbone Node* provides an overview of SNMP options.

After you have specified SNMP options and selected the remote operating mode, the Configuration Manager uses SNMP GETs to retrieve and display the BN's hardware configuration in the Wellfleet Configuration Manager Window (see Figure 2-5). Optionally, if you specify a configuration file stored in the BN's file system, the Configuration Manager uses TFTP to automatically retrieve that file for local editing. In Figure 2-5, the BN administrator has retrieved the configuration file *config.Q1*, which contains configuration information for the transceivers in slot 2. If you do not specify a BN configuration file, you must create a new configuration file based on the retrieved hardware.

When you save a new or updated configuration file in remote mode, the Configuration Manager automatically TFTPs the file to the BN. You must then use the Site Manager to reboot the BN with that configuration file in order to implement the configuration. *Performing Remote Configuration* lists the configuration steps you must follow when operating the Configuration Manager in remote mode.

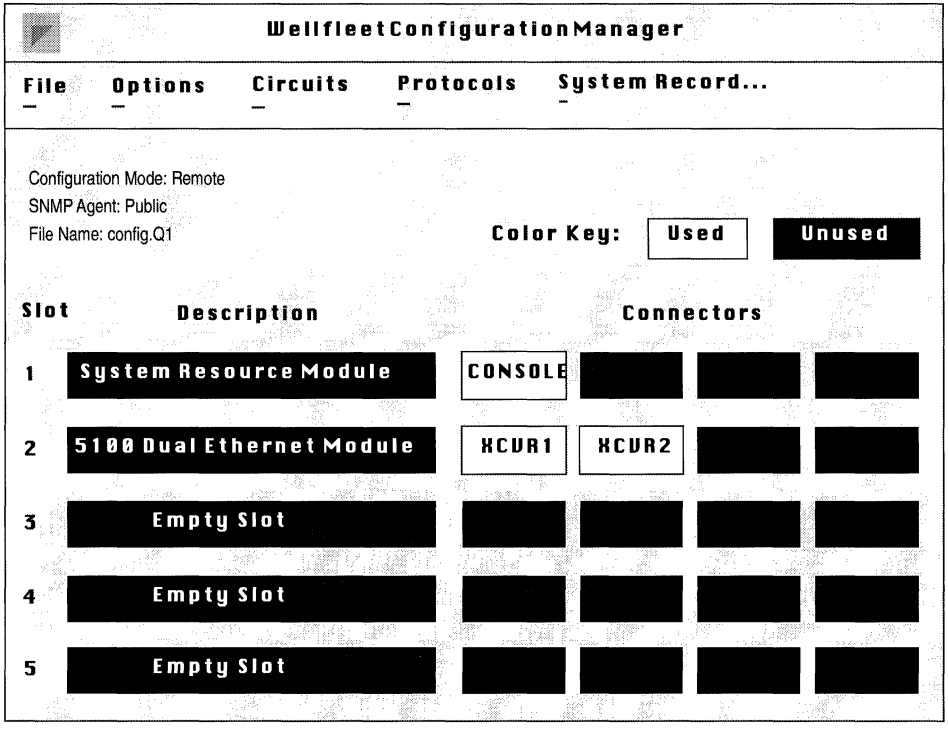


**Figure 2-4. BN Hardware Configuration in Wellfleet Configuration Manager Window**

## Dynamic Mode

You use dynamic mode if you can access the BN over the network and want to configure the system in real-time. In order to run the Configuration Manager in dynamic mode, you must first configure SNMP options that identify the BN you wish to configure and that provide the Site Manager with an SNMP community that has read/write access to the BN. *Enabling SNMP Access to the Backbone Node* provides an overview of SNMP options.

When running in dynamic mode, the Configuration Manager uses SNMP GETs to retrieve and display the BN's real-time hardware and software configuration, and uses SNMP SETs to write changes directly to the BN's memory — dynamically reconfiguring the BN.

For example, when you dynamically reconfigure a BN's IP Global parameters, the Configuration Manager uses SNMP GETs to retrieve the active configuration in the BN's memory and displays the real-time parameter settings in the IP Global Parameters Window (see Figure 2-6). When you reconfigure a parameter in the window and click on the Save button, the Configuration Manager uses an SNMP SET to write the new value directly to BN memory and dynamically reconfigures the node.

**Note:** While the reconfigured interface may incur a temporary disruption in service, the rest of the BN system continues to operate normally. Also, depending upon the configuration procedure you perform, there may be network ramifications.

The dynamically made change is lost at reboot if you do not use the Configuration Manager to save the active configuration to a configuration file.Therefore, when operating in dynamic mode, the Configuration Manager allows you to save a BN's active configuration to a named configuration file in the BN's file system. *Performing Dynamic Configuration* lists the configuration steps you must follow when operating the Configuration Manager in dynamic mode.

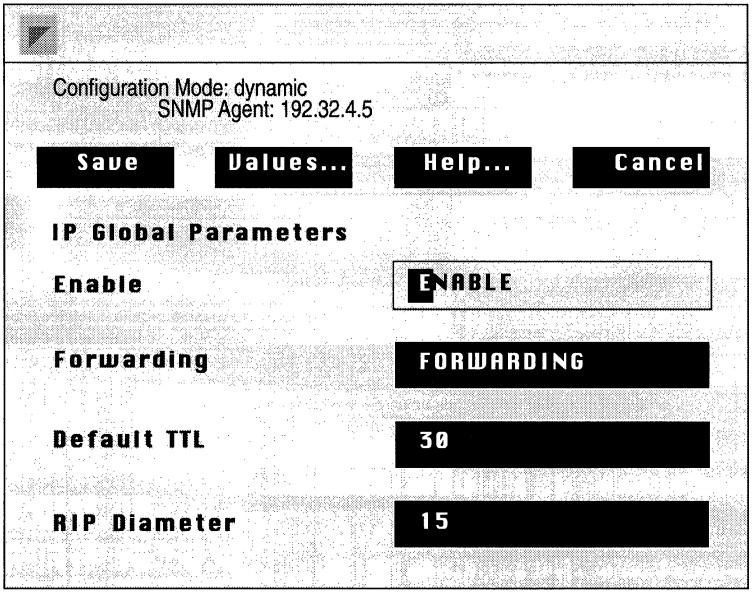


**Figure 2-5. IP Global Parameters Window in Dynamic Mode**

### Enabling SNMP Access to the Backbone Node

When you operate the Configuration Manager in dynamic or remote mode, you must configure SNMP options that identify:

- The BN you wish to configure.
- An SNMP community that enables the Configuration Manager to access the BN.

You can also specify the number of seconds the Configuration Manager waits for a response from the BN after it issues an SNMP SET or GET before the Configuration Manager reissues the command, as well as the number of times it will reissue the command.

You configure SNMP options in the SNMP Options Window (see Figure 2-6). You access this window by selecting the Options/SNMP option in the Wellfleet Site Manager Window.

**Note:** Once you have entered the Configuration Manager application, you can reconfigure SNMP options by selecting the Options/ SNMP option in the Wellfleet Configuration Manager Window; however, you may not specify another BN once you have entered the Configuration Manager application.

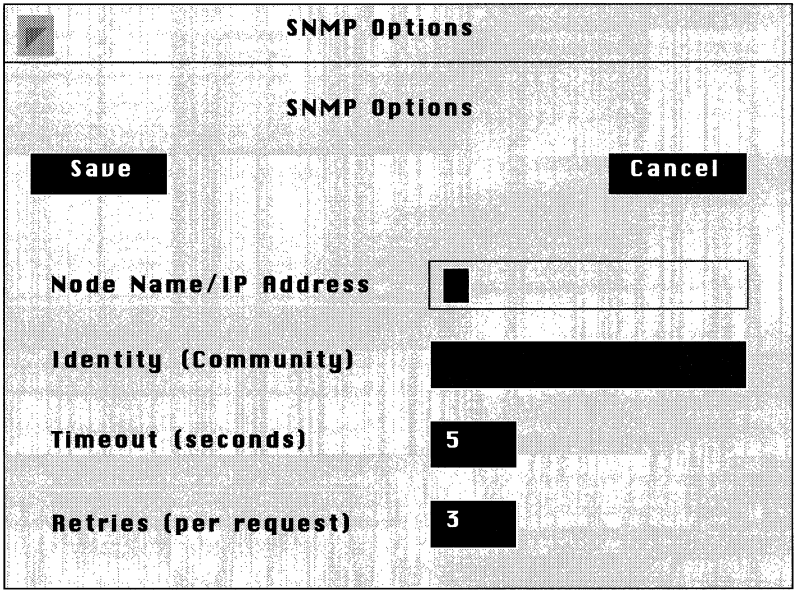


**Figure 2-6. SNMP Options Window**

This section provides information you need to edit each parameter in the SNMP Option Window. Refer to this information to edit the parameters you wish to change. For each parameter, it provides the following:

- Wellfleet default
- Range of valid settings
- Parameter's function
- Instructions for setting the parameter

When you are done editing parameters, click on the Save button to exit the window and save your changes.

| **Parameter :** | **Node Name/IP Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | Valid host name or valid IP address |
| Function: | Specifies the host name or IP address of the BN you wish to configure. |
| Instructions: | Enter the BN's IP address or host name. You may only specify a host name if you have defined a host name for the BN in the host file of your Site Manager workstation. |

**Note:** If you displayed the SNMP Options Window from the Wellfleet Configuration Manager Window, the SNMP Options window will not display the Node Name/IP Address parameter. Each instance of the Configuration Manager application allows you to access one BN only. To configure two BNs simultaneously, you must display the SNMP Options Window from the Wellfleet Site Manager Window, specify the new BN, and then run another instance of the Configuration Manager application, which automatically communicates with that BN.

| | |
|---|---|
| **Parameter:** | **Identity (Community)** |
| Wellfleet Default: | public |
| Options: | Any valid SNMP community name. |
| Function: | Specifies the name of the SNMP community that you wish the Configuration Manager to use to access the BN. |
| Instructions: | Enter the SNMP community name. The community must have read/write access to the specified BN, if you wish to use the Configuration Manager to reconfigure the BN. Community names are case sensitive and consist of up to 31 characters. |

| | |
|---|---|
| **Parameter:** | **Timeout (seconds)** |
| Wellfleet Default: | 5 seconds |
| Options: | 1 to 300 seconds |
| Function: | Specifies the number of seconds the Configuration Manager waits for a response from the BN, after it issues an SNMP SET or GET before reissuing the command. |
| Instructions: | Enter the number of seconds. |

| | |
|---|---|
| **Parameter:** | **Retries (per request)** |
| Wellfleet Default: | 3 |
| Options: | 0 to 30 |
| Function: | Specifies the number of times the Configuration Manager will reissue a command when the BN does not respond. |
| Instructions: | Enter the number of times. |

# Configuration Steps for each Operating Mode

The following sections list the configuration steps required when configuring the BN in each of the Configuration Manager's operating modes. Each step points to the appropriate section or chapter in this guide that describes the configuration procedure.

## Performing Local Configuration

Local configuration consists of the following steps:

1. Specifying the local operating mode and a configuration file (see *Specifying the Local Operating Mode* in this chapter.
2. Specifying hardware (this step may or may not be required). See *Specifying Hardware* in this chapter.
3. Configuring circuits (see the *Configuring Circuits* chapter).
4. Configuring routing/bridging protocols. Based on the protocol(s) you configure, see the corresponding configuration chapter.
5. Configuring traffic filters (see the *Configuring Filters* chapter).
6. Configuring priority filters (see the *Configuring Protocol Prioritization* chapter).
7. Saving the configuration file to the Site Manager workstation (see the *Booting the BN with the Config File* chapter).
8. Using the File System Manager to TFTP the configuration file to the BN (see the *Booting the BN with the Config File* chapter).
9. Using the Site Manager to reboot the BN with the configuration file (see the *Booting the BN with the Config File* chapter).

## Performing Remote Configuration

Remote configuration consists of the following steps:

1. Specifying the BN you wish to configure (see *Configuring SNMP Options* in this chapter.
2. Specifying the remote operating mode and a configuration file (see *Specifying the Remote Operating Mode* in this chapter).
3. Configuring circuits (see the *Configuring Circuits* chapter).
4. Configuring routing/bridging protocols.Based on the protocol(s) you configure, see the corresponding configuration chapter.
5. Configuring the SNMP agent (see the *Configuring SNMP* chapter).
6. Configuring filters (see the *Configuring Filters* chapter).
7. Configuring priority filters (see the *Configuring Protocol Prioritization* chapter).
8. Saving the configuration file to the BN's file system (see the *Booting the BN with the Config File* chapter).
9. Using the Site Manager to reboot the BN with the configuration file (see the *Booting the BN with the Config File* chapter).

## Performing Dynamic Configuration

Dynamic configuration consists of the following steps:

1. Specifying the BN you wish to configure (see *Configuring SNMP options* in this chapter).
2. Specifying the dynamic operating mode by selecting the Config/dynamic option in the Wellfleet Site Manager Window.

   The Wellfleet Configuration Manager Window appears displaying the real-time BN hardware and software configuration.
3. Configuring circuits (see the *Configuring Circuits* chapter).
4. Configuring routing/bridging protocols. Based on the protocol(s) you configure, see the corresponding configuration chapter.
5. Configuring the SNMP agent (see the *Configuring SNMP* chapter).
6. Configuring filters (see the *Configuring Filters* chapter).
7. Configuring priority filters (see the *Configuring Protocol Prioritization* chapter).
8. Saving the configuration file to the BN's file system if you wish to maintain a record of the changes you made (see the *Booting the BN with the Config File* chapter).

### Specifying the Local Operating Mode

You specify the local operating mode from the Wellfleet Site Manager Window, as follows:

1. Select the Config/local option to display the Edit Local Configuration Window (see Figure 2-7).
2. Enter the configuration file name, as follows:
   - If you wish to edit a local file, enter the configuration file name and click on the Open button.

     By default the Configuration Manager retrieves the file from the directory in which you are running the Site Manager application. If the file resides in a different directory, you must specify it.

     The Configuration Manager retrieves the specified filed and displays its contents in the Wellfleet Configuration Manager Window. If there is no local hardware configuration file associated with that configuration file, the Wellfleet Configuration Manager Window will remain empty until after you specify the hardware, at which point, the application will automatically display configured circuits.
   - If you wish to create a new configuration file, enter a file name and click on the Open button.

     The Configuration Manager displays the Create a New File Message Pop-Up Window. Click on the Ok button to display the Wellfleet Configuration Manager Window.

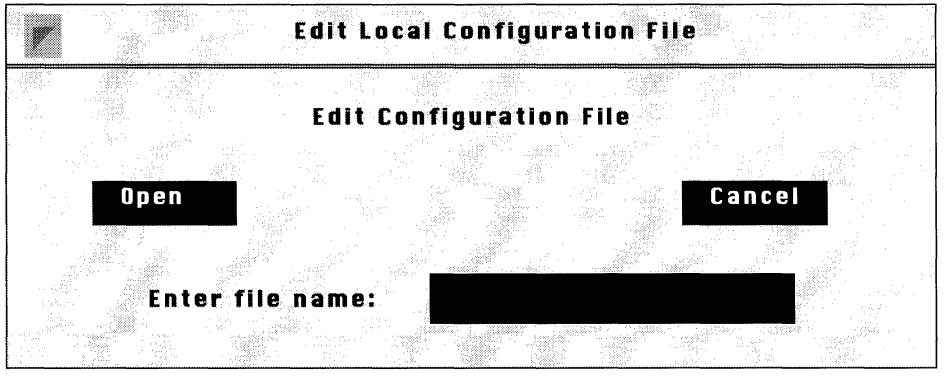


**Figure 2-7. Edit Local Configuration FIle Window**

### Specifying the Remote Operating Mode

You specify the remote operating mode from the Wellfleet Site Manager Window, as follows:

1. Select the Config/remote option in the Wellfleet Site Manager Window to display the Edit Remote Configuration Window (see Figure 2-8).
2. Enter the configuration file name, as follows:
   - If you wish to retrieve a file from the BN for local editing, select the number of the volume (or slot) containing the configuration file (by default, the Configuration Manager displays volume 2). Simply click on the Volume box to display the available volumes. The directory of the chosen volume will appear in the directory box. Then enter configuration file name and click on the Open button.

     The Configuration Manager uses TFTP to retrieve and display the specified filed in the Wellfleet Configuration Manager Window; the file name you specified is displayed in the upper-left corner of the window. If the specified file does not exist, the Wellfleet Configuration Manager Window displays the hardware configuration and no circuits.

— If you wish to create a new configuration file based on the current hardware configuration, select the number of the volume (or slot) to which you will later save the configuration file (by default, the Configuration Manager displays volume 2). Simply click on the Volume box to display the available volumes. The directory of the chosen volume will appear in the directory box. Then enter the file name and click on the Open button.

The Configuration Manager displays the Wellfleet Configuration Manager Window with the BN's current hardware configuration. The file name you specified is displayed in the upper-left corner of the window.

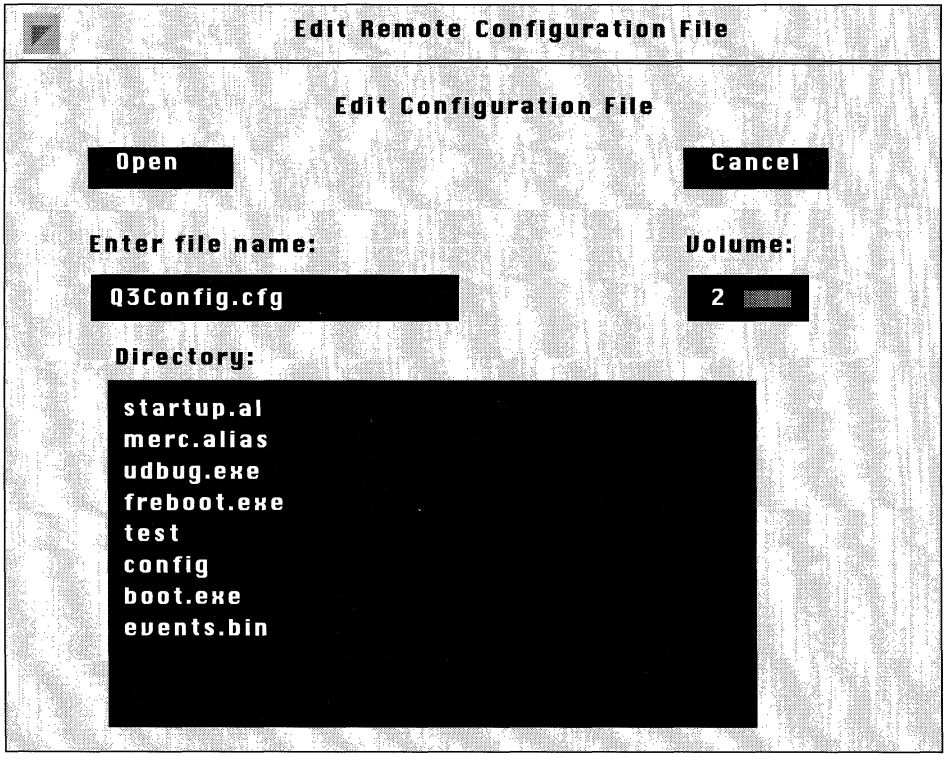


**Figure 2-8. Edit Remote Configuration File Window**

### Specifying Hardware

You can only specify hardware in local mode. The Configuration Manager allows you to add hardware to empty slots and change the hardware in occupied slots. The procedure for both is the same.

**Note:** When you change hardware in an occupied slot for which there are circuits configured, the Configuration Manager automatically deletes the circuits.

You specify hardware from the Wellfleet Configuration Manager Window, as follows:

1. Under Description, click on the slot for which you wish to specify hardware.

   The Configuration Manager displays the Module List Window (see Figure 2-9) which lists Wellfleet's Link Modules and their corresponding MIB identifiers.

2. Drag the scroll bar to see all Link Modules.

3. Select the proper Link Module for the slot and click on the Change Module button.

   The Wellfleet Configuration Manager Window reappears displaying the Link Module you selected in step 3 in the slot you selected in step 1.

Repeat this procedure for each slot for which you wish to specify hardware.

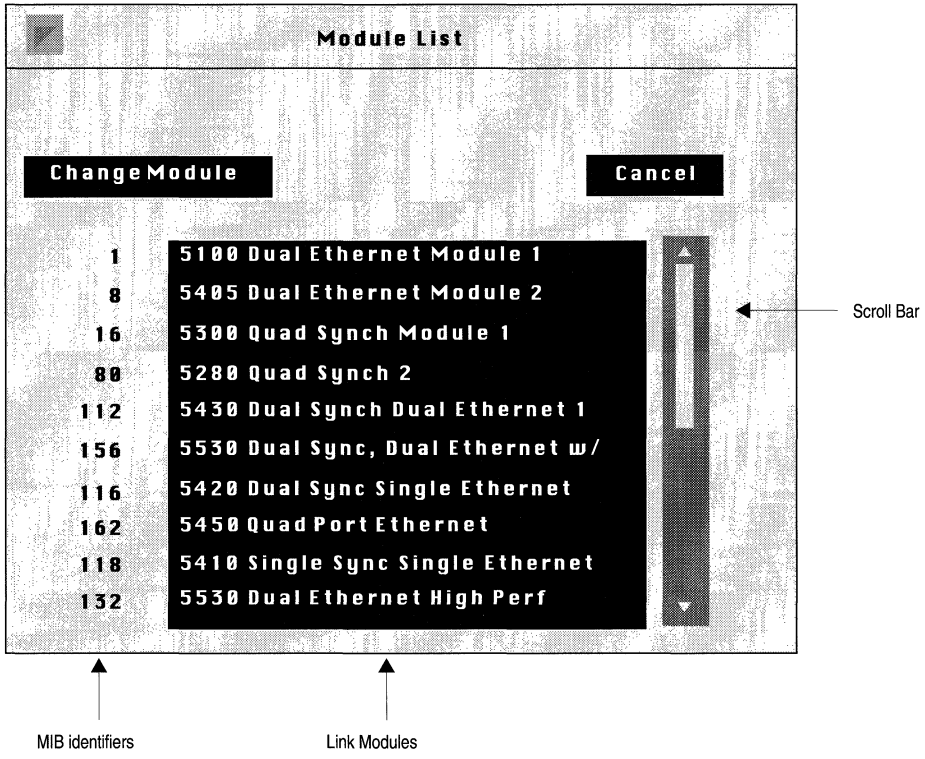


**Figure 2-9. Module List Window**

## Technician Interface (TI) Console Connection and Administrative Information

The following sections describe how to configure:

- Technician Interface (TI) console connection

  The console port on the System Resources Link Module (SRM-L) connects the BN to the TI. You use the TI during BN installation to enable the BN's start-up configuration and as an emergency interface when the Site Manager is unavailable. The console port allows you to connect the BN to the TI either directly or via a modem.

- BN administrative information

  The Configuration Manager allows you to specify a system name, contact, and location for the BN. When you access the BN from the Wellfleet Site Manager Window, the application automatically retrieves this administrative information and displays it in the Wellfleet Site Manager Window.

### Editing the Technician Interface (TI) Console Parameters

You edit console parameters in the Console Window (see Figure 2-10). You display this window from the Wellfleet Configuration Manager Window by clicking on the console port displayed in slot 1.

This section describes how to edit the parameters displayed in the Console Window. For each parameter, it provides the following:

- Wellfleet default
- Range of valid settings
- Parameter's function
- Instructions for setting the parameter

Refer to this information to edit the parameters you wish to change. When you are done, click on the Save button to exit the window and save your changes.

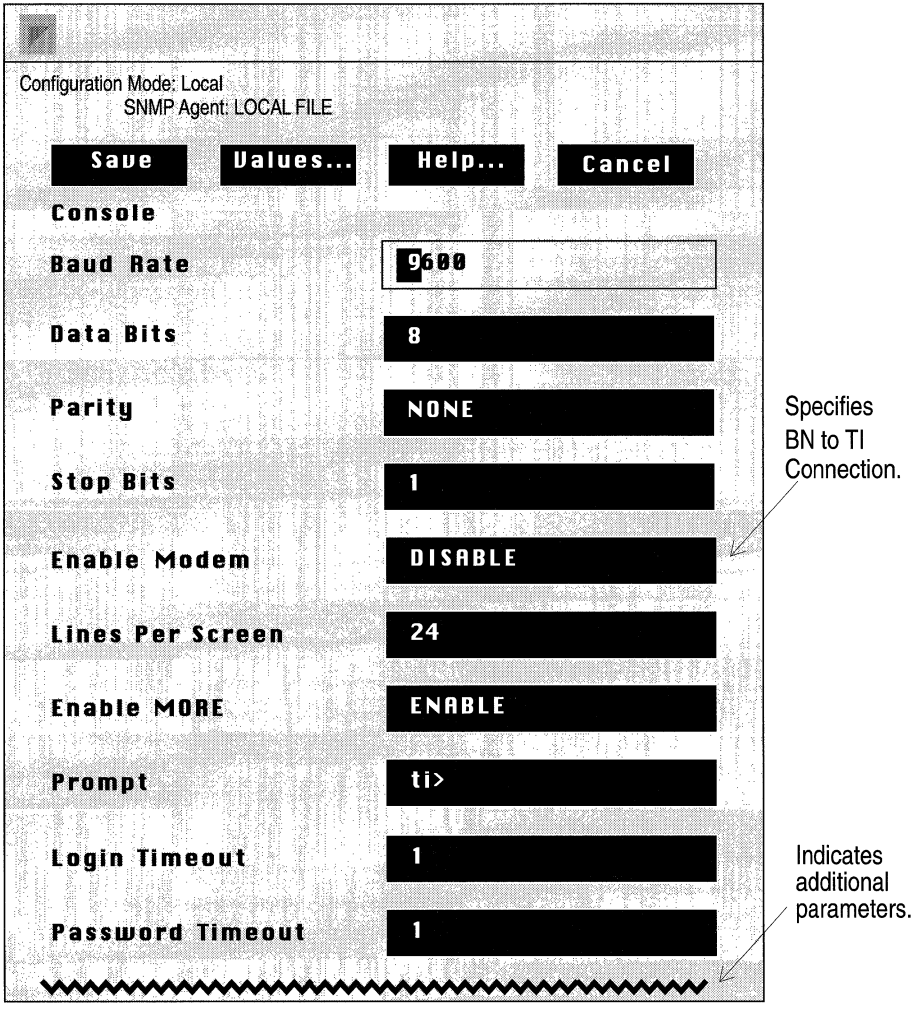


**Figure 2-10. Console Window**

| | |
|---|---|
| **Parameter :** | **Baud Rate** |
| Wellfleet Default: | 9600 |
| Options: | 9600, 4800, 1200, 600, 300 |
| Function: | Specifies the rate of data transfer between the console and the router. |
| Instructions: | Set according to your console requirements. |

| | |
|---|---|
| **Parameter :** | **Data Bit** |
| Wellfleet Default: | 8 |
| Options: | 7 or 8 |
| Function: | Specifies the number of bits in each ASCII character received or transmitted by the router. |

| | |
|---|---|
| **Parameter :** | Set according to your console requirements. |

| | |
|---|---|
| **Parameter :** | **Parity** |
| Wellfleet Default: | None |
| Options: | None, Odd, Even |
| Instructions: | Enables or disables data error detection for each character transmitted or received. |
| Instructions: | Set according to your console requirements.Odd or Even enables data error detection. None disables data error detection. |

| | |
|---|---|
| **Parameter :** | **Stop Bits** |
| Wellfleet Default: | 1 |
| Options: | 1, 1.5, 2 |
| Function: | Specifies the number of bits that follow each ASCII character received or transmitted by the router. |
| Instructions: | Set according to your console requirements. |

| | |
|---|---|
| **Parameter :** | **Enable Modem** |
| Wellfleet Default: | Disable |
| Options: | Enable/Disable |
| Function: | Specifies whether the terminal is connected directly or via a modem to the TI. |
| Instructions: | Select Enable if the terminal is connect via a modem to the TI.<br>Select Disable if the terminal is connected directly to the TI. |

| | |
|---|---|
| **Parameter :** | **Lines Per Screen** |
| Wellfleet Default: | 24 |
| Options: | 0 to 512 |
| Function: | Specifies the maximum number of lines displayed on the console screen. |
| Instructions: | Set according to your console requirements. |

| | |
|---|---|
| **Parameter :** | **Enable More** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Specifies whether the TI pauses after the screen fills with data. |
| Instructions: | Select Enable to configure the TI to pause after the screen fills with data.<br>Select Disable to configure the TI not to pause after the screen fills with data. |

| | |
|---|---|
| **Parameter :** | **Prompt** |
| Wellfleet Default: | ti> |
| Options: | Any string of up to 19 keyboard characters except for control key sequences. |
| Function: | Specifies the text used as a prompt on your console screen. |
| Instructions: | Accept the default or enter a different text string. |

| | |
|---|---|
| **Parameter :** | **Login Timeout** |
| Wellfleet Default: | 1 |
| Options: | 1 to 99 (99 indicates infinity) |
| Function: | Specifies the number of minutes to time out when the Enter key has not been pressed after the *Login:* prompt. This parameter is valid only when Enable Modem is set to Enable. |
| Instructions: | Accept the default or enter a new timeout value. |

| | |
|---|---|
| **Parameter :** | **Password Timeout** |
| Wellfleet Default: | 1 |
| Options: | 1 to 99 (99 indicates infinity) |
| Function: | Specifies the number of minutes to time out when the enter key has not been pressed after the *Password*: prompt. This parameter is valid only when Enable Modem is set to Enable. |
| Instructions: | Accept the default or enter a new timeout value. |

| | |
|---|---|
| **Parameter :** | **Command Timeout** |
| Wellfleet Default: | 15 |
| Options: | 1 to 99 (99 indicates infinity) |
| Function: | Specifies the number of minutes to time out when no one has pressed the Enter key after the prompt determined by the Prompt parameter. This parameter is valid only when Enable Modem is set to Enable. |
| Instructions: | Accept the default or enter a new timeout value. |

| | |
|---|---|
| **Parameter :** | **Login Retries** |
| Wellfleet Default: | 3 |
| Options: | 1 to 99 (99 indicates infinity) |
| Function: | Specifies the maximum number of login attempts when Enable Modem is set to Enable. |
| Instructions: | Accept the default or enter a new retry value. |

## Specifying Administrative Information

You specify BN administrative information in the System Parameters Window (see Figure 2-11). To display this window, select System Record in the menu bar of the Wellfleet Configuration Manager Window. Enter the required information and then click on the Save button to exit the window and save the administrative information.

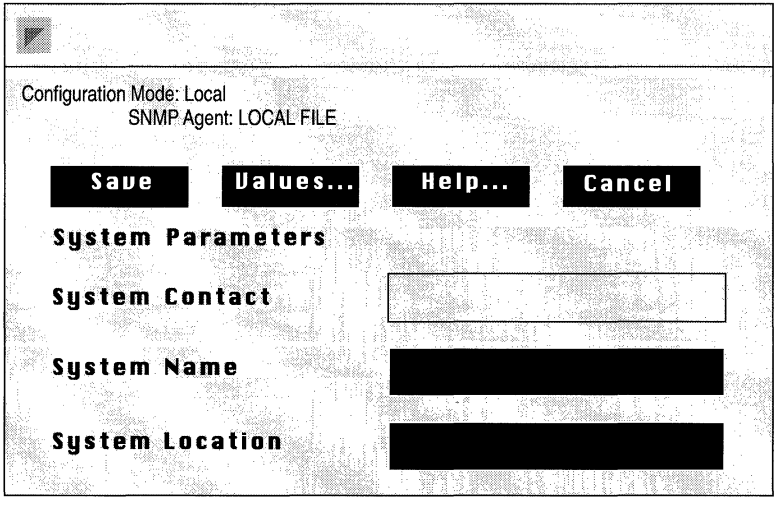


**Figure 2-11. System Parameters Window**

| **Parameter :** | **System Contact** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any person(s) |
| Function: | Provides the name of the person to contact regarding issues with this BN. |
| Instructions: | Enter name of contact, and possibly a way to contact that person. |

| | |
|---|---|
| **Parameter :** | **System Name** |
| Wellfleet Default: | None |
| Options: | Any name |
| Function: | Identifies this BN. |
| Instructions: | Enter the appropriate name for this BN. |

| | |
|---|---|
| **Parameter :** | **System Location** |
| Wellfleet Default: | None |
| Options: | Any place |
| Function: | Identifies the physical location of this BN. |
| Instructions: | Enter a location description for this BN. |

# Chapter 3

## Configuring Circuits

## List of Figures

## List of Tables

# Configuring Circuits

## About This Chapter

This chapter explains how to use the Configuration Manager first to configure circuits on the BN, and second to establish default network layer bridging and routing on each configured circuit. This procedure is easy since you only need to select a circuit connector and specify a few parameters for each protocol you add to a circuit. The Configuration Manager automatically provides the rest of the necessary circuit information. That is, once you select a connector to which a circuit will interface, the Configuration Manager provides both a default circuit name (which you may accept or change) and appropriate data-link layer parameter defaults (depending on the connector type). When you designate routing or bridging protocols to run across the circuit, the Configuration Manager queries for required protocol values and then provides a set of network layer protocol-specific parameter defaults. Data link layer and network layer default values are suitable for most networks; however, you can change them if they are inappropriate for your network topology.

You should use this chapter if you are enhancing the pilot configuration (described later in this chapter) you created using the Quick-Start procedure described in the *Quick-Start Guide*, or if you are configuring circuits, which includes:

- Adding circuits to the BN
- Editing circuit parameters, which includes:
  - Deleting a circuit from a BN
  - Renaming a circuit

— Adding/deleting routing or bridging protocols to a circuit

— Assigning additional IP addresses to a circuit

— Editing data link or network layer parameters

# Enhancing the Pilot Configuration

When you finish with the Quick-Start procedure (described in the *Quick-Start Guide*), the pilot configuration running on your BN consists of two IP interfaces: the initial interface you configured using the TI, and the pilot IP interface you configured using the Site Manager.

At this point, you probably want to enhance the pilot configuration so that it matches your actual network requirements.

**Note:** For instructions on creating new configuration files for the BN from scratch, refer to the section *Configuration Steps for Each Operating Mode* in the chapter *Configuration Manager Overview.*

In order to enhance the pilot configuration, you must first retrieve it from the BN. Beginning at the Wellfleet Site Manager window, retrieve the pilot configuration file as follows:

1. Select the Config/Remote option in the Wellfleet Site Manager Window to display the Edit Remote Configuration window (Figure 3-1).
2. Enter the pilot configuration file name, as follows:
   — Select the number of the volume (or slot) containing the configuration file (by default, the Configuration Manager displays volume 2). Simply click on the Volume box to display the available volumes. The directory of the chosen volume will appear in the directory box.

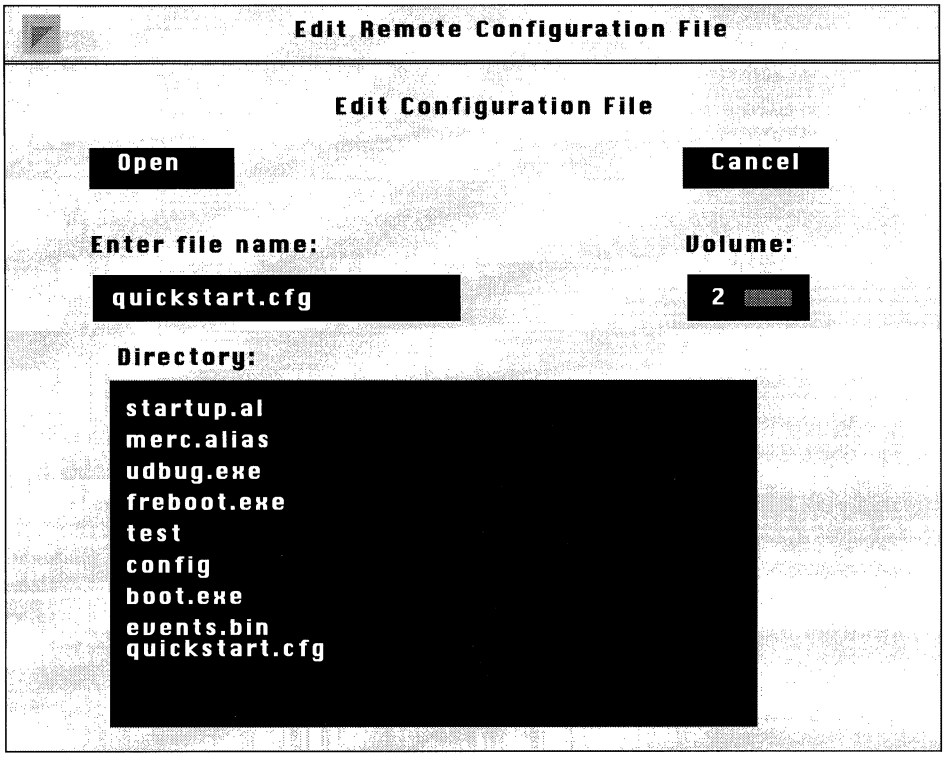


**Figure 3-1. Edit Remote Configuration File Window**

— Enter the configuration file name (in this example *quickstart.cfg*) and click on the Open button.

The Configuration Manager uses TFTP to retrieve and display the specified file in the Wellfleet Configuration Manager Window; the file name you specified is displayed in the upper-left corner of the window. If the specified file does not exist, the Wellfleet Configuration Manager Window displays the hardware configuration and no circuits.

Depending on how you wish to enhance the file, proceed to one of the following sections:

— To add new circuits to the pilot configuration, proceed to the section *Adding a Circuit to the BN*.

— To add additional bridging or routing protocols to the pilot interfaces, proceed to the section *Editing Circuits*.

— To edit the line (data link level) details for the pilot interfaces, proceed to the section *Editing Line Details for a Circuit*.

— To edit the default IP parameters that you accepted when you configured the pilot IP interfaces, proceed to the chapter *Configuring IP*.

**Note:** Before you enable bridging and routing protocols on a circuit, you may wish to refer to the chapter that describes the protocol you are enabling. For example, if you have questions about the VINES routing protocol, see the chapter *Configuring VINES* for an overview of the protocol, and implementation guidelines for enabling it on your network. (Each protocol chapter also describes how to reset the default parameters for the protocol after you have enabled it on a circuit).

# Adding a Circuit to the BN

To add a circuit to the BN, you first select a connector (port) on the BN. Then, you name the circuit that interfaces with the selected connector and save the circuit. Finally, you enable the circuit with routing or bridging protocols.

Begin at the Wellfleet Configuration Manager Window, and complete the following steps:

1. Select the Circuits/Add Circuit option.

   The Add Circuit Window appears (see Figure 3-2).

2. Click on a BN connector.

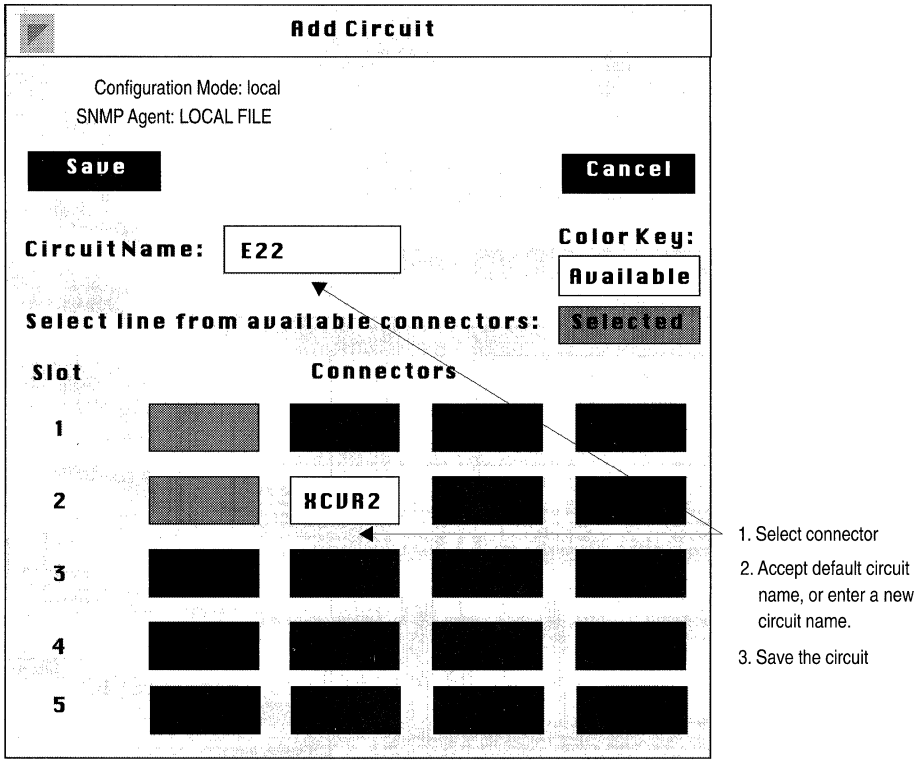


**Figure 3-2. Add Circuit Window**

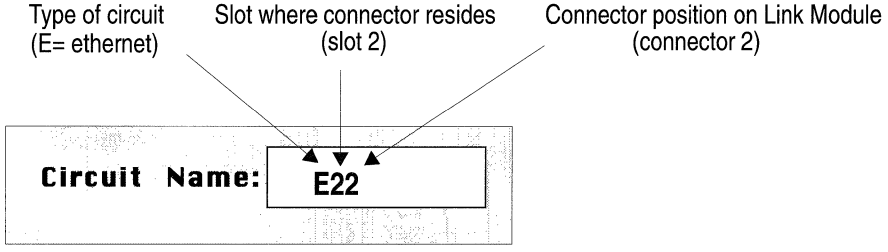


**Figure 3-3. Configuration Manager Default Circuit Name**

After you click on a connector, the Configuration Manager names the circuit that interfaces with this connector in the circuit name box. The default circuit name describes the type of circuit and the location (slot and number) of the connector with which this circuit interfaces (see Figure 3-3).

**Note:** Default circuit designators are E for Ethernet and E1circuits, F for FDDI circuits, H for HSSI circuits, O for Token Ring circuits, S for synchronous (point to point) circuits, and T for T1 circuits. As the defaults for Ethernet and E1 are identical, you may want to assign new names to E1 circuits to avoid possible confusion.

If you wish, you can change the circuit name by clicking on the circuit name box and typing in a new circuit name which may consist of any printable character.

3. Click on the Save button.

**Note:** Save each circuit after you associate it with a connector.

After you save the circuit, the Configuration Manager displays the Select Protocols Window (Figure 3-4). You enable the circuit with routing or bridging protocols from this window. In addition, you can enable the Spanning Tree algorithm for the Bridge, RIP or OSPF for IP, RIP for IPX, RIP for XNS, and Protocol Prioritization for synchronous, E1, and T1 lines.

**Note:** The Select Protocols Window differs slightly according to circuit type. For point to point circuits (E1, HSSI, synchronous, or T1), the window provides access to all routing and bridging protocols as well as to protocol prioritization functions (as shown in Figure 3-4). For LAN (Ethernet, Token Ring, or FDDI) circuits, the window provides access to all protocols except for SMDS, Frame Relay, and Protocol Prioritization which are unsupported by these media.

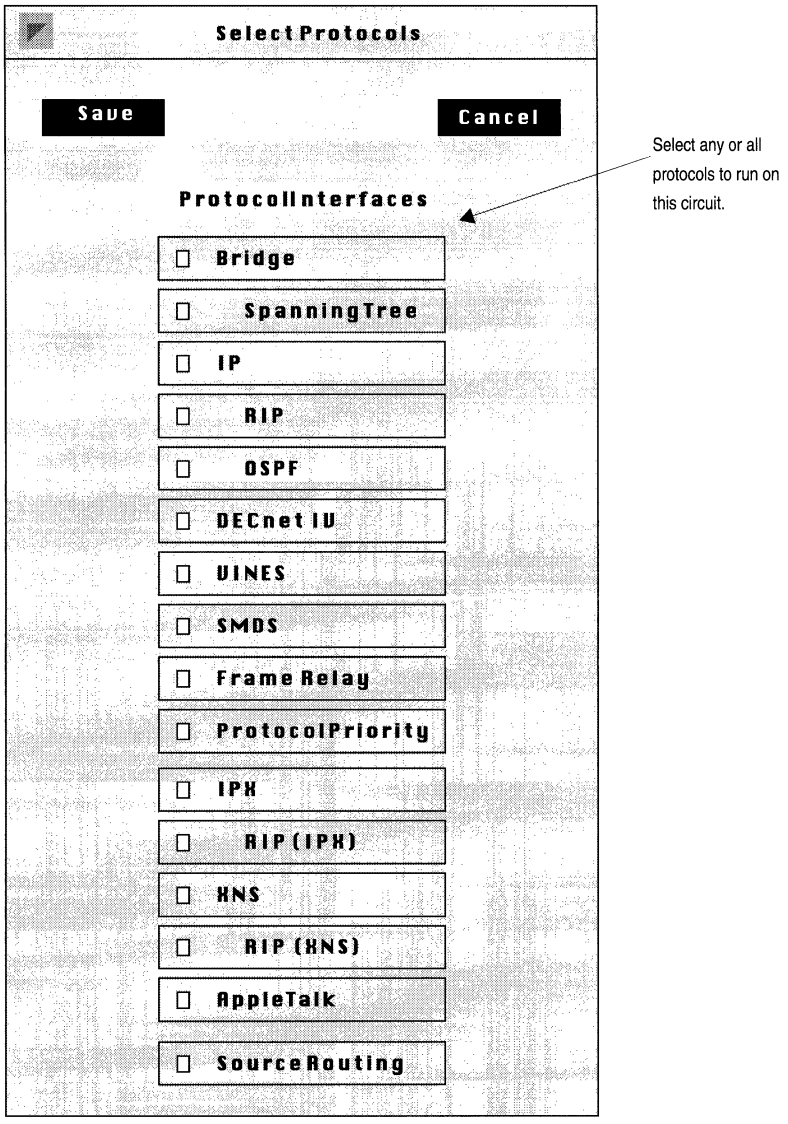


**Figure 3-4. Select Protocols Window**

4. Select all routing or bridging protocols that you want enabled on this circuit; then click on the Save button.

   If you select Spanning Tree, the Bridge is automatically enabled; similarly, if you select RIP or OSPF, IP is automatically enabled; if you select RIP (IPX) or RIP (XNS), IPX and/or XNS is automatically enabled.

   For each protocol you enable, the Configuration Manager generally displays a protocol-specific configuration window prompting for additional required information. Some protocols (for example, Frame Relay) require no additional information to provide default service. In such cases, you can skip to step 6.

5. Define each protocol you added to the circuit.

   The following sections describe how to define all routing and bridging protocols for a circuit. You only need to refer to the sections that correspond to the protocols and services you enabled.

   On each protocol window, you'll first specify the required configuration information. Then you have the option to either accept or edit default settings for the protocol. Finally, you'll save the protocol information. The Configuration Manager then displays the configuration window for the next protocol enabled on the circuit.

   Once you have defined all protocols enabled on the circuit, the Configuration Manager returns to the Wellfleet Configuration Manager Window. The connector box is now highlighted to indicate that the circuit has been added.

6. Repeat steps 1 to 5 until all BN circuits are added and defined.

## Defining the Bridge

If you enabled the Bridge on the circuit, but did not enable the Spanning Tree algorithm, you do not specify any configuration information for the Bridge. Instead, the Configuration Manager returns to the Wellfleet Configuration Manager Window (if the Bridge alone is enabled on the circuit) or the configuration window for another routing protocol.

The Configuration Manager solicits additional information only if you enabled the Spanning Tree algorithm. With the algorithm enabled, the Configuration Manager then displays the Spanning Tree Configuration Window (Figure 3-5).

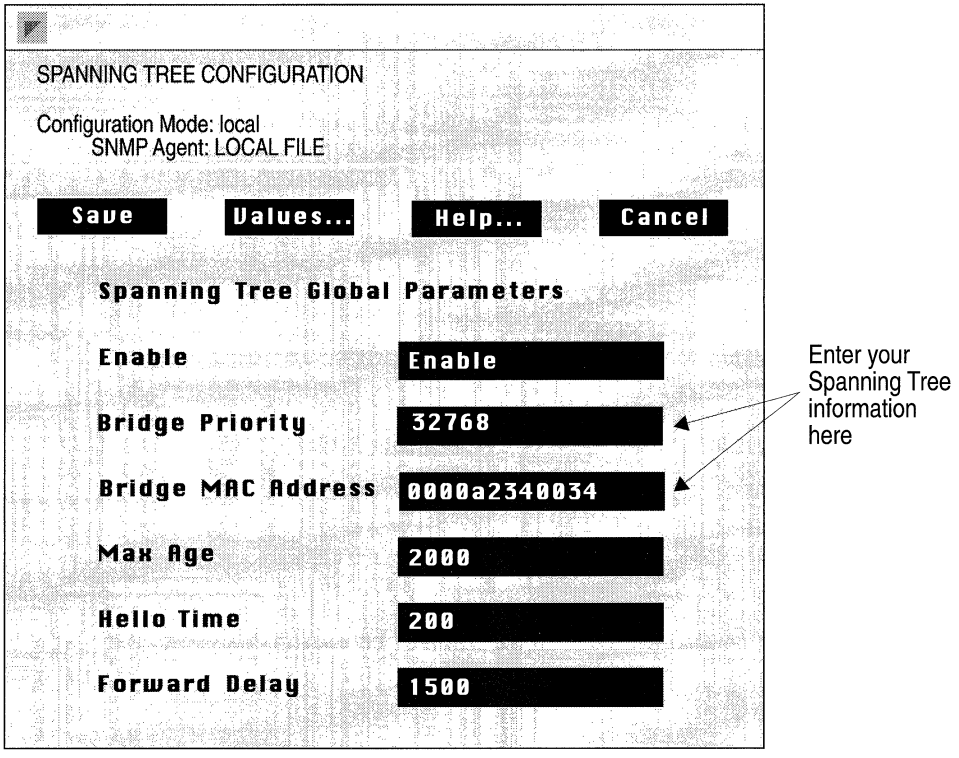


**Figure 3-5. Spanning Tree Configuration Window**

Complete the following steps:

1. Assign values to the two required (Bridge Priority and Bridge MAC Address) Spanning Tree parameters as described below.
2. Click on the Save button.
3. When the screen prompts *Do you want to edit the Spanning Tree Interface Parameters?* either
   - — click Cancel to enable default Spanning Tree service
   - — click OK, and then proceed to the section *Editing Spanning Tree Interface Parameters* in the chapter *Configuring the Bridge*.

Once you click the Save button, you cannot access the Spanning Tree Configuration Window again. To edit any Spanning Tree global parameter, refer to the section *Editing Spanning Tree Global Parameters* in the chapter, *Configuring the Bridge*.

**Note:** Because the Spanning Tree is global (that is, it runs across all Bridge circuits), the Configuration Manager only displays the Spanning Tree Configuration Window the *first* time you specify Spanning Tree for the Bridge. After that, you never again specify global spanning tree configuration information.

| | |
|---|---|
| **Parameter :** | **Bridge Priority** |
| Wellfleet Default: | None |
| Options: | 0 - 65535 |
| Function: | In conjunction with the Bridge MAC Address parameter, assigns a 64-bit Bridge ID to the BN. Bridge Priority supplies the most significant 16 bits of the Bridge ID, while Bridge MAC Address supplies the remaining (least significant) 48 bits.<br><br>The Bridge ID is used by the Spanning Tree in the selection of the Root Bridge. In selecting the Root Bridge, the Spanning Tree chooses the bridge with the lowest number Bridge ID. Thus, the lower the Bridge Priority, the more likely the BN will be selected as the Root Bridge. |
| Instructions: | Enter a decimal value from 0 to 65536. |

| | |
|---|---|
| **Parameter :** | **Bridge MAC Address** |
| Wellfleet Default: | None |
| Options: | Any valid 48-bit MAC-level address |
| Function: | In conjunction with the Bridge Priority parameter, assigns a 64-bit Bridge ID to the BN. Bridge Priority supplies the most significant 16 bits of the Bridge ID, while Bridge MAC Address supplies the remaining (least significant) 48 bits.<br><br>The Bridge ID is used by the Spanning Tree in the selection of the Root Bridge. In selecting the Root Bridge, the Spanning Tree chooses the bridge with the lowest number Bridge ID. Thus, the lower the Bridge Priority, the more likely the BN will be selected as the Root Bridge. In the event of equal Bridge Priority values, the Bridge MAC Address value determines the bridge's priority. |
| Instructions: | Enter a 48-bit MAC address expressed as a 12-digit hexadecimal value. Wellfleet recommends that you set Bridge MAC Address to the MAC address of one of the BN's Spanning Tree ports, preferably the one with the lowest priority. |

After you assign values to Bridge Priority and Bridge MAC Address, the BN provides default bridging services as described in the section *Bridge Parameters* in Appendix A. To alter this default service refer to the chapter *Configuring the Bridge*.

## Defining IP

If you enabled IP on the circuit, the Configuration Manager displays the IP Configuration Window (Figure 3-6).

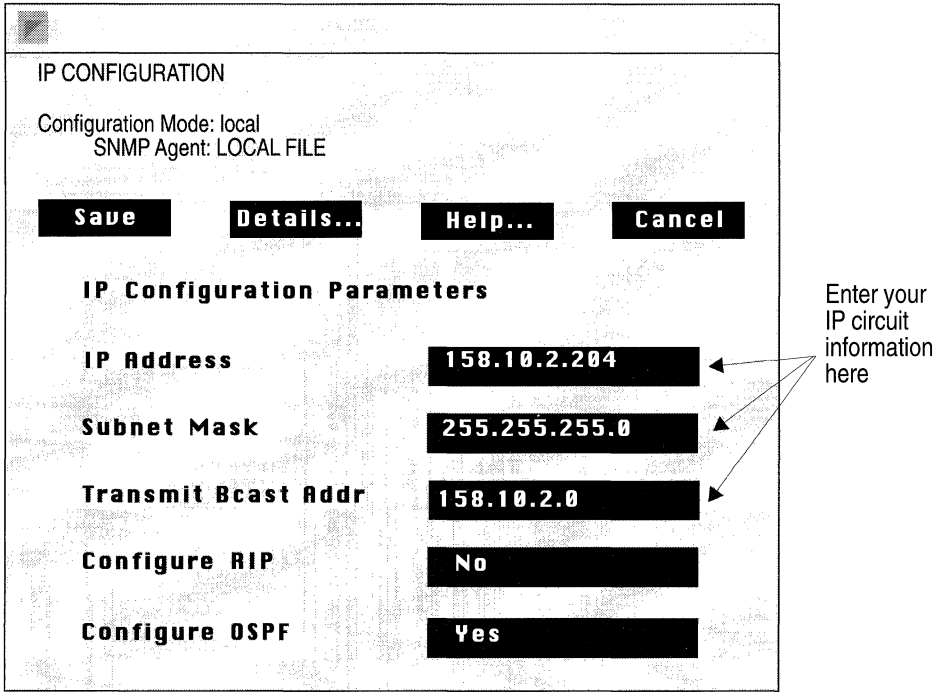


**Figure 3-6. IP Configuration Window**

Complete the following steps:

1. Assign values to the three required (IP Address, Subnet Mask, and Transmit Bcast Addr) IP interface parameters, and one optional (Area Address) OSPF parameter as described below.
2. Click on the Save button to enable default IP service.

   Or, alternatively

   Click on the Details button to access and edit IP interface parameters. For instructions on editing these parameters refer

to the section *Editing IP Interface Parameters* in the chapter *Configuring IP.*

Once you click the Save button in the IP Configuration Window you cannot access this window for this IP interface again. To edit any IP interface parameter, refer to the section *Editing IP Interfaces* in the chapter, *Configuring IP.*

| | |
|---|---|
| **Parameter :** | **IP Address** |
| Wellfleet Default: | None |
| Options: | Any valid IP address |
| Function: | Assigns a 32-bit IP address to the interface. |
| Instructions: | Enter the IP address of the interface in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Subnet Mask** |
| Wellfleet Default: | None |
| Options: | Depends on the class of the network to which the interface connects. |
| Function: | Specifies the network and subnetwork portion of the 32-bit IP address. |
| Instructions: | Enter the subnet mask in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Transmit Bcast Addr** |
| Wellfleet Default: | None |
| Options: | 0 or any valid IP broadcast address |
| Function: | Specifies the broadcast address that the IP router uses to broadcast packets across this interface. |
| Instructions: | Enter 0 to configure the IP router to use an all-1s address for broadcasting packets; or enter the broadcast address in dotted decimal notation. |

**Note:** The Wellfleet Configuration Manager enables RIP and/or OSPF as determined by your selections at the Select Protocols Window.

If you configured OSPF, the BN displays the interface-specific OSPF Area Address Window (Figure 3-7).

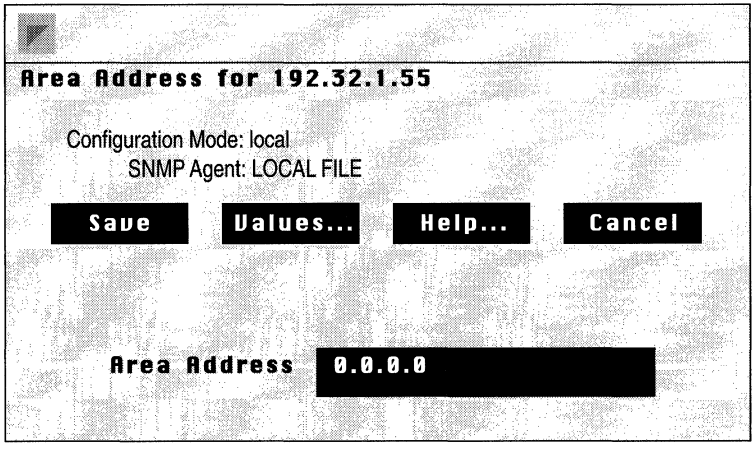


**Figure 3-7. OSPF Area Address Window**

Complete the following steps:

1. Designate the OSPF area to which the interface connects.
2. Click on the Save button.

| **Parameter :** | **Area Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any four octet number in dotted decimal notation |
| Function: | Identifies the OSPF area to which this interface belongs. |
| Instructions: | Enter the appropriate area ID in dotted decimal notation. |

**Note:** The backbone area ID is always 0.0.0.0.

After you assign values to IP Address, Subnet Mask, Transmit Bcast Addr and Area Address (if you configured OSPF), the BN provides default IP and OSPF services as described in the section *IP Parameters* and *OSPF Parameters* in Appendix A. To alter this default service refer to the sections *Configuring IP* and *Configuring OSPF.*

## Defining DECnet Phase IV

If you enable DECnet Phase IV on the circuit, the Configuration Manager displays the DECnet Phase IV Configuration Window (Figure 3-8).

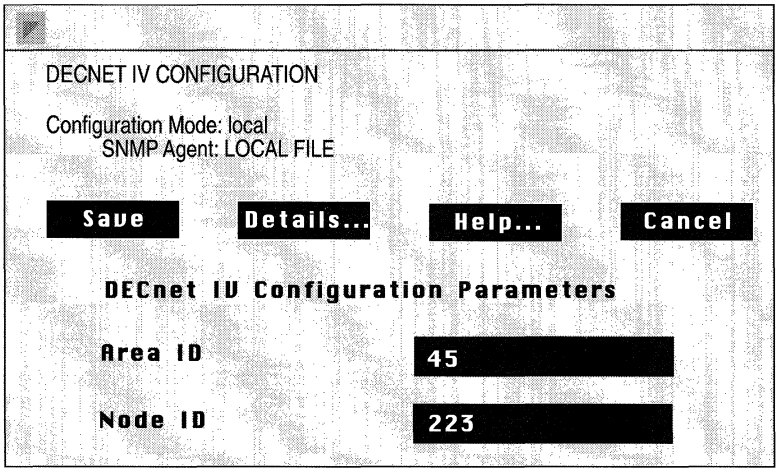


**Figure 3-8. DECnet Phase IV Configuration Window**

Complete the following steps:

1. Assign values to the two required (Area ID and Node ID) DECnet Phase IV interface parameters as described below.
2. Click on the Save button to enable default DECnet Phase IV service.

Or, alternatively

Click on the Details button to access and edit DECnet interface parameters. For instructions on editing these parameters refer to the section *Editing DECnet Interface Parameters* in the chapter *Configuring DECnet Phase IV.*

Once you click the Save button in the DECnet Phase IV Configuration Window you cannot access this window for this circuit again. To edit any DECnet Phase IV parameter, refer to the section *Editing DECnet Interface Parameters* in the chapter *Configuring DECnet Phase IV.*

| | |
|---|---|
| **Parameter :** | **Area ID** |
| Wellfleet Default: | None |
| Options: | 1 -63 |
| Function: | Specifies a unique DECnet Phase IV Area ID for this circuit.<br><br>The Area ID is the first six bits of a DECnet Phase IV node address. You specify the Area ID on a circuit-by-circuit basis; that is a single router may have individual circuits residing in different areas. |
| Instructions: | Enter the Area ID assigned to this circuit. |

| | |
|---|---|
| **Parameter :** | **Node ID** |
| Wellfleet Default: | None |
| Options: | 1-1024 |
| Function: | Specifies a unique intra-area DECnet Phase IV Node ID for this circuit.<br>The Node ID is the last 10 bits of a DECnet Phase IV node address.<br>Note that if individual circuits on a router reside in different areas, then each circuit may have a different node address. |
| Instructions: | Enter the Node ID assigned to the BN. |

After you assign values to Area ID and Node ID, the BN provides default DECnet Phase IV service as described in the section *DECnet Phase IV Router Parameters* in Appendix A. To alter this default service refer to the chapter *Configuring DECNET Phase IV.*

## Defining VINES

If you enabled VINES on the circuit, you need not specify any configuration information. In response to the *Do you want to edit the VINES interface details?* query either:

— Click on the Cancel button to enable default VINES service.

Or, alternatively

— Click on the OK button to access and edit VINES interface parameters. For instructions on editing these parameters refer to the section *Editing VINES Interface Parameters* in the chapter *Configuring VINES*.

The BN provides default VINES service as described in the section *VINES Parameters* in Appendix A. To alter this default service refer to the chapter *Configuring VINES*.

## Defining SMDS

If you enable SMDS (Switched Multi-Megabit Data Service) on the circuit, the Configuration Manager displays the SMDS Configuration Window (Figure 3-9).

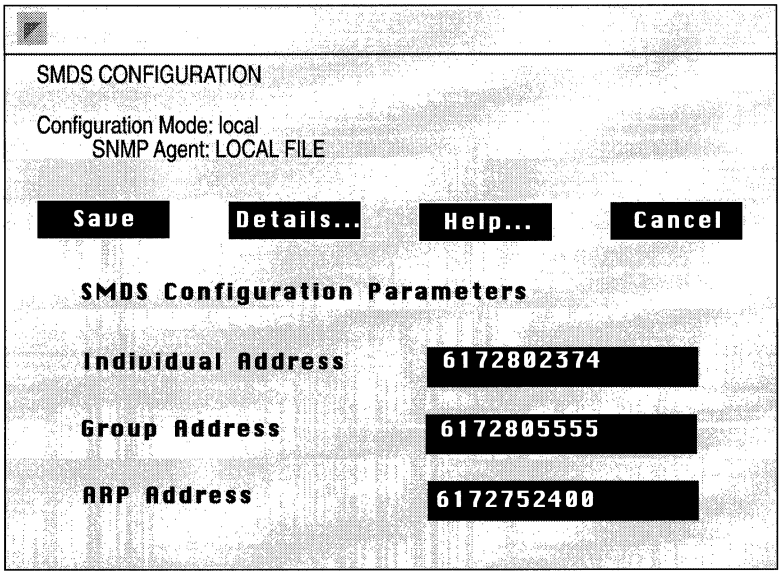


**Figure 3-9. SMDS Configuration Window**

Complete the following steps:

1. Assign values to the three required (Individual Address, Group Address, and ARP Address) SMDS interface parameters as described below.
2. Click on the Save button to enable default SMDS service.

   Or, alternatively

   Click on the Details button to access and edit SMDS interface parameters. For instructions on editing these parameters refer to the section *Editing SMDS Interface Parameters* in the chapter *Configuring SMDS*.

| | |
|---|---|
| **Parameter :** | **Individual Address** |
| Wellfleet Default: | None |
| Options: | Any valid 10-digit North American Numbering Plan (NANP) telephone number. |
| Function: | Provides a local MAC-layer address. |
| Instructions: | Enter the 10-digit local address (telephone number) as provided by the SMDS subscription agreement. |

| | |
|---|---|
| **Parameter :** | **Group Address** |
| Wellfleet Default: | None |
| Options: | Any valid 10-digit North American Numbering Plan (NANP) telephone number. |
| Function: | Provides a MAC-layer multicast address. |
| Instructions: | Enter the 10-digit multicast address as provided by the SMDS subscription agreement. |

| | |
|---|---|
| **Parameter :** | **ARP Address** |
| Wellfleet Default: | None |
| Options: | Any valid 10-digit North American Numbering Plan (NANP) telephone number. |
| Function: | Provides an address resolution multicast address. |
| Instructions: | Enter the 10-digit multicast address as provided by the SMDS subscription agreement. |

After you assign values to Individual Address, Group Address, and ARP Address, the BN provides default SMDS service as described in the section *SMDS Parameters* in Appendix A. To alter this default service refer to the chapter *Configuring SMDS*.

## Defining Frame Relay

If you enabled Frame Relay on the circuit, you need not specify any configuration information.

The BN provides default Frame Relay service as described in the section *Frame Relay Parameters* in Appendix A. To alter this default service refer to the section *Configuring Frame Relay.*

## Defining Protocol Priority

Each BN point to point interface is serviced by three types of priority queues. These queues (in actuality a series of buffers) are designated High, Normal, and Low priority. As the designations suggest, the BN implements a dequeuing algorithm which grants transmit precedence first to high, then to normal, and finally to low precedence traffic.

**Note:** Because of the high bandwidth availability of the media, Protocol Priority is not supported over HSSI.

The BN also provides a mechanism to assign priority to traffic types of your choosing. This mechanism allows for *content-based* and *length-based* priority designation. Content-based prioritization assigns priority based on the contents of pre-defined or user-defined fields within either the data link or IP header. Content-based prioritization is fully explained in the chapter *Configuring Protocol Prioritization.* Length-based prioritization assigns priority based on packet length. In contrast with content-based prioritization, length-based prioritization must be implemented at the circuit level.

Protocol Priority enables you to construct a *length-based priority filter* associated with a specific circuit. A priority filter consists of two elements: (1) a set of conditions against which a packet presented for transmission across a specific interface is compared, and (2) an action that is applied to packets which meet the specified conditions. The conditions associated with a length-based filter specify packets by encapsulation type, protocol, and packet length. Packets which meet the specified conditions are queued for normal transmission; packets

which fail to meet the conditions are queued for low priority transmission.

If you enable Protocol Prioritization on the circuit, the Configuration Manager displays the Protocol Prioritization Configuration Window (Figure 3-10).

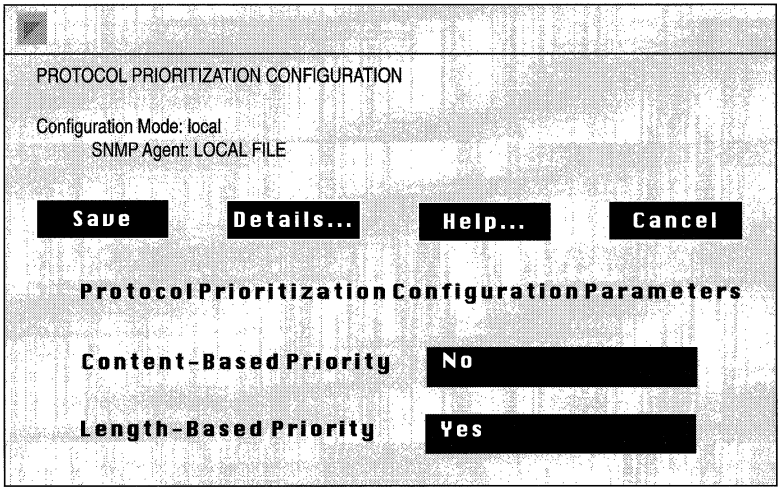


**Figure 3-10. Protocol Prioritization Configuration Window**

Complete the following steps:

1. Assign values to the two required (Content-Based Priority and Length-Based Priority) Protocol Priority interface parameters.
2. Click on the Save button.
3. Optionally, configure content-based filters (refer to chapter *Configuring Protocol Prioritization* for detailed information of content-based filters).

| | |
|---|---|
| **Parameter :** | **Content-Based Priority** |
| Wellfleet Default: | None |
| Options: | Yes or No |
| Function: | Implements priority queueing based on contents of specified fields within the data link or IP header. |
| Instructions: | Set to No if you do not wish to implement content-based prioritization. Set to Yes if you wish to implement content-based prioritization.<br><br>If you set to Yes, the Content-Based Priority Configuration Window will appear when you exit this screen. You may now configure a content-based priority filter by referring the chapter *Configuring Protocol Prioritization*; or you may click Cancel to defer content-based configuration until a later time. |

| | |
|---|---|
| **Parameter :** | **Length-Based Priority** |
| Wellfleet Default: | None |
| Options: | Yes or No |
| Function: | Implements priority queueing based on packet length. |
| Instructions: | Set to Yes to enable length-based prioritization. |

If you specify length-based prioritization, the BN displays the Length Based Priority Interface Window. Click Add to display the Length Based Priority Configuration Window (Figure 3-11).

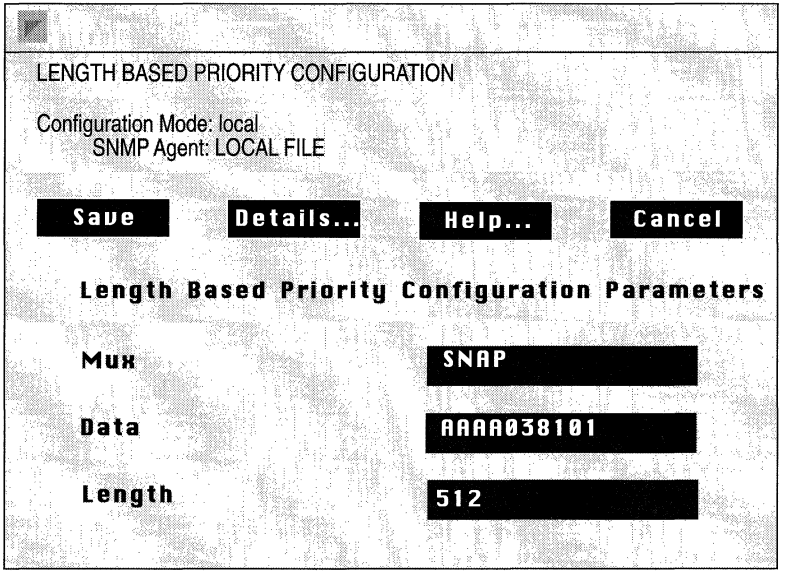


**Figure 3-11. Length Based Priority Configuration Window**

Complete the following steps:

1. Assign values to the three required (Mux, Data, Length) length-based priority configuration parameters described below.
2. Click on the Save button.

| **Parameter :** | **Mux** |
|---|---|
| Wellfleet Default: | None |
| Options: | ETYPE, LSAP, SNAP |
| Function: | Specifies one of the conditions against which traffic presented for transmission will be evaluated, namely the encapsulation method. |
| Instructions: | Select the encapsulation method. |

| **Parameter:** | **Data** |
|---|---|
| Wellfleet Default: | None |
| Options: | A protocol identifier expressed as a 5-byte (10-digit) hexadecimal number. |
| Function: | Specifies one of the conditions against which traffic presented for transmission will be evaluated, namely the protocol type.<br><br>If MUX is equal to ETYPE, enter the IEEE-assigned Ethernet Type of the encapsulated protocol traffic followed by 000000. For example, to specify DECnet Phase IV traffic enter 6003000000.<br><br>If MUX is equal to LSAP (generally used with VINES and Spanning Tree), enter the LLC1 identifier followed by 0000. For example, to specify Spanning Tree traffic enter 4242030000.<br><br>If MUX is equal to SNAP, enter the IEEE-assigned Ethernet Type of the encapsulated protocol traffic preceded by the OUI (generally, but not always 0). For example, to specify DECnet Phase IV traffic enter 0000006003; to specify AppleTalk ARP, enter 0000F880F3. |
| Instructions: | Enter the appropriate protocol identifier. |

| **Parameter :** | **Length** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any packet length |
| Function: | Specifies one of the conditions against which traffic presented for transmission will be evaluated, namely the packet length. Packets of the specified encapsulation and protocol types which are equal to or shorter than the value specified by this parameter are queued for normal transmission. Packets of the specified encapsulation and protocol types which are longer than the value specified by this parameter are queued for low priority transmission. |
| Instructions: | Enter the packet length. |

After you assign values to Mux, Data, and Length, the BN provides default Length-Based Protocol Prioritization service as described in the section *Protocol Prioritization Parameters* in Appendix A. To alter this default service or to configure content-based prioritization, refer to the chapter *Configuring Protocol Prioritization*.

## Defining IPX

If you enabled IPX on the circuit, the Configuration Manager displays the IPX Configuration Window (Figure 3-12).

Complete the following steps:

1. Assign values to the required (Network Address) IPX interface parameter as described below.
2. Click on the Save button to enable default IPX service.

   Or, alternatively

   Click on the Details button to access and edit IPX interface parameters. For instructions on editing these parameters refer to the section *Editing IPX Interface Parameters* in the chapter *Configuring IPX*.

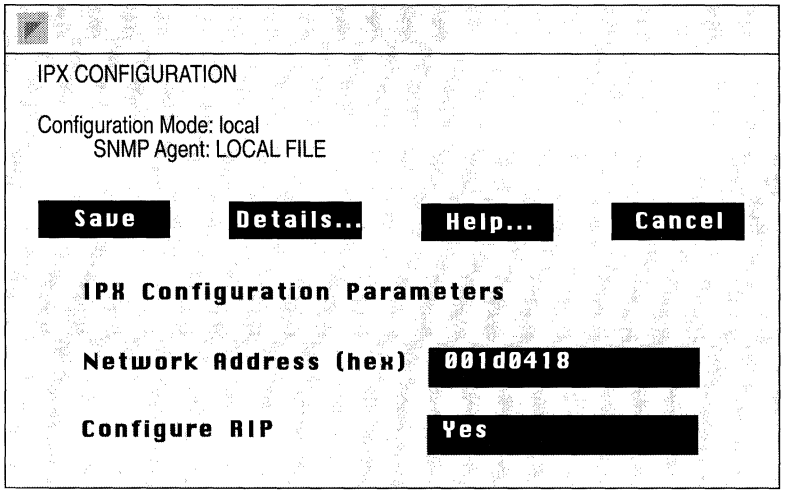


**Figure 3-12. IPX Configuration Window**

Once you click the Save or Details button in the IPX Configuration Window you cannot access this window for this IPX interface again. To edit any IPX interface parameters, refer to the section *Editing IPX Interfaces* in the chapter *Configuring IPX*.

| **Parameter :** | **Network Address (hex)** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid IPX network address |
| Function: | Assigns an IPX address to the interface. |
| Instructions: | Enter the IPX address of the interface in hexadecimal notation. |

**Note:** The Wellfleet Configuration Manager sets the Configure RIP parameter as determined by your selection in the Select Protocols Window. You can, however, change the value of this parameter in the IPX Configuration Window. Once you save this window you cannot change its settings.

After you assign a value to Network Address (hex), the BN provides default IPX service as described in the section *IPX Parameters* in Appendix A. To alter this default service refer to chapter *Configuring IPX*.

## Defining XNS

If you enabled XNS on the circuit, the Configuration Manager displays the XNS Configuration Window (Figure 3-13).

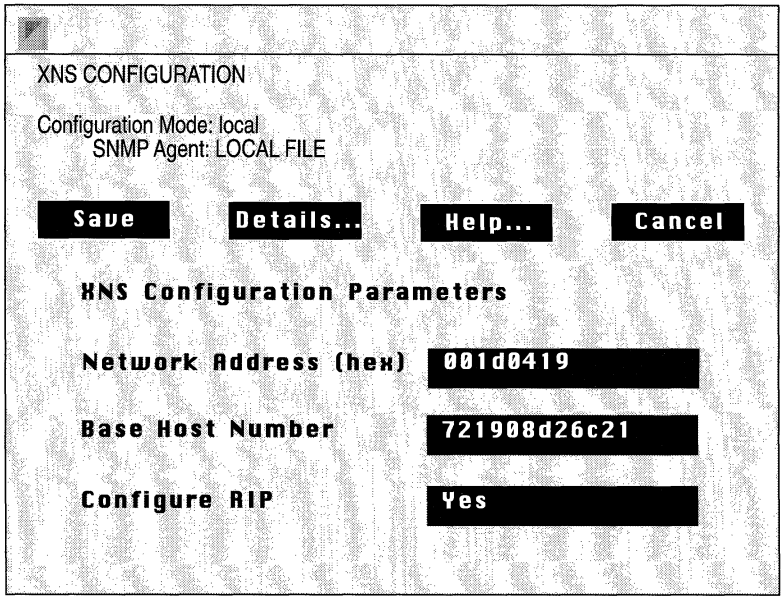


**Figure 3-13. XNS Configuration Window**

Complete the following steps:

1. Assign values to the two required (Network Address and Base Host Number) XNS interface parameters as described below.
2. Click on the Save button to enable default XNS service.

   Or, alternatively

   Click on the Details button to access and edit XNS interface parameters. For instructions on editing these parameters refer to the section *Editing XNS Interface Parameters* in the chapter *Configuring XNS*.

Once you click the Save button in the XNS Configuration Window you cannot access this window for this XNS interface again. To edit any XNS interface parameters, refer to the section *Editing XNS Interfaces* in the chapter *Configuring XNS*.

| **Parameter :** | **Network Address (hex)** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid XNS network address |
| Function: | Assigns an XNS address to the interface. |
| Instructions: | Enter the XNS address of the interface in hexadecimal notation. |

| **Parameter :** | **Base Host Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid XNS host address |
| Function: | Assigns an XNS host address to the BN and provides a default MAC-level address. |
| Instructions: | Enter the XNS host address in hexadecimal notation. |

**Note:** The Wellfleet Configuration Manager sets the Configure RIP parameter as determined by your selection in the Select Protocols Window. You can, however, change the value of this parameter in the XNS Configuration Window. Once you save this window you cannot change its settings.

After you assign values to Network Address (hex) and Base Host Address, the BN provides default XNS service as described in the section *XNS Parameters* in Appendix A. To alter this default service refer to the chapter *Configuring XNS*.

## Defining AppleTalk

If you enabled AppleTalk on the circuit, the Configuration Manager displays the AppleTalk Configuration Window (Figure 3-14).

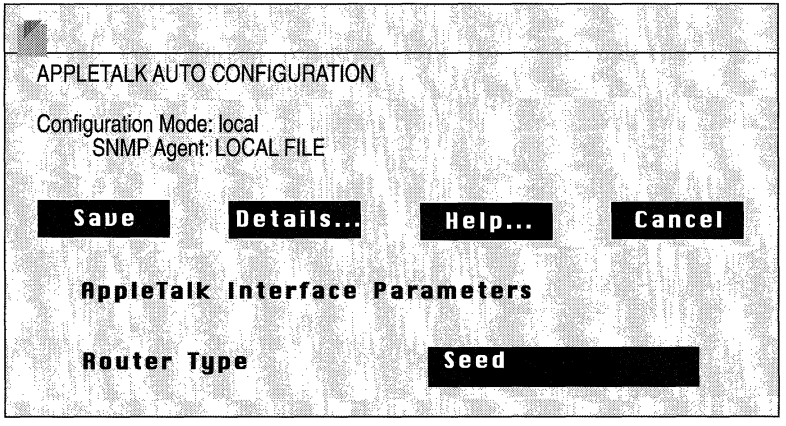


**Figure 3-14. AppleTalk Configuration Window**

Complete the following steps:

1. Assign a value, Seed or Nonseed, to the required (Router Type) AppleTalk interface parameter.

   If you need more information on how you should define the router type for this interface, refer to the section *When Should I*

*Configure My Wellfleet AppleTalk Router as a Seed Router?* in the *Configuring AppleTalk* chapter of this guide.

2. Click on the Save button.

   If you designate the interface as seed, the AppleTalk Configuration window appears (see Figure 3-15). You now specify seed router information as described in the following sections.

   If you designate the interface as nonseed, the BN provides default AppleTalk service as described in the section *AppleTalk Parameters* in Appendix A. To alter this default service refer to the chapter *Configuring AppleTalk*.

| **Parameter :** | **Router Type** |
|---|---|
| Wellfleet Default: | None |
| Options: | Seed on Nonseed |
| Function: | Designates whether the interface functions as an AppleTalk Seed or Nonseed router.<br><br>A Seed router provides other AppleTalk routers with configuration information (specifically the range of network numbers contained within the AppleTalk internet and zone names within the AppleTalk internet).<br><br>Every AppleTalk network must contain at least one Seed router. If the network contains multiple Seed routers, each such router must be configured with identical Network Start, Network End, and Default Zone information.<br><br>Nonseed routers receive network range and zone name information from Seed routers. |
| Instructions: | Select Seed or Nonseed. |

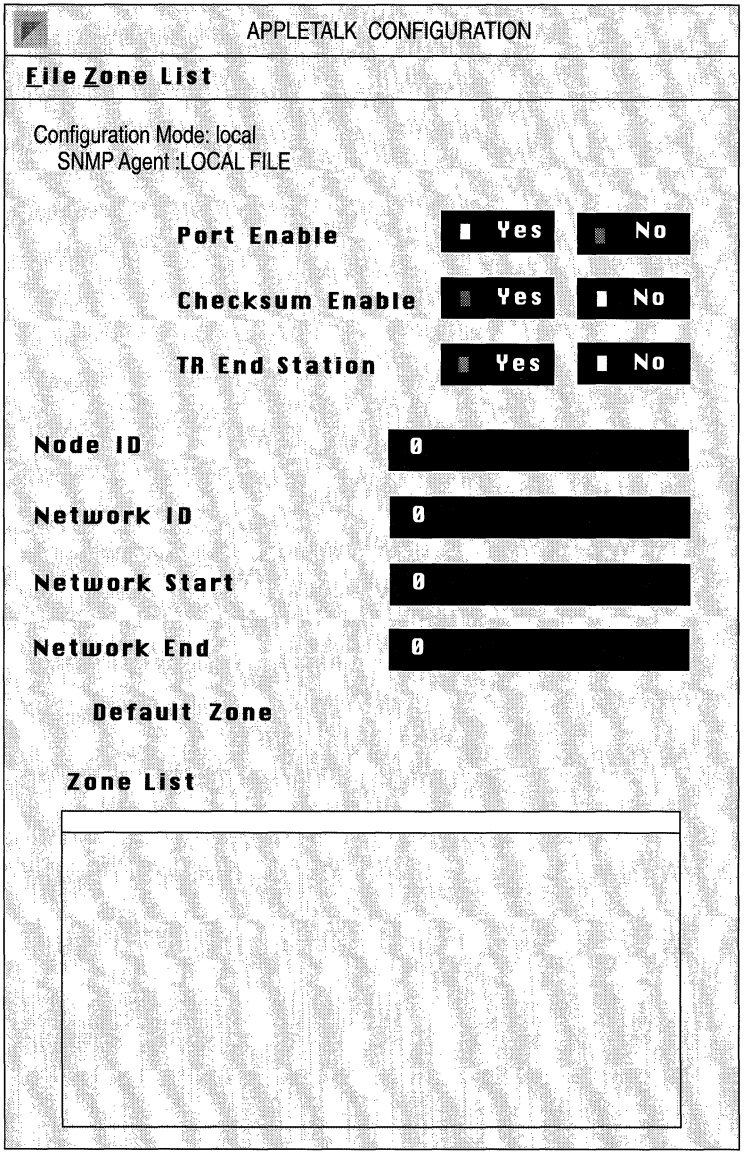


**Figure 3-15. AppleTalk Configuration Window**

### Specifying Seed Router Information

If you designate the BN as a Seed router, the Configuration Manager displays the AppleTalk Configuration Window (Figure 3-15).

Complete the following steps:

1. If you are configuring the AppleTalk interface in Dynamic configuration mode, reset the Port Enable parameter to Enable.

   If you are configuring the AppleTalk interface in either Local or Remote configuration mode, the Port Enable parameter is set to Enable by default. Accept the default setting.

2. Accept the default settings for Checksum Enable and TR End Station.

3. Assign values to the Node ID and Network ID parameters.

**Note:** Wellfleet strongly recommends accepting the default value 0 for the Node ID and Network ID parameters. This allows the BN to dynamically assign a unique AppleTalk address to the circuit. The only exception to this is if you are configuring AppleTalk on a synchronous (point-to-point) circuit, in which case you must specify a different AppleTalk address for each end point of the circuit.

4. Assign values to the Network Start, Network End, and Default Zone parameters.

   You must define these three parameters in order for the interface to act as a seed router. *If you do not assign values to these parameters, the interface becomes nonseed by default.*

5. If you are configuring a seed router for a network containing multiple zones, then proceed to the section *Defining a Zone List* to specify the zone list for the seed router.

   If the default zone is the only zone assigned to this network, then simply select File/Done.

Once you save the seed router information, the BN provides default AppleTalk service as described in the section *AppleTalk Parameters* in Appendix A. To alter this default service refer to the chapter *Configuring AppleTalk*.

| | |
|---|---|
| **Parameter :** | **Node ID** |
| Wellfleet Default: | None |
| Options: | Any decimal value from 0 to 253. |
| Function: | Specifies the node number assigned to this AppleTalk interface. Each AppleTalk node must have a network-unique Node Id. |
| Instructions: | Wellfleet recommends accepting the default 0. When you accept the default 0, the BN dynamically acquires a Node ID for the interface during startup - thus ensuring that the AppleTalk address for this interface is unique within the network.<br><br>If you change the default value, the router uses the Node ID that you specify. |

| | |
|---|---|
| **Parameter :** | **Network ID** |
| Wellfleet Default: | None |
| Options: | Any decimal value from 0 to 65279. |
| Function: | Specifies the network number assigned to this AppleTalk interface. |
| Instructions: | Wellfleet recommends accepting the default 0. When you accept the default 0, the router dynamically acquires a Network ID for the interface during startup - thus ensuring that the AppleTalk address for this interface is unique within the network.<br><br>If you change the default value, make certain that the number you specify is within the correct network range. The router then uses the Network ID that you specify. |

**Note:** If the AppleTalk address (Node ID/Network ID) matches that of any other node on the network, the interface automatically disables.

**Parameter :** **Network Start**

Wellfleet Default: 0

Options: Any decimal value from 1 to 65279.

Function: Specifies the lowest boundary (minimum) of the range of network numbers that are available for use by nodes on the network to which this interface connects.

This parameter's setting (together with the Router Type parameter) determines whether or not this interface functions as a seed router. A seed router supplies the Network Start, Network End, Default Zone and Zone List information for all other nonseed routers on this network. A nonseed router acquires this information from the other seed routers on the network.

Each network must contain at least one seed router.

Instructions: To configure this interface as a seed router, specify the Network Start as follows:

— If this is the only seed router on the network, check your network topology map to see how your network is divided and enter the lowest boundary network number here.

— If there are already seed routers on the network, enter the *same* Network Start value that is configured on all other seed routers.

**Note:** If you specify a Network Start other than the default 0, then 1) the router becomes a seed router automatically and 2) you must also specify values for the Network End and Default Zone parameters.

| | |
|---|---|
| **Parameter :** | **Network End** |
| Wellfleet Default: | None |
| Options: | Any decimal value from 1 to 65279. |
| Function: | Specifies the upper boundary (maximum) of the range of network numbers that are available for use by nodes on the network to which this interface connects.<br><br>This parameter is used in conjunction with the Network Start parameter to help define a seed router. *If you have not specified a Network Start, this parameter is ignored.* |
| Instructions: | To configure this interface as a seed router, specify the Network End as follows:<br><br>— If this is the only seed router on the network, check your network topology map to see how your network is divided and enter the upper boundary network number here.<br><br>— If there are already seed routers on the network, enter the *same* Network End value that is configured on all other seed routers. |

| | |
|---|---|
| **Parameter :** | **Default Zone** |
| Wellfleet Default: | None |
| Options: | Any valid zone name. |
| Function: | Specifies the name of the default zone where all new nodes are assigned when they first start up on this network.<br><br>This parameter is used in conjunction with the Network Start and Network End parameters to help define a seed router. *If you have not specified a Network Start, this parameter is ignored.* |
| Instructions: | To configure this interface as a seed router, then specify the Default Zone as follows:<br><br>— If this is the only seed router on the network, enter any valid Default Zone name.<br><br>— If there are already seed routers on the network, enter the *same* Default Zone name as is configured on all other seed routers.<br><br>A valid zone name can consist of 1 to 32 characters and can include any keyboard character (except the * character). |

### Defining a Zone List

This section describes how to define the zone list of an AppleTalk interface that you are *initially* configuring.

**Note:** If this AppleTalk interface has already been enabled on your network, and you are modifying the zone list to add or delete a zone name from the list, then do *not* follow these instructions. Instead, refer to the section *Adding or Removing an AppleTalk Zone* in the chapter *Configuring AppleTalk* .

Beginning from the AppleTalk Configuration Window, you add a zone name to the interface's zone list as follows:

1. Select the Zone List/Add option to display the AppleTalk Zone List window.
2. Enter the zone name you wish to add to the list.

   The zone name can consist of 1 to 32 characters and include any keyboard character (except for the * character).
3. Click the Add Zone button.

   The zone name you specified has been added to the zone list.
4. Repeat steps 1-3 to add another zone name to the zone list, or select the File/Done option to save your changes and exit from the window.

Once you save the seed router information, the BN provides default AppleTalk service as described in the section *AppleTalk Parameters* in Appendix A. To alter this default service refer to the chapter *Configuring AppleTalk*.

## Defining Source Routing

If you enabled Source Routing on the circuit, you need not specify any configuration information.

The BN provides default Source Routing service as described in the section *Source Routing Parameters* in Appendix A. To alter this default service refer to the chapter *Configuring Source Routing*.

# Editing Circuits

The remaining sections describe how to edit circuit parameters. The Configuration Manager allows you to access all parameters associated with each circuit.

For each of the procedures described in the following sections, there are two ways to access the necessary windows. Both ways require you to begin at the Wellfleet Configuration Manager Window where you have two options:

- Selecting the Circuits/Edit Circuits option

  This is the procedure documented throughout the following sections.

- Clicking a connector

  This option calls the Editing Options Window for a *specific* circuit. From this window, you may select either the Edit Circuits button, which brings you to the Circuit Definitions Window; or the Edit Line button, which brings you to the appropriate line details window. This option does not allow for deleting a circuit from the BN.

## Deleting a Circuit from the BN

To delete a circuit from the BN, you begin at the Wellfleet Configuration Manager Window and complete the following steps:

1. Select the Circuits/Edit Circuits option.

   The Circuit List Window appears (Figure 3-16).

2. Select the circuit you want to delete from list of the circuits.

   In this example, circuit E21 is selected.

3. Click the Delete button; the Delete Circuit Window appears (Figure 3-17).

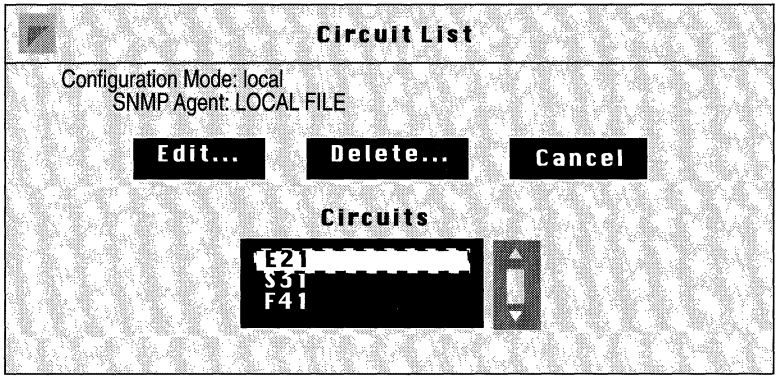


**Figure 3-16. Circuit List Window**

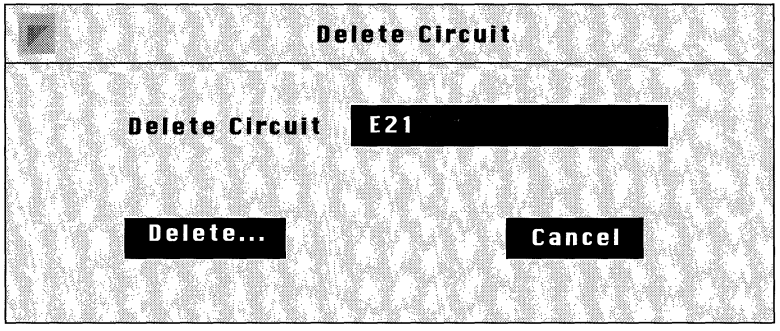


**Figure 3-17. Delete Circuit Window**

4. Click the Delete button in the Delete Circuit Window.

   The circuit is deleted from the BN, and no longer appears in the Circuit List Window.

Repeat steps 1 through 4 for each circuit that you want to delete from the BN.

## Renaming a Circuit

To rename a circuit on the BN, you begin at the Wellfleet Configuration Manager Window and complete the following steps:

1. Select the Circuits/Edit Circuits option from the Circuits menu.

   The Circuit List Window appears (see Figure 3-18).

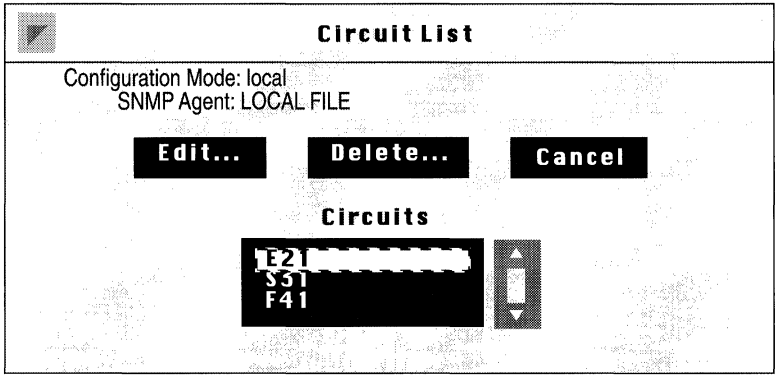


**Figure 3-18. Circuit List Window**

2. Select the circuit you want to rename from the list of circuits.

   In this example, circuit E21 is selected.

3. Click the Edit button; the Circuit Definition Window appears (see Figure 3-19).
4. Enter a new name for this circuit in the Circuit Name box.
5. Select the File/Save Lines button; the circuit's new name is saved.

Follow steps 1 through 5 for each circuit that you wish to rename.

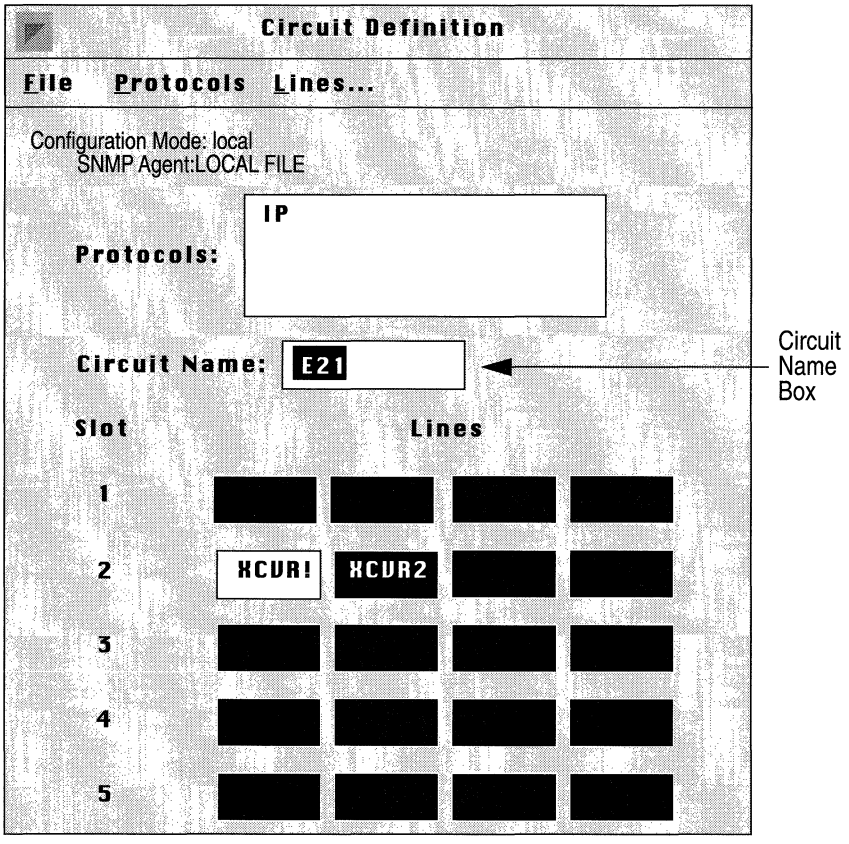


**Figure 3-19. Circuit Definition Window**

## Adding Protocols to a Circuit

To add protocols to a circuit, begin at the Wellfleet Configuration Manager Window and complete the following steps:

1. Select the Circuits/Edit Circuits option.

   The Circuit List Window appears (see Figure 3-20).

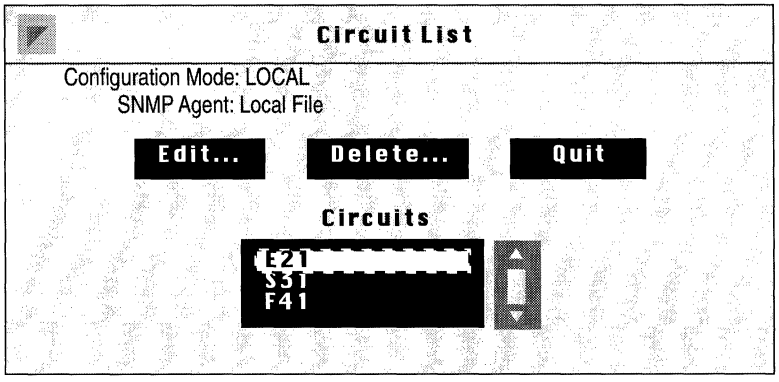


**Figure 3-20. Circuit List Window**

2. Select the circuit to which you want to add protocols.

   In this example, circuit E21 is selected.

3. Click the Edit button; the Circuit Definition Window appears (Figure 3-21).
4. Select the Protocols/Add/Delete option; the Select Protocols Window appears (previously shown in Figure 3-4).
5. Select the protocols that you want add to this circuit; then click on the Save button.

   For each protocol you add, the Configuration Manager displays a configuration window prompting you to define each protocol you enabled on the circuit.

6. Define each protocol you added to the circuit. To do this, refer to the corresponding section(s) earlier in this chapter.

Repeat steps 1-6 for each circuit to which you want to add protocols.

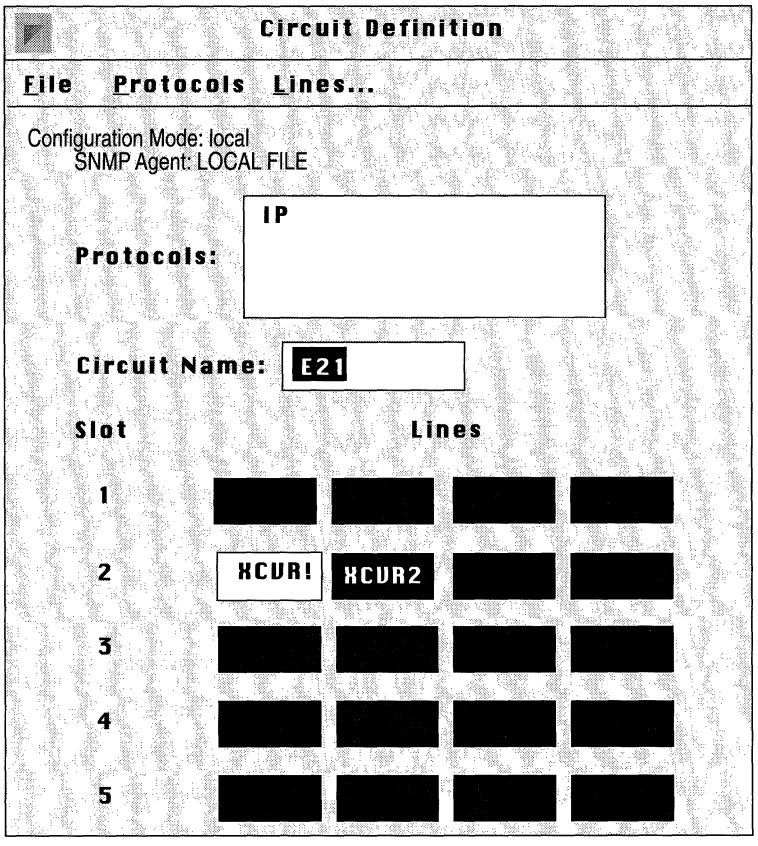


**Figure 3-21. Circuit Definition Window**

## Moving a Circuit

Once you have configured a circuit on a network interface, you can move the circuit to another network interface. When you move a circuit to a different type of network interface (for example, when you move an Ethernet circuit to an FDDI network interface connector), the Configuration Manager notes the type of connector to which the circuit now interfaces, and automatically provides the appropriate line detail parameters.

To move a circuit, begin at the Wellfleet Configuration Manager Window and complete these steps.

1. Select the Circuits/Edit Circuit option.

   The Circuit List Window appears (Figure 3-22).

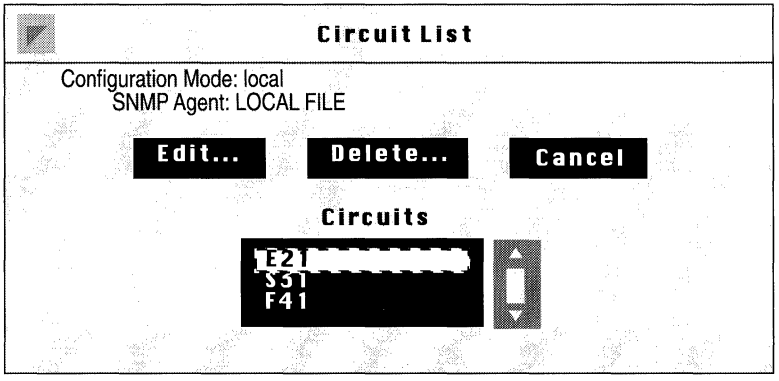


**Figure 3-22. Circuit List Window**

2. Select the circuit you want to move from list of the circuits.

   In this example, circuit E21 is selected.

3. Click the Edit button.

   The Circuit Definition Window appears (Figure 3-23). The name of the circuit you wish to move appears in the Circuit Name box, and the connector to which it interfaces is highlighted.

4. Click the circuit's connector.

   The circuit is removed from the connector, and the connector is no longer highlighted.

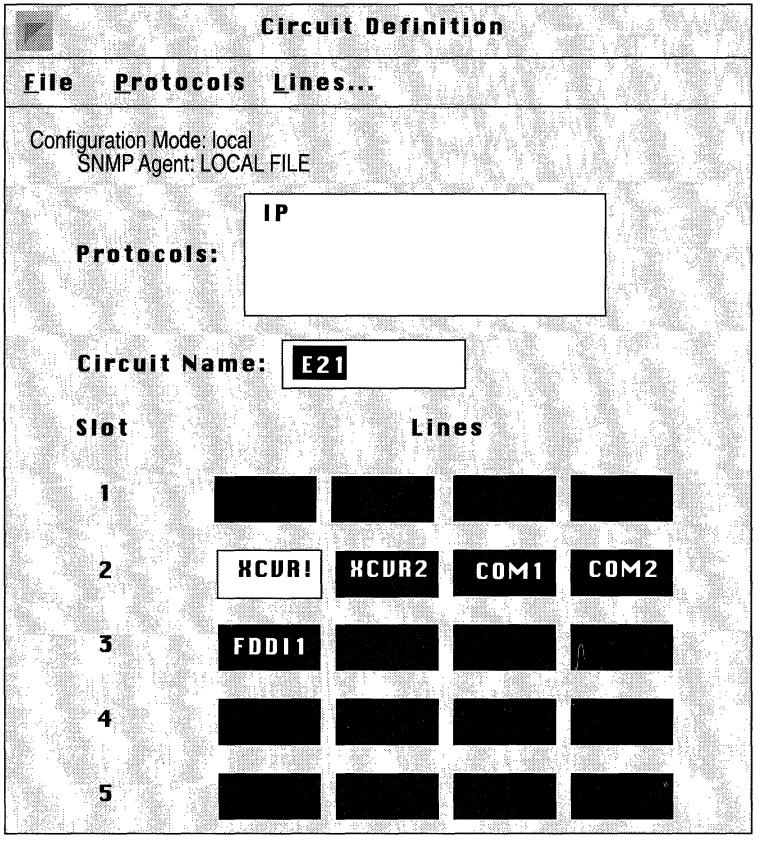


**Figure 3-23. Circuit Definition Window**

5. Click the connector with which you want this circuit to interface.

   The connector you chose is now highlighted, indicating that the circuit now connects to it.

   At this point, you may want to rename the circuit if you think the old circuit name may cause some confusion. To do this, simply enter a new name in the Circuit Name box. Select the File/Save Lines option.

6. Select the File/Cancel option to exit this window and return to the Wellfleet Configuration Manager Window.

   All circuit moves are reflected in this window.

Follow steps 1 through 6 for each circuit that you wish to move.

## Assigning an Additional IP Address to a Circuit

Wellfleet IP routing supports multinet. Multinet allows you to assign multiple IP addresses to a single circuit; thus, one circuit can support multiple IP network interfaces. For information about Multinet, see *Multinet* in the chapter *Configuring IP* .

You can assign as many IP addresses as you desire to a circuit. To assign an additional IP addresses to a circuit, begin at the Wellfleet Configuration Manager Window and complete the following steps:

1. Select the Circuits/Edit Circuits option.

   The Circuit List Window appears (see Figure 3-24).

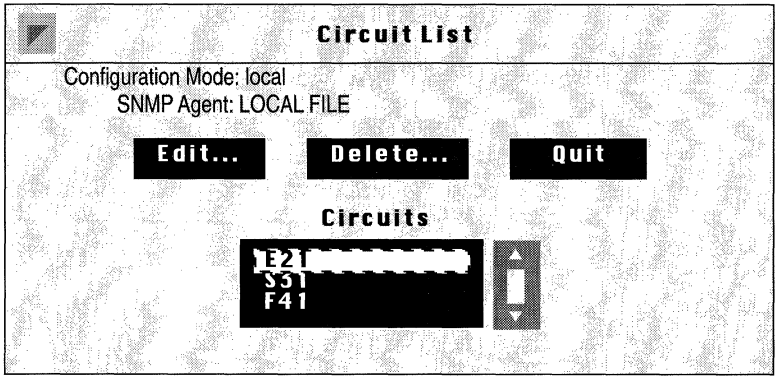


**Figure 3-24. Circuit List Window**

2. Select the circuit to which you want to assign an additional IP address, then click the Edit button.

   The Circuit Definition Window appears (see Figure 3-25).

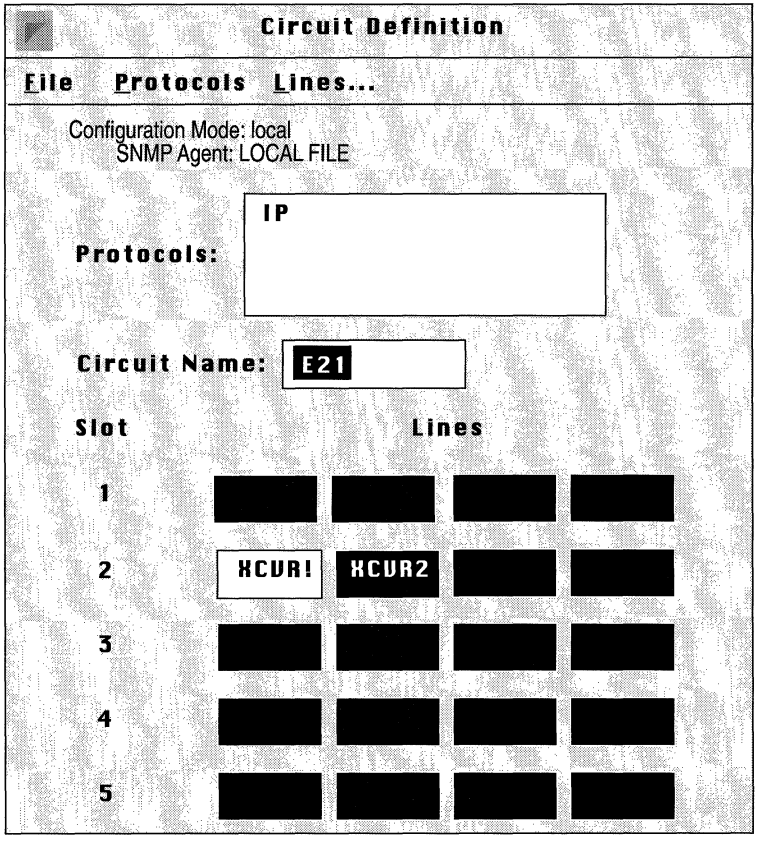


**Figure 3-25. The Circuit Definition Window**

3. Select the Protocols/Edit IP/Interfaces option.

   The IP Interface Window appears (see Figure 3-26).

4. Click the Add button.

   The IP Configuration Window appears (see Figure 3-27).

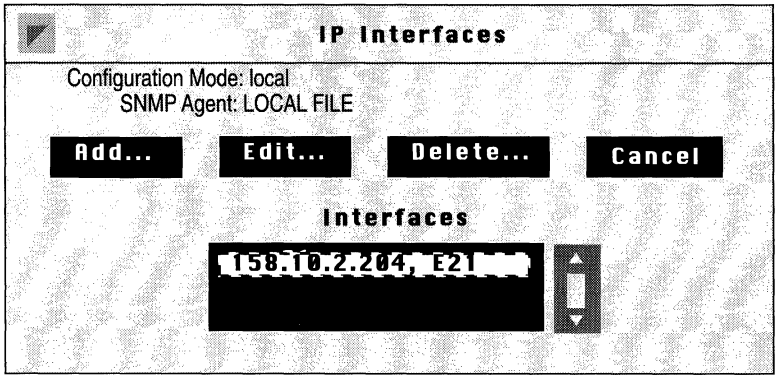


**Figure 3-26. IP Interfaces Window**

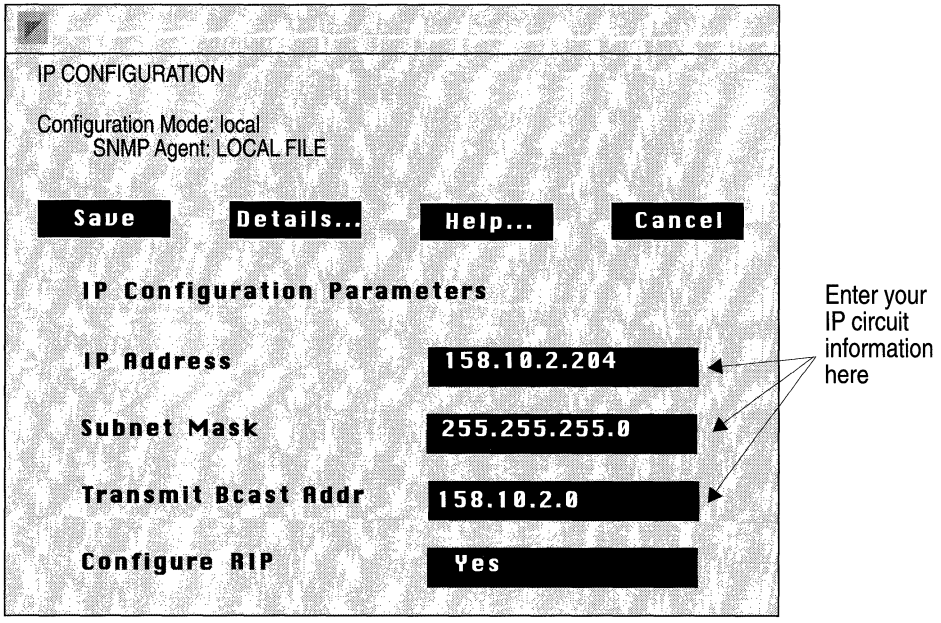


**Figure 3-27. IP Configuration Window**

5. Enter the IP address you wish to assign to this circuit in the IP Address box, then enter the Subnet Mask and Transmit Bcast Address for this circuit in the appropriate boxes.
6. Click the Save button.

   You are returned to the IP Interfaces Window. The address you just assigned to the specified circuit appears in the Interfaces scroll box.

Follow steps 1 through 6 for each IP address you wish to add to a circuit.

## Deleting Protocols from a Circuit

To delete protocols from a circuit, begin at the Wellfleet Configuration Manager Window and complete the following steps:

1. Select the Circuit/Edit Circuits option.

   The Circuit List Window appears (see Figure 3-28).

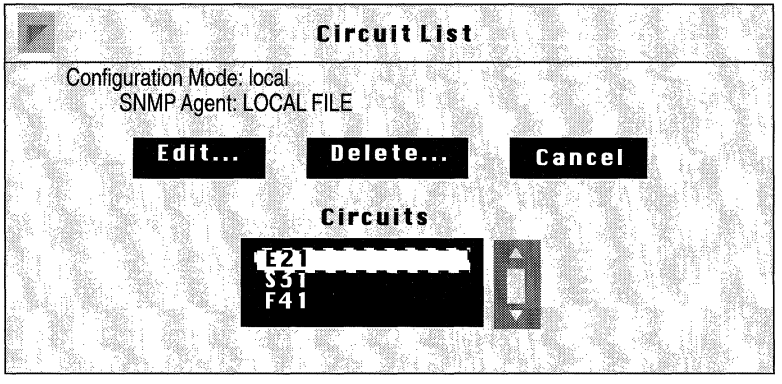


**Figure 3-28. Circuit List Window**

2. Select the circuit from which you want to delete protocols.

   In this example, circuit E21 is selected.
3. Click the Edit button; the Circuit Definition Window appears (see Figure 3-29).

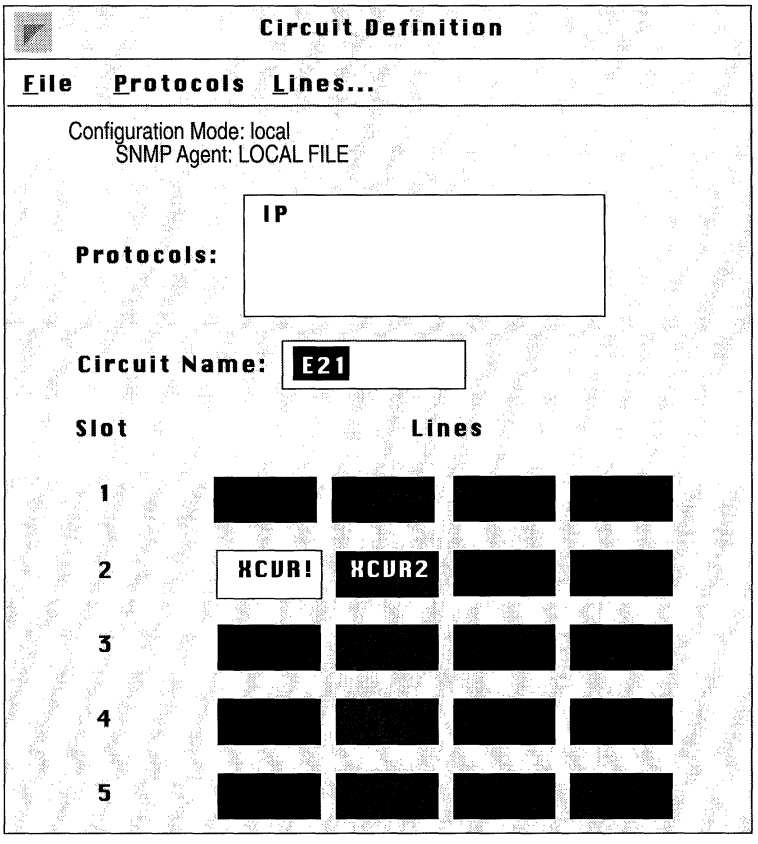


**Figure 3-29. Circuit Definition Window**

4. Select the Protocols/Add/Delete option; the Select Protocols Window appears (previously shown in Figure 3-4).
5. Select the protocols that you want delete from this circuit.
6. Click on the Save button.

   You are returned to the Circuit Definition Window. The protocol(s) you just deleted no longer appear in the Protocols scroll box.

Follow steps 1 through 6 for each circuit from which you want to delete protocols.

## Editing Line Details for a Circuit

To edit line details, start at the Wellfleet Configuration Manager Window and complete the following steps:

1. Select the Circuits/Edit Circuits option.

   The Circuit List Window appears (see Figure 3-30).

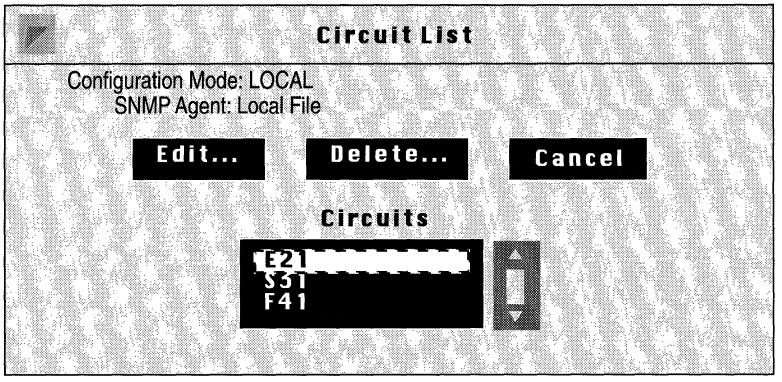


**Figure 3-30. Circuit List Window**

2. Select the circuit you want to edit.

   In this example, circuit E21 is selected.

3. Click the Edit button; the Circuit Definition Window appears (see Figure 3-31).

4. Select the Lines option; the Edit Lines Window appears (see Figure 3-32).

   This window lists the existing lines by slot number and connector.

5. Select the circuit for which you wish to edit line details and click the Edit button.

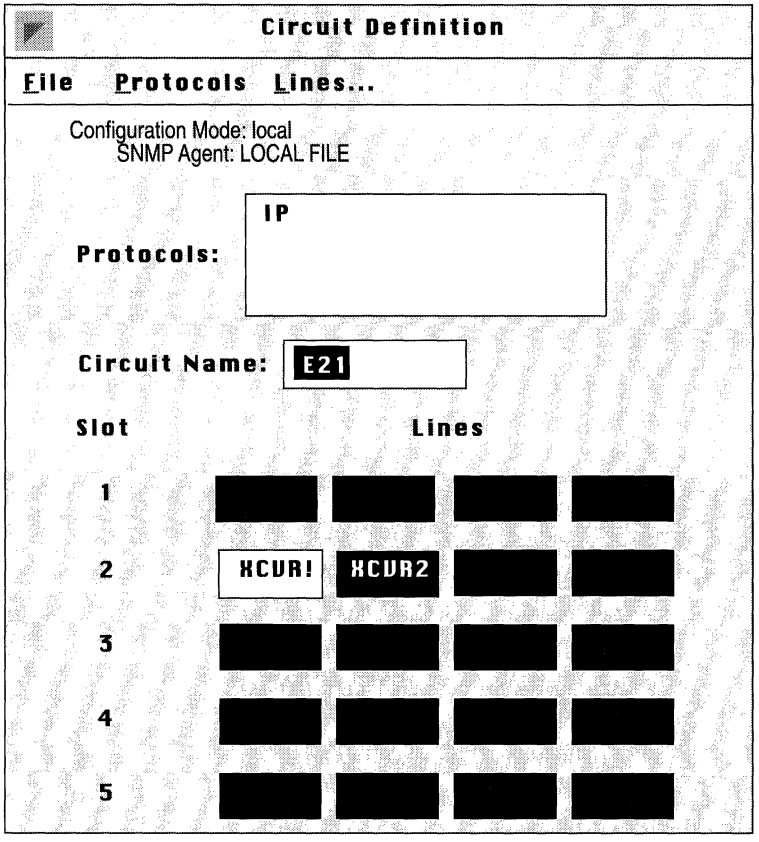


**Figure 3-31. Circuit Definition Window**

In this example, the line connected to XCVR1 in slot 2 is selected. Depending on the type of circuit that you selected in Step 2, the Configuration Manager displays a window that allows you to modify the circuit's line details.

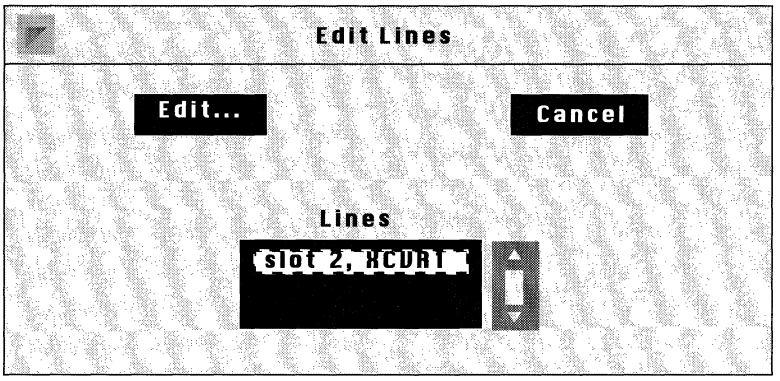


**Figure 3-32. Edit Lines Window**

6. Set each of the parameters as appropriate for the circuit on the line detail window that appears.

   The following sections describe how to edit line detail parameters. Refer only to the section that corresponds to the type of circuit you are editing.

7. Repeat steps 1 through 6 for each circuit that needs line detail modification.

### Editing E1 Line Details

If the circuit you wish to edit is an E1 circuit, the Configuration Manager now displays the E1 Window (see Figure 3-33).

Complete the following steps:

1. Enter or select new values for the E1 service parameters you want to edit.
2. If you wish to edit synchronous line parameters, click the Configure Sync button. Refer to the section *Editing Synchronous Line Details* in this chapter for necessary details.
3. Click the Save button.

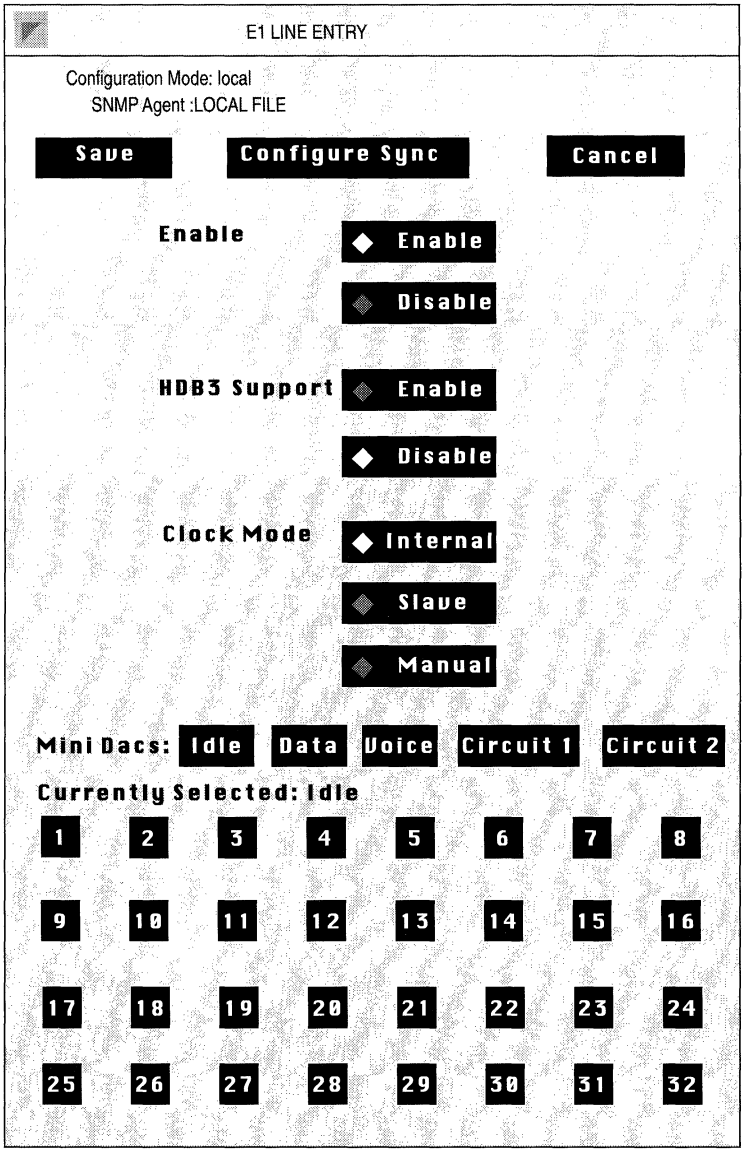


**Figure 3-33. E1 Window**

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or Disables the E1 line. |
| Instructions: | Set to Disable if you want to disable the E1 line. |

| | |
|---|---|
| **Parameter :** | **HDB3S Support** |
| Wellfleet Default: | Disable |
| Options: | Disable/Enable |
| Function: | Enables or disables High Density Bipolar Coding (a mechanism to maintain sufficient ones-density within the E1 data stream). |
| Instructions: | Enable or disable based on the ability of the associated E1 equipment to support HDB3. |

| | |
|---|---|
| **Parameter :** | **Clock Mode** |
| Wellfleet Default: | Internal |
| Options: | Internal/Slave/Manual |
| Function: | Specifies the source of the E1 transmit clock.<br>Internal specifies that the E1 transmit clock is internally generated; Slave specifies that the E1 transmit clock is externally generated (that is, the transmit clock is derived from the incoming data stream); Manual specifies that the clock source is hardware configured (that is, the source of the transmit clock is determined by jumpers on the E1 link module). Refer to the *Maintenance Guide* for information on link module hardware configuration. |
| Instructions: | Select the clocking mode making certain that the associated E1 equipment is configured in a complementary fashion. |

| **Parameter :** | **Mini Dacs** |
|---|---|
| Wellfleet Default: | Idle |
| Options: | Idle/Data/Voice/Circuit 1/Circuit 2 |
| Function: | Assigns each E1 channel to a specific function. |
| Instructions: | Assign each of the 32 E1 channels as required. |

I idles the channel

**Note:** The first E1 channel is reserved for signalling and should be set to Idle.

D assigns the channel to data pass through (E1 connector to E1 connector)

V assigns the channel to voice pass through (E1 connector to E1 connector)

**Note:** Data and/or voice pass through requires that identical channels be assigned to data or voice on both E1 connectors. For example if the first E1 connector allocates channels 2 through 8 to voice pass through and channels 9 through 16 to data pass through, the second E1 connector must also allocate channels 2 through 8 to voice and 9 through 16 to data pass through.

Circuit 1 assigns the channel to the first E1 connector

Circuit 2 assigns the channel to the second E1 connector

**Note:** E1 channels cannot be allocated to both E1 circuits. If, on the first E1 connector, channels 17 through 25 are allocated to circuit 1, the second E1 connector must idle these channels or allocate them to circuit 2.

### Editing Ethernet Line Details

If the circuit you wish to edit is an Ethernet circuit, the Configuration Manager now displays the XCVR Window (see Figure 3-34).

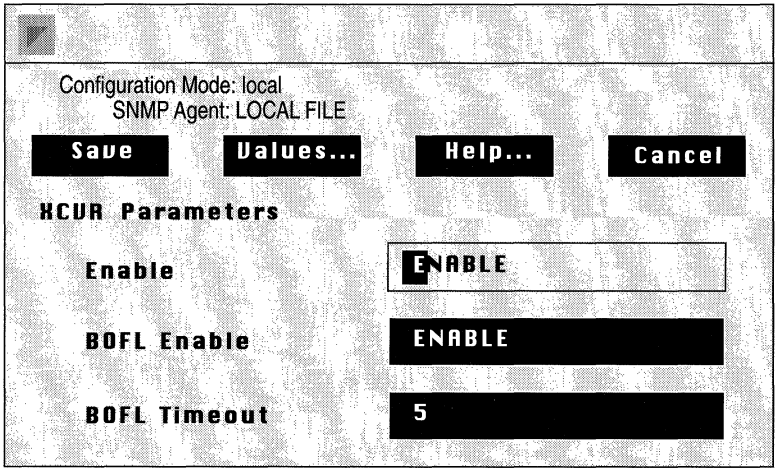


**Figure 3-34. XCVR Window**

Complete the following steps:

1. Enter or select new values for the Ethernet line detail parameters you want to edit.
2. Click the Save button.

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables and disables this Ethernet line. |
| Instructions: | Set this parameter to either Enable or Disable for this line. |

| | |
|---|---|
| **Parameter :** | **BOFL (Breath of Life) Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | When set to Enable, BOFL specifies the sending of Breath of Life messages from this system to everyone on the local network. These messages signify that the Ethernet line is up and functioning normally. |
| Instructions: | Set to Enable or Disable depending on whether you want this system to issue Breath of Life messages over this line. |

**Note:** Wellfleet recommends that BOFL be enabled.

| | |
|---|---|
| **Parameter :** | **BOFL Timeout** |
| Wellfleet Default: | 5 seconds |
| Options: | 1 - 60 seconds |
| Function: | Specifies the maximum amount of time that can elapse between the successful transmission of Breath of Life messages from this system. If this time is exceeded, the Ethernet line will go down, and then come back up.<br>This parameter is valid only if BOFL is set to Enable. |
| Instructions: | Either accept the default BOFL Timeout of 5 seconds, or specify a new value. |

## Editing FDDI Line Details

If the circuit you wish to edit is an FDDI circuit, the Configuration Manager now displays the FDDI Window (see Figure 3-35).

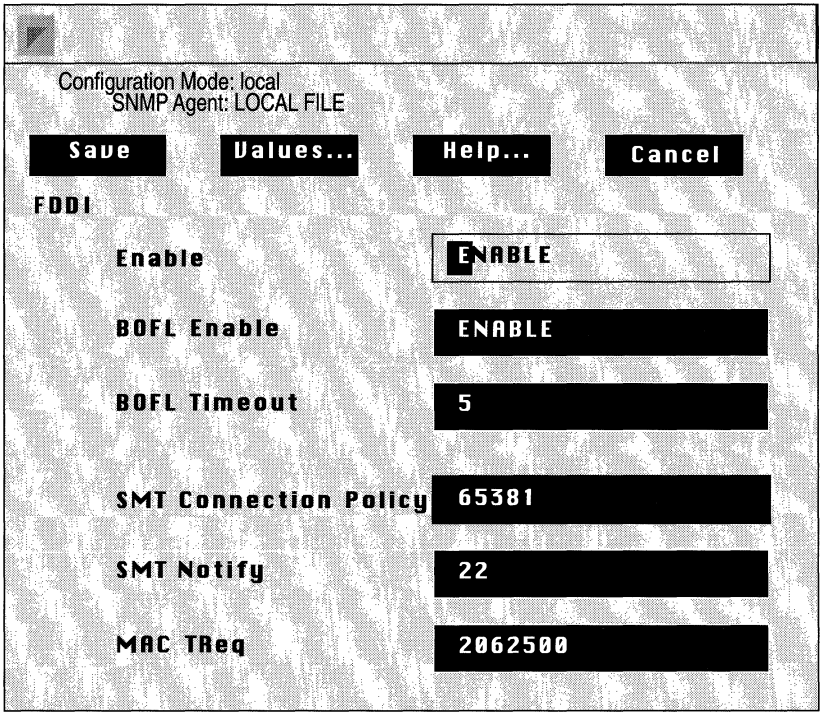


**Figure 3-35. FDDI Window**

Complete the following steps:

1. Enter or select new values for the FDDI line detail parameters you want to edit.
2. Click the Save button.

**Parameter : Enable**

Wellfleet Default: Enable

Options: Enable/Disable

Function: Enables or Disables FDDI on this interface.

Instructions: Set to Disable if you want to disable FDDI on this interface.

**Parameter : BofL Enable**

**Note:** Wellfleet recommends that BofL be enabled.

Wellfleet Default: Enable

Options: Enable/Disable

Function: Enables or disables "breath of life" polling.

When set to Enable, this parameter enables the periodic verification of SMT ring management (to confirm that the FDDI connection is functional).

Instructions: Set to Enable or Disable depending on whether you want the FDDI station to verify the link state.

**Parameter : BofL Timeout**

**Note:** BofL Timeout is valid only if BofL Enable is set to Enable.

Wellfleet Default: 5

Options: 1 - 60 seconds

Function: Specifies the maximum amount of time that can elapse between the successful polling of the link state.

Instructions: Either accept the default value of 5 seconds or specify a new value.

**Parameter :** **SMT Connection Policy**

Wellfleet Default: 65381

Options: 0 - 65535

Function: Provides a decimal equivalent of a sixteen-bit status word that specifies the connection policies requested at the FDDI station.

A station sets the corresponding policy for each of the connection types that it *wishes to reject*. The policy descriptor takes the form "rejectX-Y" where X denotes the Physical Connection (PC) type of the local port, and Y denotes the PC type of the neighbor.

**Note:** The setting of a particular connection does not necessarily mean that the connection will be rejected. The SMT standard requires the *both* sides of the connection must agree to reject, else both sides *must accept* the connection.

X and Y can take the following values:

A -- indicating that the port is a dual attachment station or concentrator that attaches to the primary IN and the secondary OUT when attaching to the dual FDDI ring

B -- indicating that the port is s dual attachment station or concentrator that attaches to the secondary IN and the primary OUT when attaching to the dual FDDI ring

M -- indicating a port in a concentrator that serves as a Master to a connected station or concentrator

S -- indicating a port in a single attachment station or concentrator

The value is a sum that initially takes the value zero, then for each of the connection policies requested by the node, 2 raised to a power is added to the sum. The powers are according to the following table.

**Table 3-1. SMT Connection Policy Values**

| Policy | Power |
|---|---|
| rejectA-A | 0 |
| rejectA-B | 1 |
| rejectS-S | 2 |
| rejectA-M | 3 |
| rejectB-M | 4 |
| rejectA-S | 5 |
| rejectB-S | 6 |
| rejectA-B | 7 |
| rejectB-A | 8 |
| rejectS-A | 9 |
| rejectS-B | 10 |
| rejectS-M | 11 |
| rejectM-A | 12 |
| rejectM-B | 13 |
| rejectM-S | 14 |
| rejectM-M | 15 |

**Note:** The SMT standard requires that Bit 15 (rejectM-M) *must* be set.

Default connection policy is illustrated in the following illustration.

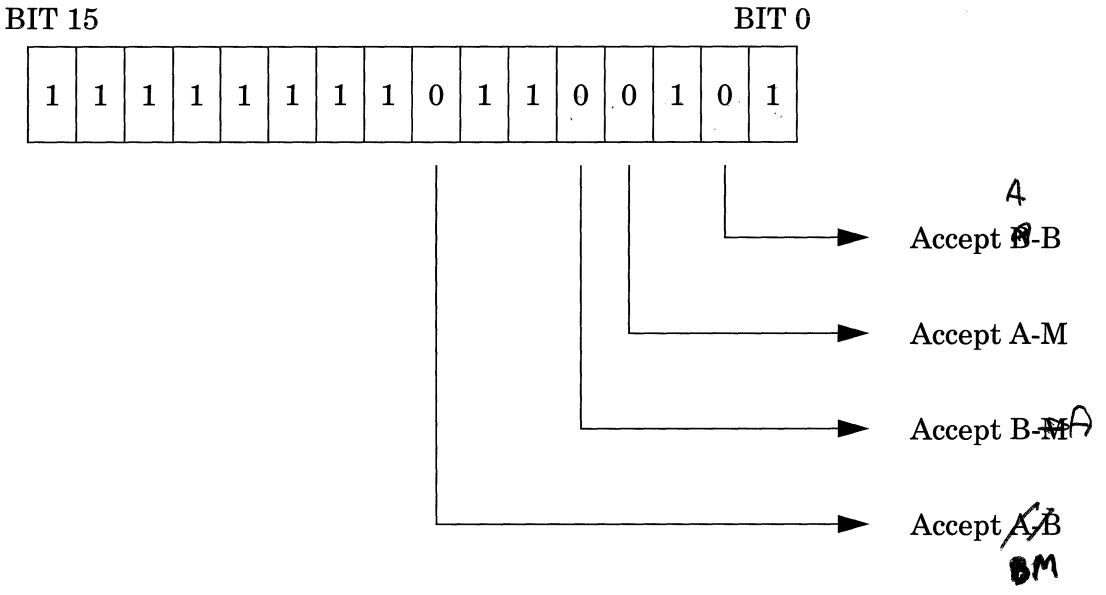


**Figure 3-36. Connection Policy Status Word**

Instructions: Set the status word value to reflect local connection policies.

**Parameter :** **SMT TNotify**

Wellfleet Default: 22 (seconds)

Options: 2 -30 seconds

Function: Specifies the interval between successful iterations of the Neighbor Notification Protocol.

The Neighbor Notification protocol (1) determines the MAC address of the FDDI upstream and downstream neighbor, (2) detects duplicate MAC addresses on the ring, and (3) generates periodic "keep-alive" traffic that verifies the local MAC transmit and receive paths.

Instructions: Accept the default value of 22 seconds or specify a new value.

| | |
|---|---|
| **Parameter:** | **MAC TReq** |
| Wellfleet Default: | 2062500 (octet units - 80 nanoseconds, or 165 milliseconds) |
| Options: | Any value |
| Function: | Specifies the TTRT (target token rotation time) carried in claim frames issued by the FDDI connection. |
| Instructions: | Enter the claim value in 80-nanosecond increments. |

### Editing HSSI Line Details

If the circuit you wish to edit is a HSSI circuit, the Configuration Manager now displays the HSSI Parameters Window (see Figure 3-37).

Complete the following steps:

1. Enter or select new values for the HSSI line detail parameters you want to edit.
2. Click the Save button.

| | |
|---|---|
| **Parameter:** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables and disables this HSSI line. |
| Instructions: | Set this parameter to either Enable or Disable for this line. |

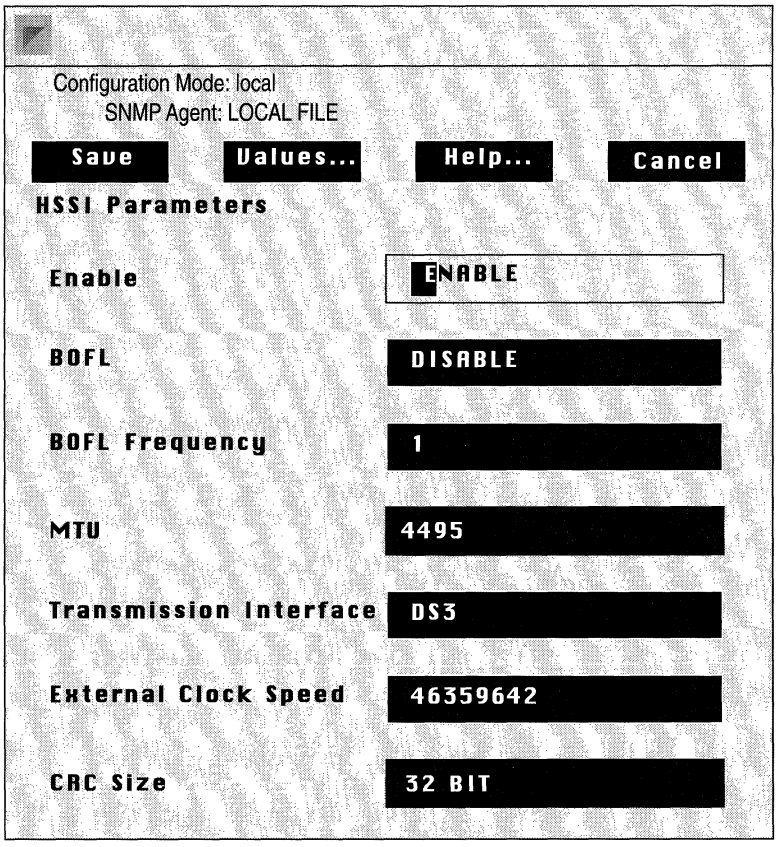


**Figure 3-37. HSSI Parameters Window**

**Note:** Wellfleet recommends that BOFL be enabled for point-to-point connections between Wellfleet peers. However, if such a connection is accomplished through a wide-area transport service such as Frame Relay or SMDS, BofL *must* be disabled.

**Parameter :** **BOFL**

Wellfleet Default: Enable

Options: Enable/Disable

Function: BOFL enables the transmission of proprietary SNAP-encapsulated Breath of Life messages *over a point-to-point connection* between the local BN and a remote peer.

The exchange of BofL messages provides a level of confidence in the point-to-point connection. With BofL enabled, the BN sends periodic *keep-alive* messages to the remote peer.

Instructions: Set to Enable or Disable depending on whether you want to transmit BOFL messages over this HSSI interface. If you enable BofL locally, the remote peer must also be configured to enable BOFL Frequency

**Parameter :** **BofL Frequency**

Wellfleet Default: 1 (second)

Options: 1 - 60 seconds

Function: Specifies the interval in seconds between BOFL transmissions. This parameter is valid only if BOFL is set to Enable.

After sending a BofL message, the BN starts a timer which has a value equal to 5 times the BofL Frequency. If a BofL is not received from the remote peer before the expiration of the timer, the BN takes the HSSI circuit down, and then attempts to restart it.

Instructions: Either accept the default BOFL Frequency of 1 second or specify a new value, making certain that both ends of the point-to-point connection are configured with the same value.

| | |
|---|---|
| **Parameter :** | **MTU (Maximum Transfer Unit)** |
| Wellfleet Default: | 4495 bytes |
| Options: | 3 - 4500 bytes |
| Function: | Specifies the buffer size for the HSSI port and, by extension, the largest frame that can be transmitted or received across the HSSI port. |
| Instructions: | Set this parameter to a value appropriate for your network. |

| | |
|---|---|
| **Parameter :** | **Transmission Interface** |
| Wellfleet Default: | DS3 |
| Options: | DS1/DS3 |
| Function: | When SMDS and/or Frame Relay is configured on this HSSI circuit, specifies whether the local management interface (LMI) employs a DS1 MIB (specified by RFC 1232) or a DS3 MIB (specified by RFC 1233). |
| Instructions: | Select on the basis of the carrier services (DS1 at 1.54M bps or DS3 at 44.736M bps) provided by the attached DCE device. |

**Note:** This parameter is meaningful only when SMDS and/or Frame Relay is configured across the HSSI interface and when LMI is enabled. The HSSI driver provides no support for either the DS1 or DS3 MIB. Rather the external DCE (for example, a DL3200 SMDS CSU/DSU from Digital Link) may provide MIB support. In instances where Frame Relay/SMDS and LMI are configured across the HSSI interface, this parameter specifies the appropriated MIB for use by the LMI.

| | |
|---|---|
| **Parameter :** | **External Clock Speed** |
| Wellfleet Default: | 46359642 (44.736M bps) |
| Options: | Any value within the range 307200 to 52638515 |
| Function: | Specifies the bandwidth provided by the HSSI channel. |
| Instructions: | Set the parameter to the value that equals or approximates the data transmission rate across the HSSI. |

**Note:** The HSSI specification requires that the DCE provide a transmit clock which times data transfer across the DTE/DCE interface. External Clock Speed (currently non-functional) will be used in the future by various routing protocols in performing route selection algorithms.

| | |
|---|---|
| **Parameter :** | **CRC Size** |
| Wellfleet Default: | 32-bit |
| Options: | 16-bit or 32-bit |
| Function: | Specifies an error detection scheme. You may choose either 16-bit (standard CCITT) or 32-bit (extended) to detect errors in the packet. |
| Instructions: | Set this parameter to either 16-bit or 32-bit making certain that the remote end of the HSSI connection is configured for the same value. |

### Editing Synchronous Line Details

If the circuit you wish to edit is an Synchronous circuit, the Configuration Manager now displays the SYNC Window (see Figure 3-38).

This Window contains 17 synchronous line parameters; however, when the window first appears, you will see only the first 10 parameters. To view the other parameters, click in one of the visible parameter boxes, then use the up and down arrow keys to scroll through the window.

Complete the following steps:

1. Enter or select new values for the line detail parameters you want to edit.
2. Click the Save button.

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables and disables this synchronous line. |
| Instructions: | Set this parameter to either Enable or Disable for this line. |

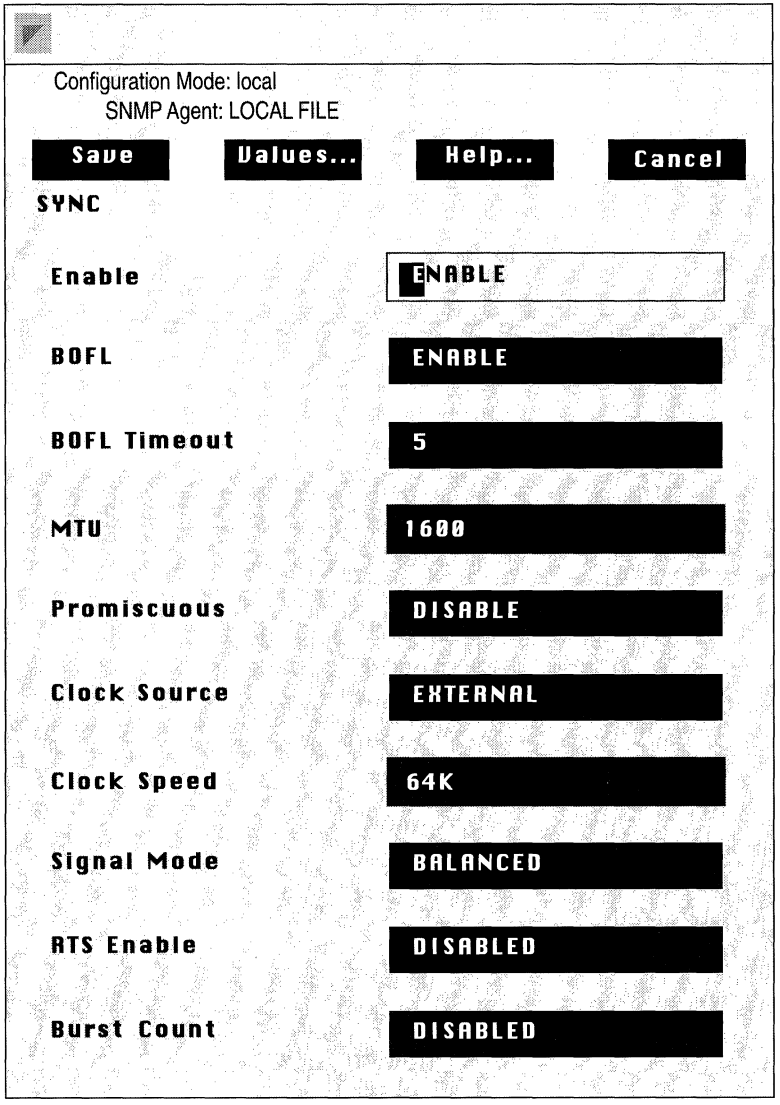


**Figure 3-38. SYNC Window**

**Parameter:** **BOFL**

Wellfleet Default: Enable

Options: Enable/Disable

Function: BOFL enables the transmission of proprietary SNAP-encapsulated Breath of Life messages *over a point-to-point connection* between the local BN and a remote peer.

The exchange of BOFL messages provides a level of confidence in the point-to-point connection. With BOFL enabled, the BN sends periodic *keep-alive* messages to the remote peer.

Instructions: Set to Enable or Disable depending on whether you want to transmit BOFL messages over this synchronous interface. If you enable BOFL locally, the remote peer must also be configured to enable BOFL.

**Note:** Wellfleet recommends that BOFL be enabled for point-to-point connections between Wellfleet peers. However, if such a connection is accomplished through a wide-area transport service such as Frame Relay or SMDS, BOFL *must* be disabled.

**Parameter:** **BOFL Timeout**

Wellfleet Default: 5 seconds

Options: 1 - 60 seconds

Function: Specifies the maximum amount of time that can elapse between the successful transmission of Breath of Life messages from this system. If this time is exceeded, the synchronous line will go down, and then come back up.

This parameter is valid only if BOFL is set to Enable.

Instructions: Either accept the default BOFL Timeout of 5 seconds, or specify a new value.

| | |
|---|---|
| **Parameter :** | **MTU (Maximum Transfer Unit)** |
| Wellfleet Default: | 1600 bytes |
| Options: | 3 - 4500 bytes |
| Function: | Specifies the largest amount of data that can be transferred across this network in one frame. |
| Instructions: | Set this parameter to a value appropriate for your network. |

| | |
|---|---|
| **Parameter :** | **Promiscuous** |
| Wellfleet Default: | Disable |
| Options: | Enable/Disable |
| Function: | Specifies whether address filtering based on the local and remote address is enabled. If Promiscuous is set to Enable, all frames are received. If Promiscuous is set to Disable, only frames destined for this local address are received. |
| Instructions: | Set this parameter to Enable or Disable. |

| | |
|---|---|
| **Parameter :** | **Clock Source** |
| Wellfleet Default: | External |
| Options: | External/Internal |
| Function: | Specifies the origin of the synchronous timing signals. If Clock Source is set to Internal, this router supplies the required timing signals. If Clock Source is set to External, an external network device supplies the required timing signals. In most cases, this parameter should be set to External. |
| Instructions: | Set this parameter to either Internal or External, as appropriate for your network. |

| | |
|---|---|
| **Parameter :** | **Clock Speed** |
| Wellfleet Default: | 64K |
| Options: | 1200B, 2400B, 4800B, 7200B, 9600B, 19200B, 32000B, 38400B, 56K, 64K, 125K, 230K, 420K, 625K, 833K, 1.25Mb, 2.5Mb, or 5Mb |
| Function: | Specifies the speed of the internal clock. Clock Speed is valid only when Clock Source is set to Internal. |
| Instructions: | Set the parameter to the value that equals the data transmission rate across the synchronous line. |

| | |
|---|---|
| **Parameter :** | **Signal Mode** |
| Wellfleet Default: | Balanced |
| Options: | Balanced/Unbalanced |
| Function: | Specifies balanced or unbalanced transmission. Balanced transmission uses two conductors to carry signals; unbalanced uses one conductor to carry a signal, with a ground providing the return path. |
| Instructions: | Set this parameter to either Balanced or Unbalanced, based on the signalling mode of the connected device. |

| | |
|---|---|
| **Parameter :** | **RTS Enable** |
| Wellfleet Default: | Disable |
| Options: | Enable/Disable |
| Function: | Enables or disables the detection of RTS signals on this interface. |
| Instructions: | Set this parameter to Enable if the connected device (for example, a modem) uses RTS/CTS flow control. |

| | |
|---|---|
| **Parameter:** | **Burst Count** |
| Wellfleet Default: | Disable |
| Options: | Enable/Disable |
| Function: | Specifies single or multiple DMA burst cycles. If Burst Count is set to Enable, the chip performs 8-word bursts. If Burst Count is set to Disable, single word cycles are performed. |
| Instructions: | Set Burst Count to either Enable or Disable. |

**Note:** The Burst Count parameter should be enabled for the following modules: DSDE2, SSE, DSE, and QSYNC.

| | |
|---|---|
| **Parameter:** | **Service** |
| Wellfleet Default: | LLC1 |
| Options: | Transparent/LLC1 |
| Function: | Specifies the link-level protocol for this circuit. If Service is set to Transparent, then raw HDLC (high-level data-link control) mode is in effect. LLC1 specifies connectionless datagram service; it prefixes the HDLC address and control fields to the frame. |
| Instructions: | Set this parameter as appropriate for this circuit. |

| | |
|---|---|
| **Parameter:** | **Minimum Frame Spacing** |
| Wellfleet Default: | 1 (flag) |
| Options: | 1 - 32 flags |
| Function: | Specifies the number of flags that are transmitted between adjacent frames. |
| Instructions: | Set this parameter to the appropriate number of flags. |

The following two parameters have to do with the point-to-point address. According to convention, one end of a point-to-point circuit is designated DCE and is assigned an address of 01; the other end of the circuit is designated DTE and is assigned an address of 03.

Conventional addressing; however, is inadequate in the case of multiple communication channels enabled by a common satellite link (see Figure 3-39). As shown, a common satellite relay-link provides a virtual point-to-point link between routers A and X, B and Y, and C and Z.

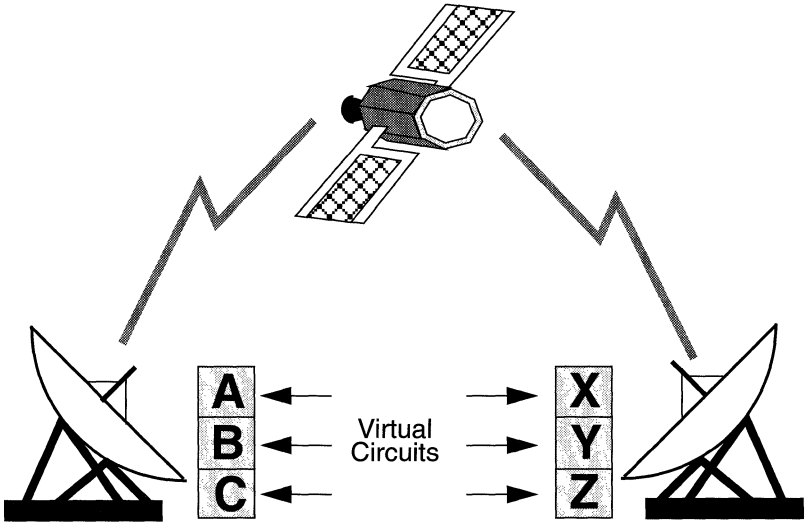


**Figure 3-39. Satellite Broadcast (Sample Topology)**

The worst case scenario consists of routers A, B and C being designated as DCE (address = 01), and routers X, Y, and Z being designated DTE (address=03). If A transmits a frame across the virtual point-to-point circuit to X, the satellite broadcast is monitored not only by X (the intended recipient), but also by Y and Z. Because X, Y and Z all perceive a properly addressed frame, all three accept delivery and attempt to process the frame contents with unpredictable results.

Explicit addressing avoids such confusion by enabling the assignment of unique addresses to each end of a point-to-point circuit. If you set Local Address and Remote Address to Explicit, you must then also specify an address for each end of the circuit. Make sure to reverse the local and remote address when you configure the other end of the point-to-point circuit.

| **Parameter :** | **Local Address** |
|---|---|
| Wellfleet Default: | Explicit |
| Options: | DTE, DCE, and Explicit |
| Function: | Specifies the 1-byte value, which is used in the address field of the HDLC packet. It may be extended to two octets if the Extended Address is set to Enable. |
| Instructions: | Set this parameter to DTE, DCE, or Explicit.<br><br>If you choose Explicit, a parameter box will automatically appear below this one. Enter a unique decimal value between 00 and 99 (avoiding the conventional address values of 01 or 03), that you wish to assign to the local end of this point-to-point circuit. |

**Note:** Make certain to reverse local and remote address values when you configure the device at the other end of the point-to-point circuit.

| | |
|---|---|
| **Parameter :** | **Remote Address** |
| Wellfleet Default: | Explicit |
| Options: | DTE, DCE, and Explicit |
| Function: | Specifies the 1-byte value, which is used in the address field of the HDLC packet. It may be extended to two octets if the Extended Address is set to Enable. |
| Instructions: | Set this parameter to DTE, DCE, or Explicit.<br><br>If you choose Explicit, a parameter box will automatically appear below this one. Enter a unique decimal value between 00 and 99 (avoiding the conventional address values of 01 or 03), that you wish to assign to the remote end of this point-to-point circuit. |

**Note:** Make certain to reverse local and remote address values when you configure the device at the other end of the point-to-point circuit.

| | |
|---|---|
| **Parameter :** | **WAN Protocol** |
| Wellfleet Default: | Standard |
| Options: | Standard/Frame Relay/SMDS |
| Function: | Indicates the WAN Protocol for this interface.<br><br>If the synchronous line supports either Frame Relay or SMDS service, the BN automatically provides the appropriate option which need not be changed. |
| Instructions: | Set to match the wide area protocol in use across the synchronous line. |

| | |
|---|---|
| **Parameter:** | **CRC Size** |
| Wellfleet Default: | 16-bit |
| Options: | 16-bit or 32-bit |
| Function: | Specifies an error detection scheme. You may chose either 16-bit (standard) or 32-bit (extended) frame check sequence (FCS) to detect errors in the packet. |
| Instructions: | Set this parameter to either 16 bit or 32 bit. |

| | |
|---|---|
| **Parameter:** | **Sync Media Type** |
| Wellfleet Default: | 1 |
| Options: | 1 (default), 2 (T1), or 3(E1) |
| Function: | Specifies the media type. |
| Instructions: | Set this parameter 2 for T1 lines or 3 for E1 lines. Otherwise accept the default value of 1. |

### Editing T1 Line Details

If the circuit you wish to edit is an T1 circuit, the Configuration Manager now displays the T1 Window (see Figure 3-40).

Complete the following steps:

1. Enter or select new values for the T1 service parameters you want to edit.
2. If you wish to edit synchronous line parameters, click the Configure Sync button. Refer to the section *Editing Synchronous Line Details* in this chapter for necessary details.
3. Click the Save button.

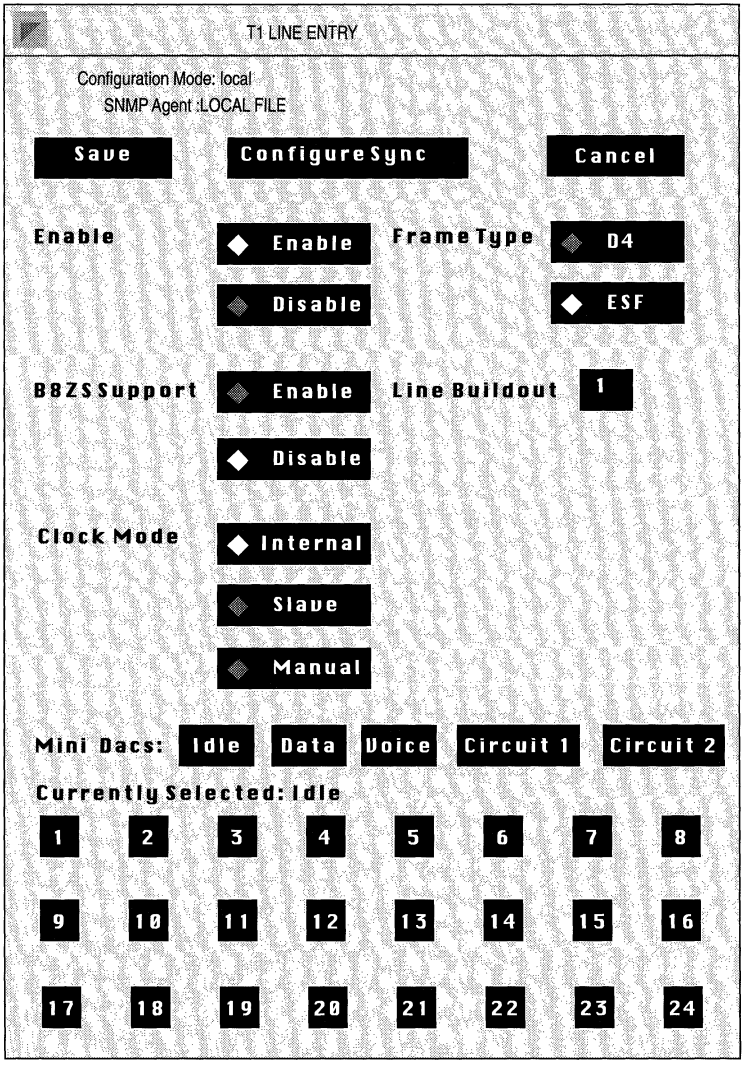


**Figure 3-40. T1 Window**

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or Disables the T1 line. |
| Instructions: | Set to Disable if you want to disable the T1 line. |

| | |
|---|---|
| **Parameter :** | **Frame Type** |
| Wellfleet Default: | ESF |
| Options: | ESF/D4 |
| Function: | Selects either ESF (extended super frame) or D4 framing formats. D4 transmits super frames consisting of 12 individual super frames. ESF, in contrast, transmits super frames consisting of 24 individual D4 frames and provides enhanced signalling and synchronization. |
| Instructions: | Select ESF or D4 based on the frame format required by the associated T1 equipment. |

| | |
|---|---|
| **Parameter :** | **B8ZS Support** |
| Wellfleet Default: | Disable |
| Options: | Disable/Enable |
| Function: | Enables or disables Binary 8 Zeros Suppression (a mechanism to maintain sufficient ones-density within the T1 data stream). |
| Instructions: | Enable or disable based on the ability of the associated T1 equipment to support B8ZS. |

| | |
|---|---|
| **Parameter :** | **Line Buildout** |
| Wellfleet Default: | 1 (foot) |
| Options: | 1 to 655 (feet) |
| Function: | Conditions BN signals to mitigate attenuation. |
| Instructions: | As signal attenuation correlates with the physical length of the T1 line, enter the approximate length of the cable connecting the BN and the associated T1 equipment. |

| | |
|---|---|
| **Parameter :** | **Clock Mode** |
| Wellfleet Default: | Internal |
| Options: | Internal/Slave/Manual |
| Function: | Specifies the source of the T1 transmit clock.<br>Internal specifies that the T1 transmit clock is internally generated; Slave specifies that the T1 transmit clock is externally generated (that is, the transmit clock is derived from the incoming data stream); Manual specifies that the clock source is hardware configured (that is, the source of the transmit clock is determined by jumpers on the T1 link module). Refer to the *Maintenance Guide* for information on link module hardware configuration. |
| Instructions: | Select the clocking mode making certain that the associated T1 equipment is configured in a complementary fashion. |

**Parameter :** **Mini Dacs**

Wellfleet Default: Idle

Options: Idle/Data/Voice/Circuit 1/Circuit 2

Function: Assigns each T1 channel to a specific function.

Instructions: Assign each of the 24 DS1 channels as required.

- I idles the channel
- D assigns the channel to data pass through (T1 connector to T1 connector)
- V assigns the channel to voice pass through (T1 connector to T1 connector)

**Note:** Data and/or voice pass through requires that identical channels be assigned to data or voice on both T1 connectors. For example if the first T1 connector allocates channels 1 through 8 to voice pass through and channels 9 through 16 to data pass through, the second T1 connector must also allocate channels 2 through 8 to voice and 9 through 16 to data pass through.

- Circuit 1 assigns the channel to the first T1 circuit
- Circuit 2 assigns the channel to the second T1 circuit

**Note:** T1 channels cannot be allocated to both T1 circuits. If, on the first T1 connector, channels 17 through 24 are allocated to circuit 1, the second T1 connector must idle these channels or allocate them to circuit 2.

### Editing Token Ring Line Details

If the circuit you wish to edit is a Token Ring circuit, the Configuration Manager displays the Token Ring Window (Figure 3-41).

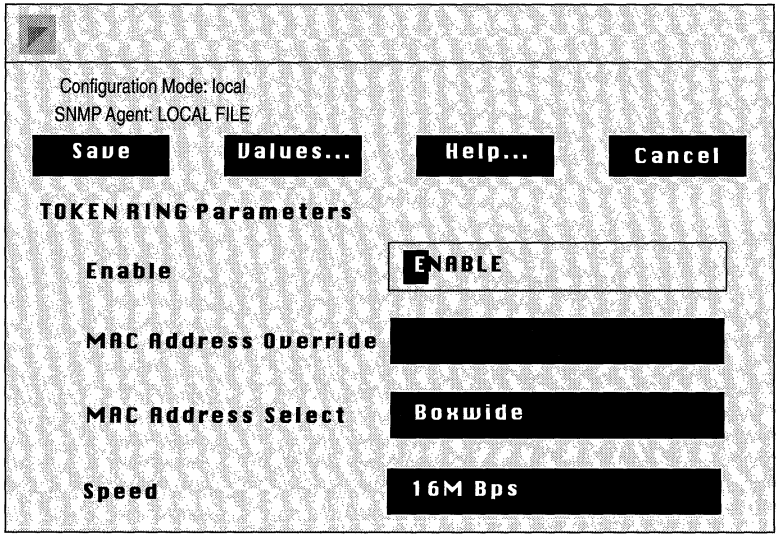


**Figure 3-41. Token Ring Window**

Complete the following steps:

1. Enter or select new values for the Token Ring parameters you want to edit.
2. Click the Save button.

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or Disables the Token Ring circuit. |
| Instructions: | Set to Disable if you want to disable the Token Ring circuit. |

| | |
|---|---|
| **Parameter :** | **MAC Address Override** |
| Wellfleet Default: | None |
| Options: | Any valid 48-bit MAC level address. |
| Function: | Enables the assignment of a user-specified MAC address. |
| Instructions: | If you want to specify a MAC level address (for example, to avoid host number conflicts on a directly connected IPX or XNS network), enter the 48-bit MAC address. |

**Note:** If you enter a MAC address at MAC Address Override, you must set MAC Address Select to CNFG.

If you want the BN to generate a MAC level address for this Token Ring interface, ignore this parameter.

| | |
|---|---|
| **Parameter :** | **MAC Address Select** |
| Wellfleet Default: | BOXWIDE |
| Options: | BOXWIDE, PROM, CNFG |
| Function: | Determines the source of the MAC address. |
| Instructions: | Enter BOXWIDE if you want the Configuration Manager to generate a MAC address automatically from the BN's serial number.<br><br>Enter PROM if you want the Configuration Manager to generate a host number automatically from programmable read-only memory on the Token Ring link module.<br><br>Enter CNFG if you explicitly assigned a MAC address at the MAC Address Override parameter. |

| | |
|---|---|
| **Parameter :** | **Speed** |
| Wellfleet Default: | 16M Bps |
| Options: | 16M Bps/4M Bps |
| Function: | Specifies the speed of the Token Ring media. |
| Instructions: | Enter the ring speed. |

**Note:** If you select 16M Bps, the BN enables the Early Token Release protocol which is used extensively on 16M Bps media. In the unlikely event that you want to disable ETR over 16M Bps Token media, you can do so through the TI.

# Editing Protocol-Specific Parameters

Once you enable a protocol on a circuit, you can edit any protocol-specific parameters. To do this, start at the Wellfleet Configuration Manager Window and complete the following steps:

1. Select the Circuits/Edit Circuits option.

   The Circuit List Window appears (Figure 3-42).

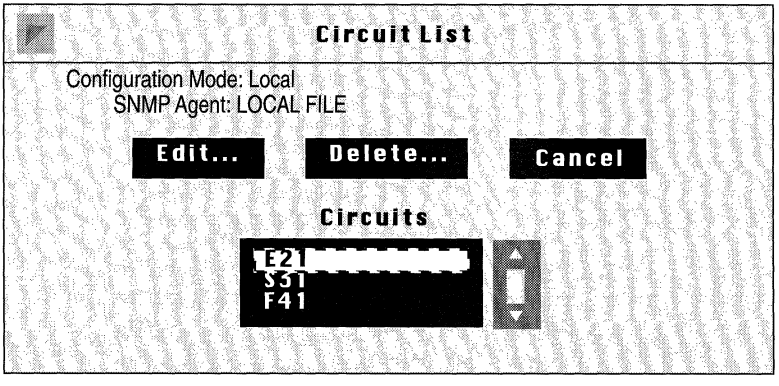


**Figure 3-42. Circuit List Window**

2. Select the circuit for which you want to edit protocol-specific parameters.

   In this example, circuit E21 is selected.

3. Click the Edit button; the Circuit Definition Window appears (Figure 3-43).

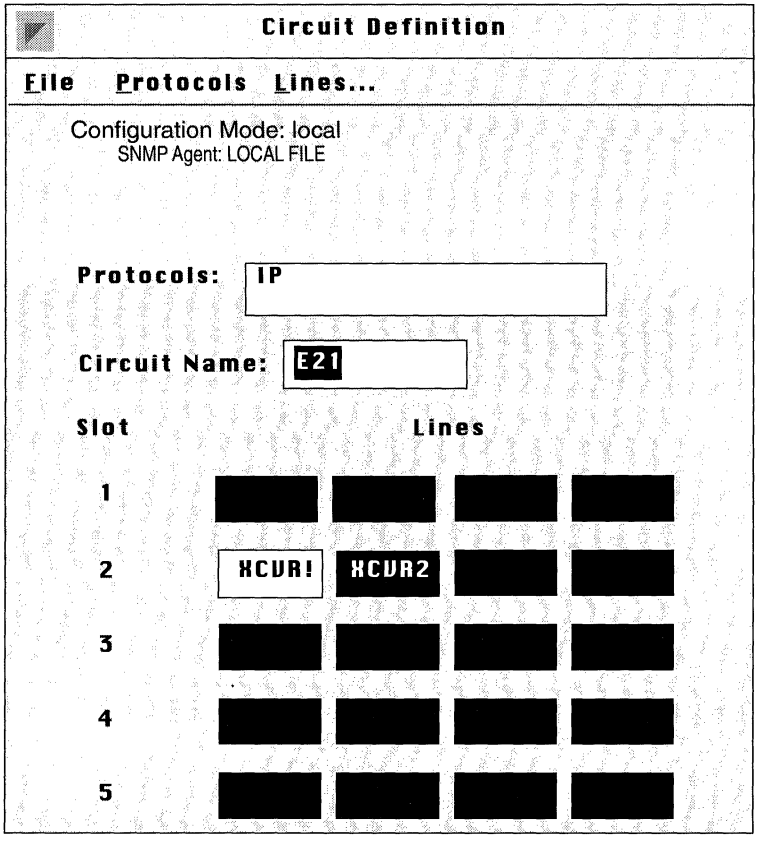


**Figure 3-43. Circuit Definition Window**

4. Select the Protocols menu.

   Depending on the protocols enabled for the circuit, certain edit-protocol options will be available.

   Select the protocol you wish to edit, and then refer to the corresponding chapter of this guide for protocol-specific information.

# Chapter 4

## Configuring Frame Relay

## List of Figures

# Configuring Frame Relay

## About This Chapter

This chapter describes how to configure Frame Relay. Users who are already familiar with Frame Relay may wish to proceed directly to the section entitled *Editing Frame Relay Parameters*. Users desiring some background material may find it useful to read the sections entitled *Frame Relay Overview* and *Frame Relay Implementation Note* in addition to consulting some of the documents listed in the *Frame Relay Bibliography.*

## Frame Relay Overview

Frame Relay is a public connection-oriented packet service that provides facilities for the interconnection of LANs. Based on a subset (the so called "core aspects") of ISDN, Frame Relay eliminates all processing at the network level and greatly restricts data-link layer processing. Such simplified processing is made possible by the availability of virtually error-free physical connectivity and the presence of intelligent protocols at the end-user which can detect and retransmit faulty packets.

Frame Relay encapsulates the protocol payload within the simple header structure (Figure 4-1). As shown in Figure 4-1 the standard Frame Relay Header allocates ten bits for a Data Link Connection Identifier (DLCI). The DLCI is a Frame Relay Permanent Virtual Circuit (PVC) number that corresponds with a particular destination.

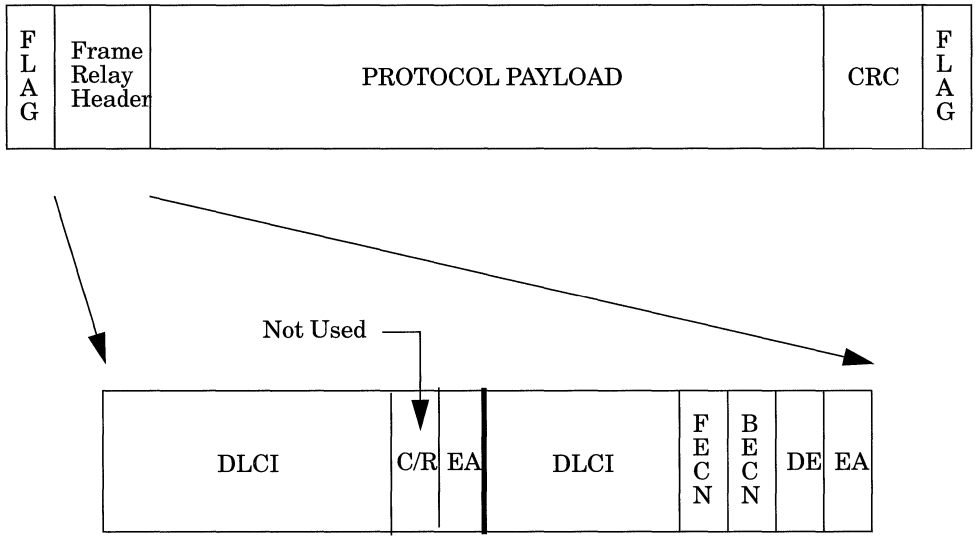


**Figure 4-1. Frame Relay Header Format**

The provision of 10-bits in the standard Frame Relay header provides a range of DLCI values from 0 to 1023. Not all DLCIs, however, are available for the identification of PVCs; some numbers are reserved for network use. The address extension (EA) bits are used to signify the presence of optional address extension bytes, which in theory provides support for upwards of 8,000,000 DLCIs. Regardless of the length of the address field, many Frame Relay services support both unicast and multicast addressing. Under the most common implementation of multicasting the Frame Relay network maps multiple addresses (an address group) to a single DLCI and delivers copies of a single Frame Relay frame to each member of the group. As the frame passes through the switch fabric, the network manipulates the DLCI so that the frame recipient receives a DLCI indicating the actual packet source, not the multicast address.

The forward explicit congestion notification (FECN) and backward explicit congestion notification (BECN) bits are used by Frame Relay

network to inform message originators and recipients of congestion within the network fabric. The BN maintains a count of such bits to allow user visibility into possible congestion conditions. Another network congestion tool, the discard eligibility indicator (DE) bit, signifies that the frame may be dropped by the network in the light of congestion.

As the DLCI identifies the destination, the Frame Relay network performs only a quick three-step process as it processes incoming frames.

1. Verify the CRC; if an error is indicated the frame is dropped.
2. Perform a table lookup for the DLCI; if the DLCI is invalid or unknown, the frame is dropped.
3. If the frame is valid, forward it towards its destination.

At present, all Frame Relay virtual circuits are permanent; the ends of the connection are defined by the network operator. Unlike X.25 virtual circuits, for instance, they are not established on a per-call basis. They are either manually configured or administratively configured by a network management system. As PVCs generally define a LAN-to-LAN connection, new PVCs are needed only when a new LAN is connected to the internet.

The interface between a router and the Frame Relay network is usually, although not always, governed by a management protocol, such as ANSI T1 617D, CCITT Annex A, or LMI. Either protocol enables the router and the Frame Relay carrier to exchange information about the interface state, the PVCs associated with the interface, and the status of each PVC.

# Frame Relay Bibliography

The following documents provide technical detail on Frame Relay protocol design and implementation.

American National Standards Institute. T1.617-1991. *Integrated Services Digital Network (ISDN) - Digital Subscriber Signalling System No 1 (DSS1) - Signalling Specification for Frame Relay Bearer Service*. June 1991.

American National Standards Institute. T1.617 Annex D-1991. *Additional Procedures for Permanent Virtual Connections (PVCs) Using Unnumbered Information Frames*. June 1991.

American National Standards Institute. T1.618-1991. *Integrated Services Digital Network (ISDN) - Core Aspects of Frame Protocol with Frame Relay Bearer Service*. June 1991.

Bradley, T. and Brown, C. *Inverse Address Resolution Protocol*. RFC 1293, Network Information Center (NIC), SRI International, Menlo Park, CA, January 1992.

Bradley, T., Brown, C. and Malis, A. *Multiprotocol Interconnect over Frame Relay*. RFC 1294, Network Information Center (NIC), SRI International, Menlo Park, CA, January 1992.

Digital Equipment Corporation et al. *T1S1-Standards based Frame Relay Specification with Common Enhancements*. Document Number 001-208966, Revision 1.0, September 1990.

The following publications provide a less technical introduction to Frame Relay service.

Davidson, R. and Muller, N. *The Guide to SONET: Planning, Installing & Maintaining Broadband Networks*. Telecom Library, Inc., 1991.

Goldstein, F. *ISDN in Perspective*. Addison-Wesley Publishing Company, 1992.

Jennings, E., Jones, T. and Rehbehn, K. *The Buyer's Guide to Frame Relay Networking*. Netrix Corporation.

# Frame Relay Implementation Note

Frame Relay service operates over synchronous, T1, E1 or HSSI media to provide support for the following protocols: Bridge (to include the Spanning Tree), IP (to include ARP and Inverse ARP support), DECnet, IPX, XNS, and Source Routing (tunneled in IP).

**Note:** AppleTalk and VINES are not supported by Frame Relay.

Frame Relay affords the user three network access modes based upon two logical views of the Frame Relay switched fabric. Each of these three access modes is described in the following sections.

## Frame Relay as a LAN — Group Access

Frame Relay LAN/Group access (the service default) is shown in Figure 4-2.

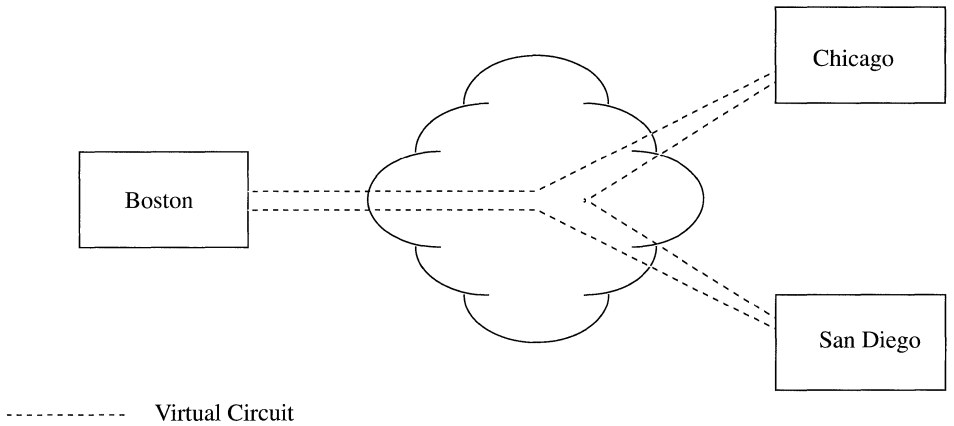


**Figure 4-2. Frame Relay — LAN/Group Access**

Group access treats the Frame Relay network interface as a single LAN-like "pipe" into the switched network. Each network interface has a single associated protocol address and supports a group of undifferentiated Permanent Virtual Circuits (PVCs). Group access

works best in a fully meshed PVC topology and provides the most efficient use of network addressing. Combined with a network management protocol such as ANSI T1 617D, Group access provides for the dynamic configuration of new PVCs as needed. Group access supports all protocols except for AppleTalk, the Bridge, and VINES. Group access is also the easiest service to configure in that you only need to define and associate a protocol with the Frame Relay interface. No explicit configuration of Frame Relay parameters or PVCs is required.

## Frame Relay as Point-to-Point Connections — Direct Access

Frame Relay Direct access is shown in Figure 4-3.

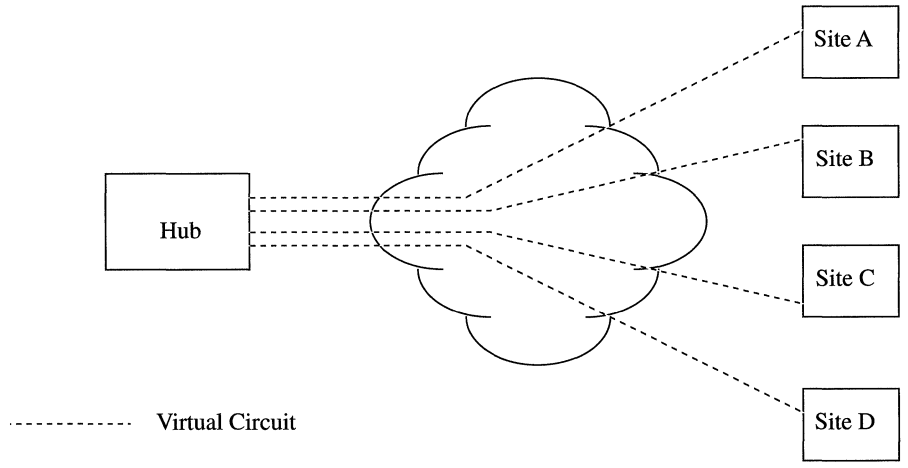


**Figure 4-3. Frame Relay — Direct Access**

Direct access views Frame Relay as a series of point-to-point connections; with Direct access each PVC is treated and configured as a network interface. Thus, in contrast to LAN/Group access, Direct access interfaces are multi-addressed in that they consist of a group of individually addressed PVCs.

Direct access allows the user to dedicate a PVC to a particular protocol, but at the cost of some configuration overhead. All direct access PVCs must be individually configured, and then associated with a particular protocol. Combined with Direct access service, a network management protocol such as ANSI T1 617D activates and inactivates PVCs as needed. Direct access supports all protocols except for AppleTalk and VINES.

### Frame Relay as Bridge and Router — Hybrid Access

Hybrid access provides for the configuration of both bridging and routing on a single Frame Relay PVC. Routing protocols view the PVC in terms of LAN/Group access, while the Bridge treats the PVC as providing Direct access. Hybrid access supports all protocols except for AppleTalk and VINES.

## Editing Frame Relay Parameters

Once you have configured a circuit to support Frame Relay , you can use the Configuration Manager to tailor the default Frame Relay service to meet your network needs.

You begin from the Wellfleet Configuration Manager Window (Figure 4-4).

### Editing Frame Relay Interface Parameters

You edit Frame Relay interface parameters from the Frame Relay Interface Parameters Window. At the Wellfleet Configuration Manager Window, select the Protocols/Frame Relay/Interfaces option to display the Frame Relay Interfaces Window (Figure 4-5). Next, select the interface you wish to edit, then click on the Edit button to display the Frame Relay DLCMI (Data Link Connection Management Interface) Parameters Window (Figure 4-6) for that interface.

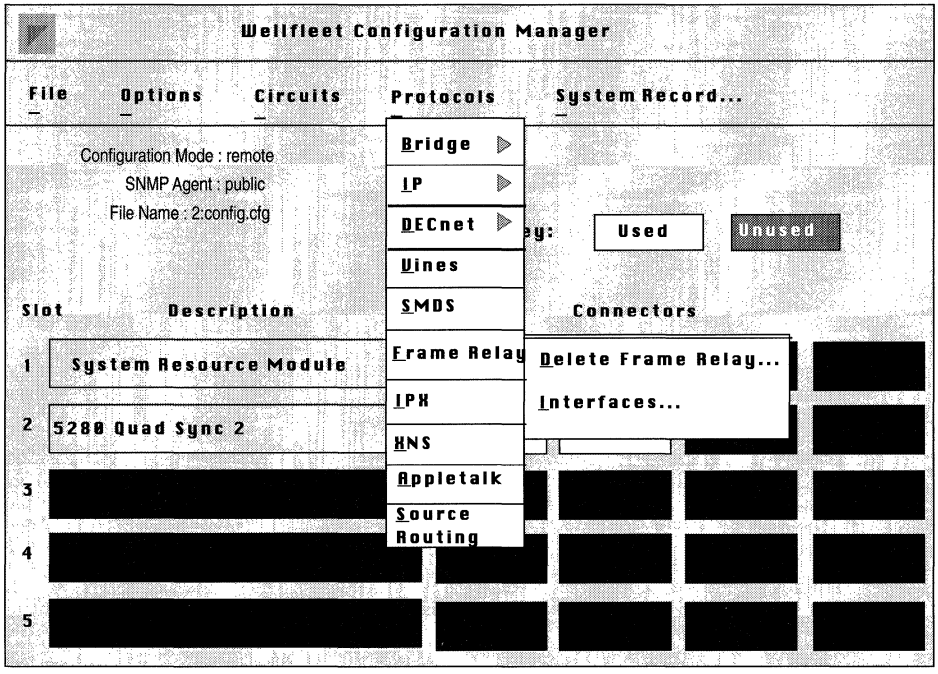


**Figure 4-4. Wellfleet Configuration Manager Window**

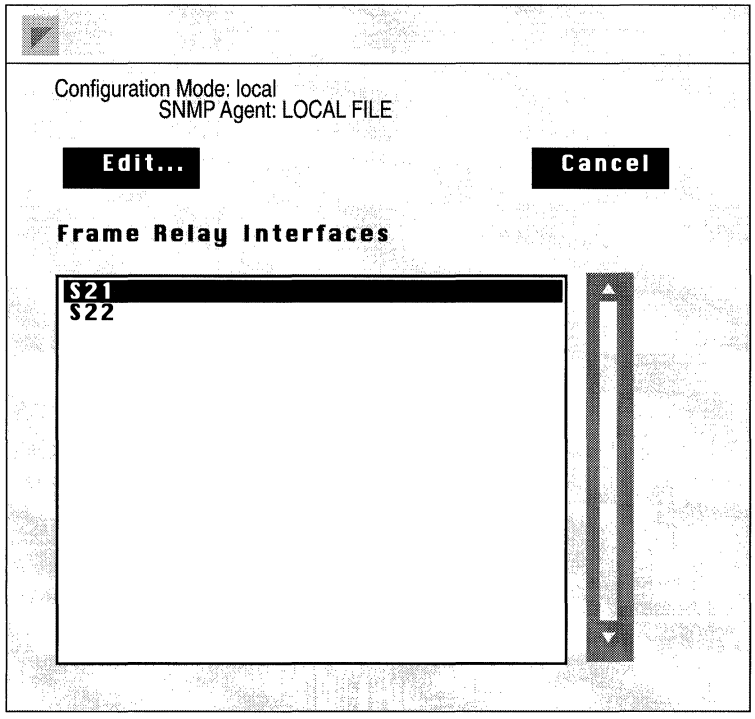


**Figure 4-5. Frame Relay Interfaces Window**

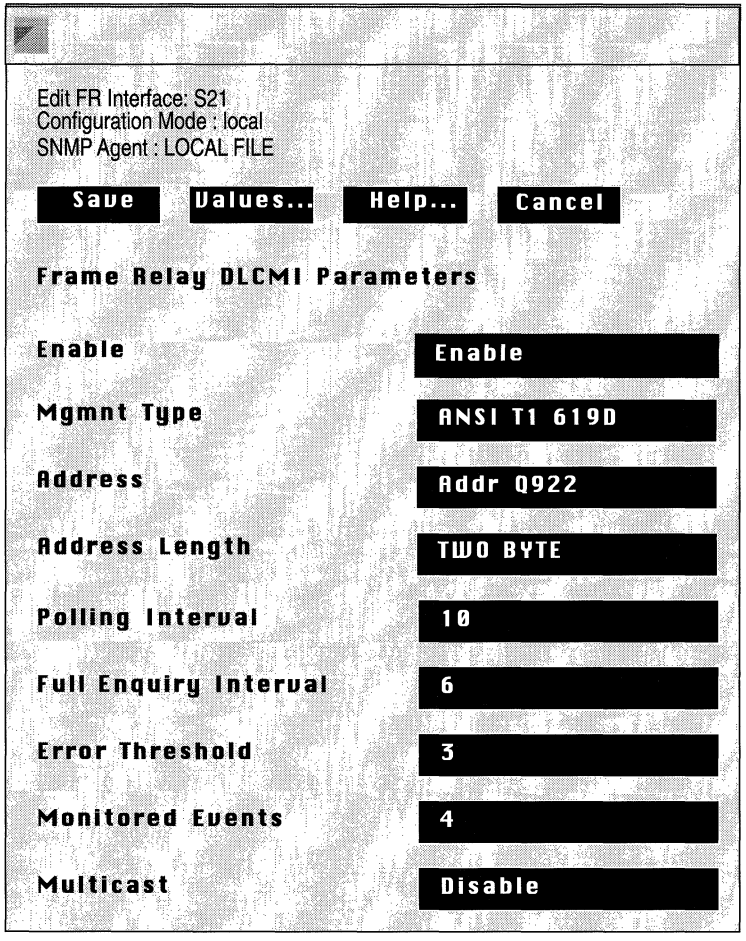


**Figure 4-6. Frame Relay DLCMI Parameters Window**

This section provides the information you need to edit each parameter in the Frame Relay DLCMI Parameters Window. Refer to this information to edit parameters you wish to change. When you are done, click the Save button to exit the window and save your changes.

| | |
|---|---|
| **Parameter:** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or Disables Frame Relay service on this interface. |
| Instructions: | Set to Disable if you want to disable Frame Relay service on this interface. |

| | |
|---|---|
| **Parameter:** | **Mgmnt Type** |
| Wellfleet Default: | ANSI T1 617D |
| Options: | ANSI T1 617D/DLCMI None/Rev 1 LMI/ CCITT Annex A |
| Function: | Selects a management protocol specifying the interface between the BN and the Frame Relay network. Each of the three supported protocols provides notification procedures for the addition or deletion of PVCs, indication of PVC status, and verification of link integrity.<br><br>ANSI TI 617D provides management services as specified in Annex D to ANSI standard TI617-1991.<br><br>Rev 1 LMI (local management interface) provides a set of vendor-generated enhancements to the original Annex D procedures.<br><br>CCITT Annex A provides management services as specified by the CCITT.<br><br>DLCMI None provides no management interface between the BN and the Frame Relay network. In the absence of management support, all PVCs must be configured manually. |
| Instructions: | Select the management protocol supported by the Frame Relay network. |

| | |
|---|---|
| **Parameter :** | **Address** |
| Wellfleet Default: | Addr Q922 |
| Options: | Addr Q922/Addr Q922 MARCH/Addr Q922 November/Addr Q921 |
| Function: | Specifies the DLCI addressing type. Three variants of addressing, all based upon CCITT draft standard Q.922, are supported.<br><br>Addr Q922 selects addressing as specified in the final version of the Q.922 standard. Q.922 provides for forward explicit congestion notification (FECN), backward explicit congestion notification (BECN), discard eligibility (DE), and address field extension (EA). While most Q.922 addresses are included within a two-octet field, the standard allows for three-octet and four-octet address fields as shown below (Figure 4-7).<br><br>Addr Q922 MARCH differs from Addr Q922 in defining an 11-bit DLCI and dropping the DE bit from the second octet of the address field.<br><br>Addr Q922 November differs from Addr Q922 in dropping the D/C bit from the "extended" (three and four byte) forms. |
| Instructions: | Select the addressing type supported by the Frame Relay network. |

| | |
|---|---|
| **Parameter :** | **Address Length** |
| Wellfleet Default: | TWO BYTE |
| Options: | TWO BYTE/THREE BYTE/FOUR BYTE |
| Function: | Specifies the length of the Frame Relay address field. |
| Instructions: | Select the address length supported by the Frame Relay network. |

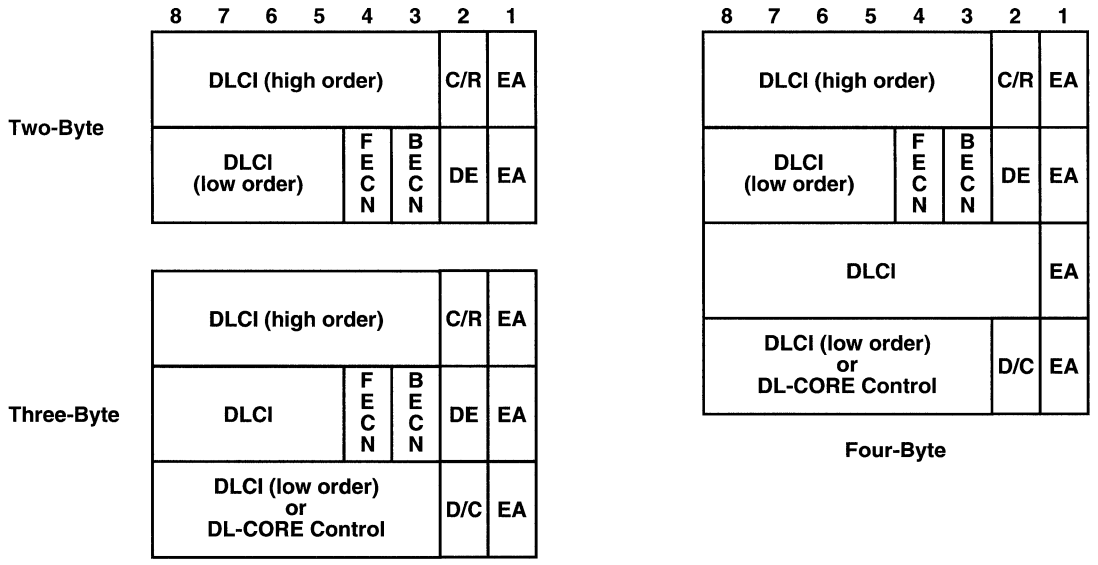


**Figure 4-7. DLCI Q.922 Encoding**

**Parameter : Polling Interval**

Wellfleet Default: 10 (seconds)

Options: 5 - 30 (seconds)

Function: Specifies the interval between *Status Enquiry* messages transmitted by the BN. *Status Enquiry* messages generate a network response in the form of a *Link Integrity Verification* message. Successful completion of the request/response "handshake" verifies the status of the BN/Frame Relay network link.

Instructions: Enter a value within the range 5 to 30 seconds.

**Note:** Polling Interval is non-functional if Mgmnt Type is set to DLCMI None.

| | |
|---|---|
| **Parameter :** | **Full Enquiry Interval** |
| Wellfleet Default: | 6 |
| Options: | 1 - 255 |
| Function: | Specifies the interval between *Full Status Enquiry* messages transmitted by the BN. *Full Status Enquiry* messages generate a network response in the form of a *Full Status Report* message which lists all PVCs, the PVC status (active or inactive), and whether the PVC is new or previously established.<br><br>The default, 6, specifies that the BN send a *Full Status Enquiry* every 6 polling intervals. For example, with a Polling Interval of 10 and a Full Enquiry Interval of 6, the BN will transmit a *Full Status Enquiry* every 60 seconds; with a Polling Interval of 20 and a Full Enquiry Interval of 30, the BN will transmit a *Full Status Enquiry* every 5 minutes (600 seconds). |
| Instructions: | Enter a value within the range 1 to 255. |

**Note:** Polling Interval is non-functional if Mgmnt Type is set to DLCMI None.

| | |
|---|---|
| **Parameter :** | **Error Threshold** |
| Wellfleet Default: | 3 |
| Options: | Any value |
| Function: | Used in conjunction with Monitored Events to establish a metric that evaluates the quality of the BN/Frame Relay network connection.<br><br>Monitored Events specifies a number of status messages (consisting of *Status Enquiries* and *Full Status Enquiries* transmitted by the BN, and *Link Integrity Verifications* and *Full Status Reports* transmitted by the Frame Relay network). Error Threshold specifies the number of faulty status messages required to bring the connection down.<br><br>For example, if default values are accepted for both Error Threshold and Monitored Events, 3 status exchange errors during a continuous sequence of 4 attempted exchanges will cause the connection to be taken down. With Error Threshold set to 5 and Monitored Events set to 10, 5 status exchange errors during a continuous sequence of 10 attempted exchanges will cause the connection to be taken down.<br><br>After a connection is taken down, status exchanges continue and the BN monitors line integrity. When the number of consecutive successful status exchanges is equal to Error Threshold, the BN restores Frame Relay service over the connection. |
| Instructions: | Enter the number of faulty status exchanges that will cause the connection to go down. |

**Note:** Error Threshold and Monitored Events are non-functional if Mgmnt Type is set to DLCMI None.

| | |
|---|---|
| **Parameter :** | **Monitored Events** |
| Wellfleet Default: | 4 |
| Options: | Any value |
| Function: | Used in conjunction with Error Threshold to establish a metric that evaluates the quality of the BN/Frame Relay network connection. |
| Instructions: | Enter the number of sequential status exchanges to be monitored. |

| | |
|---|---|
| **Parameter :** | **Multicast** |
| Wellfleet Default: | Disable |
| Options: | Disable/Enable |
| Function: | Enables or disables support for Frame Relay multicast service. |
| Instructions: | Set to enable to enable multicast service if multicast service is provided by your Frame Relay subscription service, and if this Frame Relay interface is to be a recipient of multicast messages. |

## Editing the Frame Relay Access Mode

By default Frame Relay provides LAN/Group Access to the switched network. LAN Access (described in the *Frame Relay Implementation Note* section of this chapter) is self configuring, provides for a LAN-like logical topology, and dynamically creates new DLCs. LAN Access supports all protocols except for the Bridge, AppleTalk, and VINES.

If you want to define your Frame Relay topology as a group of point-to-point connections, or if you wish to bridge traffic across a Frame Relay network, you must manually configure PVCs. Use the procedures in the following two sections to configure either Direct Access Mode or Hybrid Mode.

## Configuring Direct Access Mode

Direct Access mode (described in the *Frame Relay Implementation Note* section of this chapter) provides point-to-point connectivity between Frame Relay DCEs. It requires manual configuration of each Frame Relay PVC.

You begin from the Wellfleet Configuration Manager Window (Figure 4-2).

1. Select the Circuits/Edit Circuits option to display the Circuit List Window (Figure 4-8).

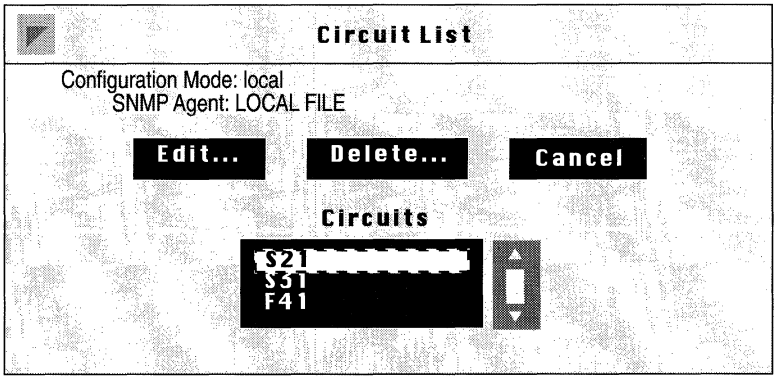


**Figure 4-8. Circuit List Window**

2. Select the Frame Relay interface you want to edit from list of the circuits.

   In this example, circuit S21 is selected.

3. Click the Edit button to display the Circuit Definition Window (Figure 4-9).
4. Select the Protocols/Edit Frame Relay/Permanent Virtual Circuits option to display the Frame Relay PVC Window (Figure 4-10).

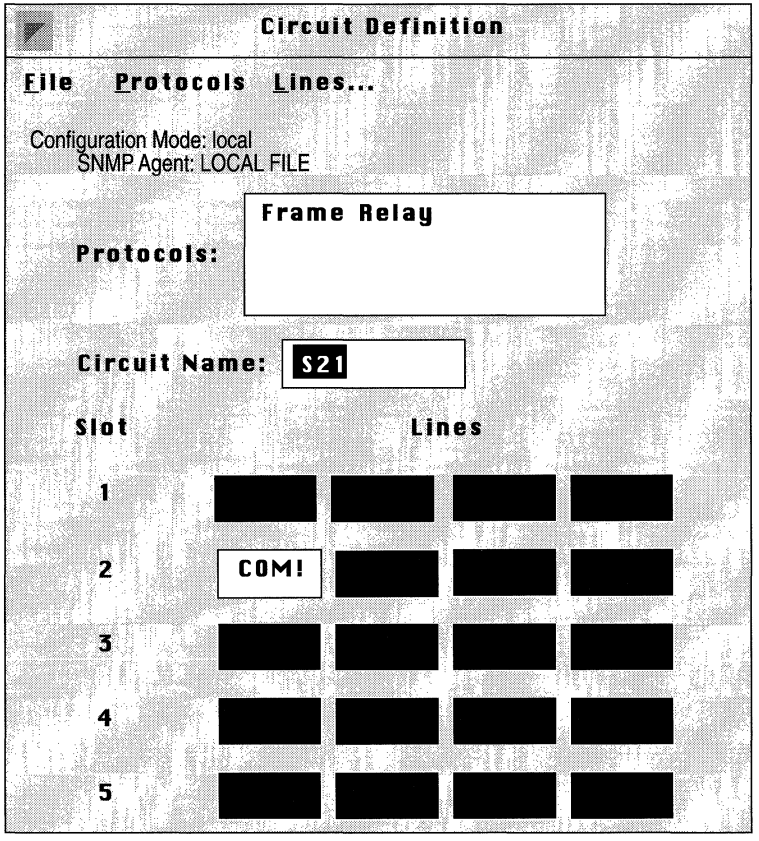


**Figure 4-9. Circuit Definition Window**

5. Click the Add button to display the Frame Relay DLCI Window (Figure 4-11).
6. Assign a DLCI (data link connection identifier) to the PVC.

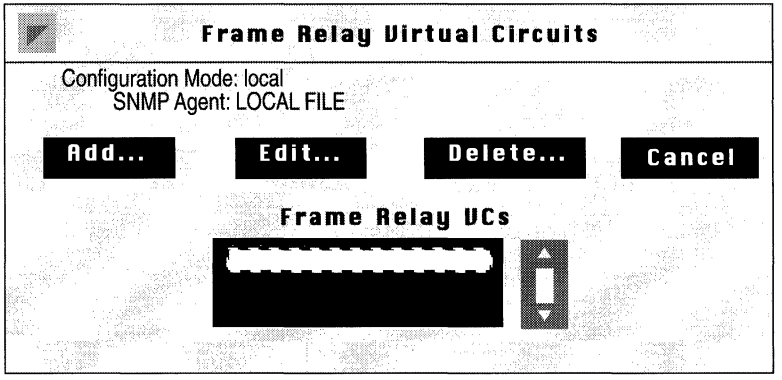


**Figure 4-10. Frame Relay PVC Window**

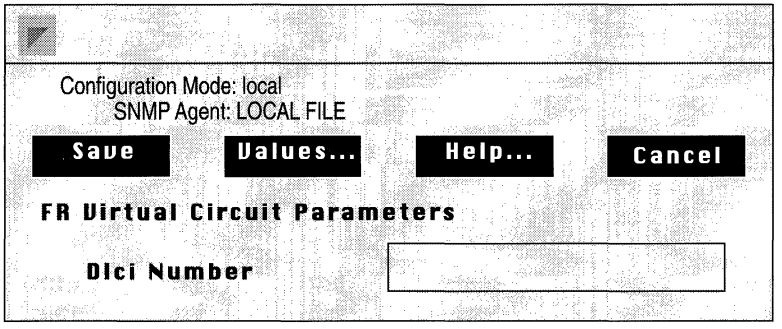


**Figure 4-11. Frame Relay DLCI Window**

| **Parameter :** | **Dlci Number** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid DLCI number |
| Function: | Assigns a DLCI to this direct access PVC. |
| Instructions: | Enter the Frame Relay service provider-assigned DLCI in decimal format. |

7. Click the Save button to return to the Frame Relay PVC Window (see Figure 4-10).
8. Select the PVC you just added from the displayed list.
9. Click the Edit button to display the Frame Relay PVC Parameter Window (Figure 4-12).

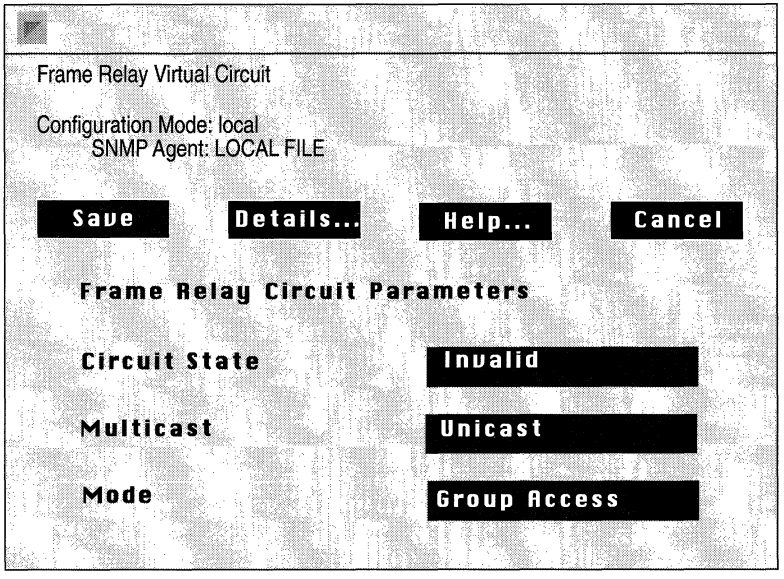


**Figure 4-12. Frame Relay PVC Parameter Window**

10. Assign values to the three required (Circuit State, Multicast, and Mode) Frame Relay PVC parameters as described below.

| | |
|---|---|
| **Parameter :** | **Circuit State** |
| Wellfleet Default: | Invalid |
| Options: | Invalid/Active/Inactive |
| Function: | Specifies the state of this direct access PVC. |
| Instructions: | Set Circuit State to either Active (indicating that the PVC is available for use) or Inactive (indicating that the PVC is configured, but not available for use). |

| | |
|---|---|
| **Parameter :** | **Multicast** |
| Wellfleet Default: | Unicast |
| Options: | Unicast/Multicast |
| Function: | Indicates whether this PVC is multicast or unicast. |
| Instructions: | Set to unicast or multicast according to the PVC type. |

| | |
|---|---|
| **Parameter :** | **Mode** |
| Wellfleet Default: | Group Access |
| Options: | Group Access/Direct Access/Hybrid Access |
| Function: | Specifies the network access mode. |
| Instructions: | Set to Direct Access. |

11. Click the Save button to display a Frame Relay Bearer Protocol Selection Window (Figure 4-13).

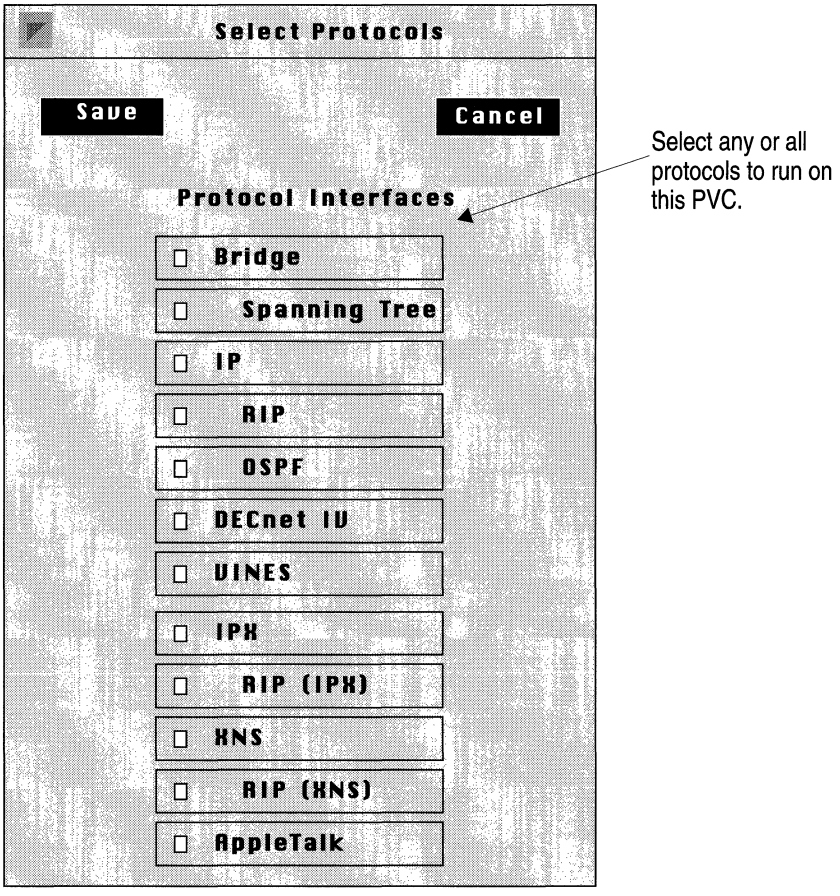


**Figure 4-13. Frame Relay Bearer Protocol Selection Window**

12. Select a protocol or protocols to be carried on this PVC, and click the Save button.

    For each protocol you enable, the Configuration Manager displays a protocol-specific configuration window prompting for additional required information. If you need assistance in responding to any queries, consult the appropriate protocol-specific chapter of this guide.

When you complete protocol definition, the Configuration Manager returns you to the Frame Relay PVC Window (see Figure 4-8).

13. Click the Cancel button to complete the configuration of this direct access PVC.
14. Repeat steps 1 through 13 to configure additional direct access PVCs.

### Configuring Hybrid Access Mode

Hybrid Access mode (described in the *Frame Relay Implementation Note* section of this chapter) allows the configuration of both routing and bridging on a single PVC. Routing treats PVCs as LAN connections while the Bridge treats PVCs as point-to-point (direct access) connections. Like Direct Access mode, Hybrid Access requires manual configuration of each Frame Relay PVC.

You begin from the Wellfleet Configuration Manager Window (Figure 4-2).

1. Select the Circuits/Edit Circuits option to display the Circuit List Window (Figure 4-14).

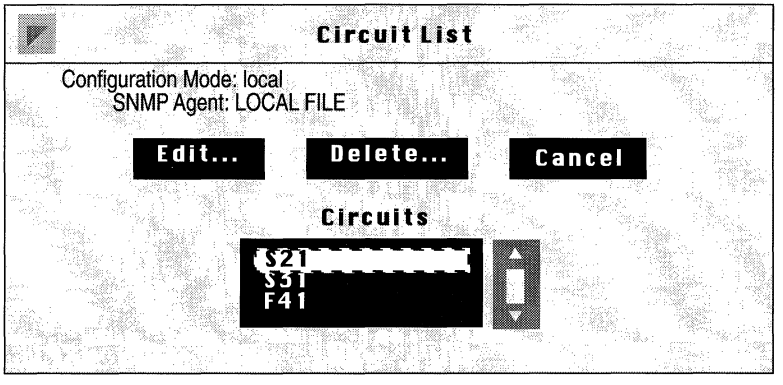


**Figure 4-14. Circuit List Window**

2. Select the Frame Relay interface you want to edit from list of the circuits.

   In this example, circuit S21 is selected.

3. Click the Edit button to display the Circuit Definition Window (Figure 4-15).

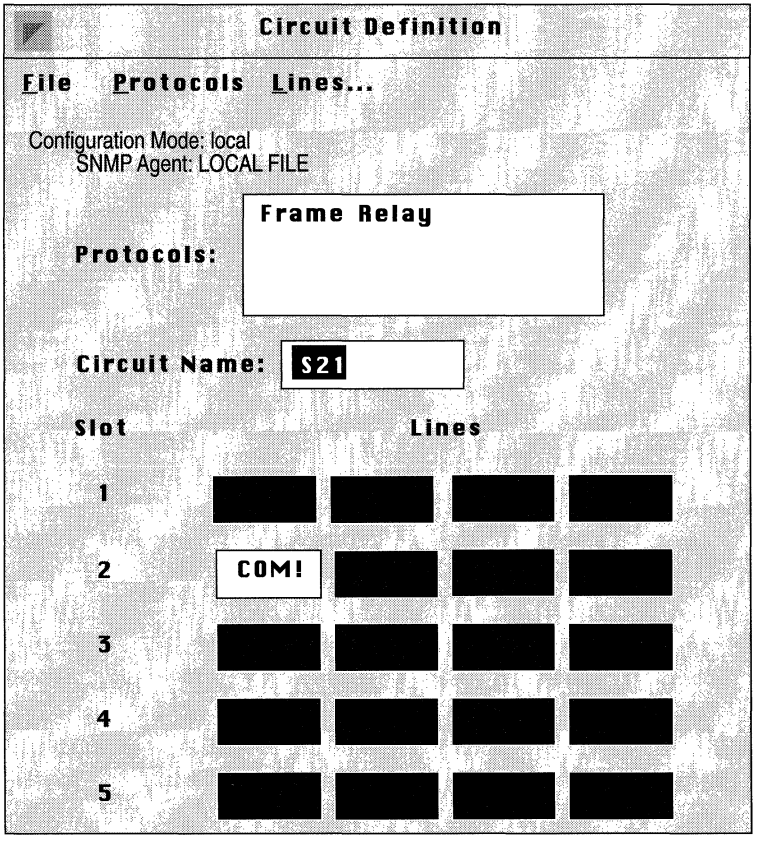


**Figure 4-15. Circuit Definition Window**

4. Select the Protocols/Edit Frame Relay/Permanent Virtual Circuits option to display the Frame Relay PVC Window (Figure 4-16).

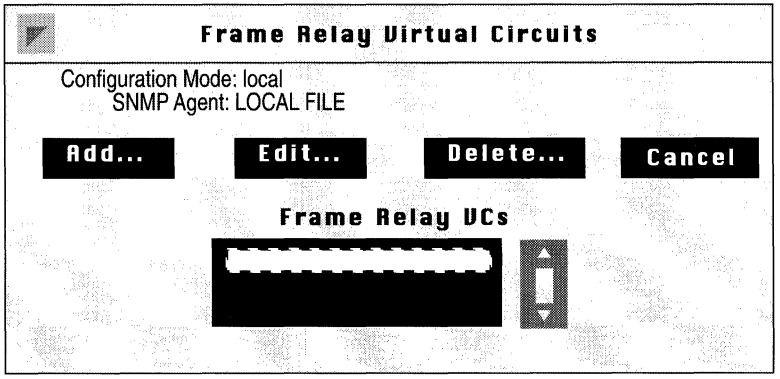


**Figure 4-16. Frame Relay PVC Window**

5. Click the Add button to display the Frame Relay DLCI Window (Figure 4-17).

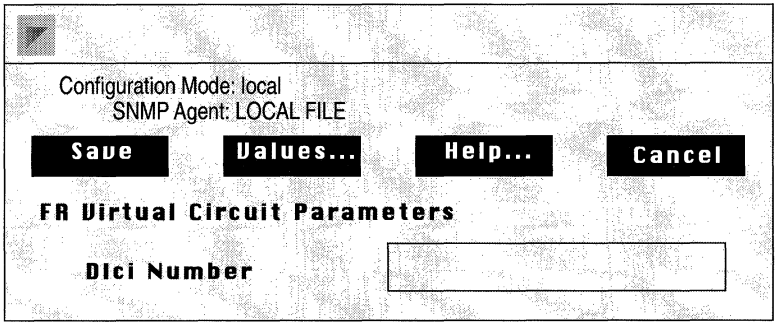


**Figure 4-17. Frame Relay DLCI Window**

6. Assign a DLCI (data link connection identifier) to the PVC.

| **Parameter :** | **Dlci Number** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid DLCI number |
| Function: | Assigns a DLCI to this hybrid access PVC. |
| Instructions: | Enter the Frame Relay service provider-assigned DLCI in decimal format. |

7. Click the Save button to return to the Frame Relay PVC Window (see Figure 4-14).
8. Select the PVC you just added from the displayed list.
9. Click the Edit button to display the Frame Relay PVC Parameter Window (Figure 4-18).

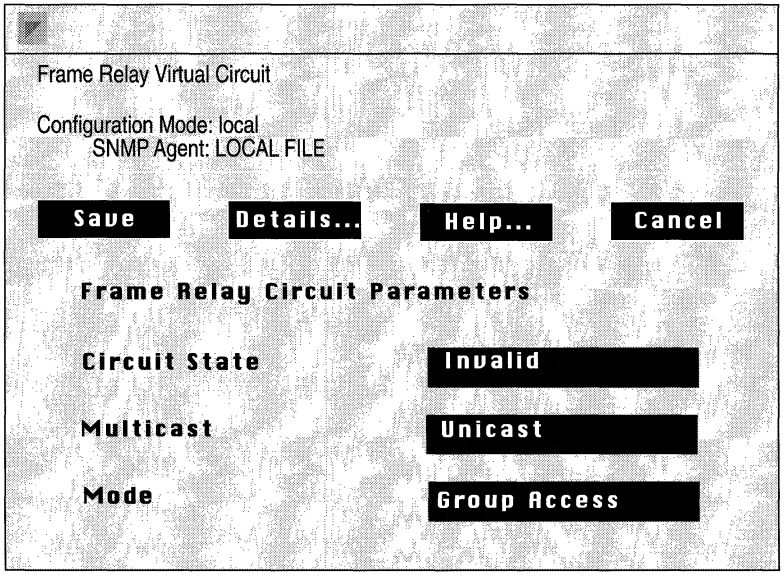


**Figure 4-18. Frame Relay PVC Parameter Window**

10. Assign values to the three required (Circuit State, Multicast, and Mode) Frame Relay PVC parameters as described below.

| | |
|---|---|
| **Parameter :** | **Circuit State** |
| Wellfleet Default: | Invalid |
| Options: | Invalid/Active/Inactive |
| Function: | Specifies the state of this hybrid access PVC. |
| Instructions: | Set Circuit State to either Active (indicating that the PVC is available for use) or Inactive (indicating that the PVC is configured, but not available for use). |

| | |
|---|---|
| **Parameter :** | **Multicast** |
| Wellfleet Default: | Unicast |
| Options: | Unicast/Multicast |
| Function: | Determines whether this PVC is multicast or unicast. |
| Instructions: | Set to unicast or multicast according to the PVC type. |

| | |
|---|---|
| **Parameter :** | **Mode** |
| Wellfleet Default: | Group Access |
| Options: | Group Access/Direct Access/Hybrid Access |
| Function: | Specifies the network access mode. |
| Instructions: | Set to Hybrid Access. |

11. Click the Save button to display the Frame Relay Hybrid Access Bridge Window (Figure 4-19).

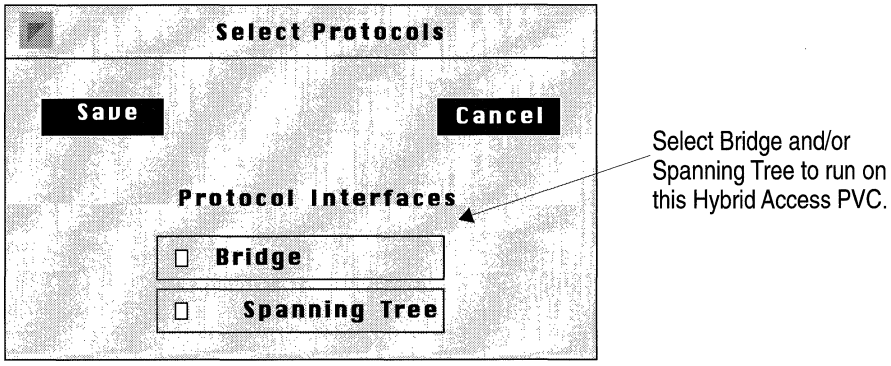


**Figure 4-19. Frame Relay Hybrid Access Bridge Window**

12. Specify the Bridge and/or Spanning Tree to be carried on this hybrid access PVC, and click the Save button.

**Note:** Wellfleet strongly recommends that the Spanning Tree be enabled on all hybrid access PVCs to detect possible loops within the network.

13. Click the Cancel button to complete the configuration of this hybrid access PVC.
14. Repeat steps 1 through 13 to configure additional hybrid access PVCs.

## Editing Frame Relay PVCs

The Frame Relay PVC Window allows you to add, edit, and delete manually configured Frame Relay PVCs. You begin from the Wellfleet Configuration Manager Window (Figure 4-2).

1. Select the Circuits/Edit Circuits option to display the Circuit List Window.
2. Select the Frame Relay interface you want to edit from list of the circuits.
3. Click the Edit button to display the Circuit Definition Window.
4. Select the Protocols/Edit Frame Relay/Permanent Virtual Circuits option to display the Frame Relay PVC Window (Figure 4-20).

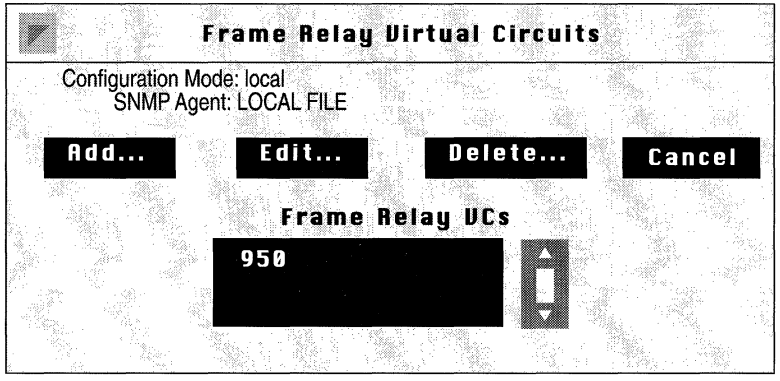


**Figure 4-20. Frame Relay PVC Window**

### Adding a PVC

To add a manually configured PVC to a Frame Relay interface click the Add button to display the Frame Relay DLCI Window (Figure 4-21).

1. Assign a DLCI (data link connection identifier) to the PVC.

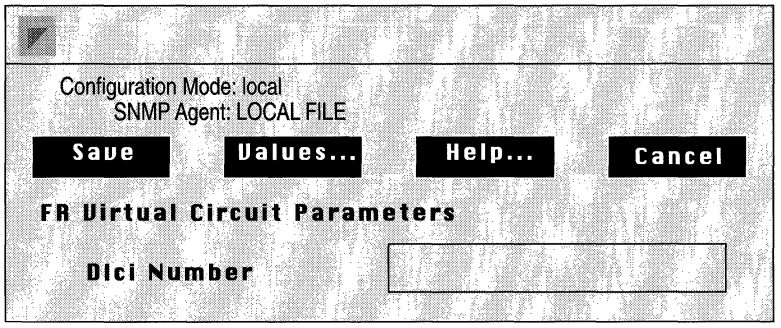


**Figure 4-21. Frame Relay DLCI Window**

| **Parameter :** | **Dlci Number** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid DLCI number |
| Function: | Assigns a DLCI to this hybrid access PVC. |
| Instructions: | Enter the Frame Relay service provider-assigned DLCI in decimal format. |

2. Click the Save button to return to the Frame Relay PVC Window.

### Editing a PVC

To edit PVC parameters select the PVC you wish to edit in the Frame Relay VCs section of the Frame Relay PVC Window, and then click the Edit button to display the Frame Relay PVC Parameter Window (Figure 4-22).

1. Assign values to the Frame Relay PVC parameters as necessary.

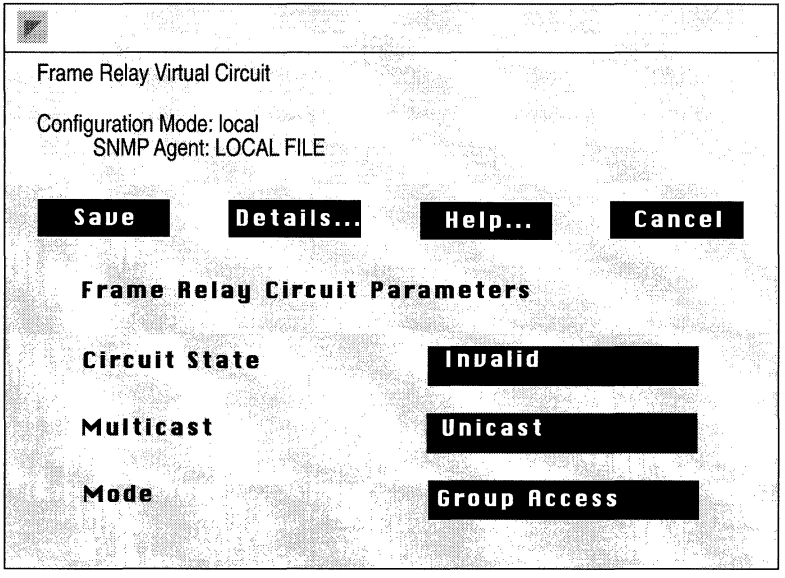


**Figure 4-22. Frame Relay PVC Parameter Window**

| **Parameter :** | **Circuit State** |
|---|---|
| Wellfleet Default: | Invalid |
| Options: | Invalid/Active/Inactive |
| Function: | Specifies the state of this hybrid access PVC. |
| Instructions: | Set Circuit State to either Active (indicating that the PVC is available for use) or Inactive (indicating that the PVC is configured, but not available for use). |

| | |
|---|---|
| **Parameter :** | **Multicast** |
| Wellfleet Default: | Unicast |
| Options: | Unicast/Multicast |
| Function: | Determines whether this PVC is multicast or unicast. |
| Instructions: | Set to enable to enable multicast service if multicast service is provided by your Frame Relay subscription service, and if this Frame Relay PVC is to be a recipient of multicast messages. |

| | |
|---|---|
| **Parameter :** | **Mode** |
| Wellfleet Default: | Group Access |
| Options: | Group Access/Direct Access/Hybrid Access |
| Function: | Specifies the network access mode. |
| Instructions: | Set to match the Frame Relay network view. |

2. Click the Save button.
3. If you selected Direct or Hybrid Access assign routing and bridging protocols as described in the sections *Configuring Direct Access Mode* and *Configuring Hybrid Access Mode* in this chapter.

### Deleting a PVC

To delete a PVC select the PVC you wish to delete in the Frame Relay VCs section of the Frame Relay PVC Window, and then click the Delete button. A Popup Window asks you to confirm the deletion. Select OK to delete the PVC or Cancel to retain the PVC.

# Deleting Frame Relay from the BN

You can delete Frame Relay service from *all* BN circuits on which it is currently configured in two steps.

Begin from the Wellfleet Configuration Manager Window (see Figure 4-4):

1. Select the Protocols/Frame Relay/Delete Frame Relay option.

   A window pops up and prompts "Do you REALLY want to delete Frame Relay?".

2. Select Ok.

   You are returned the Wellfleet Configuration Manager window. Frame Relay is no longer configured on the BN.

# Chapter 5

## Configuring SMDS

## List of Figures

# Configuring SMDS

## About This Chapter

This chapter describes how to configure Switched Multi-Megabit Data Service (SMDS). Users who are already familiar with SMDS may wish to proceed directly to the section entitled *Editing SMDS Parameters*. Users desiring some background material may find it useful to read the sections entitled *SMDS Overview* and *SMDS Implementation Note* in addition to consulting some of the documents listed in the *SMDS Bibliography*.

## SMDS Overview

SMDS is a public, high-speed, packet switched network based on cell relay technology and the IEEE 802.6 MAN standard. SMDS extends LAN-like performance beyond the subscriber's premises across a metropolitan or wide area. SMDS is well suited for bursty traffic that consumes high bandwidth for short periods of time. SMDS currently offers six network access speeds (or classes): 1.2 Mbps, 4 Mbps, 10 Mbps, 16 Mbps, 25 Mbps, and 34 Mbps. Within the SMDS fabric, network switches are generally linked by T3 trunks.

SMDS provides a connectionless transfer mode with no establishment of a logical end-to-end connection. Each SMDS packet contains a MAC-level address (referred to as an E.164 address) used by network switches to route the packet to its destination. The E.164 address format provides for both individual and multicast addresses.

SMDS requirements are spelled out in various Bellcore technical documents listed in the SMDS bibliography. *Generic System Requirements in Support of Switched Multi-Megabit Data Service* defines an SMDS Interface Protocol (SIP) which specifies SMDS addressing, formatting, framing, and error detection requirements. The SIP has three levels or layers. Level 3 specifies the addressing, formatting, and encapsulation of packetized data, referred to as L3PDU packets. Level 2specifies the segmentation of L3PDUs into short fixed length SMDS cells, referred to as L2PDU packets. Level 1 specifies the physical connectivity.

In an effort to hasten SMDS market introduction the SMDS Special Interest Group (SIG) chose to divide required SIP functionality between devices which handle local network packets (for example, a BN) and devices which interface with the digital services provided by common carriers (a CSU/DSU). This division is specified in the Data eXchange Interface (DXI) protocol which describes the router/CSU interface. Figure 5-1 illustrates the assembly of SMDS cells as specified by the DXI.

SMDS packet assembly commences when the BN SMDS service receives a network-generated packet. The BN takes the entire packet and encapsulates it within a 36-byte header (containing addressing, length and control information) and a 4-byte trailer (containing a CRC value) thus creating an SMDS L3PDU. The BN next prepares the L3PDU for transmission to the CSU by encapsulating it within an HDLC-like DXI header and trailer which provide control information. Upon receiving the DXI packet, the CSU/DSU strips the DXI header/ trailer and segments the L3PDU into fixed length (44-byte) units called cells. These cells are then encapsulated within a 7-byte header and 2-byte trailer to forms L2PDUs. Finally the CSU/DSU inserts an additional four bytes of framing information between each L2PDU and transmits the framed cells across a DS1 or DS3 connection to the SMDS network.

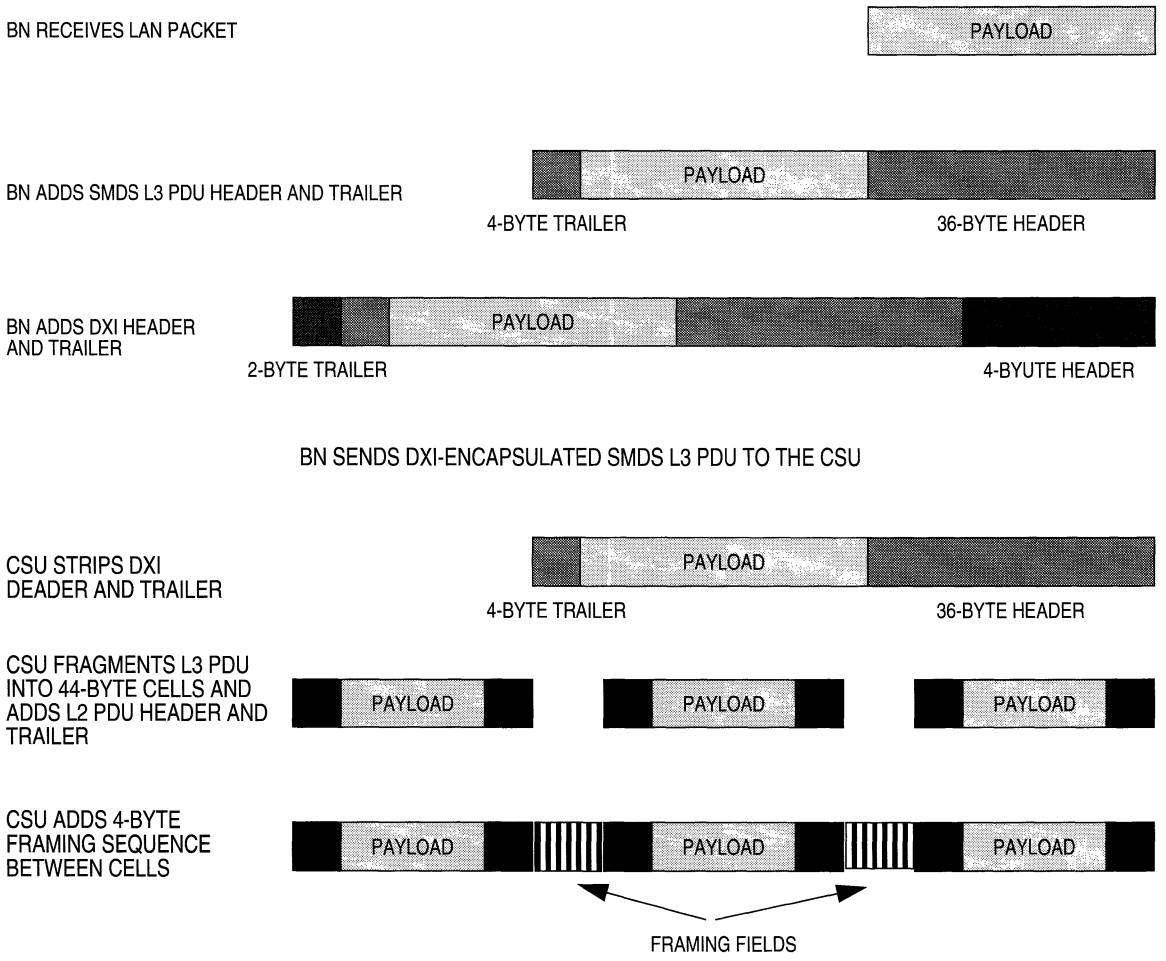


**Figure 5-1. DXI Packet Assembly**

# SMDS Bibliography

The following documents provide technical detail on SMDS and DXI protocol design and implementation.

Baker, F. and Kolb, C. *Definitions of Managed Objects for the DS1 Interface Type*. RFC 1232, Network Information Center (NIC), SRI International, Menlo Park, CA, May 1991.

Cox, T. and Tesink, K. *Definitions of Managed Objects for the DS3 Interface Type*. RFC 1233, Network Information Center (NIC), SRI International, Menlo Park, CA, May 1991.

Bellcore. *Generic System Requirements in Support of Switched Multi-Megabit Data Service*. Technical Reference TR-TSV-000772, Issue 1, May 1991.

Bellcore. *Local Access System Generic Requirements, Objectives, and Interfaces in Support of Switched Multi-Megabit Data Service*. Technical Reference TR-TSV-000773, Issue 1, June 1991.

Bellcore. *Generic Requirements for SMDS Customer Network Management Service*. Technical Advisory TA-TSV-001062, Issue 2, February 1992.

Piscitello, D. and Lawrence, J. *The Transmission of IP Datagrams over the SMDS Service*. RFC 1209, Network Information Center (NIC), SRI International, Menlo Park, CA, March 1991.

SMDS Interest Group. *SMDS Data Exchange Interface Protocol* (Revision 3.2). Technical Specification SIG-TS-001/1991, October 1991.

SMDS Interest Group. *SMDS DXI Local Management Interface*. Technical Specification SIG-TS-002/1992, May 1992.

The following publications provide a less technical introduction to SMDS service.

Davidson, R. and Muller, N. *The Guide to SONET: Planning, Installing & Maintaining Broadband Networks*. Telecom Library, Inc., 1991.

Goldstein, F. *ISDN in Perspective*. Addison-Wesley Publishing Company, 1992.

# SMDS Implementation Note

Implementation of SMDS requires the BN and an SMDS CSU/DSU which provides DS1 or DS3-based access to the switched SMDS network. The physical connection between the BN and the CSU/DSU is provided by a synchronous or HSSI connection.

**Note:** Connectivity between the BN and the CSU/DSU cannot be accomplished with either a T1 or E1 circuit.

Data exchange between the BN and the CSU/DSU is managed by Version 3.2 of the DXI protocol. As earlier DXI versions are not supported by the BN, it is imperative that the CSU/DSU support DXI Version 3.2. DXI Version 3.2 provides an optional "Heartbeat Poll" mechanism to verify the BN and CSU/DSU connection on a periodic basis. Before enabling heartbeat polling, you should ensure that CSU/DSU supports this feature and that it has also been enabled on the CSU/DSU.

As DXI provides support for both 16-bit and 32-bit CRCs you must ensure that the CRC values match for both the BN and the CSU/DSU. If necessary, CRC values (16-bit or 32-bit) can be modified for both synchronous and HSSI connections.

The Local Management Interface (LMI) protocol works in conjunction with DXI to facilitate the exchange of management information between the BN and the CSU/DSU. Implemented by the BN-resident SNMP proxy agent, LMI uses a subset of SNMP to provide for BN management queries, CSU/DSU responses to queries, and the generation of asynchronous trap event by the CSU/DSU. Before enabling LMI, you should ensure that the CSU/DSU supports this feature and that it has also been enabled on the CSU/DSU.

SMDS service provides support for the following protocols: Bridge (to include the Spanning Tree), IP (to include ARP support), DECnet, IPX, XNS, and Source Routing (with Wellfleet 8101 encapsulation).

**Note:** AppleTalk and VINES are not supported by SMDS.

# Editing SMDS Parameters

Once you have configured a circuit to support SMDS (by specifying circuit-specific individual, group and ARP addresses), you can use the Configuration Manager to tailor the default SMDS service to meet your network needs.

You begin from the Wellfleet Configuration Manager Window (Figure 5-2).

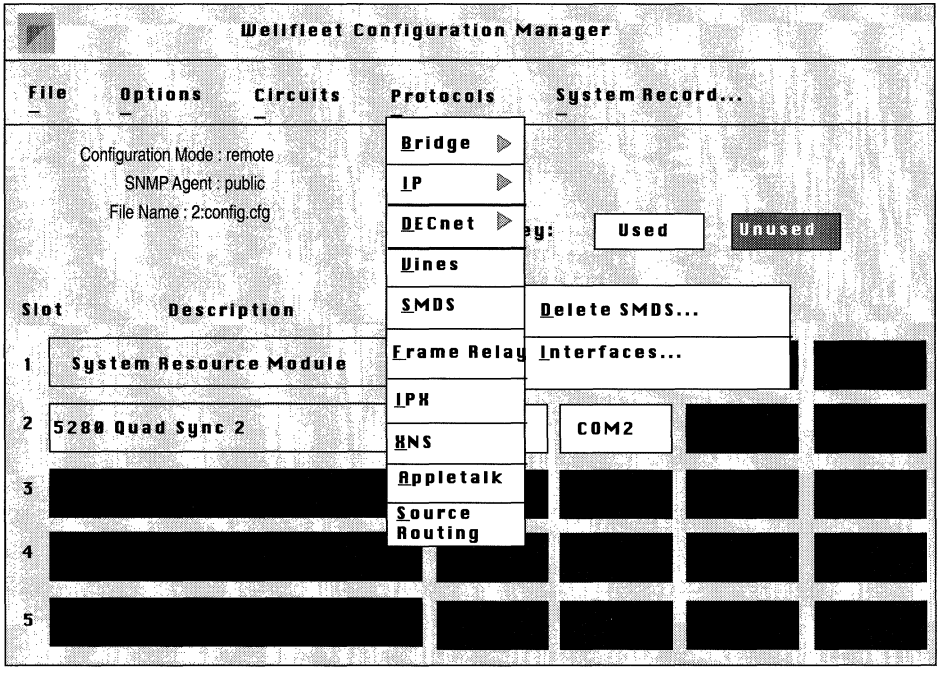


**Figure 5-2. Wellfleet Configuration Manager Window**

## Editing SMDS Interface Parameters

You edit SMDS interface parameters from the SMDS Interface Parameters Window. Begin at the Wellfleet Configuration Manager Window, then select the Protocols/SMDS/Interfaces option to display the SMDS Interfaces Window (Figure 5-3). Next, select the interface you wish to edit, then click on the Edit button to display the SMDS Interface Parameters Window (Figure 5-4) for that interface.

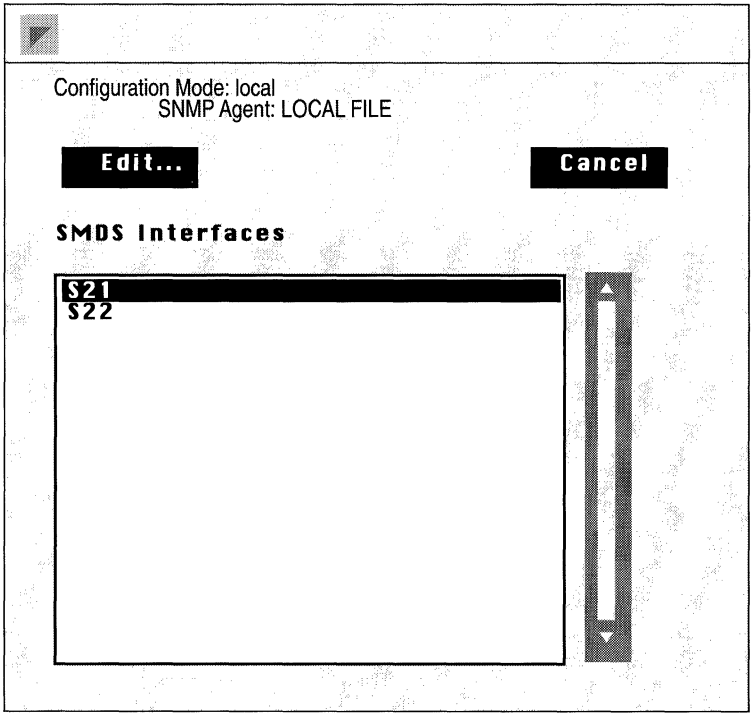


**Figure 5-3. SMDS Interfaces Window**

This section provides the information you need to edit each parameter in the SMDS Interface Parameters Window. Refer to this information to edit parameters you wish to change. When you are done, click the Save button to exit the window and save your changes.

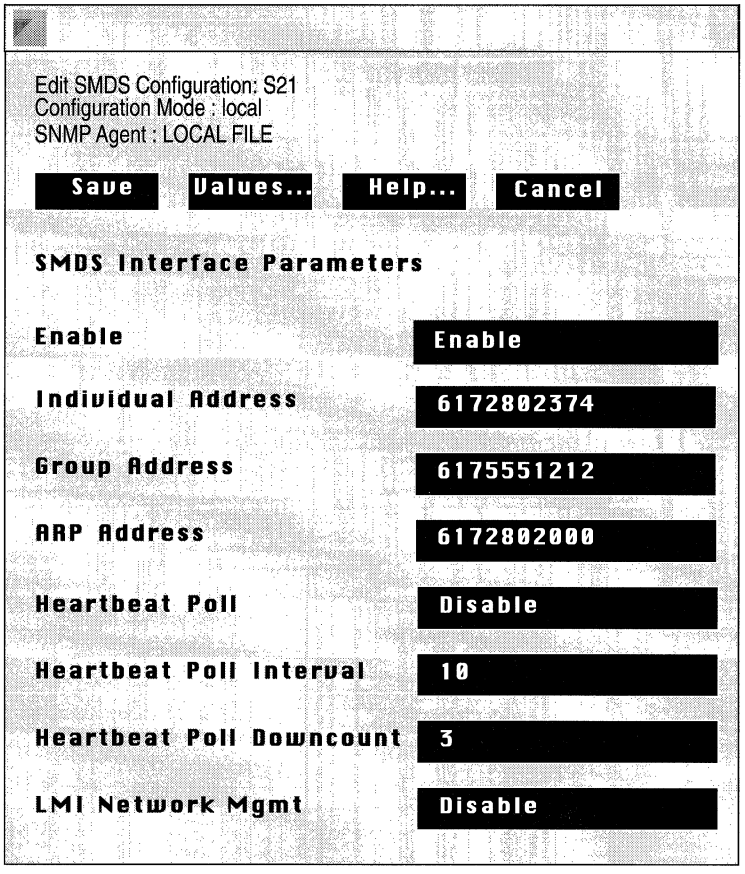


**Figure 5-4. SMDS Interface Parameters Window**

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or Disables SMDS service on this interface. |
| Instructions: | Set to Disable if you want to disable SMDS service on this interface. |

| **Parameter :** | **Individual Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid 10-digit North American Numbering Plan (NANP) telephone number. |
| Function: | Provides a local MAC-layer address. |
| Instructions: | Enter the 10-digit local address (telephone number) as provided by the SMDS subscription agreement. |

**Note:** SMDS standards specify the use of 8-octet E.164 addresses. The BN converts the Individual Address to E.164 format by prepending hexadecimal C1 and appending hexadecimal FF FF to the NANP number. For example, the individual address 6172802374 becomes C1 61 72 80 23 74 FF FF.

| **Parameter :** | **Group Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid 10-digit North American Numbering Plan (NANP) telephone number. |
| Function: | Provides a MAC-layer multicast address. |
| Instructions: | Enter the 10-digit multicast address as provided by the SMDS subscription agreement. |

**Note:** SMDS standards specify the use of 8-octet E.164 addresses. The BN converts the Group Address and ARP Address to E.164 format by prepending hexadecimal E1 and appending hexadecimal FF FF to the NANP number. For example, the group address 6175551212 becomes E1 61 75 55 12 12 FF FF.

| | |
|---|---|
| **Parameter :** | **ARP Address** |
| Wellfleet Default: | None |
| Options: | Any valid 10-digit North American Numbering Plan (NANP) telephone number. |
| Function: | Provides an address resolution multicast address. |
| Instructions: | Enter the 10-digit multicast address as provided by the SMDS subscription agreement. |

**Note:** SMDS standards specify the use of 8-octet E.164 addresses. The BN converts the ARP Address to E.164 format by prepending hexadecimal C1 and appending hexadecimal FF FF to the NANP number. For example, the ARP address 6172802000 becomes C1 61 72 80 20 00 FF FF.

| | |
|---|---|
| **Parameter :** | **Heartbeat Poll** |
| Wellfleet Default: | Disable |
| Options: | Enable/Disable |
| Function: | Enables or disables DXI heartbeat polling.<br>DXI Version 3.2 provides a heartbeat polling mechanism, a simple way to verify the integrity of the BN/DSU connection. Heartbeat polling is implemented by the regular transmission of keepalive messages from the BN to the DSU and the receipt of an acknowledgment from the DSU. |
| Instructions: | Set to Enable to enable heartbeat polling. |

| | |
|---|---|
| **Parameter :** | **Heartbeat Poll Interval** |
| Wellfleet Default: | 10 |
| Options: | 6 - 1023 seconds |
| Function: | Specifies the time interval between the transmission of heartbeat poll messages by the BN. If heartbeat polling is disabled, this parameter is non-functional. |
| Instructions: | Set to the number of seconds between BN transmission of heartbeat poll messages. |

**Note:** Heartbeat Poll Interval must be set to a value greater the 5 seconds, the length of the heartbeat poll acknowledgment timer.

| | |
|---|---|
| **Parameter :** | **Heartbeat Poll Downcount** |
| Wellfleet Default: | 3 |
| Options: | 1 - 1023 |
| Function: | Specifies the number of heartbeat poll messages that can remain unacknowledged before the BN declares the BN/DSU connection down. If heartbeat polling is disabled, this parameter is non-functional. |
| Instructions: | Set to the number of unacknowledged heartbeat poll messages that the BN will tolerate before taking the BN/DSU connection down. |

| | |
|---|---|
| **Parameter :** | **LMI Network Mgmt** |
| Wellfleet Default: | Disable |
| Options: | Enable/Disable |
| Function: | Enables or disables LMI network management.<br>LMI (Local Management Interface) works in conjunction with DXI Version 3.2. LMI is an SNMP-based protocol that enables the exchange of management information between the BN and the DSU. |
| Instructions: | Set to Enable to enable LMI protocol. |

# Deleting SMDS from the BN

You can delete SMDS service from *all* BN circuits on which it is currently configured in two steps.

Begin from the Wellfleet Configuration Manager Window (see Figure 5-2):

1. Select the Protocols/SMDS/Delete SMDS option.

   A window pops up and prompts "Do you REALLY want to delete SMDS?".

2. Select Ok.

   You are returned the Wellfleet Configuration Manager window. SMDS is no longer configured on the BN.

# Chapter 6

## Configuring AppleTalk

## List of Figures



## List of Tables

# Configuring AppleTalk

## About this Chapter

This chapter describes how to configure the Wellfleet AppleTalk router. This release of AppleTalk software supports AppleTalk Phase 2.

This chapter begins with an overview of AppleTalk technology. The second section lists additional AppleTalk reference material. The third section describes how the Wellfleet AppleTalk router works. The fourth section describes implementation guidelines for adding Wellfleet AppleTalk routers to your network. The final section describes how to use the Configuration Manager to edit AppleTalk parameters and how to delete AppleTalk from the BN.

## AppleTalk Overview

The AppleTalk network system was developed by Apple Computer Inc. to allow users of Apple and non-Apple computers on a network to communicate and share resources (such as printers, file servers etc.). AppleTalk's dynamic addressing scheme allows a user to plug a computer into an AppleTalk network and gain access to the network immediately, without first performing any complicated configuration procedures.

An AppleTalk internet is physically divided into distinct networks consisting of AppleTalk end nodes connected by routers. AppleTalk end nodes can send and receive messages. Routers (like the Wellfleet AppleTalk router) can send and receive messages, and can route messages throughout the AppleTalk network in datagram format.

## AppleTalk Addressing

There are two types of AppleTalk networks: *extended* and *nonextended*.

An extended network can support up to 16 million nodes, in theory, and has the following characteristics:

- The network is assigned a *range* of 16-bit network numbers.
- Each node within the network is dynamically assigned a 24-bit AppleTalk address (see Figure 6-1) that consists of a network number chosen from within the range assigned combined with an 8-bit node number. The numbers 0, $FF, and $FE are reserved.

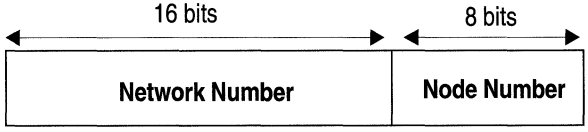


**Figure 6-1. AppleTalk Address**

A nonextended network can support up to 254 nodes and has the following characteristics:

- The entire network is assigned a *single* 16-bit network number.
- Each node within the network is dynamically assigned a 24-bit AppleTalk address that consists of the assigned network number combined with a unique, 8-bit node number. The numbers 0 and $FF are reserved.

**Note:** The Wellfleet AppleTalk router will not route packets to directly connected nonextended networks.

Figure 6-2 shows both types of AppleTalk networks: networks 2-3, 4-5, 6-7, and 8-11 are extended networks; network 1 is a nonextended network.

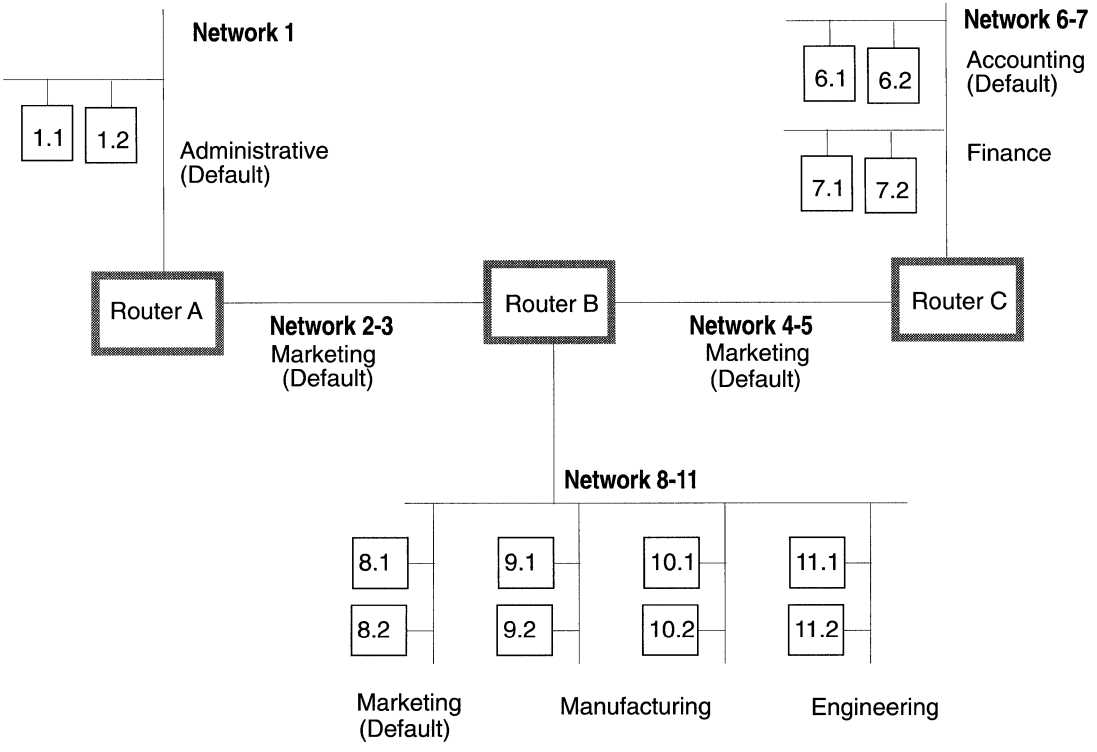


| Network Number Range | Network Type | Default Zone | Zone List |
|---|---|---|---|
| 1 | Nonextended | Administrative | Administrative |
| 2-3 | Extended | Marketing | Marketing |
| 4-5 | Extended | Marketing | Marketing |
| 6-7 | Extended | Accounting | Accounting<br>Finance |
| 8-11 | Extended | Marketing | Marketing<br>Manufacturing<br>Engineering |

**Figure 6-2. Addressing on an AppleTalk Network**

## AppleTalk Zones

Each AppleTalk network is logically divided into areas called *zones*. Zones simply help a user identify where a network entity, or service, can be found. Similar network services are usually assigned to the same zone.

A network's *zone list* contains all the zones assigned to the network. A nonextended network consists of only a single zone. An extended network can be divided up into 255 zones, one of which is designated as the *default zone*. The same zone can be part of many different networks. When a new node first starts up on the network, it is assigned to the default zone. Later, it can be reassigned to any valid zone on the zone list.

## Seed Routers and Nonseed Routers

Each AppleTalk network must have at least one router designated as the *seed* router. A seed router is configured with the following information:

- Network number start range
- Network number end range
- Default zone name
- Zone list for the network

The seed router provides this network configuration information for all other nonseed routers on the network. Multiple seed routers can reside on the same network, however, they all must be configured with the same network ranges, default zone name and zone list.

# AppleTalk Bibliography

The following document provides technical detail on AppleTalk protocol implementation.

Sidhu, Andrews, Oppenheimer. *Inside Appletalk* .Addison-Wesley Publishing Company. Second Edition. ISBN 0-201-55021-0

# How the Wellfleet AppleTalk Router Works

This section is *optional* reading. It describes the following AppleTalk protocols that the Wellfleet AppleTalk router uses to route packets across an AppleTalk network, and notes any Wellfleet divergence from these standards:

- TokenTalk Link Access Protocol (TLAP)
- EtherTalk Link Access Protocol (ELAP)
- AppleTalk Address Resolution Protocol (AARP)
- Datagram Delivery Protocol (DDP)
- Routing Table Maintenance Protocol (RTMP)
- Zone Information Protocol (ZIP)
- Name Binding Protocol (NBP)
- AppleTalk Echo Protocol (AEP)

This section also describes how a Wellfleet AppleTalk interface initializes on the network.

## AppleTalk Link Access Protocols

The Wellfleet AppleTalk router uses AppleTalk link access protocols to transmit packets between nodes on the *same* physical network. (TokenTalk controls data transmission on token ring networks, EtherTalk controls data transmission on Ethernet networks). The router also supports FDDI and Wellfleet proprietary synchronous encapsulation.

**Note:** The Wellfleet AppleTalk router does not support LocalTalk or AppleTalk Phase 1 routing. Thus, the Wellfleet AppleTalk router cannot directly attach to nonextended networks.

## AppleTalk Address Resolution Protocol

In order for a Wellfleet AppleTalk router to forward a packet to a directly connected AppleTalk node, it needs to know both the packet's AppleTalk address, and the corresponding hardware address of the node on which the AppleTalk address is found. The router uses the AppleTalk Address Resolution Protocol (AARP) to map AppleTalk addresses to their equivalent hardware addresses. The router saves this information in its Address Mapping Table (AMT), which lists all known AppleTalk addresses, corresponding hardware addresses, and the circuit/port on which the address resolution is in effect.

The router maintains and updates its AMT by broadcasting AARP *Request* packets and receiving AARP *Response* packets to all AppleTalk nodes on the network when necessary (see Figure 6-3).

When the router needs to send a packet to a given AppleTalk address, it scans its AMT to find the address. If the address is not found, the router broadcasts a single Request packet in order to find out which node is using the address. If the address exists, the node whose address matches that specified in the Request packet sends back a Response, which identifies the hardware address that maps to the AppleTalk address. The router then updates its AMT with this new information. The router will wait 2 seconds for a Response.

AARP is also responsible for generating a unique AppleTalk address for each of the router's AppleTalk interfaces that have not been explicitly assigned. This process is called *Probing*. The Wellfleet AppleTalk router implements this by first generating a tentative AppleTalk address for the interface in the format:

<start_network_number>.1

where:

<*start_network_number*> is the lowest end of the network number range assigned to the network to which this interface connects.

*1* is the first possible node number that could be assigned to this interface.

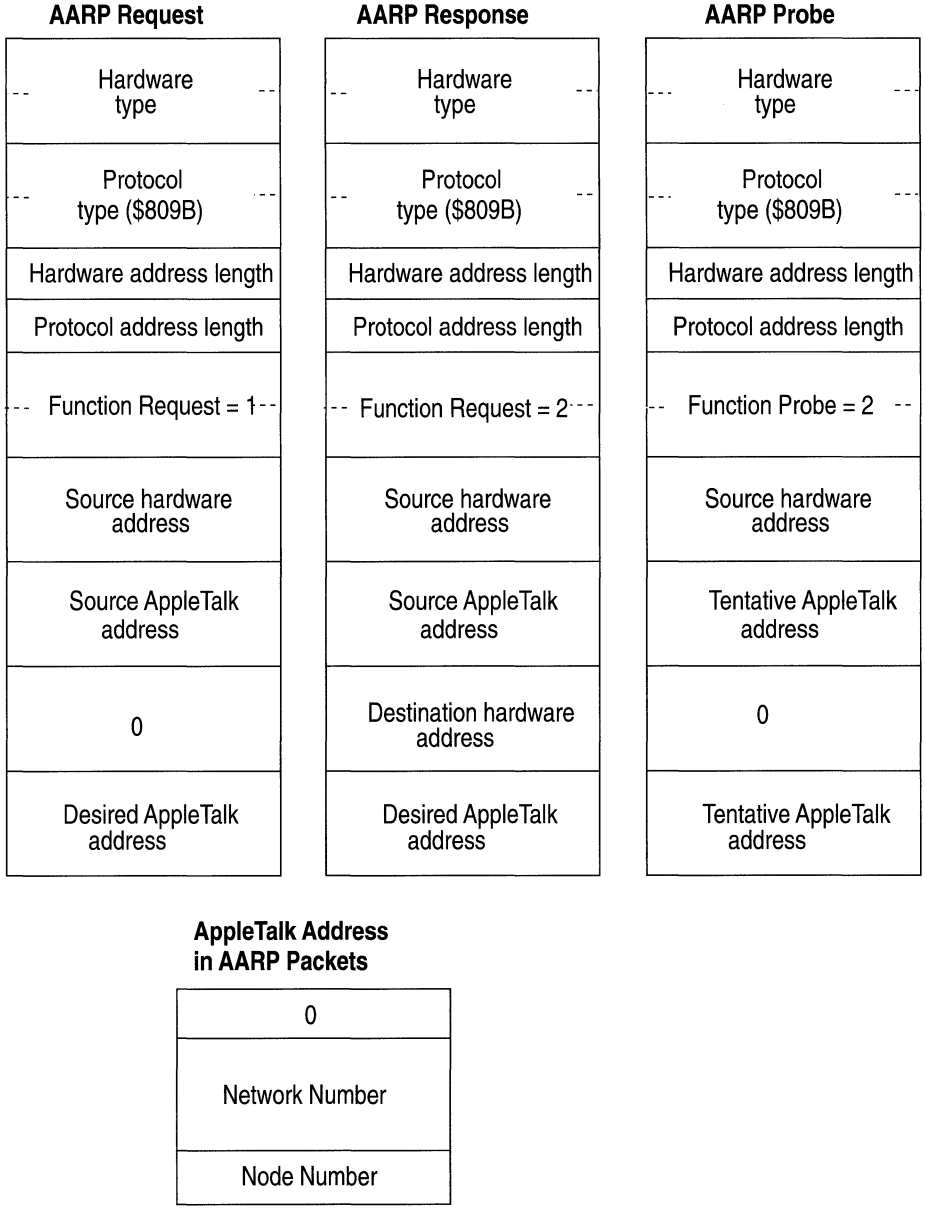


**Figure 6-3. AARP Packets**

Next, the router checks the validity of the address by broadcasting 10 AARP Probe packets containing this address at 0.2 second intervals. AARP Probe packets inquire if any other node on the network is already using this address. If the router doesn't receive a Response, then it knows that the address is unique on the network and assigns the address to the interface. If the router receives a Response (or a Probe for the same address), it knows the address is already in use. So, the router increments the node number by 1, then sends out 10 more Probes. It repeats this process until it does not receive a Response, or runs out of all possible node numbers.

If the router runs out of possible node numbers, it increments the start_network_number by 1 and repeats the entire Probe process. Finally, if the router is still unable to generate a unique address, it logs the error and shuts down the interface (see the section of this chapter entitled *AppleTalk Interface Startup* for a state diagram of this process).

## Datagram Delivery Protocol

The Wellfleet AppleTalk router uses the Datagram Delivery Protocol (DDP) to transmit packets between nodes on the network. The Datagram Delivery Protocol is an unreliable network layer protocol.

An AppleTalk datagram consists of the DDP header, immediately followed by the data. The Wellfleet router encapsulates all packets in an extended 13-byte DDP header (see Figure 6-4).

The source node sets the hop count field to zero before sending a packet out onto the network. Each router that receives the packet increments the hop count by one until it either reaches the destination end node, or reaches the maximum hop count (15), in which case it is discarded.

When the Wellfleet AppleTalk router receives a packet, it checks to see if the packet's destination network number is the local network. If it is, the router passes it down to the data link layer which forwards the packet toward the destination node. If the destination network number is a different network, the router refers to its routing tables to determine the next hop on the shortest path toward the destination.

Finally, the router increments the hop count by one, and forwards the packet toward the next hop. The router's routing tables are maintained using the Routing Table Maintenance Protocol.

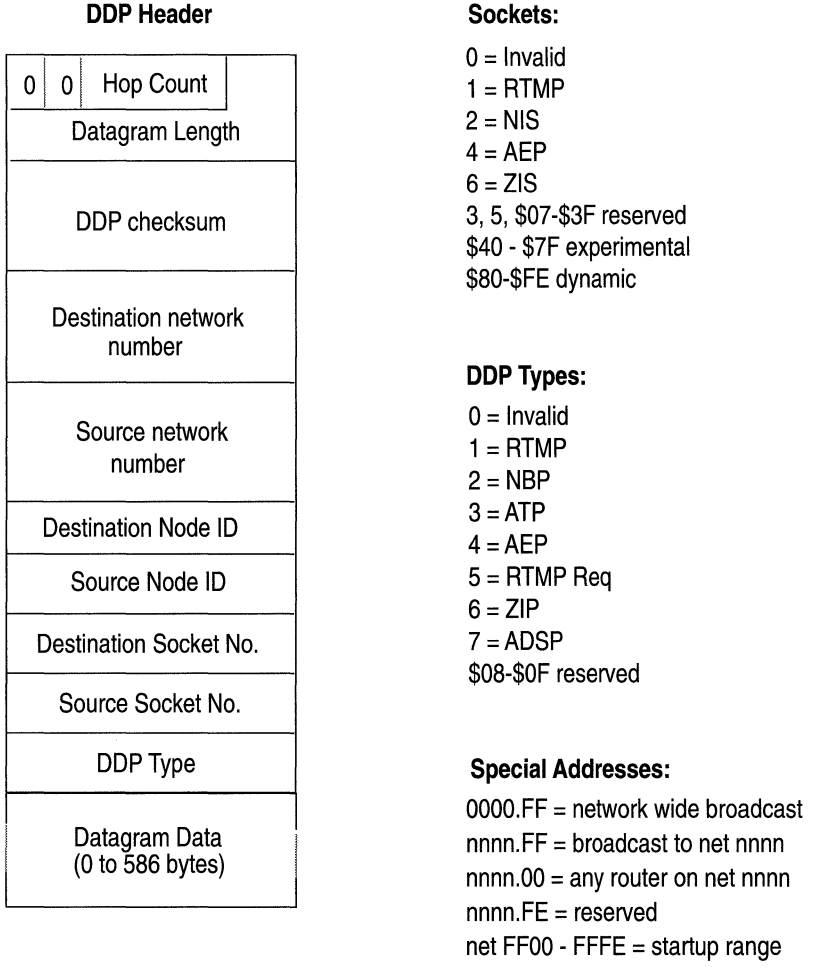


**Figure 6-4. DDP Packet**

## Routing Table Maintenance Protocol

The Wellfleet AppleTalk router uses the Routing Table Maintenance Protocol (RTMP) to create and maintain the routing information DDP uses to transmit packets across an internet. Routing information is contained in the AppleTalk routing table.

Each table entry includes a destination network range, the AppleTalk protocol address (network number and node number) through which the destination is reached, the number of router hops to the destination, and the route status.

Routers create and update their routing tables by periodically constructing and broadcasting RTMP data packets to all other routers on directly connected links. An RTMP data packet contains the source address and the information stored in the originating router's routing table. All routers receiving the data packet use this information to update their own routing tables (see Figure 6-5).

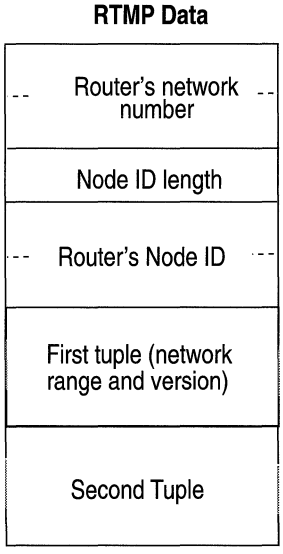


**Figure 6-5. RTMP Packet**

## Zone Information Protocol

The Wellfleet AppleTalk router uses the Zone Information Protocol (ZIP) to map networks to zone names on the internet. The router stores this information in its zone information table (ZIT). The zone information table contains one entry for each network on the internet. The entry is in the format:

<network_start, zone list>

The zone list field consists of a number of text strings that identify the zone names that are specified for that network. The router maintains and updates the ZIT by broadcasting ZIP *Query* packets for zone list information to all other routers in the network. Other routers respond with ZIP *Response* packets with the desired zone lists (see Figure 6-6).

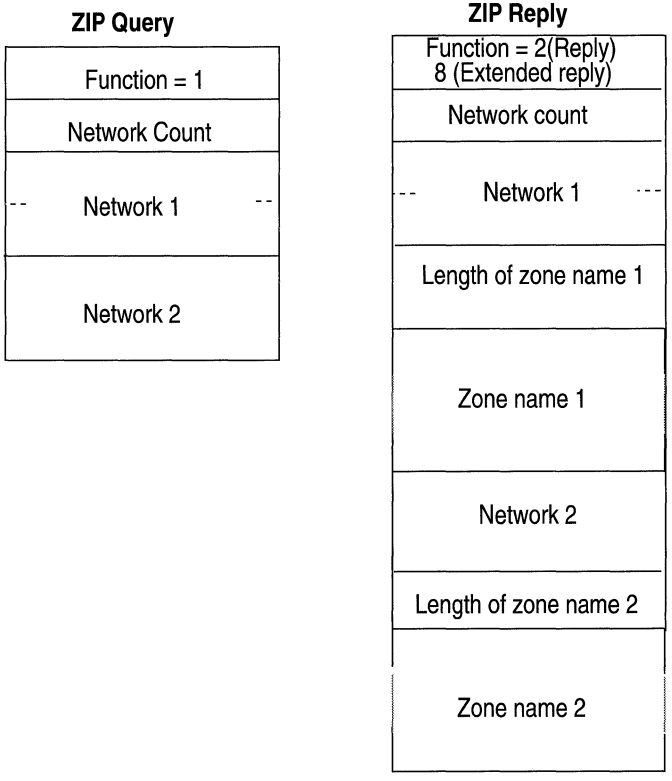


**Figure 6-6. ZIP Packets**

When an AppleTalk node starts up on the network, it requests zone and network information from all routers on the circuit by broadcasting ZIP *GetNetInfo Request* packets (see Figure 6-7). Upon receiving a GetNetInfo Request, the router replies with a *GetNetInfo Response* packet that confirms or denies the validity of the zone name and supplies the correct network range. If the stored zone name is invalid, the node uses the default zone name that is configured for the network instead.

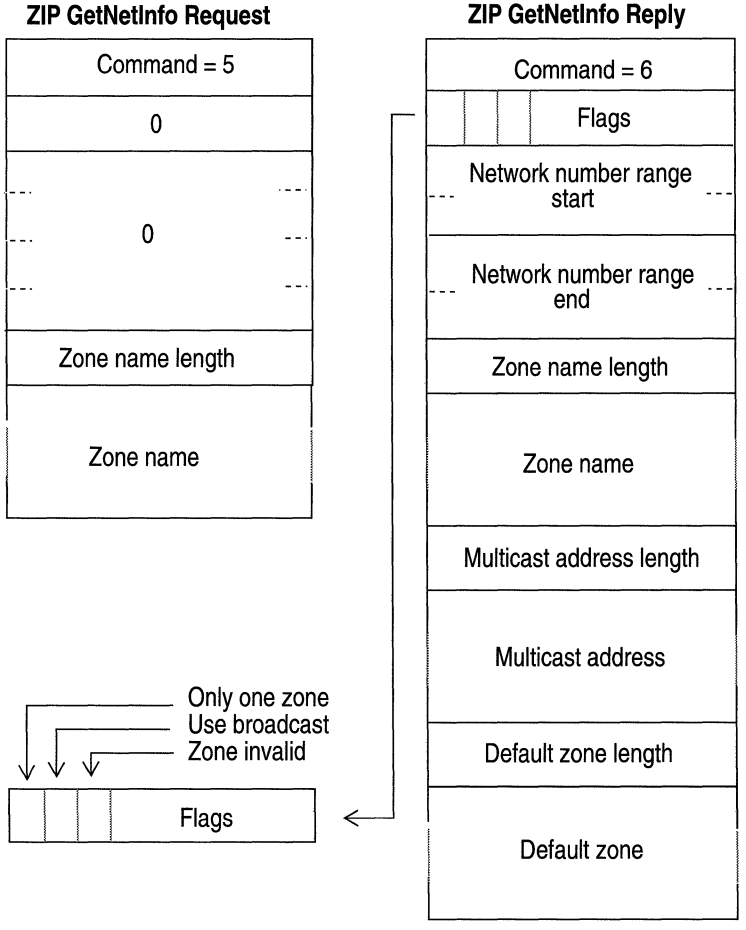


**Figure 6-7. ZIP GetNetInfo Packets**

## Name Binding Protocol

Macintosh users on an AppleTalk network refer to services by name, rather than by physical location on the network. For example, if you wished to locate a printer, you would check the Chooser and select the printer by name, rather than by its AppleTalk address.

The Wellfleet AppleTalk router uses the Name Binding Protocol (NBP), in conjunction with the Zone Information Protocol, to map names of services to AppleTalk addresses. When the router receives an NBP *Broadcast Request* packet for a named entity from an AppleTalk node, it refers to its ZIT to see in which network the requested entity's zone is located. If the zone is on the local network, it broadcasts an NBP *Lookup Request* packet. If the zone is located on a different network, it sends out an NBP *Forward Request* packet toward the router connected to the network with the destination zone. Upon receiving a Forward Request, that router then broadcasts an NBP Lookup Request (see Figure 6-8).

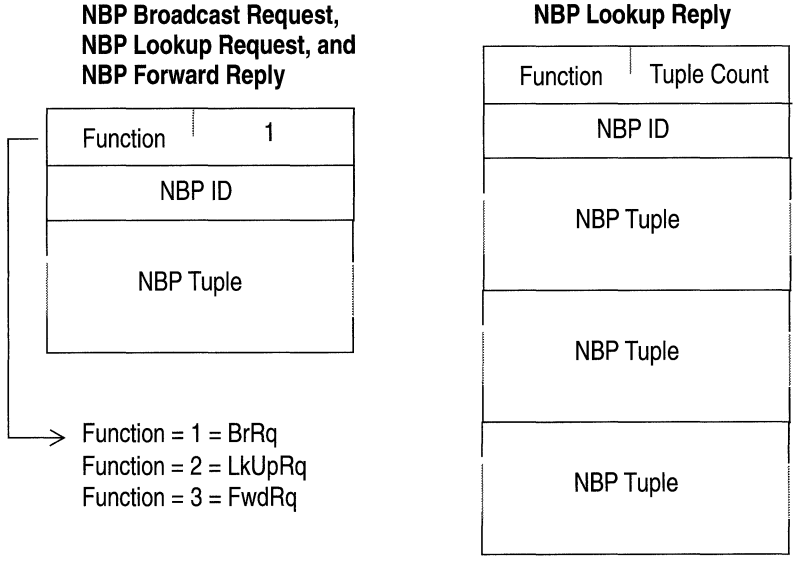


**Figure 6-8. NBP Packets**

An NBP named entity consists of a name, type, and zone name, where:

- *name* is the name assigned to the device which identifies it on the network.
- *type* specifies the device type; for example, the type specified for a LaserWriter printer is identified by the type LaserWriter.
- *zonename* specifies the AppleTalk zone in which the device resides.

## AppleTalk Echo Protocol

The Wellfleet AppleTalk router supports the AppleTalk Echo Protocol, which allows the router to respond to echo packets sent to it by other nodes on the network.

## AppleTalk Interface Startup

The following two pages show a state machine table that describes the events that occur when a Wellfleet AppleTalk interface initializes (that is, moves from a "down" state to an "up" state).

Starting from a down state, both seed and nonseed interfaces begin the initialization process by sending out AARP Probes to verify the tentative AppleTalk address that the router assigned to the interface. Then, depending on the response (event), the interface proceeds to the "next state" indicated in the state machine table, and the router performs the "action" specified by the table. This process continues until the action specifies that the interface should initialize, or shut down.

During a successful initialization, the router performs all of the actions shown in Figure 6-9.

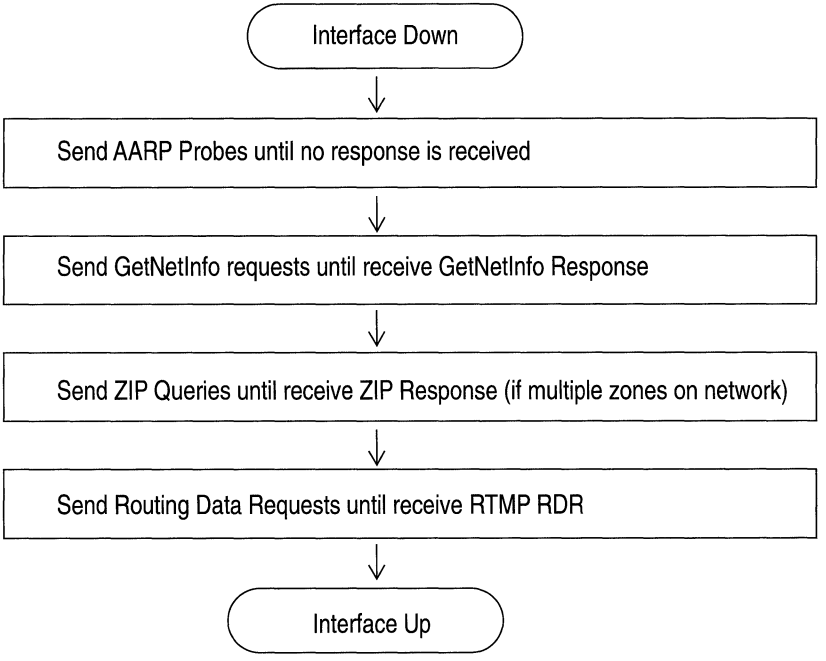


**Figure 6-9. AppleTalk Interface Initialization Process**

State →

| Event ↓ | Down | Wait Probe 1 | Wait GNI 1 | Wait Probe 2 | Wait GNI 2 | Wait ZIP | Wait RTMP | Up |
|---|---|---|---|---|---|---|---|---|
| Start Nonseed | NS=Wait Probe 1 **A=Send Probe** | | | | | | | |
| Start Seed | NS=Wait Probe 2 **A=Send Probe** | | | | | | | |
| T/O with retries > max-retry | | NS=Wait GNI 1 **A=Send GNI** | NS=Wait GNI 1 **A=Send GNI*** | NS=Wait GNI 2 **A=Send GNI** | NS=Down **A=Log**** | NS=Down **A=Log**** | NS=Down **A=Log**** | |
| T/O with retries < max-retry | | NS=Wait Probe 1 **A=Send Probe** | NS=Wait GNI 1 **A=Send GNI*** | NS=Wait Probe 2 **A=Send Probe** | NS=Wait GNI 2 **A=Send GNI** | NS=Wait ZIP **A=Send ZQ** | NS=Wait RTMP **A=Send RDR** | |
| Recv AARP Response with collision | | NS=Down **A=Log** | | NS=Down **A=Log** | | | | |
| Recv Probe with collision | | NS=Wait Probe 1 **A=Send New Probe** | NS=Wait GNI 1 **A=Send AARP Response** | NS=Wait Probe 2 **A=Send New Probe** | NS=Wait GNI 2 **A=Send AARP Response** | NS=Wait ZIP **A=Send AARP Response** | NS=Wait RTMP **A=Send AARP Response** | |
| Recv GNI with no errors | | | NS=Wait Probe 2 **A=Send Probe** | | NS=Wait ZQ **A=Send ZQ or** | | | |
| Recv GNI with errors | | | NS=Down **A=Log** | | NS=Down **A=Log** | | | |

* If a nonseed router does not receive a GNI Response, it will send out GNI Requests forever.

** If this is a seed router, then the interface will initialize (move to the "up" state).

State →

| Event ↓ | Down | Wait Probe 1 | Wait GNI 1 | Wait Probe 2 | Wait GNI 2 | Wait ZIP | Wait RTMP | Up |
|---|---|---|---|---|---|---|---|---|
| Recv ZIP reply with no errors | | | | | | NS=Wait RTMP **A=Send RDR** | | |
| Recv ZIP reply with errors | | | | | | NS=Down **A=Log** | | |
| Recv RTMP data with no errors | | | | | | | NS=Up **A=Init Port** | |
| Recv RTMP data with errors | | | | | | | NS=Down **A=Log** | |
| Recv AARP Request | | | NS=Wait GNI 1 **A=Send AARP Response** | | NS=Wait GNI 2 **A=Send AARP Response** | NS=Wait ZQ **A=Send AARP Response** | NS=Wait RTMP **A=Send AARP Response** | |
| Recv GNI Request | | | | | | NS=Wait ZQ **A=Send GNI Response** | NS=Wait RTMP **A=Send GNI Response** | |

**Key**

NS = Next State
A = Action
T/O = Time Out
ZIP = Zone Information Protocol
GNI = GetNetInfo
RDR = Routing Data Response
AARP= Address Resolution Protocol
Init Port = Intialize interface

**Maximum Retry Counts**

Probes = 10 every .2 seconds
GetNetInfos = 3 every .2 seconds
ZIP Quiries = 3 every .2 seconds
RTMP RDRs = 3 every .2 seconds

= Interface remains in the current state until event occurs that moves it into a different state

# AppleTalk Implementation Notes

This section contains some basic guidelines on adding Wellfleet AppleTalk routers to your network. It also addresses special configuration features that may match your network requirements.

## When Should I Configure My Wellfleet AppleTalk Router as a Seed Router?

When you enable an AppleTalk interface, you must specify the interface router type as either seed or nonseed. Basically, seed routers provide configuration information (network ranges and zone names) to the other AppleTalk routers on the network. When you configure a seed router, you specify the network number range, default zone, and zone list that reflects your network configuration. If there are multiple seed routers on the network, all must be configured with the same values.

**Note:** You *must* configure the router as a seed router if it is the only AppleTalk router on the network.

If a seed router already is configured on this network, you can simply configure this router to be a nonseed router.

For instructions on initially enabling AppleTalk on a circuit, see the section entitled *Defining AppleTalk* in the *Configuring Circuits* chapter of this guide. For instructions on editing an existing AppleTalk interface, see the section entitled *Editing AppleTalk Interface Parameters* later in this chapter.

## Enabling Both AppleTalk Routing and Bridging on the Same Interface

If you enable both AppleTalk routing and bridging on the same interface, then you should configure a bridge filter on the interface that will drop all AppleTalk Phase 2 AARP and DDP packets. If you do not configure a filter, then you may experience problems when you reboot the router. (This is because the bridge enables before the AppleTalk router, causing packets to be bridged before AppleTalk routing starts.)

Wellfleet suggests creating a filter that specifies a Drop (with no log) action on all SNAP packets with Ethertype ranges of 0x809B - 0x809B and 0x80F3 - 0x80F3. The hierarchial menu for the field is as follows:

Data link - 802.2 SNAP - 802.2 SNAP Ethertype

See the chapter of this guide entitled *Configuring Filters* for more information about how to construct filters for the bridge.

## How Can I Reduce Excess Routing Traffic on a Large AppleTalk Network?

If you are adding Wellfleet AppleTalk routers to a large sized internet, (for example, one that contains more than 200 routers and networks), you may wish to consider the following network configuration tips in order to reduce the amount of routing traffic on your network:

- Try to keep the physical network topology as hierarchial as possible. For example, using a functional (organizational) hierarchy isolates groups on the network and reduces the amount of excess broadcast traffic.
- Try to reduce the number of devices on a single physical network (or bridged network) by dividing the internet into a greater numbers of networks. This also helps to reduces the amount of traffic that is broadcast to all nodes on the network.
- Limit amount of Name Binding Protocol traffic on network; avoid configuring the same zone on multiple networks. Instead, keep the ratio of zones to networks as close to 1 - 1 as possible.

## Adding or Removing AppleTalk Zones

When a router initially learns about a new network, it sends out query packets requesting the network's zone information. The router uses this information to update its Zone Information Table.

However, because routers only query for this information the *first* time they learn about a network, any changes later made to the network's zone list will not be propagated to other routers on the internet.

So, to add or remove an AppleTalk zone from an AppleTalk interface's zone list, follow one of the following procedures.

### Editing Both Network and Zone List Information for an Interface

If you are editing both the network parameters for an interface (the Network ID, Network Start, and Network End) and the interface's zone list, then you can simply do the following:

1. Disable the interface by setting the Port Enable parameter to No.
2. Edit the Network ID, Network Start, and Network End parameters for the interface.
3. Edit the interface's Zone List parameter.
4. Reenable the interface by setting the Port Enable parameter to Yes.

See the section of this chapter entitled *Editing AppleTalk Interface Parameters* for instructions.

### Editing Only Zone List Information for an Interface

If you are only editing an interface's zone list, then in order to ensure that all routers on the internet update their ZITs, you must do the following:

1. Disable *all* AppleTalk router interfaces (both Wellfleet and non-Wellfleet) that connect to the network for which the zone name change is being applied.

For example, in Figure 6-10, to add the new zone name Blue to the zone list for Network 1-3, AppleTalk interfaces 1.1, 1.2, and 1.3 are disabled.

You can disable each Wellfleet AppleTalk interface by setting its Port Enable parameter to No.

2. Edit the interface's Zone List parameter.
3. Wait at least 10 minutes.

   This allows all routers on the AppleTalk internet to age-out the network and zone information for the network from their ZITs.
4. Reenable the router interfaces that connect to the network.

   You can reenable each Wellfleet AppleTalk interface by setting its Port Enable parameter to Yes.

See the section of this chapter entitled *Editing AppleTalk Interface Parameters* for instructions.

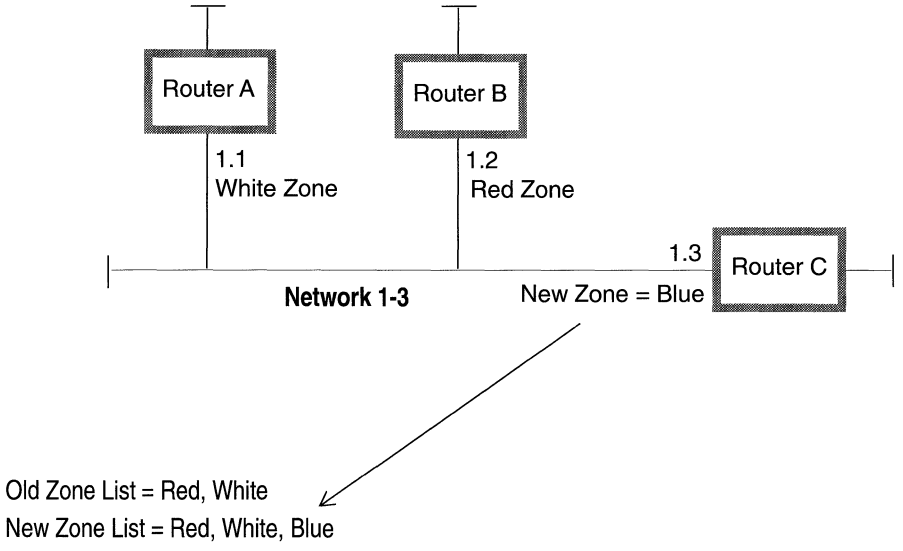


**Figure 6-10. Updating an AppleTalk Router's Zone List**

## Can I Add My Wellfleet AppleTalk Router to a Transition Network?

You can add the Wellfleet AppleTalk router to a transition network (one that generates both AppleTalk Phase 1 and AppleTalk Phase 2 traffic), however, for performance purposes, Wellfleet does not recommend it.

If you chose to do so, you must consider the following configuration requirements:

- An AppleTalk Phase 1 to Phase 2 transition router must reside on the internet.
- All defined network ranges must consist of a single number; for example, 1-1 is a valid network range, while 1-5 is invalid.
- You can only specify a single zone name for each network.

## Configuring the AppleTalk Router to Source Route Over Token Ring Networks

The Wellfleet AppleTalk router supports routing over token ring networks that contain one or more source routing bridges.

In a source routing network, every end station that sends out a frame supplies the frame with the necessary route descriptors so that it can be source routed across the network. Thus, in order for routers to route packets across a source routing network, *they must act like end stations*; supplying route descriptors for each packet before they send it out onto the network.

With end node support enabled, whenever a Wellfleet AppleTalk router receives a packet and determines that the packet's next hop is located across a source routing network, the router does the following:

- Adds the necessary RIF information to the packet's MAC header.
- Sends the packet out onto the network where it is source routed toward the next hop.

Upon receiving the packet from the token ring network, the peer router strips off the RIF field and continues to route the packet toward the destination network address (see Figure 6-11).

You configure source route end node support for each individual routing protocol on a per-circuit basis. See the section entitled *Editing AppleTalk Interface Parameters* in this chapter for instructions.

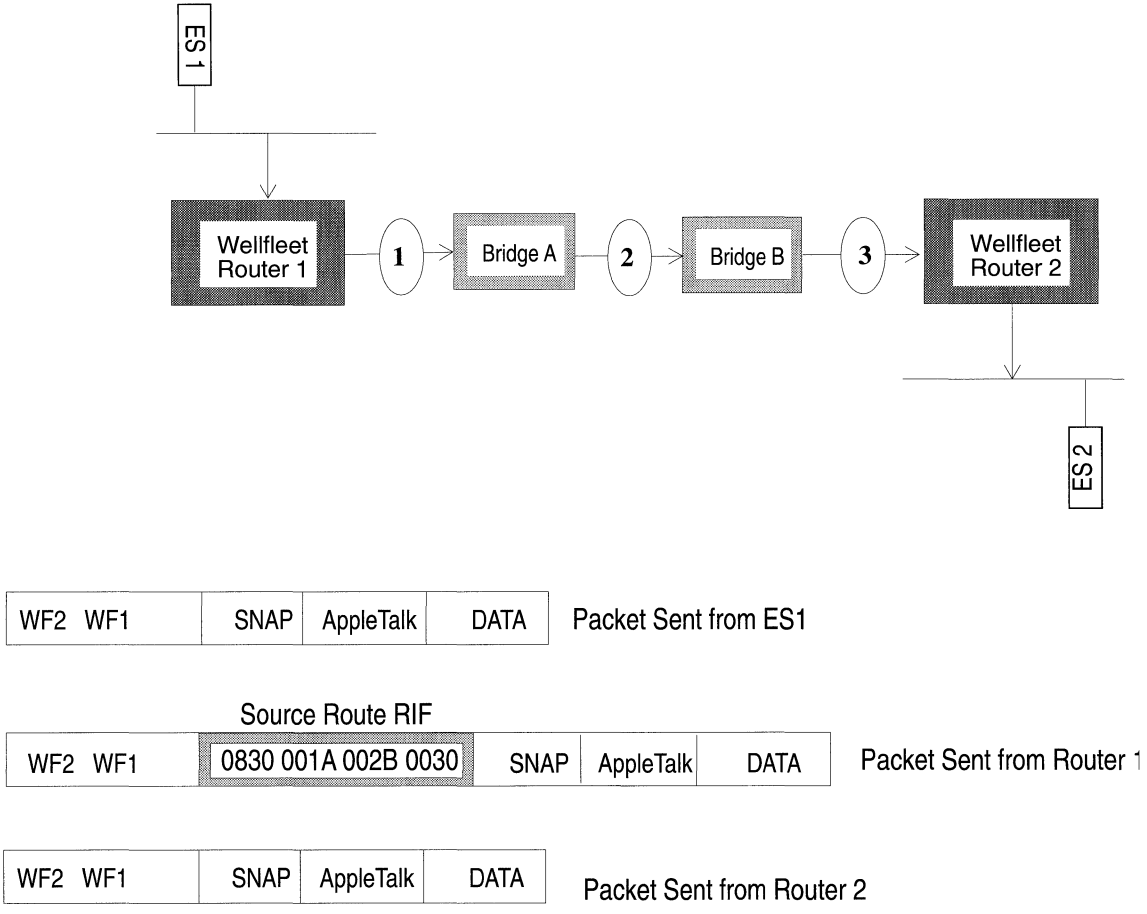


**Figure 6-11. AppleTalk Routers Source Routing Across a Token Ring Network**

# Editing AppleTalk Parameters

Once you have configured a circuit to support AppleTalk, you can use the Configuration Manager to edit AppleTalk parameters. The configuration function you wish to perform determines the type of parameters you must edit. Table 6-1 lists each configuration function and the section that describes how to perform the function.

**Table 6-1. AppleTalk Parameters and Configuration Functions**

| To Do the Following: | See this Section: |
|---|---|
| Change the state of the AppleTalk router software | *Editing AppleTalk Global Parameters* |
| Reconfigure AppleTalk on a particular circuit. | *Editing AppleTalk Interface Parameters* |

For each AppleTalk parameter, this section provides the following:

- Wellfleet default
- Valid setting options
- Parameter function
- Instructions for setting the parameter

You begin from the Wellfleet Configuration Manager window (see Figure 6-12).

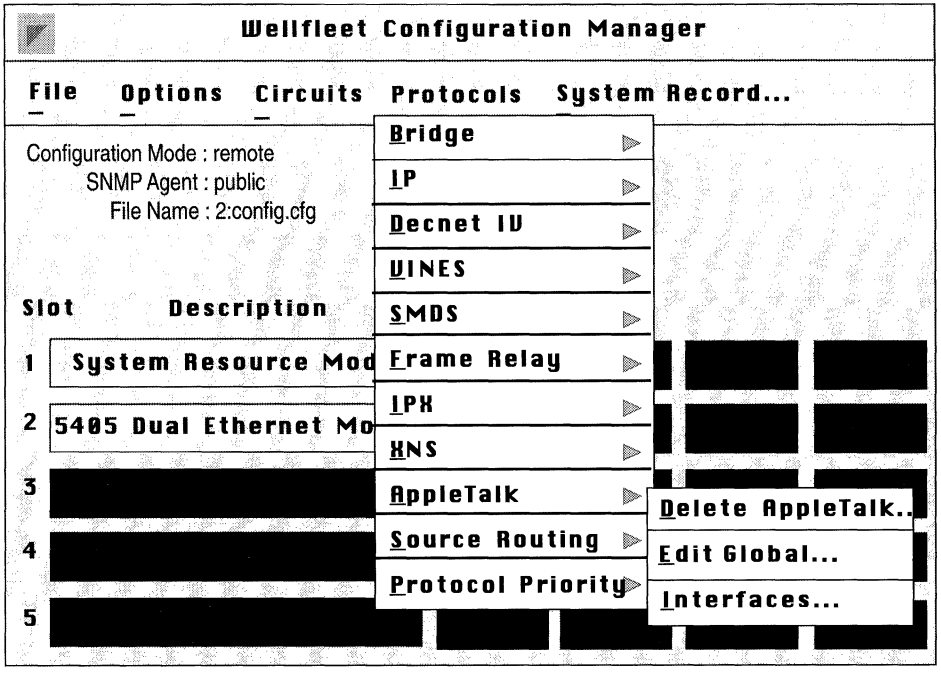


**Figure 6-12. Wellfleet Configuration Manager Window**

## Editing AppleTalk Global Parameters

To edit AppleTalk Global parameters, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols/AppleTalk/Edit Global option.

   The AppleTalk Global Parameters window appears (see Figure 6-13).

2. Edit those parameters you wish to change.

   Click the Save button to exit the window and save your changes when you are finished.

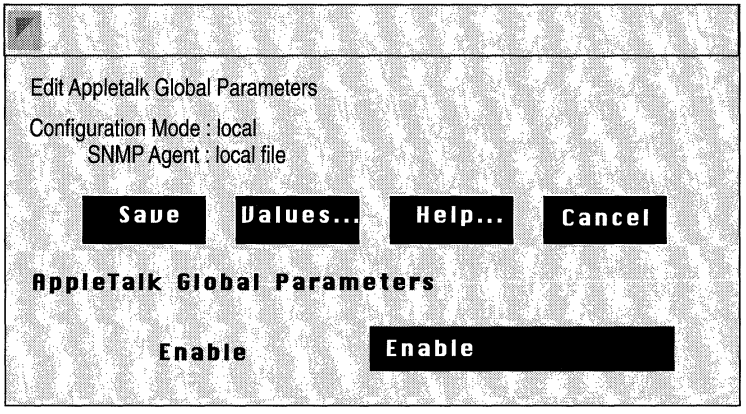


**Figure 6-13. AppleTalk Global Parameters Window**

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or Disables the AppleTalk router on the entire BN. |
| Instructions: | Set to Disable if you want to disable AppleTalk. |

## Editing AppleTalk Interface Parameters

To edit an AppleTalk interface, begin at the Wellfleet Configuration Manager window then proceed as follows:

1. Select the Protocols/AppleTalk/Interfaces option to display the AppleTalk Interfaces window (see Figure 6-14).

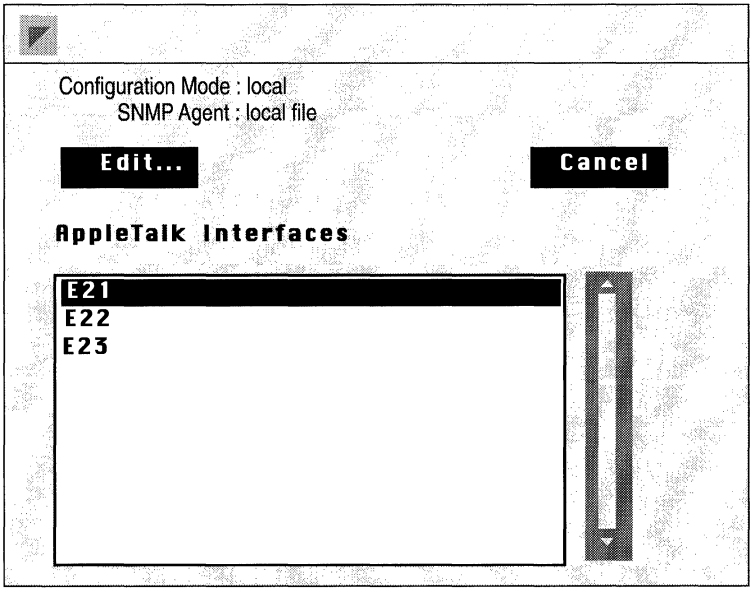


**Figure 6-14. AppleTalk Interfaces Window**

2. Select the interface you wish to edit.
3. Click on the Edit button to display the AppleTalk Interface Parameters window for that interface (see Figure 6-15).
4. Edit those parameters you wish to change.
5. Select the File/Done option save your changes and exit the window when you are finished.

**Note:** When you reconfigure an interface in dynamic configuration mode, AppleTalk restarts on that interface.

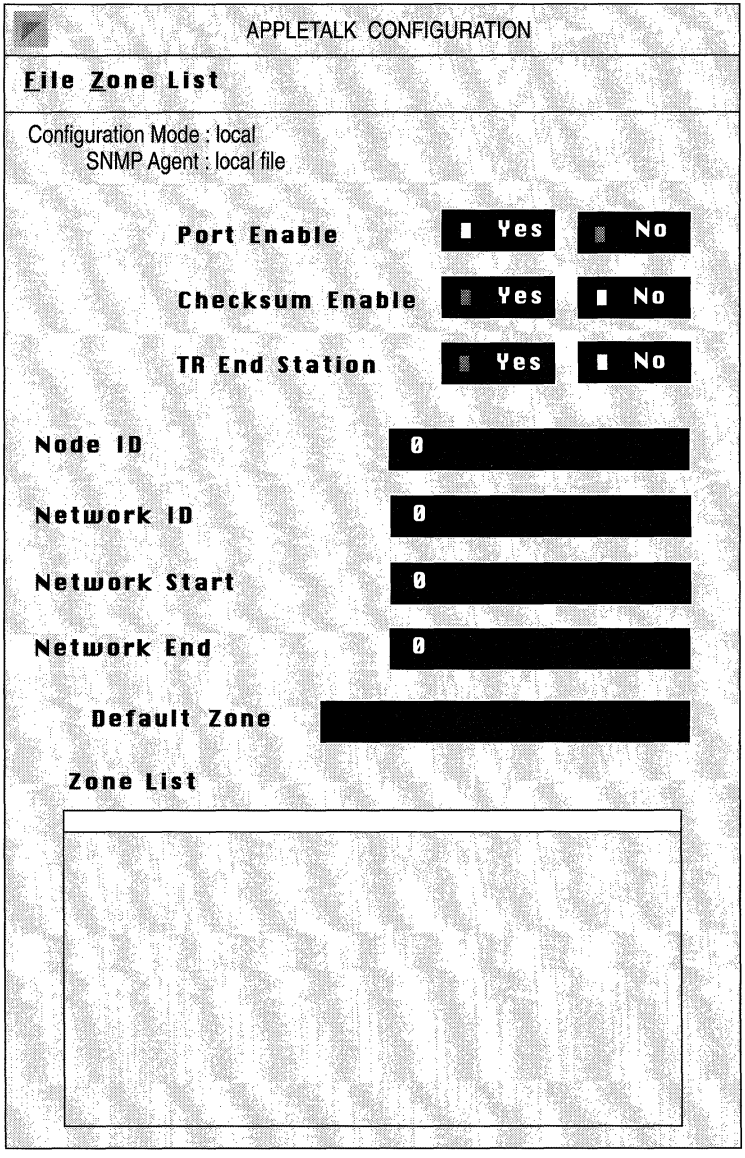


**Figure 6-15. AppleTalk Interface Parameters Window**

| | |
|---|---|
| **Parameter :** | **Port Enable** |
| Wellfleet Default: | Yes |
| Options: | Yes/No |
| Function: | Enables or disables AppleTalk over this interface. |
| Instructions: | Set to No only if you want AppleTalk disabled over this interface. (For example, if you are editing the interface's zone list). |

| | |
|---|---|
| **Parameter :** | **Checksum Enable** |
| Wellfleet Default: | No |
| Options: | Yes/No |
| Function: | Enable or disables the calculation of the DDP checksum for packets *originated* by the AppleTalk router on this interface.<br><br>With checksumming enabled, the AppleTalk router calculates and includes a 16-bit checksum in the header of DDP datagrams that it originates and sends out onto the network. With checksumming disabled, the router does not calculate a checksum, but instead inserts a value of 0 into DDP datagrams that it originates.<br><br>This parameter has no affect on incoming packets (those packets *received* over this interface). If the router receives a packet containing a checksum, it verifies the checksum value. |
| Instructions: | Wellfleet recommends accepting the default value No. However, if this circuit contains excess noise, you may chose to enable checksumming to insure packet integrity. |

| | |
|---|---|
| **Parameter :** | **TR End Station** |
| Wellfleet Default: | No |
| Options: | Yes/No |
| Function: | Specifies if this interface is enabled for source routing end station support. |
| Instructions: | Set to Yes if this interface is 1) a token ring circuit and 2) if source routing is enabled on the AppleTalk routers in this ring. |

| | |
|---|---|
| **Parameter:** | **Node ID** |
| Wellfleet Default: | 0 |
| Options: | 0-253 |
| Function: | Identifies the node number assigned to this interface. |
| Instructions: | Wellfleet recommends accepting the default 0. When you accept the default 0, the router dynamically acquires a Node ID for the interface during startup - thus ensuring that the AppleTalk address for this interface is unique within the network.<br><br>If change the default value, the router uses the Node ID you specify. |

| | |
|---|---|
| **Parameter:** | **Network ID** |
| Wellfleet Default: | 0 |
| Options: | 0- 65279 |
| Function: | Identifies the network number of the network to which this interface connects. |
| Instructions: | Wellfleet recommends accepting the default 0. When you accept the default 0, the router dynamically acquires a Network ID for the interface during startup - thus ensuring that the AppleTalk address for this interface is unique within the network.<br><br>If you change the default value, make certain that the number you specify is within the correct network range. The router then uses the Network ID that you specify. |

**Note:** If the AppleTalk address (Node ID.Network ID) matches that of any other node on the internet, the interface automatically disables.

**Parameter :** **Network Start**

Wellfleet Default: 0

Options: 0 - 65279

Function: Specifies the lowest boundary (minimum) of the range of network numbers that are available for use by nodes on the network to which this interface connects.

This parameter's setting determines whether or not this interface functions as a seed or nonseed router. A seed router supplies the Network Start, Network End, Default Zone and Zone List information for all other nonseed routers on this network. A nonseed router acquires its Network Start, Network End, Default Zone and Zone List information from the other seed routers on the network.

Each network must contain at least one seed router.

Instructions: To configure this interface as a nonseed router, simply accept the default 0.

To configure this interface as a seed router, specify the Network Start as follows:

— If this is the only seed router on the network, check your network topology map to see how your network is divided and enter the lowest boundary network number here.

— If there are already seed routers on the network, enter the *same* Network Start value that is configured on all other seed routers.

**Note:** If you specify a Network Start other than the default 0, then 1) the router becomes a seed router automatically and 2) you must also specify values for the Network End and Default Zone parameters.

| | |
|---|---|
| **Parameter :** | **Network End** |
| Wellfleet Default: | 0 |
| Options: | 0 - 65279 |
| Function: | Specifies the upper boundary (maximum) of the range of network numbers that are available for use by nodes on the network to which this interface connects.<br><br>This parameter is used in conjunction with the Network Start parameter to help define a seed router. *If you have not specified a Network Start, this parameter is ignored.* |
| Instructions: | If this interface is configured as a nonseed router, then simply accept the default 0.<br><br>If this interface is configured as a seed router, then specify the Network End as follows:<br><br>— If this is the only seed router on the network, check your network topology map to see how your network is divided and enter the upper boundary network number here.<br><br>— If there are already seed routers on the network, enter the *same* Network End value that is configured on all other seed routers. |

| | |
|---|---|
| **Parameter :** | **Default Zone** |
| Wellfleet Default: | None. |
| Options: | Any valid zone name. |
| Function: | Specifies the name of the default zone where all new nodes are assigned when they first start up on this network.<br><br>This parameter is used in conjunction with the Network Start and Network End parameters to help define a seed router. *If you have not specified a Network Start, this parameter is ignored.* |
| Instructions: | If this interface is configured as a nonseed router, then simply leave this field blank.<br><br>If this interface is configured as a seed router, then specify the Default Zone as follows:<br><br>— If this is the only seed router on the network, enter any valid Default Zone name.<br><br>— If there are already seed routers on the network, enter the *same* Default Zone name as is configured on all other seed routers.<br><br>A valid zone name can consist of up to a maximum of 32 characters and can include any keyboard character (except the * character). |

| | |
|---|---|
| **Parameter :** | **Zone List** |
| Wellfleet Default: | None |
| Options: | Any valid zone name. A valid zone name contains from 1-32 characters and can include any character except for the * character. |
| Function: | Identifies the other zones (besides the default zone) that are specified for this network. |
| Instructions: | To add a zone name to the interface's zone list, select the Zone List/Add option. Enter the zone name, then click on the Add Zone button. Repeat this for each zone you wish to add. When you are finished, select the File/Done option to save your changes.<br><br>To delete a zone name from the interface's zone list, first select a zone from those displayed, then select the Zone List/Delete option. Click on the Delete button. Repeat this for each zone you wish to delete. When you are finished, select the File/Done option to save your changes. |

**Warning:** Before you edit the AppleTalk zone list for an interface, read the section of this chapter entitled *Adding or Removing AppleTalk Zones*. It describes the impact of changing an interface's zone list, and notes when you need to wait at least 10 minutes before reenabling the interface.

# Deleting AppleTalk from the BN

You can delete AppleTalk routing protocol from all BN circuits on which it is currently enabled in two steps.

You begin from the Wellfleet Configuration Manager window as follows:

1. Select the Protocols/AppleTalk/Delete AppleTalk option.

   A window pops up and prompts "Do you REALLY want to delete AppleTalk?".

2. Select OK.

   You are returned the Wellfleet Configuration Manager window. AppleTalk is no longer configured on the BN.

If you examine the Wellfleet Configuration Manager window, you will see that the connectors for circuits on which AppleTalk was the *only* protocol enabled are no longer highlighted. Circuits must be reconfigured for these connectors; see chapter 3, entitled *Configuring Circuits* for instructions.

# Chapter 7

## Configuring the Bridge

## List of Figures



## List of Tables

# Configuring the Bridge

## About This Chapter

This chapter tells you how to configure parameters for the Transparent/Translating Bridge and for the Spanning Tree Algorithm.

Use this chapter if you want to:

- Enable or disable bridging for the entire BN, or for specific interfaces
- Enable or disable the spanning tree algorithm for the entire BN, or for specific interfaces
- Edit spanning tree parameters

The first part of this chapter provides you with information about bridging, types of bridges, the spanning tree algorithm, and filtering. If you prefer not to read the overview, go directly to the section entitled *Editing Parameters*.

# Bridge Overview

Bridges are data-link layer relay devices that connect two or more networks and use Media Access Control (MAC) source and destination addresses to relay frames between connected networks. Release 7.00 provides two types of service: transparent bridging and translating bridging. The Transparent/Translating Bridge is described in the following sections.

**Note:** For simplicity, the term *bridge* is sometimes used in the place of the term *Transparent / Translating Bridge*.

## The Transparent/Translating Bridge

The Transparent/Translating bridge provides network interconnection and/or extension services for LANs at the data-link and physical layers.

The Transparent/Translating bridge provides two types of service: transparent bridging and translating bridging. When the bridge connects LANs using the *same* data-link level protocol, it performs *transparent* bridging. When the bridge connects LANs using *different* data-link level protocols, it performs *translating* bridging. The bridge automatically determines which type of service is necessary.

Neither bridging service places any burden on end-stations. From their point of view, it appears that all end-stations are resident on a single extended network with each station identified by a unique MAC-level address.

When the bridge performs *transparent* bridging, it provides three primary services:

- Learning the addresses of end-stations on connected networks
- Forwarding or dropping frames based on the acquired knowledge of end-station addresses or user-configured filters
- Ensuring a loop-free topology throughout the extended network through the use of the spanning tree algorithm

When the bridge performs *translating* bridging, it provides the same three primary services. However, it also provides one additional service; it *translates* into the appropriate frame format all frames that are to be forwarded to a LAN that uses a different data-link level protocol. This translation process is discussed in a subsequent section.

### How the Bridge Works

The bridge receives and examines every frame transmitted on the networks to which it is attached. It learns end-station addresses by reading the source address of the end-station that transmitted the frame, and noting from which LAN the frame was received. The bridge enters this information into a data structure called the forwarding table, which the bridge constantly updates. For example, in Figure 7-1, if Node 2 transmits a frame onto the network, the bridge receives the frame and makes an update to its forwarding table to indicate that Node 2 is in the direction of LAN A.

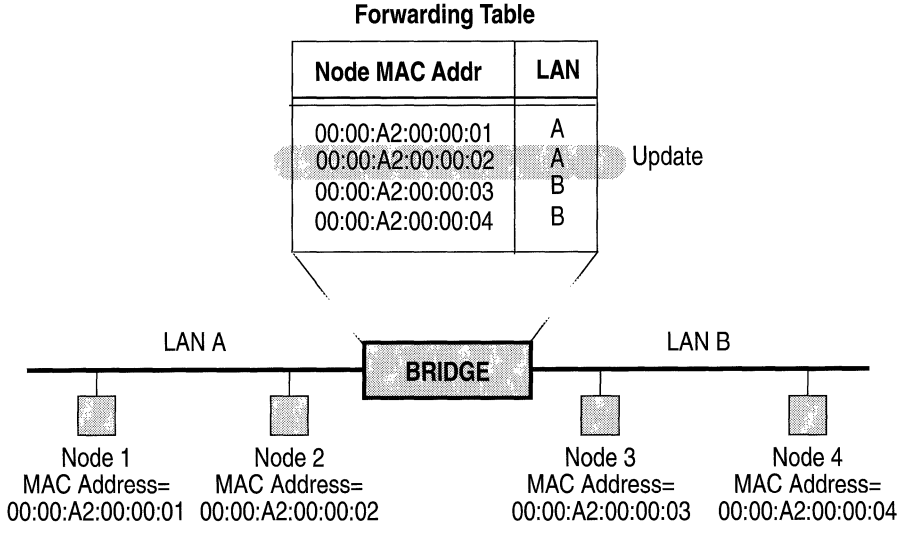


**Figure 7-1. Forwarding Table Update**

The bridge either forwards (relays to another network) or drops (discards) a frame on the basis of the forwarding table entries. When the bridge receives a frame, it compares the frame's destination address with addresses in the forwarding table, and one of the following situations results:

- If the frame's destination address is on the same LAN as its source address, the frame is discarded, since all nodes on that LAN have already received this frame.
- If the frame's destination address is on a different LAN than its source address, the frame is forwarded to that LAN.
- If there is no match for the frame's destination address in the forwarding table, the bridge forwards the frame onto all networks except for the one from which the frame was received. This is called flooding.
- If the frame is destined for the bridge (for example, a spanning tree frame), it is forwarded to the appropriate bridge entity and processed internally.

If a frame is to be flooded or forwarded (as in the two middle cases) to a LAN using a *different* data-link level protocol, the bridge translates the frame to the appropriate frame format before transmission.

### The Translation Process

If the bridge is forwarding a frame between LANs that do not use the same data-link level protocol (for example, FDDI and Ethernet), it converts the outgoing frame to the MAC frame format of the LAN onto which it will be transmitted; this is called translation.

Ethernet, 802.3 and FDDI have different MAC frame formats. When Ethernet or 802.3 frames are bridged to an FDDI LAN, they are reformatted to the FDDI MAC frame format. In this case, the original MAC type of the frame is set in the LLC header. Therefore, if the frame passes through a second FDDI-to-Ethernet/802.3 bridge, it can be translated back to its original format.

The Transparent/Translating Bridge translates frames for bridging between Ethernet/802.3 LANs and FDDI LANs.When bridging from an Ethernet/802.3 LAN to an FDDI LAN, a frame will fall into one of two categories: Ethernet MAC frames or IEEE 802.2 LLC frames.

- All Ethernet MAC frames are translated to FDDI MAC and IEEE 802.2 LLC/SNAP encapsulation as specified by RFC 1042. Protocols in this category include: DECnet Phase IV, Novell, Apple Talk Phase I and II, XNS, IPX, and IP. Figure 7-2 illustrates the format and values of the LLC and SNAP headers of these frames after translation.

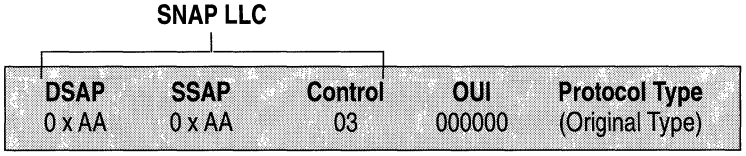


**Figure 7-2. RFC 1042 Encapsulation**

The one exception to this rule is Ethernet frames with a protocol type equal to 80F3 (Apple Talk ARP). Apple Talk ARP frames require special translation through a service called the Bridge Tunnel Service.

Figure 7-3 illustrates the LLC and SNAP headers of the Apple Talk ARP outbound frame after Bridge Tunnel Service translation:

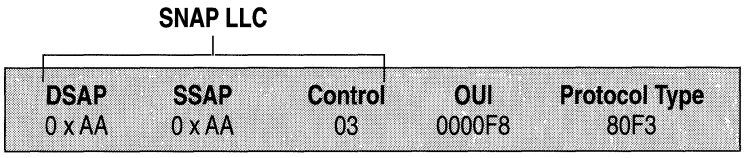


**Figure 7-3. Bridge Tunnel Service Encapsulation**

- All IEEE 802.2 LLC frames will be translated *simply by removing the length field*. Protocols in this category include the following: Apple Talk Phase 2, Novell Proprietary, and IP.

  Figures 7-4 and 7-5 illustrate disparate LAN configurations and show how different types of frames originating on LAN 1 are translated by Bridge A for transmission across the FDDI LAN.

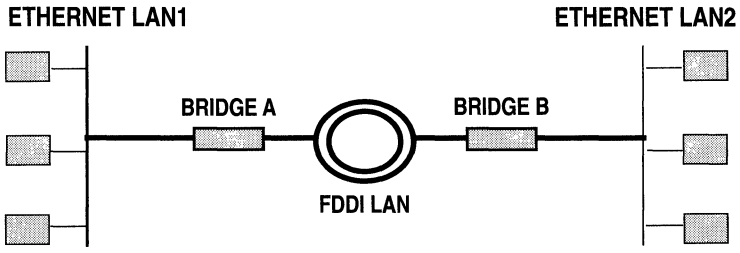


### Ethernet To FDDI Translation

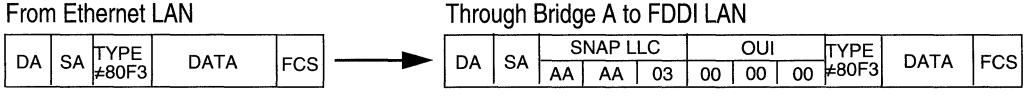


Bridge A:

- extracts addressing information from the Ethernet header
- incorporates address information into newly generated FDDI MAC header
- encapsulates Ethernet data according to RFC 1042
- recalculates FCS

### Ethernet (Apple Talk ARP) To FDDI Translation

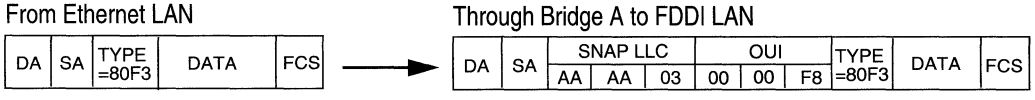


Bridge A:

- extracts addressing information from the Ethernet header
- incorporates address information into newly generated FDDI MAC header
- encapsulates Ethernet within a SNAP header with an Organizationally Unique Identifier (OUI) of 0000F8. This identifies a tunnel-encapsulated frame, specified in IEEE Draft Recommended Practice 802.1H.
- recalculates FCS

**Figure 7-4. Ethernet to FDDI Translation**

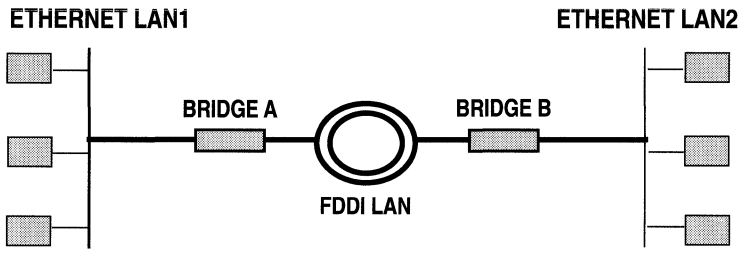


**802.3 to FDDI Translation**

From 802.3 LAN

| DA | SA | LEN | DSAP | SSAP | CTL | DATA | FCS |
|---|---|---|---|---|---|---|---|

Trough Bridge a to FDDI LAN

| DA | SA | DSAP | SSAP | CTL | DATA | FCS |
|---|---|---|---|---|---|---|

Bridge A:

- extracts addressing information from the 802.3 header
- incorporates address information into newly generated FDDI MAC header (with no length field)
- encapsulates 802.3 data within FDDI frame
- recalculates FCS

**Figure 7-5. 802.3 to FDDI Translation**

The translation process from the FDDI LAN to Ethernet/802.3 LAN (the process Bridge B in Figures 7-4 and 7-5 performs) is a mirror image of the translation process occurring on Bridge A, with one exception: an Apple Talk ARP frame *that originated on the FDDI LAN*, and is destined for an 802.3 LAN, will be translated as illustrated in Figure 7-6.

Apple Talk ARP frame originating on FDDI LAN

| DA | SA | SNAP LLC | | | OUI | TYPE | DATA | FCS |
|---|---|---|---|---|---|---|---|---|
| | | AA | AA | 03 | = 00 00 00 | =80F3 | | |

Through Bridge B to 802.3 LAN

| DA | SA | LEN | SNAP LLC | | | OUI | TYPE | DATA | FCS |
|---|---|---|---|---|---|---|---|---|---|
| | | | AA | AA | 03 | = 00 00 00 | =80F3 | | |

Bridge B:

- extracts addressing information from the FDDI MAC header
- incorporates address information into newly generated 802.3 header
- adds a length field
- encapsulates FDDI data according to RFC 1042
- recalculates FCS

**Figure 7-6. Apple Talk ARP (Originating on FDDI) to 802.3 Translation**

## Spanning Tree Algorithm

The spanning tree algorithm (fully described in *IEEE 802.1 D MAC Bridges*) ensures the existence of a loop-free topology in networks that contain parallel bridges. A loop occurs when there are alternate routes between hosts. If there is a loop in an extended network, traffic can be forwarded indefinitely. This can result in increased traffic and degradation in network performance.

Figure 7-7 shows an example of a network containing a loop: LANs A and B are connected by two parallel bridges, Bridge 1 and Bridge 2.

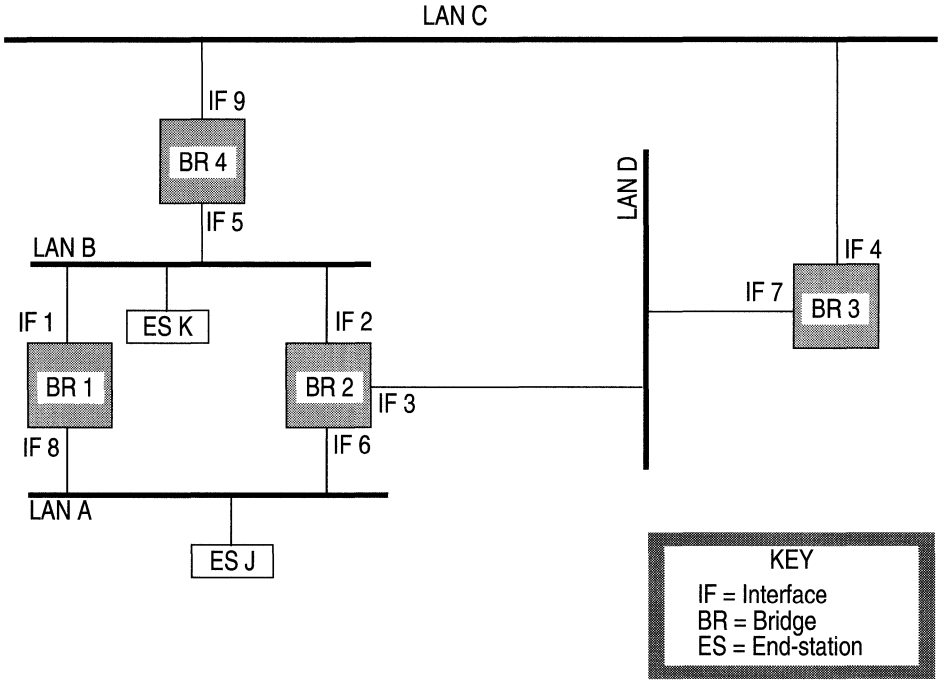


**Figure 7-7. Parallel Bridge Topology**

Consider the chain of events when End-Station J first sends a frame to End-Station K. The frame is read by both Bridge 1 and Bridge 2. Since this is the first frame between J and K, there is no forwarding table entry for J or K on either of the bridges.

Each bridge updates its forwarding table to indicate that End-Station J is in the direction of LAN A. Then, each bridge floods the frame: Bridge 1 forwards the frame over Interface 1 and Bridge 2 forwards the frame over Interface 2. Bridge 2 also forwards the frame over Interface 3; however, to simplify the example, this frame will not be traced.

Next, End-Station K receives two copies of the frame. The reception of duplicate frames by an end-station is not generally fatal. At best, such duplication represents an inefficient use of available bandwidth.

More serious, however, is the effect of duplicate frames on the two bridges. The frame flooded by Bridge 1 onto Interface 1 is ultimately read by Bridge 2 on Interface 2. When Bridge 2 reads this frame, it updates its forwarding table to indicate that End-Station J is in the direction of LAN B. Similarly, Bridge 1 reads the frame flooded by Bridge 2, and updates its forwarding table to show that End-Station J is in the direction of LAN B. Consequently, the forwarding tables of both bridges are now corrupted and neither bridge can properly forward a frame to End-Station J.

This problem can be solved by implementing the spanning tree algorithm which produces a logical tree topology out of any arrangement of bridges. The result is that only a single path exists between any two end-stations on an extended network. It also provides a high degree of fault tolerance by allowing for the automatic reconfiguration of the spanning tree topology in the face of bridge or data-path failure. Five values are required for derivation of the spanning tree topology. The first, a multicast address specifying all bridges on the extended network, is media-dependent and automatically determined by the software. The remaining four are management-assigned values:

- a network-unique identifier for each bridge on the extended network
- a unique identifier for each bridge/LAN interface (called a port)
- a priority specifying the relative priority of each port
- a cost for each port

With these values assigned, bridges multicast and process formatted frames (called Bridge Protocol Data Units or BPDUs) to derive a single loop-free topology throughout the extended network. BPDU frame exchange is accomplished quickly, thus minimizing the time during which service is unavailable between hosts.

In constructing a loop-free topology, the bridges within the extended network follow these steps:

1. The bridges first elect a root bridge.

   The bridge with the lowest priority value becomes the root bridge and serves as the root of the loop-free topology. If priority values are equal, the bridge with the lowest Bridge MAC Address becomes the root bridge.

2. The bridges determine path costs.

   The path cost is the cost of the path to the root bridge offered by each bridge port.

3. The bridges determine a root port and elect a designated bridge on each LAN.

   Each bridge designates the port that offers the lowest-cost path to the root bridge as the root port. In the event of equal path costs, the bridge examines these path's interfaces to the root bridge. The path with the lowest interface priority *to the root bridge* is the path whose port (interface) to the bridge becomes root port. For example, see Figure 7-8 which illustrates how Bridge A determines its root port.

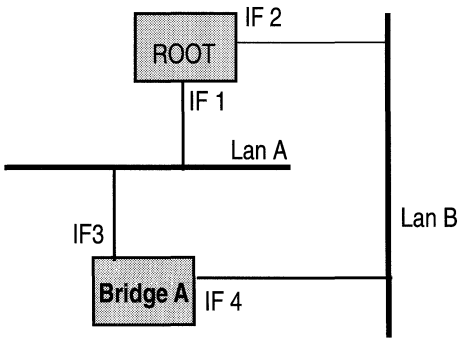


**Given:**

Path costs from IF 3* to Root and from IF 4 to Root are equal.

---------AND--------

IF 1 = priority 1
IF 2 = priority 2
IF 3 = priority 3
IF 4 = priority 4

**Then:**

IF 3 becomes root port because IF1's priority is lower than IF 2's priority.

* IF = Interface

**Figure 7-8. Root Port Determination (Equal Path Costs)**

If the paths' interfaces to the root bridge are also equal, then the root port is the port on the bridge with the lowest priority value (for example, see Figure 7-9).

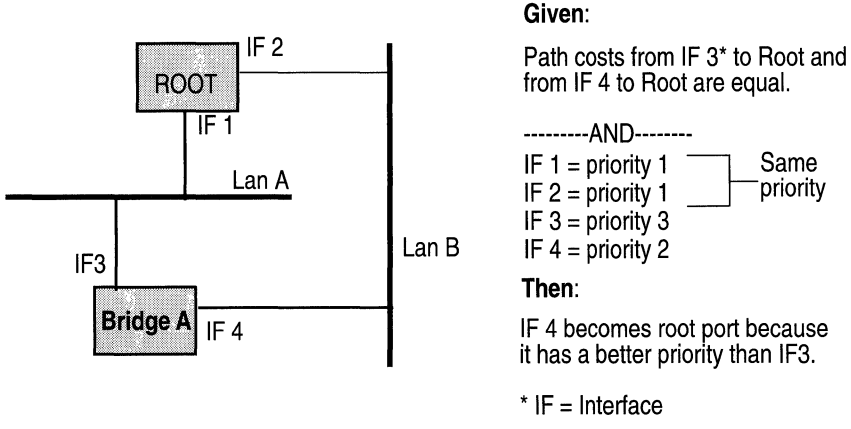


**Figure 7-9. Root Port Determination (Equal Path Costs and Root Interface Priorities)**

The one bridge on each LAN whose root port offers the lowest cost path to the root bridge is selected as the designated bridge. All bridges turn off (set to blocking state), all lines except for the single line that is the shortest cost path to the root, and any line attached to LANs for which for the bridge serves as designated bridge.

4. They elect a designated port.

   The port that connects the designated bridge to the LAN is selected as the designated port. If there is more that one such port, the port with the lowest priority is selected as designated port. This port, which carries all extended network traffic to and from the LAN, is said to be in the forwarding state.

This process ensures that all redundant ports (those providing parallel connections) are removed from service (placed in the blocking state). In

the event of a topological change, or in the event of bridge or data-path failure, the algorithm derives a new spanning tree that may move some ports from the blocking to the forwarding state.

Using Figure 7-7 as an example, if all path costs are equal, and if Bridge 2 has the lowest bridge priority, followed by Bridge 3, then Bridge 4, then Bridge 1; the spanning tree algorithm could block Bridge 1/Interface 8 and Bridge 4/Interface 9 from service. Figure 7-10 shows the resulting logical topology, which provides a loop-free topology with only a single path between any two hosts.

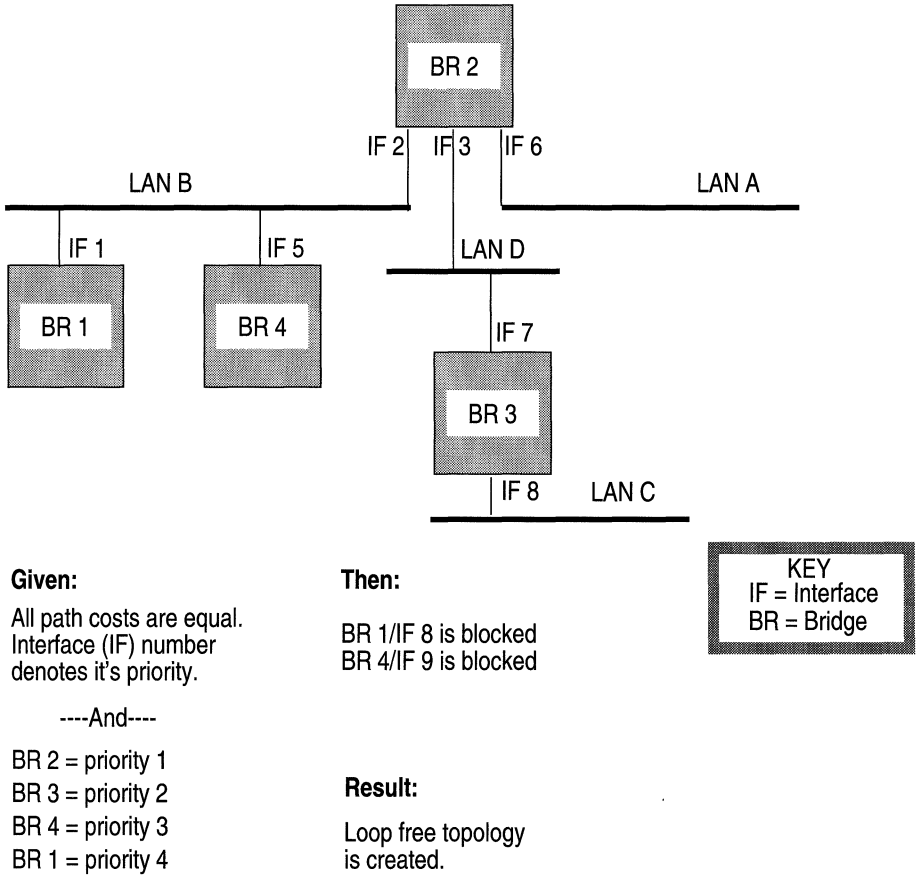


**Figure 7-10. Spanning Tree (Loop-Free) Logical Topology**

It is very important to configure the spanning tree parameters correctly. The network manager should take into account the typical flow of traffic so that the logical topology that results from the spanning tree algorithm is appropriate to the network.

If, in the network shown in Figure 7-7, a majority of traffic originates on LAN A and is destined for LAN D, it would be unwise for the network manager to set spanning tree parameters as shown in Figure 7-11. This figure exemplifies an inefficient spanning tree topology for this network because the traffic from LAN A must traverse Bridge 1, LAN B, and Bridge 2 to get to LAN D. LAN B is then congested with unnecessary traffic.

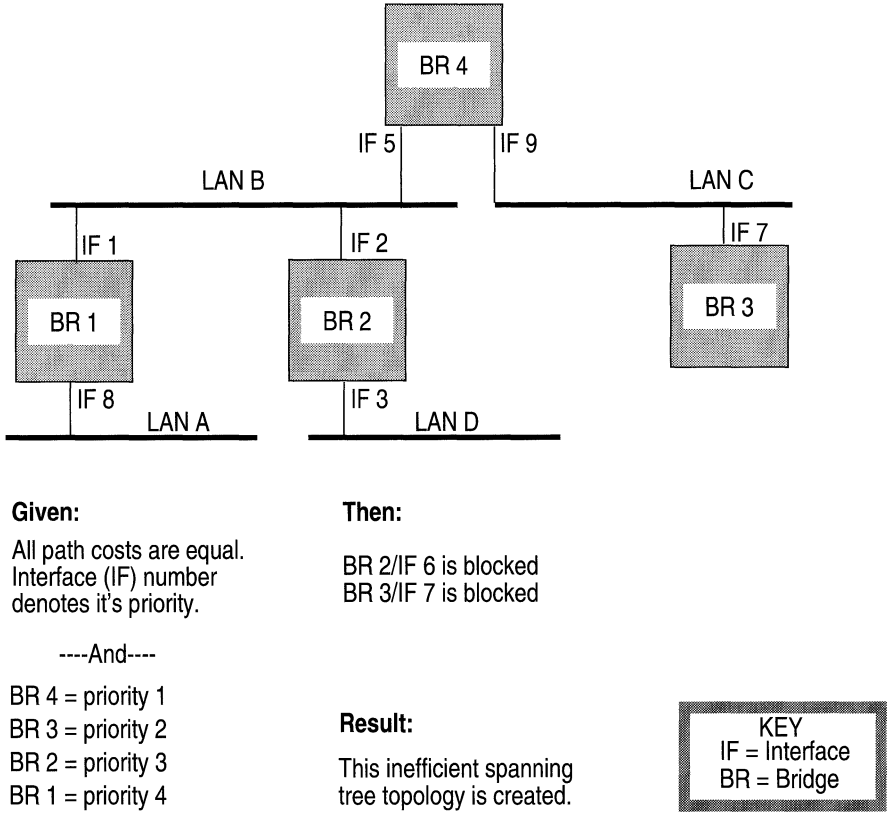


**Figure 7-11. Inefficient Spanning Tree Topology**

## Filtering

Filters are used mainly for security reasons. They enable the bridge to relay or drop a particular frame on the basis of user-selectable fields within each of the four encapsulation methods supported by the bridge; these encapsulation methods are:

- Ethernet
- IEEE 802.2 logical link control
- IEEE 802.2 LLC with SNAP header
- Novell proprietary

For more information about filters, and for instructions on how to configure filters for the bridge, refer to *Configuring Filters*.

# Editing Parameters

Once you have configured a circuit to support the Bridge and, optionally, the Spanning Tree Algorithm, you can use the Configuration Manager to edit Bridge or Spanning Tree parameters. The configuration function you wish to perform determines the type of parameters you must edit (see Table 7-1).

**Table 7-1. Bridge and Spanning Tree Parameters Configuration Functions**

| To Do the Following: | Edit these Parameters: |
|---|---|
| Enable or Disable Bridge for the entire BN. | Bridge Global Parameters |
| Enable or Disable Bridge on a particular circuit | Bridge Interface Parameters |
| Change the state of the Spanning Tree software. | Spanning Tree Global Parameters |
| Reconfigure Spanning Tree on a particular circuit. | Spanning Tree Interface Parameters |
| Configure Bridge Filters | Refer to the *Configuring Filters* chapter. |

**Note:** If the spanning tree algorithm is enabled for your network, dynamically changing *any* of the parameters described in the following sections will cause the spanning tree algorithm to reconverge.

This section describes how to access and edit the Bridge and Spanning Tree parameters listed in Table 7-1. For each parameter, it provides the following:

- Wellfleet default
- Valid options
- Parameter's function
- Instructions for setting the parameter

You begin from the Configuration Manager Window (see Figure 7-12); the first window displayed when you enter the Configuration Manager application.

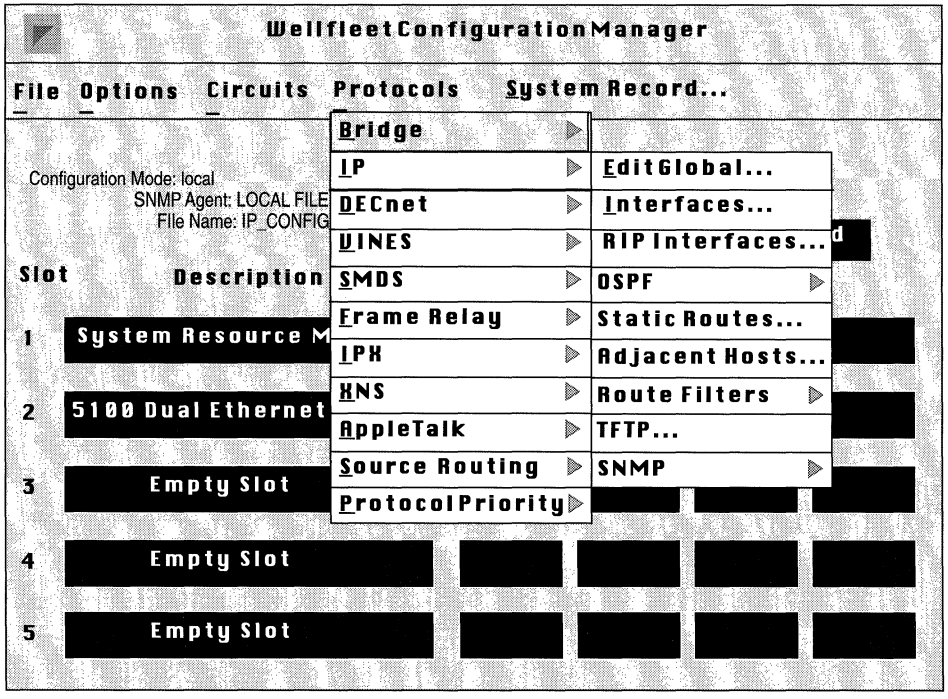


**Figure 7-12. Configuration Manager Window**

## Editing Bridge Global Parameters

For the bridge, there is only one global parameter: Enable. You edit this parameter in the Bridge Global Parameters Window. To locate this window, select the Protocols/Bridge option from the Wellfleet Configuration Manager Window to display the Bridge submenu, and then select the Edit Global option. The Bridge Global Parameters Window appears (see Figure 7-13).

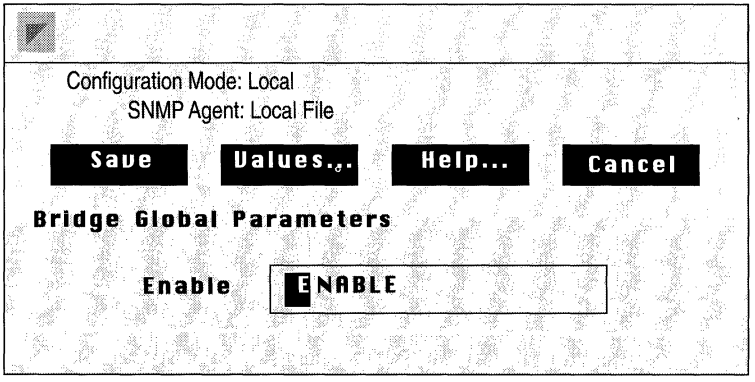


**Figure 7-13. Bridge Global Parameters Window**

When you are done editing this parameter, click the Save button.

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or disables bridging on the entire BN. |
| Instructions: | Set to Disable if you want to disable bridging for *all* circuits on the BN. |

## Editing Bridge Interface Parameters

For the bridge, there is only one interface parameter: Enable. You edit a Bridge Interface in the Bridge Interface Parameters Window for that interface. To locate this window, select the Protocols/Bridge option in the configuration Manager Window to display the Bridge submenu, and then select the Interfaces option, to display the Bridge Interfaces Window (see Figure 7-14).

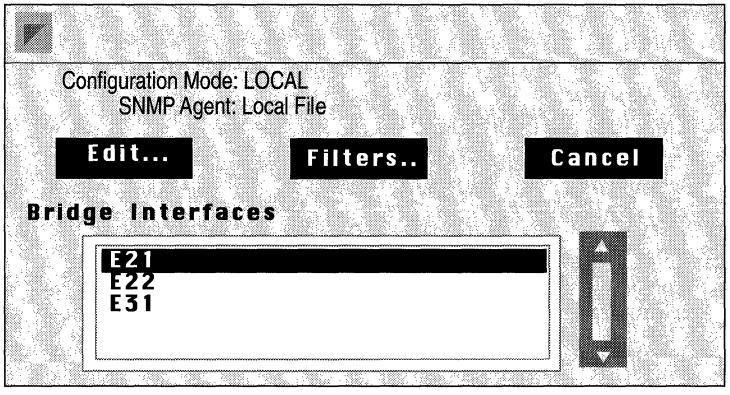


**Figure 7-14. Bridge Interfaces Window**

Next, select the interface you wish to edit, and click on the Edit button to display the Bridge Interface Parameters Window for that interface (see Figure 7-15).

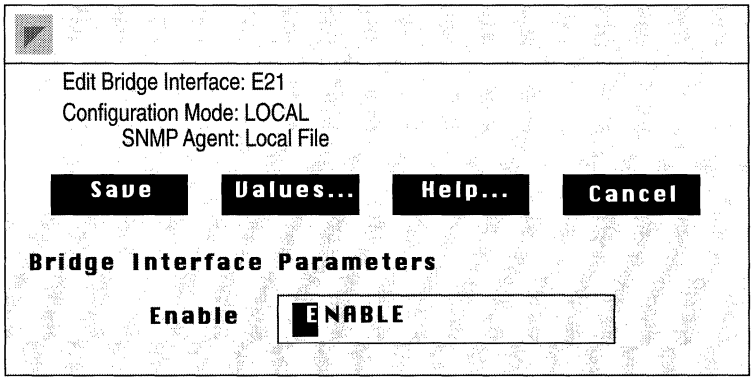


**Figure 7-15. Bridge Interface Parameters Window**

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | If you added bridging either using the Quick Start procedure or the configuring circuits procedure, this parameter defaults to Enable. If you previously used this parameter to disable bridging on this circuit, the parameter defaults to Disable. |
| Options: | Enable/Disable |
| Function: | Toggles bridging on and off for this circuit only.<br>This parameter does not allow you to add bridging to this circuit. To add the bridging protocol to this circuit, you must use the Configuration Manager (see *Editing Circuits).* |
| Instructions: | Set this parameter to either Enable or Disable. |

## Editing Spanning Tree Global Parameters

You edit global parameters in the Spanning Tree Global Parameters Window. To locate this window, select the Protocols/Spanning Tree option from the Wellfleet Configuration Manager Window to display the Spanning Tree submenu, and then select the Edit Global option to display the Spanning Tree Global Parameters Window appears (see Figure 7-16).

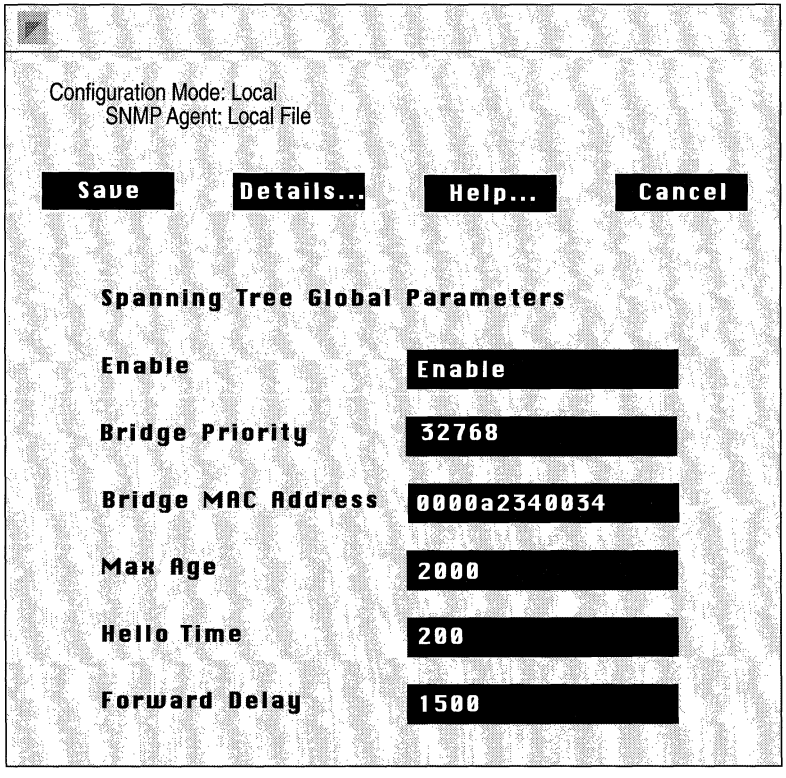


**Figure 7-16. Spanning Tree Global Parameters Window**

This section provides information you need to edit each parameter in the Spanning Tree Global Parameters Window. Refer to this

information to edit the parameters you wish to change. When you are done, click the Save button to exit the window and save your changes.

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or disables spanning tree on the entire BN. |
| Instructions: | Set to Disable if you want to disable spanning tree for the entire BN, or to Enable if you want to re-enable spanning tree for the entire BN. |

| | |
|---|---|
| **Parameter :** | **Bridge Priority** |
| Wellfleet Default: | Defaults to the Bridge Priority that was set when spanning tree was added to this circuit. |
| Options: | 0 - 65535 (expressed as a decimal value) |
| Function: | Together with the Bridge Priority, the Bridge MAC Address forms a bridge ID. The Bridge ID provides a unit of comparison for use with the Spanning Tree Protocol. The Bridge ID with the lowest value will be selected as the Root Bridge. Due to its positional significance in the Bridge ID, The Bridge priority strongly influences Root Bridge selection. The lower the bridge priority, the more likely that the bridge will be chosen as the Root Bridge. In the case of equal Bridge Priority values, the Bridge MAC Address determines the bridge's priority. The lower the address, the higher the priority. |
| Instructions: | Either accept the current Bridge Priority value, or assign a new one. |

| | |
|---|---|
| **Parameter :** | **Bridge MAC Address** |
| Wellfleet Default: | Defaults to the Bridge MAC Address that was set when spanning tree was added to this circuit. |
| Options: | Any MAC Address unique to the network (expressed as a 12 digit hexadecimal value) |
| Function: | Together with the Bridge Priority, the Bridge MAC Address forms a bridge ID. The Bridge ID provides a unit of comparison for use with the Spanning Tree Protocol. The Bridge ID with the lowest value will be selected as the Root Bridge. Due to its positional significance in the Bridge ID, The Bridge priority strongly influences Root Bridge selection. The lower the bridge priority, the more likely that the bridge will be chosen as the Root Bridge. In the case of equal Bridge Priority values, the Bridge MAC Address determines the bridge's priority. The lower the address, the higher the priority.<br><br>Wellfleet recommends that you set this parameter to the MAC Address of one of this bridge's spanning tree ports, preferably the one with the lowest priority. |
| Instructions: | Either accept the current Bridge MAC Address, or enter a new one. |

**Note:** Wellfleet recommends that you do *not* change the following three spanning tree parameters (Max Age, Hello Time, and Forward Delay) unless absolutely necessary. However, if you do decide to change them, you must adhere to the following guidelines.

$$2 \times (\text{Forward Delay} - 1 \text{ second}) \geq \text{Max Age}$$

$$\text{Max Age} \geq 2 \times (\text{Hello Time} + 1 \text{ second})$$

**Note:** If the values for Max Age, Hello Time, and Forward Delay are not the same for each bridge in your network, the Root Bridge parameters rule the entire topology

| **Parameter :** | **Max Age** |
|---|---|
| Wellfleet Default: | 20 seconds (expressed in hundredths of a second: 2000) |
| Options: | 6 - 40 seconds |
| Function: | Specifies the maximum number of seconds that protocol information (BPDUs) is considered valid. After this specified amount of time, the information is timed out and discarded.<br><br>Wellfleet recommends that you accept this default value; however, if you do decide to change it, you must adhere to the guidelines specified in the previous note. |
| Instructions: | Either accept the default value, or specify a new value in the Max Age box. Make sure to express any new value in hundredths of a second. |

| | |
|---|---|
| **Parameter :** | **Hello Time** |
| Wellfleet Default: | 2 seconds (expressed in hundredths of a second: 200) |
| Options: | 1 - 10 seconds |
| Function: | Specifies the interval in seconds between BPDUs transmitted by the bridge. BPDUs are periodic transmissions exchanged between bridges in the network to convey configuration and topology change data.<br><br>Wellfleet recommends that you accept the default value; however, if you do decide to change it, you must adhere to the guidelines specified in the previous note. |
| Instructions: | Either accept the default value, or enter a new value in the Hello Time box. Make sure to enter any new value in hundredths of a second. |

| | |
|---|---|
| **Parameter :** | **Forward Delay** |
| Wellfleet Default: | 15 seconds (expressed in hundredths of a second: 1500) |
| Options: | 4 - 30 seconds |
| Function: | Specifies the time in seconds that a circuit spends in the Listening and Learning states.<br><br>Setting Forward Delay to the minimum value causes the spanning tree to converge at its fastest rate.<br><br>As the algorithm operates, it eventually places all circuits in either a Forwarding (enabled) or Blocking (disabled) state. Later, in response to network topology changes, the algorithm may change the state of specific circuits. To prevent network looping caused by sudden state changes, the algorithm does not transition circuits directly |

from Blocking to Forwarding. Rather, it places them in two inter-mediate states, Listening and Learning.

In the Listening (stand-by) state, the circuit listens for network-generated BPDUs. It receives and drops all other traffic. When the Forward Delay Timer expires, the circuit is placed in the Learning state.

In Learning state, the circuit receives both network-generated BPDUs, and end-station-generated traffic that is subjected to the learning process but not relayed. When the Forward Delay Timer again expires, the circuit is placed in the Forwarding state.

Wellfleet recommends that you accept this default value; however, if you do decide to change this value, you must adhere to the guidelines specified in the previous note.

Instructions: Either accept the default value, or enter a new value in the Forward Delay box. Make sure to enter any new value in hundredths of a second.

## Spanning Tree Interface Parameters

You edit a Spanning Tree interface in the Spanning Tree Interface Parameters Window for that interface. To locate this window, select the Protocols/Spanning Tree option in the Wellfleet Configuration Manager Window to display the Spanning Tree submenu, and then select the Interfaces option to display the Spanning Tree Interfaces Window (see Figure 7-17).

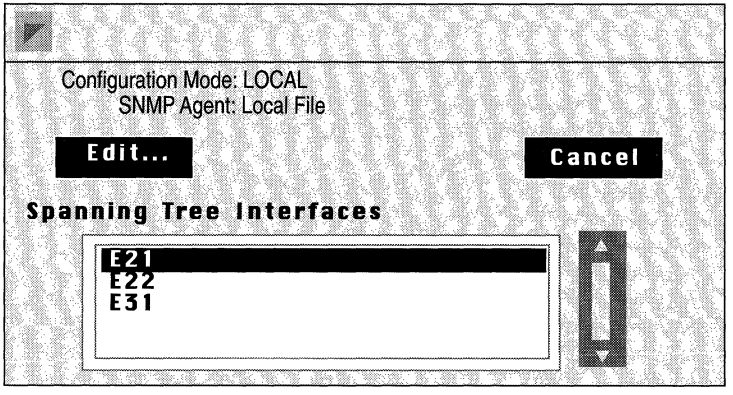


**Figure 7-17. Spanning Tree Interfaces Window**

Next, select the interface you wish to edit and click on the Edit button to display the Spanning Tree Interface Parameters Window for that interface (see Figure 7-18).

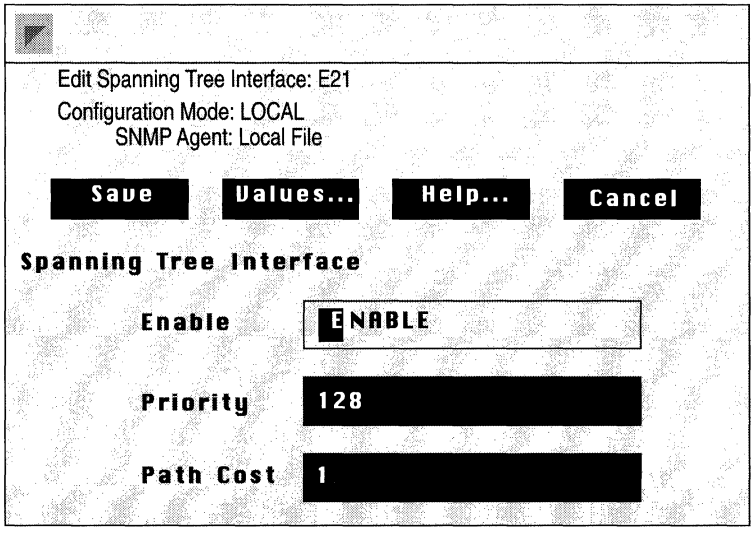


**Figure 7-18. Spanning Tree Interface Parameters Window**

This section provides information you need to edit each parameter in the Spanning Tree Interface Parameters Window. Refer to this information to edit the parameters you wish to change. When you are done, click the Save button to exit the window and save your changes.

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | If you added spanning tree either using the Quick Start procedure or the configuring circuits procedure, this parameter defaults to Enable. If you previously used this parameter to disable spanning tree on this circuit, the parameter defaults to Disable. |
| Options: | Enable/Disable |
| Function: | Toggles spanning tree on and off for this circuit only.<br><br>This parameter does not allow you add spanning tree to this circuit. To add the spanning tree to this circuit, you must use the Configuration Manager (see *Editing Circuits).* |
| Instructions: | Set this parameter to either Enable or Disable. |

| | |
|---|---|
| **Parameter :** | **Priority** |
| Wellfleet Default: | 128 |
| Options: | 0 - 255 |
| Function: | Assigns a priority to a bridge port. This interface priority value, together with the bridge ID (Bridge Priority + Bridge MAC Address), determines whether or not this port will become designated port when the spanning tree algorithm converges. The lower the priority value, the higher the priority, and the more likely that this port will be designated port. |
| Instructions: | Either accept the default value, or enter a new value in the Priority box. |

| | |
|---|---|
| **Parameter :** | **Path Cost** |
| Wellfleet Default: | 1 |
| Options: | 1 - 65535 |
| Function: | When this port *is* the root port, this Path Cost is the contribution of the path through this port to the total cost of the path to the root for this bridge. When this port is *not* the root port, this Path Cost is added to the designated cost for the root port and is used as the value of the root path cost offered in all configuration BPDUs transmitted by the bridge.<br><br>When determining Path Cost, Wellfleet recommends that you use this formula:<br><br>*Interface Path Cost = 1000/Attached LAN speed in Mb/s* |
| Instructions: | Enter a path cost value for this interface in the Path Cost box. For example, you would enter 100 if the attached LAN was Ethernet (1000/10 = 100). |

# Deleting the Bridge and Spanning Tree from the BN

You can delete the Bridge and/or the Spanning Tree from all BN circuits on which they are currently enabled in two steps.

You begin from the Wellfleet Configuration Manager Window and complete the following steps:

1. To delete the Bridge and Spanning Tree (if it is enabled), select the Protocols/Bridge/Delete Bridge option.

   Or, if you want to delete just the Spanning Tree, select the Protocols/Bridge/Spanning Tree/Delete Spanning Tree option.

   A pop-up window appears prompting whether you really want to delete.

2. Select Ok.

   You are returned the Wellfleet Configuration Manager window. The Bridge and/or Spanning Tree is no longer configured on the BN.

If you deleted the Bridge, the connectors for those circuits on which the Bridge was the *only* protocol enabled are no longer highlighted in the Wellfleet Configuration Manager Window. Circuits must be reconfigured for these connectors; see *Configuring Circuits* for instructions.

# Chapter 8

## Configuring Source Routing

## List of Figures

## List of Tables

# Configuring Source Routing

## About this Chapter

This chapter describes how to configure the Wellfleet Source Routing Bridge.

This chapter begins with an overview of source routing technology. The second section describes how the Wellfleet Source Routing Bridge works. The third section lists additional source routing reference material. The fourth section describes implementation guidelines for adding source routing bridges to your network. The final sections describes how to use the Configuration Manager to edit source routing parameters and how to delete source routing from the BN.

## Source Routing Overview

The source routing bridge is used to route frames across token ring networks. Source routing bridges do not contain routing tables; instead they depend upon end stations to provide all route descriptors for frames that are sent out on the network.

Source routing networks consist of LAN segments interconnected by source routing bridges (see Figure 8-1). Each LAN segment has a unique network-wide identification number (Ring ID). Each source routing bridge has an identification number (Bridge ID) and a unique network-wide internal or virtual LAN identification number (called Internal LAN ID). The Internal LAN ID is required because source routing relies on features of the token ring chipset to capture source routed packets from the LAN.

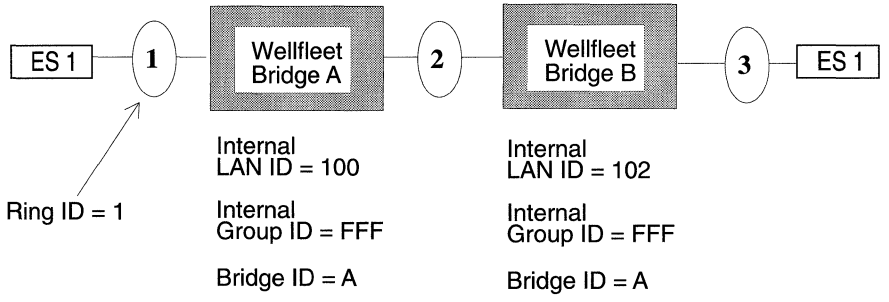


**Figure 8-1. Source Routing Network**

Each Wellfleet Source Routing Bridge is assigned an additional routing identifier called a Group LAN ID. The Group LAN ID serves as a RIF placeholder and Wellfleet identifier.

## How Source Routing Differs From Transparent/Translating Bridging

The Wellfleet Source Routing Bridge differs from the Wellfleet Transparent/Translating bridge in two ways:

- Source routing bridges can tolerate multiple paths between end stations in an extended network; Transparent/Translating bridges require loop-free topologies.
- Source routing bridges require end stations to supply the bridging information needed to deliver a frame to a destination.

The end stations on a source routing network use a process called route discovery to gather bridging information, as described in the next section.

## How End Stations on a Source Routing Network Discover Routes

End stations discover the routes to all other reachable destinations using the following routing directives:

- *All paths broadcast routing*:

  An end station configured for all paths broadcast routing generates multiple frames that traverse all paths between source and destination end-stations. Such frames are called all-paths explorer (APE) frames. Upon receiving an APE frame, a bridge within the source routing network appends a routing designator that identifies the incoming Ring ID, Internal LAN ID/Bridge ID and the outgoing Ring ID as shown in Figure 8-2.

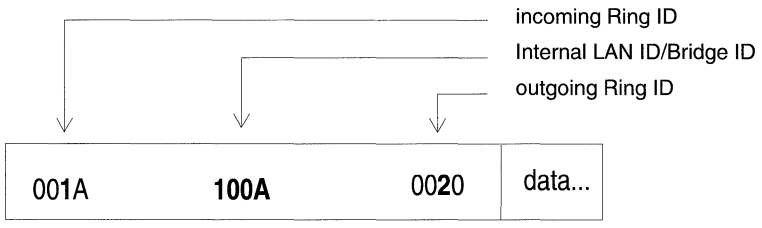


**Figure 8-2. Source Routing Designator**

After adding a routing designator, each bridge floods the frame. As a consequence, multiple copies of the same APE frame can appear on a LAN, and the frame recipient can receive multiple copies of the frame (one copy for each possible path through the extended network). Each APE frame received by the destination end station contains a unique sequenced list of routing designators tracing the frame's path through the source routing network.

**Note:** In the case of a looped topology, the originating bridge can receive the APE frame, in which case it will be discarded.

- *Spanning Tree broadcast routing*:

  An end station configured for spanning tree broadcast routing generates a single frame that follows a loop-free path from source to destination. Such frames are called spanning tree explorer frames (STEs). Upon receiving an STE, each bridge on the spanning tree forwards the frame onto all active (non-blocked) ports save the port on which the frame was received. With spanning tree broadcast routing, one copy of the STE appears on each LAN, and the frame recipient receives only a single copy of the frame.

- *Specific Routing*

  When the destination end station receives the APE or STE, it generates a single frame, called a specifically routed frame (SRF), that traverses a specific path back to the source end-station. The SRF contains a list of routing designators that maps a path through the extended network from source to destination. Upon receiving an SRF, each bridge between the source and destination examines the list of routing designators. It forwards the SRF only if it is on the specified path, otherwise it ignores the frame. Once it reaches the original source end station, the station removes the routing information and stores it in its internal routing table.

Once the route between end stations is discovered and stored in the end stations' routing tables, the end stations can begin sending specifically routed frames across the source route network (see Figure 8-3).

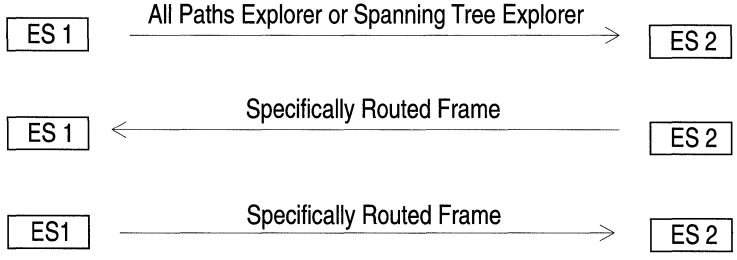


**Figure 8-3. Route Discovery**

## Source Routing Over IP Networks

The Wellfleet Source Routing Bridge now supports *IP encapsulation*. IP encapsulation allows the source routing bridge to route frames to end stations located across an IP backbone network.

Frames that are source routed across IP networks use standard IP transmission services, as well as a proprietary sequence maintenance protocol that ensures error-free, in-sequence delivery of IP encapsulated frames. The IP network can consist of any standard IP equipment and media.

### How IP Encapsulation Works

When you enable IP encapsulation on a source routing bridge (thus making it an IP encapsulating bridge), you assign a *single* Ring ID to the entire IP backbone network. Regardless of the size of the IP network that a frame traverses, only a single route descriptor is added to the frame's RIF to describe the entire internet. Thus, frames source routed over large IP networks can remain within maximum hop count restrictions.

Whenever a Wellfleet Source Routing Bridge receives an explorer frame, it sends it toward an IP encapsulating bridge residing at the edge of the IP backbone network (for example, in Figure 8-4, Bridges A, B and C are IP encapsulating bridges). The IP encapsulating bridge encapsulates the source routed frame with an IP header before sending the frame out onto the network. When the frame reaches a peer IP encapsulating bridge, the bridge decapsulates the frame and sends it out the appropriate source routing interfaces.

Each IP encapsulating bridge maintains a dynamic mapping of destination IP addresses to the ring IDs of their directly attached rings. When an IP encapsulating bridge receives a source routed frame, it checks the frame's RIF for the next Ring ID that *immediately follows* the IP network Ring ID in the RIF. Then it looks up the IP address that corresponds with this Ring ID and encapsulates the frame in an IP packet with this destination IP address. Finally, it send the frame out onto the IP network.

IP Explorer List Bridge A
1.1.2.2
1.1.3.3
1.1.5.5

IP Explorer List Bridge B
1.1.1.1
1.1.5.5

IP Explorer List Bridge C
1.1.1.1
1.1.2.2
1.1.3.3

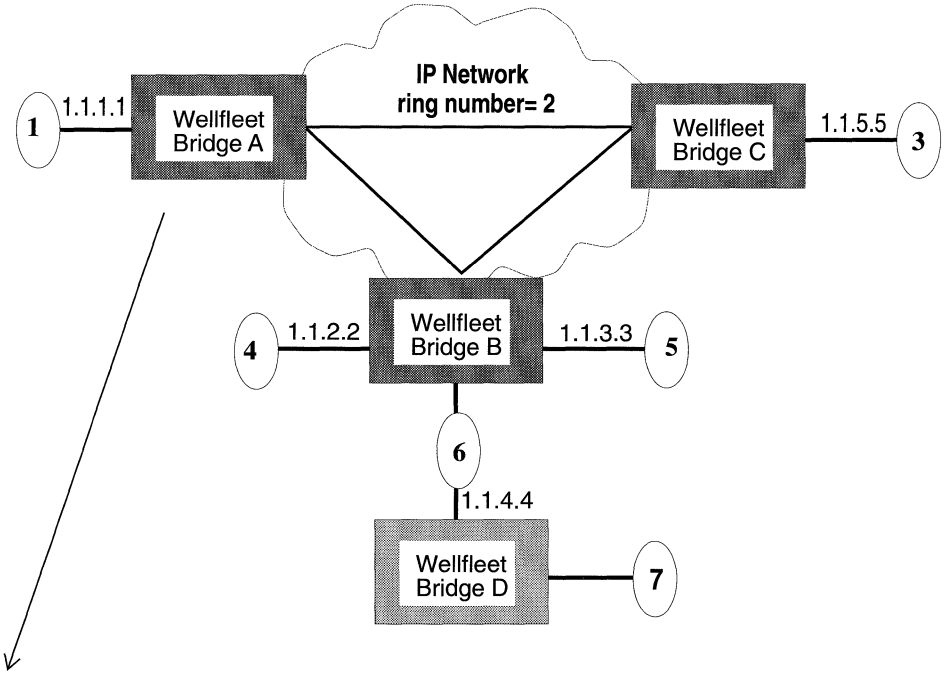


**IP Mapping Table Bridge A**

| Ring ID | IP Address |
|---|---|
| 3 | 1.1.5.5 |
| 4 | 1.1.2.2 |
| 5 | 1.1.3.3 |
| 6 | (Outgoing IP interface address on WF Bridge B) |

**Figure 8-4. Source Routing Over an IP Network**

For example, Figure 8-4 shows the IP mapping table for IP encapsulating Bridge A. When Bridge A receives a source routed frame destined for an end system on ring 4, it examines the frame's RIF (See Figure 8-5) and locates the next Ring ID (4) that immediately follows the IP network Ring ID (2). Then, it checks its mapping table for the IP address that corresponds to Ring ID 4 and encapsulates the frame with an IP header. The IP header specifies its own source address, and the proper IP destination address (1.1.2.2). Finally, it forwards the packet onto the IP network.

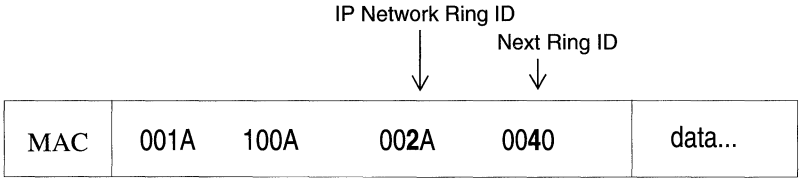


**RIF of a Source Routed Frame**

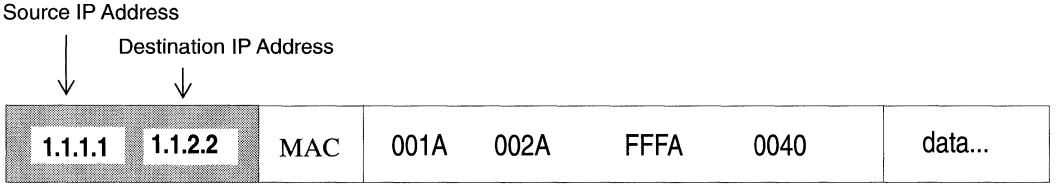


**RIF of an IP Encapsulated Frame**

**Figure 8-5. Examining the RIF Field of an SRF**

(See the section entitled *How the Wellfleet Source Routing Bridge Works* later in this chapter for more source routing RIF information).

You control which IP interfaces receive explorer frames by defining an *IP Explorer list* for each IP encapsulating bridge. For example, all of the IP encapsulating bridges that border the IP cloud in Figure 8-4 have IP Explorer lists defined. Each bridge will only forward explorer packets toward the IP addresses that are included in its individual list. Note that the IP explorer lists for each bridge can vary. This allows you to control which IP networks receive explorer traffic.

### IP Encapsulation Features

Wellfleet's implementation of IP encapsulation allows you to do the following:

- Configure redundant IP interfaces.

  If you wish, you can configure redundant IP interfaces on the same BN for critical network connections (for example, interfaces 1.1.2.2 and 1.1.3.3 on Bridge B). That way, if one of the interfaces is disabled, the other interface can still accept IP traffic for the network. (When you enable redundant IP interfaces, you also increase explorer traffic on the network. Therefore, enable redundant interfaces selectively to reduce the impact on your network performance).

- Expand your IP backbone network.

  You can expand your IP backbone to include any Wellfleet IP router on the network simply by specifying its IP address in the IP explorer list for each bridge. For example, Bridge A currently forwards all traffic destined for ring 7 to IP interface 1.1.3.3. That router then forwards the traffic toward ring 6 so it can be source routed to ring 7. If you added IP address 1.1.4.4. to the IP Explorer list for the bridge A, then bridge A would forward all traffic destined for ring 7 directly to IP interface 1.1.4.4. By expanding your IP backbone, the source routing bridge can route through more stations but still only add a single hop to a packet's RIF.

- Reduce excess broadcast traffic on your IP network.

  You can reduce the number of broadcast and explorer packets that traverse the network by constructing *directed explorer filters (*see the chapter entitled *Configuring Filters*).

- Configure both IP encapsulation support and source route end station support on the same interface.

  IP encapsulation support works independently of source route end station support (see the section entitled *Source Route End Station Support* for more information). However, both can be enabled on the same circuit.

## Source Route End Station Support

The Wellfleet IP, IPX, XNS, AppleTalk and VINES routers now support routing over token ring networks that contain one or more source routing bridges. This feature is called *source route end station support*.

In a source routing network, every end station that sends out a frame supplies the frame with the necessary route descriptors so that it can be source routed across the network. Thus, in order for routers to route packets across a source routing network, *they must act like end stations*; supplying route descriptors for each packet before they send it out onto the network.

With end node support enabled, whenever a Wellfleet router running IP, IPX, XNS, AppleTalk or VINES receives a packet and determines that the packet's next hop is located across a source routing network, the router does the following:

- Adds the necessary RIF information to the packet's MAC header.
- Sends the packet out onto the network where it is source routed toward the next hop.

Upon receiving the packet from the token ring network, the peer router strips off the RIF field and continues to route the packet toward the destination network address (see Figure 8-6).

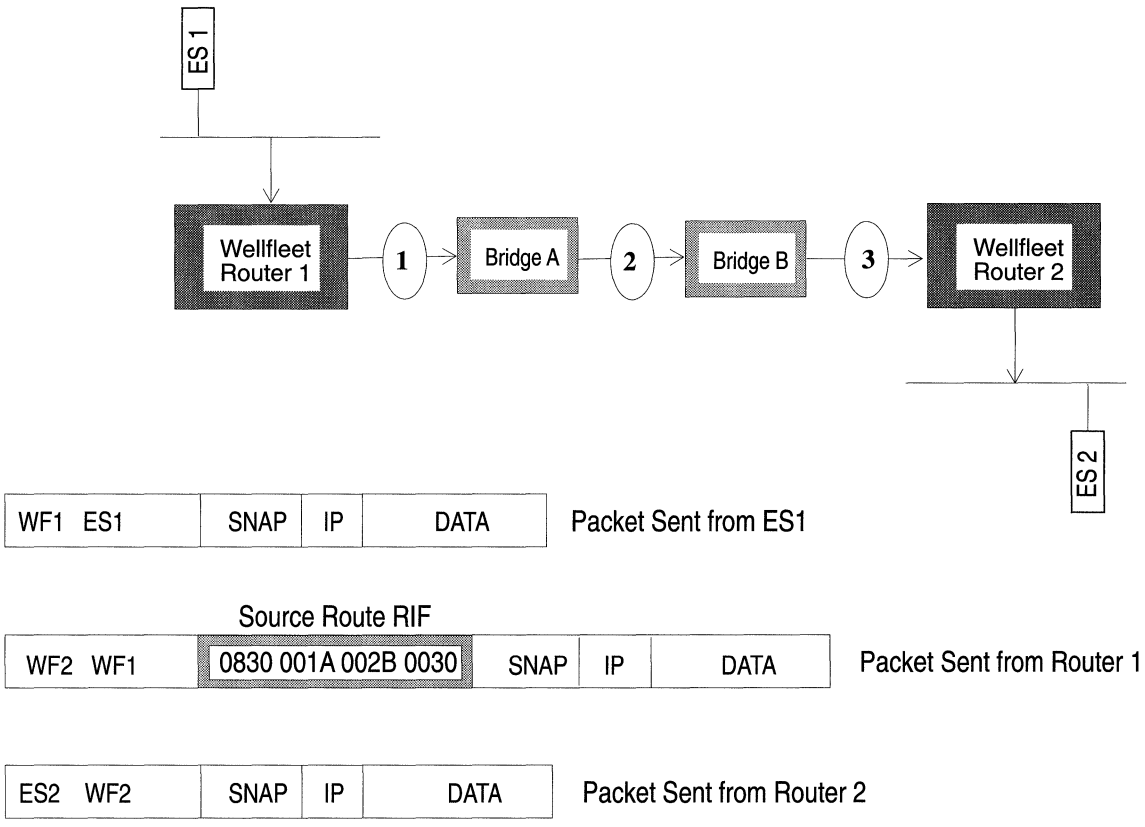


**Figure 8-6. IP Routers Source Routing Across a Token Ring Network**

You configure source route end node support for each individual routing protocol on a per-circuit basis. For more information, refer to the section entitled *Source Route End Station Support* in one of the following chapters of this guide: *Configuring IP*, *Configuring IPX*, *Configuring XNS*, *Configuring AppleTalk*, or *Configuring VINES*.

# How the Wellfleet Source Routing Bridge Works

This section is *optional* reading. First, it describes how the Wellfleet Source Routing Bridge routes frames through a token ring network. Then, it describes how the Wellfleet Source Routing Bridge routes frames across an IP backbone network (called *IP encapsulation*).

## Source Routing Across a Token Ring Network

The Wellfleet Source Routing Bridge handles incoming packets differently depending on its position in the token ring network. To demonstrate all possible scenarios, the following sections describe the Routing Information Field (RIF) of a frame as it moves back and forth between ES1 and ES2 as follows (see Figure 8-7):

- First, we track the RIF of an explorer frame sent from ES1 to ES2.
- Next, we track the RIF of a specifically routed frame sent back from ES2 to ES1.
- Finally, we track the RIF of a specifically routed frame sent from ES1 to ES2.

The size of the RIF is variable. It contains the routing information required to transmit the frame across the network.

Although the following examples show only Wellfleet Source Routing Bridge, other IBM compatible source routing bridges can reside in the same network.

### How the Wellfleet SR Bridge Handles Explorer Frames

This section describes how the Wellfleet Source Routing Bridge handles explorer frames (APEs or STEs) sent from ES1 to ES2 (see Figure 8-7). Each bridge's Internal LAN ID, Group LAN ID, and Bridge ID is listed in hexadecimal format below the bridge.

- Scenario A describes the actions of the first Wellfleet Bridge.
- Scenario B and C describe the actions of other Wellfleet Bridges in the explorer frame's path.

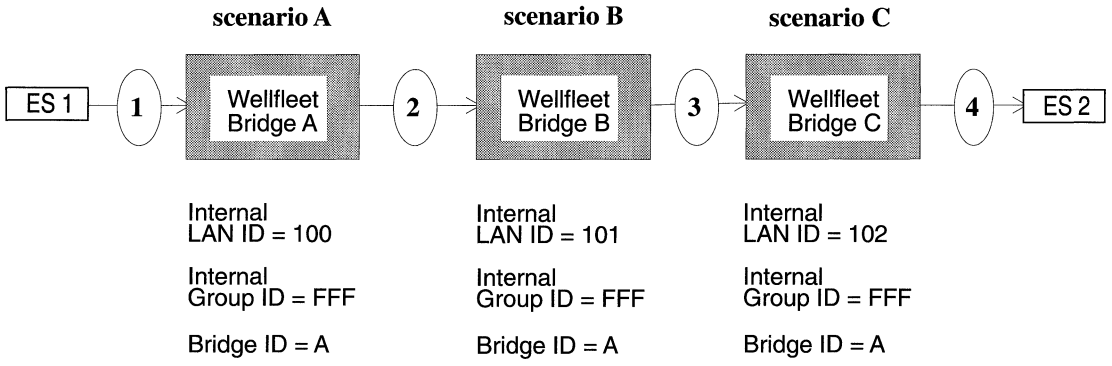


**Figure 8-7. Tracking an Explorer Frame**

### Scenario A: First Wellfleet Bridge to receive the explorer frame.

The frame received by Bridge A from ring 1 has not traversed any other bridges. This bridge simply adds the following to the RIF before transmitting the frame toward ring 2 (see Figure 8-8):

1. Incoming Ring ID/Bridge ID.
2. Internal LAN ID/Bridge ID.
3. Outgoing Ring ID/Bridge ID of 0.

### Scenario B, C: Other Wellfleet Bridges that receive the explorer frame.

The explorer frame received by Bridge B and C contains Internal LAN IDs and the Wellfleet Bridge number (s) (thus indicating that this frame has traversed at least one other Wellfleet Bridge). These bridges do the following to the RIF before transmitting the frame toward rings 3 and 4 (see Figure 8-8):

1. Removes the last Wellfleet internal LAN ID.
1. Replaces the incoming Ring ID/Bridge ID of 0 with its own Bridge ID.
2. Adds its own Internal LAN ID/Bridge ID.
3. Adds the outgoing Ring ID/Bridge ID of 0.

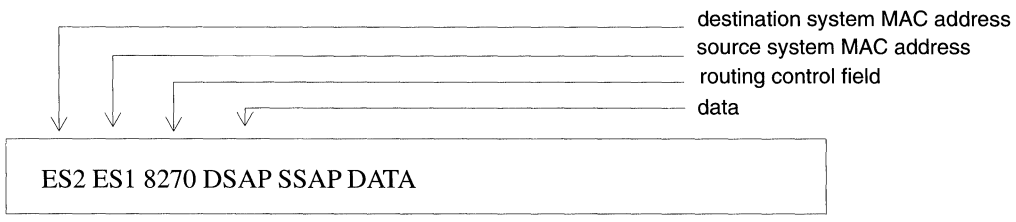


**Frame Received by Bridge A**

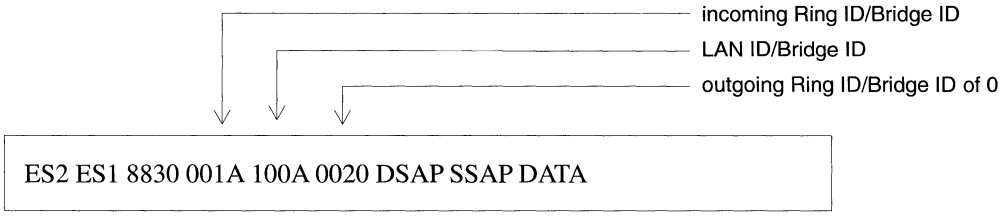


**Frame sent out by Bridge A onto ring 2**

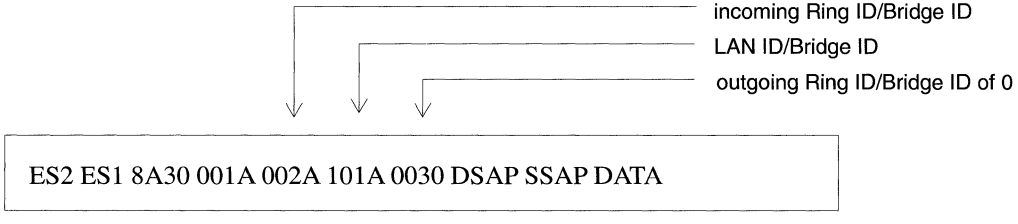


**Frame sent out by Bridge B onto ring 3**

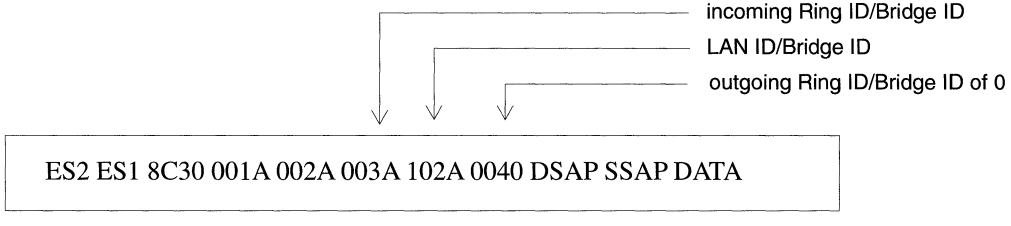


**Frame sent out by Bridge C onto ring 4**

**Figure 8-8. Structure of an Explorer Frame**

### How the Wellfleet SR Bridge Handles Specifically Routed Frames from ES2 to ES1

This section describes how the Wellfleet Source Routing Bridge handles Specifically Routed Frames (SRFs) sent from ES2 to ES1. Once again, depending on the bridge's position in the network, SRFs are handled differently (see Figure 8-9).

If there is only a single Wellfleet Bridge ID in the RIF, then the Wellfleet Bridge simply transmits the frame to the outgoing circuit without making any modification. This is only true when the frame only has to traverse a single Wellfleet Bridge (or any combination of non-Wellfleet bridges) between the source and destination end system.

If there are multiple multiple Wellfleet Bridge IDs in the RIF, then:

- Scenario A describes the actions of a Wellfleet Bridge if it is the first Wellfleet Bridge to handle the SRF.
- Scenario B describes the actions of a Wellfleet Bridge if there are multiple Wellfleet Bridge IDs in the RIF, and this Wellfleet Bridge is in between the first and the last Wellfleet Bridge that will handle the SRF.
- Scenario C describes the actions of a Wellfleet Bridge if there are multiple Wellfleet Bridge IDs in the RIF, and this Wellfleet Bridge is the last Wellfleet Bridge to handle the SRF.

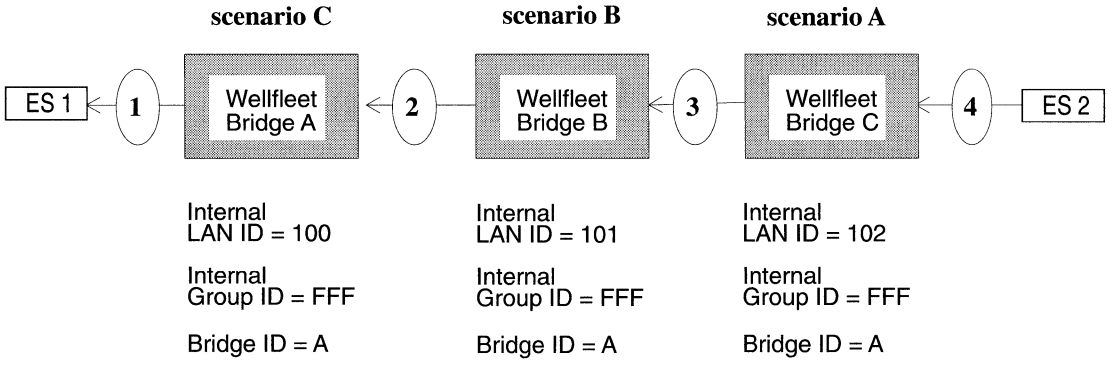


**Figure 8-9. Tracking a Specifically Routed Frame from ES2 to ES1**

**Scenario A: First of several Wellfleet Bridges to receive the SRF.**

The frame received by Bridge C from ring 4 has not traversed any other Wellfleet Bridges yet (but there are multiple Wellfleet Bridge IDs in the RIF, so it will). This bridge does the following to the RIF before transmitting the frame toward ring 3 (see Figure 8-10):

1. Changes the destination's system's MAC address at the beginning of the frame to a Wellfleet Group Address. This address appears as C000A2FFFFFFx, where "x" is the Bridge ID of the next Wellfleet Bridge specified by the RIF.
2. Removes its own Internal LAN ID and inserts the Group LAN ID (before the last Incoming Ring/Bridge ID listed in the RIF, see Figure 8-10). Eventually, the Group LAN ID will be replaced with the Internal LAN ID of the last Wellfleet Bridge along the frame's path.
3. Copies the destination system's MAC address into the data portion of the frame.

**Scenario B: Between the first and last Wellfleet Bridge to receive the SRF.**

The frame received by Bridge B from ring 3 has traversed at least one other Wellfleet Bridge. However, this is not the last Wellfleet Bridge that the frame must traverse. This bridge does the following to the RIF before transmitting the frame toward ring 2 (see Figure 8-10):

1. Locates the Bridge ID located at the end of the Group Address.
2. Changes the Bridge ID at the end of the Group Address to the Bridge ID of the next Wellfleet Bridge in the RIF. (Only if it differs from the value already in place. In this example, all Bridge IDs are the same, so the frame is not modified).

**Scenario C: Last of several Wellfleet Bridges to receive the SRF.**

The frame received by Bridge A is the last of several Wellfleet Bridges traversed by the frame. This bridge does the following to the RIF before transmitting the frame toward ring 1 (see Figure 8-10):

1. Replaces the Wellfleet Group address with the destination MAC address that was saved to the data field.
2. Replaces the Group LAN ID with its own Internal LAN ID.

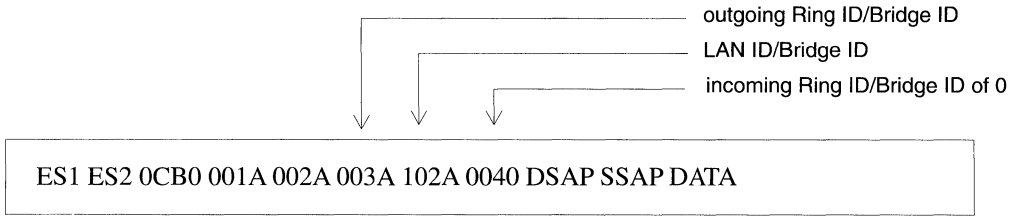


**Frame received by Bridge C**

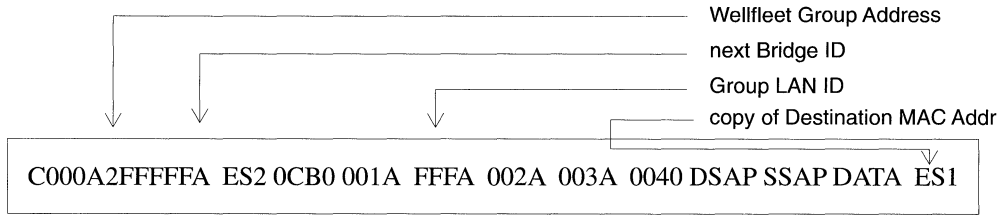


**Frame sent out by Bridge C onto ring 3**

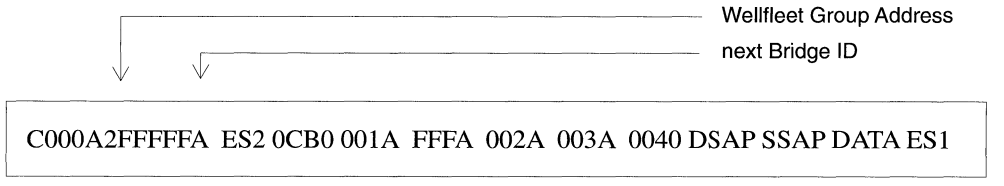


**Frame sent out by Bridge B onto ring 2**

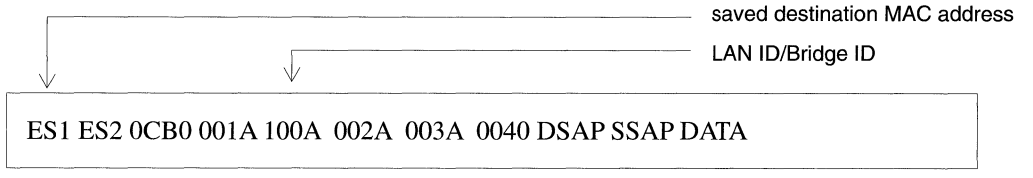


**Frame sent out by Bridge A onto ring 1**

**Figure 8-10. Structure of a Specifically Routed Frame from ES2 to ES1**

### How the Wellfleet SR Bridge Handles Specifically Routed Frames from ES1 to ES2

This section describes how the Wellfleet source routing Bridge routes specifically routed frames from ES2 to ES1. Although the algorithm works the same as when ES2 routed a specifically routed frame to ES1, this final section should clarify any questions about how the new source routing algorithm works (see Figure 8-11).

If there is only a single Wellfleet Bridge ID in the RIF, then the Wellfleet Bridge simply transmits the frame to the outgoing circuit without making any modification. This is only true when the frame only has to traverse a single Wellfleet Bridge between the source and destination end system. Because of the simplicity of this case, it is not described in any further detail here.

If there are multiple multiple Wellfleet Bridge IDs in the RIF, then:

- Scenario A describes the actions of the first Wellfleet Bridge to handle the SRF.
- Scenario B describes the actions of a Wellfleet Bridge that is in between the first and the last Wellfleet Bridge that will handle the SRF.
- Scenario C describes the actions of the last Wellfleet Bridge of several Wellfleet Bridges to handle the SRF.

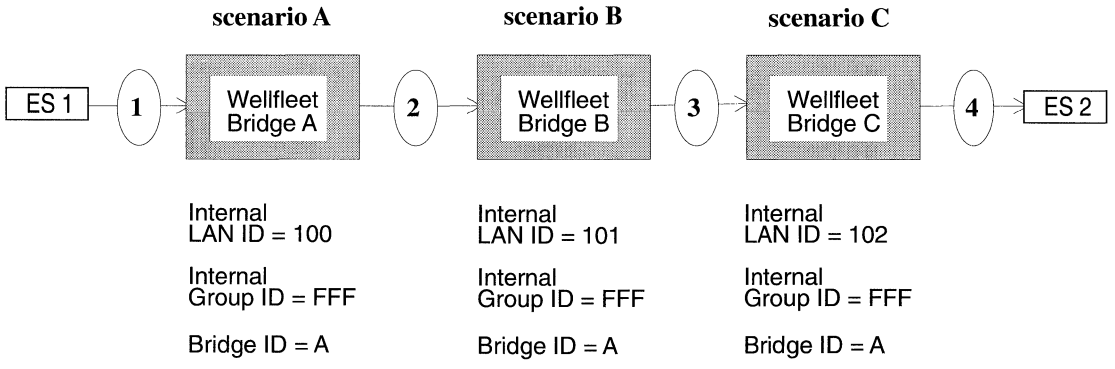


**Figure 8-11. Tracking a Specifically Routed Frame from ES1 to ES2**

**Scenario A: First of several Wellfleet Bridges to receive the SRF.**

Wellfleet Bridge A is the first to receive the frame from ring 1 and does the following to the RIF before transmitting the frame toward ring 2: (see Figure 8-12):

1. Changes the destination's system's MAC address at the beginning of the frame to a Wellfleet Group Address. This address appears as C000A2FFFFFFx, where "x" is the Bridge ID of the next Wellfleet Bridge specified by the RIF.
2. Removes its own Internal LAN ID.
3. Inserts the Group LAN ID (before the last Incoming Ring/ Bridge ID listed in the RIF, see Figure 8-12). Eventually, the Group LAN ID will be replaced with the Internal LAN ID of the last Wellfleet Bridge along the frame's path.
4. Copies the destination system's MAC address into the data portion of the frame.

**Scenario B: Between the first and last Wellfleet Bridge to receive the SRF.**

The next Wellfleet Bridge B from ring 2 does the following to the RIF before transmitting the frame toward ring 3 (see Figure 8-12):

1. Locates the Bridge ID that is located at the end of the Group Address.
2. Changes the Bridge ID at the end of the Group Address to the Bridge ID of the next Wellfleet Bridge in the RIF. (Only if it differs from the value already in place. In this example, all Bridge IDs are the same, so the frame is not modified).

**Scenario C: Last of several Wellfleet Bridges to receive the SRF.**

The last Wellfleet Bridge to receive the frame does the following to the RIF before transmitting the frame toward ring 1 (see Figure 8-12):

1. Replaces the Wellfleet Group address with the destination MAC address that was saved to the data field.
2. Replaces the Group LAN ID with its own Internal LAN ID.

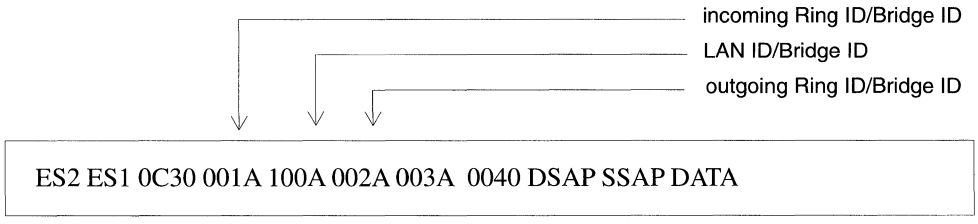


**Frame received by Bridge A**

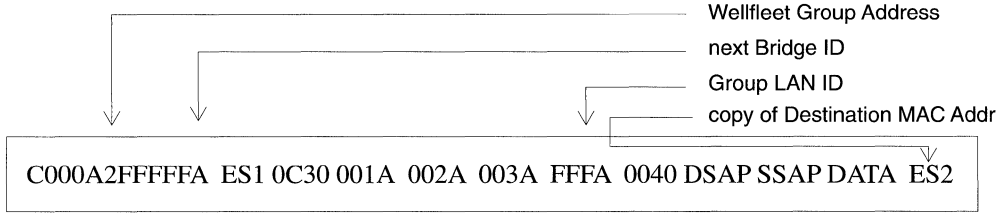


**Frame sent out by Bridge A onto ring 2**

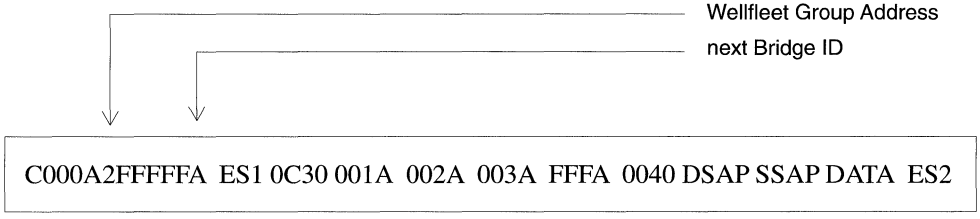


**Frame sent out by Bridge B onto ring 3**

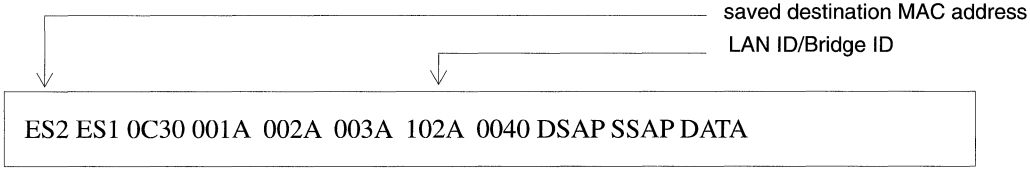


**Frame sent out by Bridge C onto ring 1**

**Figure 8-12. Structure of a Specifically Routed Frame from ES1 to ES2**

## Source Routing Across an IP Network

This section is describes how IP encapsulation works by tracing a specifically routed frame as it is sent out from End Station 1, traverses several Wellfleet Bridges and an IP network, and finally arrives at End Station 2 (see Figure 8-13).

**Note:** At this point, assume that explorer packets have traversed the network and identified the paths to all reachable interfaces.

To demonstrate IP encapsulation, this section traces a specifically routed frame as it is sent out from End Station 1 and finally arrives at End Station 2 as follows:

- Scenario A describes the actions of the first Wellfleet Bridge to handle the SRF. This bridge encapsulates the frame with an IP header before sending it out onto the IP network .
- Scenario B describes the actions of a Wellfleet Bridge that is in between the first and last Wellfleet Bridge to handle the SRF. This bridge removes the IP header from the frame before source routing it to the next bridge.
- Scenario C describes the actions of the last Wellfleet Bridge of several Wellfleet Bridges to handle the SRF.

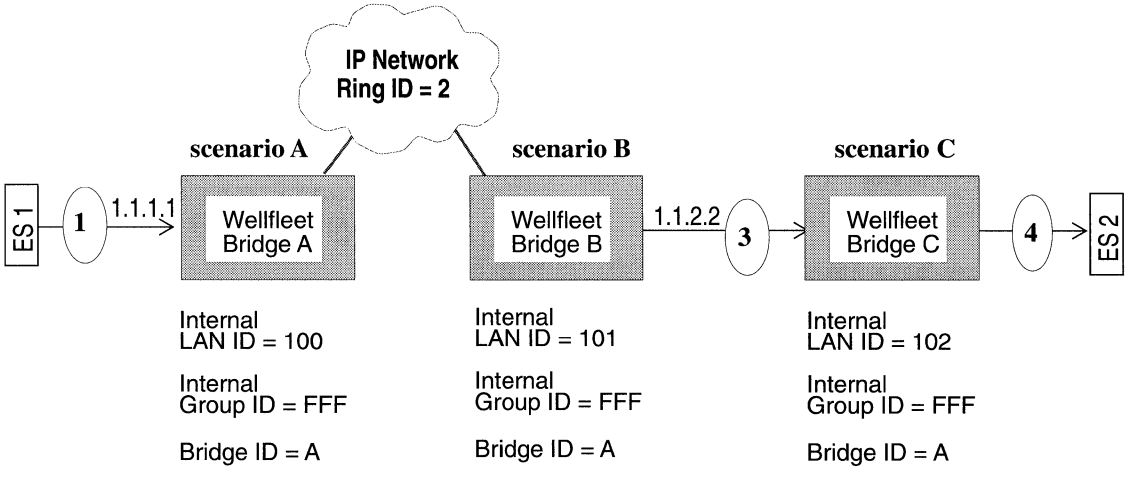


**Figure 8-13. Tracking an IP Encapsulated Frame From ES1 to ES2**

#### Scenario A: First of several Wellfleet Bridges to receive the SRF.

Wellfleet Bridge A is the first to receive the frame from ring 1 and does the following to the RIF before transmitting the frame onto the IP network (see Figure 8-14):

1. Removes its own Internal LAN ID.
2. Inserts the Group LAN ID (before the last Incoming Ring/ Bridge ID listed in the RIF, see Figure 8-14). Eventually, the Group LAN ID will be replaced with the Internal LAN ID of the last Wellfleet Bridge along the frame's path.
3. Adds an IP header containing the destination address 1.1.2.2 onto the frame and sends it toward the IP network.

#### Scenario B: Between the first and last Wellfleet Bridge to receive the SRF.

The next Wellfleet Bridge B receives the frame from the IP network (ring 2) and does the following to the RIF before transmitting the frame toward ring 3 (see Figure 8-14):

1. Strips the IP header from the packet.
2. Changes the destination system's MAC address at the beginning of the frame to a Wellfleet Group Address. This address appears as C000A2FFFFFFx, where "x" is the Bridge ID of the next Wellfleet Bridge specified by the RIF.
3. Copies the destination system's MAC address into the data portion of the frame.
4. Locates the Bridge ID that is located at the end of the Group Address, and sends it to the Bridge ID of the next Wellfleet Bridge in the RIF.

#### Scenario C: Last of several Wellfleet Bridges to receive the SRF.

The last Wellfleet Bridge to receive the frame does the following to the RIF before transmitting the frame toward ring 4 (see Figure 8-14):

1. Replaces the Wellfleet Group address with the destination MAC address that was saved to the data field.
2. Replaces the Group LAN ID with its own Internal LAN ID.

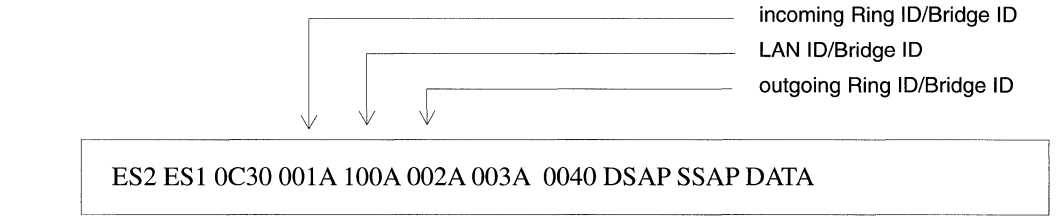


**Frame received by Bridge A**

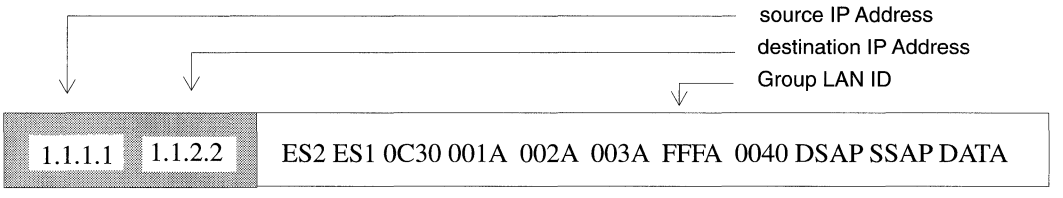


**Frame sent out by Bridge A onto the IP network (Ring 002)**

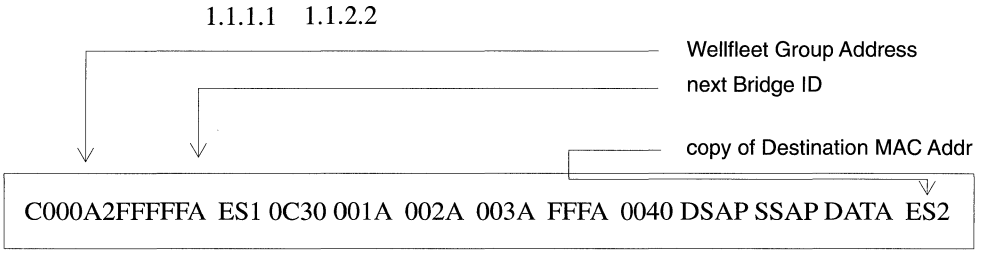


**Frame sent out by Bridge B onto ring 3**

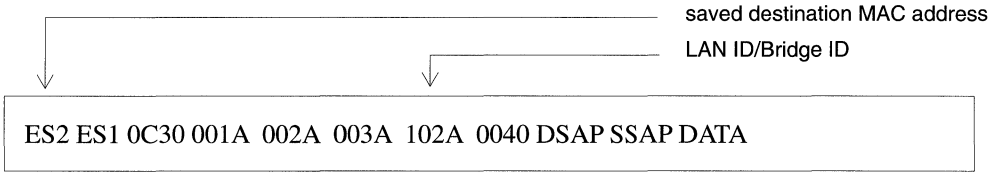


**Frame sent out by Bridge C onto ring 4**

**Figure 8-14. Structure of an IP Encapsulated Frame from ES1 to ES2**

# Source Routing Bibliography

The following documentation provides technical detail on Source Routing protocol implementation.

*IBM Token Ring Network Architecture Reference. Third Edition.* (IBM Corporation, September, 1989. SC3D-3374-D2).

Project 802.5m - Draft 7. *Source Routing Appendix to IEEE Standard 802.1d. Media Access Control (MAC) Bridges.* 1991.

Perlman, Radia. *Interconnections:Bridges and Routers.*, Addison - Wesley Publishing Company, Reading MA. First printing, May, 1992.

# Source Routing Implementation Notes

This section provides you with some basic guidelines on adding Wellfleet Source Routing Bridges to your network. It also addresses some of the special configuration features that may match your network requirements.

## Assigning Bridge IDs, Internal LAN IDs, and Group LAN IDs

When you enable the source routing bridge on the BN, you must specify its Bridge ID and Internal LAN ID. The source routing bridge uses these routing designators, together with another called the Group LAN ID, to source route packets through the network.

The Bridge ID is a standard source routing designator that identifies a bridge in the network. When you assign the Bridge ID, note the following guidelines:

- You must assign the same Bridge ID to all other Wellfleet Source Routing Bridges on the network that run Release 7.50 (*unless the bridges operate in parallel; see below).
- The Bridge IDs you assign to Wellfleet Bridges *must be unique* among all bridges on the network.

*If two or more Wellfleet Bridges operate in parallel, then you must assign *different* Bridge IDs to these bridges. In addition, you must specify the other Wellfleet Bridge ID in the Bridge Entry list for each bridge.

For example, in Figure 8-15, Wellfleet Bridges A, C and D are all assigned the same Bridge ID (A). However, because Bridge B and C operate in parallel, Bridge B is assigned a different Bridge ID (B). This Bridge ID (B) is then specified at the Bridge Entry list for Bridges A, C, and D so that they know that Bridge B is active on the network. (Similarly, the Bridge ID (A) is specified in the Bridge Entry list for Bridge B so that it knows Bridges A, C, and D are active on the network). Note that both Bridge IDs assigned to Wellfleet Bridges are unique within the network.

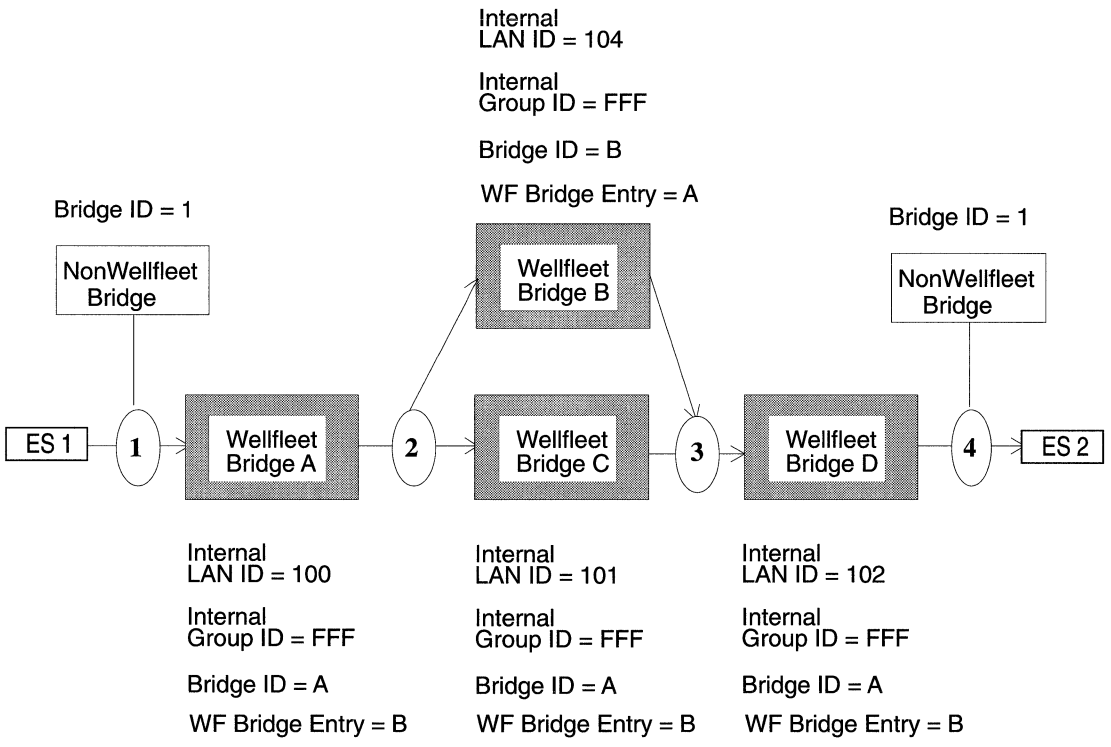


**Figure 8-15. Source Routing Packets Across a Token Ring Network**

The Internal LAN ID is a source routing designator that identifies the virtual LAN on which frames traverse. When you configure the Wellfleet Source Routing Bridge, assign a globally unique Internal LAN ID to each bridge on which source routing is enabled. For example, each source routing bridge A, B, C and D have all been assigned unique Internal LAN IDs (see Figure 8-15).

The Group LAN ID is a Wellfleet proprietary routing designator that helps the Wellfleet Source Routing Bridges in your network identify the last Wellfleet Bridge in its path. When you configure the Wellfleet Source Routing Bridge, note the following configuration guidelines:

- You must assign the *same* Group LAN ID to all Wellfleet Source Routing Bridges in the network.
- The value of the Group LAN ID you assign to all Wellfleet Bridges in your network must be *different* than the LAN ID value assigned to any other bridges in the network.

## Configuring IP Encapsulation Support

This release includes IP encapsulation support, which allows you to source route frames between Wellfleet Bridges over IP networks.

When you enable IP encapsulation on the source routing bridge, note the following guidelines:

- You must enable at least one IP interface on those BNs that you wish to source route packets through. You can enable IP on any circuit on any slot on the BN; it does not have to be the same circuit on which you have enabled source routing. See the chapter entitled *Configuring Circuits* for instructions on enabling IP on a circuit.
- If you enable redundant IP interfaces on the same BN for backup purposes, make certain that the circuits reside on different slots on the BN.
- You must enable source routing on the circuits of those Wellfleet Bridges through which you wish to source route frames. See the chapter entitled *Configuring Circuits* for instructions on enabling source routing on a circuit.
- You must enable IP encapsulation on the source routing bridges that connect to the IP backbone.

  To do this, reset the IP Encaps parameter to Enable for each source routing bridge. See the section entitled *Configuring Source Routing Global Parameters* for instructions.

- You must specify a Ring ID for the backbone IP network to which the source routing interface connects.

  To do this, specify the IP network's Ring ID at the Conn. IP NTWK Ring Number parameter for each source routing interface. Specify the same IP network Ring ID for each Wellfleet Source Routing Bridge that connects to the network. See the section entitled *Configuring Source Routing Interface Parameters* for instructions.

- You must specify the IP Explorer list for each source routing interface that connects to the IP Backbone.

  The IP Explorer List defines the IP addresses that will receive explorer frames from the bridge. See the section entitled *Adding or Deleting and IP Address from the IP Explorer Address List* for instructions.

- To reduce excess broadcast traffic on the network, you can create directed explorer filters. See the chapter entitled *Configuring Filters* for instructions.

# Editing Source Routing Parameters

Once you have configured a circuit to support source routing, you can use the Configuration Manager to edit source routing parameters. The configuration function you wish to perform determines the type of parameters you must edit. Table 8-1 lists each configuration function and the section that describes how to perform the function.

**Table 8-1. Source Routing Bridge Parameters and Configuration Functions**

| To Do the Following: | See this Section: |
|---|---|
| Change the state of the source routing software. | *Editing Source Routing Global Parameters* |
| Reconfigure source routing on a particular circuit. | *Editing Source Routing Interface Parameters* |
| Add or delete a Wellfleet Bridge ID from the Bridge Entry list. | *Adding or Deleting a Bridge ID from the Bridge Entry List* |
| Add or delete an IP address from the list of IP Explorer addresses. | *Adding or Deleting an IP Address from the IP Explorer Address List* |
| Configure source routing filters. | The chapter *Configuring Filters* |

For each source routing parameter, this section provides the following:

- Wellfleet default
- Valid setting options
- Parameter function
- Instructions for setting the parameter

You begin from the Wellfleet Configuration Manager window (see Figure 8-16).

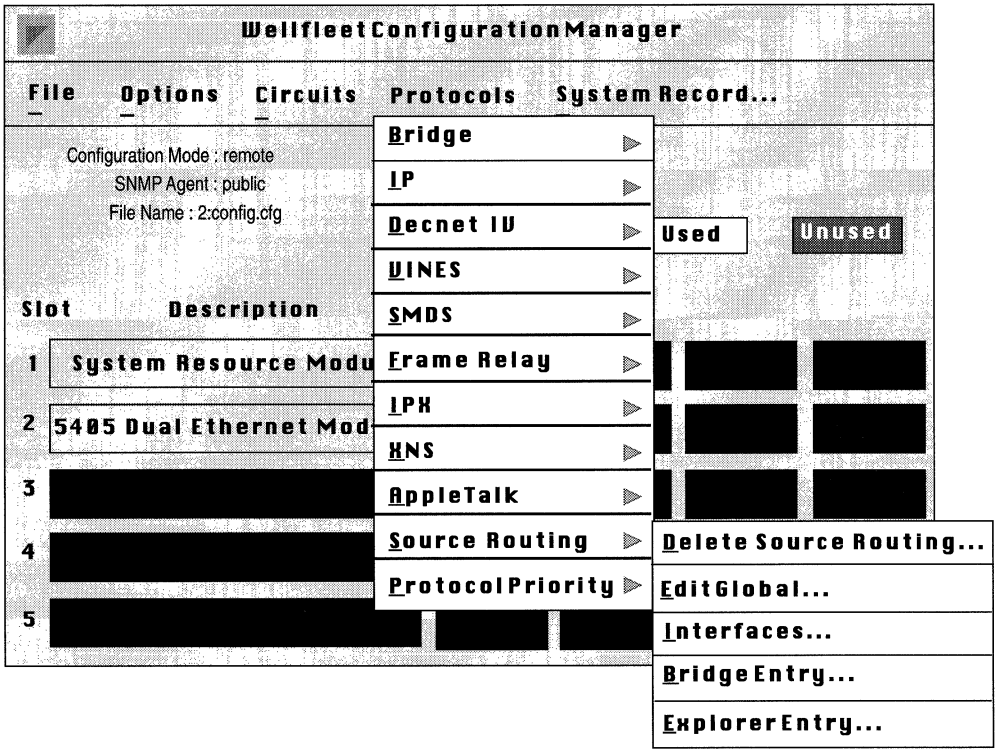


**Figure 8-16. Wellfleet Configuration Manager Window**

## Editing Source Routing Global Parameters

To edit source routing global parameters, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols/Source Routing/Edit Global option.

   The Source Routing Global Parameters window appears (see Figure 8-17).

2. Edit those parameters you wish to change.
3. Click the Save button to exit the window and save your changes when you are finished.

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or Disables the source routing on the entire BN. |
| Instructions: | Set to Disable if you want to disable source routing. |

| **Parameter :** | **SR Bridge Internal LAN ID** |
|---|---|
| Wellfleet Default: | None |
| Options: | 1 - 4095 |
| Function: | Specifies this bridge's Internal LAN ID. |
| Instructions: | Assign an Internal LAN ID that is unique among all other Internal LAN IDs and Ring IDs in the network. |

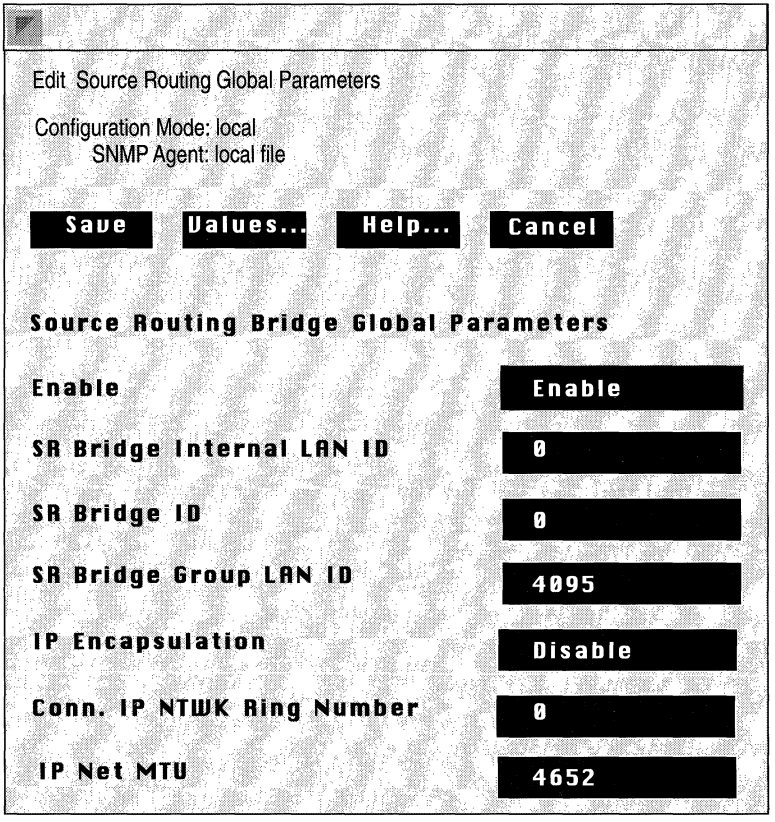


**Figure 8-17. Source Routing Global Parameters Window**

**Parameter :** **SR Bridge ID**

Wellfleet Default: None

Options: 1 - 15

Function: Specifies this bridge's Bridge ID and identifies the Wellfleet Source Routing Bridges in the network.

Instructions: Assign the same SR Bridge ID to all Wellfleet Source Routing Bridges in the network (unless two bridges operate in parallel; see note below). The SR Bridge ID must be unique among any other non-Wellfleet Bridge IDs in the network.

**Note:** If two Wellfleet Source Routing Bridges operate in parallel, then you must assign a different SR Bridge ID to one of the bridges. In addition, you must specify the SR Bridge ID in the Bridge Entry list for all other Wellfleet Source Route Bridges in the network. (See the section entitled *Assigning Bridge IDs, Internal LAN IDs and Group LAN IDs* in this chapter for more information).

**Parameter :** **SR Group LAN ID**

Wellfleet Default: 4095

Options: 1 - 4095

Function: Specifies this bridge's Group LAN ID. The bridge uses the Group LAN ID when transmitting Specifically Routed Frames (SRF) between Wellfleet Bridges. Together with the other routing designators, the Group LAN ID helps bridges manipulate the RIF.

Instructions: Assign the same Group LAN ID to all Wellfleet Source Routing Bridges in the network. The Group LAN ID must be unique among any other Group LAN IDs, Ring IDs or Internal LAN IDs in the network

**Parameter :** **IP Encapsulation**

Wellfleet Default: Disable

Options: Enable/Disable

Function: Enables IP Encapsulation for those packets destined for an IP network.

Instructions: Enable this parameter if this bridge borders an IP network cloud and you want to source route frames across this IP network. If you enable this parameter, you must also configure the Conn. IP NTWK Ring Number parameter below in order for IP encapsulation to occur.

See the section entitled *Configuring IP Encapsulation Support* for more information.

**Parameter :** **Conn. IP NTWK Ring Number**

Wellfleet Default: None

Options: 1 - 4095

Function: Identifies the Ring ID of the IP network to which this bridge connects.

Instructions: Assign the same Conn. IP NTWK Ring Number to all Wellfleet Source Routing Bridges that border the IP network cloud. In addition, make certain that the Conn. IP NTWK Ring Number is unique among any other Ring IDs, Group LAN IDs, or Internal LAN IDs in the network.

See the section entitled *Configuring IP Encapsulation Support* for more information.

| | |
|---|---|
| **Parameter :** | **IP Net Mtu** |
| Wellfleet Default: | 4562 |
| Options: | 1 - 4562 |
| Function: | Specifies the maximum Mtu size for the IP network. |
| Instructions: | Select a value that equals the smallest Mtu size of any of the links in the IP network. This allows the largest frame negotiation in the source routing exploration process to account for any link inside the IP cloud.<br><br>You can simply accept the default value 4562, however, if you have links in your IP network with smaller Mtu sizes than the default value, the IP entity may fragment packets. For maximum performance, refer to your network configuration and calculate this value based on actual Mtu sizes. |

## Editing Source Routing Interface Parameters

To edit a source routing interface, begin at the Wellfleet Configuration Manager window then proceed as follows:

1. Select the Protocols/Source Routing/Interfaces option to display the Source Routing Interfaces window (see Figure 8-18).

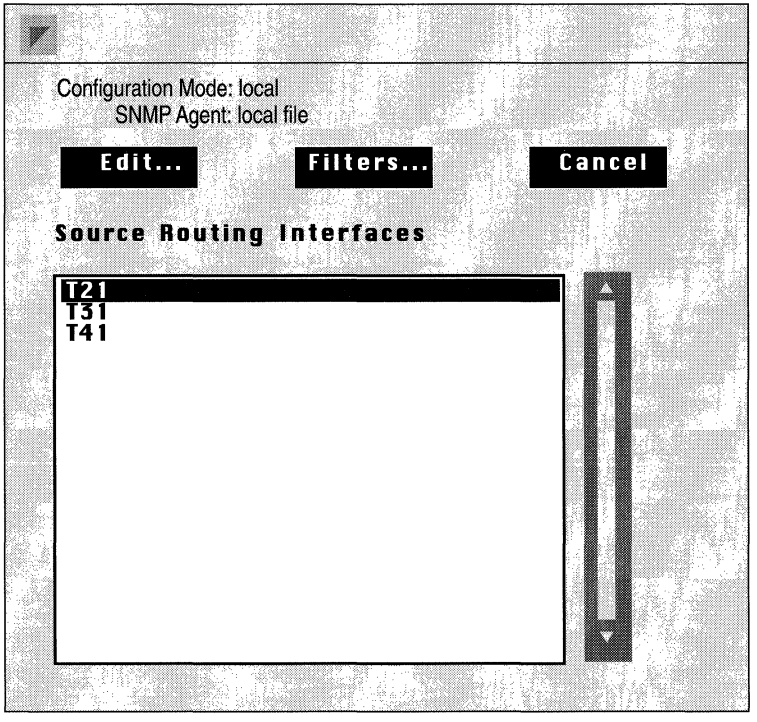


**Figure 8-18. Source Routing Interfaces Window**

2. Select the interface you wish to edit.
3. Click on the Edit button to display the Source Routing Interface Parameters window for that interface (see Figure 8-19).

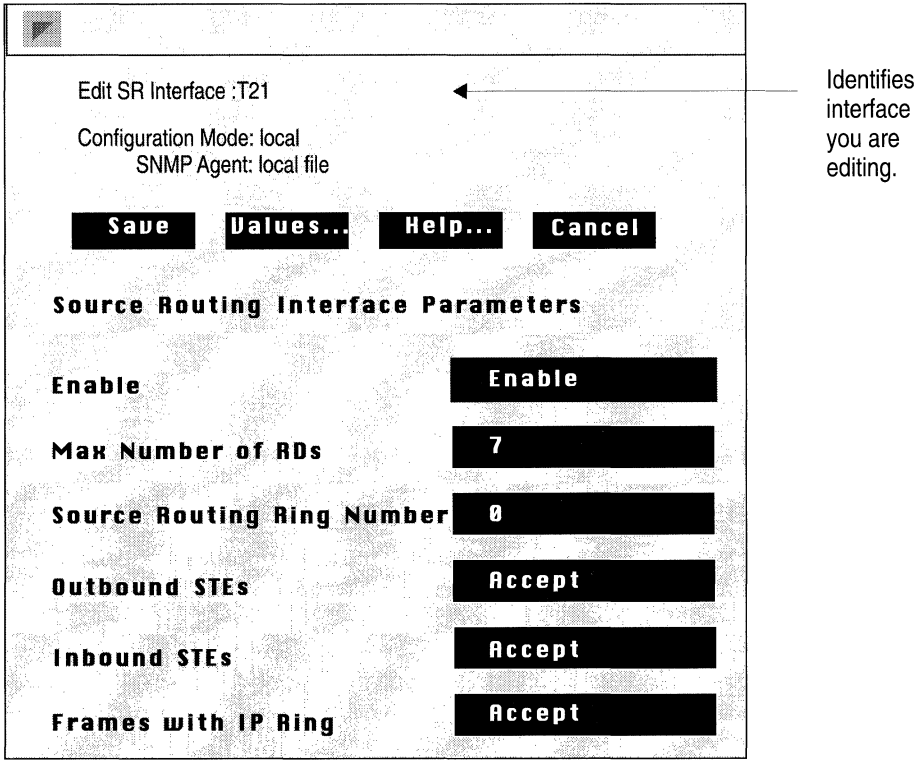


**Figure 8-19. Source Routing Interface Parameters Window**

4. Edit those parameters you wish to change.
5. Click the Save button to exit the window and save your changes when you are finished.

**Note:** When you reconfigure an interface in Dynamic mode, source routing restarts on that interface.

| | |
|---|---|
| **Parameter:** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or source routing over this circuit. |
| Instructions: | Set this parameter to Enable to enable source routing over this circuit. |

| | |
|---|---|
| **Parameter:** | **Max number of RDs** |
| Wellfleet Default: | 7 |
| Options: | 2-7 |
| Function: | Specifies the maximum number of route descriptors allowed in incoming explorer frames. Any explorer frames received that contain more route descriptors than specified here are dropped. |
| Instructions: | Accept the default unless you wish to limit which networks can be reached via this router. |

| | |
|---|---|
| **Parameter:** | **Source Routing Ring Number** |
| Wellfleet Default: | 0 |
| Options: | 1 - 4095 |
| Function: | Identifies the ring number (Ring ID) of this source routing circuit. |
| Instructions: | Assign a ring number (Ring ID) to this source routing circuit that is unique among any other Ring IDs, Group LAN IDs, or Internal LAN IDs in the network. |

| | |
|---|---|
| **Parameter :** | **Outbound STEs** |
| Wellfleet Default: | Accept |
| Options: | Accept/Block |
| Function: | When set to Block, drops STEs outbound on this circuit. This can be used to configure a static spanning tree for spanning tree explorer packets. |
| Instructions: | Set to Block only if you do not want STEs to be forwarded on this circuit. |

| | |
|---|---|
| **Parameter :** | **Inbound STEs** |
| Wellfleet Default: | Accept |
| Options: | Accept/Block |
| Function: | Specifies if the bridge should drop single route explorer frames received on this circuit. This option will not stop single route explorer frames from being transmitted on this circuit. |
| Instructions: | Set to Block only if you want spanning tree explorer packets to be dropped by this circuit. |

| | |
|---|---|
| **Parameter :** | **Frames with IP Ring** |
| Wellfleet Default: | Accept |
| Options: | Accept/Block |
| Function: | Specifies if the bridge should block inbound explorer frames received on this circuit that have already traversed the IP network via IP encapsulation. |
| Instructions: | Set to Block only if you want to limit the route selection possibilities using IP encapsulation. |

## Adding or Deleting a Bridge ID from the Bridge Entry List

You specify a Wellfleet Bridge ID for the Bridge Entry list if you are configuring two or more Wellfleet Source Routing Bridges that operate in parallel. (See the section entitled *Assigning Bridge IDs, LAN IDs and Group LAN IDs* in this chapter for more information).

Beginning at the Wellfleet Configuration Manager window, you add or delete a Wellfleet Bridge ID from the list as follows:

1. Select the Protocols/Source Routing/Bridge Entry option.

   The Source Routing Bridge IDs window appears (see Figure 8-20). It lists the other Wellfleet Bridge IDs that are currently configured on the network.

2. Add or Delete a Bridge ID from the list as described in the following sections.

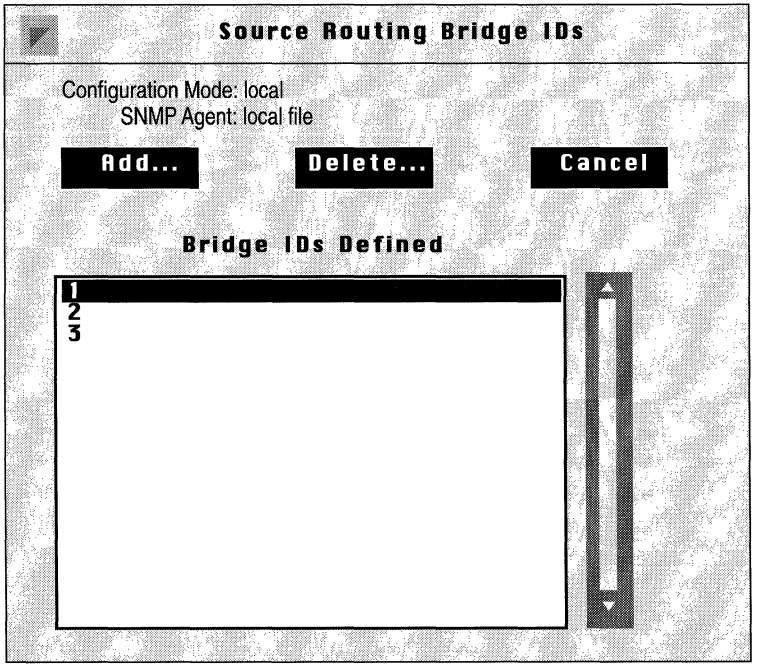


**Figure 8-20. Source Routing Bridge IDs Window**

### Adding a Wellfleet Bridge Entry

Beginning at the Source Routing Bridge IDs window, you add a Wellfleet Bridge ID to the Bridge Entry list as follows:

1. Click on the Add button.
2. Enter the new Bridge ID at the New Source Routing Bridge ID field (see the parameter description below).
3. Click on the Add SR BR ID button.

   The Source Routing Bridge IDs window appears. The Bridge ID you configured is added to the list.

| **Parameter :** | **New Source Routing Bridge ID** |
|---|---|
| Wellfleet Default: | None |
| Options: | 1 - 15 |
| Function: | Specifies the other active Wellfleet Bridge IDs that exist on the network. This parameter is only necessary if you have parallel Wellfleet Source Routing Bridges configured on your network. |
| Instructions: | Enter the Bridge ID assigned to the parallel Wellfleet Bridge on your network. See the section entitled *Assigning Bridge IDs, Internal LAN IDs, and Group LAN IDs* for more information about parallel source routing bridges. |

### Deleting a Wellfleet Bridge Entry

Beginning at the Source Routing Bridge IDs window, you delete a Wellfleet Bridge ID from the Bridge Entry list as follows:

1. Select a Wellfleet Bridge ID from the list.
2. Click on the Delete button.

   The Source Routing Bridge IDs window reappears. The Bridge ID you deleted has been removed from the list.

## Adding or Deleting an IP Address from the IP Explorer Address List

You specify an IP address for the IP Explorer Address list if you are source routing across an IP network. See the section entitled *Configuring IP Encapsulation Support* in this chapter for more information.

Beginning at the Wellfleet Configuration Manager window, you add or delete an IP Address from the list as follows:

1. Select the Protocols/Source Routing/Explorer Entry option.

   The IP Explorer Address window appears (see Figure 8-21). It lists the IP Explorer addresses defined.

2. Add or delete an IP Address from the list as described in the following sections.

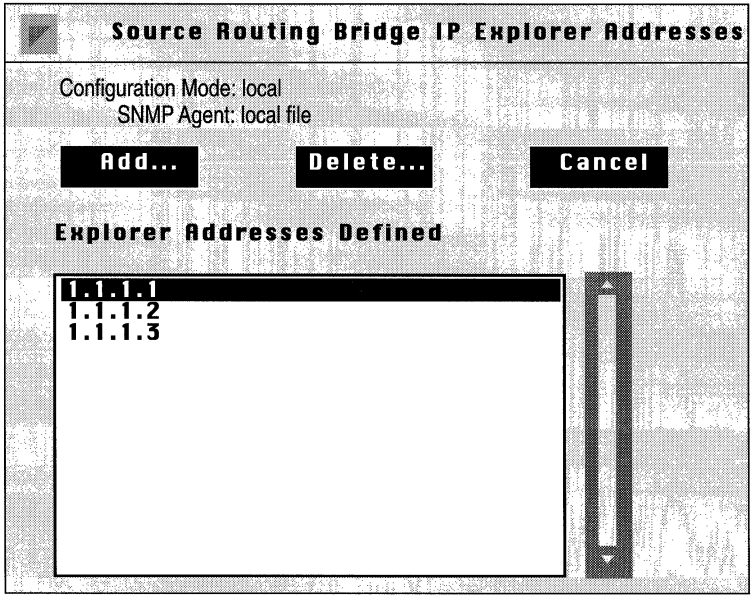


**Figure 8-21. Source Routing IP Explorer Address Window**

### Adding an IP Explorer Address

Beginning at the IP Explorer Address window, you add an IP Address to the list as follows:

1. Click on the Add button.
2. Enter the new IP Address at the New SR Bridge Explorer IP Address field (see the parameter description below).
3. Click on the Add SR BR EXP button.

   The IP Explorer Address window reappears. The IP Address you configured is added to the list.

| **Parameter :** | **New SR Bridge Explorer IP Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid IP address. |
| Function: | Specifies a destination IP address that this bridge can use to source route packets across an IP network. |
| Instructions: | Enter a valid destination IP address. See the section entitled *Source Routing Over IP Networks* in this chapter for more information about defining an IP explorer list. |

### Deleting an IP Explorer Address

Beginning at the IP Explorer Address window, you delete an IP Address from the list as follows:

1. Select an IP Address from the list.
2. Click on Delete button.

   The IP Explorer Address window reappears. The IP Address has been removed from the list.

# Deleting Source Routing from the BN

You can delete the source routing protocol from all BN circuits on which it is currently enabled in two steps.

Beginning from the Wellfleet Configuration Manager window, you delete source routing from the BN as follows:

1. Select the Protocols/Source Routing/Delete Source Routing option.

   A window pops up and prompts "Do you REALLY want to delete Source Routing?".

2. Select OK.

   You are returned the Wellfleet Configuration Manager window. Source Routing is no longer configured on the BN.

If you examine the Wellfleet Configuration Manager window, you will see that the connectors for circuits on which source routing was the *only* protocol enabled are no longer highlighted. Circuits must be reconfigured for these connectors; see chapter 3, entitled *Configuring Circuits* for instructions.

# Chapter 9

# Configuring DECnet Phase IV

## List of Figures



## List of Tables

# Configuring DECnet Phase IV

## About this Chapter

This chapter describes how to configure the DECnet Phase IV router. The first section provides an overview of DECnet Phase IV routing technology and describes how the Wellfleet DECnet Phase IV router works. The second section lists additional DECnet Phase IV reference material. The third section describes how to use the Configuration Manager to edit DECnet Phase IV parameters. The last section describes how to delete DECnet Phase IV from the BN.

## DECnet Phase IV Overview

This section describes the DECnet Phase IV network architecture. First, it tells you how end systems and routers in a DECnet network are organized. Next, it describes how the DECnet Phase IV router routes messages through the network. Finally, it describes how DECnet routing decisions are made.

## DECnet Network Organization

A DECnet network contains two types of nodes: end nodes and routers. End nodes send and receive messages. Routers (like the Wellfleet DECnet Phase IV Router) route messages to end nodes and other routers on the network.

DECnet Phase IV is designed to support large networks (in theory, over 64,000 nodes). Each network is divided into distinct areas (up to 63); each area contains up to 1023 nodes. Each area is assigned an Area ID, unique to the network. Each node within an area is assigned a Node ID, unique to the area.

Together, the Area ID and Node ID identify a 16 bit DECnet Phase IV address (see Figure 9-1); the first six bits identify the area in which the node resides, the last ten bits identify the node itself. Each DECnet address must be unique within the network.

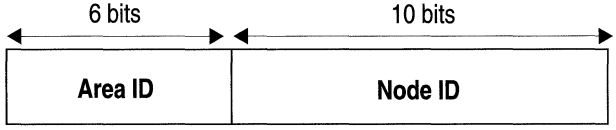


**Figure 9-1. DECnet Phase IV Address**

A Wellfleet DECnet Phase IV router can service multiple areas; that is, you assign addresses to each of the router's individual interfaces, rather to the entire router. Not all of the router's interfaces have to use the same address. For example, if a router resides in multiple areas (has circuits connecting to more than one area), it is assigned an Area ID and Node ID for each area to which it connects. As shown in Figure 9-2, circuit E21 and circuit E22 were assigned different DECnet addresses, even though they connect to the same router. Note that each circuit's address is still unique within its area and within the DECnet network.

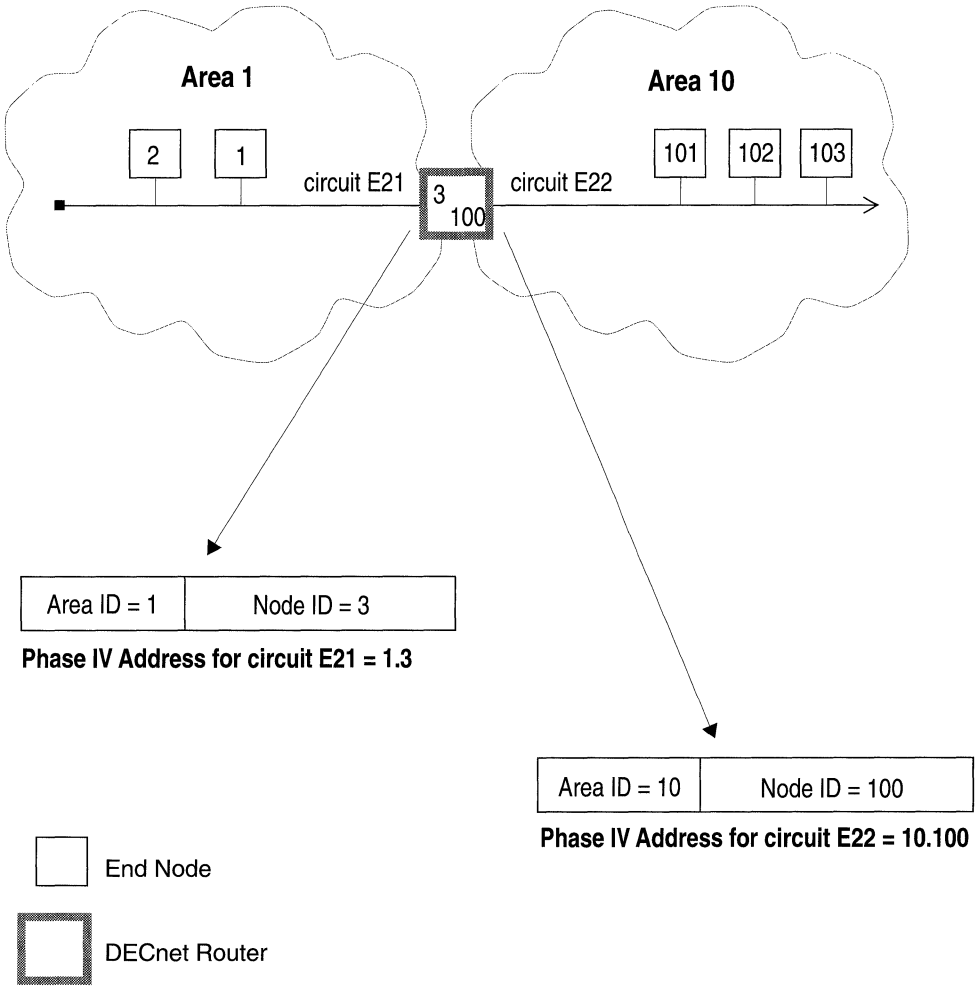


**Figure 9-2. Single DECnet Phase IV Router with Multiple Addresses**

## How the Wellfleet DECnet Router Works

The DECnet router uses a datagram service to route packets through the network. Routing within an area is called level 1 routing, or intra-area routing; routing between areas is called level 2 routing, or inter-area routing. The Wellfleet DECnet Router performs both type of routing services. That is, as a level 1 router, it maintains paths to systems within its local area. As a level 2 router, it maintains paths to all other areas within the DECnet network.

When a DECnet router receives a packet, it examines the destination address contained in the packet header. If the destination address is local, the router forwards the packet toward the destination system using the least cost path. If the destination address is to another area, the router forwards the packet toward the destination area; again, using the least cost path.

The router decides on the least cost path based on network topology and assigned circuit costs. If the least cost path is disabled, or a node fails, the router will find a different path, if one exists.

## How Routing Decisions Are Made

The DECnet Phase IV router uses the following four processes to make routing decisions: Update process, Listening process, Decision process, Forwarding process.

### Update Process

The router continually monitors the circuits directly attached to it. It periodically receives routing control messages from its adjacent neighbors. These routing updates inform the router of the current network topology. For example, if a circuit on a router fails, or another circuit is added, the network topology changes. The router then generates and transmits routing updates to all adjacent routers informing them of the changes. Routing updates describe which nodes in the local area are reachable (node updates), and which other areas in the network are reachable (area updates). Routers use this information to update their routing tables.

Timers control how often routing updates are sent out.

### Listening Process

The router periodically receives *Hello* messages from its adjacent neighbors. Hello messages inform the router of the identity of adjacent nodes, and identify the circuits that the router can use to reach the adjacent nodes. The router stores this information in an adjacency table; thus creating a database of "next hops" that it uses to forward data packets.

### Decision Process

Using information it retrieves from its routing tables, the router calculates the least cost paths from itself to all other systems that it can reach.

Every circuit on a DECnet network is assigned a default cost by the network manager. During the decision process, the DECnet router calculates the total path cost of forwarding a datagram along each possible path toward the destination. The total path cost is the sum of the costs of the circuits that make up the path. The least cost path is the one preferred by the router.

When deciding between multiple paths to a destination, the router will choose the path that is assigned a lower path cost over one assigned a higher cost, even if the lower cost path is longer (see Figure 9-3). If there is a tie between two paths, the router chooses the path whose next hop has the highest address. The amount of traffic on a circuit does not affect the path selected by the router.

Once the router determines the least cost path to a destination, it stores the identity of the corresponding adjacent router into its forwarding database. The adjacent router is the next hop on the path toward the destination. The decision process is executed separately for each routing level; the router keeps separate forwarding databases for intra-area and inter-area routing.

The complete distance (or number of hops) that a packet travels from the source to the destination is the path length. The maximum number of hops the routing algorithm will forward a packet to is called the maximum hops. If the distance between the source and destination exceeds the maximum hops, the packet is returned or discarded.

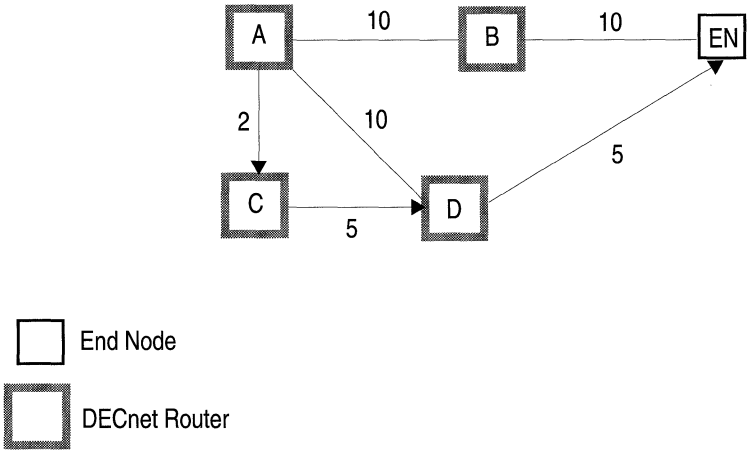


Router A wants to route a packet to the end node. Three different paths are available.
Router A forwards the packet along path choice 3 because it has the lowest Total Path Cost (12)

| Path Choice | Path Length | Path Cost |
|---|---|---|
| 1. A to B, B to EN | 2 hops | 20 |
| 2. A to D, D to EN` | 2 hops | 15 |
| 3. A to C, C to D, D to EN | 3 hops | 12 |

**Figure 9-3. Lowest Cost Path to a Destination**

## Forwarding Process

Whenever a router receives a packet, it checks the destination address to see if the packet needs to be routed locally (intra-area), or to another destination area (inter-area). If the destination is not known, or unreachable (for example, if the maximum hops value is exceeded), the router either returns the packet or discards it. Otherwise, it forwards the packet to the adjacent node specified in its forwarding database.

## Static Adjacency Support

The Wellfleet DECnet router creates and maintains its adjacency database by periodically broadcasting Hello messages to its neighbors. In DECnet terminology, an adjacency is a directly connected circuit-neighbor pair toward which packets are forwarded by the router (a neighbor is analogous to an adjacent host.) The neighbor may be a level 1 router, level 2 router or an end node.

This release of the DECnet Phase IV router allows you to configure *static adjacencies* for the router. Static adjacencies specify the DECnet address of the neighbor, the data link layer address of the neighbor, and the circuit used to reach the neighbor. Static adjacencies do not "age-out" of the router's adjacency database - even if the router never receives Hello messages from the neighbor. Thus, by configuring static adjacencies and disabling Hello message generation, you reduce the Hello message traffic traversing between the router and its neighbors.

For example, DECnet Routers A and B reside in the same area (see Figure 9-4).

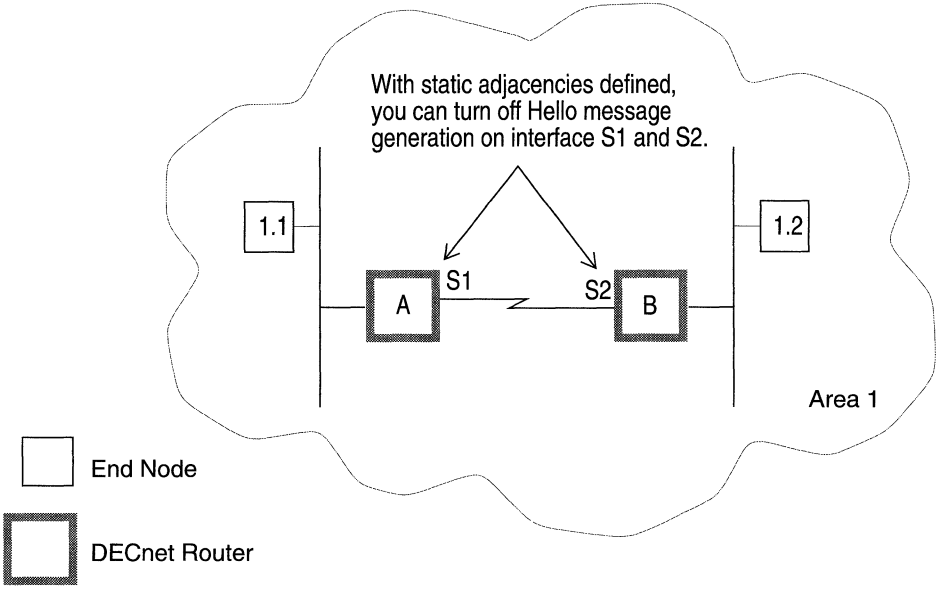


**Figure 9-4. Static Adjacencies Defined for Routers Residing in the Same Area**

The only network information they need to exchange is level 1 topology information. So in order to reduce traffic overhead, the network administrator did the following:

- Configured a static adjacency for interface S1 that specifies its neighbor - interface S2 on Router B.
- Configured a static adjacency for interface S2 that specifies its neighbor - interface S1 on Router A.
- Set the Routing Hello parameter (which enables or disables the generation of Hello packets by the interface) to Disable for both interfaces.

As a result, the routers know about each other's existence, even though no Hello messages are exchanged.

If two level 2 routers reside in different areas, you can reduce both Hello message traffic and level 1 topology traffic traversing between the two systems by disabling the Routing Hello parameter and enabling the Topology Update parameter (which keeps the DECnet router from sending level 1 topology update packets to other routers) for the connecting interfaces (see Figure 9-5).

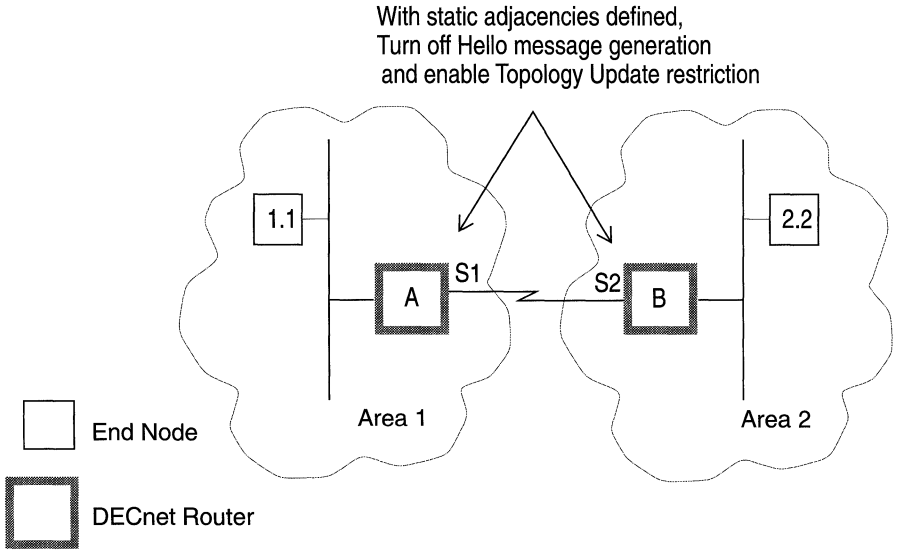


**Figure 9-5. Static Adjacencies Defined for Routers Residing in Different Areas**

See the section entitled *Adding, Editing or Deleting Static Adjacencies* for instructions on configuring static adjacencies. See the section entitled *Configuring DECnet IV Interface Parameters* for instructions on resetting the Router Hello and Topology Update parameters.

# DECnet Phase IV Bibliography

The following document provides technical detail on DECnet Phase IV protocol implementation.

*DECnet Digital Network Architecture Phase IV Routing Layer Functional Specification, Version 2.0* (Digital Equipment Corporation, December 1983.).

# Editing DECnet Phase IV Parameters

Once you have configured a circuit to support DECnet Phase IV, you can use the Configuration Manager to edit DECnet parameters. The configuration function you wish to perform determines the type of parameters you must edit. Table 9-1 lists each configuration function and the section that describes how to perform the function.

**Table 9-1. DECnet Phase IV Parameters and Configuration Functions**

| To Do the Following: | See this Section: |
|---|---|
| Change the state of the DECnet Phase IV router software. | *Editing DECnet Phase IV Global Parameters* |
| Reconfigure DECnet Phase IV on a particular circuit. | *Editing DECnet Phase IV Interface Parameters* |
| Specify a static adjacency for the DECnet Phase IV router's static adjacency list. | *Adding Editing or Deleting a Static Adjacency* |
| Configure DECnet Phase IV filters. | The chapter *Configuring Filters* |

For each DECnet Phase IV parameter, this section provides the following:

- Wellfleet default
- Valid setting options
- Parameter function
- Instructions for setting the parameter

You begin from the Wellfleet Configuration Manager window (see Figure 9-6).

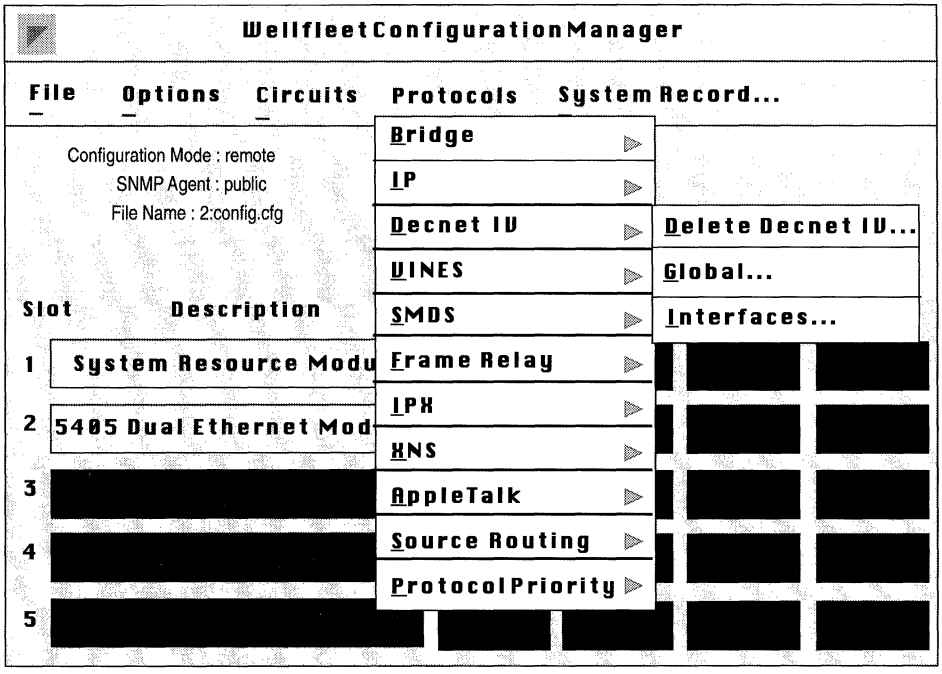


**Figure 9-6. Wellfleet Configuration Manager Window**

## Editing DECnet Phase IV Global Parameters

To edit DECnet Phase IV global parameters, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols/DECnet IV/Edit Global option to display the DECnet Phase IV Global Parameters window (see Figure 9-7).

   This section provides the information you need to edit each parameter.

2. Click the Save button to save your changes and exit the window.

| **Parameter :** | **Route Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or Disables the DECnet Phase IV router on the entire BN. |
| Instructions: | Set to Disable if you want to disable DECnet Phase IV. |

| **Parameter :** | **BroadCast Route Timer** |
|---|---|
| Wellfleet Default: | 180 |
| Options: | 1-65535 |
| Function: | Determines the maximum amount of time in seconds between routing updates issued by the router. If no routing update occurs before this timer expires, then a routing update is automatically generated. |
| Instructions: | Enter a number between 1 and 65535. If you want the BN to generate routing updates more frequently, then set this to a lower number. |

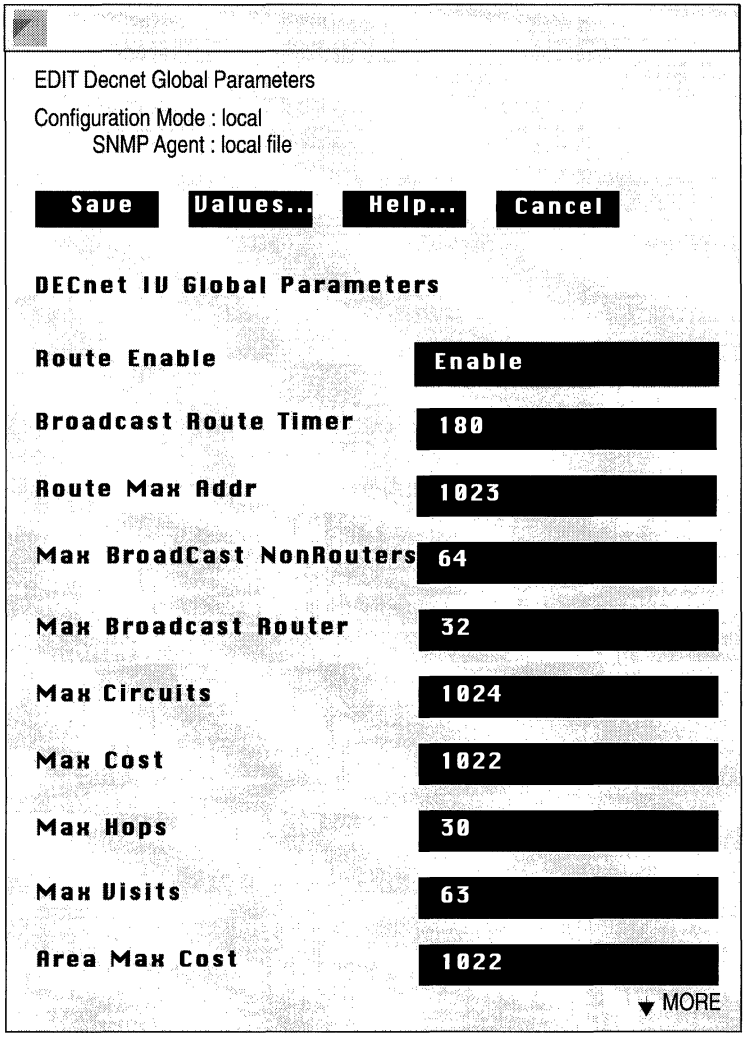


**Figure 9-7. DECnet Phase IV Global Parameters Window**

| | |
|---|---|
| **Parameter :** | **Route Max Addr** |
| Wellfleet Default: | 1023 |
| Options: | 1-1023 |
| Function: | Specifies the highest Node ID contained within all areas. |
| Instructions: | Wellfleet recommends accepting the default value. If you change the default, make certain to use the same value for each router in the network. |

| | |
|---|---|
| **Parameter :** | **Max BroadCast Non Routers** |
| Wellfleet Default: | 64 |
| Options: | 1-1023 |
| Function: | Specifies the maximum number of end node adjacencies residing on all circuits for a single slot. The higher the number of adjacent end nodes, the greater the impact on the router's performance and memory utilization. |
| Instructions: | Consult your network topology drawing. If there are more than 64 end node adjacencies on any of the router's slots, then increase this number to reflect your network topology. |

| | |
|---|---|
| **Parameter :** | **Max Broadcast Routers** |
| Wellfleet Default: | 32 |
| Options: | 1-1023 |
| Function: | Specifies the maximum number of router adjacencies on all circuits for a single slot. |
| Instructions: | Consult your network topology drawing. If there are more than 32 router adjacencies on any of the router's slots, then increase this number to reflect your network topology. |

| | |
|---|---|
| **Parameter :** | **Max Circuits** |
| Wellfleet Default: | 1024 |
| Options: | 1-1024 |
| Function: | Specifies the maximum number of circuits that this router can know about. |
| Instructions: | Wellfleet recommends accepting the default value. |

| | |
|---|---|
| **Parameter :** | **Max Cost** |
| Wellfleet Default: | 1022 |
| Options: | 1-1022 |
| Function: | Specifies the maximum path cost from this router to any destination node in the local area. The path cost is the sum of the individual circuit costs between this router and the destination node.<br>The router will declare a destination node unreachable if the least cost path to the destination node exceeds this number. |
| Instructions: | Determine the maximum path cost between this router and any node in the area, and enter it. |

| | |
|---|---|
| **Parameter :** | **Max Hops** |
| Wellfleet Default: | 30 |
| Options: | 1-30 |
| Function: | Specifies the maximum path length in hops between this router and any other destination node in the local area. A hop is the logical distance between two nodes. |
| Instructions: | Calculate the maximum path length in hops from this router to any other destination node in the area. Double the number you derive, and enter it. |

| | |
|---|---|
| **Parameter :** | **Max Visits** |
| Wellfleet Default: | 63 |
| Options: | 1-63 |
| Function: | Used to detect routing loops. That is, it enables the packet lifetime control, which limits the number of times a packet can pass through a router. If the router receives a packet that 1) is not destined for the router and 2) whose MaxVisits value is exceeded, the router will discard the packet because it has traversed too many nodes. |
| Instructions: | Determine the maximum path length (in hops) of between the two nodes furthest separated on the network. Enter a number that is at least as large as this value. |

| | |
|---|---|
| **Parameter :** | **Area Max Cost** |
| Wellfleet Default: | 1022 |
| Options: | 1-1022 |
| Function: | Specifies the maximum path cost from this router to any other area in the network. The router will declare a destination area unreachable if the least cost path to the destination area exceeds this number. |
| Instructions: | Determine the total path cost of the worst-case longest path between this router and any other area in the network and enter it here. |

**Parameter :** **Area Max Hops**

Wellfleet Default: 30

Options: 1-30

Function: Specifies the maximum path length in hops from this router to any other destination area in the network. A hop is the logical distance between two routers. The router will declare a destination area unreachable if the path length to the destination area exceeds this number.

Instructions: Determine the maximum path length in hops from this router to any other destination area in the network. Double the number you derive, and enter it here.

**Parameter :** **Max Area**

Wellfleet Default: 63

Options: 1-63

Function: Specifies the number of local areas in your DECnet network.

Instructions: Refer to your network topology map, then enter the number of areas in your network.

## Editing DECnet Phase IV Interface Parameters

You edit a DECnet Phase IV interface from the DECnet Phase IV Interface Parameters window for that interface. To edit a DECnet Phase IV interface, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols/DECnet IV/Interfaces option to display the DECnet Phase IV Interfaces window (see Figure 9-8).

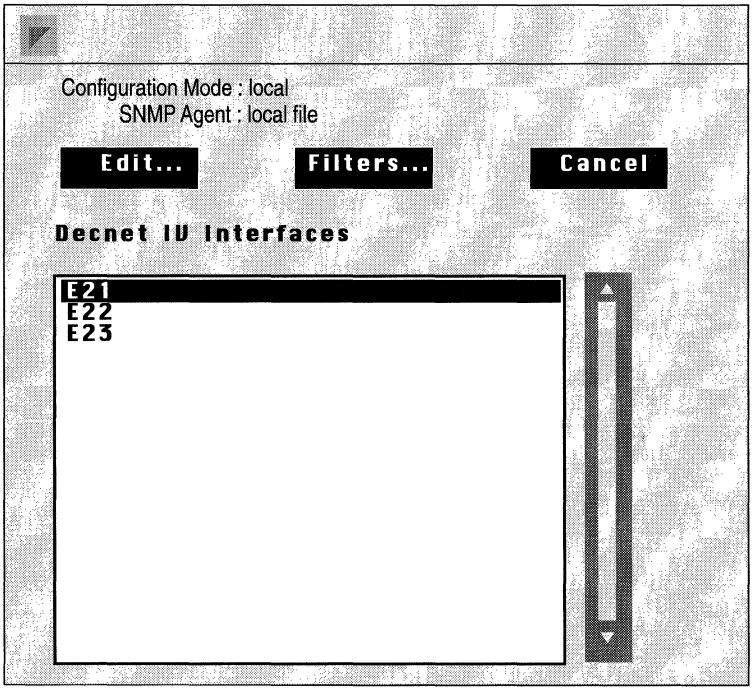


**Figure 9-8. DECnet Phase IV Interfaces Window**

2. Select the interface you wish to edit, then click on the Edit button to display the DECnet Phase IV Interface Parameters window for that interface (see Figure 9-9).

   This section provides the information you need to edit each parameter.

3. Click the Save button to exit the window and save your changes.

**Note:** When you reconfigure an interface in Dynamic mode, DECnet Phase IV restarts on that interface.

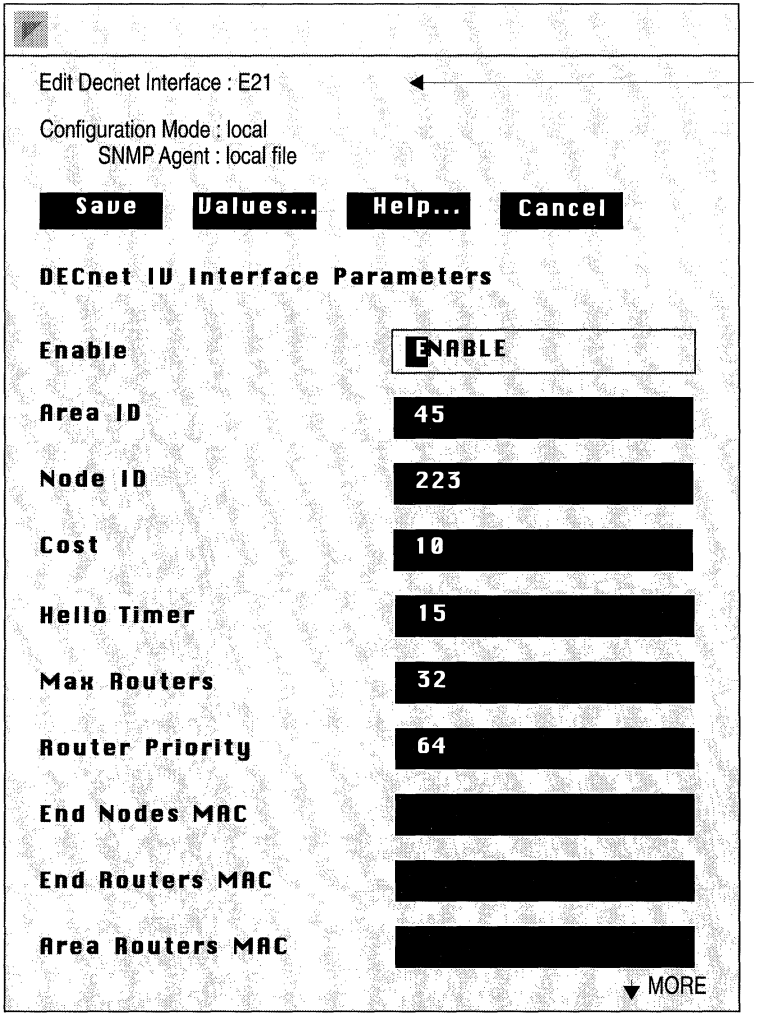


**Figure 9-9. DECnet Phase IV Interface Parameters Window**

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or Disables DECnet Phase IV over this circuit. |
| Instructions: | Disable only if you want DECnet disabled over this circuit. |

| | |
|---|---|
| **Parameter :** | **Area ID** |
| Wellfleet Default: | None |
| Options: | 1-63 |
| Function: | Specifies a DECnet Phase IV Area ID for this circuit.<br>The Area ID is the first six bits of a DECnet Phase IV node address.<br>You specify the Area ID on a circuit by circuit basis; that is a single BN may have individual circuits residing in different areas. |
| Instructions: | Enter the Area ID assigned to this circuit. |

| | |
|---|---|
| **Parameter :** | **Node ID** |
| Wellfleet Default: | None |
| Options: | 1-1024 |
| Function: | Specifies a unique DECnet Phase IV Node ID for this circuit.<br>The Node ID is the last 10 bits of a DECnet Phase IV node address. Note that individual circuits on a BN may have different Node IDs. |
| Instructions: | Enter the Node ID assigned to this circuit |

| | |
|---|---|
| **Parameter :** | **Cost** |
| Wellfleet Default: | 10 |
| Options: | 1-63 |
| Function: | Specifies the relative cost of routing over this circuit.<br><br>The sum of the individual circuit costs from a source node to a destination node is the *total path cost*. When the router receives a data packet, it decides which circuit to forwards the packet over based on *least cost path* toward the destination. |
| Instructions: | If you want the circuit to be used on a regular basis, then assign it a low cost; similarly, assign the circuit a high cost if you don't want it used on a regular basis. |

| | |
|---|---|
| **Parameter :** | **Hello Timer** |
| Wellfleet Default: | 15 |
| Options: | 1-8191 |
| Function: | Specifies in seconds how often the router broadcasts DECnet hello messages to all nodes on this Ethernet circuit. |
| Instructions: | Increase the Hello Timer if you want to reduce the amount of traffic traversing a slow line. |

**Note:** Inconsistent HelloTimer settings can cause confusion between DECnet routers and end nodes when rerouting occurs.

| | |
|---|---|
| **Parameter :** | **Max Routers** |
| Wellfleet Default: | 33 |
| Options: | 1-33 |
| Function: | Specifies the maximum number of routers attached to this Ethernet circuit, including the router itself. |
| Instructions: | Wellfleet recommends accepting the default value. If you change the default, refer to your network topology drawing to determine the number, then enter it here. |

| | |
|---|---|
| **Parameter :** | **Router Priority** |
| Wellfleet Default: | 64 |
| Options: | 0-127 |
| Function: | Dictates which router becomes the designated router on an Ethernet circuit. The designated router performs additional services for other nodes attached to the Ethernet circuit; all which know the address of the designated router.<br><br>When an end node attempts to send a package to a destination node that either 1) is not in its destination address cache, or 2) does not reside on the circuit, it sends the packet to the designated router, which forwards the packet towards the destination. |
| Instructions: | If you want this node to be the designated router on this circuit, then assign it the highest priority value among all routers on the circuit. (If you don't chose a designated router, the router with highest ID becomes the designated router by default). |

| | |
|---|---|
| **Parameter:** | **End Nodes MAC** |
| Wellfleet Default: | None |
| Options: | 0-127 |
| Instructions: | Assigns the AllEndnodes multicast MAC address value to this circuit for use over Frame Relay and SMDS WAN connections. The multicast MAC address value you specify here determines the destination address for DECnet end node broadcast traffic across the Frame Relay or SMDS cloud. This parameter is only valid if this circuit is a Frame Relay or SMDS circuit. |
| Instructions: | Enter the multicasting address (Frame Relay circuits) or group address (SMDS circuits) that you obtain from your service provider. If you do not specify a value for this parameter, then broadcast traffic will be sent to *all* VCs configured on the interface. |

| | |
|---|---|
| **Parameter:** | **End Routers MAC** |
| Wellfleet Default: | None |
| Options: | 0-127 |
| Instructions: | Assigns the AllRouters multicast MAC address value to this circuit for use over Frame Relay and SMDS WAN connections. This allows you to specify a destination address for DECnet Level 1 broadcast traffic across the Frame Relay or SMDS cloud. This parameter is only valid if this circuit is a Frame Relay or SMDS circuit. |
| Instructions: | Enter the multicasting address (Frame Relay circuits) or group address (SMDS circuits) that you obtain from your service provider. If you do not specify a value for this parameter, then broadcast traffic will be sent to *all* VCs configured on the interface. |

**Parameter :** **Area Routers MAC**

Wellfleet Default: none

Function: Assigns the AllAreaRouters multicast MAC address value to this circuit for use over Frame Relay and SMDS WAN connections. This allows you to specify a destination address for DECnet Level 2 broadcast traffic across the Frame Relay or SMDS cloud. This parameter is only valid if this circuit is a Frame Relay or SMDS circuit.

Instructions: Enter the multicasting address (Frame Relay circuits) or group address (SMDS circuits) that you obtain from your service provider. If you do not specify a value for this parameter, then broadcast traffic will be sent to *all* VCs configured on the interface.

**Parameter :** **Node Hello**

Wellfleet Default: Enable

Options: Enable/Disable

Function: When enabled, keeps the DECnet router from sending Hello packets to end nodes. This parameter is used to limit the amount of traffic sent over a WAN connection.

Instructions: To limit the amount of traffic traversing the WAN connection, Wellfleet recommends accepting the default Enable.

| | |
|---|---|
| **Parameter :** | **Router Hello** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | When enabled, keeps the DECnet router from sending Hello packets to other routers. This parameter is used to limit the amount of traffic sent over a WAN connection. |
| Instructions: | To limit the amount of traffic traversing the WAN connection, Wellfleet recommends accepting the default Enable. |

| | |
|---|---|
| **Parameter :** | **Topology Update** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | When enabled, keeps the DECnet router from sending Level 1 topology update packets to other routers. This parameter is used to limit the amount of traffic sent over a WAN connection. |
| Instructions: | To limit the amount of traffic traversing the WAN connection, Wellfleet recommends accepting the default Enable. |

## Adding, Editing, or Deleting a Static Adjacency

The Wellfleet DECnet router allows you to specify static transmission paths to adjacent hosts. The static adjacency may or may not be another router.

Beginning at the Wellfleet Configuration Manager window , you add, edit or delete a DECnet static adjacency as follows:

1. Select the Circuits/Edit option.

   The Circuit List window appears. It lists circuits that are configured on the router.

2. Select a DECnet circuit from this list and click on Edit.

   The Circuit Definition window appears.

3. Select the Protocols/Edit DECnet IV/Static Adjacency option.

   The DECnet Static Adjacencies window appears. It lists the static adjacencies that are currently configured for the circuit (see Figure 9-10).

Add, edit or delete a static adjacency from the list as described in the following sections.

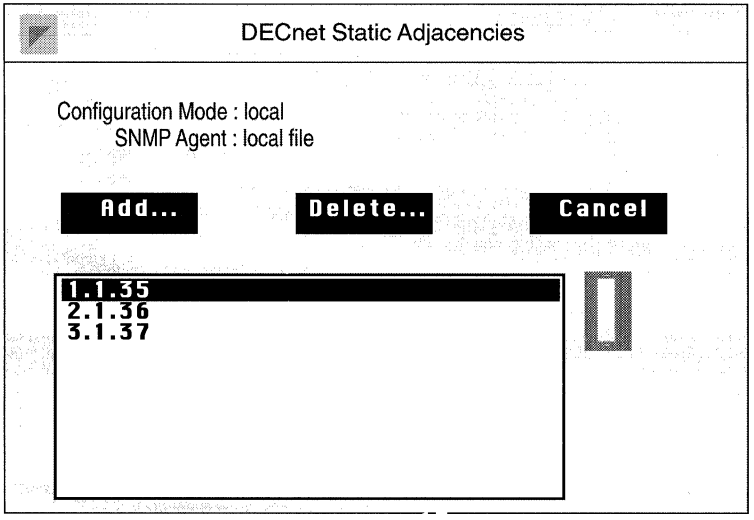


**Figure 9-10. DECnet Phase IV Static Adjacencies Window**

### Adding a Static Adjacency

Beginning at the DECnet Static Adjacencies window, you add a static adjacency to the list as follows:

1. Click on the Add button.
2. Enter the static adjacency information described in the following static adjacency parameters.
3. Click on the Add button.

   The Static Adjacencies window reappears. The static adjacency you configured is added to the list.

| | |
|---|---|
| **Parameter :** | **Adjacent Area ID** |
| Wellfleet Default: | None |
| Function: | 1-63 |
| Instructions: | Specifies the Area ID portion of the static adjacency's DECnet address. |
| Instructions: | Enter the Area ID assigned to the static adjacency. |

| | |
|---|---|
| **Parameter :** | **Adjacent Node ID** |
| Wellfleet Default: | None |
| Options: | 1-1024 |
| Function: | Specifies the Node ID portion of the static adjacency's DECnet address. |
| Instructions: | Enter the Node ID assigned to the static adjacency. |

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | The Configuration Manager automatically sets this parameter to Enable when you click on the Add button shown the Adjacent Host window. |
| Function: | Enable/Disable |
| Instructions: | Specifies the state of the static adjacency in the DECnet router's routing tables.<br><br>Select Disable to make the static adjacency record inactive in the DECnet routing table; the router will not consider this a static adjacency.<br><br>Select Enable to make this static adjacency record active again. |

| | |
|---|---|
| **Parameter :** | **Adjacent Type** |
| Wellfleet Default: | Area |
| Options: | Area, Routing IV, Non-Routing IV |
| Function: | Specifies whether the static adjacency is another router or is an end node. |
| Instructions: | Select Area if the static adjacency is a level 2 router, select Routing IV if the static adjacency is a level 1 router, or select Non-Routing IV if the static adjacency is an end node. |

| | |
|---|---|
| **Parameter :** | **Adjacent Priority** |
| Wellfleet Default: | None |
| Options: | 0-127 |
| Function: | If the static adjacency is another router, this parameter specifies the router's priority for becoming the designated router on the network. The designated router performs additional services for other nodes attached to the Ethernet circuit; all which know the address of the designated router. |
| Instructions: | If you want this node to be the designated router on this circuit, then assign it the highest priority value among all routers on the circuit. (If you don't chose a designated router, the router with highest ID becomes the designated router by default). |

| **Parameter :** | **Destination MAC Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid MAC address |
| Function: | Specifies the 48-bit Ethernet address of the static adjacency. |
| Instructions: | Enter the MAC address as a 12-digit hexadecimal number. |

### Editing a Static Adjacency

Beginning at the DECnet Static Adjacencies window, you edit a static adjacency that is already configured as follows:

1. Select a static adjacency from the list, then click on the Edit button.
2. Edit the static adjacency parameters you wish to change.
3. Click on the Save button.

   The DECnet Static Adjacencies window reappears.

### Deleting a Static Adjacency

Beginning at the DECnet Static Adjacencies window, you delete a static adjacency from the list as follows:

1. Select the static adjacency you wish to delete then click on the Delete button.

   The Delete DECnet Static Adjacency window appears.
2. Click on the Delete button to delete the static adjacency.

# Deleting DECnet Phase IV from the BN

You can delete DECnet Phase IV routing protocol from all BN circuits on which it is currently enabled in two steps.

You begin from the Wellfleet Configuration Manager window (see Figure 9-6) as follows:

1. Select the Protocols/Decnet IV/Delete Decnet IV option.

   A window pops up and prompts "Do you REALLY want to delete Decnet IV?".

2. Select Ok.

   You are returned the Wellfleet Configuration Manager window. DECnet Phase IV is no longer configured on the BN.

If you examine the Wellfleet Configuration Manager window, you will see that the connectors for circuits on which DECnet Phase IV was the *only* protocol enabled are no longer highlighted. Circuits must be reconfigured for these connectors; see chapter 3, entitled *Configuring Circuits* for instructions.

# Chapter 10

## Configuring IP

## List of Figures

## List of Tables

# Configuring IP

## About this Chapter

This chapter describes how to configure IP on your router. The first section provides an overview of IP routing technology and describes the features of the Wellfleet IP router. The second section describes how to use the Configuration Manager to edit IP parameters.

## IP Overview

IP routers interconnect networks that run the TCP/IP (Transmission Control Protocol/Internet Protocol) protocol suite. The Advanced Research Projects Agency (ARPA) of the U.S. Department of Defense (DoD), otherwise, known as DARPA, developed TCP/IP to enable computers manufactured by different vendors to interoperate.

In the early seventies, DARPA developed a large packet-switched network, called the ARPAnet. By the late seventies, the number of defense computers manufactured by different vendors and "speaking" different protocols had grown so high, that the DoD required all defense equipment to implement military-standard protocols (the TCP/IP protocol suite developed by DARPA) to ensure interoperability on the ARPAnet. At that time, no computer vendors implemented international standards.

With the success of TCP/IP, the original ARPAnet evolved into a large collection of networks called the Internet system. Today, the Internet system is both an operational and engineering-research network that carries real traffic and supports extensive protocol experimentation.

TCP/IP is now popular in the non-military market place. Numerous vendors have introduced TCP/IP-based products, and expenditures for such products have risen steadily. Many commercial users find TCP/IP a mature (although still evolving), robust, and reliable means of achieving multi-vendor interoperability. Some commercial users have adopted TCP/IP as an interim step while awaiting the widespread availability of OSI products.

## IP Router Functions

Both local-area networks (LANs) and wide-area networks (WANs) can run TCP/IP. Computer end-systems residing on networks that run TCP/IP are called hosts. A host can connect to more than one network, and, for reliability, can connect to the same network more than once.

IP routers enable hosts on networks that run TCP/IP to communicate. In doing so, an IP router performs three basic functions:

- Acquires knowledge of other routers and hosts on the network.

  IP routers use routing protocols to learn transmission paths (or routes) to other networks and to hosts residing on networks directly connected to the router.

- Builds the network topology based on the acquired information.

  IP routers store network-topology information in internal routing tables.

- Selects the best path, based on the information in its routing tables, for a particular datagram (a self-contained unit of data) to follow to reach its destination.

  IP Routers process each datagram individually. The datagram header provides the router with the destination IP address, as well as other routing information. Routers select a transmission path based on the IP address of the destination network, not of the destination host.

In short, routing is the process of identifying the routers that connect networks, and determining a transmission path to the destination network based on that information.

In performing the routing function, the IP router uses portions of the TCP/IP protocol suite. Table 10-1 lists the Internet Requests for Comments (RFCs) with which the Wellfleet IP router complies. This chapter assumes you are familiar with these RFCs.

**Table 10-1. Wellfleet IP Router RFC Support**

| RFC | Specifies |
|---|---|
| 768 | User Datagram Protocol (UDP) |
| 783 | Trivial File Transfer Protocol (TFTP) |
| 791 | Internet Protocol (IP) |
| 792 | Internet Control Message Protocol (ICMP) |
| 826 | Address Resolution Protocol (ARP) |
| 950 | Internet subnetting procedures |
| 1009 | Internet Gateways |
| 1058 | Routing Information Protocol (RIP) |
| 1063 | Maximum Transmission Unit (MTU) discovery option |
| 1247 | Open Shortest Path First (OSPF) Protocol Version 2 |
| 1157 | Simple Network Management Protocol (SNMP) |
| 1188 | IP over FDDI networks |
| 1042 | IP over IEEE 802.x networks |

## IP Datagrams

IP datagrams contain fields providing the following routing information:

- Type of Service

  The Type of Service field indicates the quality of service the datagram requires. The IP router inspects the Type of Service

field to obtain information about the datagram's precedence and expected delay characteristics.

- Time to Live

  The Time to Live field determines the datagram's lifetime in the internet system. Each time an IP router processes the datagram header, it decrements the value in the Time to Live field by at least one. When the value reaches zero, the IP router discards the datagram, unless it is destined for the router itself; thus, preventing undeliverable datagrams from looping endlessly through the network, consuming internet resources.

- Options

  The Options field may or may not be present in a datagram; thus, IP datagrams vary in length. There are three classes of Options:

  — Security

    Specifies security level and distribution restrictions.

  — Timestamps

    A 32-bit value measured in milliseconds since midnight universal time, or any other value if the high-order bit is set to 1.

  — Special Routing

    Specifies host-discovered paths to other hosts, or a specific path for the datagram to take.

- Header Checksum

  The Header Checksum field contains a value that the IP router calculates each time it processes a datagram's IP header. The algorithm used to calculate the checksum value is a 16 bit one's complement addition of the 16-bit words contained only within the IP header. The IP router discards datagrams received with an incorrect IP header checksum.

## Network Interfaces

Depending on the complexity of your network topology, the IP router is connected to at least two, and in most instances, more than two, TCP/IP networks. Each connection is an IP interface and has its own unique IP address, which identifies the network it connects. For example, in Figure 10-1, the IP router has three network interfaces.

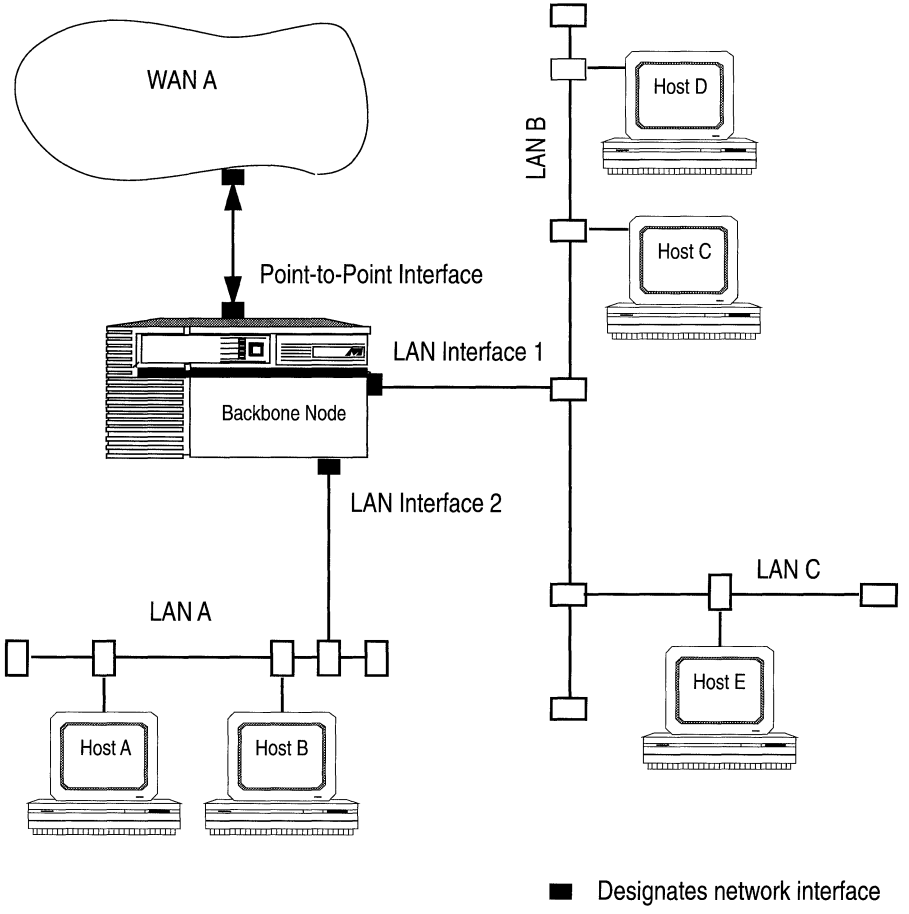


**Figure 10-1. IP Interfaces**

As shown in the figure, there are two types of interfaces. A LAN interface connects the router to an Ethernet or FDDI local-area medium. A point-to-point interface connects the router to a single long-haul medium terminated by a host or another router. Regardless of the interface type, each network interface requires a specific IP address.

The IP router can support multiple logical networks on a single physical network interface. For example, in Figure 10-1, LAN interface 1 provides a connection to both LAN B and LAN C.

## IP Address

An IP address consists of 32 bits having the form *network.host*. The network portion is a network number, ranging from 8 to 24 bits. The host portion is the remaining 24 to 8 bits identifying a specific host on the network. The Internet Network Information Center (NIC) assigns the network portion of the IP address. Your network administrator assigns the host portion.

The NIC recognizes three primary classes of networks (A, B, and C). In addition, the NIC has recently identified two other classes: Class D for networks that support multicasting, which allows an IP datagram to be transmitted to a single multicast group consisting of hosts spread across separate physical networks; and Class E for experimental networks. The Wellfleet IP router does not fully support Class D or Class E networks.

Based on the size of the network, the NIC classifies a network as either Class A, B, or C (the most common). The network class determines the number of bits assigned to the network and host portions of the IP address, as follows:

| Network Contains: | Class Is: | Network Portion Is: | Host Portion Is: |
|---|---|---|---|
| Over 65,534 hosts | A | 8 bits | 24 bits |
| 254 to 65,533 hosts | B | 16 bits | 16 bits |
| Under 254 hosts | C | 24 bits | 8 bits |

The position of the first bit set to 0 (whether it is the first, second, third, or fourth bit) in the first octet of an IP address indicates the network class (A, B, C, or D). If no bit is set to 0, it is a Class E network. Figure 10-2 shows the placement of the first bit set to 0 for Class A, B, and C networks. The figure also shows how a network's class affects the network and host portion of the IP address.

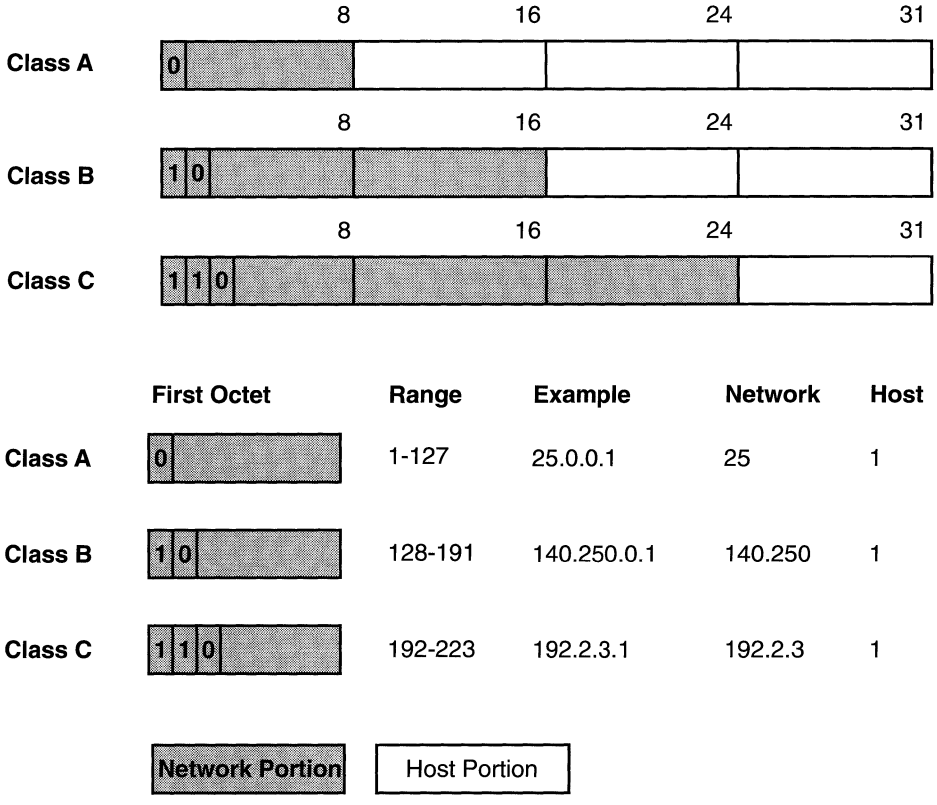


**Figure 10-2. Network and Host Portions of IP Address**

You specify IP addresses in dotted decimal notation. To specify an IP address in dotted decimal notation, you convert each 8-bit octet of the IP address to a decimal number, and separate the numbers by decimal points.

For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted decimal notation as 128.32.10.167. The most-significant two bits (10) in the first octet indicate that the network is Class B; therefore, the first 16 bits compose the NIC-assigned network portion field. The third octet (00001010) and fourth octet (10100111) compose the host field.

## Subnets

The concept of subnetworks (or subnets) extend the IP addressing scheme. Subnets are two or more physical networks that share a common network-identification field (the NIC-assigned network portion of the 32-bit IP address). Subnets allow an IP router to hide the complexity of multiple LANs from the rest of the internet.

With subnets, you partition the host portion of an IP address into a subnet number and a "real" host number on that subnet. The IP address is then defined by *network.subnet.host*. Routers outside the network do not interpret separately the subnet and host portions of the IP address.

Routers inside a network containing subnets use a 32-bit subnet mask which identifies the extension bits. In *network.subnet.host*, the *subnet.host* portion (or the local portion) contains an arbitrary number of bits. The network administrator allocates bits within the local portion to subnet and host, and then assigns values to subnet and host.

For example, the following is the IP address of a network that contains subnets: 10000000 00100000 00001010 10100111. You specify this address in dotted decimal notation as 128.32.10.167.

The second bit of the first octet is set to 0 indicating that the network is a Class B network. Therefore, the NIC-assigned network portion contains 16 bits, and the locally assigned local portion contains 16 bits.

The network administrator allocates the 16 bits in the local portion field, as follows:

- Allocates the upper-eight bits (00001010) with a value of 10 to the subnet portion.
- Allocates the lower-eight bits (10100111) with a value of 167 to the host portion.

In other words, the 16-bit local portion field, together with the 16-bit network field, specify host 167 on Subnet 10 of network 128.32.

You now need a subnet mask to identify those bits in the 32-bit IP address that specify the network field and those bits that specify the subnet field. Like the IP address, you specify the subnet mask in dotted decimal notation.

You construct a subnet mask, as follows:

- Assign a value of 1 to each of the 8, 16, or 24 bits in the network field.
- Assign a value of 1 to each bit in the subnet field.
- Assign a value of 0 to each bit in the host field.
- Convert the resulting 32-bit string to dotted decimal notation.

For example, to construct a subnet mask for the IP address described earlier (10000000 00100000 00001010 10100111), do the following:

1. Assign a value of 1 to each bit in the network field.

   The position of the first bit set to 0 in the first octet of the IP address indicates that the network is Class B; therefore the network field contains 16 bits: 11111111 11111111.

2. Assign a value of 1 to each bit in the subnet field.

   The network administrator allocated the upper-eight bits of the local portion to the subnet portion, as follows: 11111111.

3. Assign a value of 0 to each bit in the host field.

   The network administrator allocated the lower-eight bits of the local portion field to the host identification, as follows: 00000000.

4. Convert the resulting 32-bit string (11111111 11111111 11111111 00000000) to dotted decimal notation, as follows: 255.255.255.000.

## Multinet

The Wellfleet IP router supports multinet. As mentioned earlier, the network-layer connection to an IP network, called an IP interface, is assigned a unique IP address. Under each IP interface lies a data-link layer connection to the physical network, called a circuit. Multinet allows you to assign multiple IP addresses to a single circuit; thus, one circuit can support multiple IP network interfaces.

Multinet is commonly used in IP networks containing hosts that do not understand subnetting. For example, in Figure 10-3, hosts A, B, and C are connected by a router. Because the hosts do not understand subnetting, A, B, and C operate as if they are all on the same network. While A and C are on the same network, B is not. To facilitate connectivity between the three hosts, the router is configured with interfaces that connect three distinct subnets, as defined by the mask pair 255.255.255.0. In the figure, A and C are on a multinet interface.

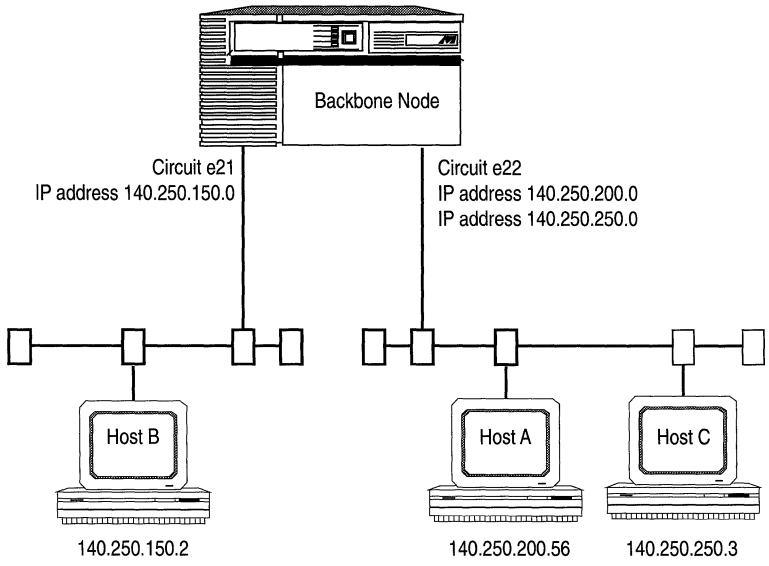


**Figure 10-3. Multinet Configuration**

## Broadcast Addresses

Broadcasting is when the IP router transmits a single packet to every host on an attached network. To do so, it uses a broadcast address that refers to all hosts on the network. A broadcast address is simply an IP address that contains all 1s or all 0s in the host portion.

For example, if you have an IP network with IP address 10.3.45.12, you can configure a broadcast address for that network, as follows:

- Because the address is for a Class A network (the network portion is 1 byte), the host portion contains three bytes.
- Because the host portion of a broadcast address consists of all 1s or all 0s, the broadcast address for that network can be one of the following: 10.255.255.255, 10.0.0.0, 255.255.255.255, or 0.0.0.0.

Some networks do not support broadcasts; thus, configuring an IP broadcast address does not guarantee efficient broadcast delivery.

### Subnet Broadcasts

The way you configure a broadcast address for a subnet is different from the way you configure a broadcast address for a network. Because you extend the network portion of the IP address when you create subnets, you automatically take away from the host portion of the IP address. To configure a subnet broadcast, you take the subnet mask for that subnet and invert it. For example, if the subnet's IP address is 10.4.2.3, and the mask is 255.255.0.0, then the subnet broadcast address is either 10.4.255.255. or 10.4.0.0.

## Routing Protocols

LANs and WANs interconnected by IP routers form a network of networks called an internet. For administrative purposes, an internet is segmented into autonomous systems. An autonomous system is simply a collection of routers and hosts. Figure 10-4 depicts a sample internet segmented into three autonomous systems.

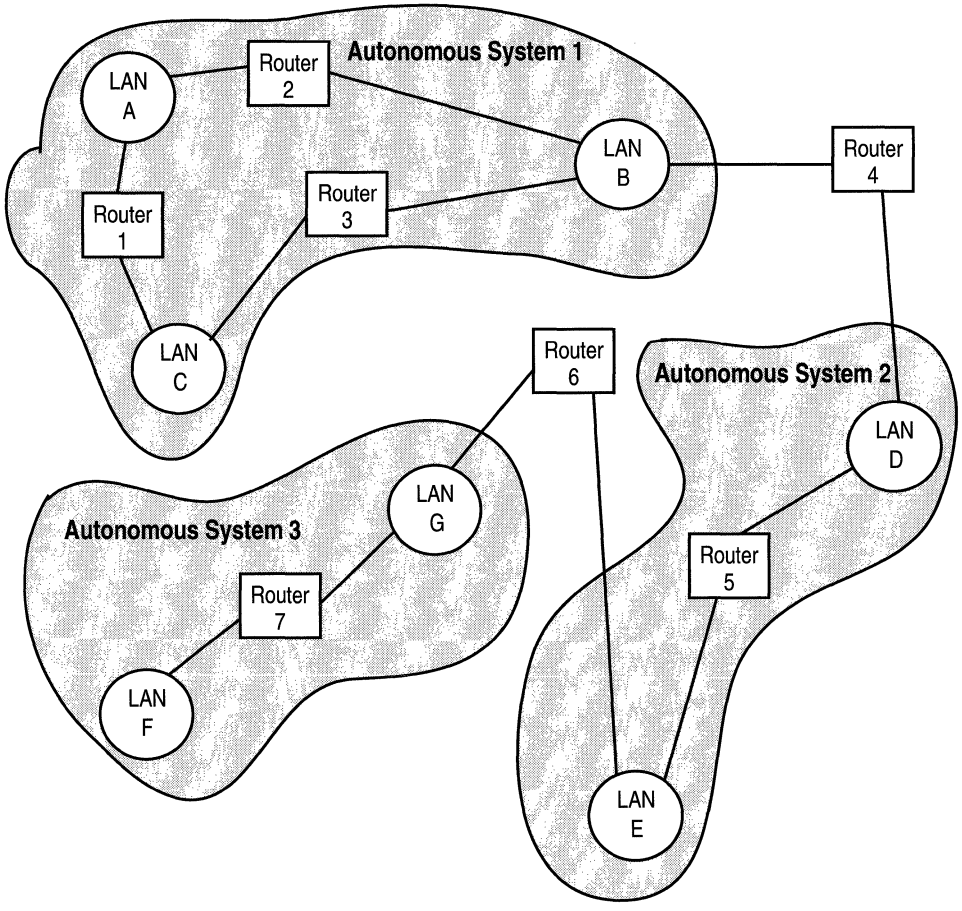


**Figure 10-4. Internet Segmented into Autonomous Systems**

Routers inside an autonomous system use an interior gateway protocol (IGP) to communicate network-topology changes to each other. Routers in separate autonomous systems use an exterior gateway protocol (EGP) to communicate. The Wellfleet IP router implements two dynamic IGPs: the Routing Information Protocol (RIP) and the Open Shortest Path First (OSPF) protocol. You enable RIP or OSPF for your BN.

### Routing Information Protocol (RIP)

The Routing Information Protocol (RIP) is a distance-vector protocol that enables routers in the same autonomous system to exchange routing information by means of periodic RIP updates. Routers transmit their own RIP updates to neighboring networks and listen for RIP updates from the routers on those neighboring networks. Routers use the information in the RIP updates to update their internal routing tables. For RIP, the "best" path to a destination is the shortest path (the path with the fewest hops). RIP computes distance as a metric, usually the number of hops (or routers) from the origin network to the target network.

### Open Shortest Path First (OSPF) Protocol

The Open Shortest Path First (OSPF) Protocol is an IGP intended for use in large networks. It exchanges routing information, using a link state algorithm, between routers in an Autonomous System. Routers synchronize their topological databases. Once the routers are synchronized and the routing tables are built, the routers will flood topology information only in response to some topological change. For OSPF, the "best" path to a destination is the path that offers the least cost metric delay. In OSPF, cost metrics are configurable; thus, allowing you to specify preferred paths. A more detailed overview, and OSPF configuration information is provided in *Configuring OSPF.*

## Static Routes

Static routes are manually configured routes that specify the transmission path a datagram must follow based on the datagram's destination address. A static route specifies a transmission path to another *network*. You configure a static route if you want to restrict the paths that datagrams follow to paths you specifically configure. In this situation, you would disable all dynamic routing capabilities, as well as disable all RIP and default-route supply and listen activities. You would then configure all routes statically. Unlike routes learned through RIP, static routes remain in IP routing tables until you remove them.

## Adjacent Hosts

An adjacent host specifies a transmission path to a *network device*, which may or may not be a router, but which must reside on a locally attached network. You configure an adjacent host if your topology includes a network or hosts that do not implement ARP. In this situation, you need to configure an adjacent host to each non-ARP device that resides on a network directly connected to the router.

Also, if a local network does implement ARP, you may wish to configure adjacent hosts to preempt the ARP process, and thereby, pre-resolve the Ethernet address.

## Routing Table Hierarchy

The IP router maintains internal routing tables. When making a decision on how to route a datagram, the IP router consults these tables to determine the specific route a datagram should take. However, it is possible for the routing tables to contain multiple routes to the same destination, because the routing tables can contain direct routes for the IP router's network interfaces, static routes (adjacent hosts are maintained in a separate table), as well as the routes the router learned from RIP (if you enabled RIP) or OSPF (if you enabled OSPF). By default, the IP router selects routes in the following order:

- direct
- OSPF internal
- static
- OSPF external
- RIP

## Routing Table Management

The IP router allows you to control the flow of routing data to and from the routing tables. This control is provided by user-configured import and export route filters:

- Import route filters
- Export route filters

### Import Route Filters

Import route filters govern the addition of new RIP-derived or OSPF-derived routes to the routing tables. When RIP or OSPF receives a new routing update, it consults its import route filter(s) to validate the information before entering the update into the routing tables. Import route filters contain search information (to match fields in incoming routing updates) and action information (to specify the action to take with matched fields).

### Export Route Filters

Export route filters govern the propagation of RIP or OSPF routing information. When preparing a routing advertisement, RIP or OSPF consults its export route filter(s) to determine whether the routes to specific networks are to be advertised and how they are to be propagated. Export route filters contain network numbers (to associate a filter with a specific network) and action information (to specify a route propagation procedure).

## Address Resolution Protocol (ARP)

The IP router needs both a physical address and an IP address to transmit a datagram. The Address Resolution Protocol (ARP) enables a router to determine a network host's physical address, when all it knows is the network host's IP address, by binding a 32-bit IP address to a 48-bit Ethernet address. A router can use ARP across a single network only, and the network hardware must support physical broadcasts.

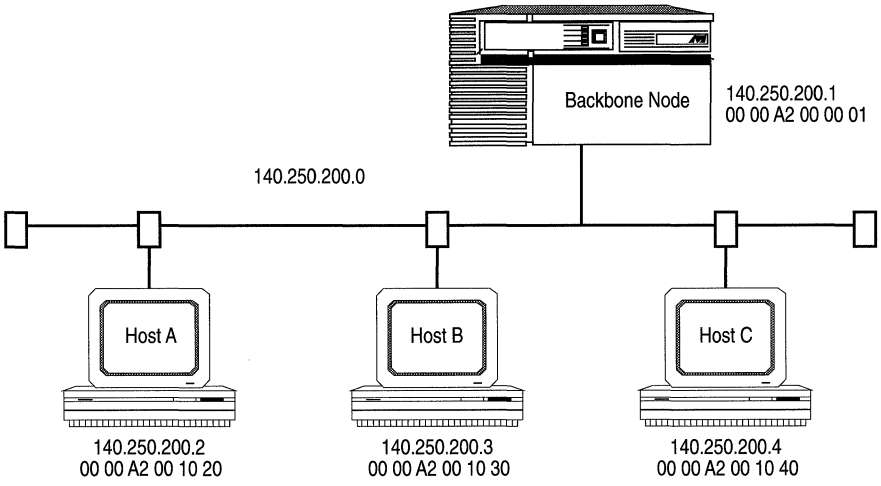


**Figure 10-5. ARP Example**

For example, in Figure 10-5, the BN and Host C are on the same physical network. Both devices have an assigned IP address (the BN's is 140.250.200.1 and Host C's is 140.250.200.4) and both devices have an assigned physical address (the BN's is 00 00 A2 00 00 01 and Host C's is 00 00 A2 00 10 40).

In the figure, the BN wants to send a packet to Host C, but only knows Host C's IP address. The BN uses ARP to determine Host C's physical address, as follows:

- The BN broadcasts a special packet, called an ARP request, that asks IP address 140.250.200.4 to respond with its physical address.
- All network hosts receive the broadcast request.
- Only Host C responds with its hardware address.
- The BN maps Host C's IP address (140.250.200.4) to its physical address (00 00 A2 00 10 40) and saves the results in an address-resolution cache for future use.

## Proxy ARP

Proxy ARP allows a router to answer a local ARP request for a remote destination. For example, in Figure 10-6, Hosts B and C are on the same internet, but are on separate subnetworks. Hosts B and C do not understand subnetworking. The BN connecting the two physical networks knows which host resides on which network. The address mask is 255.255.255.000. In this example, one subnet is a remote network compared to the other subnet.

Host B wants to talk to Host C, so Host B broadcasts an ARP request, which asks IP address 140.250.250.2 to respond with its physical address. The BN captures Host B's ARP request and responds with its hardware address 00 00 A2 00 00 01 and Host C's IP address 140.250.250.2. Host B maps Host C's IP address 140.250.250.2 to the BN's hardware address 00 00 A2 00 00 01.

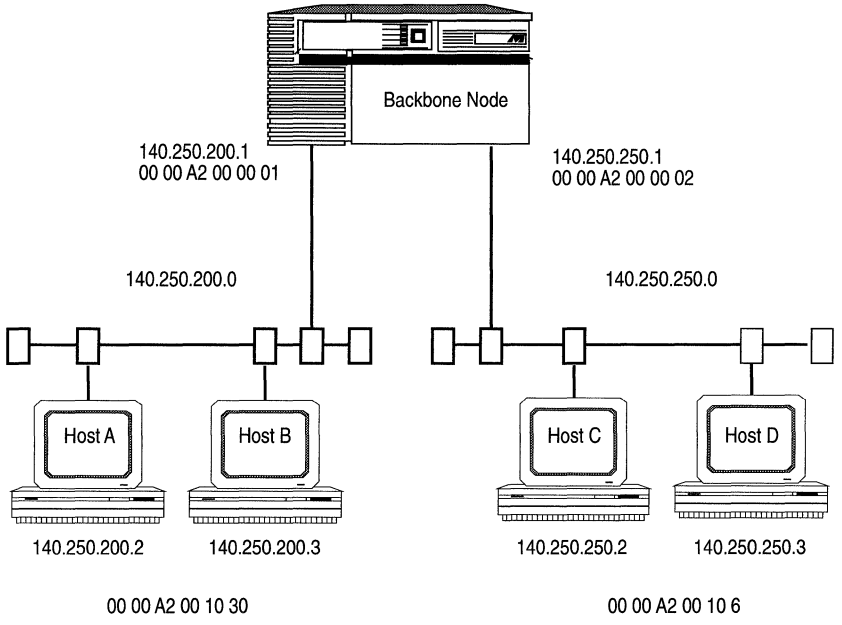


**Figure 10-6. Proxy ARP Example**

## Configuring the IP Router to Source Route Over Token Ring Networks

The Wellfleet IP router can route over token ring networks that contain one or more source routing bridges.

In a source routing network, every end station that sends out a frame supplies the frame with the necessary route descriptors so that it can be source routed across the network. Thus, in order for IP routers to route packets across a source routing network, *they must act like end stations*; supplying route descriptors for each packet before they send it out onto the network.

With end node support enabled, whenever a Wellfleet IP router receives a packet and determines that the packet's next hop is located across a source routing network, the router does the following:

- Adds the necessary RIF information to the packet's MAC header.
- Sends the packet out onto the network where it is source routed toward the next hop.

Upon receiving the packet from the token ring network, the peer router strips off the RIF field and continues to route the packet toward the destination network address (see Figure 10-7).

You configure source route end node support on a per-circuit basis by setting the TR End Station parameter to Enable. See the section entitled *Editing IP Interface Parameters* for instructions on enabling this parameter.

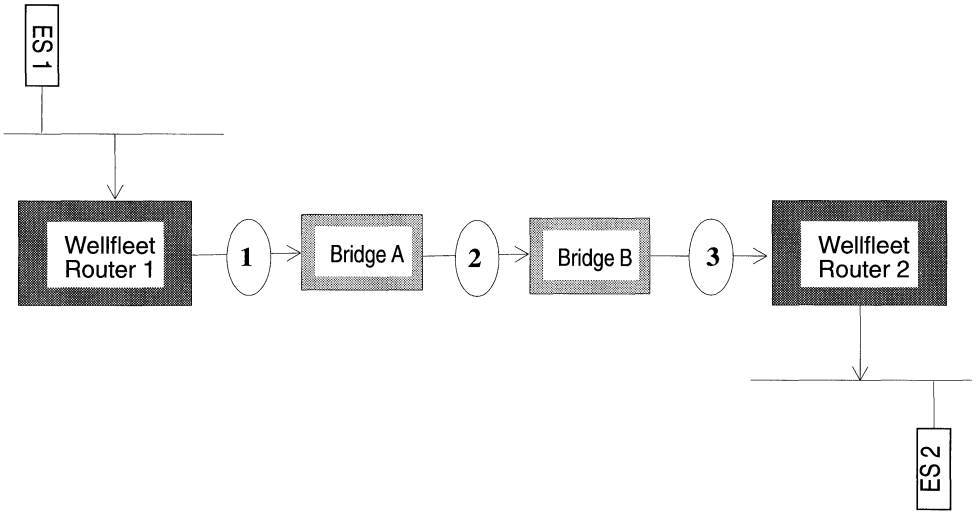


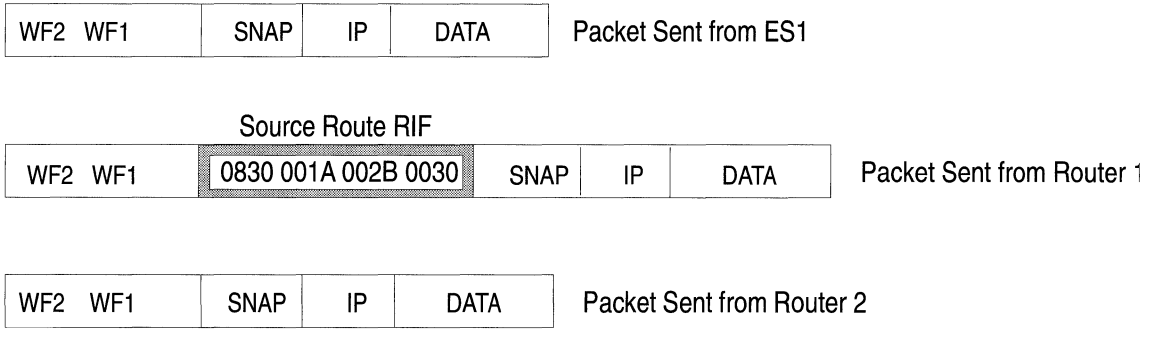


**Figure 10-7. IP Routers Source Routing Across a Token Ring Network**

## Trivial File Transfer Protocol (TFTP)

The Trivial File Transfer Protocol (TFTP) is a TCP/IP standard protocol for transferring files with minimum capability and minimal overhead. TFTP is implemented on top of the unreliable, connectionless datagram delivery service (UDP) and is used to move files between network devices that implement UDP.

TFTP was designed to be small and easy to implement. Because it is small, it is more restrictive, lacking most of the features of the File Transfer Protocol (FTP). TFTP provides inexpensive, unsophisticated file-transfer service only. It cannot list directories and provides no authentication.

TFTP runs on top of UDP and uses time-out and retransmission to ensure that data arrives. Each file transfer begins with a request to

read or write a file; this request also serves to request a connection. If the server grants the request, the connection is opened and the file is sent in fixed-length blocks (data packets) of 512 bytes. Each data packet contains one block of data and must be acknowledged by an acknowledgment packet before the next packet is sent. A data packet of less than 512 bytes terminates the transfer.

If a packet gets lost in the network, the intended recipient will time-out and may retransmit its last packet (which can be data or an acknowledgment), causing the sender of the lost packet to retransmit the packet. Because the lock-step acknowledgment guarantees that all older packets have been received, the sender keeps one packet only on hand for transmission.

Both devices involved in a TFTP transfer are senders and receivers. One device sends data and receives acknowledgments; the other device sends acknowledgments and receives data.

The IP router includes a client and server implementation of the Trivial File Transfer Protocol (TFTP) enabling the router to transmit and receive files across an internet.

# Editing IP Parameters

Once you have configured a circuit to support IP, you can use the Configuration Manager to edit IP parameters. The configuration function you wish to perform, determines the type of parameters you must edit. Table 10-2 lists each configuration function and the section that describes how to perform the function.

**Table 10-2. IP Parameters and Configuration Functions**

| To Do the Following: | See this Section: |
|---|---|
| Change the state of the IP router software.<br>Change whether or not the IP router routes IP traffic.<br>Change the TTL counter or RIP Diameter value that the IP router uses. | *Editing IP Global Parameters* |
| Reconfigure IP on a particular interface. | *Editing IP Interfaces* |
| Reconfigure the Routing Information Protocol (RIP) on an interface. | *Editing Routing Information Protocol (RIP) Interfaces* |
| Reconfigure OSPF on a particular interface. | The chapter *Configuring OSPF* |
| Add, edit, and delete static routes. | *Editing Static Route Parameters* |
| Add, edit, and delete adjacent hosts. | *Editing Adjacent Host Parameters* |
| Reconfigure the Trivial File Transfer Protocol (TFTP). | *Editing TFTP Parameters* |
| Configure IP Filters. | The chapter *Configuring Filters* |
| Configure a single circuit with multiple IP addresses. | The chapter *Configuring Circuits* |
| Add RIP to a circuit. | The chapter *Configuring Circuits* |
| Add OSPF to a circuit. | The chapter *Configuring Circuits* |

This section describes how to access and edit IP parameters. For each parameter, it provides the following:

- Wellfleet default
- Valid options
- Parameter's function
- Instructions for setting the parameter

You begin from the Wellfleet Configuration Manager Window (see Figure 10-8); the first window displayed when you enter the Configuration Manager application.

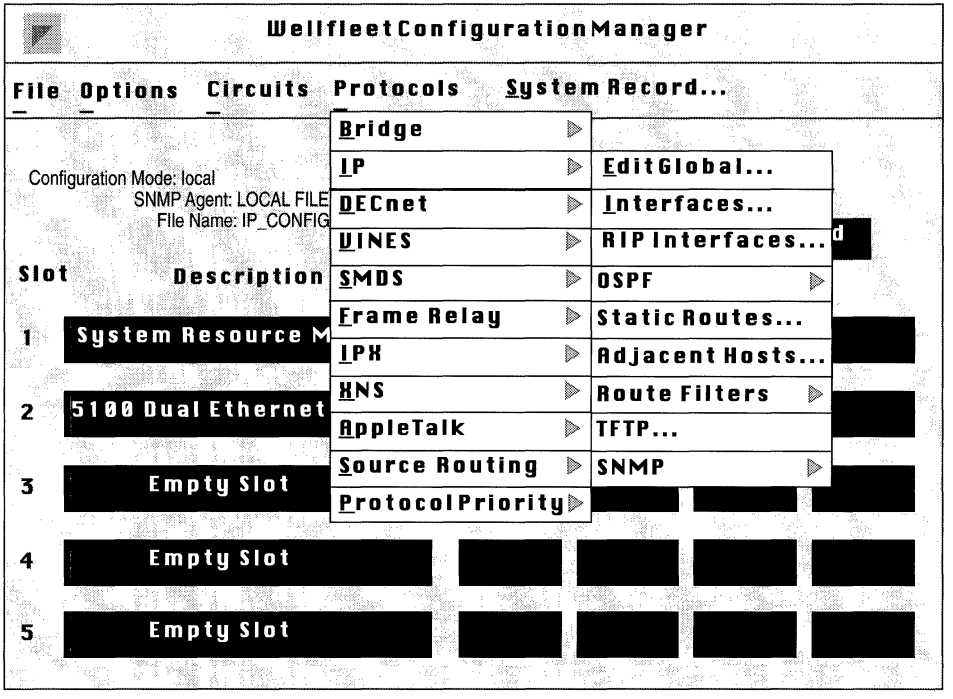


**Figure 10-8. Wellfleet Configuration Manager Window**

## Editing IP Global Parameters

You edit global parameters in the IP Global Parameters Window (see Figure 10-9). To display the window, select the Protocols/IP/Edit Global option in the Wellfleet Configuration Manager Window.

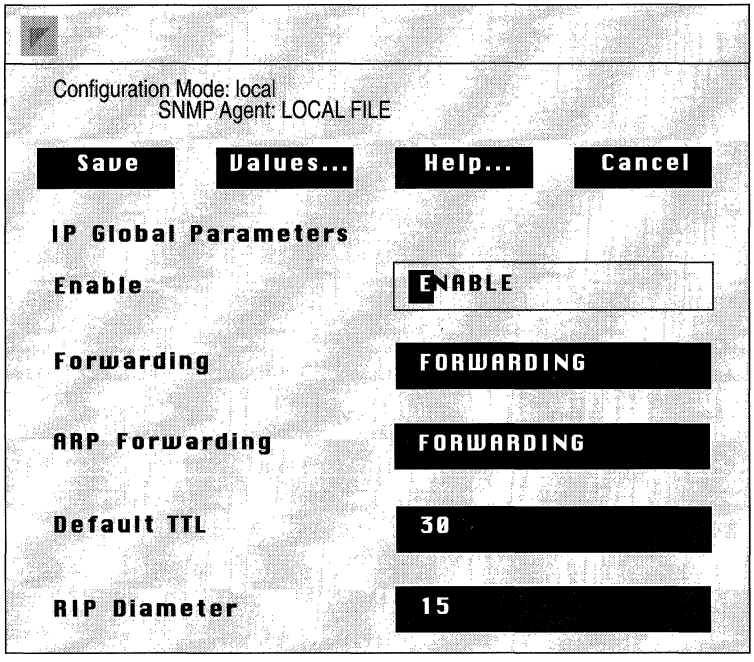


**Figure 10-9. IP Global Parameters Window**

This section provides information you need to edit each parameter in the IP Global Parameters Window. Refer to this information to edit the parameters you wish to change. When you are done, click on the Save button to exit the window and save your changes.

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | The Configuration Manager automatically sets the global Enable parameter to Enable when you add IP support to a circuit. |
| Options: | Enable/Disable |
| Function: | Specifies the state of the IP router software. |
| Instructions: | Select Enable if you have previously disabled the IP router software and now wish to re-enable it.<br>Select Disable to disable the IP router software. |

**Warning:** In dynamic mode, when you set the global Enable parameter to Disable, you immediately prohibit all Site Manager communication with the BN.

**Parameter :** **Forwarding**

Wellfleet Default: Forwarding

Options: Forwarding/Not Forwarding

Function: Specifies whether the IP router forwards IP traffic that is not explicitly addressed to it.

Instructions: Select Forwarding if you want the IP router to route (forward) IP traffic. Forwarding configures the IP router to process all broadcast packets and all IP packets explicitly addressed to it, and to *route* all other IP packets.

Select Not Forwarding if you want to provide IP management access (by means of TFTP and SNMP) to all active IP interfaces, but want to prohibit the IP router from forwarding IP traffic. You must specify an identical IP address and mask combination for each active IP interface that will provide management access. Not Forwarding configures the IP router to act as an IP host; it does not forward IP traffic, but still processes packets explicitly addressed to it. In Not Forwarding mode, only static routes and adjacent-host routes are allowed. No routing protocols are initiated.

Because the IP router does not forward IP traffic in Not Forwarding mode, you must configure the BN to *bridge* IP traffic not explicitly addressed to it. You must configure the Bridge for each circuit that conveys IP datagrams. The Bridge will then forward all IP datagrams that are not explicitly addressed to the BN.

**Warning:** When you reset this parameter in dynamic mode, the IP router restarts, causing the Site Manager to lose its BN connection temporarily, and to display a warning message. To verify that the change took effect, redisplay the IP Global Parameters Window and inspect the setting.

| | |
|---|---|
| **Parameter :** | **ARP Forwarding** |
| Wellfleet Default: | Forwarding |
| Options: | Forwarding/Not Forwarding |
| Function: | Specifies whether ARP forwards IP traffic that is not explicitly addressed to it. If this parameter is set to Forwarding, then ARP packets are either consumed, if destined for the BN, or dropped. If this parameter is set to Not Forwarding, ARP packets are consumed, if destined for the BN, or bridged onto remaining ARP interfaces. |
| Instructions: | Either accept the default, Forwarding, or select Not Forwarding. |

**Note:** When you reset this parameter in dynamic mode, the IP router restarts, causing the Site Manager to lose its BN connection temporarily, and to display a warning message. To verify that the change took effect, redisplay the IP Global Parameters Window and inspect the setting.

| | |
|---|---|
| **Parameter :** | **Default TTL** |
| Wellfleet Default: | 30 |
| Options: | 1 to 255 hops |
| Function: | Specifies the starting time of the Time to Live (TTL) counter for each packet which originates at the BN, called a source packet, that the BN transmits. When a source packet is transmitted, the TTL counter starts to decrement. Each router, or hop, that the packet traverses decrements the TTL counter by one. When the counter reaches zero, the router discards the packet unless it is destined for a locally attached network. The TTL counter prevents packets from looping endlessly through the network. |
| Instructions: | Enter the maximum number of hops a source packet can traverse. |

**Warning:** When you reset this parameter in dynamic mode, the IP router restarts, causing the Site Manager to lose temporarily its BN connection, and to display a warning message. To verify that the change took effect, redisplay the IP Global Parameters Window and inspect the setting.

**Parameter :** **RIP Diameter**

Wellfleet Default: 15

Options: 1 to 127

Function: Specifies the value, or hop count, the Routing Information Protocol (RIP) uses to denote infinity. In order for RIP to operate properly, every router within the network must be configured with an identical RIP Diameter value. If RIP is not enabled, the RIP Diameter parameter specifies the maximum number of hops within the autonomous system; if RIP is not enabled, the IP router still must understand network width.

Instructions: You must set the RIP Diameter parameter so that none of the interface cost, static cost, or route filter cost parameters exceed the RIP Diameter parameter. Wellfleet recommends that you accept the default RIP Diameter value of 15.

**Warning:** When you reset this parameter in dynamic mode, the IP router restarts, causing the Site Manager to lose temporarily its BN connection, and to display a warning message. To verify that the change took effect, redisplay the IP Global Parameters Window and inspect the setting.

## Editing IP Interfaces

You edit an IP interface in the IP Interface Parameters Window for that interface. To display this window, first select the Protocols/IP/ Interface option in the Wellfleet Configuration Manager Window to display the IP Interfaces Window (see Figure 10-10). Next, select the interface you wish to edit and click on the Edit button to display the IP Interface Parameters Window for that interface (see Figure 10-11).

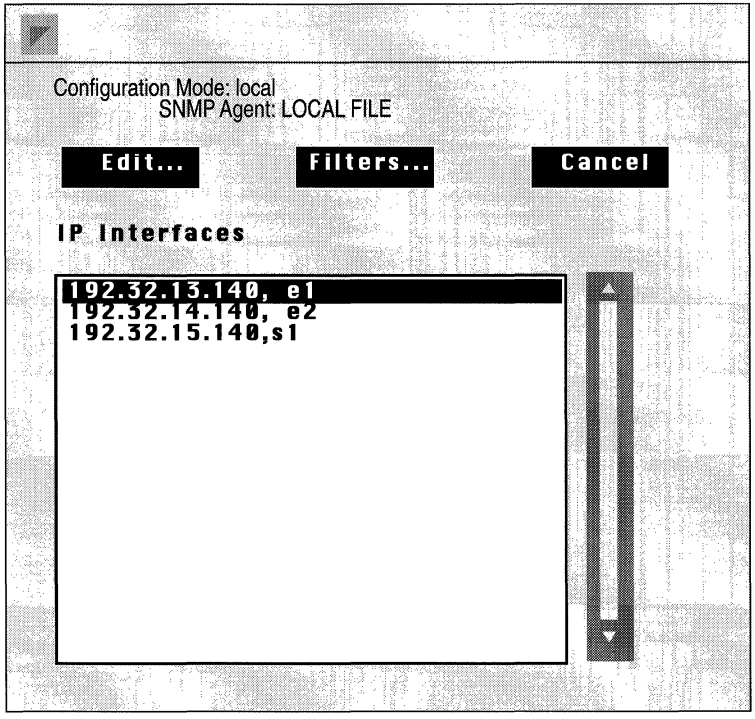


**Figure 10-10. IP Interfaces Window**

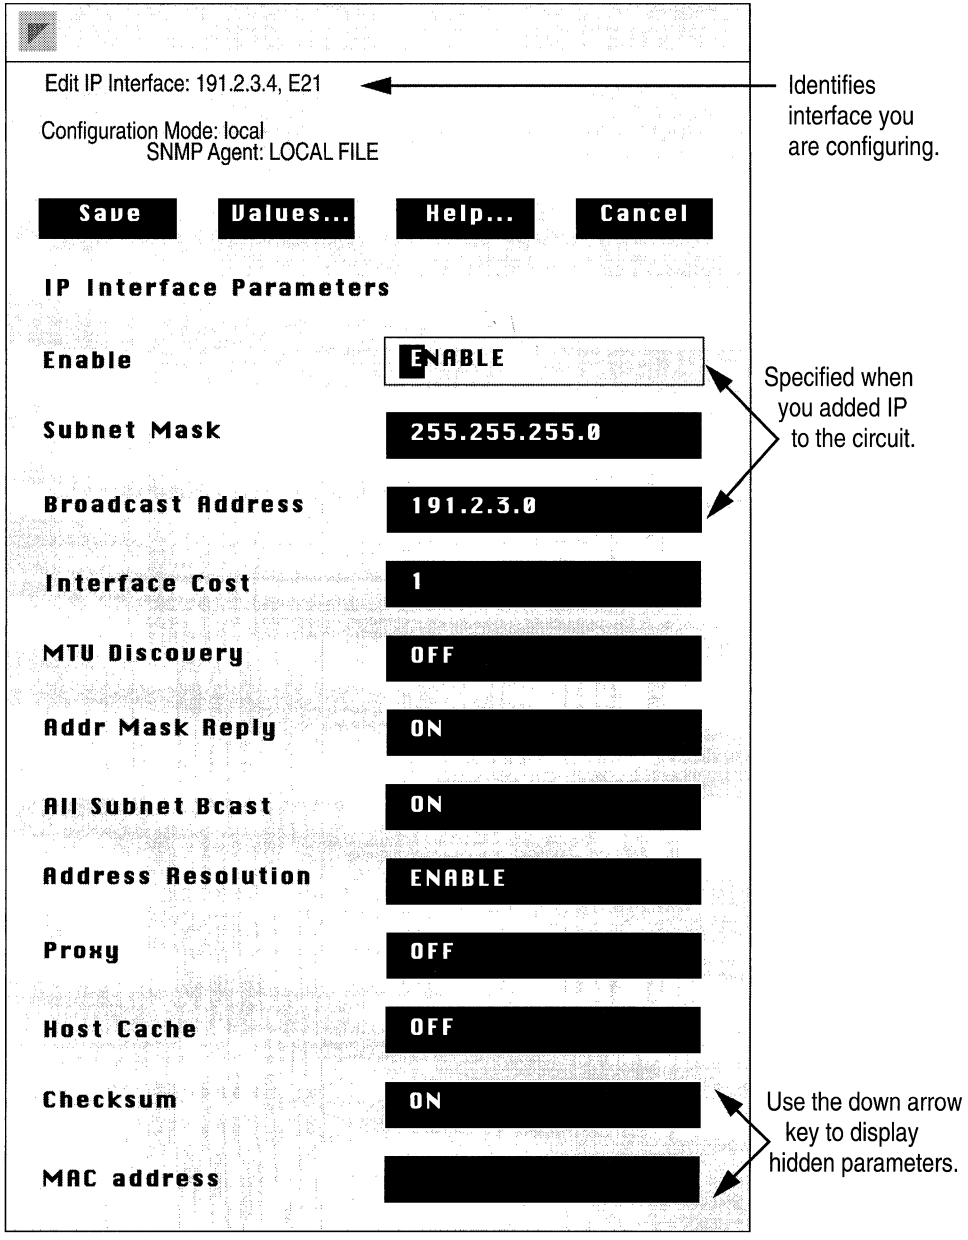


**Figure 10-11. IP Interface Parameters Window**

This section provides information you need to edit each parameter in the IP Interface Parameters Window. Refer to this information to edit the parameters you wish to change. When you are done, click on the Save button to exit the window and save your changes.

**Warning:** When you reconfigure an interface in dynamic mode, IP on that interface restarts. Thus, if the interface you reconfigure is the interface which supports the Site Manager's SNMP connection to the BN, restarting IP on that interface will cause the Site Manager to lose temporarily its BN connection and to display a warning message. To verify that the change took effect, redisplay the IP Global Parameters Window and inspect the setting.

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | The Configuration Manager automatically sets this interface-specific Enable parameter to Enable when you add IP support to this circuit. |
| Options: | Enable/Disable |
| Function: | Enables or disables IP routing on this interface. |
| Instructions: | Select Enable if you previously set this parameter to Disable and now wish the circuit to support IP routing.<br><br>Select Disable only if you wish to disable IP routing over this circuit. |

| | |
|---|---|
| **Parameter :** | **Subnet Mask** |
| Wellfleet Default: | You specified the Subnet Mask parameter when you added IP to the circuit. |
| Options: | Depends on the class of the network to which the interface connects. |
| Function: | Specifies the network and subnetwork portion of the 32-bit IP address. |
| Instructions: | Enter the subnet mask in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Broadcast Address** |
| Wellfleet Default: | You specified the Broadcast Address parameter when you added IP to the circuit. |
| Options: | 0 or any IP address |
| Function: | Specifies the broadcast address that the IP router uses to broadcast packets. |
| Instructions: | Select 0 to configure the IP router to use an all-1s address for broadcasting packets; or enter the broadcast address in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Interface Cost** |
| Wellfleet Default: | 1 |
| Options: | 1 to the value of RIP Diameter |
| Function: | Sets the cost of this interface. The interface cost is added to routes learned on this interface through RIP and is specified in subsequent RIP packets out other interfaces. |
| Instructions: | Enter the interface cost value (standard RIP implementation assigns a cost of 1), however, keep in mind that increasing this value causes the upper bound set by RIP Network Diameter to be attained more rapidly. |

| | |
|---|---|
| **Parameter :** | **MTU Discovery** |
| Wellfleet Default: | Off |
| Options: | On/Off |
| Function: | Specifies whether the Reply MTU option (option 11 in RFC 1063) is enabled on this interface. When the option is enabled, this interface responds to Probe MTUs (option 12 in RFC 1063). A Probe MTU requests the minimum MTU (Maximum Transmission Unit) of all networks an IP datagram must traverse from source to destination. By enabling this interface to respond to Probe MTUs, you eliminate transit fragmentation and destination reassembly for datagrams destined for this interface, and therefore, decrease network load. |
| Instructions: | Select On to enable the Reply MTU option on this option; select Off to disable the option on this interface. |

| | |
|---|---|
| **Parameter :** | **Addr Mask Reply** |
| Wellfleet Default: | On |
| Options: | On/Off |
| Function: | Specifies whether this interface generates ICMP (Internet Control Message Protocol) address-mask-reply messages in response to valid address-mask-request messages. The interface generates ICMP address-mask-reply messages in compliance with the relevant sections of RFCs 950 and 1009. |
| Instructions: | Select On to enable ICMP address-mask-reply message generation on this interface.<br><br>Select Off to disable ICMP address-mask-reply message generation on this interface. |

| | |
|---|---|
| **Parameter :** | **All Subnet Bcast** |
| Wellfleet Default: | On |
| Options: | On/Off |
| Function: | Specifies whether or not the IP router floods received ASB datagrams across other router interfaces.<br><br>An ASB datagram has a destination address equal to the broadcast address for an entire subnet. For example, if a network interface serves the subnet 128.10.2.1, with a subnet mask of 255.255.255.0, the IP router considers any datagram with a destination address of 128.10.255.255 or 128.10.0.0 to be an ASB datagram. |
| Instructions: | Select On to configure the IP router to flood ASBs received on this interface to all interfaces configured with ASB set to On, which service the same subnet. Similarly, these other interfaces flood received ASBs to this interface.<br><br>Select Off to configure the IP router to prohibit ASB flooding on this interface. |

| | |
|---|---|
| **Parameter :** | **Address Resolution** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Specifies whether this interface uses ARP to map 32-bit IP addresses to 48-bit Ethernet addresses (refer to the *Address Resolution Protocol* in the *IP Overview* for more information). |
| Instructions: | Select Enable to enable ARP on this interface.<br><br>Select Disable to disable ARP on this interface. |

**Parameter :** **Proxy**

Wellfleet Default: Off

Options: On/Off

Function: Specifies whether this interface uses Proxy ARP to respond to ARPs for a remote network (refer to the *Proxy ARP* in the *IP Overview* for more information).

Instructions: Select On to enable Proxy ARP on this interface. In order to enable Proxy ARP, you must have set ARP to Enable for this interface. When you enable Proxy ARP, the IP router assumes responsibility for IP datagrams destined for the remote network.

Select Off to disable Proxy ARP on this interface.

**Parameter :** **Host Cache**

Wellfleet Default: 1 (Off)

Options: 1 (Off), 120, 180, 240, 300, 600, 900 or 1200 seconds

Function: Specifies whether the IP router ages entries in the address-resolution cache for this interface, and specifies the aging interval in seconds if the interface does age entries. The address-resolution cache contains host physical addresses learned by means of ARP or Proxy ARP. A host entry is aged (deleted), if the IP router sends no traffic destined for that host within the specified aging period.

Instructions: Select 1 to disable aging on this interface; the IP router does not age out address-resolution cache entries. Select one of the other values to enable aging with an aging interval equal to the value you select (120, 180, 240, 300, 600, 900, or 1200 seconds); the IP router removes address-resolution cache entries that have not been accessed within the specified number of seconds (for example, 120 seconds, 180 seconds, etc.). Once an entry is removed, the IP router must use ARP to re-acquire the physical-level address.

| | |
|---|---|
| **Parameter :** | **Checksum** |
| Wellfleet Default: | On |
| Options: | On/Off |
| Function: | Specifies whether UDP checksum processing is enabled on this interface. |
| Instructions: | Select On to enable UDP checksum processing for the interface; all outgoing and incoming UDP datagrams are subject to checksumming. You should select On in virtually all instances.<br><br>Select Off to disable UDP checksum processing and provides backward compatibility with UNIX BSD 4.1 |

| | |
|---|---|
| **Parameter :** | **MAC Address** |
| Wellfleet Default: | None |
| Options: | 0 or a User-Specified MAC Address |
| Function: | Specifies a MAC (media access control) address for this IP interface. The IP router will use its IP address and this MAC address when transmitting and receiving packets on this interface. |
| Instructions: | Select 0 to configure the IP router to use its IP address and the circuit's MAC address when transmitting packets on this interface.<br><br>Enter a your own MAC address to configure the IP router to use its IP address and the specified MAC address when transmitting packets on this interface. |

| | |
|---|---|
| **Parameter :** | **TR End Station** |
| Wellfleet Default: | Disable |
| Options: | Enable/Disable |
| Function: | Specifies if this interface is enabled for source routing end station support. |
| Instructions: | Enable if this interface is 1) of type token ring and 2) if source routing is enabled on the IP routers in this ring. |

| | |
|---|---|
| **Parameter :** | **SMDS Group Address** |
| Wellfleet Default: | None |
| Options: | Any valid 10-digit North American Numbering Plan (NANP) telephone number. |
| Function: | Provides a MAC-layer multicast address for this IP interface in an SMDS network. |
| Instructions: | Enter the 10-digit multicast address as provided by the SMDS subscription agreement. |

**Note:** SMDS standards specify the use of 8-octet E.164 addresses. The BN converts the SMDS Group Address and SMDS ARP Address to E.164 format by prepending hexadecimal E1 and appending hexadecimal FF FF to the NANP number. For example, the group address 6175551212 becomes E1 61 75 55 12 12 FF FF.

| | |
|---|---|
| **Parameter :** | **SMDS Arp Req Address** |
| Wellfleet Default: | None |
| Options: | Any valid 10-digit North American Numbering Plan (NANP) telephone number. |
| Function: | Provides an address resolution multicast address for this IP interface in an SMDS network. |
| Instructions: | Enter the 10-digit multicast address as provided by the SMDS subscription agreement. |

| | |
|---|---|
| **Parameter :** | **FR Broadcast DLCI** |
| Wellfleet Default: | 0 |
| Options: | Any decimal number |
| Function: | Provides a broadcast address for this IP interface in a Frame Relay network. If a value is entered for FR Broadcast DLCI, it means that the Frame Relay switch, rather than the BN, will broadcast the message. |
| Instructions: | Enter the broadcast address as provided by the Frame Relay subscription agreement. |

**Parameter :** **FR Multicast DLCI #1**

Wellfleet Default: 0

Options: Any decimal number

Provides a multicast address for this IP interface that will send messages to all OSPF routers in a Frame Relay network. If a value is entered for FR Multicast DLCI #1, it means that the Frame Relay switch, rather than the BN, will send the message to all OSPF routers.

This parameter has meaning only if OSPF has been added to this interface.

Instructions: Enter the multicast address for all OSPF routers as provided by the Frame Relay subscription agreement.

**Parameter :** **FR Multicast DLCI #2**

Wellfleet Default: 0

Options: Any decimal number

Provides a multicast address for this IP interface that will send messages to all OSPF designated routers in a Frame Relay network. If a value is entered for FR Multicast DLCI #2, it means that the Frame Relay switch, rather than the BN, will send the message to all OSPF designated routers.

This parameter has meaning only if OSPF has been added to this interface.

Instructions: Enter the multicast address for all OSPF designated routers as provided by the Frame Relay subscription agreement.

| | |
|---|---|
| **Parameter :** | **Redirects** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Indicates whether or not this interface sends out ICMP redirects.<br><br>ICMP redirects are messages sent by the BN to alert a source end station that the destination to which it has just sent a message resides on the same network as the source end station. In an Ethernet network, it would not be necessary for that source end station to send messages through the BN when sending to that same destination. However, on a Frame Relay network two stations on the same network may not be directly connected if the network is not fully meshed. In this case, set Redirects to Disabled. For more information about Frame Relay, see *Configuring Frame Relay*. |
| Instructions: | Either accept the default value, Enabled, or select Disabled if you do not want this interface to send out redirects. |

| | |
|---|---|
| **Parameter :** | **Encapsulation** |
| Wellfleet Default: | Ethernet |
| Options: | Ethernet, SNAP, both |
| Function: | Defines the datalink encapsulation to use for ARP packets generated at this interface if the underlying media is Ethernet. This parameter has no meaning if the underlying media is anything other than Ethernet. |
| Instructions: | Either accept the default value, Ethernet, or select an alternate option. |

## Editing Routing Information Protocol (RIP) Interface Parameters

Once you have enabled RIP on an interface, you can edit that interface in the RIP Interface Parameters Window for that interface. For instructions on how to enable RIP on an interface, see the *Configuring Circuits* chapter. Also, if you wish to change the RIP Diameter Value for the IP Router, refer to *Editing IP Global Parameters* in this chapter.

You display the RIP Interface Parameters Window for a RIP interface, by first selecting the Protocols/IP/RIP option in the Wellfleet Configuration Manager Window to display the RIP Interfaces Window (see Figure 10-12).

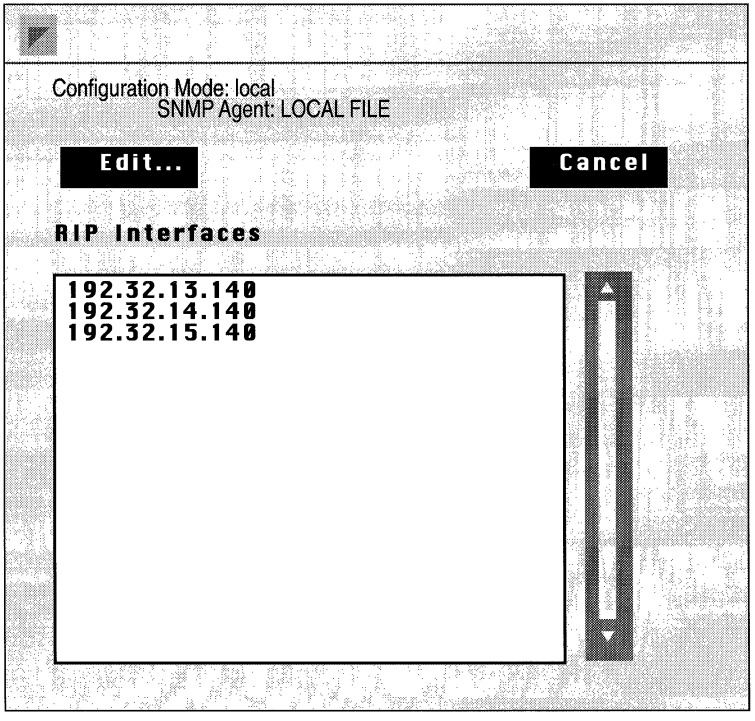


**Figure 10-12. RIP Interfaces Window**

In the RIP Interfaces Window, select the RIP interface you wish to edit and then click on the Edit button to display the RIP Interface Parameters Window for that interface (see Figure 10-13).

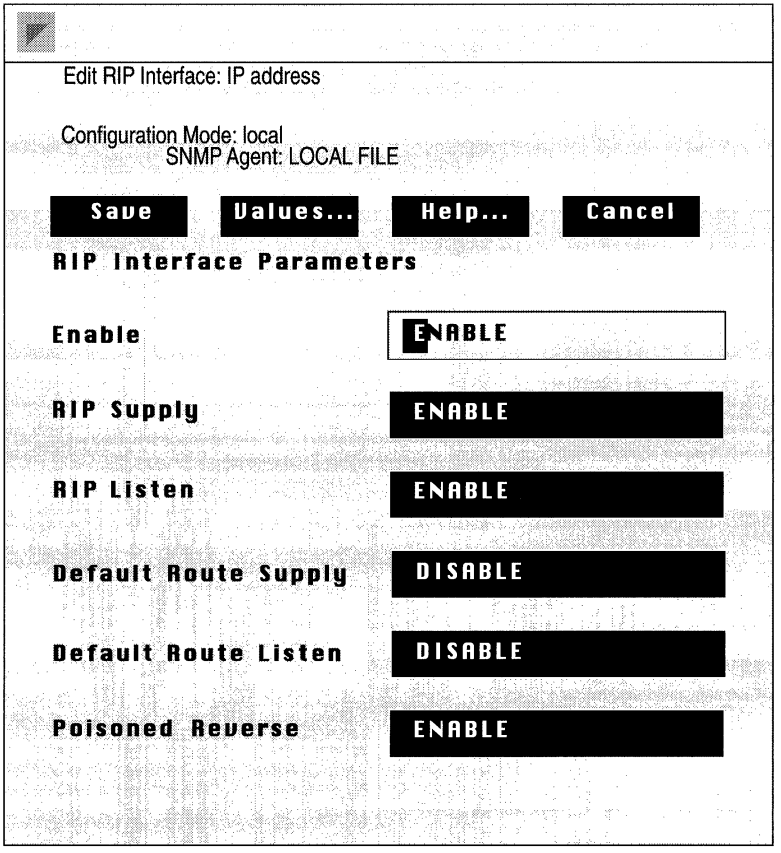


**Figure 10-13. RIP Interface Parameters Window**

This section provides information you need to edit each parameter in the RIP Interface Parameters Window. Refer to this information to edit the parameters you wish to change. When you are done, click on the Save button to exit the window and save your changes.

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | If you enabled RIP when you added the circuit or if you edited this circuit to support RIP, the Configuration Manager automatically sets this interface-specific RIP Enable parameter to Enable; otherwise, it is set to Disable. |
| Options: | Enable/Disable |
| Function: | Specifies whether the Routing Information Protocol (RIP) is enabled on this interface. |
| Instructions: | Select Enable to enable RIP on this interface.<br><br>Select Disable to disable RIP on this interface. |

| | |
|---|---|
| **Parameter :** | **RIP Supply** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Specifies whether the interface transmits periodic RIP updates to neighboring networks. |
| Instructions: | Select Enable to configure the interface to transmit RIP updates. You must select Enable if you want the interface to supply default route information.<br><br>Select Disable to prohibit the interface from transmitting RIP updates. |

| | |
|---|---|
| **Parameter :** | **RIP Listen** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Specifies whether this interface listens to RIP updates from neighboring networks. |
| Instructions: | Select Enable to configure this interface to listen to RIP updates, and thus, add received routing information to its internal routing table.<br><br>If you set RIP Listen to Enable, a route filter can still prohibit the interface from updating its internal routing tables.<br><br>Select Disable to configure the interface to ignore RIP updates from neighboring routers. Thus, the interface does not add received routing information to its internal routing table. |

| | |
|---|---|
| **Parameter :** | **Default Route Supply** |
| Wellfleet Default: | Disable |
| Options: | Enable/Disable |
| Function: | Specifies whether or not the interface advertises a default route in RIP updates sent to neighboring networks. In order to advertise a default route, you must have either statically configured a default route, or the router must have learned the default route (0.0.0.0). When a router does not know the direction of a particular address, it uses the default route as the destination. |
| Instructions: | Select Enable to configure the interface to advertise the default route. If you set Default Route Supply to Enable, you must also set RIP Supply to Enable.<br><br>Select Disable to configure the interface not to advertise the default route. |

| | |
|---|---|
| **Parameter :** | **Default Route Listen** |
| Wellfleet Default: | Disable |
| Options: | Enable/Disable |
| Function: | Specifies whether or not IP adds default route information to its internal routing table. |
| Instructions: | Select Enable to configure the RIP interface to listen for and potentially add the default route (0.0.0.0) information to its internal routing table. If you select Enable, you must also enable RIP Listen on this interface.<br><br>Select Disable to prohibit the RIP interface from adding the default route (0.0.0.0) information to its internal routing table. |

| | |
|---|---|
| **Parameter :** | **Poisoned Reverse** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Specifies how the RIP interface advertises routes it learns from an adjacent network in periodic updates subsequently sent to that network. |
| Instructions: | Select Enable to configure this RIP interface to implement poisoned reverse, which means the RIP interface advertises routes learned from an adjacent network in RIP updates subsequently sent to that network with a hop count of RIP Network Diameter plus one; thus declaring the destination unreachable. Poisoned Reverse can speed up the convergence of the network routing tables.<br><br>Select Disable to configure this RIP interface to advertise routes with the learned cost.<br><br>Select Split Horizon to configure this RIP interface to implement a split-horizon, which means the RIP interface omits routes learned from a neighbor in RIP updates subsequently sent to that neighbor. |

## Editing Static Route Parameters

The IP Static Routes Window allows you to add, edit, and delete static routes. The following sections describe each procedure. To display the IP Static Routes Window (see Figure 10-14), select the Protocols/IP/ Static Routes option in the Wellfleet Configuration Manager Window.

### Adding a Static Route

To add a static route, click on the Add button in the IP Static Routes Window to display the Add IP Static Route Window (see Figure 10-15).

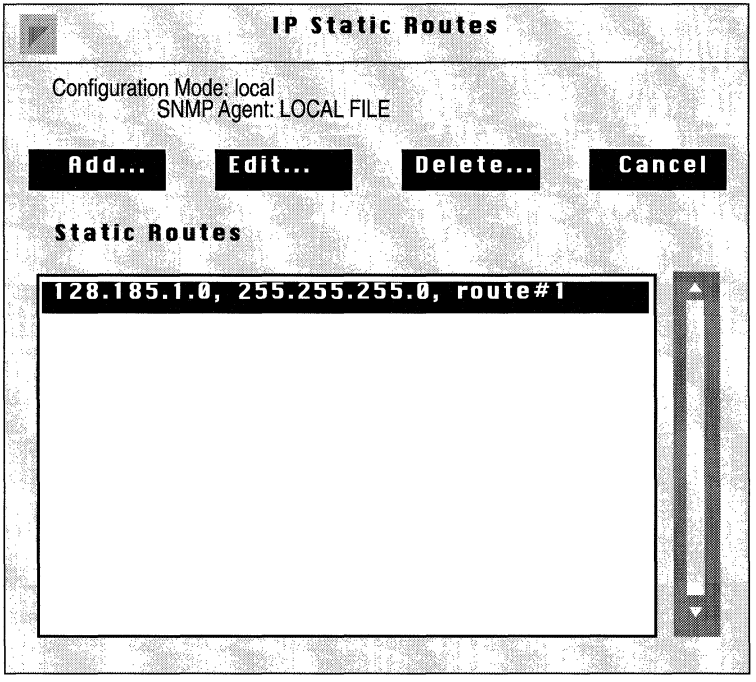


**Figure 10-14. IP Static Routes Window**

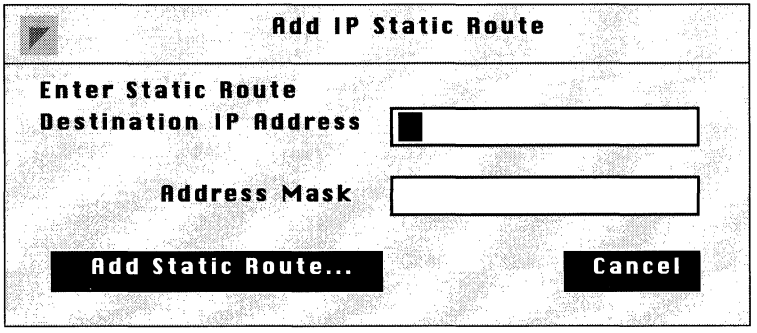


**Figure 10-15. Add IP Static Route Window**

Refer to the following parameter descriptions to enter the required information and then click on the Add Static Route button to display the Static Route Window. *Editing Parameters in the Static Route Window* describes how to configure a static route.

| **Parameter :** | **Destination IP Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any valid IP network address |
| Function: | Specifies the IP address of the network to which you wish to configure the static route. |
| Instructions: | Enter the destination IP address in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Address Mask** |
| Wellfleet Default: | None |
| Options: | Based on the network class of the IP address you specified at the Destination IP Address parameter. |
| Function: | Specifies the subnet mask of the destination network. |
| Instructions: | Enter the subnet mask in dotted decimal notation. |

### Editing a Static Route

The Configuration Manager does not allow you to reconfigure the Destination IP Address or Address Mask parameters for a static route. To change these parameters, you must delete the static route and add a new route with the proper information. Otherwise, you can reconfigure all other parameters associated with a static route.

To edit a static route, first select the static route you wish to edit in the IP Static Routes Window (see Figure 10-14); then click on the Edit button to display that static route's configuration in the Static Route Window (see Figure 10-16). *Editing Parameters in the Static Route Window* describes how to reconfigure the static route.

### Editing Parameters in the Static Route Window

This section provides information you need to edit each parameter in the Static Route Window (see Figure 10-16). Refer to this information to edit the parameters you wish to change. When you are done, click on the Save button to exit the window and save your changes.

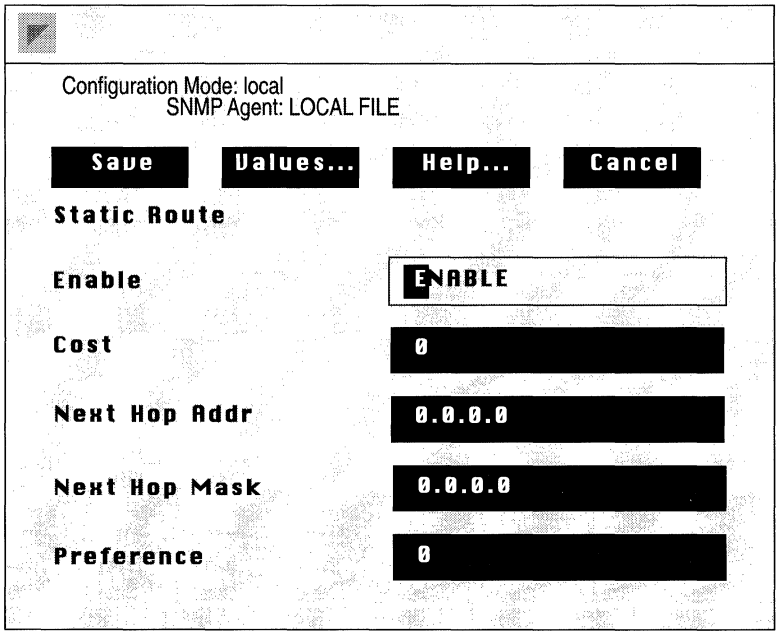


**Figure 10-16. Static Route Window**

**Parameter : Enable**

Wellfleet Default: The Configuration Manager automatically sets this parameter to Enable when you click on the Add Static Route button in the Add IP Static Route Window.

Options: Enable/Disable

Function: Specifies the state (active or inactive) of the static route record in the IP routing tables.

Instructions: Select Disable to make the static route record inactive in the IP routing table; the IP router will not consider this static route.

Select Enable to make the static route record active again the IP routing table.

| | |
|---|---|
| **Parameter :** | **Cost** |
| Wellfleet Default: | 1 |
| Options: | 1 to the value of the RIP Diameter parameter |
| Function: | Specifies the number of router hops a datagram can traverse before reaching the destination IP address. The IP router uses Cost when determining the best route for a datagram to follow. The Cost is also propagated through RIP, etc. |
| Instructions: | Enter the number of router hops. |

| | |
|---|---|
| **Parameter :** | **Next Hop Addr** |
| Wellfleet Default: | None |
| Options: | Any valid IP address |
| Function: | Specifies the IP address of the next hop router. |
| Instructions: | Enter the IP address in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Next Hop Mask** |
| Wellfleet Default: | None |
| Options: | Not Applicable |
| Function: | Specifies the subnet mask of the next hop router. |
| Instructions: | Enter the Subnet mask in dotted decimal notation. |

| Parameter : | Preference |
|---|---|
| Wellfleet Default: | 1 |
| Options: | 1 to 16 |
| Function: | Specifies a weighted value (from 1 to 16, with 16 being the most preferred) that the IP router uses to select a route when its routing tables contain multiple routes to the same destination. |
| Instructions: | Enter a value from 1 to 16 for this static route. |

### Deleting a Static Route

To delete a static route, first select the static route you wish to delete in the IP Static Routes Window (see Figure 10-14), and then click on the Delete button to display the Delete IP Static Route Window (see Figure 10-17). Click on the Delete button to delete the static route.

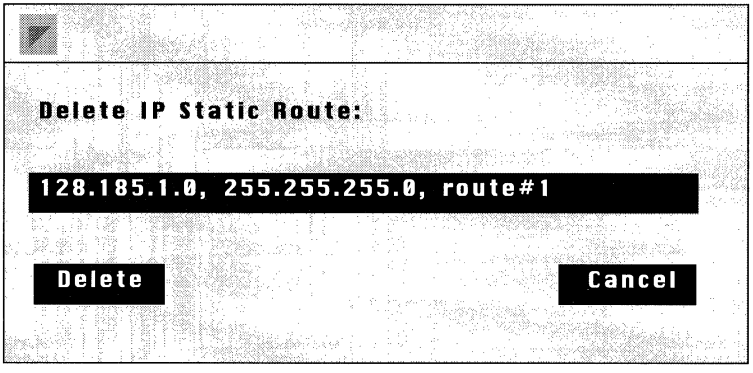


**Figure 10-17. Delete IP Static Route Window**

## Editing Adjacent Host Parameters

The IP Adjacent Hosts Window (Figure 10-18) allows you to add, edit, and delete adjacent host routes. The following sections describe each procedure. To display the IP Adjacent Hosts Window, select the Protocols/IP/Adjacent Hosts option in the Wellfleet Configuration Manager Window.

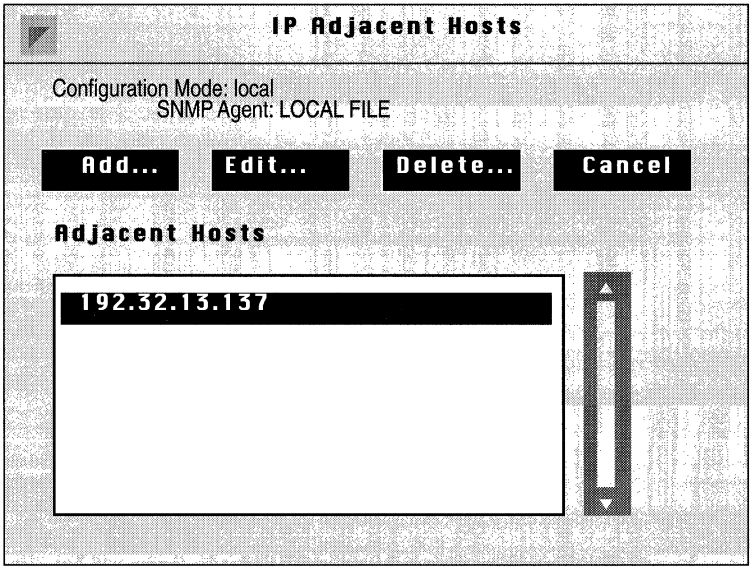


**Figure 10-18. IP Adjacent Hosts Window**

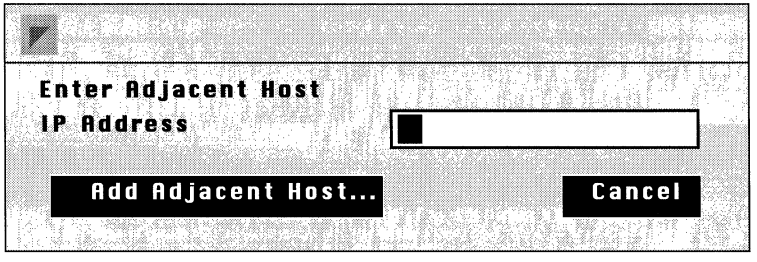


**Figure 10-19. Enter Adjacent Host Window**

### Adding an Adjacent Host

To add an adjacent host, click on the Add button in the IP Adjacent Hosts Window to display the Enter Adjacent Host Window (see Figure 10-19). Enter the required information, as follows:

| **Parameter :** | **IP Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | Valid IP address. |
| Function: | Specifies the IP address of the device for which you wish to configure an adjacent host. |
| Instructions: | Enter the IP address in dotted decimal notation. |

Click on the Add Adjacent Host button to display the Adjacent Host Window (Figure 10-20). *Editing Parameters in the Adjacent Host Window* describes how to configure an adjacent host route.

### Editing an Adjacent Host

The Configuration Manager does not allow you to change the adjacent host's IP address. If you wish to change this parameter, you must delete the adjacent host and configure a new adjacent host with the proper IP address. Otherwise, you can reconfigure all other parameters associated with an adjacent host.

To edit an adjacent host, first select the adjacent host you wish to edit in the IP Adjacent Hosts Window (see Figure 10-18), and then click on the Edit button to display that adjacent host's configuration in the Adjacent Host Window (see Figure 10-20). Refer to *Editing Parameters in the Adjacent Host Window* for instructions on how to reconfigure an adjacent host route.

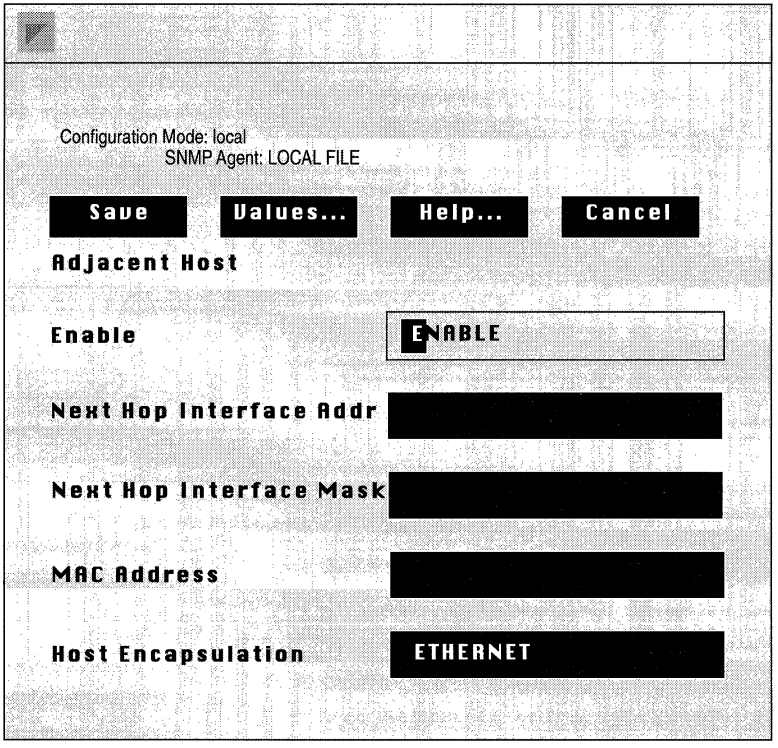


**Figure 10-20. Adjacent Host Window**

## Editing Parameters in the Adjacent Host Window

This section provides information you need to edit each parameter in the Adjacent Host Window (see Figure 10-20). Refer to this information to edit the parameters you wish to change. When you are done, click on the Save button to exit the window and save your changes.

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | The Configuration Manager automatically sets this parameter to Enable when you click on the Add Adjacent Host button in the Enter Adjacent Host Window. |
| Options: | Enable/Disable |
| Function: | Specifies the state (active or inactive) of the adjacent host in the IP routing tables. |
| Instructions: | Select Disable to make the adjacent host record inactive in the IP routing table; the IP router will not consider this adjacent host.<br><br>Select Enable to make the adjacent host record active again in the IP routing table. |

| | |
|---|---|
| **Parameter :** | **Next Hop Interface Addr** |
| Wellfleet Default: | None |
| Options: | Valid IP address |
| Function: | Specifies the IP address of the router's network interface to the adjacent host. |
| Instructions: | Enter the IP address in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Next Hop Interface Mask** |
| Wellfleet Default: | None |
| Options: | Based on the network class of the IP address specified at the Next Hop Interface Addrs parameter. |
| Function: | Specifies the subnet mask of the IP address specified for the Next Hop Addrs parameter. |
| Instructions: | Enter the subnet mask in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **MAC Address** |
| Wellfleet Default: | None |
| Options: | Not Applicable |
| Function: | Specifies the 48-bit Ethernet (or 64-bit for SMDS), address of the adjacent host. |
| Instructions: | Enter the MAC address as a 12-digit hexadecimal number. |

| | |
|---|---|
| **Parameter :** | **Host Encapsulation** |
| Wellfleet Default: | Ethernet |
| Options: | Ethernet/SNAP |
| Function: | Specifies the encapsulation method of the adjacent host. |
| Instructions: | Select Ethernet if you are defining a point-to-point network interface, or if the adjacent host resides on an Ethernet.<br><br>Select SNAP (Service Network Access Point) for all other instances. |

### Deleting an Adjacent Host

To delete an adjacent host, select the adjacent host you wish to delete in the IP Adjacent Hosts Window (see Figure 10-18), and click on the Delete button to display the Delete IP Adjacent Host Window (see Figure 10-21). Click on the Delete button to delete the adjacent host.

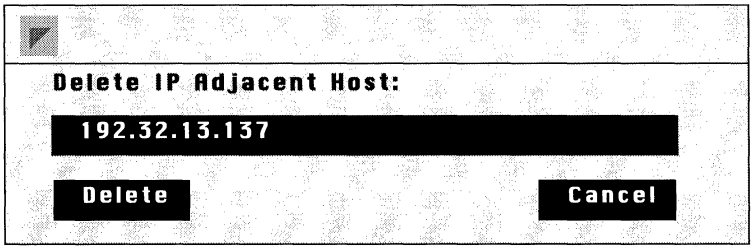


**Figure 10-21. Delete IP Adjacent Host Window**

## Editing RIP Import Route Filters

The RIP Import Route Filters List Window (see Figure 10-22), allows you to add, edit, and delete RIP import route filters. The following sections describe each procedure. To begin, display the RIP Import Route Filters List Window, by selecting the Protocols/IP/Route Filters/ RIP/Import Filters option in the Wellfleet Configuration Manager Window.

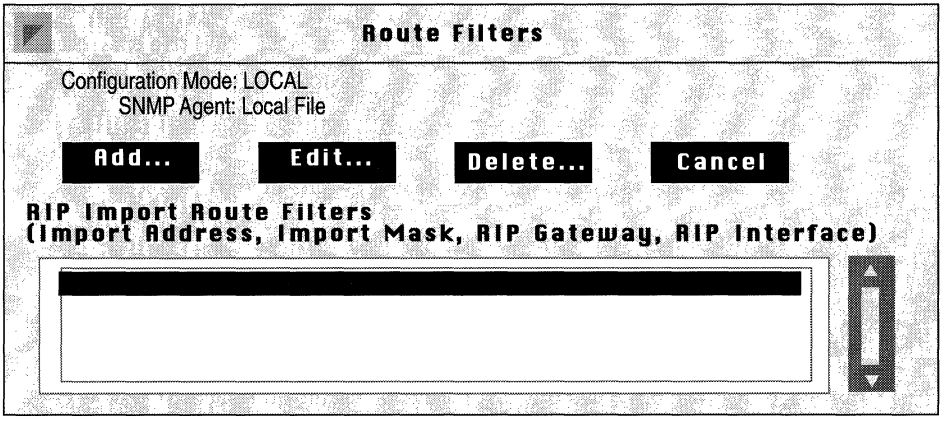


**Figure 10-22. RIP Import Route Filters List Window**

### Adding a RIP Import Route Filter

To add an import route filter, click on the Add button in the RIP Import Route Filters Window (see Figure 10-22), to display the RIP Import Route Filter Configuration Window (see Figure 10-23).

This section provides information you need to set each parameter in the RIP Import Route Filter Configuration Window. Refer to this information as necessary. When you are done, click on the Save button to exit the window and to add the import route filter.

**Note:** When you add an import route filter, the Configuration Manager automatically sets three additional parameters (specifically, it enables the filter, sets the Action parameter to Accept, and sets the Preference parameter to 0); however, these parameters are not displayed in the RIP Import Route Filter Configuration Window. If you wish to edit these parameters, refer to *Editing an Import Route Filter* for instructions.

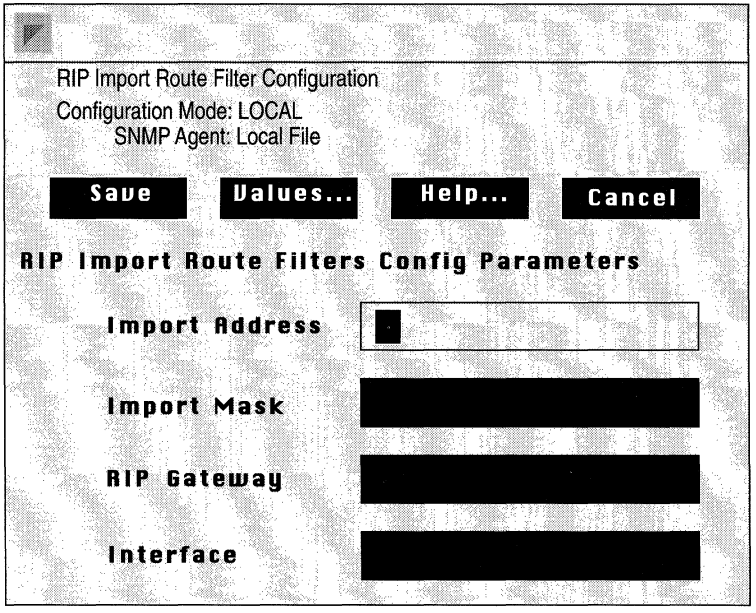


**Figure 10-23. RIP Import Route Filter Configuration Window**

**Parameter :** **Import Address**

Wellfleet Default: None

Options: Any IP network address

Function: Identifies, by IP address, the network to which this filter applies. If this field is left blank, the filter applies to all networks.

Instructions: Enter the appropriate network address in dotted decimal notation.

**Parameter :** **Import Mask**

Wellfleet Default: None

Options: Depends on the address class of the network address.

Function: Specifies the range of addresses upon which this filter acts.

For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Import Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Import Mask, the Net ID and Subnet ID portions of the address will be filtered.

If the Import Address field was left blank, the Import Mask should be left blank also.

Instructions: Enter the appropriate mask in dotted decimal notation.

| | |
|---|---|
| **Parameter :** | **RIP Gateway** |
| Wellfleet Default: | None |
| Options: | Any IP address |
| Function: | Identifies, by IP address, the router that is sending the updates. This filter will apply to updates from that router.<br><br>If left blank, this filter applies to updates from any router. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Interface** |
| Wellfleet Default: | None |
| Options: | Any IP address |
| Function: | Specifies the local IP address of the interface that connects this router to the RIP Gateway. This filter will apply only to those updates received on this interface.<br><br>If left blank, this filter applies to all interfaces. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. |

### Editing a RIP Import Route Filter

The Configuration Manager allows you to edit three parameters (Enable, Action, and Preference) associated with a RIP import route filter. You edit these parameters in the RIP Import Route Filters Window (see Figure 10-24). To display this window for a particular import route filter, do the following:

- First, select the import route filter you wish to edit in the RIP Import Route Filters scroll box in the RIP Import Route Filters List Window (see Figure 10-22).
- Second, click on the Edit button to display the RIP Import Route Filters Window for that route filter.

This section provides information you need to set each parameter in the RIP Import Route Filters Window. Refer to this information as necessary. When you are done, click on the Save button to exit the window and to save your changes.

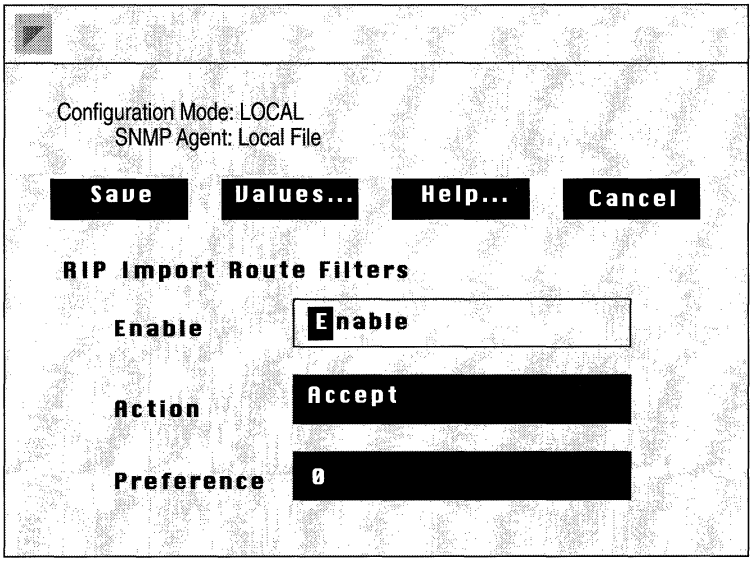


**Figure 10-24. RIP Import Route Filters Window**

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Range: | Enable/Disable |
| Function: | Enables or disable this import route filter. |
| Instructions: | Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now wish to re-enable it. |

| | |
|---|---|
| **Parameter :** | **Action** |
| Wellfleet Default: | Accept |
| Range: | Accept/Ignore |
| Function: | Specifies whether the route is transferred to the routing tables. If Action is set to Accept (default), the routing information is sent to the routing tables. If Action is set to Ignore, the routing information is dropped. |
| Instructions: | Either accept the default Accept, or select Ignore. |

**Parameter :** **Preference**

Wellfleet Default: 0

Range: 0 through 16

Function: Assigns a weighted precedence value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, OSPF external, and RIP.

If this hierarchy is acceptable, accept the default value 0 for preference. If you want to grant preference to this RIP-derived route, assign a new preference value in the range of 1 to 16 (the greater the number, the higher the preference).

**Note:** The default preference for static routes is 1, but may be set to any value between 0 and 16 (refer to *Editing Static Route Parameters* for more information). If you want to grant a RIP-derived route preference over a static route, make sure the preference value you assign to the RIP-derived route is greater than the preference value of the static route you want it to override.

Instructions: Either accept the default value 0, or enter a new value.

### Deleting an Import Route Filter

You delete an import route filter, as follows:

- First, select the import route filter you wish to delete in the RIP Import Route Filters scroll box in the Route Filters Window for import route filters (see Figure 10-22).
- Second, click on the Delete button to display the Delete Import Route Filters Window for that import route filter.
- Finally, click on the Delete button to delete the import route filter and to return to the Route Filters Window for import route filters, which no longer displays that import route filter.

## Editing RIP Export Route Filters

The RIP Export Route Filters List Window (see Figure 10-25), allows you to add, edit, and delete export route filters. The following sections describe each procedure. To begin, display the RIP Export Route Filters List Window for export route filters, by selecting the Protocols/ IP/Route Filters/RIP/Export Filters option in the Wellfleet Configuration Manager Window.

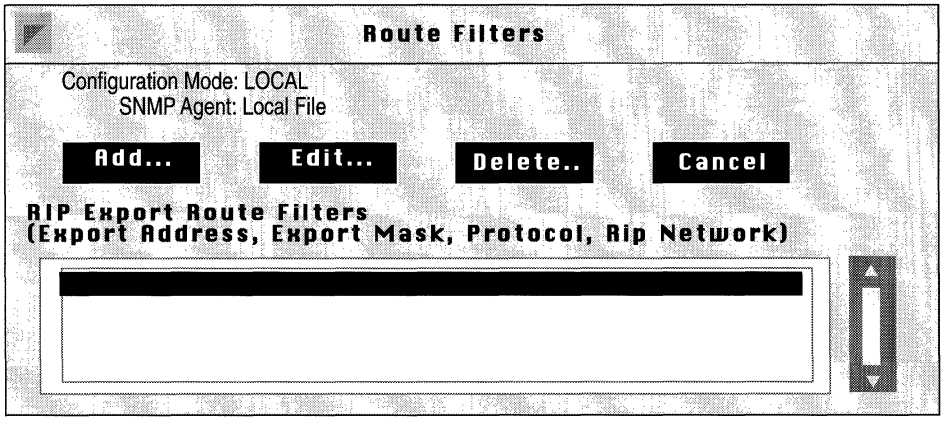


**Figure 10-25. RIP Export Route Filters List Window**

### Adding a RIP Export Route Filter

To add a RIP export route filter, click on the Add button in the RIP Export Route Filters List Window (see Figure 10-25), to display the RIP Export Route Filter Configuration Window (see Figure 10-26).

This section provides information you need to set each parameter in the RIP Export Route Filter Configuration Window. Refer to this information as necessary. When you are done, click on the Save button to exit the window and to add the export route filter.

**Note:** When you add an export route filter, the Configuration Manager automatically sets three parameters (specifically, it enables the filter, sets the Action parameter to Propagate, and sets the Metric parameter to equal the actual route cost as learned); however, these parameter are not displayed in the RIP Export Route Filter Configuration Window. If you wish to edit these parameters, see *Editing an Export Route Filter* for instructions.

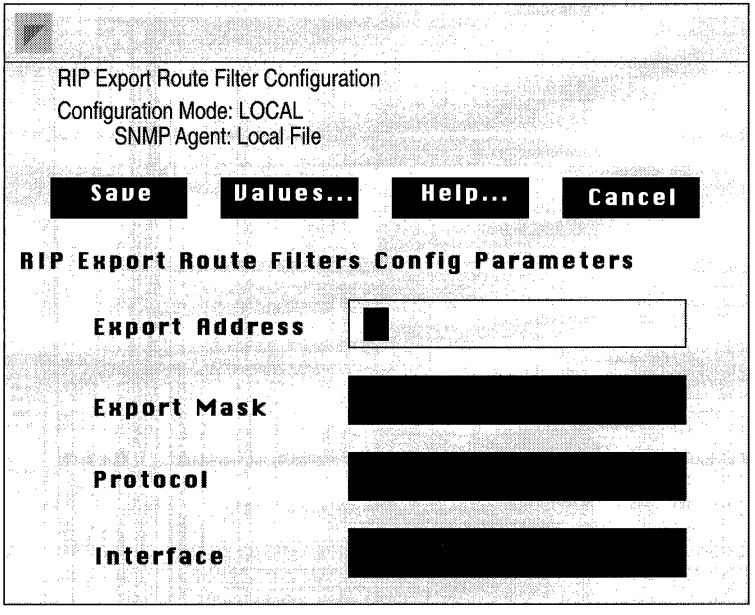


**Figure 10-26. RIP Export Route Filter Configuration Window**

| | |
|---|---|
| **Parameter :** | **Export Address** |
| Wellfleet Default: | None |
| Options: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. If left blank, the filter applies to all networks. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Export Mask** |
| Wellfleet Default: | None |
| Options: | Depends on the address class of the network address. |
| Function: | Specifies the range of addresses upon which this filter acts.<br><br>For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Export Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Export Mask, the Net ID and Subnet ID portions of the address will be filtered.<br><br>If the Export Address field was left blank, Export Mask should be left blank also. |
| Instructions: | Enter the appropriate mask in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Protocol** |
| Wellfleet Default: | None |
| Options: | RIP, Static, Direct, OSPF |
| Function: | Identifies the source of the routing information: direct connection, static route, or RIP-derived route. |
| Instructions: | Select the appropriate option. |

| | |
|---|---|
| **Parameter :** | **Interface** |
| Wellfleet Default: | None |
| Options: | Any IP address |
| Function: | Identifies the outgoing IP interface for the RIP update. This filter will only apply to this interface.<br>If the Interface field is left blank, this filter applies to all interfaces. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. |

### Editing a RIP Export Route Filter

The Configuration Manager allows you to edit three parameters (Enable, Action, and Metric) associated with an export route filter. You edit these parameters in the RIP Export Route Filters Window. To display this window for a particular export route filter, do the following:

- First, select the export route filter you wish to edit from the RIP Export Route Filters scroll box in the RIP Export Route Filters List Window (see Figure 10-25).
- Second, click on the Edit button to display the RIP Export Route Filters Window for that route filter.

This section provides information you need to see each parameter in the RIP Export Route Filters Window. Refer to this Information as necessary. When you are done, click on the Save button to exit the window and save your changes.

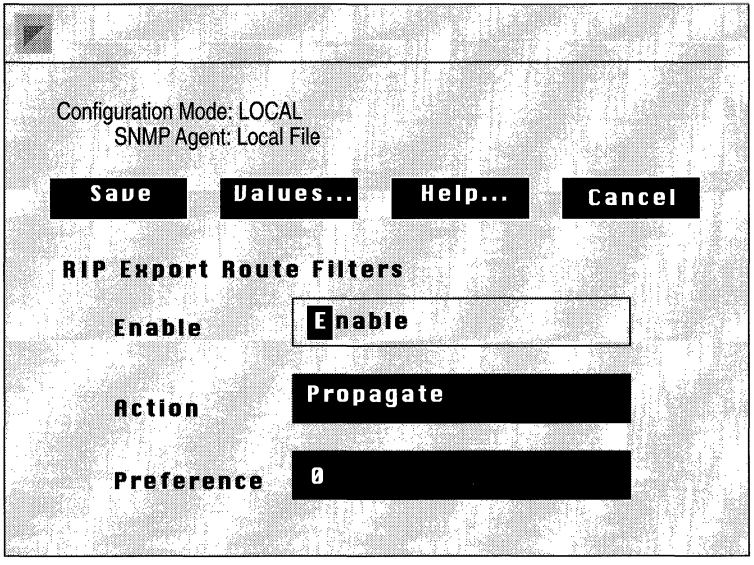


**Figure 10-27. RIP Export Route Filters Window**

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or disables this export route filter. |
| Instructions: | Set to Disable if you want to disable this export route filter.<br>Set to Enable if you previously disabled this export route filter and now want to re-enable it. |

| **Parameter :** | **Action** |
|---|---|
| Wellfleet Default: | Propagate |
| Options: | Propagate/Ignore |
| Function: | Controls the flow of routing information. If Action is set to Propagate, this route is advertised. If Action is set to Ignore, advertising of this route is suppressed. |
| Instructions: | Either accept the default Propagate, or select Ignore. |

| **Parameter :** | **Metric** |
|---|---|
| Wellfleet Default: | 0 (0 = the actual route cost as learned) |
| Options: | 0 - 15 |
| Function: | Assigns a RIP cost to the propagated route. The value 0 causes the actual route cost (as learned) to be used. |
| Instructions: | Either accept the default Metric value 0, or enter a new value. |

**Note:** Do not use a value that exceeds the diameter of the RIP network.

### Deleting a RIP Export Route Filter

You delete a RIP export route filter, as follows:

- First, select the export route filter you wish to delete in the RIP Export Route Filters scroll box in the RIP Export Route Filters List Window (see Figure 10-26).
- Second, click on the Delete button to display the Delete RIP Export Route Filter Window for the export route filter you selected.
- Finally, click on the Delete button to delete the export route filter and to return to the Route Filters Window for export route filters, which no longer displays that export route filter.

## Editing OSPF Import Route Filters

The OSPF Import Route Filters List Window (see Figure 10-28), allows you to add, edit, and delete OSPF import route filters. OSPF route filters pertain only to AS Boundary routers. The following sections describe each procedure. To begin, display the OSPF Import Route Filters List Window, by selecting the Protocols/IP/Route Filters/OSPF/ Import Filters option in the Wellfleet Configuration Manager Window.

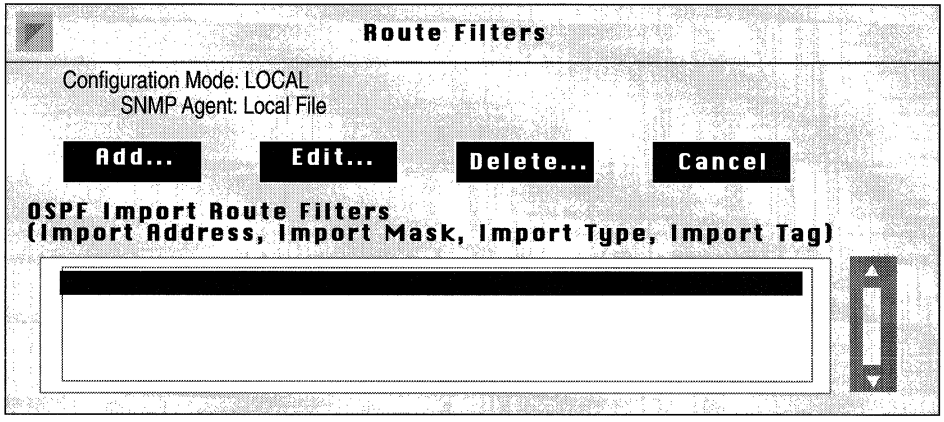


**Figure 10-28. OSPF Import Route Filters List Window**

### Adding a OSPF Import Route Filter

To add an import route filter, click on the Add button in the OSPF Import Route Filters List Window (see Figure 10-28), to display the OSPF Import Route Filter Configuration Window (see Figure 10-29).

This section provides information you need to set each parameter in the OSPF Import Route Filter Configuration Window. Refer to this information as necessary. When you are done, click on the Save button to exit the window and to add the import route filter.

**Note:** When you add an import route filter, the Configuration Manager automatically sets three additional parameters (specifically, it enables the filter, sets the Action parameter to Accept, and sets the Preference parameter to 0); however, these parameters are not displayed in the RIP Import Route Filter Configuration Window. If you wish to edit these parameters, refer to *Editing an Import Route Filter* for instructions.

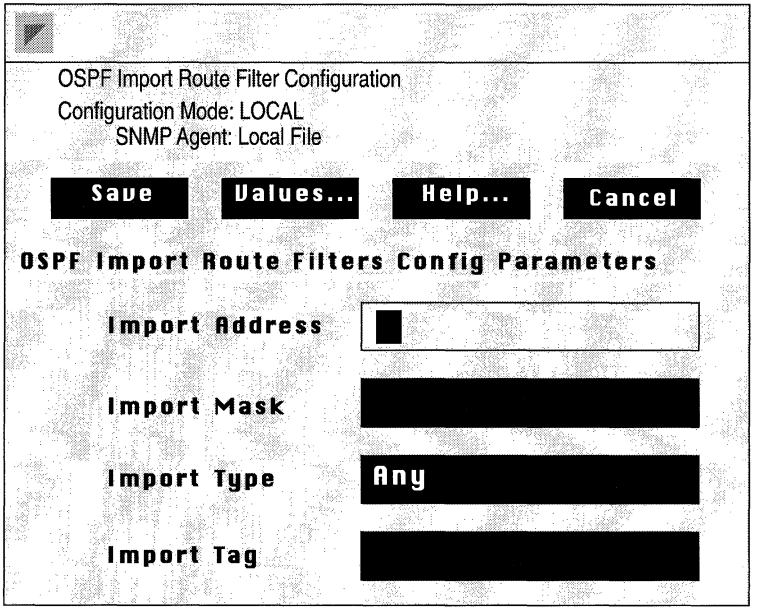


**Figure 10-29. OSPF Import Route Filter Configuration Window**

| | |
|---|---|
| **Parameter :** | **Import Address** |
| Wellfleet Default: | None |
| Options: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. If this field is left blank, the filter applies to all networks. |
| Instructions: | Enter the appropriate network address in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Import Mask** |
| Wellfleet Default: | None |
| Options: | Depends on the address class of the network address. |
| Function: | Specifies the range of addresses upon which this filter acts.<br><br>For example, consider Class B Network 172.32.0.0. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Import Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Import Mask, the Net ID and Subnet ID portions of the address will be filtered.<br><br>If the Import Address field was left blank, the Import Mask should be left blank also. |
| Instructions: | Enter the appropriate mask in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Import Type** |
| Wellfleet Default: | Any |
| Options: | Any, Type1/Type2 |
| Function: | Indicates the type of route to which this filter applies. If you choose Any, the filtering does not rely on the type of route. Type 1 indicates that only AS External Type 1 routes are to be filtered. Type 2 indicates that only AS External Type 2 routes are to be filtered. |
| Instructions: | Either accept the default, Any, or select another option. |

| | |
|---|---|
| **Parameter :** | **Import Tag** |
| Wellfleet Default: | None |
| Options: | Any decimal number |
| Function: | Indicates the tag with which this route filter is concerned. Each AS External Advertisement contains a tag field. If the tag field matches Import Tag, the appropriate action is taken, either the route is accepted or ignored.<br><br>Import Tag is pertinent to AS External Advertisements only. |
| Instructions: | Enter the appropriate tag number. |

### Editing an OSPF Import Route Filter

The Configuration Manager allows you to edit three parameters (Enable, Action, and Preference) associated with an OSPF import route filter. You edit these parameters in the OSPF Import Route Filters Window (see Figure 10-23). To display this window for a particular import route filter, do the following:

- First, select the import route filter you wish to edit in the OSPF Import Route Filters scroll box in the OSPF Import Route Filters List Window (see Figure 10-28).
- Second, click on the Edit button to display the OSPF Import Route Filters Window for that route filter.

This section provides information you need to set each parameter in the OSPF Import Route Filters Window. Refer to this information as necessary. When you are done, click on the Save button to exit the window and to save your changes.

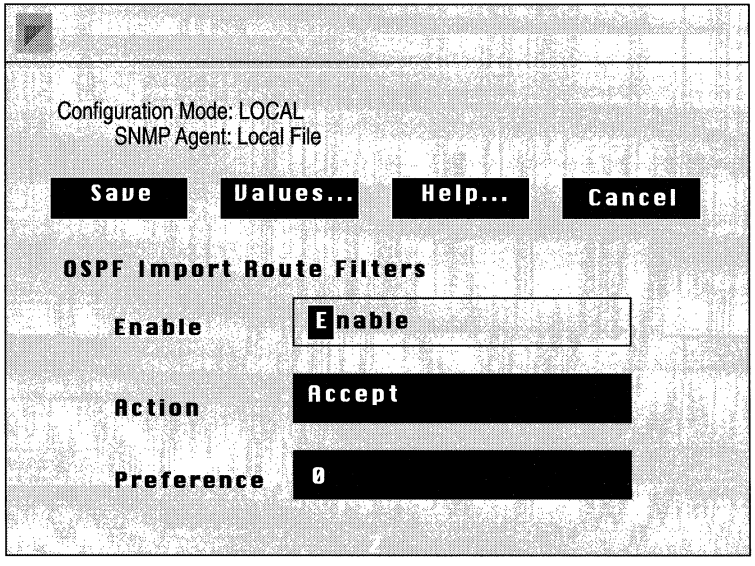


**Figure 10-30. OSPF Import Route Filters Window**

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Range: | Enable/Disable |
| Function: | Enables or disable this import route filter. |
| Instructions: | Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now wish to re-enable it. |

| | |
|---|---|
| **Parameter :** | **Action** |
| Wellfleet Default: | Accept |
| Range: | Accept/Ignore |
| Function: | Specifies whether the route is transferred to the routing tables. If Action is set to Accept (default), the routing information is sent to the routing tables. If Action is set to Ignore, the routing information is dropped. |
| Instructions: | Either accept the default Accept, or select Ignore. |

**Parameter :** **Preference**

Wellfleet Default: 0

Range: 0 through 16

Function: Assigns a weighted precedence value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, OSPF external, and RIP.

If this hierarchy is acceptable, accept the default value 0 for preference. If you want to grant preference to this OSPF-derived route, assign a new preference value in the range of 1 to 16 (the greater the number, the higher the preference).

**Note:** The default preference for static routes is 0, but may be set to any value between 0 and 16 (refer to *Editing Static Route Parameters* for more information). If you want to grant a OSPF-derived route preference over a static route, make sure the preference value you assign to the OSPF-derived route is greater than the preference value of the static route you want it to override.

Instructions: Either accept the default value 0, or enter a new value.

### Deleting an OSPF Import Route Filter

You delete an OSPF import route filter, as follows:

- First, select the OSPF import route filter you wish to delete in the OSPF Import Route Filters scroll box in the OSPF Import Route Filters Window (see Figure 10-21).
- Second, click on the Delete button to display the Delete OSPF Import Route Filters Window for the import route filter you selected.
- Finally, click on the Delete button to delete the import route filter and to return to the OSPF Import Route Filters Window, which no longer displays that import route filter.

## Editing OSPF Export Route Filters

The OSPF Export Route Filters List Window (see Figure 10-31), allows you to add, edit, and delete OSPF export route filters. The following sections describe each procedure. To begin, display the OSPF Export Route Filters List Window, by selecting the Protocols/IP/Route Filters/ OSPF/Export Filters option in the Wellfleet Configuration Manager Window.

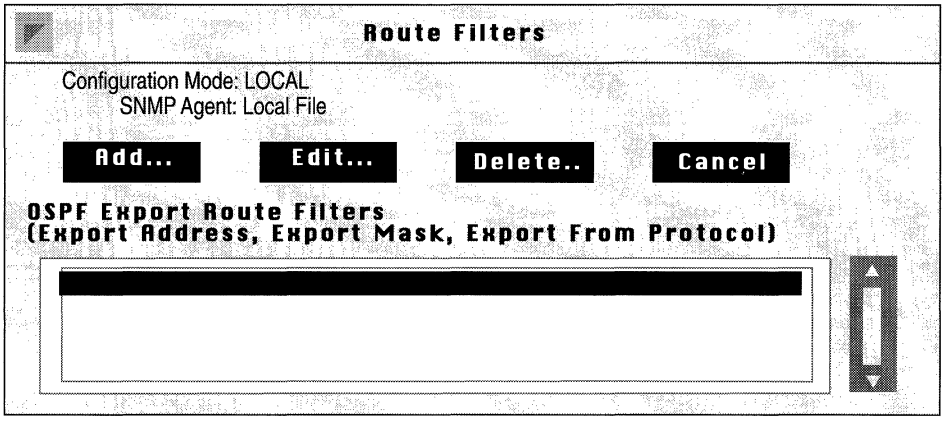


**Figure 10-31. OSPF Export Route Filters List Window**

### Adding an OSPF Export Route Filter

To add a OSPF export route filter, click on the Add button in the OSPF Export Route Filters List Window (see Figure 10-31), to display the OSPF Export Route Filter Configuration Window (see Figure 10-32).

This section provides information you need to set each parameter in the OSPF Export Route Filter Configuration Window. Refer to this information as necessary. When you are done, click on the Save button to exit the window and to add the export route filter.

**Note:** When you add an export route filter, the Configuration Manager automatically sets three parameters (specifically, it enables the filter, sets the Action parameter to Propagate, and Type and Tag); however, these parameters are not displayed in the OSPF Export Route Filter Configuration Window. If you wish to edit these parameters, see *Editing an Export Route Filter* for instructions.

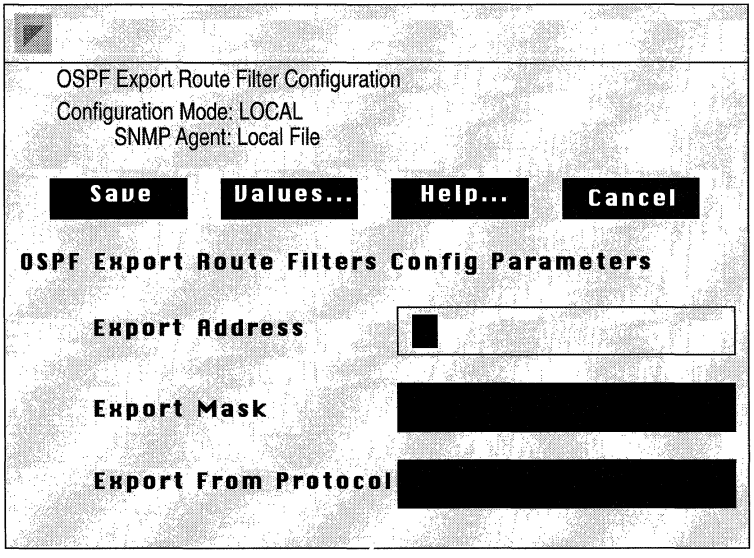


**Figure 10-32. OSPF Export Route Filter Configuration Window**

| | |
|---|---|
| **Parameter :** | **Export Address** |
| Wellfleet Default: | None |
| Options: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. If left blank, the filter applies to all networks. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Export Mask** |
| Wellfleet Default: | None |
| Options: | Depends on the address class of the network address. |
| Function: | Specifies the range of addresses upon which this filter acts.<br><br>For example, consider Class B Network 172.32.0.0. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Export Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Export Mask, the Net ID and Subnet ID portions of the address will be filtered.<br><br>If the Export Address field was left blank, Export Mask should be left blank also. |
| Instructions: | Enter the appropriate mask in dotted decimal notation. |

| Parameter : | Export From Protocol |
|---|---|
| Wellfleet Default: | None |
| Options: | RIP, Static, Direct |
| Function: | Identifies the source of the routing information: direct connection, static route, or RIP-derived route. |
| Instructions: | Select the appropriate option. |

### Editing an OSPF Export Route Filter

The Configuration Manager allows you to edit four parameters (Enable, Action, Type and Tag) associated with an OSPF export route filter. You edit these parameters in the OSPF Export Route Filters Window. To display this window for a particular export route filter, do the following:

- First, select the export route filter you wish to edit from the OSPF Export Route Filters scroll box in the OSPF Export Route Filters List Window (see Figure 10-31).
- Second, click on the Edit button to display the OSPF Export Route Filters Window for that route filter.

This section provides information you need to see each parameter in the OSPF Export Route Filters Window. Refer to this Information as necessary. When you are done, click on the Save button to exit the window and save your changes.

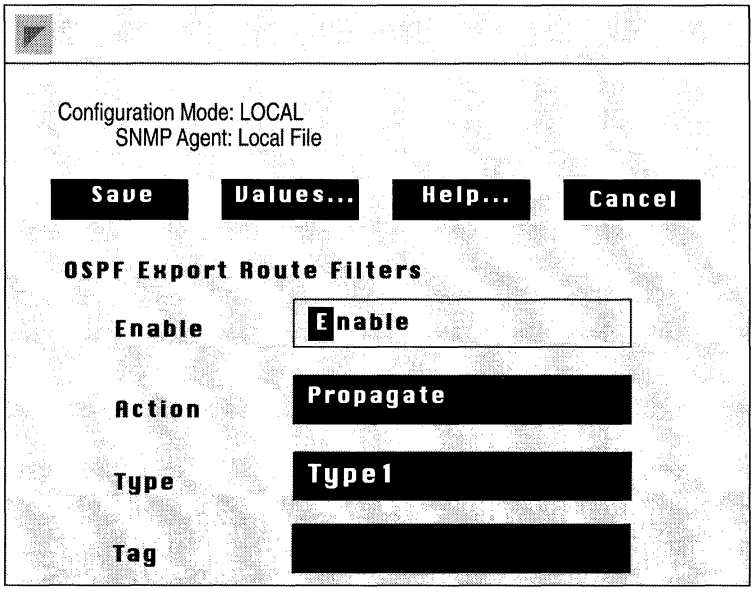


**Figure 10-33. OSPF Export Route Filters Window**

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or disables this export route filter. |
| Instructions: | Set to Disable if you want to disable this export route filter.<br><br>Set to Enable if you previously disabled this export route filter and now want to re-enable it. |

| | |
|---|---|
| **Parameter :** | **Action** |
| Wellfleet Default: | Propagate |
| Options: | Propagate/Ignore |
| Function: | Controls the flow of routing information. If Action is set to Propagate, this route is advertised. If Action is set to Ignore, advertising of this route is suppressed. |
| Instructions: | Either accept the default Propagate, or select Ignore. |

| | |
|---|---|
| **Parameter :** | **Type** |
| Wellfleet Default: | Type1 |
| Options: | Type1, Type2 |
| Function: | Indicates the type of AS External Advertisement generated for this network. If you choose Type 1, it indicates that AS External Type 1 routes are generated. Type 2 indicates that AS External Type 2 routes are generated.<br>This parameter has meaning only when Action is set to Propagate. |
| Instructions: | Either accept the default value, Type1, or select Type 2. |

| | |
|---|---|
| **Parameter :** | **Tag** |
| Wellfleet Default: | None |
| Options: | Any decimal number |
| Function: | Sets the tag value for the AS External Advertisement that is generated for this network.<br>This parameter has meaning only when Action is set to Propagate. |
| Instructions: | Enter the appropriate tag. |

### Deleting an OSPF Export Route Filter

You delete a OSPF export route filter, as follows:

- First, select the export route filter you wish to delete in the OSPF Export Route Filters scroll box in the OSPF Export Route Filters List Window (see Figure 10-31).
- Second, click on the Delete button to display the Delete OSPF Export Route Filter Window for the export route filter you selected.
- Finally, click on the Delete button to delete the export route filter and to return to the Route Filters Window for export route filters, which no longer displays that export route filter.

## Editing TFTP Parameters

The TFTP Parameters Window (see Figure 10-34) allows you to edit TFTP parameters. To display this window, select the Protocols/IP/ TFTP option in the Wellfleet Configuration Manager Window.

This section provides information you need to edit each parameter in the TFTP Parameters Window. Refer to this information to edit the parameters, you wish to change. When you are done, click on the Save button to exit the window and save your changes.

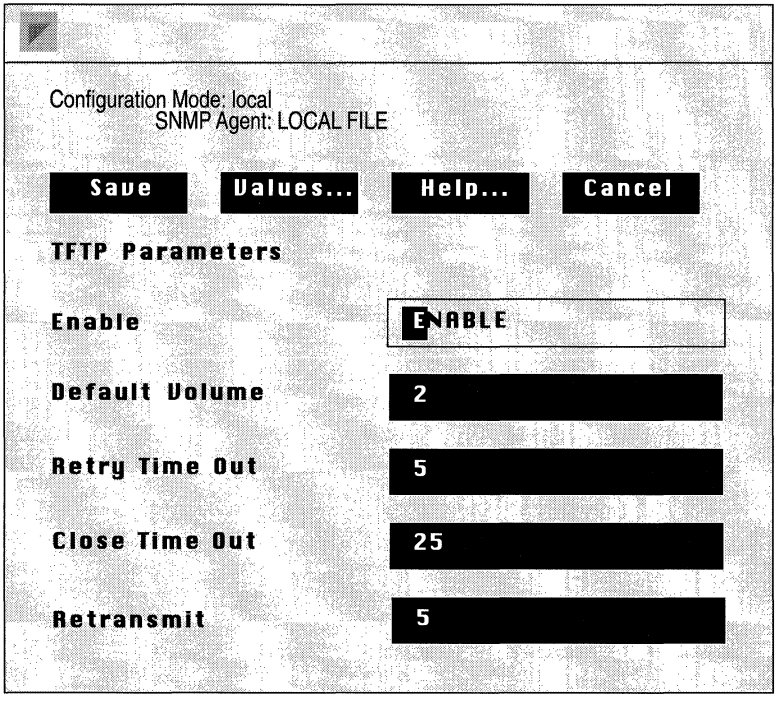


**Figure 10-34. TFTP Parameters Window**

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Specifies whether TFTP is enabled for the IP router. |
| Instructions: | Select Enable to enable TFTP for the IP router. Because TFTP allows write access to the BN's file system, Wellfleet recommends that you do not enable TFTP in network environments in which you are concerned with security.<br><br>Select Disable to disable TFTP for the IP router. |

| | |
|---|---|
| **Parameter :** | **Default Volume** |
| Wellfleet Default: | 2 |
| Options: | 2 to 5 |
| Function: | Specifies the number of the BLN slot (or volume) that will be used, by default, for all TFTP GETs and PUTs. |
| Instructions: | Specify the slot number. |

| | |
|---|---|
| **Parameter :** | **Retry Time Out** |
| Wellfleet Default: | 5 seconds |
| Options: | Any number of seconds |
| Function: | Specifies the number of seconds TFTP waits for an acknowledgment before retransmitting the last packet. |
| Instructions: | Specify a number of seconds. |

| | |
|---|---|
| **Parameter :** | **Close Time Out** |
| Wellfleet Default: | 25 seconds |
| Options: | Any number of seconds |
| Function: | Specifies the number of seconds TFTP waits, after it has successfully received a file, to make sure that the sender has received the last acknowledgment. |
| Instructions: | Specify a number of seconds. |

| | |
|---|---|
| **Parameter :** | **Retransmit** |
| Wellfleet Default: | 5 retransmissions |
| Options: | Any number of retransmissions |
| Function: | Specifies the number of times TFTP retransmits an unacknowledged message before abandoning the transfer attempt. |
| Instructions: | Specify the number of retransmissions. |

# Chapter 11

## Configuring OSPF

## List of Figures

## List of Tables

# Configuring OSPF

## About This Chapter

This chapter tells you how to edit parameters for the OSPF (Open Shortest Path First) protocol. Use this chapter if you want to:

- Learn about OSPF
- Enable or disable OSPF for the entire BN, or for specific interfaces
- Edit OSPF parameters, which includes:
  - Editing the global parameters
  - Adding, editing or deleting an OSPF area
  - Adding, editing or deleting OSPF area ranges
  - Editing an OSPF interface
  - Adding, editing or deleting OSPF interface neighbors
  - Adding, editing or deleting a virtual link interface

**Note:** This chapter does not explain how to add an OSPF interface; rather, it explains how to edit OSPF parameters relating to the entire BN and to circuits that already have been configured as OSPF circuits. If you want to add OSPF interfaces, you must do so at the circuit level; refer to the *Configuring Circuits* chapter for instructions.

When you originally added OSPF circuits (see *Configuring Circuits*), you assigned each circuit an area ID (specifying the area to which it belongs), and specified the area password. The OSPF circuits and areas you created then assumed the OSPF parameter defaults. Doing this for all of your OSPF circuits allows OSPF to run with a basic configuration. However; if you need to do any of the following things, you will need to use the instructions in this chapter.

- Configure virtual links

  You must configure virtual links for each area border router that does not reside within or directly interface to the backbone (see Figure 11-1). Every area border router must have a configured path to the backbone.

- Configure a preferred path

  Rather than just a hop count, OSPF considers the cost of a path when choosing the best path. Each interface; however, is assigned the default cost 1 for the path to which it interfaces. If you have a preferred path, you must edit the Metric Cost parameter for your interfaces. You will need to assign a higher Metric Cost for those paths which are *not* preferred paths.

- Configure timer values to match other devices

  If you have any devices in your network running OSPF, and are now adding a BN, you must make sure that the BN's timer values coincide with the timers in your other devices. Determine the timer values of the other devices, and change the BN's timer values to match them.

- Configure new OSPF area ranges/interfaces/neighbors/virtual links due to a topology change

  If there is a topology change (for example, if you add an area, combine two areas, move routers, etc.), you will have to reconfigure the appropriate OSPF elements.

The next section provides an OSPF overview; however, if do not wish to read it, go directly to *Editing Parameters.*

# OSPF Overview

OSPF is an internal gateway protocol intended for use in large IP networks. It exchanges routing information, using a link state algorithm, between routers in an Autonomous System, (a group of networks and routers which share routing information using the same routing protocol). OSPF responds quickly to topological changes, calculating new loop-free paths using only a small amount of routing protocol traffic.

OSPF supports three types of interfaces (or networks):

- *Point-to-point networks* join a single pair of OSPF routers. An example of such a network would be a network of synchronous lines.
- *Broadcast networks* support multiple routers, and can address a single physical message to all attached routers. Examples of such a network are Ethernet, FDDI and Token Ring.
- *Nonbroadcast multi-access networks* support multiple routers, and cannot address a single physical message to all routers. Examples of such a network are Frame Relay or X.25 networks.

There are several features in OSPF that contribute to routing overhead reduction, quicker convergence time, and increased security. Each of these features is discussed in the subsequent sections; they are:

- Link state protocol
- Routing areas
- Backbone
- Configurable cost metrics
- Variable length subnet masks
- Authentication
- External routes
- Neighbors, adjacencies and the Hello Protocol
- Designated and Backup Designated routers

## Link State Protocol

A link state protocol requires each router in the system to have synchronized databases, and from that, build a routing table. However, it does not require each router to send it's entire routing table to each of it's neighboring routers, as is the case with distance vector routing protocols, such as RIP. Instead, each router floods only link state change information throughout the system (a system, in this case, may be the Autonomous System, or a subset of the Autonomous System called an area). This process is referred to as the synchronization of the routers' topological databases.

With the link information, each router builds a Shortest Path Tree with itself as the root of the tree. It then can identify the shortest path from itself to each destination, and build it's routing table. Once the routers are synchronized and the routing tables are built, the routers will flood topology information only in response to some topological change (a disabled router, a downed line, etc.). It is obvious how this reduces network routing traffic when you consider that a distance

vector routing protocol requires it's routers to automatically flood their entire routing table every 30 seconds, regardless of whether there has been topological change.

## Routing Areas

OSPF allows the Autonomous System to be subdivided into areas. An area consists of groups of contiguous networks and hosts, and any router having an interface to any of the networks in the group. Each area maintains a separate copy of the basic routing algorithm, and shares an identical topological database with every other router in the area.

Dividing the Autonomous System into areas reduces the level of protocol traffic by reducing the amount of flooded topology change information. That is, when there is a topological change, the router floods this information to other routers in it's area only, rather than to every router in the Autonomous System.

Dividing the Autonomous System into areas also provides a level of security, in that the physical topology of an area is invisible to non-area residents. Conversely, routers that reside within a single area know nothing of the physical topology external to that area.

Certain routers can belong to multiple areas; these routers are called area border routers. An area border router will have a separate topological database for each area to which it interfaces. Router types are discussed later in this chapter.

## Backbone

The OSPF backbone is also an area. The backbone is the central area responsible for connecting all other areas and distributing routing information between them. It consists of the networks that are *not* included in any other area, and all area border routers (see Figure 11-1).

Area border routers are routers that interface to more than one area. Every area border router must belong to the backbone area. If an area border router does not interface directly to the backbone, a virtual link

must be configured from it to the backbone.These virtual links are configured like a point-to-point connection between one area border router and another. The area through which the virtual link is configured is called the transit area. Once configured, these virtual links are considered part of the backbone.

The backbone behaves like the other areas:

- It must be contiguous.
- It uses link state information to create an SPF tree, and from that, build a routing table.
- It's topology is invisible to all other areas, and it does not know the topology of other areas.

The backbone's main function is to distribute routing information between all of the other areas in the Autonomous System. It must always take the area id of 0.0.0.0.

### Stub Areas

Another type of area that OSPF supports is the stub area. These areas are dependent on default routes only to get out of the area; external routes are not flooded into or throughout stub areas. This not only reduces bandwidth overhead, but also the internal router's topological database size and, in turn, it's memory requirements.

Stub areas can be configured when there is just one point of exit from an area. This is usually characterized by a leaf or branch node (a router that has one connection to a LAN and one connection to the rest of the world. You may also configure a stub area when the choice of exit does not need to be made on a per-external-destination basis (you can get to all AS external routes through all of the stub's area border routers). Stub areas require the following conditions:

- Default routing must be used in a stub area (the stubs area's area border router must advertise a default route into the stub area).
- All routers in an area must agree on whether the area has been configured as a stub area.
- Virtual links cannot be configured through a stub area.
- AS boundary routers cannot be placed into a stub area.

## OSPF Router Classification

Segmentation of the AS into areas and a backbone results in four *functional* classifications of routers. They are described in Table 11-1. Note that the four classifications overlap; a router can fall into more than one classification.

**Table 11-1. OSPF Router Classifications**

| Router Type | Description/Function |
|---|---|
| Internal Router | The internal router resides within an area. All of its directly connected networks belong to the same area. Routers with only backbone interfaces also fall into this category. Each internal router runs a single copy of the basic routing algorithm. |
| Area Border Router | The area border router attaches to more than one area, and runs multiple copies of the basic routing algorithm - one copy for each area to which it is attached. An area border router distributes topological information about each of its attached areas to the backbone, then, the backbone distributes that same information to other areas. |
| Backbone Router | The backbone router is any router that has an interface to the backbone, including all routers that have an interface to more than one area (area border router). Backbone routers with all interfaces connected to the backbone are considered to be internal routers. |
| AS Boundary Router | The AS boundary router is the Autonomous System's link to other routing domains. The AS boundary router exchanges router information with routers belonging to other routing domains. Such a router has AS external routes that are advertised throughout the Autonomous System. The path to each AS boundary router is known to every other router in the Autonomous System. |

The following figure shows the different classifications of Wellfleet OSPF routers within an OSPF Autonomous System.

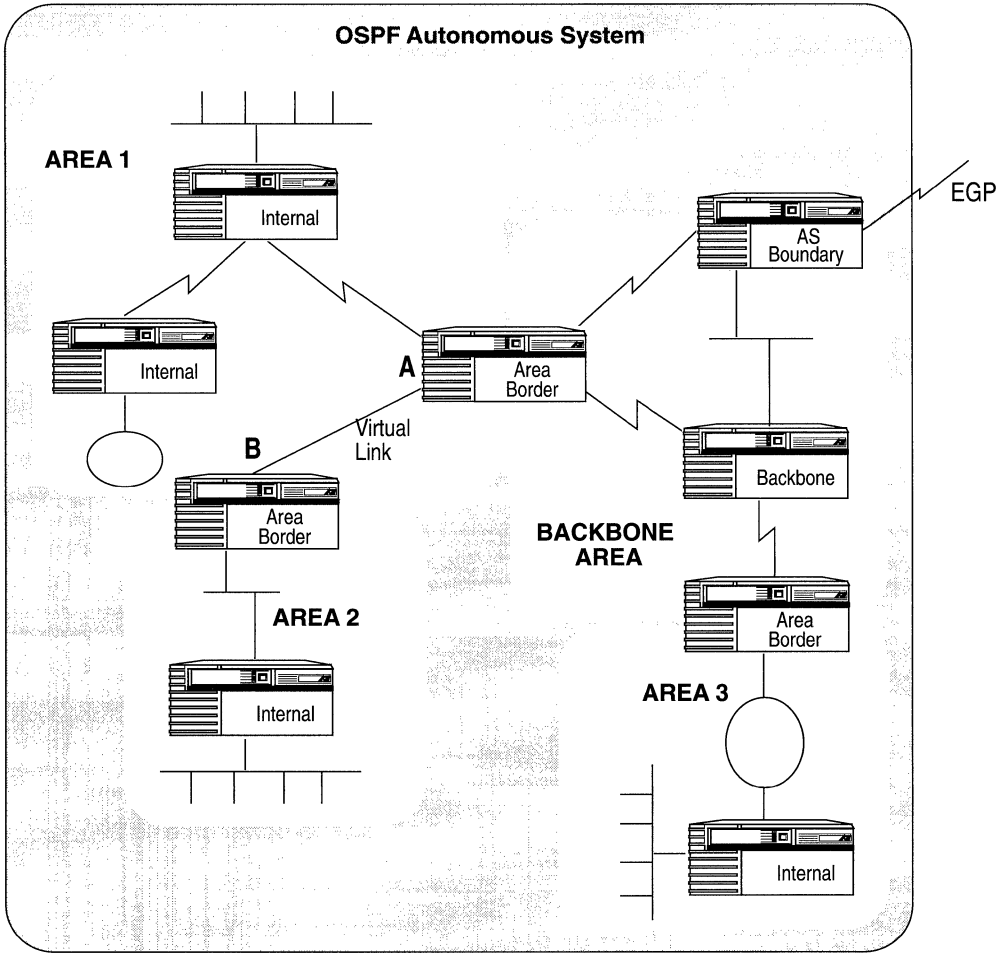


**Figure 11-1. OSPF Autonomous System**

In the network shown in Figure 11-1, a virtual link would need to be configured from area border router A, through Area1 (the transit area), to area border router B. This is necessary to restore the contiguity of the backbone, because area border router B does not have a direct physical connection to the backbone.

### Routing: Intra, Inter and External

Because the Autonomous System can be divided into areas, there are two types of routing within the Autonomous System: intra-area routing and inter-area routing. Intra-area routing takes place between a source and destination that reside in the same area. Inter-area routing takes place between a source and destination residing in different areas. Inter-area routing always involves area border routers, which are responsible for providing the source area with information about the topology of some other area(s) in the AS.

The third type of routing, external routing, takes place between a source and destination that are in different routing domains. External routing always involves AS boundary routers, which are responsible for providing the AS in which it resides information about other routing domains. It floods this information throughout its own AS, excepting all stub areas. Paths to AS boundary routers are summarized by nonstub area border routers.

## Configurable Cost Metrics

In contrast to RIP, a distance vector routing protocol, which considers only a hop count in calculating the best path, OSPF considers a cost metric that you assign to a path.

OSPF recognizes that a simple hop count takes no account of reliability, bandwidth, delay, or actual dollar cost of using a path. Passing through an extra hop to get to a 1.54 Mb T1 channel, for instance, may be more efficient than traversing a shorter, but slower route. For OSPF, the best path is the one that offers the least cost metric delay. With Wellfleet's implementation of OSPF, every path automatically takes a cost metric value of 1. You must configure cost metrics if you want to specify a preferred path. To specify a preferred path, you would allow the preferred path to retain the cost metric value of 1, and then assign higher cost metric values to the less preferred paths.

Figure 11-2 shows the benefit of using configurable cost metrics. Assigning the 56Kb line a cost metric value of 10 forces OSPF to choose the faster T1 line path as the best path, despite the extra hop, when transmitting a packet from Host A to Host B.

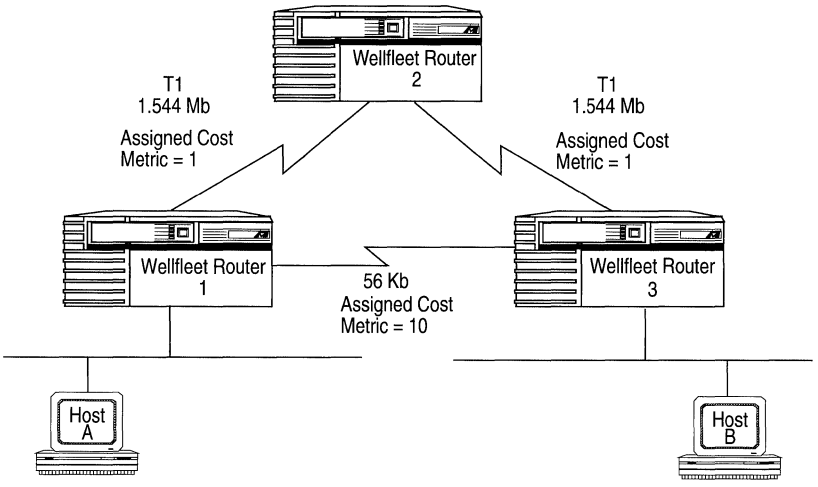


**Figure 11-2. Configurable Cost Metrics Usage Example**

## Variable Length Subnets

Subnetting is a strategy that allows a number of networks to be aggregated by address and appear to be one single network. OSPF supports IP *variable length* subnet masks. Variable length subnet masks allow IP's address space to be used more efficiently by dividing the network number into subnets of varying sizes. There is no natural address class restriction on the masks.

In OSPF, for advertised route you can configure a corresponding subnet mask that indicates an address range being described by the route. The area border router sends a summary link advertisement for the single address/mask pair, rather than sending one summary link advertisement for every network defined by the address/mask pair, thus reducing the amount of routing traffic. For example, a summary advertisement for the destination 140.191.0.0 with a mask of 255.255.0.0 actually is describing a single route to the collection of destinations 140.191.0.0 to 140.191.255.255. When a packet is forwarded, it is always forwarded to the network that is the best (longer or most specific) match for the packet's destination.

## Authentication

OSPF provides a measure of security through the use of passwords. Every router in an area must either have the same password configured, or no password at all. Two routers can communicate only if the password is properly configured. When a packet is received, the password is checked before anything is done with the packet. Unauthorized routers are not allowed to communicate with the OSPF system.

## External Routes

External routes are learned and propagated by AS boundary routers. These routers run not only OSPF (on interfaces internal to the AS), but may also run some exterior gateway protocol (on the interface that connects to another AS), such as EGP. The following list defines routes that OSPF considers external routes:

- a route to a destination outside the AS
- a static route
- a default route
- a route derived by RIP

These external routes can be tagged if you configure route filters, (see *Configuring the IP Router* for information on route filters), to identify the system that is delivering the routes to the OSPF system. The AS boundary router floods an External Links Advertisement, describing these routes, throughout the entire AS.

## Neighbors, Adjacencies and the Hello Protocol

In each OSPF network, neighbors are discovered and maintained through OSPF's Hello Protocol. Neighbors are any two routers that have an interface to the same network. On a Broadcast or Point-to-Point network, the Hello Protocol dynamically discovers neighbors; however, on a nonbroadcast multi-access network, the manual configuration of neighbors is necessary.

The Hello Protocol is responsible for ensuring that communication between neighbors is bidirectional. Periodically, OSPF routers send out Hello Packets over all interfaces. Included in these Hello packets are:

- the router's priority
- the router's hello timer value
- a list of routers that have sent this router Hello packets
- the router's choice for Designated Router and Backup Designated Router

Bidirectional communication is determined when one router sees itself listed in the neighbor's Hello packet.

Neighbors may form a relationship called an adjacency for the purpose of exchanging routing information. Not every pair of routers forms this relationship; although, all routers connected by a point-to-point network, or a virtual link will always form an adjacency. Also, every router on a multi access network forms this relationship with the Designated Router.

When two routers form an adjacency, the routers go through a process to synchronize their topological databases. When their databases are synchronized, the routers are said to be fully adjacent. From this point on, only routing change information is passed between the adjacencies, thus conserving bandwidth.

## Designated and Backup Designated Routers

To further reduce the amount of routing traffic, the Hello Protocol elects a designated router and a backup designated router on each multi-access network. Instead of neighboring routers forming adjacencies and swapping link state information with each other (which on a large network can mean a lot of routing protocol traffic), all routers on the network form an adjacency with the designated router and the backup designated router only and send link state information to it. The designated router then redistributes the information from each router to every other router in the form of a network links advertisement.

In case the designated router goes down, a backup designated router is always elected at the same time that the designated router is elected. It's responsibility is to take over all of the designated router's functions should the designated router fail.

# Summary

OSPF is a routing protocol, based on link state technology. The OSPF domain consists of an Autonomous System, which is further divided into areas (a contiguous group of networks and hosts, and routers having interfaces to those networks), and a backbone (networks not contained in any area, routers attached to those networks, and routers attached to more than one network). Both the areas and the backbone must be contiguous. If the backbone becomes physically non-contiguous, it's contiguity may be restored through the configuration of virtual links. The virtual link is configured from one area border router to another. The area through which it is configured is called the transit area.

Within the areas and the backbone reside four different types of routers: internal, backbone, area border, and AS boundary routers. These are functional classifications and can overlap.

OSPF supports three types of interfaces (networks): point-to-point, broadcast, and nonbroadcast multi-access. It also supports IP subnetting, address ranges, and stub areas. Stub areas rely on default routing.

OSPF routing can be broken into three categories: intra-area routing, which occurs when source and destination reside within the same area; inter-area routing, which occurs when source and destination reside in different areas; and external routing, which occurs when the destination resides within different Autonomous Systems, or between OSPF and RIP networks within the same Autonomous System.

Routing in OSPF depends on all routers in an area having synchronized databases for that area. First, the router in an area discovers neighbors through the Hello Protocol — the router sends periodic Hello Packets out all interfaces and checks to see themselves

listed in the Hello Packets they receive from other routers. They then form an adjacency relationship with certain neighbors, or, on a multi-access network, with the Designated Router and the Backup Designated Router. This relationship is established to facilitate the distribution of routing information, All routing protocol packets, except for the Hello Packet, are sent only over adjacencies.

Through the issuing of Link State Advertisements, the adjacent routers synchronize their area topology databases, thus facilitating routing between sources and destinations in that same area. To route beyond the confines of its area, a router depends on area border routers. Area border routers advertise topology information to the backbone, the backbone in turn, advertises this information to all other areas, thus facilitating routing between different areas. The AS boundary router exchanges information with routers from other Autonomous Systems, or with routers from RIP networks within the same Autonomous System. Each AS boundary router has AS external routes, which it advertises throughout the Autonomous System. The path to every AS boundary router is known to each router in the OSPF network, thus facilitating routing to external networks, such as RIP and EGP networks.

# Implementation Notes

This section provides some suggestions to help you when configuring your OSPF network. Wellfleet's OSPF does not restrict you to these suggestions, but is providing them as they may be helpful to you.

- Keep the same password throughout an area, or even throughout the entire OSPF AS, if possible.
- Use the default timers, unless you are running 9.6K sync lines. In this case, double the default timers on both ends of the link.
- Use address ranges if your network is a subnetted network.
- Keep all subnets within one area. If you cross areas, you cannot configure summaries.
- Make sure AS Border Router is enabled if the router has any non-OSPF interfaces.

# OSPF References

If you would like more information about OSPF, refer to the following documents:

- Moy, J. *OSPF Version 2*. RFC 1247, Network Information Center (NIC), SRI International, Menlo Park, CA, July 1991
- Comer, Douglas E. *Internetworking with TCP/IP, Volume I: Principle, Protocols, and Architecture*. Second Edition. Prentice Hall, Inc., Englewood Cliffs, New Jersey, 1991
- Perlman, Radia. *Interconnections: Bridges and Routers.*, Addison-Wesley Publishing Company, Reading, MA, First Printing, May 1992.

# Editing Parameters

Once you have configured a circuit to support OSPF, you can use the Configuration Manager to edit OSPF parameters. The configuration function you wish to perform determines the type of parameters you must edit. (see Table 11-2).

**Table 11-2. OSPF Configuration Functions**

| To Do the Following: | See this Section: |
|---|---|
| Enable or disable OSPF for the entire BN | *Editing OSPF Global Parameters* |
| Editor delete an area | *Editing OSPF Area Parameters* |
| Enable or disable OSPF on a particular circuit | *Editing OSPF Interface Parameters* and/or *Editing OSPF Virtual Interface Parameters* |
| Edit or delete an OSPF interface | *Editing OSPF Interface Parameters* |
| Add, edit or delete an OSPF Virtual link | *Editing OSPF Virtual Interfaces Parameters* |

This section describes how to access and edit the OSPF parameters listed in Table 11-2. For each parameter, it provides the following:

- Wellfleet default
- Valid options
- Parameter's function
- Instructions for setting the parameter

You begin from the Configuration Manager Window (Figure 11-3); the first window displayed when you enter the Configuration Manager application.

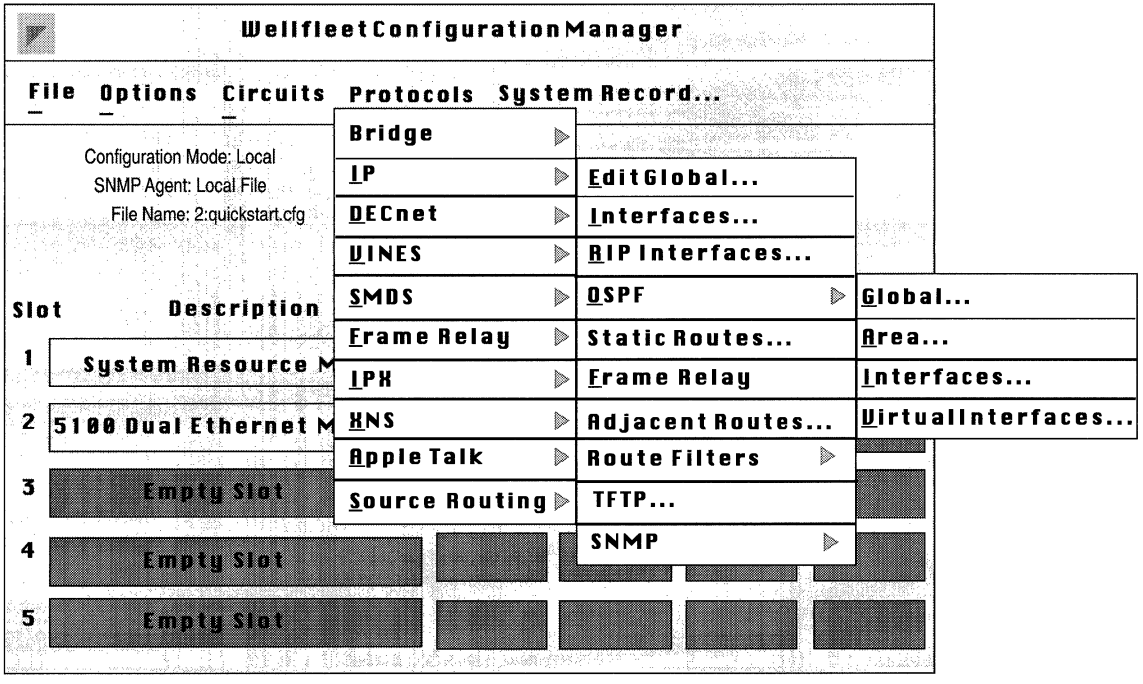


**Figure 11-3. Configuration Manager Window**

## Editing OSPF Global Parameters

When you edit OSPF global parameters, you are editing parameters that affect OSPF on the entire router. To edit OSPF Global parameters, begin at the Wellfleet Configuration Manager Window and complete the following steps:

1. Select the Protocols/IP/OSPF/Global option.

   The OSPF Global Parameters Window appears (Figure 11-4).

2. Edit those parameters you wish to change.

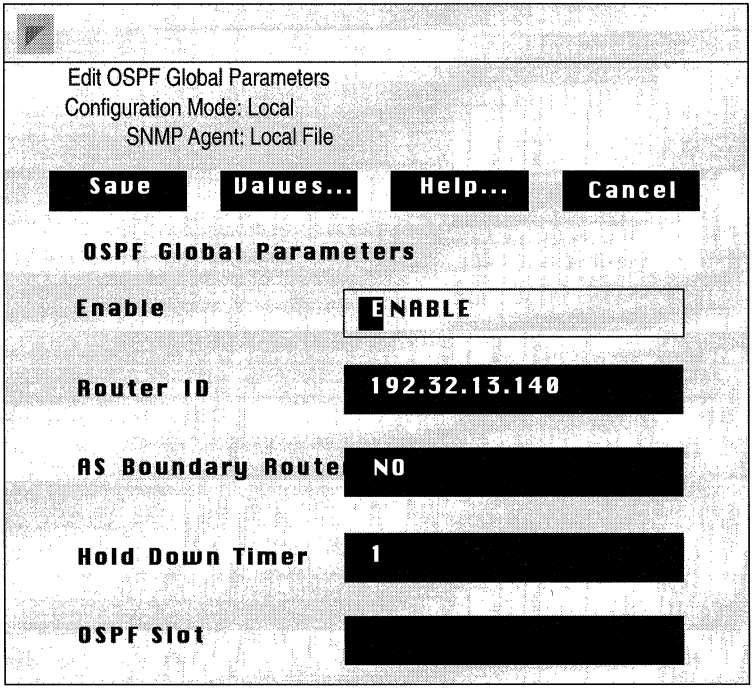


**Figure 11-4. OSPF Global Parameters Window**

3. Click the Save button to exit the window and save your changes when you are finished.

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | This parameter allows you to globally enable or disable OSPF on all BN interfaces. |
| Instructions: | Set to Disable if you want to disable OSPF for the entire router. Set to Enable if you previously disabled OSPF on the router and now wish to reenable it. |

| | |
|---|---|
| **Parameter :** | **Router ID** |
| Wellfleet Default: | The IP address of the first OSPF circuit configured on this router. |
| Options: | Any IP address, preferably, one of the router's IP interface addresses. |
| Function: | This IP address uniquely identifies this router in the OSPF domain. By convention, and to ensure uniqueness, one of the router's IP interface addresses should be used as the Router ID.<br><br>The Router ID will determine the Designated Router on a broadcast link if the priority values of the routers being considered are equal. The higher the Router ID, the greater it's priority. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **AS Boundary Router** |
| Wellfleet Default: | No |
| Options: | Yes/No |
| Function: | This parameter indicates whether or not this router will function as an AS Boundary Router.<br><br>The router can be an AS Boundary router if one or more of its interfaces is connected to a non-OSPF network (for example, RIP or EGP). |
| Instructions: | Set this parameter to Yes if this router will function as an AS Boundary Router. Otherwise; accept the default value, No. |

| | |
|---|---|
| **Parameter :** | **Hold Down Timer** |
| Wellfleet Default: | 1 |
| Options: | 0 - 10 seconds |
| Function: | Prevents the algorithm from running more than once per the value of Hold Down Timer. It's purpose is to free up the CPU. Note that a value of 0 means there is no hold down time. |
| Instructions: | Either accept the default value of 1 second, or enter a new value. |

| Parameter : | OSPF Slot |
|---|---|
| Wellfleet Default: | All slots |
| Options: | Any slot on the BN |
| Function: | Indicates which slot(s) the OSPF soloist is eligible to run on. If the slot on which the OSPF soloist is running goes down, the BN will attempt to run OSPF another slot specified at the OSPF Slot parameter. |
| Instructions: | Select the appropriate slots. |

**Note:** Use caution when selecting the slot(s) on which OSPF may run. If you choose an empty slot, and it is the only slot you choose, OSPF will not run. Also, if you choose a slot that becomes disabled, and it is the only slot you choose, OSPF wont get restarted.

## Editing OSPF Area Parameters

To edit OSPF Area Parameters, begin at the Wellfleet Configuration Manager Window and proceed as follows:

1. Select the Protocols/IP/OSPF/Areas option.
2. The OSPF Area List Window appears (Figure 11-5).

   It is from this window that you can perform any of the functions described by the subsections listed below.

   — *Editing an Area*

   — *Deleting an Area*

   — *Adding a Range to an Area*

   — *Editing an Area's Range*

   — *Deleting a Range from an Area*

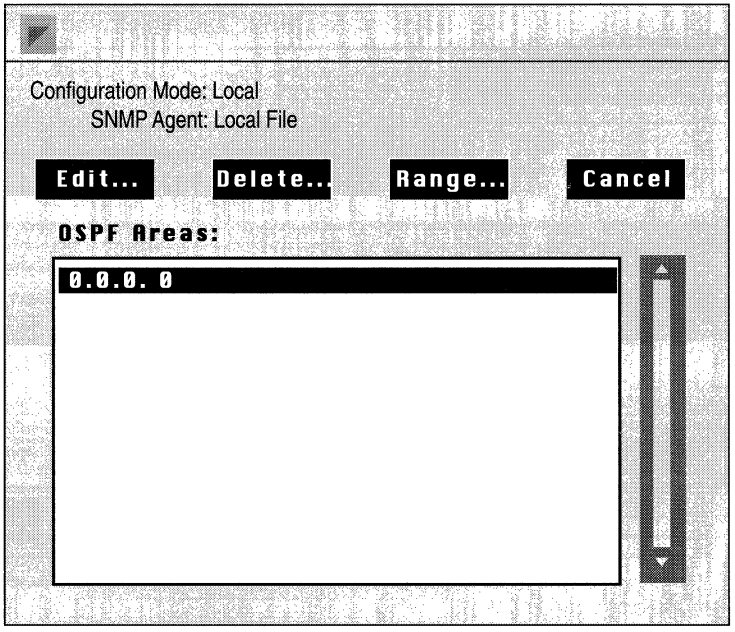


**Figure 11-5. OSPF Area List Window**

### Editing an Area

After you add areas, you may change any of your area's defaults. To edit an area, complete the following steps:

1. Select the area you wish to edit in the OSPF Areas scroll box on the OSPF Area List Window (Figure 11-5).
2. Click the Edit button.

   The OSPF Area Parameters Window appears (Figure 11-6).

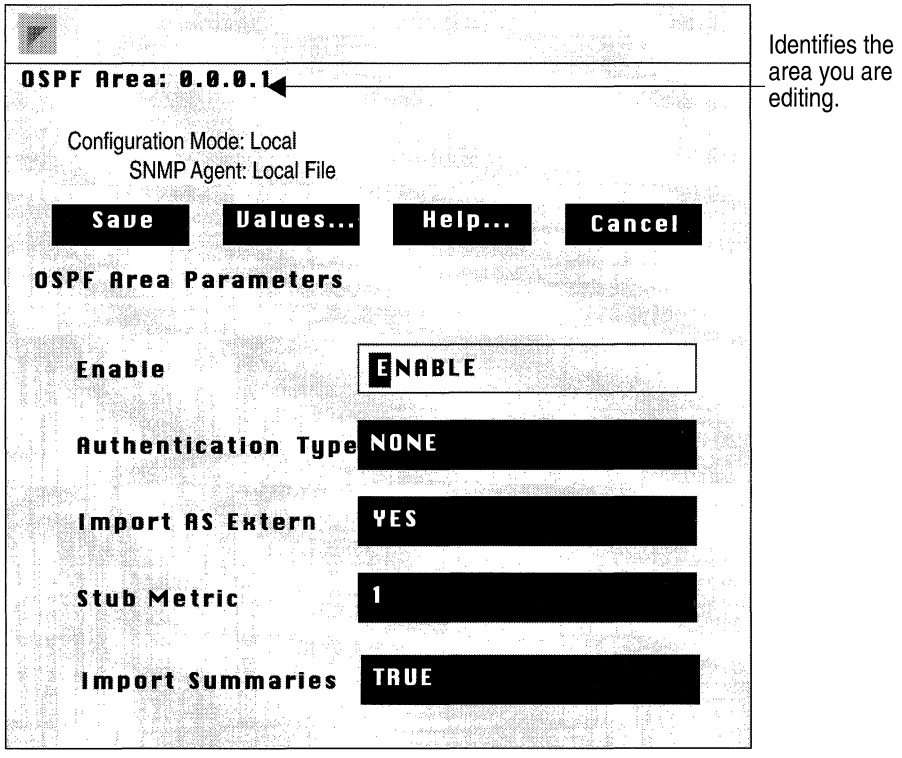


**Figure 11-6. OSPF Area Parameters Window**

3. Set the OSPF area parameters, then click the Save button to save your changes and exit the window.

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Allows you to enable and disable this area. This parameter is useful if you want to temporarily disable an area rather than delete it. |
| Instructions: | Set this parameter to Disable if you want to disable this area. Or, set this parameter to Enable if you previously disabled the area, and now wish to reenable it. This will cause OSPF to restart. |

| | |
|---|---|
| **Parameter :** | **Authentication Type** |
| Wellfleet Default: | None |
| Options: | None/Simplepassword |
| Function: | Enables or disables password authentication for the area. With Simplepassword chosen, enabling password authentication, only those routers sharing the correct password will be able to communicate with each other. If you do choose Simplepassword when you configure the interface, you will be prompted for the area password. If you accept the default None, password authentication is disabled for this area. |
| Instructions: | Either accept the default value None to disable password authentication, or select Simplepassword to enable password authentication. If you select Simplepassword, you will be prompted for the area password. |

| | |
|---|---|
| **Parameter :** | **Import AS Extern** |
| Wellfleet Default: | Yes |
| Options: | Yes/No |
| Function: | Indicates whether or not this area will import AS external link state advertisements. If this area does *not* import AS external link state advertisements, it is a Stub area. If it does import AS external link state advertisements, it is not a stub area. |
| Instructions: | Set this parameter to No if this area will function as a stub area. Otherwise; accept the default value, Yes. |

| | |
|---|---|
| **Parameter :** | **Stub Metric** |
| Wellfleet Default: | 1 |
| Options: | 1 to 255 |
| Function: | When an area border router is connected to a stub area, it generates a default link summary into the area indicating a default route. The Stub Metric indicates the cost of that route. By default, Stub Metric equals 1. This parameter has meaning only when Import AS Extern is set to No. |
| Instructions: | Either accept the Stub Metric default value 1, or supply some other Stub Metric value. |

| **Parameter :** | **Import Summaries** |
|---|---|
| Wellfleet Default: | True |
| Options: | True/False |
| Function: | Indicates whether of not network summaries should be flooded into a stub area. This variable has meaning only if Import AS Extern is set to No. |
| Instructions: | Either accept the Import Summaries default value True, or select False if Import AS Extern is set to No *and* you do not want network summaries imported into the stub area. |

### Deleting an Area

Sometimes, as the result of a topology change, you may want to delete an area. To delete an area, complete the following steps:

1. Select the area you wish to delete from the OSPF Areas scroll box in the OSPF Area List Window (see Figure 11-5).

   The Delete OSPF Area Window appears (Figure 11-7).

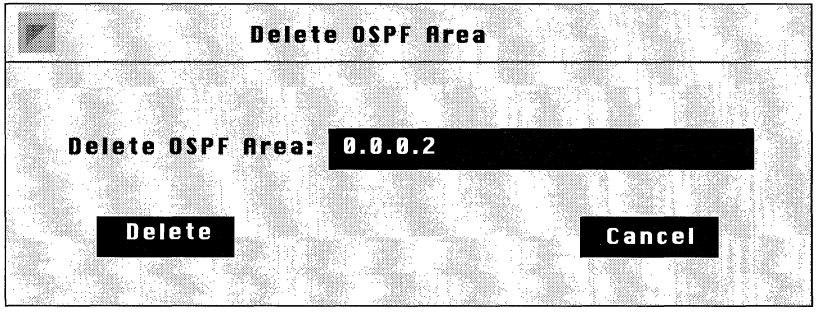


**Figure 11-7. Delete OSPF Area Window**

2. Click the Delete button if the area ID in the Delete OSPF Area box correctly identifies the area you wish to delete.

You are returned to the OSPF Area List Window; the area you just deleted no longer appears in the OSPF Areas scroll box.

### Adding a Range to an Area

You can define a range for an area. Ranges are address/mask pairs that allow you to group subnetted networks that reside in the same area, and to have that group be advertised by *one* network summary advertisement. Otherwise; a summary advertisement would be generated for each subnet in the area. To add a range to an area, complete the following steps:

1. Choose the area for which you wish to define a range from the OSPF Areas scroll box on the OSPF Area List Window (Figure 11-5).
2. Click the Ranges button.

   The OSPF Range List Window appears (Figure 11-8).

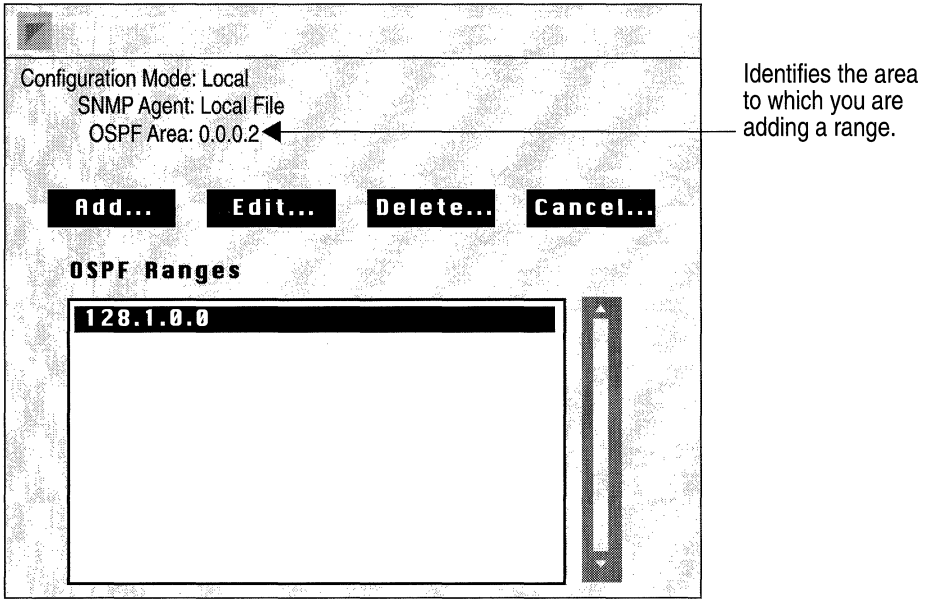


**Figure 11-8. OSPF Range List Window**

3. Click the Add button.

   The Add OSPF Range Window appears (Figure 11-9).

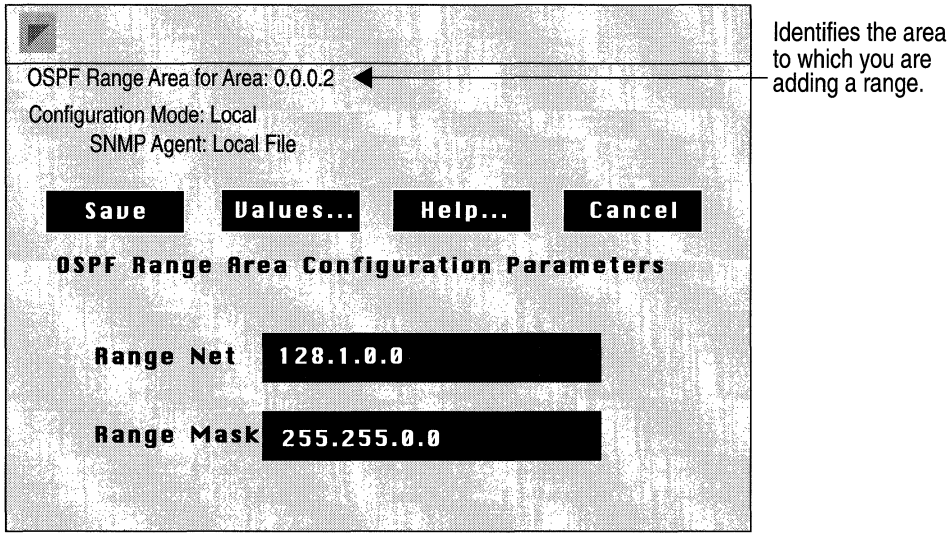


**Figure 11-9. Add OSPF Range Window**

4. Set the range parameters, then click the Save button when you are finished to save your changes and exit the window.

| | |
|---|---|
| **Parameter :** | **Range Net** |
| Wellfleet Default: | None |
| Options: | Any network number |
| Function: | This parameter allows you to assign a single network address to a group of subnets. This network address, together with the subnet mask you provide, specifies the subnets to be grouped in this area range. Just one link summary advertisement will be generated for all subnets in this range, rather than one link summary advertisement for each of the subnets included in that network. |
| Instructions: | Enter the appropriate network number in dotted decimal notation. |

| | |
|---|---|
| **Parameter :** | **Range Mask** |
| Wellfleet Default: | None |
| Options: | Any address mask |
| Function: | This parameter, together with Range Net, indicates all of the networks that belong to this range. The Range Mask is not restricted to the natural address class mask for the address supplied at Range Net.<br><br>In this example, Range Net is 128.1.0.0 and Range Mask is 255.255.0.0. That means that the link summary advertisement generated will summarize networks 128.1.0.0 to 128.1.254.254. |
| Instructions: | Enter the appropriate subnet mask in dotted decimal notation. |

**Note:** When setting up your OSPF network, keep all subnetted networks in the same area.

### Editing an Area's Range

Once you add a range to an area, you may edit the range. To edit a range, complete the following steps:

1. Choose the area for which you wish to edit a range from the OSPF Areas scroll box on the OSPF Area List Window (Figure 11-5).
2. Click the Ranges button.

   The OSPF Range List Window appears (Figure 11-10).

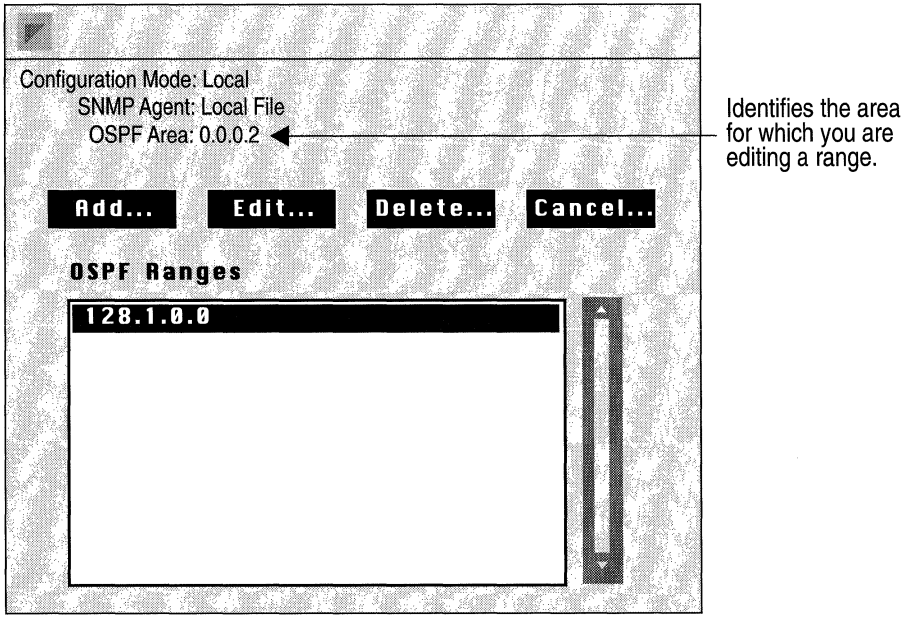


**Figure 11-10. OSPF Range List Window**

3. Choose the range that you wish to edit from the OSPF Range scroll box.
4. Click the Edit button.

   The OSPF Area Range Parameters Window appears (Figure 11-11).

5. Edit the parameters you wish to change, then click the Save button to save your changes and exit the window.

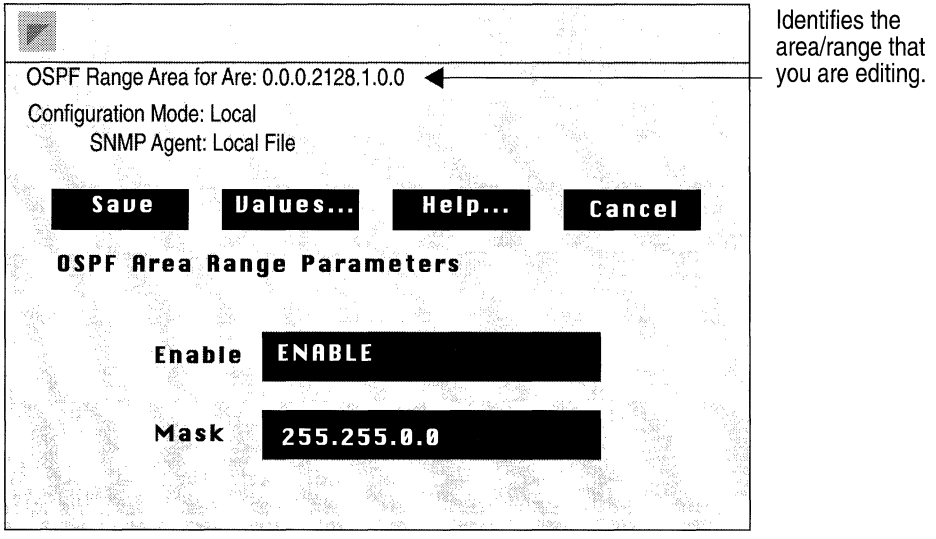


**Figure 11-11. OSPF Area Range Parameters Window**

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or disable this range for the specified area. This parameter is useful if you want to disable the range, rather than delete it. |
| Instructions: | Set this parameter to Disable if you want to disable this range. Set the parameter to Enable if you previously disabled this range and now wish to reenable it. |

| | |
|---|---|
| **Parameter :** | **Mask** |
| Wellfleet Default: | None |
| Options: | Any address mask |
| Function: | This parameter allows you to change the mask portion of this area range. Mask, together with Range Net, indicates all of the networks that belong to this range. The Mask is not restricted to the natural address class mask for the address supplied at Range Net.<br><br>In this example, Range Net is 128.1.0.0 and Range Mask is 255.255.0.0. That means that the link summary advertisement generated will summarize networks 128.1.0.0 to 128.1.255.255 |
| Instructions: | Enter the appropriate address mask in dotted decimal notation. |

### Deleting a Range from an Area

If you no longer want a range to be associated with an area, you can delete it. To delete a range, complete the following steps:

1. Select the area in the OSPF Area List Window for which you wish to delete a range (Figure 11-5).
2. Click the Ranges button.

   The OSPF Range List Window appears (Figure 11-12).

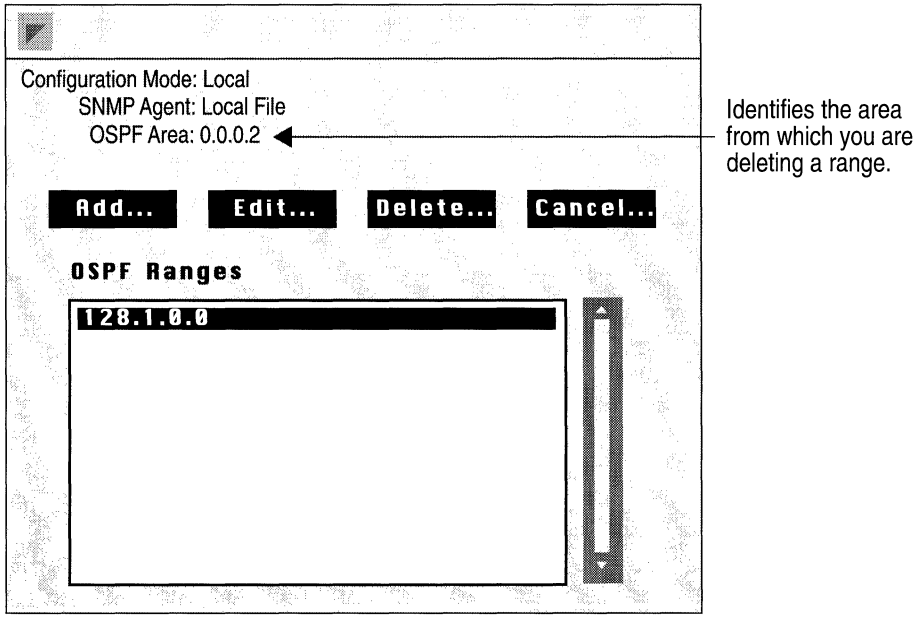


**Figure 11-12. OSPF Range List Window**

3. Select the range you wish to delete from the OSPF Ranges scroll box.
4. Click the Delete button.

   The Delete OSPF Range Window appears (Figure 11-13).

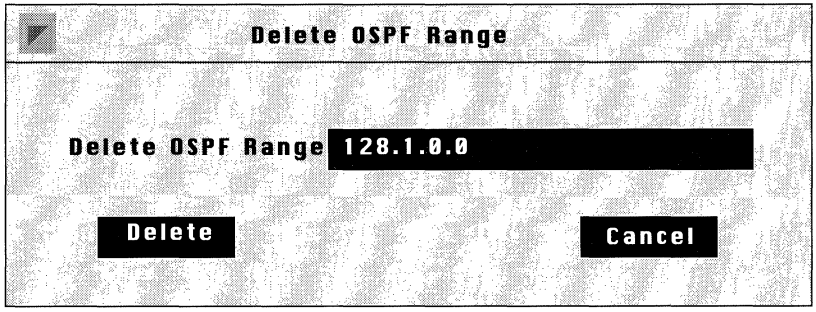


**Figure 11-13. Delete OSPF Range Window**

5. Click the Delete button if the range in the Delete OSPF Range box correctly identifies the range you wish to delete.

   This range no longer appears in the OSPF Ranges scroll box.

## Editing OSPF Interface Parameters

All OSPF interfaces assume certain default values when you first configure them. You can; however, change these defaults by editing the interface specific parameters. The changes you make affect only the interface you select.

To edit OSPF Interface Parameters, begin at the Wellfleet Configuration Manager Window and complete the following steps:

1. Select the Protocols/IP/OSPF/Interfaces option.

   The OSPF Interface List Window appears (Figure 11-14).

   It is from this window that you can perform any of the functions described by the subsections listed below.

   — *Editing an Interface*

   — *Adding a Neighbor to an Interface*

   — *Editing an Interface's Neighbor*

   — *Deleting a Neighbor from an Interface*

**Note:** You configure neighbors for NBMA interfaces only.

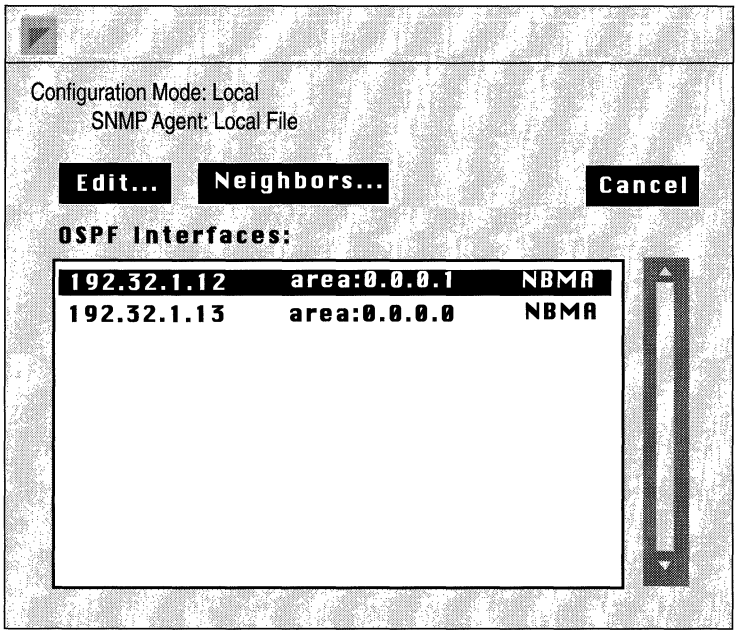


**Figure 11-14. OSPF Interface List Window**

### Editing an Interface

To edit an interface's parameter defaults, complete the following steps.

1. Select the interface you wish to edit from the OSPF Interfaces scroll box in the OSPF Interface List Window (Figure 11-12).
2. Click the Edit button.

   The OSPF Interface Parameters Window appears (Figure 11-15).
3. Edit those parameters you wish to change, and click the Save button to save your changes and exit the window.

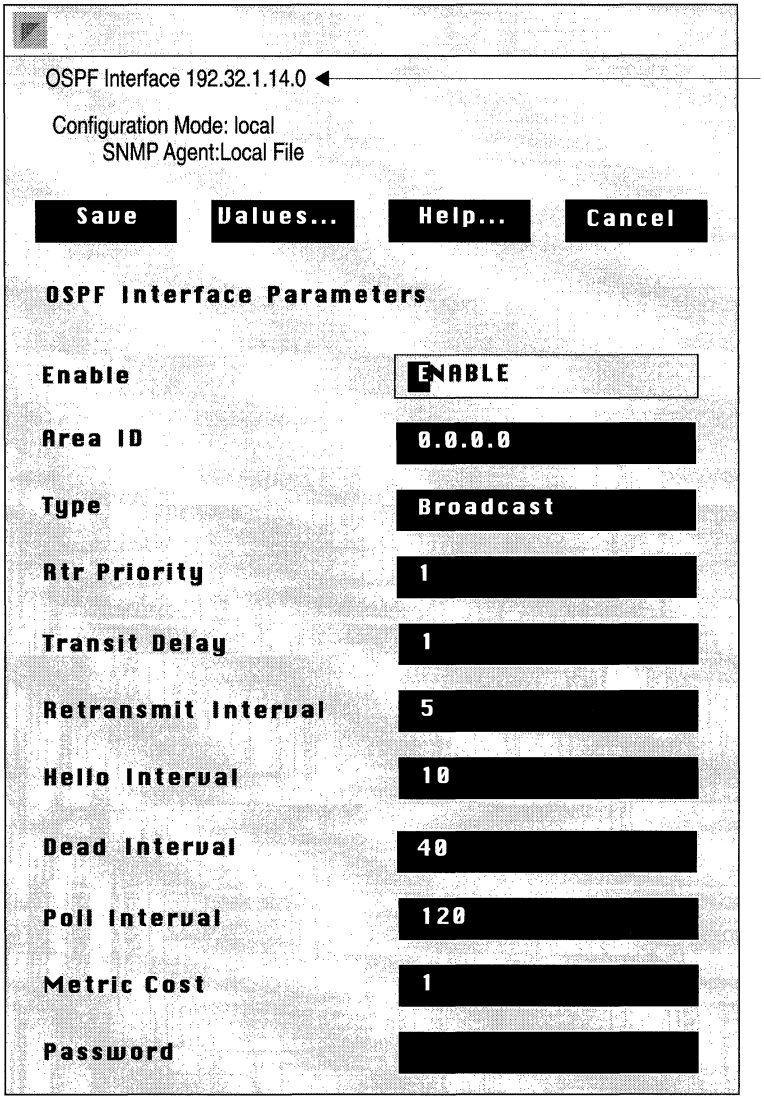


Identifies interface you are editing

**Figure 11-15. OSPF Interface Parameters Window**

**Note:** Except for when you dynamically change the Transit Delay, Hello Interval, Retransmission Interval, or Dead Interval timers, when you reconfigure an interface in dynamic mode, OSPF restarts on all interface.

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | This parameter indicates whether or not OSPF is enabled on this interface. The default value Enable indicates that neighbor relationships may be formed on this interface, and that this interface will be advertised as an internal route to some area. The value Disable indicates that this is not an OSPF interface. |
| Instructions: | Set this value to Disable if you do not want OSPF enabled on the interface. Or, set it to Enable if you previously disabled OSPF on this interface, and now wish to reenable it. |

| **Parameter :** | **Area ID** |
|---|---|
| Wellfleet Default: | 0.0.0.0 |
| Options: | Any four octet number in dotted decimal notation |
| Function: | This parameter identifies the area to which this interface is belongs. |
| Instructions: | Enter the appropriate area ID in dotted decimal notation. |

**Note:** Note that area ID 0.0.0.0 is used only for the OSPF backbone.

| | |
|---|---|
| **Parameter :** | **Type** |
| Wellfleet Default: | Broadcast |
| Options: | Broadcast, NBMA (Non-broadcast multi-access), or Point-to-point |
| Function: | Indicates this interface's type (the type of network to which it is attached). Set this parameter to Broadcast if this network is a broadcast LAN, such as Ethernet. Set it to NBMA for an X.25 or similar type of interface. Or, set it to Point-to-point for a synchronous point-to-point interface.<br><br>Note that if Type is set to NBMA, you will need to configure neighbors manually. |
| Instructions: | Set this parameter to match this interface type. |

| | |
|---|---|
| **Parameter :** | **Rtr Priority** |
| Wellfleet Default: | 1 |
| Options: | 0 to 255 |
| Function: | Indicates the priority of this interface. The Router Priority value is used in multi-access networks (Type is set to Broadcast or NBMA), for the election of the designated router. If this parameter is set to 0, this router is not eligible to become the designated router on this particular network.<br><br>In the case of equal Router Priority values, the router ID will determine which router will become designated router. However, if there already is a designated router on the network when you boot up, it will remain the designated router no matter what your priority or router ID. |
| Instructions: | Set the Router Priority to a value between 0 and 255, or accept the default value, 1. |

| | |
|---|---|
| **Parameter :** | **Transit Delay** |
| Wellfleet Default: | 1 second |
| Options: | 1 to 360 seconds |
| Function: | Indicates the estimated number of seconds it takes to route a packet over this interface. |
| Instructions: | Either accept the default value of 1 second, or enter some slightly higher number for slower speed serial lines, for example, 15 to 20 seconds for a 19.8K line. |

**Parameter :** **Retransmit Interval**

Wellfleet Default: 5 seconds

Options: 1 to 360 seconds

Function: Indicates the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. The Retransmit Interval value is also used when retransmitting OSPF packets. Although the default value is 5, Wellfleet suggests the following values for Retransmission Interval.

| Network Type | Suggested Retransmit Interval |
|---|---|
| Broadcast | 5 seconds |
| Point-to-point | 10 seconds |
| NBMA | 10 seconds |

Instructions: Either accept the default value of 5 seconds, or set the Retransmit Interval to some slightly higher number for slower speed serial lines.

**Parameter :** **Hello Interval**

Wellfleet Default: 10 seconds

Options: 1 to 360 seconds

Function: Indicates the number of seconds between the Hello Packets that the router sends on the interface. Although the default value is 10 seconds, Wellfleet suggests the following values for Hello Interval.

| Network Type | Suggested Hello Interval |
|---|---|
| Broadcast | 10 seconds |
| Point-to-point | 15 seconds |
| NBMA | 20 seconds |

Instructions: Either accept the default value of 10 seconds, or set the Hello Interval to some higher number for slower speed serial lines.

**Note:** The Hello Interval value must be the same for all routers attached to the same network.

**Parameter :** **Dead Interval**

Wellfleet Default: 40 seconds

Options: 1 to 2000 seconds

Function: Indicates the number of seconds that a router's Hello packets have not been seen before it's neighbors declare the router down. The Dead Interval value should be some multiple of the Hello Interval. Although the default value is 40 seconds, Wellfleet suggests the following values for Dead Interval.

| Network Type | Suggested Dead Interval |
|---|---|
| Broadcast | 40 seconds |
| Point-to-point | 60 seconds |
| NBMA | 80 seconds |

Instructions: Either accept the default value of 40 seconds, or set Dead Interval some higher number for slower speed serial lines.

**Note:** The Dead Interval value must be the same for all routers attached to the same network.

| | |
|---|---|
| **Parameter :** | **Poll Interval** |
| Wellfleet Default: | 120 seconds |
| Options: | 1 to 2000 seconds |
| Function: | Indicates the largest number of seconds allowed between Hello packets sent to an inactive non-broadcast multi-access neighbor. |
| Instructions: | Either accept the default value of 120 seconds, or set Poll Interval to some slightly higher number for slower speed serial lines. |

**Parameter:** **Metric Cost**

Wellfleet Default: 1

Options: 1 to 255

Function: Indicates the cost of using this type of service on this interface. Although the default value is 1, Wellfleet suggests the following values for Metric Cost.

| Service Type | Suggested Metric Cost |
|---|---|
| FDDI | 1 |
| Ethernet | 2 |
| HSSI | 2 |
| Token Ring | 2 |
| E1 | 3 |
| T1 | 4 |
| 56K | 16 |

Metric Cost is the parameter that allows you to configure preferred paths. If you do want to configure a preferred path, allow that path to retain the default value of 1, or assign it a relatively low Metric Cost. Then, assign the less preferred paths a higher Metric Cost value.

Instructions: Ether accept the default value 1, or enter a larger number for a slower path or a backup route.

| **Parameter :** | **Password** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any ASCII character string up to 8 characters long. |
| Function: | Specifies the password used for this area. You can specify a password up to eight ASCII characters in length that will appear in the authentication field of all OSPF packets across this interface. Password is valid only when Authentication Type is set to Simplepassword. |
| Instructions: | Enter the appropriate character string. |

**Note:** All routes in the same area must either have no Authentication, or have the same Password.

### Adding a Neighbor to an Interface

In an NBMA network, neighbors are not learned dynamically. For each neighbor on the network, you need to enter it's IP address. To add a neighbor to an NBMA interface, complete the following steps.

1. Select the interface to which you wish to add a neighbor from the OSPF Interfaces scroll box in the OSPF Interface List Window (Figure 11-12).
2. Click the Neighbors button.

   The OSPF Neighbor List Window appears (Figure 11-16).

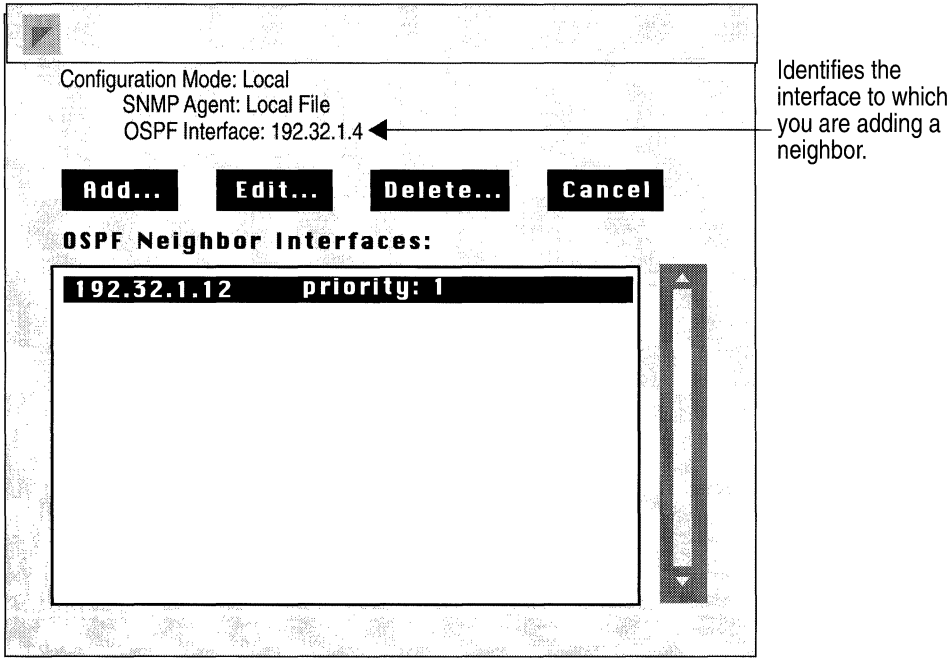


**Figure 11-16. OSPF Neighbor List Window**

3. Click the Add button.

   The Add OSPF Neighbors Window appears (Figure 11-17).

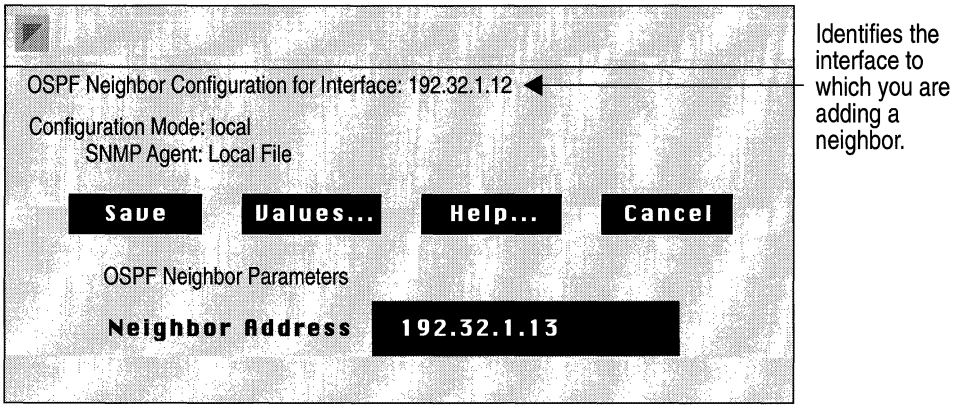


**Figure 11-17. Add OSPF Neighbors Window**

4. Supply the appropriate neighbor address (described below), then click the Save button.

| **Parameter :** | **Neighbor Address** |
|---|---|
| Wellfleet Default: | None |
| Options: | IP address of neighbor |
| Function: | Indicates by IP address a nonbroadcast multi-access neighbor for this interface. |
| Instructions: | Enter the appropriate IP address of the nonbroadcast multi-access neighbor in dotted decimal notation. |

### Editing a Neighbor

Once you have configured the neighbors for an NBMA interface, you can change them. To edit a neighbor, complete the following steps.

1. Select the interface for which you wish to edit a neighbor from the OSPF Interfaces scroll box in the OSPF Interface List Window (Figure 11-12).
2. Click the Neighbors button.

   The OSPF Neighbor List Window appears (Figure 11-18).

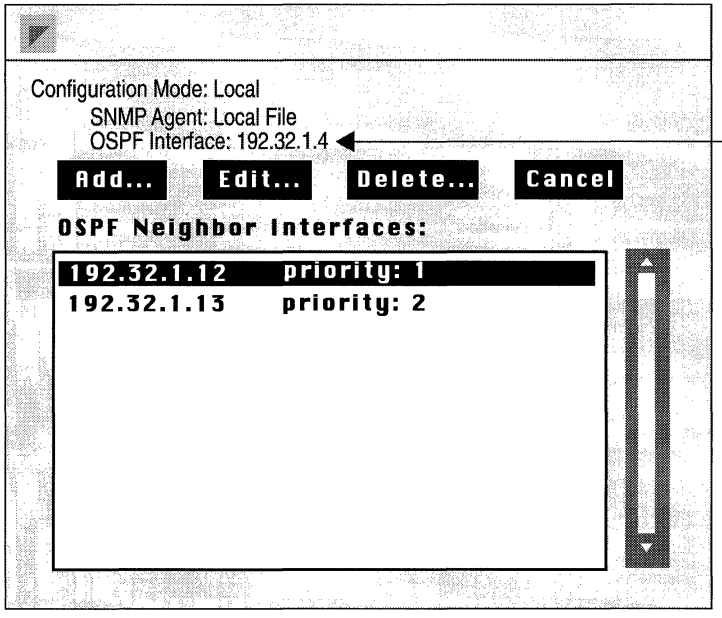


**Figure 11-18. OSPF Neighbor List Window**

3. Click the Edit button.

   The OSPF Neighbor Parameters Window appears (Figure 11-19).

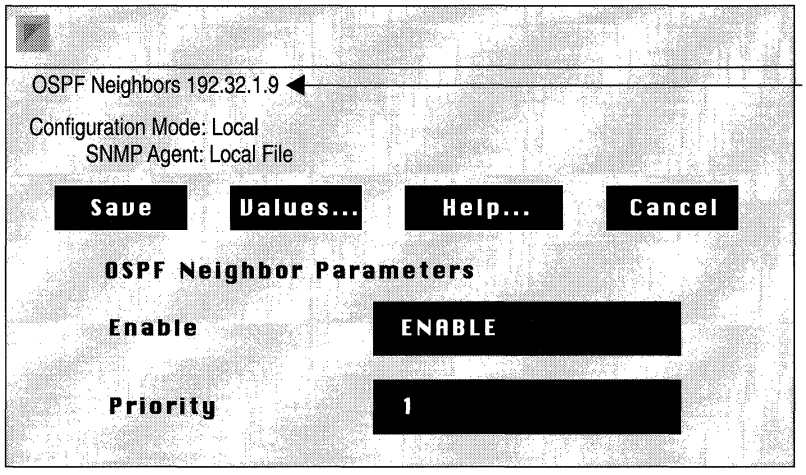


**Figure 11-19. OSPF Neighbor Parameters Window**

4. Set the OSPF neighbor parameters, then click the Save button to save your changes and exit the window.

| **Parameter :** | **Enable** |
|---|---|
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Allows you to enable and disable this neighbor configuration for this interface.This parameter is useful if you want to temporarily disable a neighbor configuration rather than delete it. |
| Instructions: | Set to Disable if you want to disable this neighbor configuration. Or, set to Enable if you previously disabled this neighbor configuration and now wish to reenable it. |

| | |
|---|---|
| **Parameter :** | **Neighbor Priority** |
| Wellfleet Default: | 1 |
| Options: | 0 to 255 |
| Function: | Indicates the priority of this neighbor, with 255 indicating the highest priority. The Neighbor Priority value is used in multi-access networks for the election of the designated router. If this parameter is set to 0, this router is not eligible to become the designated router on this particular network. |
| Instructions: | Either accept the default Neighbor priority value of 1, or enter some other value between 0 and 255. |

### Deleting a Neighbor

To delete a neighbor from an NBMA interface, complete the following steps.

1. Select the interface from which you wish to delete a neighbor from the OSPF Interfaces scroll box in the OSPF Interface List Window (Figure 11-12).
2. Click the Neighbors button.

   The OSPF Neighbor List Window appears (Figure 11-20).

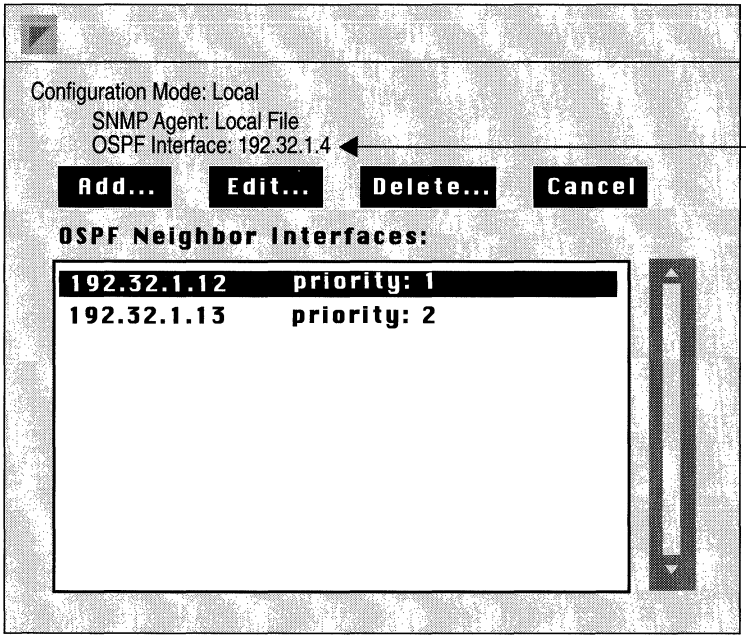


**Figure 11-20. OSPF Neighbor List Window**

3. Select the OSPF neighbor interface you wish to delete from the OSPF Neighbor Interfaces scroll box.
4. Click the Delete button.

   The Delete OSPF Neighbor Window appears (Figure 11-21).

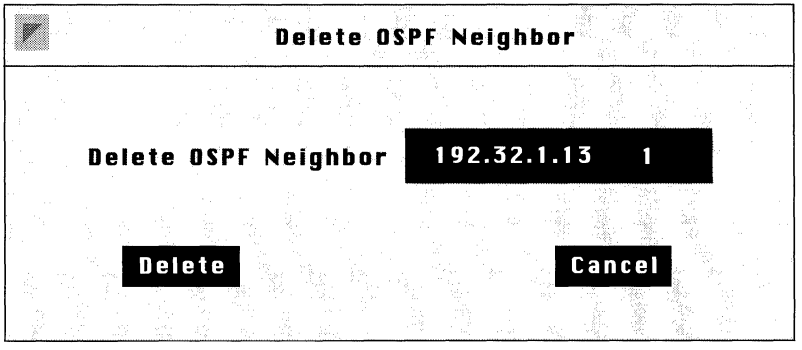


**Figure 11-21. Delete OSPF Neighbors Window**

5. Click the Delete button if the IP address in the Delete OSPF Neighbor box correctly reflects the neighbor you wish to delete.

## Editing OSPF Virtual Link Parameters

To edit OSPF Virtual Link Parameters, begin at the Wellfleet Configuration Manager Window and proceed as follows:

1. Select the Protocols/IP/OSPF/Virtual Interfaces option.
2. The OSPF Virtual Interface List Window appears (Figure 11-22).

   It is from this window that you can perform any of the functions described by the subsections listed below.

   — *Adding a Virtual Interface*

   — *Editing a Virtual Interface*

   — *Deleting a Virtual Interface*

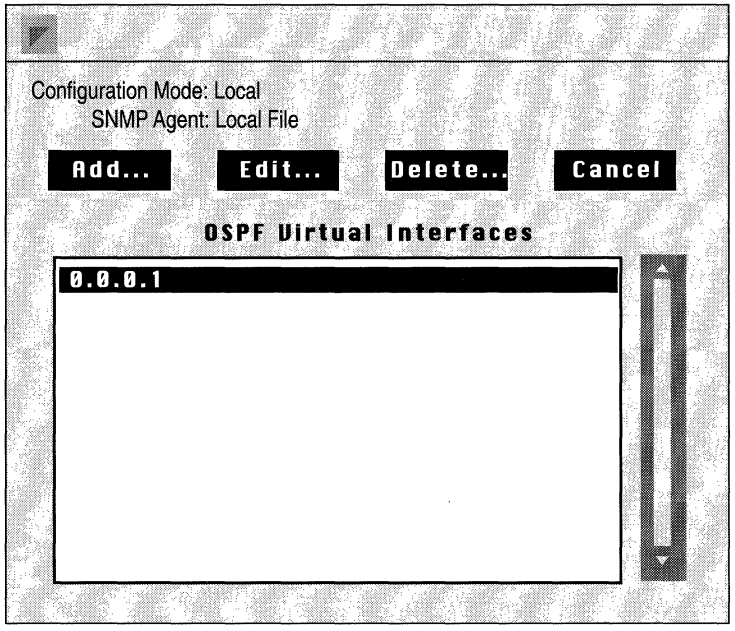


**Figure 11-22. OSPF Virtual Interface List Window**

### Adding a Virtual Interface

To add a virtual interface, complete the following steps:

1. Click the Add button in the OSPF Virtual Interface List Window (Figure 11-22).

   The Add Virtual Interfaces Window appears (Figure 11-23).

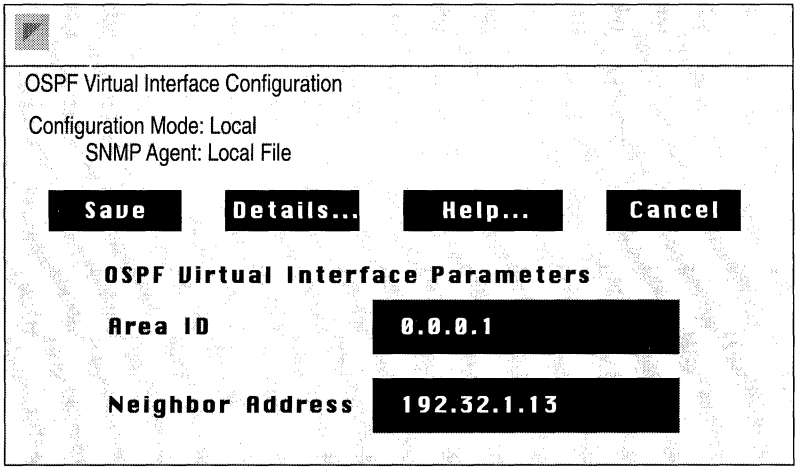


**Figure 11-23. Add Virtual Interfaces Window**

2. Set the appropriate parameters, then click the Save button.

   The Password popup window appears, prompting you for the appropriate password for the transit area.

3. Supply the appropriate password for the transit area, then click the Save button.

| **Parameter :** | **Area ID** |
|---|---|
| Wellfleet Default: | None |
| Options: | Any four octet dotted decimal number |
| Function: | Identifies the transit area through which this virtual link is configured. |
| Instructions: | Enter the appropriate area ID in dotted decimal notation. |

| Parameter : | Neighbor Address |
|---|---|
| Wellfleet Default: | None |
| Options: | Any IP address |
| Function: | Identifies the interface at the other end of this virtual link. |
| Instructions: | Enter the appropriate IP address. |

### Editing a Virtual Interface

When you first configure a virtual interface, it automatically takes on certain parameter defaults. You can; however, change these default values by editing the virtual interface parameters. To edit a virtual interface, complete these steps.

1. Select the virtual interface that you wish to edit from the OSPF Virtual Interfaces scroll box in the OSPF Virtual Interface Window (Figure 11-22).
2. Click the Edit button.

   The OSPF Virtual Interface Parameters Window appears (Figure 11-24).

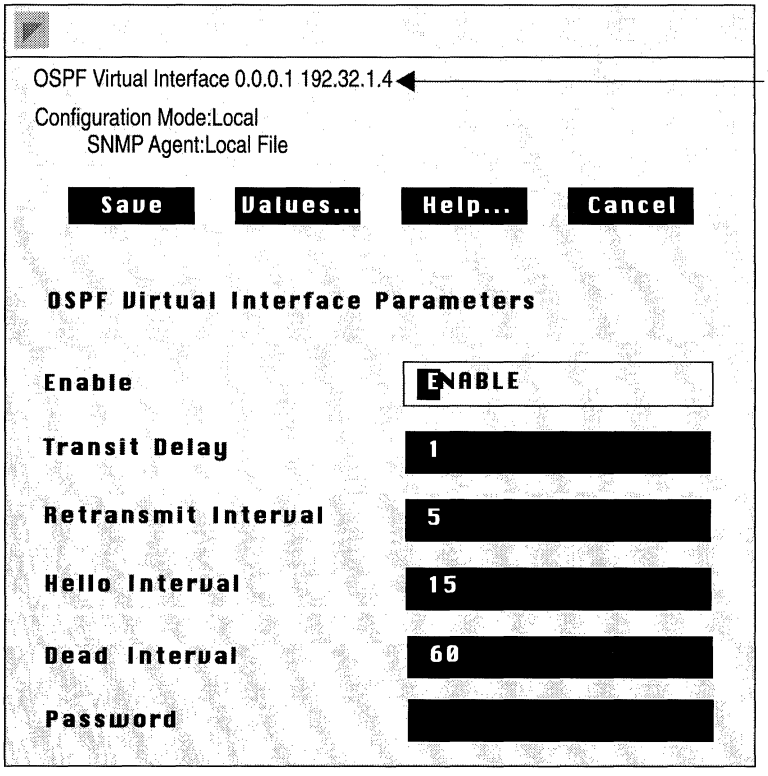


**Figure 11-24. OSPF Virtual Interface Parameters Window**

3. Edit those parameters you wish to change.
4. Click the Save button when you are finished to save your changes and exit the window.

**Note:** Except when you change the Hello Interval, Retransmission Interval, or the Dead Interval timers, when you reconfigure a virtual interface in dynamic mode, OSPF restarts on that interface.

| | |
|---|---|
| **Parameter :** | **Enable** |
| Wellfleet Default: | Enable |
| Options: | Enable/Disable |
| Function: | Enables or disable this virtual link. This parameter is useful when you want to temporarily disable a virtual link rather than delete it. |
| Instructions: | Set to Disable to turn off this virtual link. Or, set to Enable if you previously disabled this virtual link and now wish to reenable it. |

| | |
|---|---|
| **Parameter :** | **Transit Delay** |
| Wellfleet Default: | 1 second |
| Options: | 1 to 360 seconds |
| Function: | Indicates the estimated number of seconds it takes to transmit a link state update packet over this interface. |
| Instructions: | Either accept the default value of 1 second, or enter a new value between 1 and 360 seconds. |

**Parameter :** **Retransmit Interval**

Wellfleet Default: 10 seconds

Options: 1 to 360 seconds

Function: Indicates the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. The Retransmit Interval value is also used when retransmitting database description and link-state request packets. This value should be well over the expected round trip time. Although the default value is 10, Wellfleet suggests the following values for Retransmit Interval.

| Network Type | Suggested Retransmit Interval |
|---|---|
| Broadcast | 10 seconds |
| Point-to-point | 15 seconds |
| NBMA | 15 seconds |

Instructions: Either accept the default value of 10 seconds, or set the Retransmit Interval to some other value between 1 and 360 seconds.

**Parameter :** **Hello Interval**

Wellfleet Default: 15 seconds

Options: 1 to 360 seconds

Function: Indicates the number of seconds between the Hello Packets that the router sends on the interface. Although the default value is 15 seconds, Wellfleet suggests the following values for Hello Interval.

| Network Type | Suggested Hello Interval |
|---|---|
| Broadcast | 10 seconds |
| Point-to-point | 15 seconds |
| NBMA | 20 seconds |

Instructions: Either accept the default value of 15 seconds, or set the Hello Interval to some other value between 1 and 360 seconds.

**Note:** The Hello Interval value must be the same for the virtual neighbor.

**Parameter :** **Dead Interval**

Wellfleet Default: 60 seconds

Options: 1 to 2000 seconds

Function: Indicates the number of seconds that a router's Hello packets have not been seen before it's neighbors declare the router down. The Dead Interval value should be some multiple of the Hello Interval. Although the default value is 60 seconds, Wellfleet suggests the following values for Dead Interval.

| Network Type | Suggested Dead Interval |
|---|---|
| Broadcast | 40 seconds |
| Point-to-point | 60 seconds |
| NBMA | 80 seconds |

Instructions: Either accept the default value of 60 seconds, or enter some other value for Virtual Link Dead Interval.

**Note:** The Dead Interval value must be the same for all routers attached to the same network.

| | |
|---|---|
| **Parameter :** | **Password** |
| Wellfleet Default: | None |
| Options: | Any ASCII character string up to 8 characters long. |
| Function: | Specifies the password used for this area. You can specify a password up to eight ASCII characters in length that will appear in the authentication field of all OSPF packets across this interface. Password is valid only when Authentication Type is set to Simplepassword. |
| Instructions: | Enter the appropriate character string. |

**Note:** All routes in the same area must either have no Authentication, or have the same Password.

### Deleting a Virtual Interface

To delete a virtual interface, complete the following steps.

1. Select the virtual interface that you wish to delete from the OSPF Virtual Interfaces scroll box in the OSPF Virtual Interface Window (Figure 11-22).
2. Click the Delete button.

   The Delete OSPF Virtual Interface Window appears (Figure 11-25).

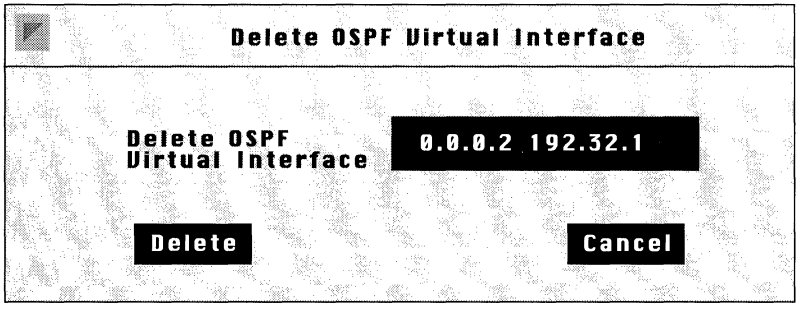


**Figure 11-25. Delete OSPF Virtual Interface Window**

3. Click the Delete button if the area/address in the Delete OSPF Virtual Interface box correctly reflects the virtual interface you wish to delete.

# Index

## A

## B

## C

## D



## E



## F

## G

## H



## I

## L



## M



## N



## O

## P

## Q



## R



## S

## T

## U



## V

## X



## Z