CompatiView 5.4
Reference Guide

Compatible Systems Corporation
4730 Walnut Street
Suite 102
Boulder, Colorado 80301

303-444-9532
800-356-0283
http://www.compatible.com

CompatiView Reference Guide, Version 5.4

Part number: A00-1087

Chapter 1 - Installation and Overview 1

Chapter 2 - IP Routing & Bridging 21

Chapter 3 - IPX Routing & Bridging 57

Chapter 4 - AppleTalk Routing & Bridging 71

Chapter 5 - DECnet Routing & Bridging 91

Chapter 6 - VPN Ports and LAN-to-LAN Tunnels 97

Chapter 7 - VPN Client Tunnels 109

Chapter 8 - IntraGuard Firewall Configuration 131

Chapter 9 - Bridging 149

Chapter 10 - WAN Link Protocols 155

Chapter 11 - TCP/IP Filtering 183

Chapter 12 - IPX Filtering 197

Chapter 13 - AppleTalk Filtering 211

Chapter 14 - General 219

Chapter 15 - OSPF 255

Chapter 16 - BGP 263

Appendices 277

Chapter 1 - Installation and Overview

CompatiView Quickstart

- **Follow** the instructions in the Installation Guide for your internetworking device to connect it to your network.
- **Install** CompatiView by running the install program included on the CD-ROM which was included with your Compatible Systems device.
- **Run** CompatiView.
- **Select** a network transport protocol using the Database menu’s Options dialog box.
- **Add** your device to CompatiView’s device view using the Open menu item under the File menu.
- **Click** on your device in the Device View to open a list of configuration section icons. The default password is “letmein.”
- **Open** configuration dialog boxes by double clicking on the protocol icons under each configuration section icon.
- **Edit** the device’s default configuration using these dialog boxes.
- **Download** your changes to the device using the Save to Device menu item in the File menu.

❖ **Note:** *Parameters and options in this manual which are marked with a* **>** *symbol must be set in order to use the associated device feature.*

❖ **Note:** *If this Quickstart section is a little too quick, don’t worry. This manual completely documents CompatiView. You can use it as a reference to learn more about any of the steps listed above.*

About this Manual

This manual documents CompatiView v5.3, which can be used to configure and manage all Compatible Systems products except the MicroRouter 900i and 1000R and the RISC Router 3000E. CompatiView v4.8x may be used to configure those devices.

CompatiView v4.8x is available in the Network Management\CompatiView\Windows directory on the CD-ROM that was included with your shipping package and in the Software Downloads section of our Web site (http://www.compatible.com).

CompatiView v5.3 is for Windows environments only. An older version of CompatiView which is Macintosh-compatible is available in the Network Management\CompatiView\Macintosh directory on the CD-ROM and on our Web site.

For the latest documentation on Compatible Systems products, including the most current version of this manual, visit the Technical Support section of our Web site.

CompatiView Installation Notes

CompatiView can be installed or updated simply by running the installation program which is located in the Network Management/CompatiView/Windows directory on the CD-ROM. The program will install CompatiView and its associated files on the drive you specify.

System Requirements for Windows

CompatiView for Windows requires a 486 machine or faster, running Microsoft Windows 95 or later, or Windows NT (version 4.0 or later).

❖ **Note:** *Windows 95, Windows 98 and Windows NT are shipped with IP and IPX protocol stacks. See your operating system documentation for instructions on setting up these stacks.*

Selecting IP or IPX Operation with Windows

CompatiView for Windows defaults to using IP as a transport protocol. The IP protocol does not provide a method for CompatiView to automatically discover the device. To initially contact the device over IP using CompatiView, you must first enter a valid IP address into the device. You can do this either on a console directly connected to the device or by setting a workstation’s IP address to 198.41.12.2 with a Class C subnet mask (255.255.255.0) so that it can communicate over Ethernet with 198.41.12.1 (the shipping default of Ethernet A/0 on all devices). After setting the device’s IP address, be sure to change the workstation’s configuration back to its original settings.

To use IPX, which will allow you to contact the device without setting any parameters over the device’s Console port, you can either set the appropriate radio button in the Database menu’s Options dialog box or click on the IP/IPX box at the bottom of the main CompatiView screen. (The status bar must be checked in the View menu for the latter to work.)

CompatiView’s Menus and Main Windows

There are four main menus and three main windows in CompatiView. The File, Database and Control menus are loosely tied to the Device View and Main Windows. The Statistics menu is directly tied to the Output Window. More information on the windows and menus follows.

- The **File** menu’s options are primarily focused on the creation, editing and saving of configuration files and device configuration files. The two types of configuration files are different in that generic configuration files have not been associated with any particular device. These files can be used as templates to speed up the configuration of multiple devices. Device configuration files are files which came directly from a particular device.
- The **Database** menu allows you to create and manage lists of devices. All of the devices on your network can be grouped together for administration in a single Device View, or they can be divided up into smaller groups. This menu also allows you to set CompatiView preferences and device properties.
- The **Control** menu allows you to update device software, do TFTP downloads and restart devices.
- The **Statistics** menu provides in-depth technical information on a device’s operation, including packet statistics and routing table listings as appropriate. Output from the Statistics menu commands will appear in the Output Window’s Command Line Output tab.

CompatiView also provides several other menus.

- The **View** menu, with options for toolbar settings, an on/off setting for the status bar, and an on/off setting for Workbook Mode, which places tabs under the configuration dialog boxes.
- The **Window** menu, which controls the placement of windows and screens and allows you to move between open windows.

- The **Help** menu, which provides standard help functions.

❖ **Note:** *Some of the menu items will be grayed out unless you are currently logged into a device. Where applicable, menu selections are put into effect for the* ***current device****. This is the device which is currently highlighted in the Device View and is shown in the title of the CompatiView screen.*

The Device View and the Main Window

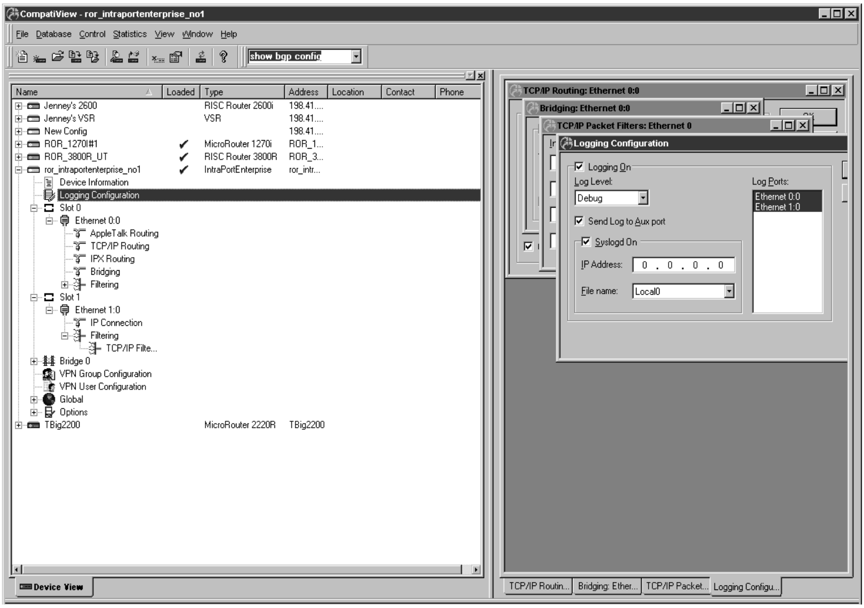
CompatiView - ror_intraportenterprise_no1
File Database Control Statistics View Window Help
show bgp config
Device View
TCP/IP Routing: Ethernet 0:0
Bridging: Ethernet 0:0
TCP/IP Packet Filters: Ethernet 0
Logging Configuration
TCP/IP Routin... Bridging: Ether... TCP/IP Packet... Logging Configu...

Name	Loaded	Type	Address	Location	Contact	Phone
Jenney's 2600		RISC Router 2600i	198.41....			
Jenney's VSR		VSR	198.41....			
New Config			198.41....			
ROR_1270I#1	✓	MicroRouter 1270i	ROR_1...			
ROR_3800R_UT	✓	RISC Router 3800R	ROR_3...			
ror_intraportenterprise_no1	✓	IntraPortEnterprise	ror_intr...			
Device Information						
Logging Configuration						
Slot 0						
Ethernet 0:0						
AppleTalk Routing						
TCP/IP Routing						
IPX Routing						
Bridging						
Filtering						
Slot 1						
Ethernet 1:0						
IP Connection						
Filtering						
TCP/IP Filte...						
Bridge 0						
VPN Group Configuration						
VPN User Configuration						
Global						
Options						
TBig2200		MicroRouter 2220R	TBig2200			

☑ Logging On
Log Level: Debug
☑ Send Log to Aux port
☑ Syslogd On
IP Address: 0 . 0 . 0 . 0
File name: Local0
Log Ports: Ethernet 0:0, Ethernet 1:0

The Device View The Main Window

The Device View displays a list of configurations. These configurations may be generic configuration files which are not associated with a particular device, or they may be a specific device's configuration file. The File menu allows you to add both types of configurations to the Device View.

Included in the window are the configuration's name, type, network address, and a checkmark if it has been loaded. Clicking on the + symbol next to a device loads the device's configuration into CompatiView's memory and brings up a list of the device's configuration section icons, such as device information, interfaces, global device settings, and options. Some of these configuration section icons contain a further list of protocol icons.

If the device is a multislot product such as a VSR or IntraPort Enterprise, both the slot number and the interface number are shown, separated by a colon (e.g., Ethernet 0:0 indicates Slot 0, Ethernet 0, while Ethernet 1:0 indicates Slot 1, Ethernet 0).

Administrative information will also be included if it has been set using the Item Properties option under the Database Menu.

The list of configuration items associated with each device is an *edit area*. To view or edit the configuration information for a specific interface and protocol, click on the protocol icon. A configuration dialog box will be opened in the Main Window.

The information in these configuration dialog boxes is used by a device's operating software to determine how it will interface with wide area communications devices, communicate on IPX subnets, filter network packets, etc.

If you determine that a device needs to use new or different configuration information, you must change the configuration file which is stored in its Flash ROM. (See the File Menu section for more information on downloading a set of configuration parameters to a device.)

If you have made changes to a configuration and then quit CompatiView without downloading those changes, they will be lost.

If the parameters in an edit area are different from the configuration which is currently in the device (because of changes you have made in the edit area), the protocol, interface and device labels in the Device View will be red.

❖ **Note:** *Compatible Systems devices are designed to require less configuration than other devices. Whenever possible, auto-configuration is used to preset parameters with working values.*

Right-Clicking in the Device View

Right-clicking when the mouse is on any item within the Device View will bring up a menu which allows you to add or delete subinterfaces and VPN ports, restart or delete the selected device, or set administrative properties, including how the device will handle Save commands (see the Save/Restart Tab under the Database Menu for more information). The other options are also available as menu items and are documented in detail under the appropriate menu section.

The File Menu

The File menu provides options which allow you to create and manage configurations in CompatiView's Device View.

New Config

This option loads default parameters for a particular type of device in the Device View. You will first be asked to select a device type from a list. This option may be useful to preconfigure a device or to use as a base configuration for multiple devices.

You can edit and view the parameter information by double-clicking on the protocol icons under each configuration section icon. This window will immediately reflect any values you change in the edit area.

> Open - Device

This option provides a way to load a device's configuration into CompatiView's Device View.

The exact method of adding a device depends on the transport protocol you are using with CompatiView.

- If you are using the **IPX** transport stack, this menu item will open a list of all the Compatible Systems devices on your network. Items which are not already entered in CompatiView's Device View are marked with an * in front of the device name.
- If you are using the **IP** transport stack, this menu item will open a window in which you can enter the IP address or domain name of a device.

Open - Config File

This option loads a previously saved configuration file from disk. This will open a browser to allow you to select a configuration file.

> Save to - Device

This option allows you to download the changes you have made to a configuration from CompatiView to a device's Flash ROM. Enter the IP address or a DNS (Domain Name Service) Name for the device to download a configuration to.

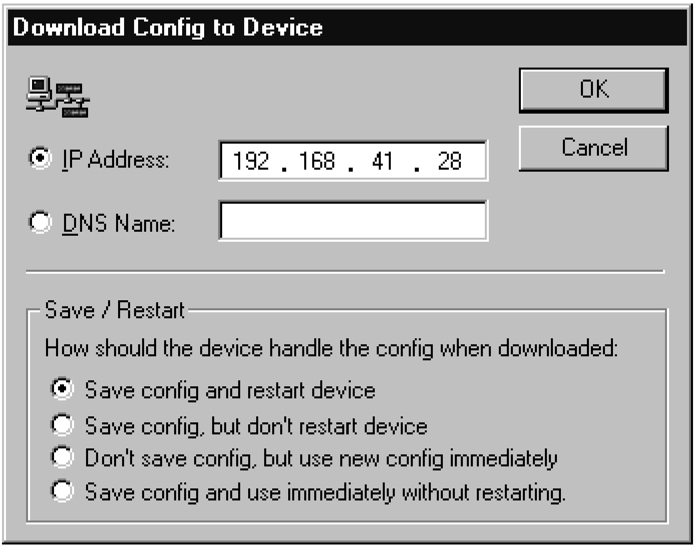

Download Config to Device

● IP Address: 192 . 168 . 41 . 28

○ DNS Name:

Save / Restart

How should the device handle the config when downloaded:

● Save config and restart device

○ Save config, but don't restart device

○ Don't save config, but use new config immediately

○ Save config and use immediately without restarting.

Download Config to Device Dialog Box

Save / Restart Options

The settings in this dialog box are specific for this device. For global Save/Restart settings use the Database menu, select options, and choose the Save/Restart tab. To change the Save/Restart mode for a particular device, modify the “Device Properties” for that device.

- **Save config and restart device.** This parameter will save an edited configuration to the device’s Flash ROM and restart the device to apply the changes. This is the equivalent of the command line’s **save** command.

- **Save config, but don’t restart device.** This parameter will save an edited configuration without restarting the device. The changes will not be applied until the device is restarted. This is the equivalent of the command line’s **write** command.

- **Don’t save config, but use new config immediately.** This parameter will apply an edited (but not saved) configuration to the device’s current operations. If a restart occurs, changes will be lost. This is the equivalent of the command line’s **apply edited** command.

- **Save config and use immediately without restarting.** This parameter will save an edited configuration and immediately apply it to the device’s current operations without restarting the device. This is the equivalent of issuing the **apply** command and then the **write** command in the command line.

While the download is taking place, arrows will move in a circular motion around the device icon in the Device View. To display the amount of time left for the download, click on the + sign next to the device icon.

💣 **Caution:** *Turning off a device in the middle of a download may cause it to **lose its operating software**. Please wait at least **5 minutes** before deciding that a download has failed to be stored in Flash ROM.*

Save To - File

This option saves a configuration as a text file. Use this option to back up the configurations you have downloaded to the devices on your network. When you select this item, you will be asked to enter a file name. The edit area which is exported will correspond to the current configuration.

❖ **Note:** *Configuration text files are useful to Compatible Systems technical support when diagnosing network problems. It is generally a good idea to keep a full set of backup copies of your device configurations in case one of your devices develops a hardware fault and must be replaced. It is not recommended that a text file be used to edit the configuration, since there is no syntax checker and even small mistakes can create configuration errors.*

If any changes are made to a configuration text file while CompatiView has the configuration loaded, CompatiView will ask whether you wish to reload the text file or keep CompatiView's version. If you keep CompatiView's version, any externally made changes will be lost.

Subinterface

This option allows you to add or delete an IP subinterface to one of the device's current interfaces. **Add** opens a dialog box which allows you to specify a port and the subinterface number to create. **Delete** opens a confirmation prompt to delete the subinterface. You must have a subinterface selected to enable the Delete option.

VPN Port

This option allows you to add or delete VPN ports for the device. **Add** opens a dialog box which allows you to specify the VPN port number to create. **Delete** opens a confirmation prompt to delete the port. You must have a VPN port selected to enable the Delete option.

Firewall Path

This option allows you to add or delete firewall paths for an IntraGuard Firewall. **Add** opens a dialog box which allows you to name the firewall path.

Delete opens a confirmation prompt to delete the path. You must have a firewall path selected to enable the Delete option.

View

This menu item brings up the Local Config View tab in the Output Window, which displays the configuration text file for the current device.

Print

This menu item prints the configuration text file for the current device.

Recent File

This menu item holds a list of files that have recently been saved.

> Exit

Exiting takes you out of CompatiView. If you made changes to the information in one or more edit areas (which will now appear in red) and have not saved or downloaded them, you will be given an opportunity to do so.

The Database Menu

New Device Database

This option allows you to create configuration database files. If no other database files have been created, CompatiView automatically saves a database file, "MASTER.INI," every time you close. When you use this option, an empty configuration database will be created to which you may add new devices and configurations.

Open Device Database

This option allows you to open existing configuration database files. When you use this option, a list of files will be opened. Select a file from the list, or browse through the files to find the one you want.

Delete Device

Use this menu option to delete a configuration from CompatiView's Device View.

First, mark the configuration in the list you wish to delete by clicking on it. When you select the Delete menu option, you will be asked whether you wish to remove the configuration from the Device View.

Device Properties

Use this menu option to add administrative information for a particular device. You can enter a device's physical location, a contact name for the device, and a phone number for the contact. This information is maintained in CompatiView and is not downloaded into the device.

Options

This menu item brings up a dialog box which lets you set a variety of options having to do with CompatiView's operation.

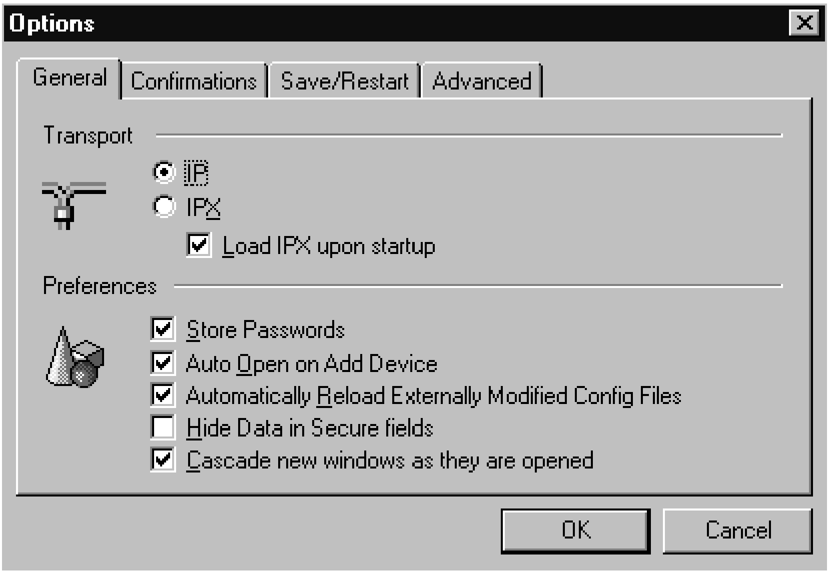

Options

● IP
○ IPX
☑ Load IPX upon startup

☑ Store Passwords
☑ Auto Open on Add Device
☑ Automatically Reload Externally Modified Config Files
□ Hide Data in Secure fields
☑ Cascade new windows as they are opened

Database Options Dialog Box

General Tab

- **IPX Transport - IP Transport.** This set of radio buttons determines whether CompatiView for Windows will use IPX or IP as a transport.
- **Load IPX upon startup.** CompatiView runs IPX behind the scenes to generate IPX tables. I f you do not have IPX on your system, you may want to leave this box unchecked so that CompatiView will not load IPX upon startup.
- **Store Passwords.** This checkbox controls whether CompatiView saves device passwords in its Device View. If you store passwords, you will not need to enter them each time you log into a device.
- **Auto Open on Add Device.** This checkbox controls whether a device configuration will be opened when it is added to the Device View.
- **Automatically Reload Externally Modified Config Files.** If this box is checked, all changes made to the configuration files will automatically

be loaded to the file on disk. If left unchecked, you will be prompted each time the config files are changed and not loaded to disk.

- **Hide Data in Secure fields.** This checkbox will not show passwords in display dialogs or edit boxes, or the text configuration of the current device in Local Config View at the bottom of the screen. If this box is not checked, passwords will be displayed in the clear.
- **Cascade new windows as they are opened.** This checkbox specifies how the dialog boxes in the Main Window are displayed.

Confirmations Tab

- **Confirm before deleting devices from the database.** This checkbox controls whether a confirmation prompt will appear before a device is deleted from the Device View.
- **Confirm before deleting subinterfaces.** This checkbox controls whether a confirmation prompt will appear before an IP subinterface is deleted.
- **Confirm before deleting VPN Ports.** This checkbox controls whether a confirmation prompt will appear before a VPN port is deleted.
- **Confirm before deleting Firewall Paths.** This checkbox controls whether a confirmation prompt will appear before a firewall path is deleted.
- **Confirm configuration download.** This checkbox controls whether a confirmation prompt will appear before a configuration is downloaded to a device.
- **Confirm before restarting devices.** This checkbox controls whether a confirmation prompt will appear before a device is restarted.
- **Confirm before resetting device statistics.** This checkbox controls whether a confirmation prompt will appear before resetting device statistics.

Save/Restart Tab

❖ **Note:** *These selections are global and only sets the "default" for a device* ***when it is added to the database****. They do not change the mode for a device. To change the Save/Restart mode for a particular device, modify the "Device Properties" for that device.*

- **Save config and restart device.** This parameter will save an edited configuration to the device's Flash ROM and restart the device to apply the changes. This is the equivalent of the command line's **save** command.

- **Save config, but don’t restart device.** This parameter will save an edited configuration without restarting the device. The changes will not be applied until the device is restarted. This is the equivalent of the command line’s **write** command.
- **Don’t save config, but use new config immediately.** This parameter will apply an edited (but not saved) configuration to the device’s current operations. If a restart occurs, changes will be lost. This is the equivalent of the command line’s **apply edited** command.
- **Save config and use immediately without restarting.** This parameter will save an edited configuration and immediately apply it to the device’s current operations without restarting the device. This is the equivalent of issuing the **apply** command and then the **write** command in the command line.

❖ **Note:** *Some of these options are not yet available for all Compatible Systems products. To find out whether your device supports them, you must right-click on any configuration item for that device in the Device View and select Properties from the popup menu, then click on the Save/Restart tab.*

Advanced Tab

- **Packet Retry Interval.** This parameter determines how long CompatiView will wait for a response from a device before resending a packet. The default value is 10 seconds.
- **Maximum Connection Timeout.** This parameter determines how long CompatiView will continue retrying before giving up. The default value is 40 seconds.
- **SAP Update Interval.** This parameter determines how frequently CompatiView will retrieve SAP packets. When IPX is in use, lowering this number may make devices appear more quickly when adding new devices to the Device View. The default value is 20 seconds.

❖ **Note:** *The default value of 40 seconds for the Maximum Connection Timeout is long enough to bring up a modem-based dial-on-demand link.*

The Control Menu

The Control menu is primarily concerned with operations on physical devices.

Compatible Systems products use Flash ROM technology to store their operating software and configuration parameters. Flash Rooms can be rewritten tens of thousands of times and will maintain the information which has been written in them regardless of whether they are powered on or not.

The Control menu lets you update the software contained in the Flash ROM of a device.

Download Software

When new features are added to the operating software for a particular type of device, you may wish to update a device with the new version.

When you are using **IPX** transport protocols and select this option, a window listing all eligible devices will appear. You will first be asked to select one or more devices (which must all be of the same type). To select multiple devices, hold down the Control key on your keyboard while clicking on devices.

When you are using **IP** transport protocols and select this option, you will be asked to enter an IP address (the IP address of the current device will be provided as a hint when the window opens).

Once you select one or more devices, CompatiView will log in to the first device in the list (requesting a password from you if it isn’t stored in CompatiView), and then will ask you to select a download file from disk. This file will be downloaded into Flash ROM in the device(s).

Although the old software stored in Flash ROM will be overwritten, the device will maintain any configuration information (addresses, device name, password, etc.) you had previously loaded.

❖ **Note:** *Whenever the Flash ROM in a device is downloaded, whether with new software or with a new configuration, the device will automatically be restarted. The download/restart process will take from 1 to 2 minutes, depending on the amount of memory in the device.*

TFTP Download

This menu option allows you to use the Trivial File Transfer Protocol (TFTP) to download software to a device. This feature is generally only useful if you have erased the operating software in a device’s Flash ROM and are attempting to reload it.

When you select the option, you will be asked for an IP address. CompatiView will then provide a file dialog to allow you to choose the download file.

❖ **Note:** *TFTP can also be used to download operating software into a device which is running standard software from Flash ROM.*

Restart Device

Use this menu option to restart a device in CompatiView’s Device View.

Mark the device in the list you wish to restart by clicking on it. The device you select will be restarted after you select this menu item.

The Output Window

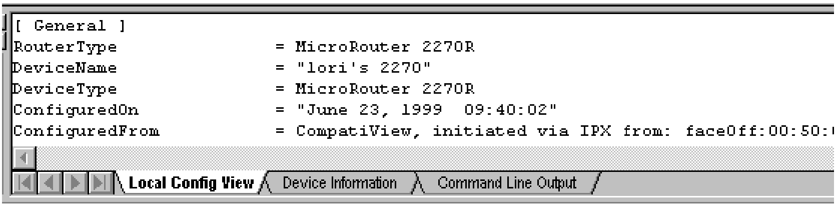

```
[ General ]
RouterType              = MicroRouter 2270R
DeviceName              = "lori's 2270"
DeviceType              = MicroRouter 2270R
ConfiguredOn            = "June 23, 1999  09:40:02"
ConfiguredFrom          = CompatiView, initiated via IPX from: faceOff:00:50:
```

The CompatiView Output Window

There is an Output Window at the bottom of the Device View which lets you quickly check the current status of the selected configuration parameter or review the device configuration. The tabs show different types of parameter values. In some cases, these parameters may be different than those stored in the device’s Flash ROM due to auto-configuration.

The Output Window is broken up into three tabbed sections.

- The **Local Config View** tab displays the complete device configuration and will reflect any changes you have made in the edit area for a device.
- The **Device Information** tab displays the hardware configuration of the device.
- The **Command Line Output** tab is where output from the Statistics menu options will appear. This tab also displays information currently in effect on the device.

The Statistics Menu

This menu allows you to display protocol routing tables and other information for a device. The output from these options is displayed in the Command Line Output tab in the Output Window. The specific menu options available depend on the current device type.

The first set of menu items displays the same information that is available when using certain commands within the command line interface. Refer to the section in the *Text-Based Configuration and Command Line Reference Guide* as indicated for a detailed description of the output from these menu items.

❖ **Note:** *If you are experienced with internetworking devices, the information in these windows will be familiar to you. If you are not, this information can* *be used by Compatible Systems technical support to determine the cause of many problems.*

Ethernet

This menu item displays ethernet port statistics and is the equivalent of the command line’s **show ethernet statistics** command. (See the **ethernet(show)** section.)

WAN State

This menu item displays WAN port status and connection statistics and is the equivalent of the command line’s **show wan state** command. (See the **wan(show)** section.)

Serial Statistics

This menu item displays packet and physical layer statistics for the WAN ports and is the equivalent of the command line’s **show wan serial statistics** command. (See the **wan(show)** section.)

RADIUS

This menu item displays packet statistics for the RADIUS client and is the equivalent of the command line’s **show radius statistics** command. (See the **radius(show)** section.)

PPP Statistics

This menu item displays packet statistics for WAN interfaces set for PPP and is the equivalent of the command line’s **show ppp statistics** command. (See the **ppp(show)** section.)

Frame Relay Statistics

This menu item displays packet statistics for WAN interfaces set for Frame Relay and is the equivalent of the command line’s **show frelay statistics** command. (See the **frelay(show)** section.)

Frame Relay State

This menu item displays the status of the PVCs (Permanent Virtual Circuits) on WAN interfaces set for Frame Relay and is the equivalent of the command line’s **show frelay pvc** command. (See the **frelay(show)** section.)

ARP Cache

This menu item displays the ARP cache, which is the mapping between high level protocol addresses and physical addresses. This command is the equivalent of the command line’s **show arp** command. (See the **arp(show)** section.)

IP Route Table

This menu item displays the IP route table and is the equivalent of the command line's **show ip routing** command. (See the **ip(show)** section.)

IP Routing

This menu item displays IP statistics and is the equivalent of the command line's **show ip statistics** command. (See the **ip(show)** section.)

IPX Route Table

This menu item displays the IPX route table, and is the equivalent of the command line's **show ipx routing** command. (See the **ipx(show)** section.)

IPX SAP Table

This menu item displays the IPX server table, and is the equivalent of the command line's **show ipx sap** command. (See the **ipx(show)** section.)

AppleTalk Route Table

This menu item displays the AppleTalk route table and is the equivalent of the command line's **show appletalk routing** command. (See the **appletalk(show)** section.)

AppleTalk Routing

This menu item displays AppleTalk statistics and is the equivalent of the command line's **show appletalk statistics** command. (See the **appletalk(show)** section.)

OSPF Configuration

This menu item displays user-configured values that are currently being used by the OSPF protocol and is the equivalent of the command line's **show ospf config** command. (See the **ospf(show)** section).

OSPF Packet Statistics

This menu item displays how many of each of the five types of OSPF packets (Hello, Database Description, Link State Request, Link State Update, and Link State Acknowledgement) have been received and sent. This is the equivalent of the command line's **show ospf stats** command. (See the **ospf(show)** section).

OSPF Interface Database

This menu item displays the OSPF interface database and is the equivalent of the command line's **show ospf if** command. (See the **ospf(show)** section).

OSPF Neighbors

This menu item displays an abbreviated list of current neighbors an their state. This is equivalent to the command line’s **show ospf nbr** command. (See the **ospf(show)** section)

Buffer

This menu item displays detailed information on the current status of the device’s memory allocation and is the equivalent of the command line’s **show os memory** command. (See the **os(show)** section.)

Show Restart Info

This menu item displays detailed information about the status of the device when the last restart event occurred, and is the equivalent of the command line’s **show os resevent** command. (See the **os(show)** section.)

Device Log

This menu item displays the log buffer, and is the equivalent of the command line’s **show system log buffer** command. (See the **system(show)** section.)

Command Line Interface

This menu item allows you to enter other **show** commands in the Command Line entry box, as described below.

Reset Statistics

This menu item sends a command to the current device which causes it to reset all of its statistic counters.

The Command Line Edit Box

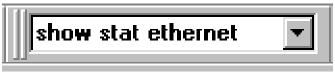

show stat ethernet

The Command Line Edit Box

This box is both a pull-down list and an edit box which allows you to enter command line **show** commands. Any Statistics menu item you use will be added to this pull-down menu. To enter other **show** commands which are not included in the Statistics menu, choose the Statistics menu’s Command Line Interface option to enter the command in the edit box. Press the Return key to send the command to the device.

❖ **Note:** *Other types of commands (e.g.,* ***reset****,* ***add****, etc.) are not fully supported by CompatiView. Only* ***show*** *commands should be used.*

Moving and Customizing the Windows

Right-clicking in the area between windows brings up a popup menu which controls the placement of the windows.

- **Allow Docking.** This menu option, when checked, allows the window to be docked in a firm place within the main window.
- **Hide.** This menu option will hide the selected window. Use the **Window** menu to view a hidden window again.

Clicking and dragging the double bars at the top or side of a window allows you to move the window around on the screen, according to the options described above. Pressing the Control key as you click and drag will disable docking, and the window can be placed anywhere on the screen, including outside the Main window.

The View menu

Use this menu option to view your display in full screen or in workbook mode. You can also change the size of the window or move the window around the screen by clicking and dragging the double bars at the top of the window.

Customize

To customize the display windows, select Customize in the View menu. This dialog box gives options for customizing the toolbars and command icons.

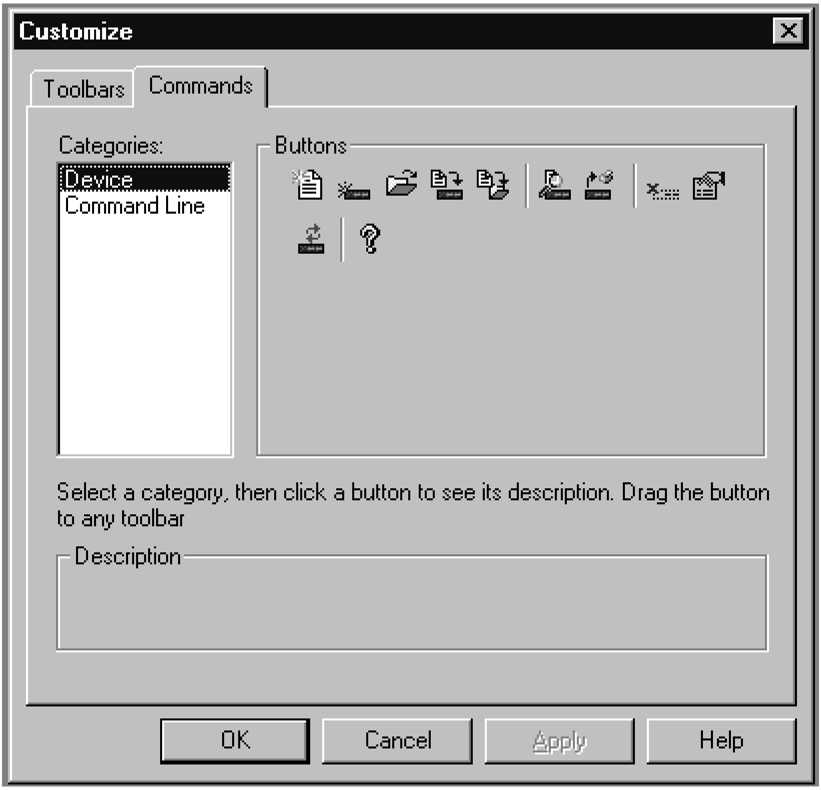

Categories:

Device

Command Line

Select a category, then click a button to see its description. Drag the button to any toolbar

Description

Customize Window View Dialog Box

Toolbars

This tab allows you to choose the toolbars that you want in your display window.

Commands

This tab allows you to create your own toolbar by placing device commands or command line buttons onto any toolbar.

The Window Menu

This menu allows you to toggle the database workspace (device view) and the output window. You can also choose how your windows will be displayed in the workspace.

Chapter 2 - IP Routing & Bridging

TCP/IP Routing: Ethernet Dialog Box

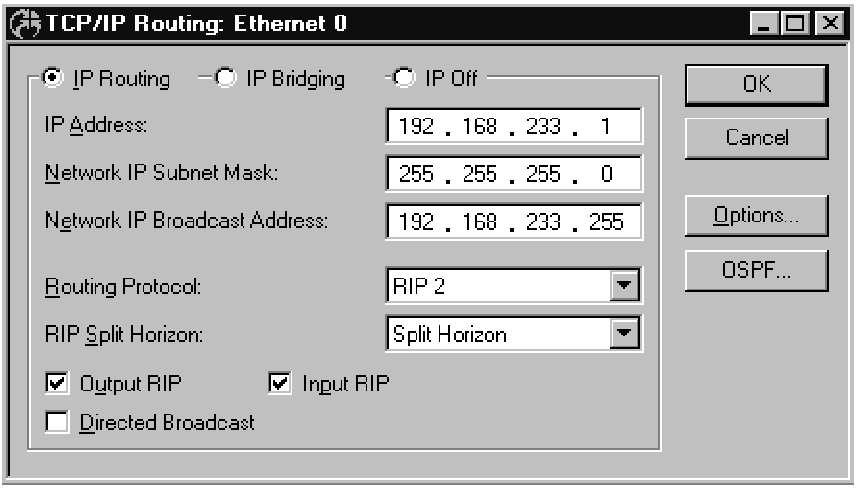

TCP/IP Routing: Ethernet 0

● IP Routing ○ IP Bridging ○ IP Off

IP Address: 192 . 168 . 233 . 1

Network IP Subnet Mask: 255 . 255 . 255 . 0

Network IP Broadcast Address: 192 . 168 . 233 . 255

Routing Protocol: RIP 2

RIP Split Horizon: Split Horizon

☑ Output RIP ☑ Input RIP

□ Directed Broadcast

TCP/IP Routing: Ethernet Configuration Dialog Box

❖ **Note:** *If you need more information about the IP protocol, see “IP 101” in the Appendices to this manual.*

To access this dialog box, select Ethernet/TCP/IP Routing from the Device View.

> IP Routing/Bridging/Off

This set of radio buttons controls how IP packets are handled for this interface.

- If set to **IP Routing**, then IP packets received on this interface are routed to the correct interface on the router.
- If set to **IP Bridging**, then any IP packets received on this interface are forwarded to the router’s internal bridge. This setting makes this Ethernet interface a member of the “IP Bridge Group” for this router.

❖ **Note:** *The IP Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Bridging: Ethernet Dialog Box (under Ethernet/Bridging).*

- If set to **IP Off**, then any IP packets received on this interface are discarded.

> IP Address

Every network interface on an IP internetwork must have a unique IP address that identifies that interface to other devices on the internetwork. Part of this address identifies the network segment the router interface is connected to, and the remainder uniquely identifies the router interface itself.

This address should be entered as four decimal numbers separated by periods -- for example 198.238.9.1

❖ **Note:** *The single most common problem encountered in IP networking is the use of a duplicate IP address. You must carefully track the network numbers you have assigned to various devices in order to avoid hard-to-diagnose problems.*

> Network IP Subnet Mask

Most IP networks use “subnetting” in order to subdivide a large network into smaller logical sub-networks. The subnet mask value is used to tell the router what part of the IP address identifies the network segment (the “network” portion), and what part identifies individual interfaces (the “host” portion).

There are three generally used “classes” of subnetted IP networks: A, B and C. Each class uses a different amount of the IP address for the network and host portions. These classes may also be further divided by correctly setting the subnet mask.

If you do not enter a number in the Subnet Mask field, CompatiView will derive a default value from the IP Address number you entered just above. This default assumes you want a single subnet for all of the available host addresses. You must manually set the field if you want to further divide the address range.

To have CompatiView calculate a default mask, make sure that the Subnet Mask field is empty, position the cursor in the IP Address field, then just tab through the Subnet Mask field.

> Network IP Broadcast Address

The router will use this address to send any IP broadcast messages. The standard broadcast address is all 255’s (hexadecimal FFs) in the host portion of the address. A few networks use all zeroes in this field. If you are unsure which type your network uses, check with your network administrator.

To have CompatiView calculate a default broadcast address, make sure that the Broadcast Address field is empty, position the cursor in the Subnet Mask field, then just tab through the Broadcast Address field.

> Routing Protocol

Routers exchange information about the most effective path for packet transfer between various end points. There are a number of different protocols which have been defined to facilitate the exchange of this information.

Routing Information Protocol (RIP) 1 is the most widely used routing protocol on IP networks. All gateways and routers that support RIP 1 periodically broadcast routing information packets. These RIP 1 packets contain information concerning the networks that the routers and gateways can reach as well as the number of routers/gateways that a packet must travel through to reach the receiving address.

RIP 2 is an enhancement of RIP 1 which allows IP subnet information to be shared among routers, and provides for authentication of routing updates. When this protocol is chosen, the router will use the multicast address 224.0.0.9 to send and/or receive RIP 2 packets for this network interface. As with RIP 1, the router’s routing table will be periodically updated with information received in these packets.

RIP 2 is more useful in a variety of environments and allows the use of variable subnet masks on your network. It is also necessary for implementation of “classless” addressing as accomplished with CIDR (Classless Inter Domain Routing).

It is recommended that RIP 2 be used on any segment where all routers can use the same IP routing protocol. If one or more routers on a segment must use RIP 1, then all other routers on that segment should also be set to use RIP 1.

- If **RIP 2** is selected with this pull-down menu, the router will send and/or accept RIP 2 packets over this interface, and will then periodically update its routing table with the information provided from these packets. On a large network, an up-to-date routing table will enhance network performance since the router will always be aware of the optimal path to use when sending packets.
- If **RIP 1** is selected with this pull-down menu, the router will send and/or accept RIP 1 packets, and will then periodically update its routing table with the information provided from these packets.

- If **None** is selected with this pull-down menu, the router will not be able to update its routing table and will always direct traffic for addresses it does not have a route for (addresses not on one of the networks connected to its interfaces) to the “gateway/port” defined in its IP Static Route Dialog Box. It will then be the responsibility of the default router to direct the packets to the correct address. For information on setting the default router see the discussion of the IP Static Route Dialog Box later in this chapter.

❖ **Note:** *Some routers, in particular those designed to create very large corporate backbones, may use other routing protocols such as OSPF (Open Shortest Path First). These routers can simultaneously use RIP 1 (and in some cases RIP 2) to communicate with smaller routers, or each of the smaller routers can be set to use one of these backbone routers as their default router.*

RIP Split Horizon

Normally, RIP uses a technique called split horizon to avoid routing loops and allow smaller update packets. This technique specifies that when the router sends a RIP update out a particular network interface, it should never include routing information acquired over that same interface.

There is a variation of the split horizon technique called “poison reverse” which specifies that all routes should be included in an update out a particular interface, but that the metric should be set to infinity for those routes acquired over that interface. One drawback is that routing update packet sizes will be increased when using poison reverse.

- If **Split Horizon** is selected with this pull-down menu, the router will apply the split horizon technique to routes being output over this interface.
- If **No Split Horizon** is selected with this pull-down menu, the router will include all routes in an output packet, regardless of which interface they were acquired over, and will use a normal metric.
- If **Poison Reverse** is selected with this pull-down menu, the router will include all routes in an output packet, but will set the metric to infinity for those routes which were acquired over this interface.

Output RIP - Input RIP

These flags control the behavior of RIP 1 and RIP 2 for this interface, allowing the router to selectively send RIP, receive RIP, or both. The default (assuming RIP 1 or RIP 2 is turned on in the Routing Protocol popup) is to both send and receive.

Directed Broadcast

This checkbox sets whether the interface will forward network-prefix-directed broadcasts. This is a security feature which can help prevent your network from being used as an intermediary in certain kinds of attacks which use ICMP echo traffic (pings) or UDP echo packets with fake (i.e., “spoofed”) source addresses to inundate a victim with erroneous traffic.

Options

The options button brings up the Ethernet TCP/IP Options Dialog Box which provides access to Proxy ARP, UDP Relays and other configuration information. This dialog box is discussed later in this chapter.

OSPF

This option button brings up the OSPF Dialog Box which allows the OSPF routing protocol to be enabled. For more information on this dialog box and other OSPF parameters, refer to ***Chapter 15 - OSPF***.

TCP/IP Routing: WAN Configuration Dialog Box

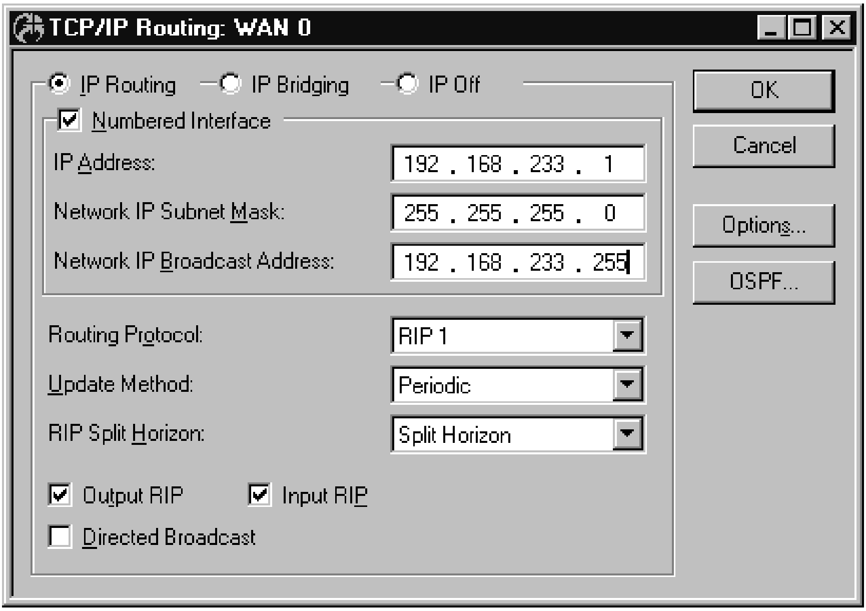

TCP/IP Routing: WAN 0

● IP Routing ○ IP Bridging ○ IP Off

☑ Numbered Interface

IP Address: 192 . 168 . 233 . 1

Network IP Subnet Mask: 255 . 255 . 255 . 0

Network IP Broadcast Address: 192 . 168 . 233 . 255

Routing Protocol: RIP 1

Update Method: Periodic

RIP Split Horizon: Split Horizon

☑ Output RIP ☑ Input RIP

☐ Directed Broadcast

TCP/IP Routing: WAN Configuration Dialog Box

❖ **Note:** *If you need more information about the IP protocol, see “IP 101” in the Appendices to this manual.*

To access this dialog box, select WAN/TCP/IP Routing from the Device View.

> IP Routing/Bridging/Off

This set of radio buttons controls how IP packets are handled for this interface.

- If set to **IP Routing**, then IP packets received on this interface are routed to the correct interface on the router.
- If set to **IP Bridging**, then any IP packets received on this interface are forwarded to the router’s internal bridge. This setting makes this WAN interface a member of the “IP Bridge Group” for this router.

❖ **Note:** *The IP Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Bridging: WAN Dialog Box (under WAN/Bridging).*

- If set to **IP Off**, then any IP packets received on this interface are discarded.

> Numbered Interface

This check box determines whether the Wide Area Network connected to this interface will have an IP network number associated with it.

Many WAN connections are simple point-to-point links. These links do not generally require a network number because there are only two devices on the link. All traffic sent from one end is, by definition, destined for the other end. You generally do not need a numbered WAN interface if you are using the PPP transport protocol.

In contrast, Frame Relay networks may have a number of participating routers connected through a single physical interface. Because of this, use of the Frame Relay transport protocol requires a numbered WAN interface.

- If **checked**, then you must set an IP Address, Subnet Mask, and Broadcast Address (as described below) for this WAN interface. The default is unchecked.

❖ **Note:** *If you are connecting the router to an Internet Service Provider using PPP, you may be required to use a numbered interface. Check with their tech support staff.*

IP Address

Every network interface on an IP internetwork must have a unique IP address that identifies that interface to other devices on the internetwork. Part of this address identifies the network segment the router interface is connected to, and the remainder uniquely identifies the router interface itself.

This address should be entered as four decimal numbers separated by periods -- for example, 198.238.9.5

❖ **Note:** *The single most common problem encountered in IP networking is the use of a duplicate IP address. You must carefully track the network numbers you have assigned to various devices in order to avoid hard-to-diagnose problems.*

Network IP Subnet Mask

Most IP networks use “subnetting” in order to subdivide a large network into smaller logical sub-networks. The subnet mask value is used to tell the router what part of the IP address identifies the network segment (the “network” portion), and what part identifies individual interfaces (the “host” portion).

There are three generally used “classes” of subnetted IP networks: A, B and C. Each class uses a different amount of the IP address for the network and host portions. These classes may also be further divided by correctly setting the subnet mask.

If you do not enter a number in the Subnet Mask field, CompatiView will derive a default value from the IP Address number you entered just above. This default assumes you want a single subnet for all of the available host addresses. You must manually set the field if you want to further divide the address range.

To have CompatiView calculate a default mask, make sure that the Subnet Mask field is empty, position the cursor in the IP Address field, then just tab through the Subnet Mask field.

Network IP Broadcast Address

The router will use this address to send any IP broadcast messages. The standard broadcast address is all 255’s (hexadecimal FFs) in the host portion of the address. A few networks use all zeroes in this field. If you are unsure which type your network uses, check with your network administrator.

To have CompatiView calculate a default broadcast address, make sure that the Broadcast Address field is empty, position the cursor in the Subnet Mask field, then just tab through the Broadcast Address field.

> Routing Protocol

Routers exchange information about the most effective path for packet transfer between various end points. There are a number of different protocols which have been defined to facilitate the exchange of this information.

Routing Information Protocol (RIP) 1 is the most widely used routing protocol on IP networks. All gateways and routers that support RIP 1 periodically broadcast routing information packets. These RIP 1 packets contain information concerning the networks that the routers and gateways can reach as well as the number of routers/gateways that a packet must travel through to reach the receiving address.

RIP 2 is an enhancement of RIP 1 which allows IP subnet information to be shared among routers, and provides for authentication of routing updates. When this protocol is chosen, the router will use the multicast address 224.0.0.9 to send and/or receive RIP 2 packets for this network interface. As with RIP 1, the router’s routing table will be periodically updated with information received in these packets.

RIP 2 is more useful in a variety of environments and allows the use of variable subnet masks on your network. It is also necessary for implementation of "classless" addressing as accomplished with CIDR (Classless Inter Domain Routing).

It is recommended that RIP 2 be used on any segment where all routers can use the same IP routing protocol. If one or more routers on a segment must use RIP 1, then all other routers on that segment should also be set to use RIP 1.

- If **RIP 2** is selected with this pull-down menu, the router will send and/or accept RIP 2 packets over this interface, and will then periodically update its routing table with the information provided from these packets. On a large network, an up-to-date routing table will enhance network performance since the router will always be aware of the optimal path to use when sending packets.
- If **RIP 1** is selected with this pull-down menu, the router will send and/or accept RIP 1 packets, and will then periodically update its routing table with the information provided from these packets.
- If **None** is selected with this pull-down menu, the router will not be able to update its routing table and will always direct traffic for addresses it does not have a route for (addresses not on one of the networks connected to its interfaces) to the "default router" defined in its IP Static Route Dialog Box. It will then be the responsibility of the default router to direct the packets to the correct address. For information on setting the default router see the discussion of the IP Static Route Dialog Box later in this chapter.

❖ **Note:** *Some routers, in particular those designed to create very large corporate backbones, may use other routing protocols such as OSPF (Open Shortest Path First). These routers can simultaneously use RIP 1 (and in some cases RIP 2) to communicate with smaller routers, or each of the smaller routers can be set to use one of these backbone routers as their default router.*

> Update Method

WAN interfaces which are configured to provide "dial-on-demand" service will bring a connection up (i.e. dial the other end) when there are network packets which must be transferred over the link. Once a dial-on-demand connection is up, network traffic passing across the link causes the inactivity timer for the link to be reset, keeping the connection up.

The RIP protocol periodically sends out update information across a link. These periodic update packets will cause a WAN interface set for dial-on-demand operation to stay up indefinitely.

- If **Triggered** is selected with this pull-down menu, the router will modify the standard RIP behavior for this interface to send RIP packets only when there has been an update to its routing table information, or when it has detected a change in the accessibility of the next hop router.
- If **Periodic** is selected with this pull-down menu, the router will use the standard RIP protocol, which sends RIP packets over the link every 30 seconds.

RIP Split Horizon

Normally, RIP uses a technique called split horizon to avoid routing loops and allow smaller update packets. This technique specifies that when the router sends a RIP update out a particular network interface, it should never include routing information acquired over that same interface.

There is a variation of the split horizon technique called “poison reverse” which specifies that all routes should be included in an update out a particular interface, but that the metric should be set to infinity for those routes acquired over that interface. One drawback is that routing update packet sizes will be increased when using poison reverse.

- If **Split Horizon** is selected with this pull-down menu, the router will apply the split horizon technique to routes being output over this interface.
- If **No Split Horizon** is selected with this pull-down menu, the router will include all routes in an output packet, regardless of which interface they were acquired over, and will use a normal metric.
- If **Poison Reverse** is selected with this pull-down menu, the router will include all routes in an output packet, but will set the metric to infinity for those routes which were acquired over this interface.

Output RIP - Input RIP

These flags control the behavior of RIP 1 and RIP 2 for this interface, allowing the router to selectively send RIP, receive RIP, or both. The default (assuming RIP 1 or RIP 2 is turned on in the Routing Protocol popup) is to both send and receive.

Directed Broadcast

This checkbox sets whether the interface will forward network-prefix-directed broadcasts. This is a security feature which can help prevent your network from being used as an intermediary in certain kinds of attacks which use ICMP echo traffic (pings) or UDP echo packets with fake (i.e., “spoofed”) source addresses to inundate a victim with erroneous traffic.

Options

The options button brings up the WAN IP Options Dialog Box which allows you to set a Remote Node IP Address, Van Jacobson Header Compression, and other configuration information. This dialog box is discussed later in this chapter.

OSPF

This option button brings up the OSPF Dialog Box which allows the OSPF routing protocol to be enabled. For more information on this dialog box and other OSPF parameters, refer to ***Chapter 15 - OSPF***.

TCP/IP Routing: VPN Configuration Dialog Box

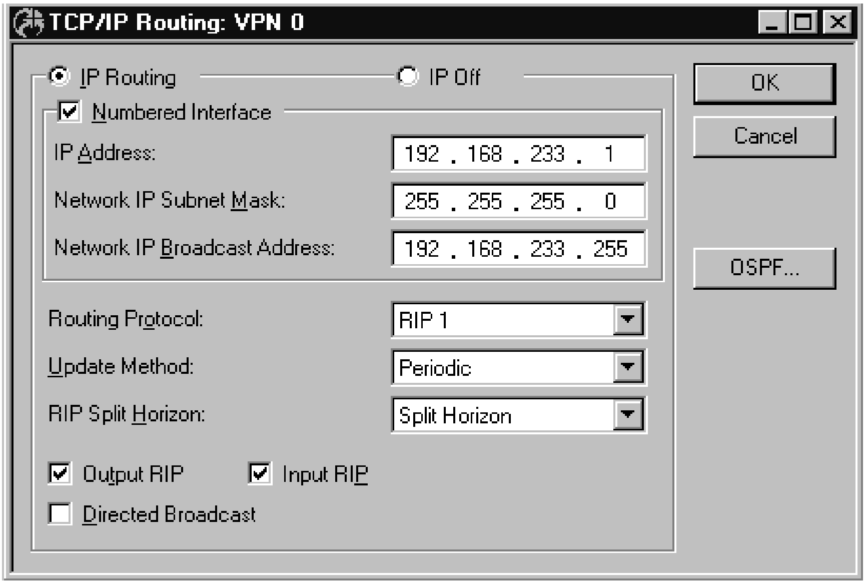

● IP Routing ○ IP Off

☑ Numbered Interface

IP Address: 192 . 168 . 233 . 1

Network IP Subnet Mask: 255 . 255 . 255 . 0

Network IP Broadcast Address: 192 . 168 . 233 . 255

Routing Protocol: RIP 1

Update Method: Periodic

RIP Split Horizon: Split Horizon

☑ Output RIP ☑ Input RIP

□ Directed Broadcast

TCP/IP Routing: VPN Configuration Dialog Box

VPN (Virtual Private Network) ports must first be added to the edit area of a device before they can be configured. For more information about adding and deleting VPN ports, see ***Chapter 6 - VPN Ports and Tunnels***.

Once you have created a VPN port, you may access the TCP/IP Routing: VPN Configuration Dialog Box by clicking TCP/IP Routing under the VPN port's icon.

A VPN port is a virtual port which handles tunneled traffic. Tunnels are virtual point-to-point connections through a public network such as the Internet. All packets sent through a VPN tunnel are IP-encapsulated packets, including AppleTalk, IPX and even IP packets. This encapsulation is added or removed, depending on the direction, by "Tunnel Partner" routers. Once a packet reaches the remote Tunnel Partner, the TCP/IP encapsulation is stripped off, leaving the original protocol. The unencapsulated packet is then handled according to the VPN port's protocol configuration settings. Networks connected via a tunnel will communicate as if they are on the same network, even though they are separated by the Internet.

❖ **Note:** *Remember that you must set up both ends of every tunnel. Therefore, you must repeat this setup with the remote router.*

> IP Routing/IP Bridging/IP Off

This set of radio buttons controls how IP packets are handled for this interface.

- If set to **IP Routing**, then IP packets received on this interface are routed to the correct interface on the device.
- If set to **IP Bridging**, then any IP packets received on this interface are forwarded to the device's internal bridge. This setting makes this VPN port a member of the "IP Bridge Group" for this device.

❖ **Note:** *The IP Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Bridging: VPN Dialog Box (under VPN/Bridging).*

- If set to **IP Off**, then any IP packets received on this interface are discarded.

Numbered Interface

This check box determines whether the VPN port will have an IP network number associated with it.

VPN tunnels are essentially point-to-point links. These links do not generally require a network number because all traffic sent from one end is, by definition, destined for the other end. However, you may wish to assign an address for network tracking purposes.

- If **checked**, then you must set an IP Address, Subnet Mask, and Broadcast Address (as described below) for this VPN port. The default is unchecked.

IP Address

If you wish to assign an IP address, it must be <u>unique</u>. Part of this address identifies the network segment the router interface is connected to, and the remainder uniquely identifies the router interface itself.

This address should be entered as four decimal numbers separated by periods -- for example, 198.238.9.5

❖ **Note:** *The single most common problem encountered in IP networking is the use of a duplicate IP address. You must carefully track the network numbers you have assigned to various devices in order to avoid hard-to-diagnose problems.*

Network IP Subnet Mask

Most IP networks use “subnetting” in order to subdivide a large network into smaller logical sub-networks. The subnet mask value is used to tell the device what part of the IP address identifies the network segment (the “network” portion), and what part identifies individual interfaces (the “host” portion).

There are three generally used “classes” of subnetted IP networks: A, B and C. Each class uses a different amount of the IP address for the network and host portions. These classes may also be further divided by correctly setting the subnet mask.

If you do not enter a number in the Subnet Mask field, CompatiView will derive a default value from the IP Address number you entered just above. This default assumes you want a single subnet for all of the available host addresses. You must manually set the field if you want to further divide the address range.

To have CompatiView calculate a default mask, make sure that the Subnet Mask field is empty, (re)position the cursor in the IP Address field, then just tab through the Subnet Mask field.

Network IP Broadcast Address

The standard broadcast address is all 255’s (hexadecimal FFs) in the host portion of the address. A few networks use all zeroes in this field. If you are unsure which type your network uses, check with your network administrator.

To have CompatiView calculate a default broadcast address, make sure that the Broadcast Address field is empty, (re)position the cursor in the Subnet Mask field, then just tab through the Broadcast Address field.

> Routing Protocol

Routers exchange information about the most effective path for packet transfer between various end points. There are a number of different protocols which have been defined to facilitate the exchange of this information.

Routing Information Protocol (RIP) 1 is the most widely used routing protocol on IP networks. All gateways and routers that support RIP 1 periodically broadcast routing information packets. These RIP 1 packets contain information concerning the networks that the routers and gateways can reach as well as the number of routers/gateways that a packet must travel through to reach the receiving address.

RIP 2 is an enhancement of RIP 1 which allows IP subnet information to be shared among routers, and provides for authentication of routing updates. When this protocol is chosen, the router will use the multicast address 224.0.0.9 to send and/or receive RIP 2 packets for this network interface. As with RIP 1, the router’s routing table will be periodically updated with information received in these packets.

RIP 2 is more useful in a variety of environments and allows the use of variable subnet masks on your network. It is also necessary for implementation of “classless” addressing as accomplished with CIDR (Classless Inter Domain Routing).

It is recommended that RIP 2 be used on any segment where all routers can use the same IP routing protocol. If one or more routers on a segment must use RIP 1, then all other routers on that segment should also be set to use RIP 1.

- If **RIP 2** is selected with this pull-down menu, the router will send and/or accept RIP 2 packets over this interface, and will then periodically update its routing table with the information provided from these packets. On a large network, an up-to-date routing table will enhance network performance since the router will always be aware of the optimal path to use when sending packets.
- If **RIP 1** is selected with this pull-down menu, the router will send and/or accept RIP 1 packets, and will then periodically update its routing table with the information provided from these packets.
- If **None** is selected with this pull-down menu, the router will not be able to update its routing table and will always direct traffic for addresses it does not have a route for (addresses not on one of the networks connected to its interfaces) to the “default router” defined in its IP Static Route Dialog Box. It will then be the responsibility of the default router to direct the packets to the correct address. For information on setting the

default router see the discussion of the IP Static Route Dialog Box later in this chapter.

❖ **Note:** *Some routers, in particular those designed to create very large corporate backbones, may use other routing protocols such as OSPF (Open Shortest Path First). These routers can simultaneously use RIP 1 (and in some cases RIP 2) to communicate with smaller routers, or each of the smaller routers can be set to use one of these backbone routers as their default router.*

> Update Method

VPN links which are configured to provide “dial-on-demand” service will bring a connection up (i.e. dial the other end) when there are network packets which must be transferred over the link. Once a dial-on-demand connection is up, network traffic passing across the link causes the inactivity timer for the link to be reset, keeping the connection up.

The RIP protocol periodically sends out update information across a link. These periodic update packets will cause a VPN link set for dial-on-demand operation to stay up indefinitely.

- If **Triggered** is selected with this pull-down menu, the router will modify the standard RIP behavior for this link to send RIP packets only when there has been an update to its routing table information, or when it has detected a change in the accessibility of the next hop router.
- If **Periodic** is selected with this pull-down menu, the router will use the standard RIP protocol, which sends RIP packets over the link every 30 seconds.

RIP Split Horizon

Normally, RIP uses a technique called split horizon to avoid routing loops and allow smaller update packets. This technique specifies that when the device sends a RIP update out a particular network interface, it should never include routing information acquired over that same interface.

There is a variation of the split horizon technique called “poison reverse” which specifies that all routes should be included in an update out a particular interface, but that the metric should be set to infinity for those routes acquired over that interface. One drawback is that routing update packet sizes will be increased when using poison reverse.

- If **Split Horizon** is selected with this pull-down menu, the device will apply the split horizon technique to routes being output over this interface.

- If **No Split Horizon** is selected with this pull-down menu, the device will include all routes in an output packet, regardless of which interface they were acquired over, and will use a normal metric.
- If **Poison Reverse** is selected with this pull-down menu, the device will include all routes in an output packet, but will set the metric to infinity for those routes which were acquired over this interface.

Output RIP - Input RIP

These flags control the behavior of RIP 1 and RIP 2 for this interface, allowing the router to selectively send RIP, receive RIP, or both. The default (assuming RIP 1 or RIP 2 is turned on in the Routing Protocol popup) is to both send and receive.

Directed Broadcast

This checkbox sets whether the interface will forward network-prefix-directed broadcasts. This is a security feature which can help prevent your network from being used as an intermediary in certain kinds of attacks which use ICMP echo traffic (pings) or UDP echo packets with fake (i.e., “spoofed”) source addresses to inundate a victim with erroneous traffic.

OSPF

This option button brings up the OSPF Dialog Box which allows the OSPF routing protocol to be enabled. For more information on this dialog box and other OSPF parameters, refer to ***Chapter 15 - OSPF***.

TCP/IP Routing: Bridge Configuration Dialog Box

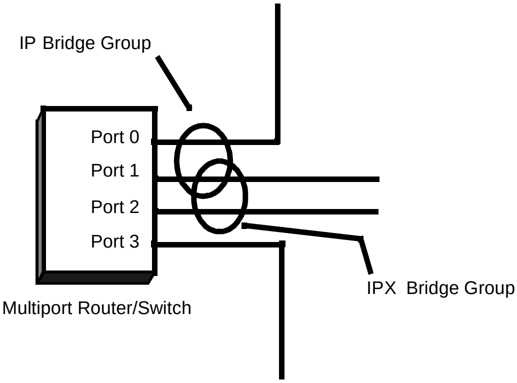

Bridge Logical Diagram

❖ **Note:** *If you need more information about bridging, see "Bridging 101" in the Appendices to this manual.*

Bridging operates on physical network addresses (such as Ethernet addresses), rather than logical addresses (such as IP addresses). From the standpoint of IP networking, interfaces which are set to bridge IP between themselves appear as a single logical entity.

Thus, a device's "IP Bridge Group" is made up of all of the physical network interfaces in a device which have been set to bridge IP. This setting can be found in the TCP/IP Routing Configuration Dialog Box for each individual physical interface. For example, see the IP Routing On/Bridge/Off radio buttons in the TCP/IP: Ethernet Routing Configuration Dialog Box.

Logically, the IP Bridge Group is treated by the device as an interface (Bridge 0). The settings in the TCP/IP Routing: Bridge 0 Configuration Dialog Box (discussed next) determine the IP parameters for all of the physical network interfaces which make up the IP Bridge Group. This is shown schematically in the diagram above.

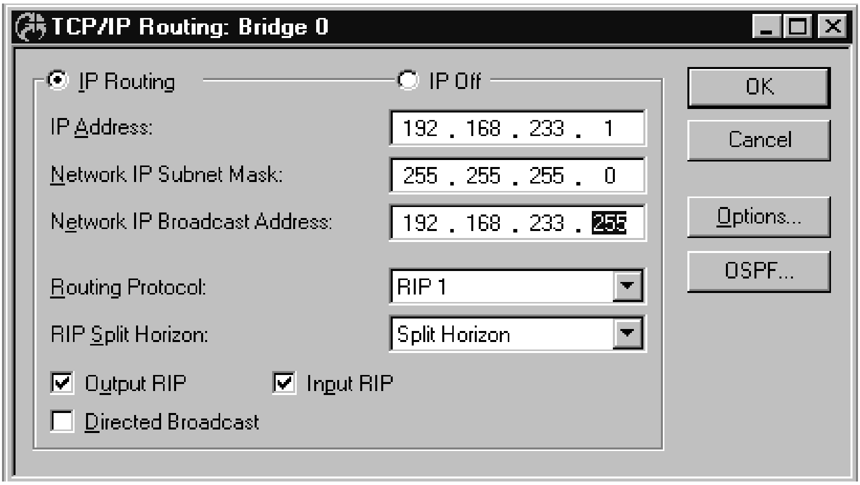

TCP/IP Routing: Bridge 0

● IP Routing ○ IP Off

IP Address: 192 . 168 . 233 . 1

Network IP Subnet Mask: 255 . 255 . 255 . 0

Network IP Broadcast Address: 192 . 168 . 233 . 255

Routing Protocol: RIP 1

RIP Split Horizon: Split Horizon

☑ Output RIP ☑ Input RIP

□ Directed Broadcast

TCP/IP Routing: Bridge 0 Configuration Dialog Box

❖ **Note:** *If you need more information about the IP protocol, see "IP 101" in the Appendices to this manual.*

To access this dialog box, select Bridge 0/TCP/IP Routing from the Device View.

> IP Routing/Off

These radio buttons control whether IP packets received by a member interface of the IP Bridge Group are passed on for IP routing.

- If set to **IP Routing**, then IP packets received on a member interface of the IP Bridge Group which cannot simply be bridged to another member interface of the group are passed on for IP routing.
- If set to **IP Off**, then IP packets received on a member interface of the IP Bridge Group which cannot be bridged to another member interface of the group are dropped. This setting means that further IP configuration information is not required for the IP Bridge Group.

> IP Address

Every network interface (including a logical interface, like the IP Bridge Group) on an IP internetwork must have a unique IP address that identifies that interface to other devices on the internetwork. Part of this address identifies the network segment(s) the IP Bridge Group is connected to, and the remainder uniquely identifies the IP Bridge Group itself.

This address should be entered as four decimal numbers separated by periods -- for example 198.238.9.5

❖ **Note:** *The single most common problem encountered in IP networking is the use of a duplicate IP address. You must carefully track the network numbers you have assigned to various devices in order to avoid hard-to-diagnose problems.*

> Network IP Subnet Mask

Most IP networks use “subnetting” in order to subdivide a large network into smaller logical sub-networks. The subnet mask value is used to tell the device what part of the IP address identifies the network segment (the “network” portion), and what part identifies individual interfaces (the “host” portion).

There are three generally used “classes” of subnetted IP networks: A, B and C. Each class uses a different amount of the IP address for the network and host portions. These classes may also be further divided by correctly setting the subnet mask.

If you do not enter a number in the Subnet Mask field, CompatiView will derive a default value from the IP Address number you entered just above. This default assumes you want a single subnet for all of the available host addresses. You must manually set the field if you want to further divide the address range.

To have CompatiView calculate a default mask, make sure that the Subnet Mask field is empty, position the cursor in the IP Address field, then just tab through the Subnet Mask field.

> Network IP Broadcast Address

The device will use this address to send any IP broadcast messages. The standard broadcast address is all 255’s (hexadecimal FFs) in the host portion of the address. A few networks use all zeroes in this field. If you are unsure which type your network uses, check with your network administrator.

To have CompatiView calculate a default broadcast address, make sure that the Broadcast Address field is empty, position the cursor in the Subnet Mask field, then just tab through the Broadcast Address field.

> Routing Protocol

Routers pass information between themselves about the most effective path for packet transfer between various end points. There are a number of different protocols which have been defined to facilitate the exchange of this information.

Routing Information Protocol (RIP) 1 is the most widely used routing protocol on IP networks. All gateways and routers that support RIP 1 periodically broadcast routing information packets. These RIP 1 packets contain information concerning the networks that the routers and gateways can reach as well as the number of routers/gateways that a packet must travel through to reach the receiving address.

RIP 2 is an enhancement of RIP 1 which allows IP subnet information to be shared among routers, and provides for authentication of routing updates. When this protocol is chosen, the router will use the multicast address 224.0.0.9 to send and/or receive RIP 2 packets for this Bridge Group's member interfaces. As with RIP 1, the router's routing table will be periodically updated with information received in these packets.

RIP 2 is more useful in a variety of environments and allows the use of variable subnet masks on your network. It is also necessary for implementation of "classless" addressing as accomplished with CIDR (Classless Inter Domain Routing).

It is recommended that RIP 2 be used on any logical network segment, including multiple physical segments which are part of a logical IP Bridge Group, where all routers can use the same IP routing protocol. If one or more routers on a segment must use RIP 1, then all other routers on that segment should also be set to use RIP 1.

- If **RIP 2** is selected with this pull-down menu, the router will send and/or accept RIP 2 packets via this Bridge Group's member interfaces, and will then periodically update its routing table with the information provided from these packets. On a large network, an up-to-date routing table will enhance network performance since the router will always be aware of the optimal path to use when sending packets.
- If **RIP 1** is selected with this pull-down menu, the router will send and/or accept RIP 1 packets, and will then periodically update its routing table with the information provided from these packets.
- If **None** is selected with this pull-down menu, the router will not be able to update its routing table and will always direct traffic for addresses it does not have a route for (addresses not on one of the networks connected to its interfaces) to the "default router" defined in its IP Static Route Dialog Box. It will then be the responsibility of the default router to direct the packets to the correct address. For information on setting the default router see the discussion of the IP Static Route Dialog Box later in this chapter.

❖ **Note:** *Some routers, in particular those designed to create very large corporate backbones, may use other routing protocols such as OSPF (Open Shortest Path First). These routers can simultaneously use RIP 1 (and in some cases RIP 2) to communicate with smaller routers, or each of the* *smaller routers can be set to use one of these backbone routers as their default router.*

RIP Split Horizon

Normally, RIP uses a technique called split horizon to avoid routing loops and allow smaller update packets. This technique specifies that when the router sends a RIP update out a particular network interface (including a Bridge Group logical interface made up of multiple physical member interfaces), it should never include routing information acquired over that same interface.

There is a variation of the split horizon technique called “poison reverse” which specifies that all routes should be included in an update out a particular interface, but that the metric should be set to infinity for those routes acquired over that interface. One drawback is that routing update packet sizes will be increased when using poison reverse.

- If **Split Horizon** is selected with this pull-down menu, the router will apply the split horizon technique to routes being output over this Bridge Group’s member interfaces.
- If **No Split Horizon** is selected with this pull-down menu, the router will include all routes in output packets sent over this Bridge Group’s member interfaces, regardless of which interface they were acquired over, and will use a normal metric.
- If **Poison Reverse** is selected with this pull-down menu, the router will include all routes in an output packet sent over this Bridge Group’s member interfaces, but will set the metric to infinity for those routes which were acquired over these interfaces.

Directed Broadcast

This checkbox sets whether the interface will forward network-prefix-directed broadcasts. This is a security feature which can help prevent your network from being used as an intermediary in certain kinds of attacks which use ICMP echo traffic (pings) or UDP echo packets with fake (i.e., “spoofed”) source addresses to inundate a victim with erroneous traffic.

Options

The options button brings up the Bridge-TCP/IP Routing Options Dialog Box which provides access to Proxy ARP, UDP Relays and other configuration information. This dialog box is discussed later in this chapter.

OSPF

This option button brings up the OSPF Dialog Box which allows the OSPF routing protocol to be enabled. For more information on this dialog box and other OSPF parameters, refer to ***Chapter 15 - OSPF***.

IP Subinterface Dialog Box

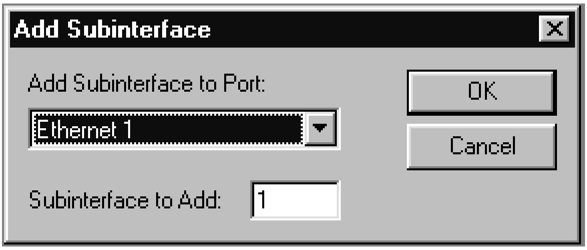

Add Subinterface

Add Subinterface to Port: Ethernet 1

Subinterface to Add: 1

Add IP Subinterface Dialog Box

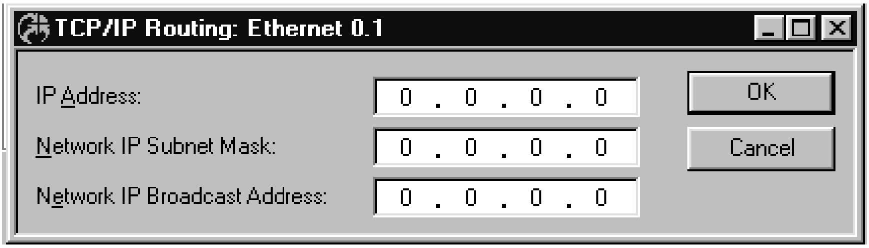

TCP/IP Routing: Ethernet 0.1

IP Address: 0 . 0 . 0 . 0

Network IP Subnet Mask: 0 . 0 . 0 . 0

Network IP Broadcast Address: 0 . 0 . 0 . 0

IP Subinterface Configuration Dialog Box

Subinterfaces are added to the edit area of a device by right-clicking on any configuration item for the device, then choosing Sub interface/Add. To delete a sub interface, right-click on the subinterface icon, then choose Subinterface/Delete. These functions are also available in the **Device** menu.

Once you have created a subinterface, you may access the IP Subinterface Configuration Dialog Box by clicking on TCP/IP under the subinterface icon.

IP subinterfaces allow the device to service more than one IP address range on a single physical network segment.

Because a routed IP packet does not contain any information regarding which networks it has passed across, the device must associate all IP packets received from a physical segment with the primary interface connected to that segment. As a result of this, the only IP parameters which can be set for subinterfaces are the IP Address, IP Subnet Mask, and IP Broadcast Address.

❖ **Note:** *Subinterfaces are only allowed on WAN ports configured for Frame Relay operation. They are not allowed on WAN ports configured for PPP. Frame Relay Glacis must be statically mapped when subinterfaces are in use, because IARP can only resolve a physical port, not a logical subinterface on that port.*

IP Connection Dialog Box

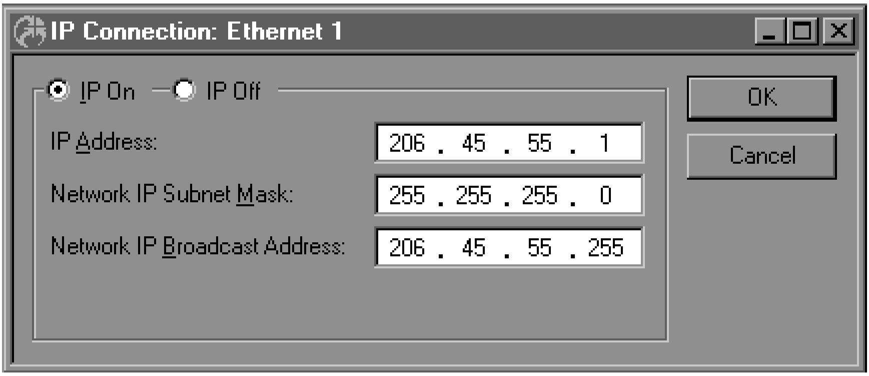
IP Connection: Ethernet 1

● IP On ○ IP Off

IP Address: 206 . 45 . 55 . 1

Network IP Subnet Mask: 255 . 255 . 255 . 0

Network IP Broadcast Address: 206 . 45 . 55 . 255

IP Connection Dialog Box

The IP Connection Dialog Box controls the IP settings for the IPSec-only port on an IntraPort VPN Access Router with two or more Ethernet interfaces. This port will only handle IPSec traffic (i.e., authenticated and/or encrypted packets).

To access this dialog box, select Ethernet/TCP/IP Routing from the Device View.

> IP On/IP Off

This set of radio buttons controls how IP packets are handled for this interface.

- If set to **IP On**, then IPSec packets received on this interface are routed to the correct interface on the router.
- If set to **IP Off**, then any IP packets received on this interface are discarded.

IP Address

This is the IP address of the IPSec port. It should be entered as four decimal numbers separated by periods -- for example, 198.238.9.5

❖ **Note:** *This IP address must be on the same IP network as the IPSec Gateway, which is configured using the IPSec Gateway Dialog Box (under Global/IPSec Gateway).*

Network IP Subnet Mask

The subnet mask value is used to tell the router what part of the IP address identifies the network segment (the “network” portion), and what part identifies individual interfaces (the “host” portion).

If you do not enter a number in the Subnet Mask field, CompatiView will derive a default value from the IP Address number you entered just above. This default assumes you want a single subnet for all of the available host addresses. You must manually set the field if you want to further divide the address range.

To have CompatiView calculate a default mask, make sure that the Subnet Mask field is empty, position the cursor in the IP Address field, then just tab through the Subnet Mask field.

Network IP Broadcast Address

The router will use this address to send any IP broadcast messages. To have CompatiView calculate a default broadcast address, make sure that the Broadcast Address field is empty, position the cursor in the Subnet Mask field, then just tab through the Broadcast Address field.

IP Static Routing Dialog Box

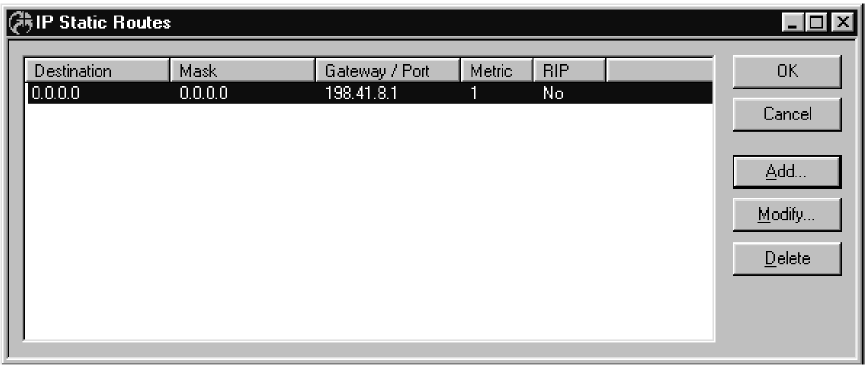

Destination	Mask	Gateway / Port	Metric	RIP
0.0.0.0	0.0.0.0	198.41.8.1	1	No

Static IP Routing Configuration Dialog Box

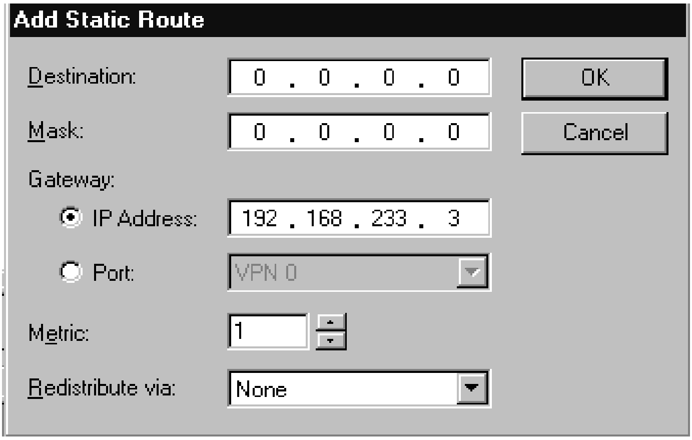

Destination: 0 . 0 . 0 . 0

Mask: 0 . 0 . 0 . 0

Gateway:

● IP Address: 192 . 168 . 233 . 3

○ Port: VPN 0

Metric: 1

Redistribute via: None

Add Static Route Dialog Box

To open the Static IP Routing Configuration Dialog Box, select Global/IP Static Routes. This dialog box displays static routes which have already been entered, but is not used to add or modify the entries.

To add or modify IP static route entries, you must access the Add Static Route Dialog Box by selecting the **Add...** or **Modify...** buttons in the Static IP Routing Configuration Dialog Box. The Add Static Route Dialog Box allows you to set a default IP router and to assign multiple static routes.

When you are finished adding entries, making changes, and marking deletions, click **OK** to store them in CompatiView's edit area for the device, for later downloading. If you click **Cancel**, CompatiView will discard any changes and additions you made in this dialog box.

❖ **Note:** *The "default router" is used as a "route of last resort" when your device cannot determine where an IP packet should be sent. In very simple routing setups, including connecting small networks to the Internet through an Internet Service Provider, a default router entry may be the only routing information required.*

Static routes are used to provide information to the device about where IP packets should be sent when the device itself has not been able to determine a correct route for them using dynamic routing information acquired through an IP routing protocol such as RIP.

In cases where the routing metrics (i.e. the number of routing hops to a destination) are equal between a static route and a dynamic route, Compatible Systems devices will use the dynamic route.

❖ **Note:** *Static routes are more difficult to maintain and are generally not as reliable as dynamically determined routes. We recommend that you use static routing only when the network does not provide adequate routing information through RIP.*

> Destination

Enter an IP address here in decimal notation for which you wish to provide static routing information. This can be a network address or an entire host address (e.g. 198.238.9).

By convention, 0.0.0.0 is used here for a default router entry.

> Mask

Enter a mask value here to tell the device how much of the Destination Address entry should be considered when determining the route for a packet. If you simply tab into this field, CompatiView will calculate a standard mask depending on the class of the Destination Address network. For instance, 255.255.255.0 tells the device to consider only the first three octets of a packet's address in determining whether it should be routed to the Gateway.

By convention, 0.0.0.0 is used here for a default router entry.

> Gateway

This section allows you to specify a gateway machine which is responsible for packets being sent to the Destination Address.

- If **IP Address** is selected, enter the IP address of the gateway.

- If **Port** is selected, use the pull-down menu to select an interface on the device you are configuring.

❖ **Note:** *The name of a physical port cannot be used when that port is configured for Frame Relay operation. This is because the Frame Relay protocol allows multiple IP addresses to be reached over a single physical port via different PVCs (permanent virtual circuits).*

> Metric

This is the number of “hops” that your device will assume exist between itself and the Gateway. It is also the number of hops that will be reported to other routers if you check the RIP box (as described below). When choosing how to forward a packet, a router will always pick a route with fewer hops over one with more. This value should be between 1 and 15.

❖ **Note:** *If you enter a smaller metric number, this route will tend to be preferred by your routers and other routers. If you enter a larger number, this route will tend to be overlooked in favor of other routes (if any exist) with lower metrics.*

> Redistribute via

This pull-down menu indicates whether a static route should be redistributed. Only one protocol can be selected for redistributing each static route.

- If **None** is specified, the static route will not be redistributed. Only one routing protocol can be selected for redistributing each static route.
- If **RIP** is specified, the static route entry will be redistributed into the RIP routing protocol which means that other routers will be able to choose this device as a way to forward packets to the destination address, depending on the metric and what other routes are available.

  Routing information received via RIP from other routers will be redistributed out other interfaces where RIP processing is enabled. When routes are rebroadcast in this fashion, the metric for this route is increased by 1, which increases the cost of the route.
- If **OSPF1** or **OSPF2** is specified, the static route entry will be redistributed into the OSPF routing protocol. The 1 or 2 refer to the two types of external metrics which may be used in OSPF.

  A type 1 cost is the sum of both the external cost and the internal cost used to reach that router. The cost of a type 2 route is simply the external cost, regardless of the interior (i.e., within OSPF) cost to reach that router.

- If **BGP** is specified, the static route entry will be redistributed into the BGP routing protocol.

Ethernet IP Options

Bridge IP Options

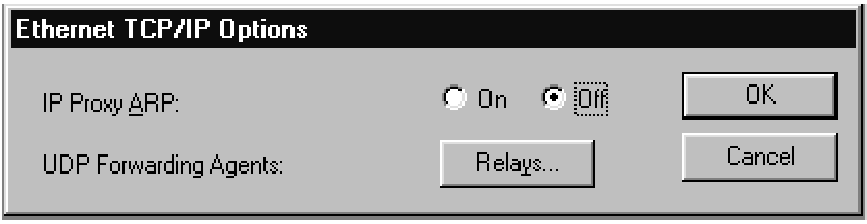

Ethernet TCP/IP Options

IP Proxy ARP: ○ On ● Off

UDP Forwarding Agents: Relays...

Ethernet or Bridge TCP/IP Options Dialog Box

To access this dialog box, select Ethernet/ or Bridge/TCP/IP Routing from the Device View, then click on the **Options** button.

This dialog box provides access to settings for IP Proxy ARP settings and the UDP Forwarding Agents Dialog Box.

IP Proxy ARP

Proxy ARP (Address Resolution Protocol) is used to allow the network portion of a group of IP addresses to be shared between several physical network segments. An example would be sharing one Class C address range between two physical Ethernets.

The ARP protocol itself provides a way for devices on an IP network to create a mapping between physical (i.e. Ethernet) addresses and logical IP addresses.

Proxy ARP makes use of this mapping feature by instructing a device to answer ARP requests as a “proxy” for the IP addresses behind one of its interfaces. The device which sent the ARP request will then correctly assume that it can reach the requested IP address by sending packets to the physical address that was returned to it.

This technique effectively hides the fact that a network has been (further) subnetted.

- If set to **On**, then any ARP request received on this interface whose IP network portion matches the network portion of the IP address on another interface of the device (as found by applying the Subnet Mask for that interface to the IP address for that interface) will be answered by the device with the physical address of this interface.
- If set to **Off**, then the device will only respond to ARP requests received for its own IP interface address. This is the default setting.

❖ **Note:** *Using Proxy ARP requires an in depth understanding of the workings of the IP protocol, along with careful manipulation of the IP subnet masks for the interfaces on a device. A more straightforward method of achieving similar results is to use Bridging (if available in your device).*

UDP Forwarding Agents (Relays)

The “Relays” button brings up a configuration dialog box that can be used to turn on a relay agent in the device for UDP (User Datagram Protocol) broadcast packets.

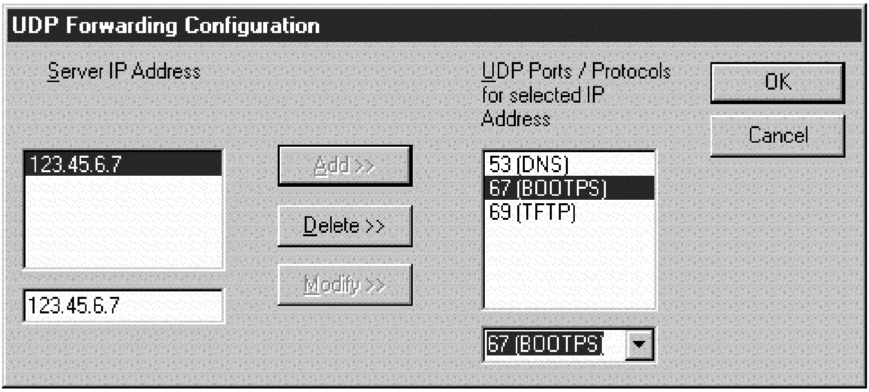

Server IP Address
123.45.6.7
123.45.6.7

UDP Ports / Protocols for selected IP Address
53 (DNS)
67 (BOOTPS)
69 (TFTP)
67 (BOOTPS)

UDP Forwarding Agents Dialog Box

Normally, a device will not forward UDP broadcast packets. However, many network applications use UDP broadcasts to configure addresses, hostnames, and other information. If hosts attempting to use these protocols are not on the same network segment as the servers which provide the information, the hosts will not receive a response unless a relay agent has been enabled in a device.

When a relay agent is enabled for an interface, the device is instructed to forward specific protocols received on that interface to a Server IP Address. The server does not need to reside on a network segment directly attached to the device.

> Server IP Address

You may enter server IP addresses in this list. When the Server IP Address edit box is selected, the Add, Delete, and Modify buttons will be activated for the list.

> UDP Ports/Protocols

This list allows you to enter the ports for which UDP relay will be performed. The list will show the services for well known ports in parentheses. When the UDP Port edit box is selected, the Add, Delete, and Modify buttons will be activated for the list.

The pull-down menu on the UDP Port edit box provides a list of well known services and automatically enters the UDP port number for a selected service into the list.

WAN IP Options

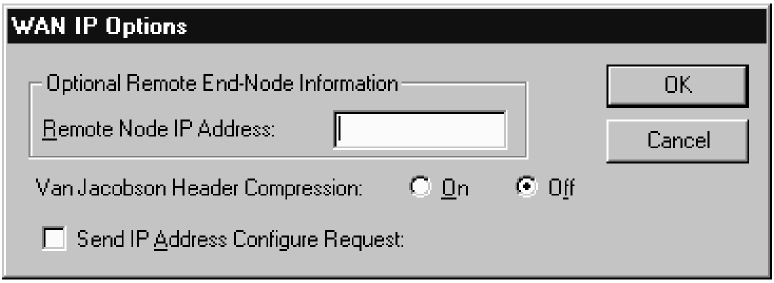

Remote Node IP Address:

Van Jacobson Header Compression: ○ On ● Off

□ Send IP Address Configure Request:

WAN IP Options Dialog Box

To access this dialog box, select WAN/TCP/IP Routing from the Device View, then click on the **Options** button.

This dialog box provides access to settings for Remote Node IP Address, Van Jacobson Header Compression, and IP Address Configure Request.

> Optional Remote End-Node Address

Besides defining a method for router-to-router communication, the PPP protocol defines a method for individual client machines to dial in to a router interface. Once a client machine has connected to a router interface in this fashion, the router provides proxy services which allow the client machine to participate as a node on one of the router’s local networks.

If remote node operation is desired, the WAN interface would usually be set up as an unnumbered interface, and the Remote Node Address would then be set to an unused IP address from the router's Ethernet network(s).

Alternatively, if the interface is set to be numbered, an unused address from the interface's host range may be used.

As always, it is imperative in either case that this IP address be unique.

The address should be entered as four decimal numbers separated by periods -- for example 198.238.10.10

> Van Jacobson Header Compression

Named for the gentleman who developed it, VJHC (Van Jacobson Header Compression) is a standard method of reducing the amount of redundant IP header information which is transferred over a wide area connection. VJHC reduces the size of the IP header to as few as three bytes.

There is a trade-off between the amount of time it takes to compress the header information, and the amount of time it would take to simply send it in native form across the WAN link.

❖ **Note:** *A general rule of thumb for Compatible Systems routers would be to use VJHC on uncompressed links at up to 56K rates, but to turn it off at higher speeds or if other means of compression (such as the V.42 compression built into modems) are in use. A few simple file copy transfer tests over your particular WAN setup will yield a more exact answer.*

Send IP Address Configure Request

A few third party routers implement the PPP specification in such a way that they require a PPP Address Configure Request to be sent when IP communications are being negotiated. This checkbox tells the router to include such a request with the IP address for this interface. Most routers do not require this information, and this checkbox should generally be left unchecked (default value).

TCP/IP Routing Options

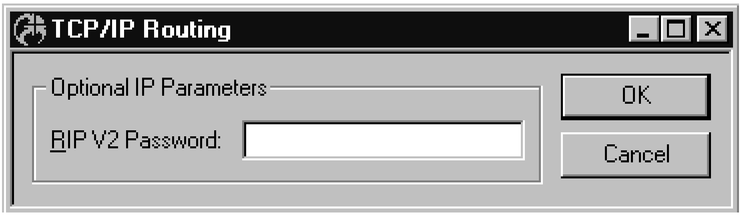

RIP V2 Password:

TCP/IP Routing Options Dialog Box

This dialog box can be brought up selecting Options/TCP/IP Routing from the Device View. These parameters are not associated with a particular interface and are global to the device.

RIP V2 Password

This password is used for authentication of RIP 2 packets received by the device. It is also included in RIP 2 packets sent by the device.

IP Multiprotocol Precedence Dialog Box

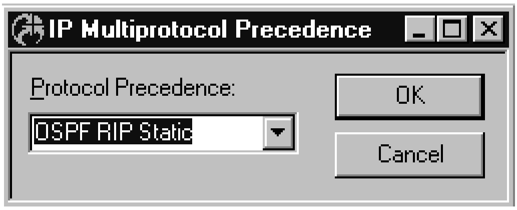

Protocol Precedence: OSPF RIP Static

IP Multiprotocol Precedence Dialog Box

This dialog box sets the precedence order the router will follow for including routes in its routing table when multiple IP routing protocols are in use on the network. To access this dialog box, select Global/IP Multiprotocol Precedence from the Device View.

Protocol Precedence

This pull-down menu sets the precedence order for including routes in the device’s IP routing table. This parameter is only relevant if there is more than one possible route to a destination. For example, if there are no OSPF or RIP routes to a destination but there is a static route, that route will be installed even if the precedence is **Ospf Rip Static**. Also, if there is a configured static route to a destination for which there was a RIP or OSPF route with greater precedence, that static route will be automatically re-installed if the RIP/OSPF route goes away.

❖ **Note:** *The BGP protocol will always be checked for first. Protocol Precedence is used to set the precedence order for RIP, Static, and OSPF protocols.*

❖ **Note:** *An exception to the precedence rule is an OSPF external (i.e., type ASE) route. OSPF external routes will be overwritten by a RIP or static route, regardless of the precedence. This is because OSPF external routes originally come from another protocol, usually RIP or static. If the router is running both RIP and OSPF, but another router on the network is redistributing RIP into OSPF, the RIP routes would be overwritten by OSPF external routes without this exception. In order to get the RIP routes via OSPF external routes, simply uncheck the* **Input RIP** *checkbox in the TCP/IP Routing Dialog Box, and it will then install the routes as OSPF externals.*

IP Route Redistribution

This section sets global configuration parameters which allow the redistribution of routes from one dynamic IP routing protocol into another. This allows RIP, OSPF, and BGP protocols to co-exist and exchange routing information. Route redistribution is global to the device.

❖ **Note:** *Redistribution of static routes can be done using the IP Multiprotocol Precedence Dialog Box.*

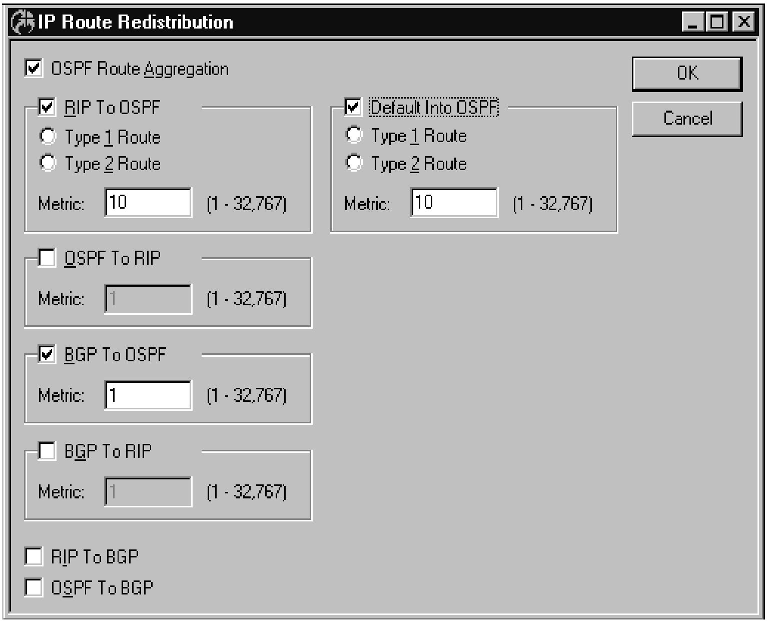

IP Route Redistribution

☑ OSPF Route Aggregation

☑ RIP To OSPF
○ Type 1 Route
○ Type 2 Route
Metric: 10 (1 - 32,767)

☑ Default Into OSPF
○ Type 1 Route
○ Type 2 Route
Metric: 10 (1 - 32,767)

□ OSPF To RIP
Metric: 1 (1 - 32,767)

☑ BGP To OSPF
Metric: 1 (1 - 32,767)

□ BGP To RIP
Metric: 1 (1 - 32,767)

□ RIP To BGP
□ OSPF To BGP

IP Route Redistribution Dialog Box

To access this dialog box, select Global/IP Route Redistribution from the device view.

OSPF Route Aggregation

This checkbox sets whether static and RIP routes will be consolidated along class boundaries before they are advertised into OSPF. If the router has a split subnet coming into the device from different interfaces, the box should be left unchecked.

❖ **Note:** *OSPF Route Aggregation is only used for importing static and RIP routes into OSPF. Aggregation of BGP routes is set in the BGP Aggregation dialog box. Refer to* ***Chapter 16 - BGP*** *for more information on configuration of BGP.*

RIP to OSPF

This checkbox sets whether the router will redistribute RIP routes into OSPF.

- **Type 1** is the sum of both the external cost and the internal cost used to reach that route.

- **Type 2** is the external cost, regardless of the interior cost to reach that route.
- The **Metric** parameter sets the external cost to be used. The value can be a number between 1 and 32,767. For a type 1 route, the internal costs along the routing path will be added to this cost to get the total cost.

Default into OSPF

This checkbox sets whether the router will redistribute default routes into OSPF. If left unchecked, a RIP or BGP default route will not be advertised into the OSPF domain even if non-default routes from that protocol are being redistributed.

- **Type 1** is the sum of both the external cost and the internal cost used to reach that route.
- **Type 2** is the external cost, regardless of the interior cost to reach that route.
- The **Metric** parameter sets the external cost to be used. The value can be a number between 1 and 32,767. For a type 1 route, the internal costs along the routing path will be added to this cost to get the total cost.

OSPF to RIP

This checkbox sets whether the router will redistribute OSPF routes in RIP. If checked, RIP will pick up the OSPF routes along with any other routes it is going to advertise.

BGP to OSPF

This checkbox sets whether the router will redistribute BGP routes into the OSPF routing domain.

❖ **Note:** *The full Internet BGP routing table cannot be redistributed into OSPF. Only up to 1,000 BGP routes will be accepted.*

BGP to RIP

This checkbox sets whether the router will redistribute BGP routes into RIP. If checked, RIP will pick up the BGP routes along with any other routes it is going to advertise.

RIP to BGP

This checkbox sets whether the router will redistribute RIP routes into the BGP routing domain.

OSPF to BGP

This checkbox sets whether the router will redistribute OSPF routes into the BGP routing domain.

❖ **Note:** *BGP will provide its own hop count in its route advertisements.*

Chapter 3 - IPX Routing & Bridging

IPX Routing: Ethernet Configuration Dialog Box

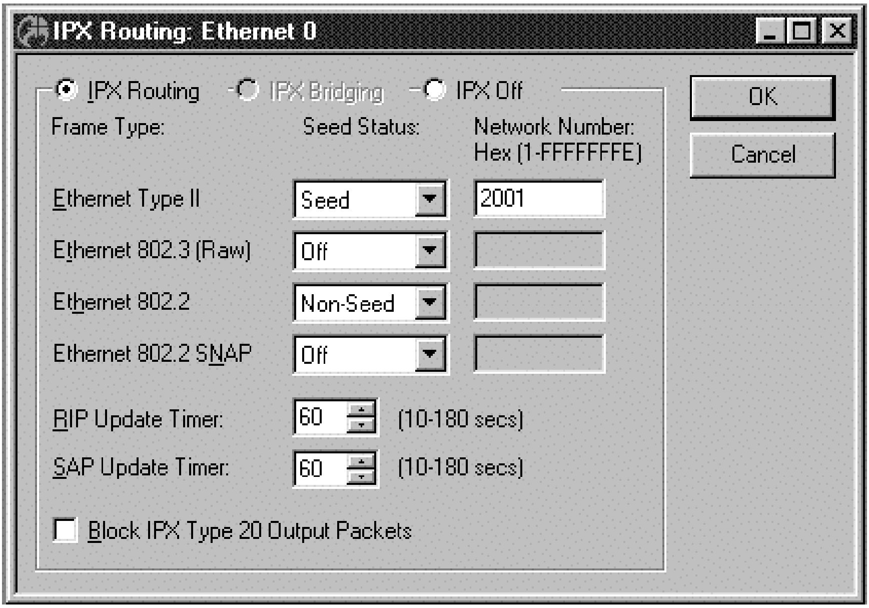

● IPX Routing ○ IPX Bridging ○ IPX Off

Frame Type:	Seed Status:	Network Number: Hex (1-FFFFFFFE)
Ethernet Type II	Seed	2001
Ethernet 802.3 (Raw)	Off	
Ethernet 802.2	Non-Seed	
Ethernet 802.2 SNAP	Off	

RIP Update Timer: 60 (10-180 secs)

SAP Update Timer: 60 (10-180 secs)

☐ Block IPX Type 20 Output Packets

IPX Routing: Ethernet Configuration Dialog Box

❖ **Note:** *If you need more information about the IPX protocol, see "IPX 101" in the Appendices to this manual.*

To access this dialog box, select Ethernet/IPX Routing in the Device View.

IPX Ethernet Frame Types

Compatible Systems devices support all four defined IPX frame types, and will perform routing between frame types as necessary. Whether each or all of these frame types are used on an individual Ethernet interface is determined by the settings for each type.

- **Ethernet Type II** is commonly used by TCP/IP and DECnet. The default seeding value is Non-Seed.
- **Ethernet 802.3 (Raw)** is the default frame type for earlier versions of Novell Netware. The default seeding value is Auto-Seed.
- **Ethernet 802.2** is a modified version of Ethernet_II and is the default frame type for Novell Netware 4. The default seeding value is Auto-Seed.
- **Ethernet 802.2 SNAP** is used by the AppleTalk protocol. The default seeding value is Non-Seed.

> IPX Routing/Bridging/Off

This set of radio buttons controls how IPX packets are handled for this interface.

- If set to **IPX Routing**, then IPX packets received on this interface are routed to the correct interface on the device.
- If set to **IPX Bridging**, then any IPX packets received on this interface are forwarded to the device's internal bridge. This setting makes this Ethernet interface a member of the "IPX Bridge Group" for this device.

❖ **Note:** *The IPX Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Ethernet-Bridging Dialog Box (under Ethernet/Bridging).*

- If it is set to **IPX Off**, then any IPX packets received on this interface are discarded.

> Seed Status (per Frame Type)

One of the functions which routers perform in IPX internetworking is setting the IPX network number for each network segment. A router which sets the network number for a segment is said to have "seeded" the network.

- **Seed** means the device will listen for an IPX network number being set by another router (including Novell software routers residing on servers) on the segment connected to this interface and use this number if it exists. If it doesn't discover a number in use, the device will use the configured IPX Network Number (discussed below) to set the network number for the segment.
- **Non-Seed** means the device will listen for an IPX network number being set by another router (including Novell software routers residing on servers) on the segment connected to this interface and use this number

if it exists. If it doesn’t discover a number in use, the device will wait indefinitely until a number is set by another router on the segment.

- **Auto-Seed** means the device will listen for an IPX network number being set by another router (including Novell software routers residing on servers) on the segment connected to this interface and use this number if it exists. If it doesn’t discover a number in use, the device will auto-generate a valid number using its routing tables.
- **Off** means the device will neither listen for, nor send packets with this frame type on this interface.

> Network Number (per Frame Type)

This is an eight-digit hexadecimal number that uniquely identifies the network segment connected to this interface. Values range from 1 to FFFFFFFE.

❖ **Note:** *Accidental selection of an IPX network number which is already in use on another network segment may cause hard-to-diagnose problems. You should carefully track which IPX network numbers are in use, and where they are used.*

RIP Update Timer

This value dictates how often the device sends out IPX RIP (Routing Information Protocol) packets on the network segment attached to this interface. The RIP packets sent out on this interface contain information about networks for which this device is responsible. RIP packets received tell the device about other networks and routers. The default is 60 seconds.

SAP Update Timer

This value dictates how often the device sends out IPX SAP (Service Access Protocol) packets on the network segment attached to this interface. The SAP packets sent out on this interface contain information about services (such as servers, printers, etc.) for which this device is responsible. SAP packets received tell this device about services available on other network segments. The default is 60 seconds.

Block IPX Type 20 Output Packets

In order for some protocols, notably NetBIOS, to function in the NetWare environment, routers must propagate a certain type of broadcast packet throughout an IPX internetwork. IPX packet type 20 is designated to perform broadcast propagation for these protocols.

When an IPX device receives a type 20 packet, it rebroadcasts it out all interfaces, except the one on which it was received. The IPX network number of the originating interface is included in the rebroadcast packets.

This checkbox allows you to control the rebroadcasting of IPX type 20 packets on this interface. This is useful for on-demand WAN links where the link may be brought up as a result of propagating this type of packet.

- If **checked**, then type 20 packets will not be propagated on this interface. The default is unchecked.

IPX Routing: WAN Configuration Dialog Box

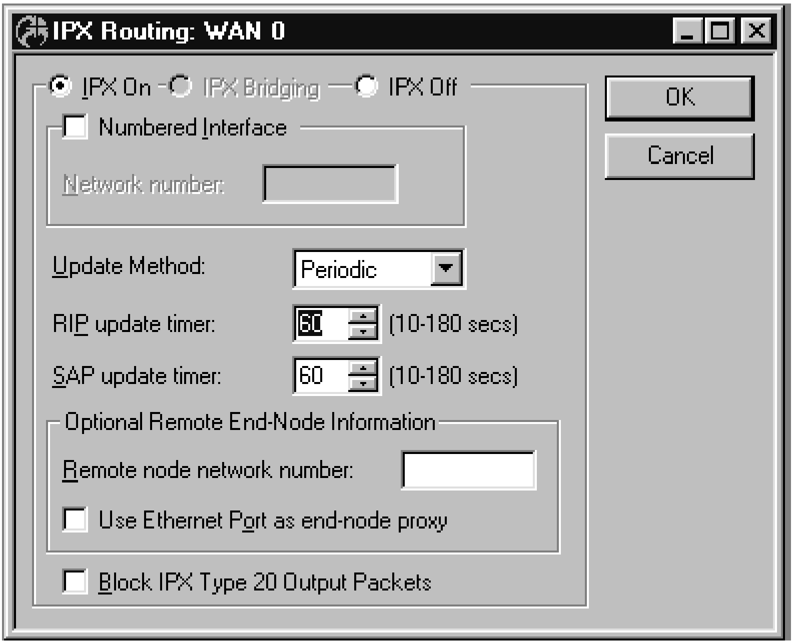

IPX Routing: WAN 0

● IPX On ○ IPX Bridging ○ IPX Off

☐ Numbered Interface

Network number:

Update Method: Periodic

RIP update timer: 60 (10-180 secs)

SAP update timer: 60 (10-180 secs)

Optional Remote End-Node Information

Remote node network number:

☐ Use Ethernet Port as end-node proxy

☐ Block IPX Type 20 Output Packets

IPX Routing: WAN Configuration Dialog Box

❖ **Note:** *If you need more information about the IPX protocol, see “IPX 101” in the Appendices to this manual.*

To access this dialog box, select WAN/IPX Routing in the Device View.

> IPX Routing/Bridging/Off

This set of radio buttons controls how IPX packets are handled for this interface.

- If set to **IPX Routing**, then IPX packets received on this interface are routed to the correct interface on the device.
- If set to **IPX Bridging**, then any IPX packets received on this interface are forwarded to the device’s internal bridge. This setting makes this interface a member of the “IPX Bridge Group” for this device.

❖ **Note:** *The IPX Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the WAN-Bridging Dialog Box (under WAN/Bridging).*

- If it is set to **IPX Off**, then any IPX packets received on this interface are discarded.

> Numbered Interface

This checkbox determines whether the Wide Area Network connected to this interface will have an IPX network number associated with it.

Many WAN connections are simple point-to-point links. These links do not generally require a network number because there are only two devices on the link. All traffic sent from one end is, by definition, destined for the other end. You generally do not need a numbered WAN interface if you are using the PPP transport protocol.

In contrast, Frame Relay networks may have a number of participating devices connected through a single physical interface. Because of this, use of the Frame Relay transport protocol requires a numbered WAN interface.

- If **checked**, then you must set an IPX Network Number (as described below) for this WAN interface. The default is unchecked.

Network Number

This is an eight-digit hexadecimal number that uniquely identifies the network segment connected to this interface. Values range from 1 to FFFFFFFE.

❖ **Note:** *Accidental selection of an IPX network number which is already in use on another network segment may cause hard-to-diagnose problems. You should carefully track which IPX network numbers are in use, and where they are used.*

> Update Method

WAN interfaces which are configured to provide “dial-on-demand” service will bring a connection up (i.e. dial the other end) when there are network packets which must be transferred over the link. Once a dial-on-demand connection is up, network traffic passing across the link causes the inactivity timer for the link to be reset, keeping the connection up.

The IPX RIP protocol periodically sends out update information across a link. These periodic update packets will cause a WAN interface set for dial-on-demand operation to either stay up indefinitely, or to continuously dial, connect, and then drop the connection.

- If **Triggered** is selected with this pull-down menu, the device will modify the standard IPX RIP behavior for this interface to send IPX RIP packets only when there has been an update to its routing table information, or when it has detected a change in the accessibility of the next hop router.
- If **Periodic** is selected with this pull-down menu, the device will use the standard IPX RIP protocol, which sends RIP packets over the link based on the RIP Update Timer value set below.

RIP Update Timer

This value dictates how often the device sends out IPX RIP (Routing Information Protocol) packets on the WAN link attached to this interface. The RIP packets sent out on this interface contain information about networks for which this device is responsible. RIP packets received tell the device about other networks and routers. The default is 60 seconds.

SAP Update Timer

This value dictates how often the device sends out IPX SAP (Service Access Protocol) packets on the WAN link attached to this interface. The SAP packets sent out on this interface contain information about services (such as servers, printers, etc.) for which this device is responsible. SAP packets received tell this device about services available on other network segments. The default is 60 seconds.

Optional Remote Node Network Number

Besides defining a method for router-to-router communication, the PPP protocol defines a method for individual client machines to dial in to a router interface. Once a client machine has connected to a router interface in this fashion, the router provides proxy services which allow the client machine to participate as a node on one of the router’s local networks.

If remote node operation is desired, the WAN interface would usually be set up as an unnumbered interface, and the Remote Node Network Number would then be set to an IPX network number from the router’s Ethernet interface(s).

Alternatively, if the interface is set to be numbered, an unused IPX network number may be used.

Use Ethernet Port as End-Node Proxy

The router can be set to dynamically reserve an IPX address for this WAN interface on an Ethernet segment. This proxy address will then be used if the remote PPP IPX implementation requests address negotiation (generally used by end-node clients).

Since the reserved address will be assigned to this interface, this checkbox can only be checked on an interface set to be unnumbered.

- If **checked**, then an IPX address will be reserved for this WAN interface on an Ethernet segment. The default is unchecked.

Block IPX Type 20 Output Packets

In order for some protocols, notably NetBIOS, to function in the NetWare environment, routers must propagate a certain type of broadcast packet throughout an IPX internetwork. IPX packet type 20 is designated to perform broadcast propagation for these protocols.

When an IPX router receives a type 20 packet, it rebroadcasts it out all interfaces, except the one on which it was received. The IPX network number of the originating interface is included in the rebroadcast packets.

This checkbox allows you to control the rebroadcasting of IPX type 20 packets on this interface. This is useful for on-demand WAN links where the link may be brought up as a result of propagating this type of packet.

- If **checked**, then type 20 packets will not be propagated on this interface. The default is unchecked.

❖ **Note:** *Novell’s router specification recommends that type 20 packets not be propagated across links with bandwidths of less than 1 megabit per second (such as asynchronous dial-up links and 56K leased lines).*

IPX Routing: VPN Configuration Dialog Box

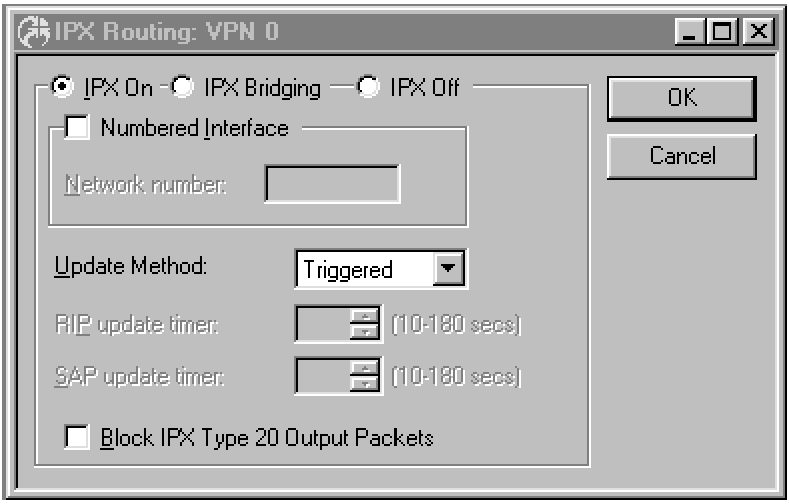

IPX Routing: VPN 0

● IPX On ○ IPX Bridging ○ IPX Off

□ Numbered Interface

Network number:

Update Method: Triggered

RIP update timer: (10-180 secs)

SAP update timer: (10-180 secs)

□ Block IPX Type 20 Output Packets

IPX Routing: VPN Configuration Dialog Box

VPN (Virtual Private Network) ports must first be added to the edit area of a device before they can be configured. For more information about adding and deleting VPN ports, see **Chapter 6 - VPN Ports and Tunnels**.

Once you have created a VPN port, you may access the IPX Routing: VPN Configuration Dialog Box by clicking on IPX Routing under the VPN port's icon.

A VPN port is a virtual port which handles tunneled traffic. Tunnels are virtual point-to-point connections through a public network such as the Internet. All packets sent through a VPN tunnel are IP-encapsulated packets, including AppleTalk, IPX and even IP packets. This encapsulation is added or removed, depending on the direction, by "Tunnel Partner" devices. Once a packet reaches the remote Tunnel Partner, the TCP/IP encapsulation is stripped off, leaving the original protocol. The unencapsulated packet is then handled according to the VPN port's protocol configuration settings. Networks connected via a tunnel will communicate as if they are on the same network, even though they are separated by the Internet.

❖ **Note:** *Remember that you must set up both ends of every tunnel. Therefore, you must repeat this setup with the remote device.*

To access this dialog box, select VPN/IPX Routing in the Device View.

> IPX Routing/Bridging/Off

This set of radio buttons controls how IPX packets are handled for this interface.

- If set to **IPX Routing**, then IPX packets received on this interface are routed to the correct interface on the device.
- If set to **IPX Bridging**, then any IPX packets received on this interface are forwarded to the device's internal bridge. This setting makes this interface a member of the "IPX Bridge Group" for this device.

❖ **Note:** *The IPX Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the VPN-Bridging Dialog Box (under VPN/Bridging).*

- If it is set to **IPX Off**, then any IPX packets received on this interface are discarded.

Numbered Interface

This checkbox determines whether the VPN port will have an IPX network number associated with it.

VPN tunnels are essentially point-to-point links. These links do not generally require a network number because all traffic sent from one end is, by definition, destined for the other end. However, you may wish to assign an address for network tracking purposes.

Network Number

This IPX Network Number is an eight-digit hexadecimal number that uniquely identifies the network segment(s) connected to this interface. Values range from 1 to FFFFFFFE.

❖ **Note:** *Accidental selection of an IPX network number which is already in use on another network segment may cause hard-to-diagnose problems. You should carefully track which IPX network numbers are in use, and where they are used.*

> Update Method

VPN links which are configured to provide "dial-on-demand" service will bring a connection up (i.e. dial the other end) when there are network packets which must be transferred over the link. Once a dial-on-demand connection is up, network traffic passing across the link causes the inactivity timer for the link to be reset, keeping the connection up.

The IPX RIP protocol periodically sends out update information across a link. These periodic update packets will cause a VPN link set for dial-on-demand operation to either stay up indefinitely, or to continuously dial, connect, and then drop the connection.

- If **Triggered** is selected with this pull-down menu, the device will modify the standard IPX RIP behavior for this link to send IPX RIP packets only when there has been an update to its routing table information, or when it has detected a change in the accessibility of the next hop router.
- If **Periodic** is selected with this pull-down menu, the device will use the standard IPX RIP protocol, which sends RIP packets over the link based on the RIP Update Timer value set below.

RIP Update Timer

This value dictates how often the device sends out IPX RIP (Routing Information Protocol) packets on the network segments attached to this interface. The RIP packets sent out on this interface contain information about networks for which this device is responsible. RIP packets received tell the device about other networks and routers. The default is 60 seconds.

SAP Update Timer

This value dictates how often the device sends out IPX SAP (Service Access Protocol) packets on the network segments attached to this interface. The SAP packets sent out on this interface contain information about services (such as servers, printers, etc.) for which this device is responsible. SAP packets received tell this device about services available on other network segments. The default is 60 seconds.

Block IPX Type 20 Output Packets

In order for some protocols, notably NetBIOS, to function in the NetWare environment, devices must propagate a certain type of broadcast packet throughout an IPX internetwork. IPX packet type 20 is designated to perform broadcast propagation for these protocols.

When an IPX device receives a type 20 packet, it rebroadcasts it out all interfaces, except the one on which it was received. The IPX network number of the originating interface is included in the rebroadcast packets.

This checkbox allows you to control the rebroadcasting of IPX type 20 packets on this interface. This is useful for on-demand links where the link may be brought up as a result of propagating this type of packet.

- If **checked**, then type 20 packets will not be propagated on this interface. The default is unchecked.

❖ **Note:** *Novell's router specification recommends that type 20 packets not be propagated across links with bandwidths of less than 1 megabit per second (such as asynchronous dial-up links and 56K leased lines).*

IPX Routing: Bridge Configuration Dialog Box

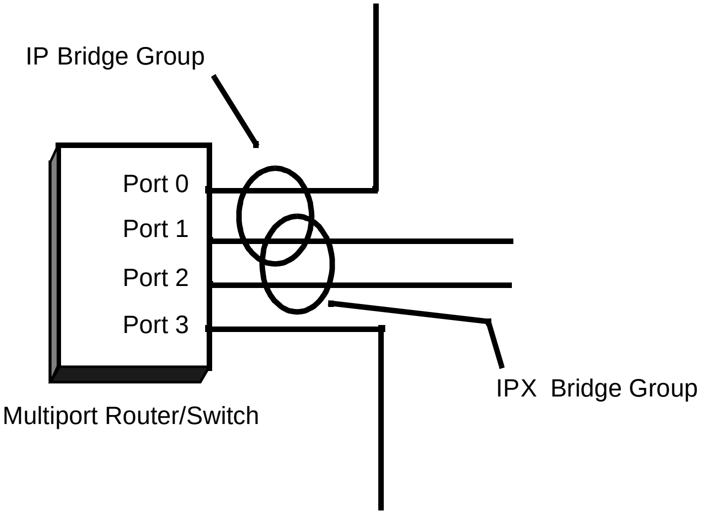

Bridge Logical Diagram

❖ **Note:** *If you need more information about bridging, see "Bridging 101" in the Appendices to this manual.*

Bridging operates on physical network addresses (such as Ethernet addresses), rather than logical addresses (such as IPX addresses). From the standpoint of IPX networking, interfaces which are set to bridge IPX between themselves appear as a single logical entity.

Thus, a device's "IPX Bridge Group" is made up of all of the physical network interfaces in a device which have been set to bridge IPX. This setting can be found in the IPX Configuration Dialog Box for each individual physical interface. For example, see the IPX Routing/Bridging/Off radio buttons in the IPX Routing: Ethernet Configuration Dialog Box.

Logically, the IPX Bridge Group is treated by the device as an interface (Bridge 0). The settings in the IPX Routing: Bridge 0 Configuration Dialog Box (discussed below) determine the IPX parameters for all of the physical network interfaces which make up the IPX Bridge Group. This is shown schematically in the diagram above.

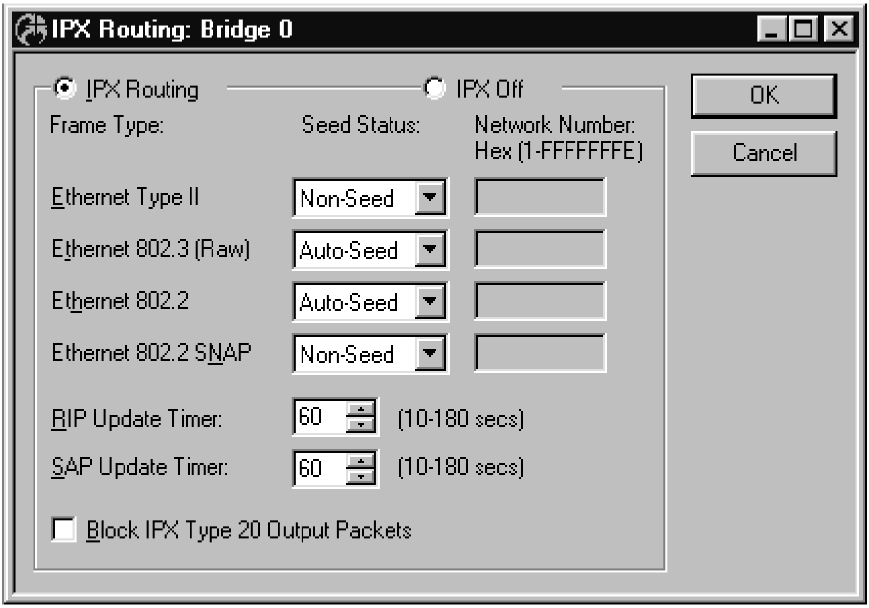

● IPX Routing ○ IPX Off

Frame Type:	Seed Status:	Network Number: Hex (1-FFFFFFFE)
Ethernet Type II	Non-Seed	
Ethernet 802.3 (Raw)	Auto-Seed	
Ethernet 802.2	Auto-Seed	
Ethernet 802.2 SNAP	Non-Seed	

RIP Update Timer: 60 (10-180 secs)

SAP Update Timer: 60 (10-180 secs)

□ Block IPX Type 20 Output Packets

IPX Routing: Bridge 0 Configuration Dialog Box

❖ **Note:** *If you need more information about the IPX protocol, see “IPX 101” in the Appendices to this manual.*

To access this dialog box, select Bridge 0/IPX Routing in the Device View.

IPX Frame Types

Compatible Systems devices support all four defined IPX frame types, and will perform routing between frame types as necessary. Whether each or all of these frame types are used on an individual Bridge interface is determined by the settings for each type.

- **Ethernet Type II** is commonly used by TCP/IP and DECnet. The default seeding value is Non-Seed.

- **Ethernet 802.3 (Raw)** is the default frame type for earlier versions of Novell Netware. The default seeding value is Auto-Seed.
- **Ethernet 802.2** is a modified version of Ethernet_II and is the default frame type for Novell Netware 4. The default seeding value is Auto-Seed.
- **Ethernet 802.2 SNAP** is used by the AppleTalk protocol. The default seeding value is Non-Seed.

> IPX Routing/Off

These radio buttons control whether IPX packets received by a member interface of the IPX Bridge Group are passed on for IPX routing.

- If set to **Routing**, then IPX packets received on a member interface of the IPX Bridge Group which cannot simply be bridged to another member interface of the group are passed on for IPX routing.
- If set to **Off**, then IPX packets received on a member interface of the IPX Bridge Group which cannot be bridged to another member interface of the group are dropped. This setting means that further IPX configuration information is not required for the IPX Bridge Group.

> Seed Status (per Frame Type)

One of the functions which routers perform in IPX internetworking is setting the IPX network number for each network segment. A device which sets the network number for a segment is said to have “seeded” the network. Remember that all segments connected to interfaces which are members of an IPX Bridge Group will appear as the same logical segment.

- **Seed** means the device will listen for an IPX network number being set by another device (including Novell software routers residing on servers) on the segment(s) connected to this interface and use this number if it exists. If it doesn’t discover a number in use, the device will use the configured IPX Network Number (discussed below) to set the network number for the segment(s)
- **Non-Seed** means the device will listen for an IPX network number being set by another router (including Novell software routers residing on servers) on the segment(s) connected to this interface and use this number if it exists. If it doesn’t discover a number in use, the device will wait indefinitely until a number is set by another router on the segment(s).
- **Auto-Seed** means the device will listen for an IPX network number being set by another router (including Novell software routers residing on servers) on the segment(s) connected to this interface and use this

number if it exists. If it doesn't discover a number in use, the device will auto-generate a valid number using its routing tables.

- **Off** means the device will neither listen for, nor send packets with this frame type on this interface.

> Network Number (per Frame Type)

This is an eight-digit hexadecimal number that uniquely identifies the network segment(s) connected to this interface. Values range from 1 to FFFFFFFE.

❖ **Note:** *Accidental selection of an IPX network number which is already in use on another network segment may cause hard-to-diagnose problems. You should carefully track which IPX network numbers are in use, and where they are used.*

RIP Update Timer

This value dictates how often the device sends out IPX \RIP (Routing Information Protocol) packets on the network segment(s) attached to this interface. The RIP packets sent out on this interface contain information about networks for which this device is responsible. RIP packets received tell the device about other networks and routers. The default is 60 seconds.

SAP Update Timer

This value dictates how often the device sends out IPX SAP (Service Access Protocol) packets on the network segment(s) attached to this interface. The SAP packets sent out on this interface contain information about services (such as servers, printers, etc.) for which this device is responsible. SAP packets received tell this device about services available on other network segments. The default is 60 seconds.

Block IPX Type 20 Output Packets

In order for some protocols, notably NetBIOS, to function in the NetWare environment, routers must propagate a certain type of broadcast packet throughout an IPX internetwork. IPX packet type 20 is designated to perform broadcast propagation for these protocols.

When an IPX device receives a type 20 packet, it rebroadcasts it out all interfaces, except the one on which it was received. The IPX network number of the originating interface is included in the rebroadcast packets.

This checkbox allows you to control the rebroadcasting of IPX type 20 packets on this interface. This is useful for on-demand WAN links where the link may be brought up as a result of propagating this type of packet.

- If **checked**, then type 20 packets will not be propagated on this interface. The default is unchecked.

Chapter 4 - AppleTalk Routing & Bridging

AppleTalk Routing: Ethernet Configuration Dialog Box

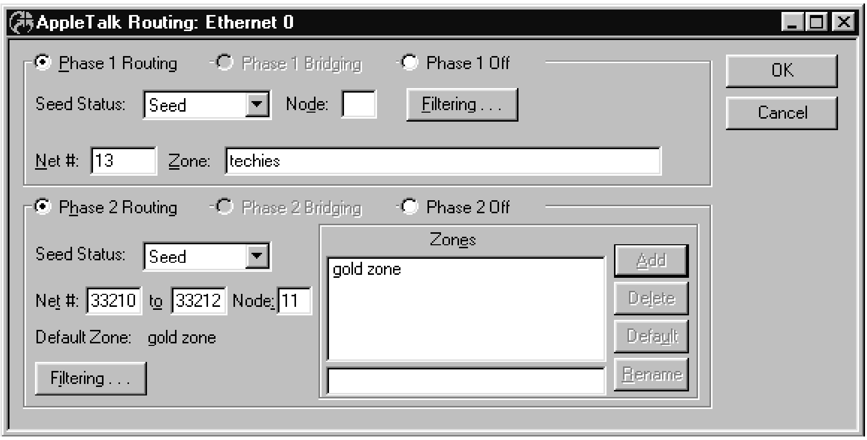

AppleTalk Routing: Ethernet 0

● Phase 1 Routing ○ Phase 1 Bridging ○ Phase 1 Off

Seed Status: Seed Node:

Net #: 13 Zone: techies

● Phase 2 Routing ○ Phase 2 Bridging ○ Phase 2 Off

Seed Status: Seed

Net #: 33210 to 33212 Node: 11

Default Zone: gold zone

Zones: gold zone

AppleTalk Routing: Ethernet Configuration Dialog Box

❖ **Note:** *If you need more information about the AppleTalk protocol, see "AppleTalk 101" in the Appendices to this manual.*

To access this dialog box, select Ethernet/AppleTalk Routing in the Device View.

AppleTalk Phase 1 Configuration

AppleTalk Phase 1 is an earlier version of the AppleTalk protocol which is still in use on some large legacy networks. Compatible Systems routers support this protocol, and "transitional routing" between it and AppleTalk Phase 2.

❖ **Note:** *Although Compatible Systems routers support AppleTalk Phase 1, we recommend that all new AppleTalk installations use AppleTalk Phase 2, which is much more capable.*

❖ **Note:** *In transitional routing installations, the same range of potential AppleTalk network numbers is shared by both Phase 1 and Phase 2. Care must be taken to avoid network number conflicts in these installations.*

> Phase 1 Routing/Bridging/Off

This set of radio buttons controls how AppleTalk Phase 1 packets are handled for this interface.

- If set to **Phase 1 Routing**, then AppleTalk Phase 1 packets received on this interface are routed to the correct interface on the router.
- If set to **Phase 1 Bridging**, then any AppleTalk Phase 1 packets received on this interface are forwarded to the router's internal bridge. This setting makes this Ethernet interface a member of the "AppleTalk Phase 1 Bridge Group" for this router.

❖ **Note:** *The Phase 1 Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Bridging: Ethernet Dialog Box (under Ethernet/Bridging).*

- If it is set to **Phase 1 Off**, then any AppleTalk Phase 1 packets received on this interface are discarded.

Phase 1 Seed Status

One of the functions which routers perform in AppleTalk internetworking is setting the AppleTalk network number for each network segment. A router which sets the network number for a segment is said to have "seeded" the network.

- **Seed** means the router will listen for an AppleTalk Phase 1 network number being set by another router on the segment connected to this interface and use this number if it exists. If it doesn't discover a number in use, the router will use the configured **AppleTalk Phase 1 Net #** (discussed below) to set the Phase 1 network number for the segment. It will also assign the configured **Phase 1 Zone** name to the segment.
- **Non-Seed** means the router will listen for an AppleTalk Phase 1 network number being set by another router on the segment connected to this interface and use this number if it exists. If it doesn't discover a number in use, the router will wait indefinitely until a number is set by another router on the segment.

- **Auto-Seed** means the router will listen for an AppleTalk Phase 1 network number being set by another router on the segment connected to this interface and use this number if it exists. If it doesn't discover a number in use, the router will auto-generate a valid number using its routing tables.

Phase 1 Net

For Ethernet interfaces which you set to **Seed** Phase 1, you must provide a network number. This is a decimal number that uniquely identifies the network segment connected to this interface, for Phase 1. Acceptable values range from 1 to 65,279.

❖ **Note:** *Accidental selection of an AppleTalk network number which is already in use on another network segment may cause hard-to-diagnose problems. You should carefully track which AppleTalk network numbers are in use, and where they are used.*

Phase 1 Zone

For Ethernet interfaces which you set to **Seed** Phase 1, you must provide a zone name. This is the name associated with the network number entered above. Zone names may be up to 32 characters in length.

Typically a name is chosen which has some significance to the physical location or the corporate purpose of the network segment. An example would be "Accounting Department."

This name will appear in the Chooser program of computers which support AppleTalk.

Phase 1 Node

You can provide a suggestion for the node number the router should use on this AppleTalk Phase 1 interface. The router will try to claim this number when it is powered up or restarted.

❖ **Note:** *The AppleTalk protocol allows network nodes to dynamically claim node numbers when they start up. Assigning known AppleTalk node numbers to router interfaces can make it easier to diagnose network problems using a network packet monitor.*

NBP Lookup Filters (Filtering)

The parameters required for NBP Filtering are contained in a configuration screen brought up by the "Filtering" button. This screen is discussed later in this chapter.

AppleTalk Phase 2 Configuration

AppleTalk Phase 2 is an updated version of the AppleTalk protocol which allows for more than 256 nodes on an Ethernet segment, and reduces the overhead required by AppleTalk RTMP (Routing Table Maintenance Protocol). AppleTalk Phase 2 should be used for all new installations.

> Phase 2 Routing/Bridging/Off

This set of radio buttons controls how AppleTalk Phase 2 packets are handled for this interface.

- If set to **Phase 2 Routing**, then AppleTalk Phase 2 packets received on this interface are routed to the correct interface on the router.
- If set to **Phase 2 Bridging**, then any AppleTalk Phase 2 packets received on this interface are forwarded to the router's internal bridge. This setting makes this Ethernet interface a member of the "AppleTalk Phase 2 Bridge Group" for this router.

❖ **Note:** *The Phase 1 Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Bridging: Ethernet Dialog Box (under Ethernet/Bridging).*

- If it is set to **Phase 2 Off**, then any AppleTalk Phase 2 packets received on this interface are discarded.

Phase 2 Seed Status

One of the functions which routers perform in AppleTalk internetworking is setting the AppleTalk network number for each network segment. A router which sets the network number for a segment is said to have "seeded" the network.

- **Seed** means the router will listen for an AppleTalk Phase 2 network range being set by another router on the segment connected to this interface and use this range if it exists. If it doesn't discover a range in use, the router will use the configured **AppleTalk Phase 2 Net #** range (discussed below) to set the Phase 2 network number(s) for the segment. It will also assign the configured **Phase 2 Zone** list to the segment.
- **Non-Seed** means the router will listen for an AppleTalk Phase 2 network range being set by another router on the segment connected to this interface and use this range if it exists. If it doesn't discover a range in use, the router will wait indefinitely until a range is set by another router on the segment.

- **Auto-Seed** means the router will listen for an AppleTalk Phase 2 network range being set by another router on the segment connected to this interface and use this range if it exists. If it doesn't discover a range in use, the router will auto-generate a valid number (a range of size 1) using its routing tables.

Phase 2 Net # Range

For Ethernet interfaces which you set to **Seed** Phase 2, you must provide a network number range. These two decimal numbers uniquely identify the range of AppleTalk network numbers for the network segment connected to this interface, for Phase 2. Acceptable values vary from 1 to 65,279. The value on the left must be smaller than the value on the right.

Each individual number in the range will support up to 253 node addresses.

❖ **Note:** *Accidental selection of an AppleTalk network number (or range of numbers) which is already in use on another network segment may cause hard-to-diagnose problems. You should carefully track which AppleTalk network numbers are in use, and where they are used.*

Phase 2 Zones

For Ethernet interfaces which you set to **Seed** Phase 2, you must provide a network number range. These are the names associated with the network number range entered above. You must specify at least one name, but it isn't necessary to specify a name for every number in the range. Zone names may be up to 32 characters in length.

Typically names are chosen which have some significance to the physical location or the corporate purpose of the network segment. Examples would be "Main Accounting," "Cost Accounting" and "Bookkeeping."

These names will appear in the Chooser program of computers which support AppleTalk. using the Network Control Panel, Macintosh computers are able to pick the zone in which they are located.

Phase 2 Default Zone

Use the Default button next to the Zone list to select which entry the router should designate as the default zone name for the segment. If you do not specify a default name, the router will designate the first name in the list.

Phase 2 Node

You can provide a suggestion for the node number the router should use on this AppleTalk Phase 2 interface.

❖ **Note:** *The AppleTalk protocol allows network nodes to dynamically claim node numbers when they start up. Assigning known AppleTalk node numbers* *to router interfaces can make it easier to diagnose network problems using a network packet monitor.*

NBP Lookup Filters (Filtering)

The parameters required for NBP Filtering are contained in a configuration screen brought up by the “Filtering” button. This screen is discussed later in this chapter.

AppleTalk Routing: WAN Configuration Dialog Box

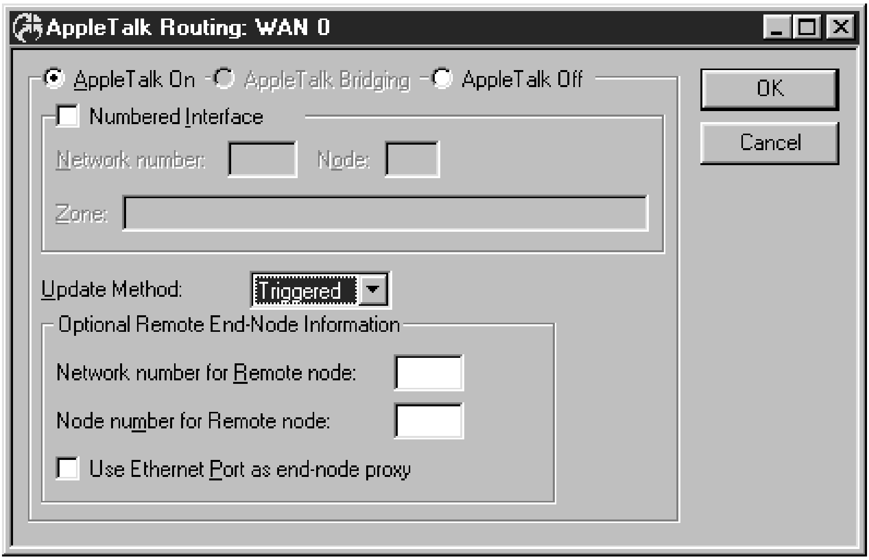

AppleTalk Routing: WAN 0

● AppleTalk On ○ AppleTalk Bridging ○ AppleTalk Off

□ Numbered Interface

Network number: Node:

Zone:

Update Method: Triggered

Optional Remote End-Node Information

Network number for Remote node:

Node number for Remote node:

□ Use Ethernet Port as end-node proxy

AppleTalk Routing: WAN Configuration Dialog Box

❖ **Note:** *If you need more information about the AppleTalk protocol, see “AppleTalk 101” in the Appendices to this manual.*

To access this dialog box, select WAN/AppleTalk Routing in the Device View.

> AppleTalk On/Bridging/Off

This set of radio buttons controls how AppleTalk packets are handled for this interface.

- If set to **AppleTalk On**, then AppleTalk packets received on this interface are routed to the correct interface on the router.
- If set to **AppleTalk Bridging**, then any AppleTalk packets received on this interface are forwarded to the router’s internal bridge. This setting makes this Ethernet interface a member of the “AppleTalk Phase 2 Bridge Group” for this router.

❖ **Note:** *The AppleTalk Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Bridging: WAN Dialog Box (under WAN/Bridging).*

- If it is set to **AppleTalk Off**, then any AppleTalk packets received on this interface are discarded.

> Numbered Interface

This check box determines whether the Wide Area Network connected to this interface will have an AppleTalk network number associated with it.

Many WAN connections are simple point-to-point links. These links do not generally require a network number because there are only two devices on the link. All traffic sent from one end is, by definition, destined for the other end. You generally do not need a numbered WAN interface if you are using the PPP transport protocol.

In contrast, Frame Relay networks may have a number of participating routers connected through a single physical interface. Because of this, use of the Frame Relay transport protocol requires a numbered WAN interface.

- If **checked**, then you must set an AppleTalk Network Number and Zone (as described below) for this WAN interface. The default is unchecked.

Network Number

If you have set this interface to be a numbered interface, you must provide a network number to identify the WAN link. This number creates a “non-extended” AppleTalk network on the WAN link. Acceptable values vary from 1 to 65,279.

❖ **Note:** *Accidental selection of an AppleTalk network number which is already in use on another network segment may cause hard-to-diagnose problems. You should carefully track which AppleTalk network numbers are in use, and where they are used.*

Zone

If you have set this interface to be a numbered interface, you must provide a zone name which will be associated with the network number entered above. Zone names may be up to 32 characters in length.

Typically a name is chosen which has some significance to the physical locations connected by the WAN link. An example would be “NYC - Chicago WAN.”

This name will appear in the Chooser program of computers which support AppleTalk, but there will be no selectable AppleTalk devices in the zone.

Node

If you have set this interface to be a numbered interface, you must provide an AppleTalk node number in this field which is unique for the network number you entered above.

❖ **Note:** *Compatible Systems routers require the assignment of a unique AppleTalk node number for numbered WAN interfaces. On Frame Relay networks in particular, you should keep a list of node number assignments to avoid conflicts.*

> Update Method

WAN interfaces which are configured to provide “dial-on-demand” service will bring a connection up (i.e. dial the other end) when there are network packets which must be transferred over the link. Once a dial-on-demand connection is up, network traffic passing across the link causes the inactivity timer for the link to be reset, keeping the connection up.

The AppleTalk RTMP protocol periodically sends out update information across a link. These periodic update packets will cause a WAN interface set for dial-on-demand operation to either stay up indefinitely or to continuously dial, connect, and then drop the connection.

- If **Triggered** is selected with this pull-down menu, the router will modify the standard AppleTalk RTMP behavior for this interface to send AppleTalk RTMP packets only when there has been an update to its routing table information, or when it has detected a change in the accessibility of the next hop router.
- If **Periodic** is selected with this pull-down menu, the router will use the standard AppleTalk RTMP protocol, which sends RTMP packets over the link every 10 seconds.

Optional Remote End-Node Network Number

Besides defining a method for router-to-router communication, the PPP protocol defines a method for individual client machines to dial in to a router interface. Once a client machine has connected to a router interface in this fashion, the router provides proxy services which allow the client machine to participate as a node on one of the router's local networks.

If remote end-node operation is desired, you must set the **AppleTalk Numbered Interface** checkbox on, and then set this network number field to the same value as you set in the **AppleTalk Network Number** field above.

Optional Remote End-Node Node Number

After setting the **Remote End-Node Network Number** above, select an unused node number for this field.

Do not use the same value you set in the **AppleTalk Node** field above.

Optional Remote End-Node Proxy

This checkbox sets the device to dynamically reserve an AppleTalk address on Ethernet for the WAN interface. This option can only be used on an unnumbered interface. If you wish to seed the proxy address to a specific network or node number, you must set the **AppleTalk Network Number** and the **AppleTalk Node** fields instead.

AppleTalk Routing: VPN Configuration Dialog Box

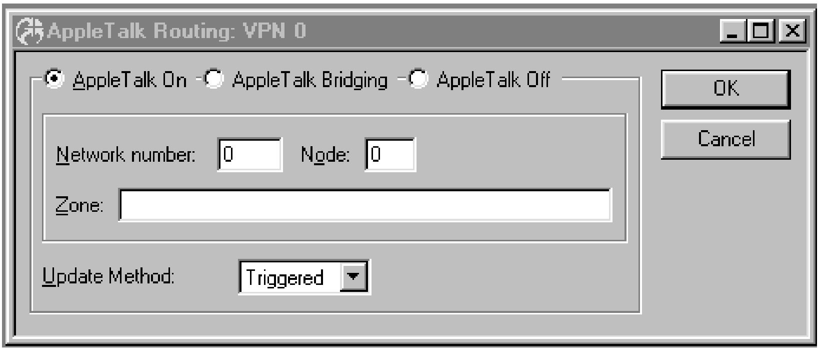

AppleTalk Routing: VPN 0

● AppleTalk On ○ AppleTalk Bridging ○ AppleTalk Off

Network number: 0 Node: 0

Zone:

Update Method: Triggered

AppleTalk Routing: VPN Configuration Dialog Box

VPN (Virtual Private Network) ports must first be added to the edit area of a device before they can be configured. For more information about adding and deleting VPN ports, see **Chapter 6: VPN Ports and Tunnels**.

Once you have created a VPN port, you may access the AppleTalk Routing: VPN Configuration Dialog Box by clicking AppleTalk Routing under the VPN port's icon.

A VPN port is a virtual port which handles tunneled traffic. Tunnels are virtual point-to-point connections through a public network such as the Internet. All packets sent through a VPN tunnel are IP-encapsulated packets, including AppleTalk, IPX and even IP packets. This encapsulation is added or removed, depending on the direction, by "Tunnel Peer" routers. Once a packet reaches the remote Tunnel Peer, the TCP/IP encapsulation is stripped off, leaving the original protocol. The unencapsulated packet is then handled according to the VPN port's protocol configuration settings. Networks connected via a tunnel will communicate as if they were on the same network, even though they are separated by the Internet.

❖ **Note:** *Remember that you must set up both ends of every tunnel. Therefore, you must repeat this setup with the remote router.*

> AppleTalk On/Bridging/Off

This set of radio buttons controls how AppleTalk packets are handled for this interface.

- If set to **AppleTalk On**, then AppleTalk packets received on this interface are routed to the correct interface on the router.
- If set to **AppleTalk Bridging**, then any AppleTalk packets received on this interface are forwarded to the router's internal bridge. This setting makes this Ethernet interface a member of the "AppleTalk Phase 2 Bridge Group" for this router.

❖ **Note:** *The AppleTalk Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Bridging: VPN Dialog Box (under VPN/Bridging).*

- If it is set to **AppleTalk Off**, then any AppleTalk packets received on this interface are discarded.

Network Number

If you have set this interface to be a numbered interface, you must provide a network number to identify the VPN port. This number creates a "non-extended" AppleTalk network on the VPN port. Acceptable values vary from 1 to 65,279.

❖ **Note:** *Accidental selection of an AppleTalk network number which is already in use on another network segment may cause hard-to-diagnose* *problems. You should carefully track which AppleTalk network numbers are in use, and where they are used.*

AppleTalk Zone

If you have set this interface to be a numbered interface, you must provide a zone name which will be associated with the network number entered above. Zone names may be up to 32 characters in length.

Typically a name is chosen which has some significance to the physical locations connected by the VPN link. An example would be “NYC - Chicago VPN.”

This name will appear in the Chooser program of computers which support AppleTalk, but there will be no selectable AppleTalk devices in the zone.

Node

If you have set this interface to be a numbered interface, you must provide an AppleTalk node number in this field which is unique for the network number you entered above.

❖ **Note:** *Compatible Systems routers require the assignment of a unique AppleTalk node number for numbered interfaces.*

> Update Method

VPN links which are configured to provide “dial-on-demand” service will bring a connection up (i.e. dial the other end) when there are network packets which must be transferred over the link. Once a dial-on-demand connection is up, network traffic passing across the link causes the inactivity timer for the link to be reset, keeping the connection up.

The AppleTalk RTMP protocol periodically sends out update information across a link. These periodic update packets will cause a VPN link set for dial-on-demand operation to either stay up indefinitely or to continuously dial, connect, and then drop the connection.

- If **Triggered** is selected with this pull-down menu, the router will modify the standard AppleTalk RTMP behavior for this interface to send AppleTalk RTMP packets only when there has been an update to its routing table information, or when it has detected a change in the accessibility of the next hop router.

- If **Periodic** is selected with this pull-down menu, the router will use the standard AppleTalk RTMP protocol, which sends RTMP packets over the link every 10 seconds.

AppleTalk Routing: Bridge Configuration Dialog Box

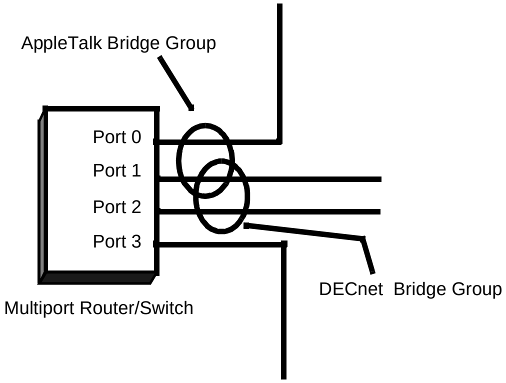

Bridge Logical Diagram

❖ **Note:** *If you need more information about bridging, see "Bridging 101" in the Appendices to this manual.*

Bridging operates on physical network addresses (such as Ethernet addresses), rather than logical addresses (such as AppleTalk Phase 2 addresses). From the standpoint of AppleTalk networking, router interfaces which are set to bridge AppleTalk Phase 2 between themselves appear as a single logical entity.

Thus, a router's "AppleTalk Phase 2 Bridge Group" is made up of all of the physical network interfaces in a router which have been set to bridge AppleTalk Phase 2. This setting can be found in the AppleTalk configuration dialog box for each individual physical interface. For example, see the AppleTalk Phase 2 Routing/Bridging/Off radio buttons in the AppleTalk Routing: Ethernet Configuration Dialog Box.

Logically, the AppleTalk Phase 2 Bridge Group is treated by the router as an interface (Bridge 0). The settings in the AppleTalk Routing: Bridge 0 Configuration Dialog Box (discussed below) determine the AppleTalk Phase 2 parameters for all of the physical network interfaces which make up the

AppleTalk Phase 2 Bridge Group. This is shown schematically in the diagram above.

❖ **Note:** *AppleTalk Phase 1 is generally treated as a distinct protocol for bridging and routing purposes, and thus will have its own "bridge group" should you decide to have a router bridge it.*

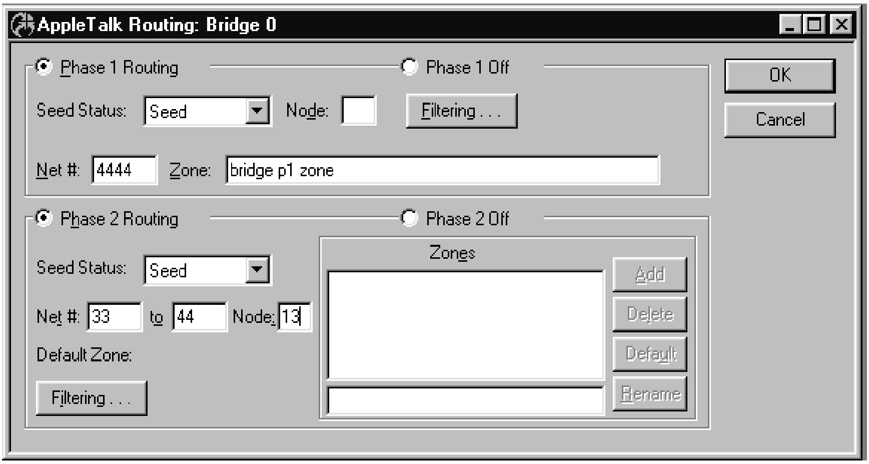

AppleTalk Routing: Bridge 0

● Phase 1 Routing ○ Phase 1 Off

Seed Status: Seed Node:

Net #: 4444 Zone: bridge p1 zone

● Phase 2 Routing ○ Phase 2 Off

Seed Status: Seed

Net #: 33 to 44 Node: 13

Default Zone:

AppleTalk Routing: Bridge 0 Configuration Dialog Box

❖ **Note:** *If you need more information about the AppleTalk protocol, see "AppleTalk 101" in the Appendices to this manual.*

To access this dialog box, select Bridge0/AppleTalk Routing in the Device View

AppleTalk Phase 1 Configuration

AppleTalk Phase 1 is an earlier version of the AppleTalk protocol which is still in use on some large legacy networks. Compatible Systems routers support this protocol, and "transitional routing" between it and AppleTalk Phase 2.

❖ **Note:** *Although Compatible Systems routers support AppleTalk Phase 1, we recommend that all new AppleTalk installations use AppleTalk Phase 2, which is much more capable.*

❖ **Note:** *In transitional routing installations, the same range of possible AppleTalk network numbers is used by both Phase 1 and Phase 2. Care must be taken to avoid network number conflicts in these installations.*

> Phase 1 Routing/Off

These radio buttons control whether AppleTalk Phase 1 packets received by a member interface of the AppleTalk Phase 1 Bridge Group are passed on for AppleTalk routing.

- If set to **Phase 1 Routing**, then AppleTalk Phase 1 packets received on a member interface of the AppleTalk Phase 1 Bridge Group which cannot simply be bridged to another member interface of the group are passed on for AppleTalk routing.
- If set to **Phase 1 Off**, then AppleTalk Phase 1 packets received on a member interface of the AppleTalk Phase 1 Bridge Group which cannot be bridged to another member interface of the group are dropped. This setting means that further AppleTalk configuration information is not required for the AppleTalk Phase 1 Bridge Group.

Phase 1 Seed Status

One of the functions which routers perform in AppleTalk internetworking is setting the AppleTalk network number for each network segment. A router which sets the network number for a segment is said to have “seeded” the network.

- **Seed** means the router will listen for an AppleTalk Phase 1 network number being set by another router on the segment(s) which are members of the AppleTalk Phase 1 Bridge Group and use this number if it exists. If it doesn’t discover a number in use, the router will use the configured **AppleTalk Phase 1 Net #** (discussed below) to set the Phase 1 network number for the segment(s). It will also assign the configured **Phase 1 Zone** name to the segment(s).
- **Non-Seed** means the router will listen for an AppleTalk Phase 1 network number being set by another router on the segment(s) which are members of the AppleTalk Phase 1 Bridge Group and use this number if it exists. If it doesn’t discover a number in use, the router will wait indefinitely until a number is set by another router on the segment(s).
- **Auto-Seed** means the router will listen for an AppleTalk Phase 1 network number being set by another router on the segment(s) which are members of the AppleTalk Phase 1 Bridge Group and use this number if it exists. If it doesn’t discover a number in use, the router will auto-generate a valid number using its routing tables.

Phase 1 Net

For an AppleTalk Phase 1 Bridge Group which you set to **Seed** Phase 1, you must provide a network number. This is a decimal number that uniquely identifies the network segment(s) which are part of the group, for Phase 1. Acceptable values range from 1 to 65,279.

❖ **Note:** *Accidental selection of an AppleTalk network number which is already in use on another network segment may cause hard-to-diagnose problems. You should carefully track which AppleTalk network numbers are in use, and where they are used.*

Phase 1 Zone

For an AppleTalk Phase 1 Bridge Group which you set to **Seed** Phase 1, you must provide a zone name. This is the name associated with the network number entered above. Zone names may be up to 32 characters in length.

Typically a name is chosen which has some significance to the physical location or the corporate purpose of the network segment(s). An example would be “Accounting Department.”

This name will appear in the Chooser program of computers which support AppleTalk.

Phase 1 Node

You can provide a suggestion for the node number the router should use on this AppleTalk Phase 1 Bridge Group. The router will try to claim this number when it is powered up or restarted.

❖ **Note:** *The AppleTalk protocol allows network nodes to dynamically claim node numbers when they start up. Assigning known AppleTalk node numbers to router interfaces can make it easier to diagnose network problems using a network packet monitor.*

NBP Lookup Filters (Filtering)

The parameters required for NBP Filtering are contained in a configuration screen brought up by the “Filtering” button. This screen is discussed later in this chapter.

AppleTalk Phase 2 Configuration

AppleTalk Phase 2 is an updated version of the AppleTalk protocol which allows for more than 256 nodes on an Ethernet segment, and reduces the overhead required by AppleTalk RTMP (Routing Table Maintenance Protocol). AppleTalk Phase 2 should be used for all new installations.

> Phase 2 Routing/Off

These radio buttons control whether AppleTalk Phase 2 packets received by a member interface of the AppleTalk Phase 2 Bridge Group are passed on for AppleTalk routing.

- If set to **Phase 2 Routing**, then AppleTalk Phase 2 packets received on a member interface of the AppleTalk Phase 2 Bridge Group which cannot simply be bridged to another member interface of the group are passed on for AppleTalk routing.
- If set to **Phase 2 Off**, then AppleTalk Phase 2 packets received on a member interface of the AppleTalk Phase 2 Bridge Group which cannot be bridged to another member interface of the group are dropped. This setting means that further AppleTalk configuration information is not required for the AppleTalk Phase 2 Bridge Group.

Phase 2 Seed Status

One of the functions which routers perform in AppleTalk internetworking is setting the AppleTalk network number for each network segment. A router which sets the network number for a segment is said to have “seeded” the network.

- **Seed** means the router will listen for an AppleTalk Phase 2 network range being set by another router on the segment(s) which are members of the AppleTalk Phase 2 Bridge Group and use this range if it exists. If it doesn’t discover a range in use, the router will use the configured **AppleTalk Phase 2 Net #** range (discussed below) to set the Phase 2 network number(s) for the segment(s). It will also assign the configured **Phase 2 Zone** list to the segment(s).
- **Non-Seed** means the router will listen for an AppleTalk Phase 2 network range being set by another router on the segment(s) which are members of the AppleTalk Phase 2 Bridge Group and use this range if it exists. If it doesn’t discover a range in use, the router will wait indefinitely until a range is set by another router on the segment(s).
- **Auto-Seed** means the router will listen for an AppleTalk Phase 2 network range being set by another router on the segment(s) which are members of the AppleTalk Phase 2 Bridge Group and use this range if it exists. If it doesn’t discover a range in use, the router will auto-generate a valid number (a range of size 1) using its routing tables.

Phase 2 Net # Range

For an AppleTalk Phase 2 Bridge Group which you set to **Seed** Phase 2, you must provide a network number range. These two decimal numbers uniquely identify the range of AppleTalk network numbers for the network segment(s) connected to this interface, for Phase 2. Acceptable values vary from 1 to 65,279. The value on the left must be smaller than the value on the right.

Each individual number in the range will support up to 253 node addresses.

❖ **Note:** *Accidental selection of an AppleTalk network number (or range of numbers) which is already in use on another network segment may cause hard-to-diagnose problems. You should carefully track which AppleTalk network numbers are in use, and where they are used.*

Phase 2 Zones

For an AppleTalk Phase 2 Bridge Group which you set to **Seed** Phase 2, you must provide a network number range. These are the names associated with the network number range entered above. You must specify at least one name, but it isn't necessary to specify a name for every number in the range. Zone names may be up to 32 characters in length.

Typically names are chosen which have some significance to the physical location or the corporate purpose of the network segment(s). Examples would be "Main Accounting," "Cost Accounting" and "Bookkeeping."

These names will appear in the Chooser program of computers which support AppleTalk. using the Network Control Panel, Macintosh computers are able to pick the zone in which they are located.

Phase 2 Default Zone

Use the Default button next to the Zone list to select which entry the router should designate as the default zone name for the segment(s) which are part of the group. If you do not specify a default name, the router will designate the first name in the list.

Phase 2 Node

You can provide a suggestion for the node number the router should use on this AppleTalk Phase 2 Bridge Group.

❖ **Note:** *The AppleTalk protocol allows network nodes to dynamically claim node numbers when they start up. Assigning known AppleTalk node numbers to router interfaces can make it easier to diagnose network problems using a network packet monitor.*

NBP Lookup Filters (Filtering)

The parameters required for NBP Filtering are contained in a configuration screen brought up by the "Filtering" button. This screen is discussed later in this chapter.

NBP Filtering

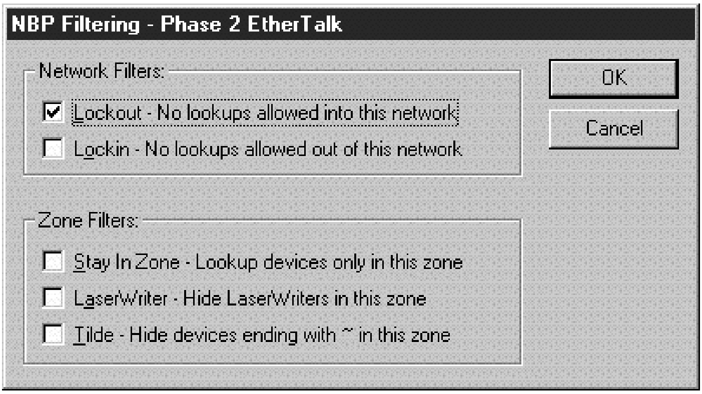

NBP Filtering - Phase 2 EtherTalk

Network Filters:

☑ Lockout - No lookups allowed into this network

□ Lockin - No lookups allowed out of this network

Zone Filters:

□ Stay In Zone - Lookup devices only in this zone

□ LaserWriter - Hide LaserWriters in this zone

□ Tilde - Hide devices ending with ~ in this zone

NBP Filtering Configuration Dialog Box

❖ **Note:** *The filtering functions discussed here are much less flexible than those discussed in the AppleTalk Filtering section of this manual. We suggest you read that section before choosing to use the filters discussed here.*

The NBP (Name Binding Protocol) Filtering Dialog Box is accessed by clicking the “Filtering” button in any Ethernet or Bridge port’s AppleTalk Configuration menu. NBP is a part of the AppleTalk protocols (both Phase 1 and Phase 2) which is used to discover the AppleTalk network number and node address of a named device on a network segment.

When the AppleTalk Chooser is opened on a computer, it causes NBP “lookup” packets for a specified device type in a selected AppleTalk zone to be sent. AppleTalk routers usually forward these NBP lookups onto any physical segments which are seeded with the selected AppleTalk zone name, and then forward any NBP replies back to the requesting computer.

NBP filters cause a router to selectively change the way it treats NBP lookup packets and NBP replies.

❖ **Note:** *These filter options can be used regardless of whether or not this router is acting as a seed router.*

Network Filters

Network filters are applied to the physical network segment connected to this interface. You may choose none, one or both of these options, depending upon how you wish to secure your network.

- Setting **Lockout** causes the router to drop any NBP lookups which are destined for this physical segment (or AppleTalk Bridge Group). This will protect devices on the segment from access by users on other segments.
- If you choose to **Lockin** lookups, the users on this network segment (or AppleTalk Bridge Group) will not have access through the router to network devices on other segments.

Zone Filters

Zone filters are applied based on logical AppleTalk zones rather than on physical segments. You may choose any or all combinations, depending on your network security requirements.

On AppleTalk Phase 1 networks and LocalTalk networks, zone filters are applied for the AppleTalk zone configured for the network segment. On AppleTalk Phase 2 networks, they are applied to the AppleTalk default zone configured for the network segment. For more information about creating a zone name on this port’s network segment, see the AppleTalk Routing configuration screen for this interface.

- **Stay In Zone** means the router will not forward NBP lookups which are directed from the AppleTalk zone configured for this port’s network segment to any other zone.
- The **LaserWriter** filter protects all LaserWriters in the AppleTalk zone configured for this port’s network segment from NBP lookup by computers in other AppleTalk zones.
- The **Tilde** filter protects all devices in the AppleTalk zone configured for this port’s network segment whose names end with a tilde (~) character from NBP lookup by computers in other AppleTalk zones.

❖ **Note:** *In order for Zone Name filters to work, the NBP lookup packets must pass through the router. This means that lookups between AppleTalk Phase 2 zones which are on the same network segment cannot be filtered in this fashion.*

AppleTalk Options Configuration Dialog Box

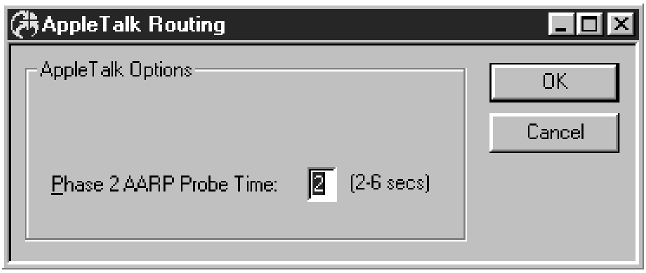

Phase 2 AARP Probe Time: 2 (2-6 secs)

AppleTalk Options Configuration Dialog Box

To access this dialog box, select Options/AppleTalk Routing from the Device View.

Phase 2 AARP Probe Time

This field allows the timeout for the AARP (Apple Address Resolution Protocol) address claim probes made at router startup time to be lengthened from the standard 2 seconds.

This may be necessary on AppleTalk networks which include WAN bridges. On these networks, it may take longer than 2 seconds for a node on the far side of a WAN bridge connection (logically still on the same AppleTalk network) to respond to an AARP address claim made by the router.

Chapter 5 - DECnet Routing & Bridging

Main DECnet Routing Configuration Dialog Box

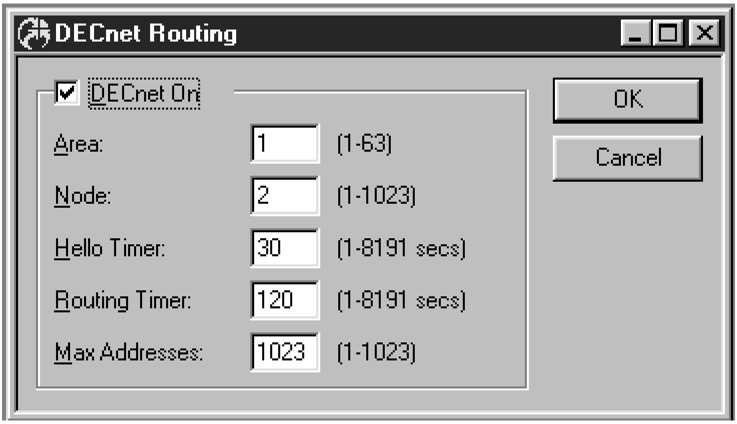

☑ DECnet On

Area: 1 (1-63)

Node: 2 (1-1023)

Hello Timer: 30 (1-8191 secs)

Routing Timer: 120 (1-8191 secs)

Max Addresses: 1023 (1-1023)

Main DECnet Routing Configuration Dialog Box

To access this dialog box, select Global/DECnet Routing in the Device View.

❖ **Note:** *Compatible Systems routers provide DECnet Phase IV Level 1 intra-area routing. All references to “DECnet” in this manual are to this set of protocols.*

> DECnet On

This checkbox controls how DECnet packets are handled for this router.

- If **checked**, then DECnet packets received on any interface in the router which has DECnet turned on will be routed to the correct interface.
- If **unchecked**, then DECnet packets received by this router will be discarded, and no DECnet packets will be sent by this router.

> Area

DECnet areas create a logical group of DECnet nodes. A DECnet area may include one or more physical network segments. The **Area** value must be within the range of 1 to 63.

The area information is specific to this individual router and, along with the **Node** number, uniquely identifies it on the network. If you are unsure what value to use here, check with your network administrator.

> Node

Each device in an area must have a unique node number. The Node value must be within the range of 1 to 1023.

The node number is specific to this individual router and, along with the **Area** number, uniquely identifies it on the network. If you are unsure what value to use here, check with your administrator.

❖ **Note:** *Using the same **Area:Node** combination as an address for two different devices can cause difficult-to-diagnose problems on your network. You should carefully track the assignment of this information for devices on your DECnet network.*

Hello Timer

DECnet hello messages tell end nodes which routers are available to route packets. This parameter tells the router how frequently it should send hello messages on its LAN interfaces.

The Hello Timer value is also inserted into the hello messages themselves. Once an end node has received a hello message from a router, it begins to track the availability of that router. If an end node does not hear an additional hello message within 3 timer periods, it assumes that this router is no longer available.

The default value for this parameter is 30 seconds.

❖ **Note:** *The **Hello Timer** values for individual WAN interfaces are set in separate windows. For more information, see the section in this chapter on the DECnet: WAN Configuration Dialog Box.*

Routing Timer

DECnet routing messages are exchanged between routers and contain routing table information including node numbers, hello timer values, hop counts and costs. This parameter tells the router how frequently it should send routing messages on its LAN interfaces.

The default value for this parameter is 120 seconds.

❖ **Note:** *The* ***Routing Timer*** *values for individual WAN interfaces are set in separate windows. For more information, see the section in this chapter on the DECnet: WAN Configuration Dialog Box.*

Max Addresses

This is the maximum number of node addresses allowed for this particular area. The default value for this parameter is 1023.

By limiting the number of addresses, a network administrator can limit the size of the internal routing table and the size of the routing messages sent to other routers.

Generally, all routers on the network should be consistent and use the same value for this parameter. This number should be at least as large as the number entered for this router's node number.

DECnet: Ethernet Configuration Dialog Box

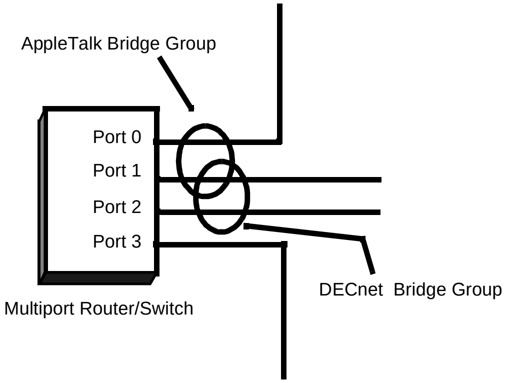

Bridge Logical Diagram

Bridging operates on physical network addresses (such as Ethernet addresses), rather than logical addresses (such as DECnet addresses). From the standpoint of DECnet networking, router interfaces which are set to bridge DECnet between themselves appear as a single logical entity.

Thus, a router's "DECnet Bridge Group" is made up of all of the physical network interfaces in a router which have been set to bridge DECnet.

Logically, the DECnet Bridge Group is treated by the router as an interface (Bridge 0). The settings in the Main DECnet Routing Configuration Dialog Box (discussed earlier in this chapter) determine the DECnet parameters for all of the physical network interfaces which make up the DECnet Bridge Group. This is shown schematically in the Bridge Logical Diagram.

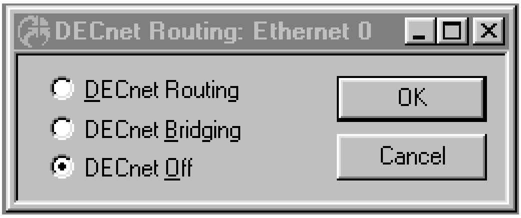

DECnet Routing: Ethernet 0

○ DECnet Routing
○ DECnet Bridging
● DECnet Off

DECnet: Ethernet Configuration Dialog Box

To access this dialog box, select Ethernet/DECnet Routing in the Device View.

❖ **Note:** *CompatiView only provides this configuration dialog box for routers which support bridging. Ethernet parameters for other routers are set globally in the Main DECnet Routing Configuration Dialog Box.*

❖ **Note:** *Compatible Systems routers provide DECnet Phase IV Level 1 intra-area routing. All references to "DECnet" in this manual are to this set of protocols.*

> DECnet Routing/Bridging/Off

This set of radio buttons controls how DECnet packets are handled for this interface.

- If set to **DECnet Routing**, then DECnet packets received on this interface are routed to the correct interface on the router.
- If set to **DECnet Bridging**, then any DECnet packets received on this interface are forwarded to the router's internal bridge. This setting makes this Ethernet interface a member of the "DECnet Bridge Group" for this router.

❖ **Note:** *The DECnet Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Bridging: Ethernet Dialog Box (under Ethernet/Bridging).*

- If it is set to **DECnet Off**, then any DECnet packets received on this interface are discarded.

DECnet: WAN Configuration Dialog Box

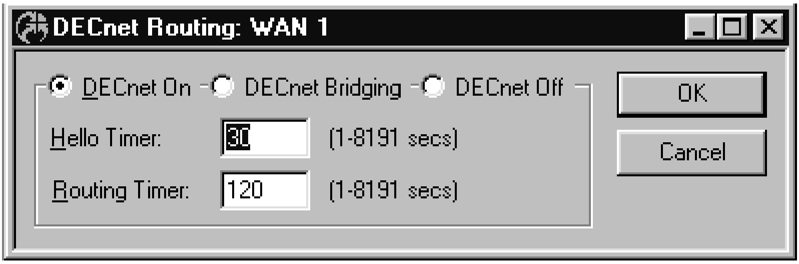
DECnet Routing: WAN 1

● DECnet On ○ DECnet Bridging ○ DECnet Off

Hello Timer: 30 (1-8191 secs)

Routing Timer: 120 (1-8191 secs)

DECnet: WAN Configuration Dialog Box

To access this dialog box, select WAN/DECnet Routing in the Device View.

❖ **Note:** *Compatible Systems routers provide DECnet Phase IV Level 1 intra-area routing. All references to “DECnet” in this manual are to this set of protocols.*

> DECnet On/Bridging/Off

This set of radio buttons controls how DECnet packets are handled for this interface.

- If set to **DECnet On**, then DECnet packets received on this interface are routed to the correct interface on the router.
- If set to **DECnet Bridging**, then any DECnet packets received on this interface are forwarded to the router’s internal bridge. This setting makes this WAN interface a member of the “DECnet Bridge Group” for this router.

❖ **Note:** *The DECnet Bridging radio button will be grayed out unless bridging has been turned on globally for the device using the Main Bridging Configuration Dialog Box (under Global/Bridging) and locally on this interface using the Bridging: WAN Dialog Box (under WAN/Bridging).*

- If it is set to **DECnet Off**, then any DECnet packets received on this interface are discarded.

Hello Timer

DECnet hello messages tell end nodes which routers are available to route packets. This parameter tells the router how frequently it should send hello messages on this interface.

The Hello Timer value is also inserted into the hello messages themselves. Once an end node has received a hello message from a router, it begins to track the availability of that router. If an end node does not hear an additional hello message within 3 timer periods, it assumes that this router is no longer available.

The default value for this parameter is 30 seconds. The maximum value is 8191 seconds (approximately 2 hours and 15 minutes).

❖ **Note:** *For dial-on-demand links, this parameter should be set to the longest period practical, since the router will dial the remote end each time one of these packets is sent.*

Routing Timer

DECnet routing messages are exchanged between routers and contain routing table information including node numbers, hello timer values, hop counts and costs. This parameter tells the router how frequently it should send routing messages on this interface.

The default value for this parameter is 120 seconds. The maximum value is 8191 seconds (approximately 2 hours and 15 minutes).

❖ **Note:** *For dial-on-demand links, this parameter should be set to the longest period practical, since the router will dial the remote end each time one of these packets is sent.*

Chapter 6 - VPN Ports and LAN-to-LAN Tunnels

Add VPN Port Dialog Box

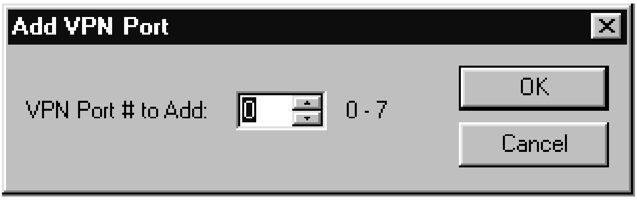

VPN Port # to Add: 0 0 - 7

Add VPN Port Dialog Box

This section configures VPN tunnel parameters and defines a virtual port for LAN-to-LAN tunnel traffic.

VPN (Virtual Private Network) ports are added to the edit area of a device by right-clicking on any configuration item for the device, then choosing VPN Port/Add VPN Port from the popup menu. The Add VPN Port Dialog Box will open in the Main Window and will allow you to select a number for the port. To delete a VPN port, right-click on the port's icon, then choose VPN Port/Delete VPN Port. These functions are also available in the **File** menu.

A VPN port is a virtual port which handles tunneled traffic. Tunnels are virtual point-to-point connections through a public network such as the Internet. All packets sent through a VPN tunnel are IP-encapsulated packets, including AppleTalk, IPX and even IP packets. This encapsulation is added or removed, depending on the direction, by "Tunnel Partner" routers. Once a packet reaches the remote Tunnel Partner, the TCP/IP encapsulation is stripped off, leaving the original protocol. The unencapsulated packet is then handled according to the VPN port's protocol configuration settings. Networks connected via a tunnel will communicate as if they are on the same network, even though they are separated by the Internet.

Tunnel Partner: VPN Configuration Dialog Box

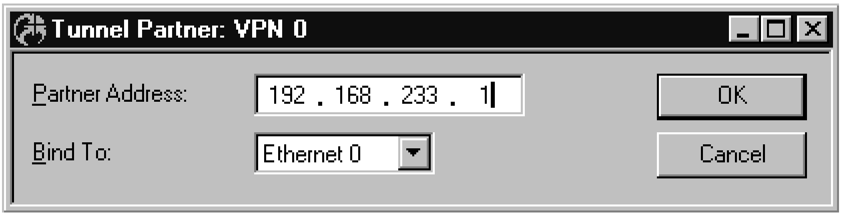

Partner Address: 192 . 168 . 233 . 1

Bind To: Ethernet 0

Tunnel Partner: VPN Configuration Dialog Box

Once you have created a VPN port, you may access the Tunnel Partner: VPN Configuration Dialog Box by clicking on the port's icon and selecting VPN Tunnel Partner.

❖ **Note:** *Remember that you must set up both ends of every tunnel. Therefore, you must repeat this setup with the remote router.*

> Partner Address

Enter the IP address of the remote Tunnel Partner with which this VPN port will communicate via the tunnel. This will be an interface on the remote router which has been set to route IP and will also be the remote VPN port's **Bind To** interface.

> Bind To Interface

Tunnel Partner devices must know each other's IP address in order to correctly address the packets destined for the far end of the tunnel. This device's tunnel end must have an IP address so that the Tunnel Partner can address packets to it. Use the pull-down menu to select an interface on this device which has been set to route IP.

❖ **Note:** *If both Ethernet ports are being used on an IntraPort 2/2+, then the Bind To port must be set to Ethernet 1.*

❖ **Note:** *You must enter the IP address for the interface you selected here into the Tunnel Partner dialog box of the Tunnel Partner routers*

IKE Key Management

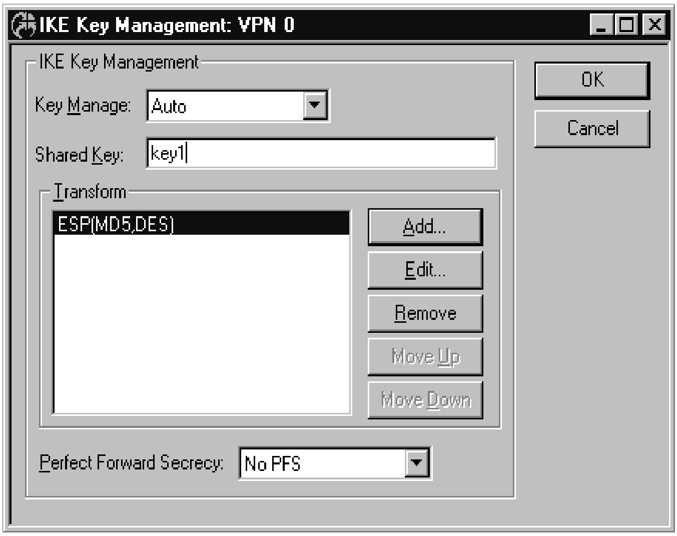

IKE Key Management: VPN 0

IKE Key Management

Key Manage: Auto

Shared Key: key1

Transform

ESP(MD5,DES)

Perfect Forward Secrecy: No PFS

IKE Key Management Dialog Box

Once you have created a VPN port, you may access the IKE Key Management Dialog Box by clicking on the port’s icon and selecting IKE Key Management.

This dialog box sets the Internet Security Association Key Management Protocol/Internet Key Exchange (ISAKMP/IKE) parameters. These settings control how each tunnel partner will identify and authenticate each other.

> Key Manage

- If **Auto** key management is selected, IKE will be used to allow two devices to negotiate between themselves what type of encryption and authentication to use for the tunnel. The **Auto** setting should only be used when the tunnel partner is another Compatible Systems VPN device.
- If **Manual** is selected, this Tunnel Partner will not use IKE, and the tunnel’s encryption and authentication parameters must be manually set in the Manual Key Management Dialog Box.
- If **Initiate** is selected, this Tunnel Partner will use IKE, but will only initiate tunnel establishment. It will not respond to tunnel establishment attempts from other devices.

- If **Respond** is selected, this Tunnel Partner will use IKE, but will only respond to tunnel establishment attempts which have been initiated by other devices. It will not initiate tunnel establishment.

Shared Key

This is a shared alphanumeric secret between 1-255 characters long. It is used to generate session keys which are used to authenticate and/or encrypt each packet received or sent through the tunnel.

Transform

This list box specifies the protection types and algorithms which will be used for tunnel sessions. Each option is a protection piece which specifies the authentication and/or encryption parameters to be used.

Use the **Move Up** and **Move Down** buttons to arrange the priority of the protection options.

> Perfect Forward Secrecy

Perfect Forward Secrecy (PFS) allows you to add an additional security parameter to tunnel sessions. PFS means that every time encryption and/or authentication key are computed, a new Diffie-Hellman Key Exchange is included.

Diffie-Hellman Key Exchange uses a complex algorithm and public and private keys to encrypt and then decrypt tunneled data. Adding PFS to a tunneled session greatly increases the difficulty of finding the session keys used to encrypt a VPN session. It also means that even if the keys are somehow cracked, only a portion of the traffic is recoverable.

- If **No PFS** is selected, this security parameter will not be added for this group configuration.

- If **Phase 1 Group** is selected, the group used in Phase 1 of the IKE negotiation is used as the group for the PFS Diffie-Hellman Key Exchange. This group is set (as G1 or G2) in the IKE Policy Dialog Box. For more information on the IKE Policy Dialog Box, refer to **Chapter 7 - VPN Client Tunnels**.

- If **DH Group 1** is selected, the Diffie-Hellman Group 1 algorithm will be used for the Diffie-Hellman Key Exchange.

- If **DH Group 2** is selected, the Diffie-Hellman Group 2 algorithm will be used for the Diffie-Hellman Key Exchange. Because larger numbers are used by the DH Group 2 algorithm, it is more secure than DH Group 1.

To add, remove, or edit a Transform, you must access the IKE Configuration Dialog Box by selecting the **Add...**, **Remove...**, or **Edit...** buttons.

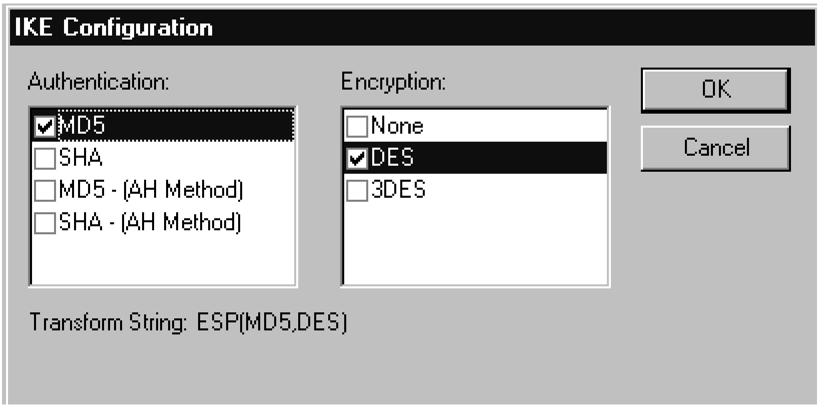

IKE Configuration

Authentication:

☑ MD5
□ SHA
□ MD5 - (AH Method)
□ SHA - (AH Method)

Encryption:

□ None
☑ DES
□ 3DES

Transform String: ESP(MD5,DES)

IKE Configuration Dialog Box

Authentication

This set of checkboxes specifies the authentication algorithm to be used for the negotiation. MD5 is the Message-Digest 5 hash algorithm. SHA is the Secure Hash Algorithm.

Choosing either of the top two checkboxes means that the Encapsulating Security Payload (ESP) header will be used to encrypt and authenticate packets.

Choosing either of the bottom two checkboxes specifies that the Authentication Header (AH) will be used to authenticate packets.

Encryption

This set of checkboxes specifies the encryption algorithm. DES (Data Encryption Standard) uses a 56-bit key to scramble the data. 3DES uses three different keys and three applications of the DES algorithm to scramble the data.

❖ **Note:** *You may choose only one authentication and one encryption method. The default setting of ESP (MD5,DES) is recommended for most setups.*

Manual Key Management

.

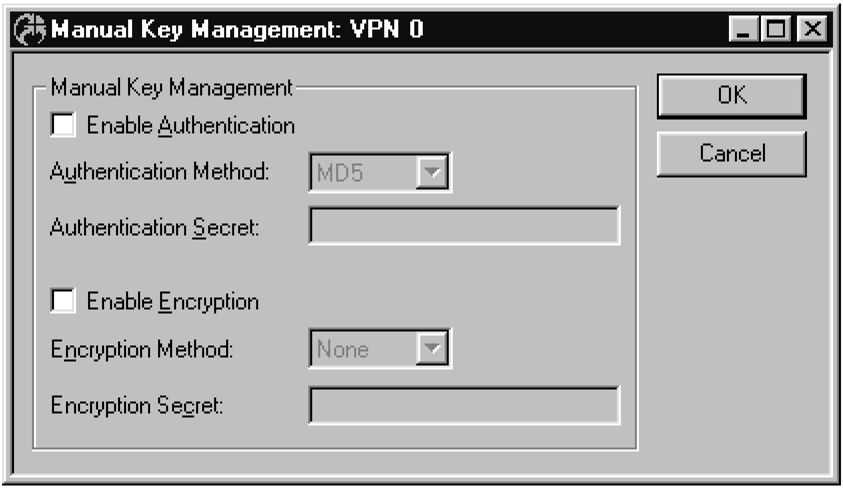

Manual Key Management: VPN 0

Manual Key Management

□ Enable Authentication

Authentication Method: MD5

Authentication Secret:

□ Enable Encryption

Encryption Method: None

Encryption Secret:

Manual Key Management Dialog Box

Once you have created a VPN port, you may access the Manual Key Management Dialog Box by clicking on the port’s icon and selecting Manual Key Management.

This dialog box sets encryption parameters for non-IKE tunnels.

Enable Authentication

This checkbox controls whether all tunnel traffic will be authenticated.

- If **checked**, then each packet will be digitally signed before sending. The receiving end of the tunnel will check the signature before allowing the traffic onto its local network.

Authentication Method

If Authentication has been enabled, **MD5** will appear here and packet-by-packet authentication will be done using the **Authentication Secret** set below.

Authentication Secret

This secret is used to generate session keys which are used to authenticate each packet received from or sent through the tunnel. The secret can be from 1 to 255 characters in length.

Enable Encryption

This checkbox controls whether all tunnel traffic will be encrypted.

- If checked, each packet will be digitally scrambled before sending. The receiving end of the tunnel will unscramble the data using a shared key before allowing the traffic onto its local network.

Encryption Method

This pull-down menu allows an encryption method to be specified.

- If **None** is selected, the tunnel session will be sent in the clear in both directions.
- If **Fixed** is selected, Personal Level Encryption will be used to scramble the data using a fixed key.
- If **PLE** is selected, Personal Level Encryption will be used to scramble the data using a key generated from the encryption secret.
- If **DES56** is selected, the DES algorithm will be used. DES provides better security than PLE, but also requires more time to operate.
- If **3DES** is selected, the Triple DES algorithm will be used. In Triple DES, the data is processed three times, each time with a different 56-bit key.

❖ **Note:** *Some VPN devices may not allow* **3DES** *as an option.*

Encryption Secret

This secret is used to generate session keys which are used to encrypt/decrypt each packet received from or sent through the tunnel. The secret can be from 1 to 255 characters in length.

❖ **Note:** **PLE**, **DES56** *and* **3DES** *all require that the same* **Encryption Secret** *be configured for each end of the tunnel.*

Interoperability Settings Dialog Box

This dialog box enables the IntraPort to interoperate with other vendors' devices. If the remote Tunnel Partner is a Compatible Systems device, it is not necessary to configure these settings. Interoperability settings are individually set for each tunnel partner.

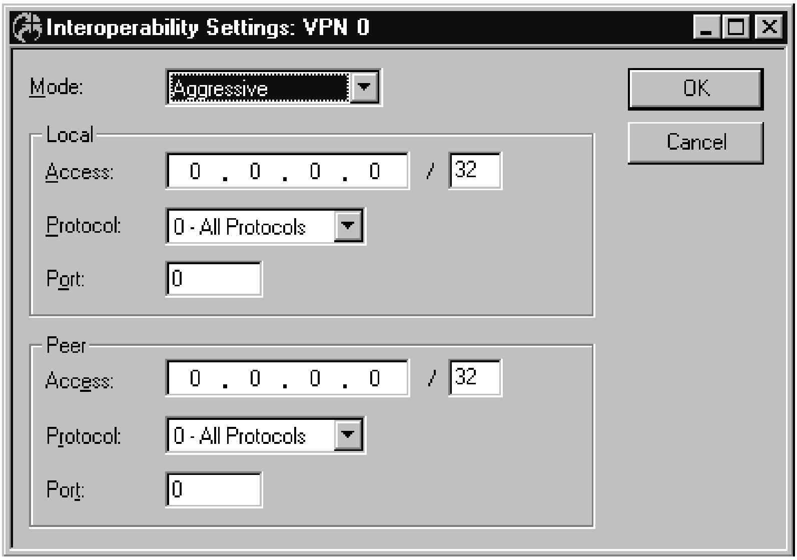

Interoperability Settings: VPN 0

Mode: Aggressive

Local

Access: 0 . 0 . 0 . 0 / 32

Protocol: 0 - All Protocols

Port: 0

Peer

Access: 0 . 0 . 0 . 0 / 32

Protocol: 0 - All Protocols

Port: 0

Interoperability Settings Dialog Box

To access this dialog box, select VPN Port #/Interoperability Settings from the device view.

Mode

This pull-down menu set the IKE Phase 1 negotiation mode between the devices. Phase 1 controls how the two devices identify and authenticate each other so that tunnel sessions can be established.

Main and **Aggressive** are the two IPSec standard methods for performing the Phase 1 negotiation. This setting must match the Phase 1 negotiation mode of the remote peer. Other vendors may support only the **Main** mode.

Local and Peer Settings

As part of their interoperability function, the following settings specify access *from* one area behind a VPN device *to* another area behind a VPN device.

The **Local** settings specify what local subnets, hosts, ports and/or protocols will be reachable via the tunnel.

The **Peer** settings specify what remote subnets, hosts, ports and/or protocols will be reachable via the tunnel. The remote tunnel partner (i.e., peer) must have a matching policy in order for traffic to be successfully tunneled.

Local / Access

This used to specify a local host or subnet which will be reachable by the tunnel. It is entered as an IP address followed by a slash followed by the number of significant bits in the entered IP address (i.e., 192.168.41.9/32). To allow access to only a single host, specify 32 in the bits portion.

Local / Protocol

The pull-down menu is used to specify an IP protocol which will accepted by this end of the tunneled. The default of 0 will allow all protocols. Accepted IP Protocol numbers are:

- •1 - ICMP (Internet Control Message Protocol)
- •6 - TCP (Transmission Control Protocol)
- •17 - UDP (User Diagram Protocol)
- •47 - GRE (Generic Routing Encapsulation)
- •50 - ESP (Encapsulating Security Protocol)
- •51 - AH (Authentication Header)
- •89 - OSPF (Open Shortest Path First)

Local / Port

The is used to specify a local port number which will be reachable via the tunnel. The default of 0 will allow all ports.

❖ **Note:** *Refer to the* **IP Filter Name** *section in the* ***Text Based Configuration and Command Line Management Reference Guide*** *for more information on commonly used ports and their numbers.*

Peer / Access

This is used to specify a host or subnet behind the remote tunnel partner which will be reachable via the tunnel. It is entered as an IP address followed by a slash followed by the number of significant bits in the entered IP address (i.e., 192.168.41.9/32). To tunnel to only a single host, specify 32 in the bits portion.

Peer / Protocol

This pull-down menu is used to specify an IP protocol which will be tunneled. If a protocol number is specified, then only traffic of that protocol type will be tunneled. The default of 0 will allow all protocols. Accepted IP Protocol numbers are:

- •1 - ICMP (Internet Control Message Protocol)
- •6 - TCP (Transmission Control Protocol)
- •17 - UDP (User Diagram Protocol)
- •47 - GRE (Generic Routing Encapsulation)

- 50 - ESP (Encapsulating Security Protocol)
- 51 - AH (Authentication Header)
- 89 - OSPF (Open Shortest Path First)

Peer / Port

This is used to specify a port number. If a Peer Port number is specified, then only traffic destined for that particular port will be tunneled. The default of 0 will allow all ports.

❖ **Note:** *Refer to the* **IP Filter Name** *section in the* ***Text Based Configuration and Command Line Management Reference Guide*** *for more information on commonly used ports and their numbers.*

Chapter 7 - VPN Client Tunnels

VPN Group Configuration Dialog Box

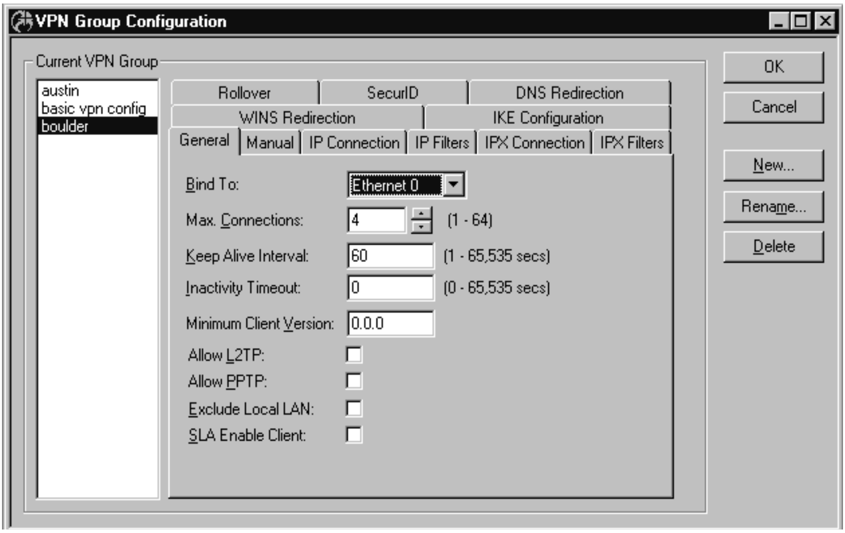

Current VPN Group: austin, basic vpn config, boulder

Bind To: Ethernet 0

Max. Connections: 4 (1 - 64)

Keep Alive Interval: 60 (1 - 65,535 secs)

Inactivity Timeout: 0 (0 - 65,535 secs)

Minimum Client Version: 0.0.0

Allow L2TP: □

Allow PPTP: □

Exclude Local LAN: □

SLA Enable Client: □

VPN Group Configuration Dialog Box and General Tab

To access this dialog box, select VPN Group Configuration from the Device View.

This dialog box displays and allows editing of all VPN Group Configurations for an IntraPort VPN Access Server. VPN group configurations define tunneling profiles for a group of one or more IntraPort users

The following table lists the maximum number of VPN group configurations allowed per device type.

Device Type	Maximum Number of VPN Groups
IntraPort 2	16
IntraPort 2+	100
IntraPort Enterprise-2 IntraPort Carrier-2 IntraPort Enterprise-8 IntraPort Carrier-8	1,000

> **Current VPN Group**

This edit box allows a VPN group configuration to be selected. Any changes made in the tab windows will be stored to the selected group configuration.

> **New**

Clicking on this button will bring up a dialog box which allows the creation of a new group configuration.

Rename

Clicking on this button will bring up a dialog box to allow the currently selected VPN group configuration to be renamed.

Delete

Clicking on this button will delete the presently selected group configuration.

VPN Group Configuration General Tab

> **Bind To**

This pulldown allows the selection of an interface on the device. The interface selected will act as the local end point for the tunnels defined by this configuration.

❖ **Note:** *Some VPN devices may only have one interface available for this function. In this case, the pull-down will not have any other choices available.*

Max Connections

This is the maximum number of simultaneous client connections using this configuration which will be allowed by the device.

This setting can be used to limit the number of connections for certain classes of users by assigning users to different group configurations. This number may not exceed the maximum number of tunnel connections supported by the device. If the sum of the **Max Connections** for all VPN Group sections exceeds the maximum number of tunnel connections supported by the device, tunnel connections will be served on a first-come, first-served basis.

Keep Alive Interval

This is the number of seconds between keep-alive packets sent to each connected client by the device.

Clients which do not answer these packets and/or generate other traffic within several keep-alive intervals will have their connections shut down.

Keep-alive packets are only sent in the case where no other traffic has been received from the client in the specified number of seconds.

Inactivity Timeout

This is the number of seconds the device will wait without receiving any traffic from a client belonging to this VPN Group configuration before ending the tunnel session.

Keep-alive packets and ICMP (ping) traffic do not affect this timeout. This prevents users from using ping to keep their tunnels up. The range is 1 to 65535 seconds. The default of 0 seconds means there is no timeout.

MinimumVersion

This places a limit on the VPN Client Software version number which will be allowed.

- A value of **0** or **1** will allow any software version number.
- A value of **2** will prevent Compatible’s older STAMP Clients from having access.
- A value of **3** will prevent both older STAMP Clients and any other Clients with version numbers less than 3.0.
- A value greater than three will prevent all clients from having access.

❖ **Note:** *The Allow L2TP and PPTP boxes should be left unchecked. These protocols are not currently implemented.*

Exclude Local LAN

This checkbox specifies whether remote client LAN traffic will be tunneled.

- If checked, remote LAN traffic will not be tunneled when a wildcard of 0.0.0.0/0 has been used as the Local IPNet. (specified on the IP Connection Tab.)

❖ **Note:** *The user login in the VPN Client software must also have the Exclude Local LAN from Tunnel checkbox checked.*

SLA Enable Client

This checkbox specifies that Service Level Agreement (SLA) information will be gathered for tunnel sessions using this VPN Group Configuration. SLA measures the speed of traffic across the tunnel and can be used to ensure that service guarantees are met.

SNMP is used to display the gathered information. This requires that SNMP be enabled in the Advanced SNMP Dialog Box. Refer to ***Chapter 14 - General*** for more information on SNMP Configuration.

VPN Group Configuration IKE Configuration Tab

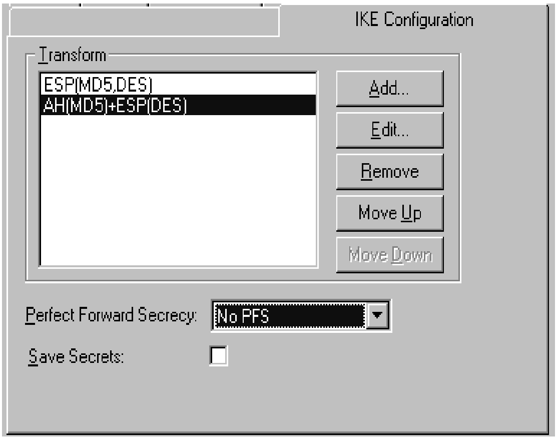

Transform:
ESP(MD5,DES)
AH(MD5)+ESP(DES)

Perfect Forward Secrecy: No PFS

Save Secrets: □

VPN Group Configuration IKE Configuration Tab

Transform

This specifies the protection types and algorithms that will be used for IKE tunnel sessions for this group configuration. Each option is a protection piece which specifies authentication and/or encryption parameters.

Use the **Move Up** and **Move Down** buttons to arrange the priority of the protection options.

> Perfect Forward Secrecy

Perfect Forward Secrecy (PFS) allows you to add an additional security parameter to tunnel sessions. PFS means that every time encryption and/or authentication key are computed, a new Diffie-Hellman Key Exchange is included.

Diffie-Hellman Key Exchange uses a complex algorithm and public and private keys to encrypt and then decrypt tunneled data. Adding PFS to a tunneled session greatly increases the difficulty of finding the session keys used to encrypt a VPN session. It also means that even if the keys are somehow cracked, only a portion of the traffic is recoverable.

- If **No PFS** is selected, this security parameter will not be added for this group configuration.
- If **Phase 1 Group** is selected, the group used in Phase 1 of the IKE negotiation is used as the group for the PFS Diffie-Hellman Key Exchange.

This group is set (as G1 or G2) in the IKE Policy Dialog Box. The IKE Policy Dialog Box is discussed later in this chapter.

- If **DH Group 1** is selected, the Diffie-Hellman Group 1 algorithm will be used for the Diffie-Hellman Key Exchange.
- If **DH Group 2** is selected, the Diffie-Hellman Group 2 algorithm will be used for the Diffie-Hellman Key Exchange. Because larger numbers are used by the DH Group 2 algorithm, it is more secure than DH Group 1.

> Save Secrets

This checkbox allows all users assigned to this particular VPN Group Configuration to save their shared secret to disk.

- If checked, users in this Group will not be prompted for their secret after the first session.

To add, edit, or remove a Transform, you must access the IKE Configuration Dialog Box by selecting the **Add...**, **Edit...**, or **Remove...** buttons in the IKE Configuration Dialog Box.

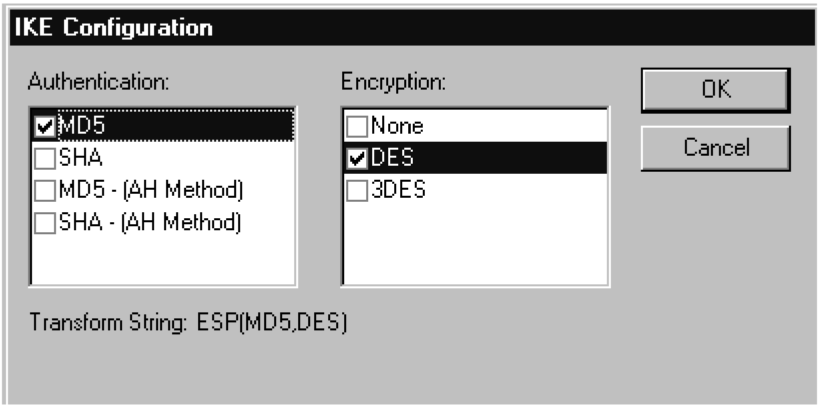

IKE Configuration

Authentication:

☑ MD5
□ SHA
□ MD5 - (AH Method)
□ SHA - (AH Method)

Encryption:

□ None
☑ DES
□ 3DES

Transform String: ESP(MD5,DES)

IKE Configuration Dialog Box

Authentication

This set of checkboxes specifies the authentication algorithm to be used for the tunnel session. MD5 is the Message-Digest 5 hash algorithm. SHA is the Secure Hash Algorithm.

Choosing either of the top two checkboxes means that the Encapsulating Security Payload (ESP) header will be used to encrypt and authenticate packets.

Choosing either of the bottom two checkboxes specifies that the Authentication Header (AH) will be used to authenticate packets.

Encryption

This set of checkboxes specifies the encryption algorithm to be used for the tunnel session.

DES (Data Encryption Standard) uses a 56-bit key to scramble the data. **3DES** uses three different keys and three applications of the DES algorithm to scramble the data.

❖ **Note:** *You may choose only one authentication and one encryption method. The default setting of ESP (MD5,DES) is recommended for most setups.*

VPN Group Configuration Manual Tab

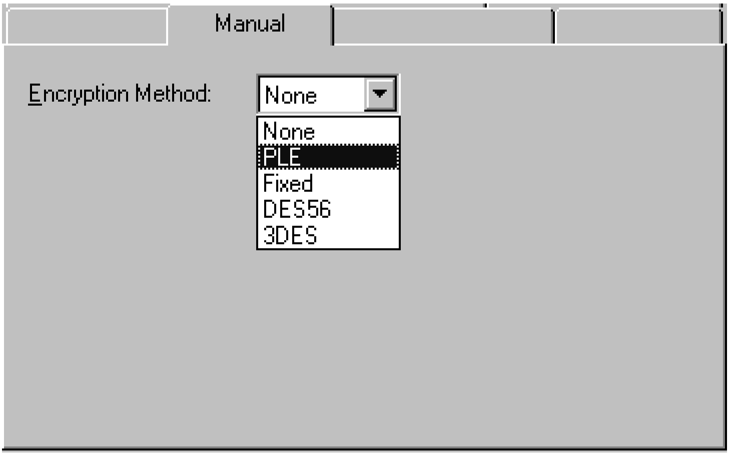

Encryption Method: None

VPN Group Configuration Manual Tab

> Encryption Method

This pull-down allows the selection of the encryption algorithm for non-IKE client sessions for this group configuration.

- If **None** is selected, the tunnel session will be sent in the clear in both directions.

- If **Fixed** is selected, Personal Level Encryption will be used to scramble the data using a fixed key.
- If **PLE** is selected, Personal Level Encryption will be used to scramble the data using a key generated from the encryption secret.
- If **DES56** is selected, the DES algorithm will be used. DES provides better security than PLE, but also requires more time to compute.
- If **3DES** is selected, the Triple DES algorithm will be used. In Triple DES, the data is processed three times, each time with a different 56-bit key.

❖ **Note:** **PLE, DES56** *and* **3DES** *all require that an Encryption Secret be configured for each user in the VPN Users dialog box. Some VPN devices may not allow* **3DES** *as an option.*

VPN Group Configuration IP Connection Tab

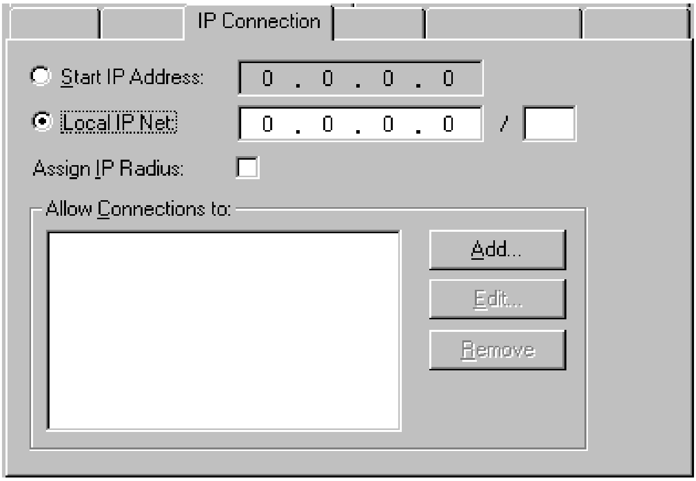

○ Start IP Address: 0 . 0 . 0 . 0

● Local IP Net: 0 . 0 . 0 . 0 /

Assign IP Radius: □

Allow Connections to:

VPN Group Configuration IP Connection Tab

> Start IP Address

The Start IP Address specifies the first IP address to be assigned to client sessions under this configuration. This start address will be incremented by one for each new client session, until the **Max Connections** limit (specified using the General Tab) is reached. The IP address is freed when the client is finished.

Each of the addresses thus generated must be a valid, unique, and *unused* IP address. Also, these addresses must not conflict with any networks specified in other VPN Group configuration or with any other IP address within the server.

These addresses must be on the *internal* TCP/IP network (i.e., for an IntraPort 2/2+, on the same network as Ethernet 0 or a subinterface thereof)

❖ **Note:** *There is no default value for the Start IP Address or Local IP Net. In order for IP-in-IP tunneling to operate with this VPN Group configuration, a group of local IP addresses must be set. Use the Start IP Address, the Local IP Net, or configure a Radius server to serve the addresses (see* ***Assign IP Radius*** *below).*

> Local IP Net

This edit box sets the local network or subnet to be assigned to client sessions under this configuration. For each new client session, an available IP address from this network or subnet is assigned to that session, until the **Max Connections** limit (specified using the General tab) is reached. The IP address is freed when the client session is finished.

This network or subnet must be *unused* and completely unique in the IP network to which the IntraPort is connected (i.e., not part of any Class C network in use) and may not conflict with address ranges specified in other group configurations. The mask may be between 8 and 30 bits.

The address should be entered as four decimal numbers separated by periods (e.g. 198.238.9.1). The part of this address which identifies the network segment is determined by the size of the mask, specified in bits.

❖ **Note:** *If Local IP Net is selected, either a dynamic routing protocol or static routes must be configured into the controlling router (e.g., the firewall) in order for traffic to find the Local IP Net.*

Assign IP Radius

This checkbox specifies whether a RADIUS server can be used to assign IP addresses to VPN users.

- If checked, communication with a RADIUS server must be configured, and be set up to serve the IP addresses.
- If left unchecked, IP addresses will be assigned using the address pool specified by either the Start IP Address or the Local IP Net.

❖ **Note:** *For more information on RADIUS configuration, see* ***Chapter 14 - General****.*

> Allow Connections To

This scrolling list displays the IP networks which the client will be told are reachable via the tunnel.

Any communications with an address which is part of one of the networks in the list will be tunneled. Communications with any other addresses will occur normally, without tunneling.

> Add

Clicking on this button will bring up a dialog box which allows an IP network address and mask size to be entered.

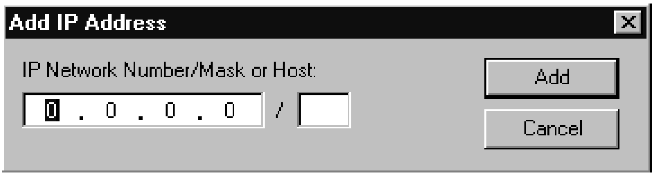

IP Network Number/Mask or Host:

0 . 0 . 0 . 0 /

Add IP Address Dialog Box

The part of this address which identifies the network segment is determined by the size of the mask, specified in bits. For example, an entry of 192.168.32.0/19 would specify that traffic with all IP addresses from 192.168.32.1 through 192.168.63.255 will be tunneled. As a special case, the entry, 0.0.0.0/0, specifies that all IP traffic should be tunneled. To tunnel to only a single host, specify 32 in the bits portion.

❖ **Note:** *Following convention, mask values in CompatiView are generally entered as decimal numbers separated by periods (e.g. 255.255.255.0) where both the network portion and the host portion of an address are significant. Entry of mask size in bits is an alternative but equivalent way of specifying the size of the network portion of the IP address when only that portion is significant to the function being performed.*

Edit

Clicking on this button brings up a dialog box which allows editing of a previously entered network and mask.

Remove

This button removes a network/mask entry from the list.

VPN Group Configuration IP Filters Tab

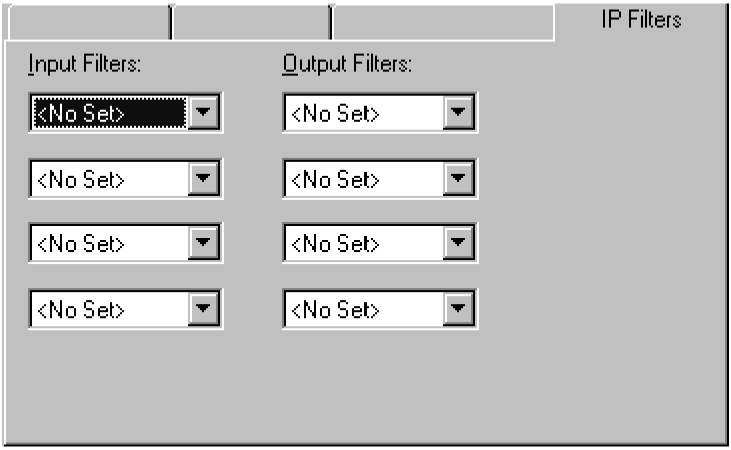

Input Filters:

<No Set>

<No Set>

<No Set>

<No Set>

Output Filters:

<No Set>

<No Set>

<No Set>

<No Set>

VPN Config IP Filters Tab

> **Input Filters**

These pulldowns allow the selection of previously created filter scripts which will be applied to tunnel packets coming into the device from users who are connected according to the selected configuration.

Up to four separate filters may be selected.

> **Output Filters**

These pulldowns allow the selection of previously created filter scripts which will be applied to tunnel packets sent out of the device to users who are connected according to the selected configuration.

Up to four separate filters may be selected.

❖ **Note:** *IP Filters are created using the TCP/IP Filter Editor. For more information on creating and editing IP Filters, refer to* ***Chapter 11 - TCP/IP Filtering****.*

VPN Group Configuration IPX Connection Tab

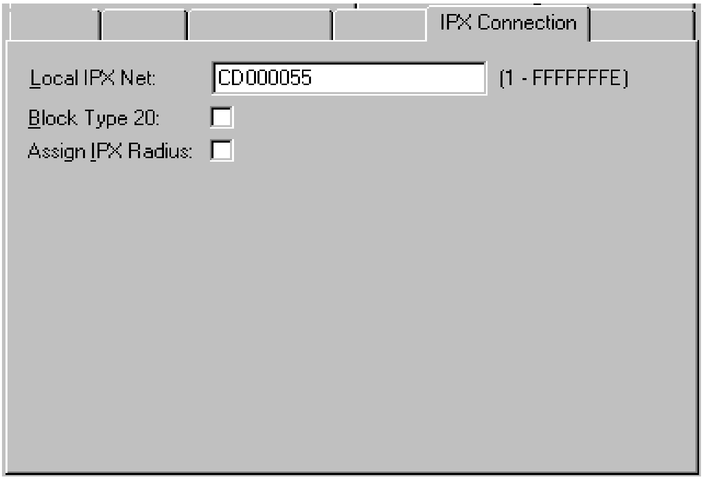

Local IPX Net: CD000055 (1 - FFFFFFFE)

Block Type 20: □

Assign IPX Radius: □

VPN Config IPX Connection Tab

> **Local IPX Net**

This edit box specifies the entry of the first local IPX network number to be assigned to client sessions under this configuration. This address will be incremented by one for each new client session, until the **Max Connection** limit (specified on the General Tab) is reached. When a client is connected to the device, the first available IPX address from this range is assigned to that session. The IPX address is freed when the client session is finished.

Each of the addresses thus generated must be a valid, unique and *unused* IPX address. Also, these addresses must not conflict with any networks specified in other VPN Group configurations or with any other IPX address within the server.

❖ **Note:** *There is no default value for the Local IPX Net. In order for IPX-in-IP tunneling to operate with this VPN Group configuration, a group of local IPX addresses must be set using either the Local IPX Net, or a RADIUS server must be configured to serve the addresses (see **Assign IPX Radius**).*

Block Type 20

In order for certain protocol implementations, like NetBIOS, to function in the NetWare environment, routers must allow a broadcast packet to be propagated throughout an internet. The IPX Packet Type 20 is designated to perform broadcast propagation for these protocols. This checkbox specifies whether IPX Packet Type 20 should be rebroadcast through the tunnel.

- If checked, IPX Packet Type 20 packets will not be rebroadcast during tunnel sessions. This is useful for reducing the bandwidth load on the tunnel.
- If left unchecked, these propagated packets will be rebroadcast during tunnel sessions.

Assign IPX Radius

This checkbox specifies whether a RADIUS server can be used to assign IPX addresses to VPN users.

- If checked, communication with a RADIUS server must be configured, and be set up to serve the IPX addresses.
- If left unchecked, IPX addresses will be assigned using the address pool specified by the Local IPX Net.

❖ **Note:** *For more information on RADIUS configuration, see* ***Chapter 14 - General****.*

VPN Group Configuration IPX Filters Tab

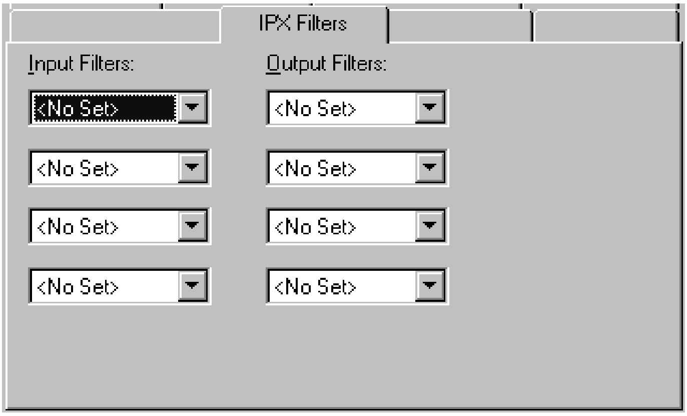

IPX Filters

Input Filters:

<No Set>

<No Set>

<No Set>

<No Set>

Output Filters:

<No Set>

<No Set>

<No Set>

<No Set>

VPN Group Configuration IPX Filters Tab

> **Input Filters**

These pulldowns allow the selection of previously created filter scripts which will be applied to tunnel packets coming into the device from users who are connected according to the selected configuration.

Up to four separate filters may be selected.

> **Output Filters**

These pulldowns allow the selection of previously created filter scripts which will be applied to tunnel packets sent out of the device to users who are connected according to the selected configuration.

Up to four separate filters may be selected.

❖ **Note:** *IPX Filters are created using the IPX Filter Editor. For more information on creating and editing IPX Filters, refer to* ***Chapter 11 - IPX Filtering****.*

VPN Group Configuration Rollover Tab

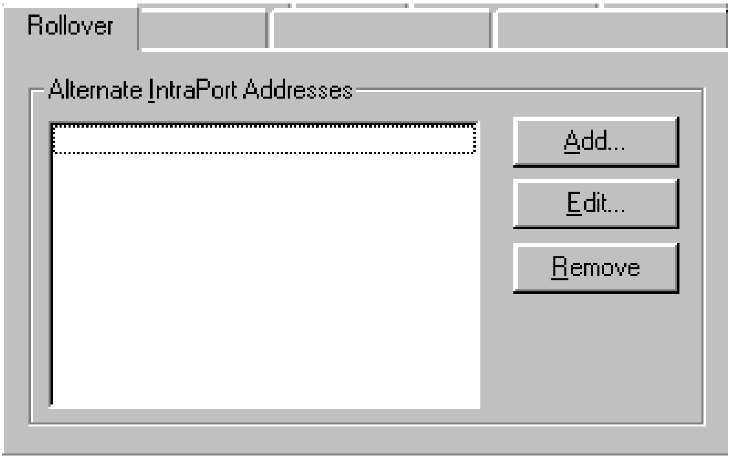

VPN Group Configuration Rollover Tab

Alternate IntraPort Addresses

This list displays all entered alternate IntraPort addresses. If an alternate address (or addresses) has been set, an IntraPort server which is full will be able to roll a client over to the specified alternate server. The IP address should be in standard dotted-decimal notation.

VPN Group Configuration SecurID Tab

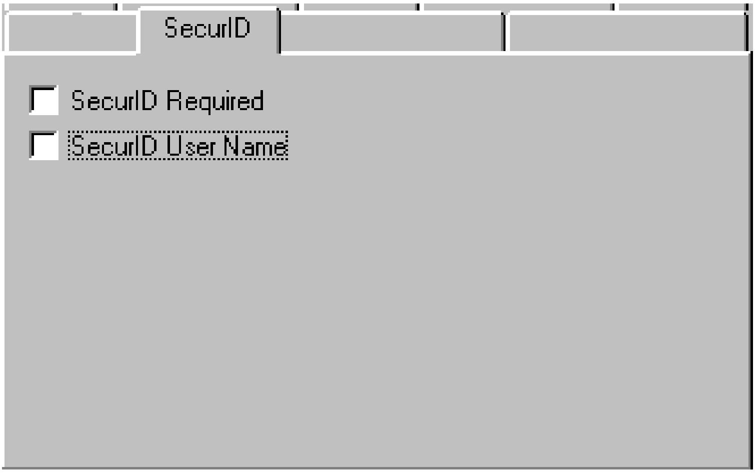

□ SecurID Required

□ SecurID User Name

VPN Group Configuration SecurID Tab

SecurID Required

Check this box to specify that all users assigned to this VPN Group configuration will undergo SecurID authentication. SecurID is Security Dynamic’s proprietary system which requires ACE/Server software and SecurID tokens to perform dynamic two-factor authentication.

SecurID User Name

Check this box if the VPN user name will also serve as the SecurID user name. If this box is checked, all users assigned to this VPN Group configuration will be prompted for their SecurID user name for authentication. If unchecked, the names for each user entered in the IntraPort and the ACE/Server must be the same.

VPN Group Configuration DNS Redirection Tab

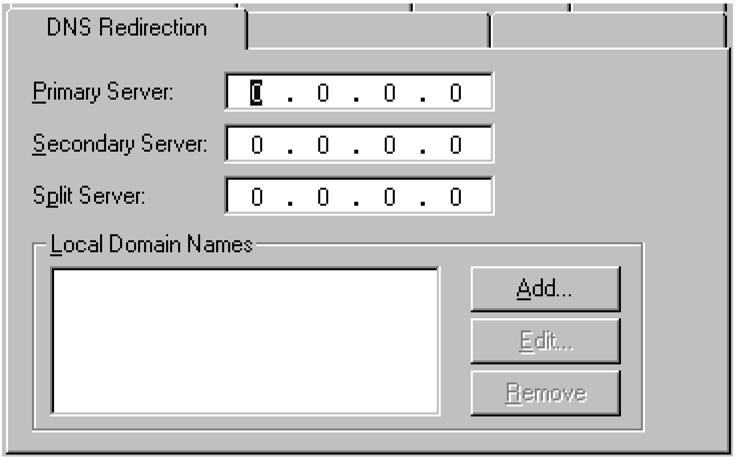

DNS Redirection

Primary Server: 0 . 0 . 0 . 0

Secondary Server: 0 . 0 . 0 . 0

Split Server: 0 . 0 . 0 . 0

VPN Group Configuration DNS Redirection Tab

Primary Server

The primary server specifies the primary IP address of a DNS server. If a Primary Server has been set, then the VPN Client will tunnel all DNS queries to the IntraPort and the IntraPort will take all DNS queries bound for the client’s primary DNS server and send them to the specified address. The IP address should be in standard dotted-decimal notation.

Secondary Server

The secondary server specifies the IP address of a backup DNS server. A primary server must be specified before a secondary server is chosen. The IP address should be in standard dotted-decimal notation.

Split Server

The split server specifies the IP address of a "split" DNS server. This is useful for setups where queries for internal names are handled by one server (the primary server) while queries for external names are handled by another server (the "split" server). The IP address should be in standard dotted-decimal notation.

Local Domain Names

This list specifies the domain names that will be compared to the name in DNS queries to the DNS server in order to determine whether the query is for an internal or external domain.

To add or modify the Local Domain Names, click on the appropriate button to access the Add Local Domain Dialog Box

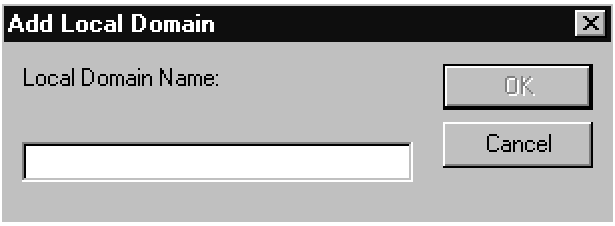

Add Local Domain Dialog Box

Local Domain Name

Local Domain Names can be between 1 and 255 characters in length.

VPN Group Configuration WINS Redirection Tab

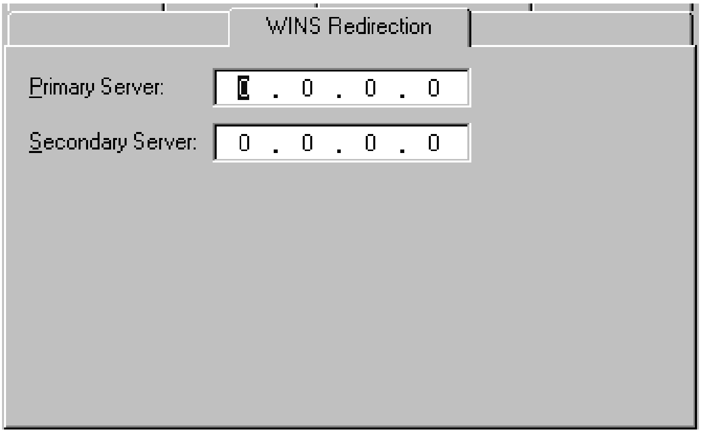

Primary Server: 0 . 0 . 0 . 0

Secondary Server: 0 . 0 . 0 . 0

VPN Group Configuration WINS Tab

Primary Server

The primary server specifies the primary IP address of a WINS server. If a Primary Server has been set, then the VPN Client software will tunnel all

WINS queries to the IntraPort and the IntraPort will take all WINS queries bound for the client's primary WINS server and send them to the specified address. The IP address should be in standard dotted-decimal notation.

Secondary Server

The secondary server specifies the IP address of a backup WINS server. A primary server must be specified before a secondary server is chosen. The IP address should be in standard dotted-decimal notation.

VPN User Configuration Dialog Box

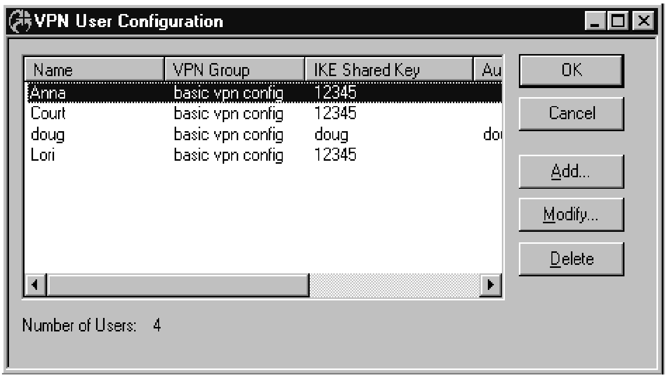

Name	VPN Group	IKE Shared Key	Au
Anna	basic vpn config	12345	
Court	basic vpn config	12345	
doug	basic vpn config	doug	do
Lori	basic vpn config	12345	

Number of Users: 4

VPN User Configuration Dialog Box

You can access the VPN User Configuration Dialog Box by selecting VPN User Configuration from the Device View. This dialog box displays all VPN users configured on an IntraPort VPN Access Server, but is not used to add or modify the entries.

To add or modify user entries, you must access the VPN User Dialog Box by selecting the **Add...** or **Modify...** buttons in the VPN User Configuration Dialog Box.

This user database is global to the device.

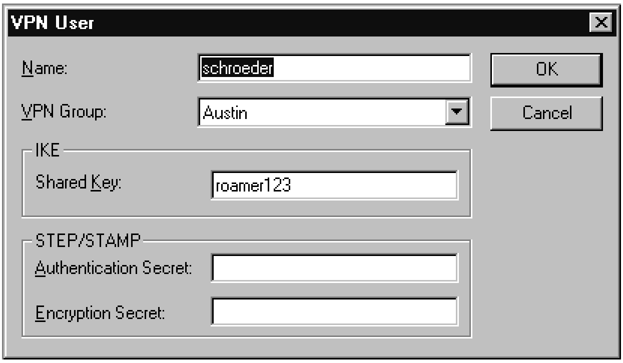

Name: schroeder

VPN Group: Austin

IKE

Shared Key: roamer123

STEP/STAMP

Authentication Secret:

Encryption Secret:

VPN User Dialog Box

> Name

This is the name of a user who will connect to the device using VPN client software.

> VPN Group

The user whose name is entered in the first column will be given the privileges and session parameters described in the specified VPN Group Configuration.

Any number of user entries may specify the same VPN Group Configuration in the database, but the VPN Group Configuration itself may allow a limited number of simultaneous users to actually have open VPN sessions with a device.

IKE Shared Key

This is a shared alphanumeric secret between 1-255 characters long. It is used to generate session keys which are used to authenticate and/or encrypt each packet received or sent through the tunnel.

STEP/STAMP Authentication Secret

This is a shared alphanumeric long term secret between 1-255 characters long. It is used to generate a series of short term keys which will authenticate traffic from this user on a packet-by-packet basis.

The same secret must be entered into the VPN client in order for authentication to succeed.

STEP/STAMP Encryption Secret

This is a shared alphanumeric long term secret between 1-255 characters long. It is used to generate a series of short term keys which will be used to encrypt/decrypt information to and from the user.

The same secret must be entered into the VPN client in order for encryption and decryption to succeed.

IKE Policy

This section is used to set the Internet Security Association Key Management Protocol/Internet Key Exchange (ISAKMP/IKE) parameters. These settings control how the IntraPort server and client will identify and authenticate each other. This initial negotiation is referred to as Phase 1.

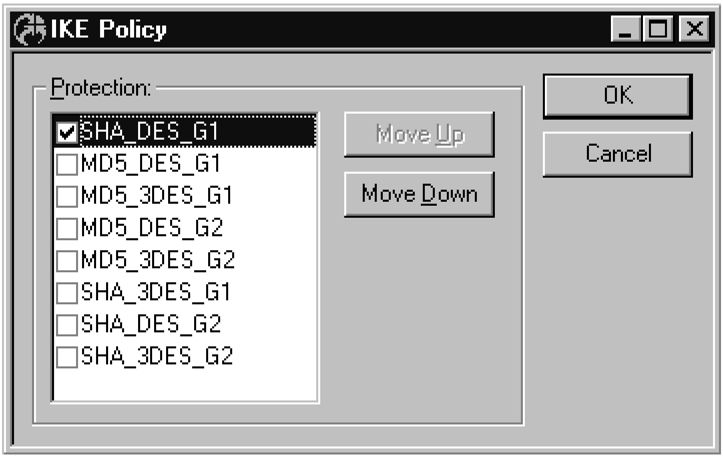

☑ SHA_DES_G1
□ MD5_DES_G1
□ MD5_3DES_G1
□ MD5_DES_G2
□ MD5_3DES_G2
□ SHA_3DES_G1
□ SHA_DES_G2
□ SHA_3DES_G2

IKE Policy Global Dialog Box

To access this dialog box, select Global/IKE Policy from the Device View.

The parameters set in this dialog box are global to the device and are not associated with a particular interface. These parameters specify a protection suite for the IKE negotiation between the IntraPort server and client. There are pieces to the IKE protection suite.

1. The first piece of each option is the authentication algorithm to be used for the negotiation. MD5 is the message-digest 5 hash algorithm. SHA is the Secure Hash Algorithm, which is considered to be somewhat more secure than MD5.

2. The second piece is the encryption algorithm. DES (Data Encryption Standard) uses a 56-bit key to scramble the data. 3DES uses three different keys and three applications of the DES algorithm to scramble the data.
3. The third piece is the Diffie-Hellman group to be used for key exchange. Because larger numbers are used by the Group 2 (G2) algorithm, it is more secure than Group 1 (G1).

Use the **Move Up** and **Move Down** buttons to arrange the priority of the protection suites.

❖ **Note:** *Phase 2 IKE negotiation sets how the IntraPort server and client will handle individual tunnel sessions. Phase 2 IKE negotiation parameters are set in the VPN Group Configuration Dialog Box, in the IKE Configuration Tab.*

IPSec Gateway Dialog Box

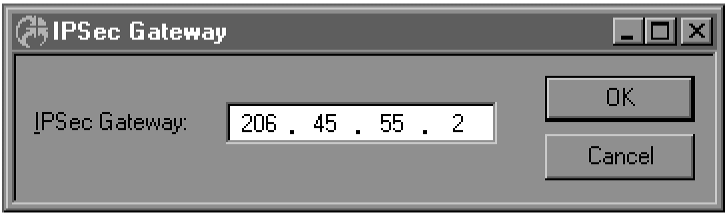

IPSec Gateway: 206 . 45 . 55 . 2

IPSec Gateway Configuration Dialog Box

To access this dialog box, select Global/IPSecGateway in the Device View.

> **IPSec Gateway**

This is the IP address that will be used as the gateway to the Internet for IPSec traffic on a dual-Ethernet IntraPort VPN Access Server. This is a required parameter only when the device is set to operate in parallel with your existing firewall (i.e. using both Ethernet ports) as the IPSec component of your security system.

The address should be entered as four decimal numbers separated by periods (e.g. 198.238.9.1).

❖ **Note:** *This IP address must be on the same IP network as the IPSec interface, which is configured using the IP Connection Dialog Box (under Ethernet/IP Connection on the IPSec port of an IntraPort VPN Access Server with two or more Ethernet interfaces).*

Chapter 8 - IntraGuard Firewall Configuration

There are three pre-set paths in the IntraGuard Firewall. A path defines a route for packets through the firewall. Each of the three paths already has a name, a security policy and interface definitions. While the names and parameters of the firewall paths can be modified, the default settings should work for many installations.

Firewall paths can be added to the edit area of a device, renamed or deleted.

To **Add** a path, right-click on any configuration item for the device, then select Firewall Path/Add Firewall Path from the popup menu.

❖ **Note:** *The IntraGuard Firewall currently supports up to three firewall paths. Any additional paths may cause configuration problems. It is recommended that you add firewall paths only if you have previously deleted a path, so that no more than three paths exist at a time.*

To **Rename** a path, right-click on the path’s icon, then choose Firewall Path/Rename Firewall Path

To **Delete** a path, right-click on the path’s icon, then choose Firewall Path/Delete Firewall Path.

These functions are also available in the **File** menu.

Settings: FirewallPath Dialog Box

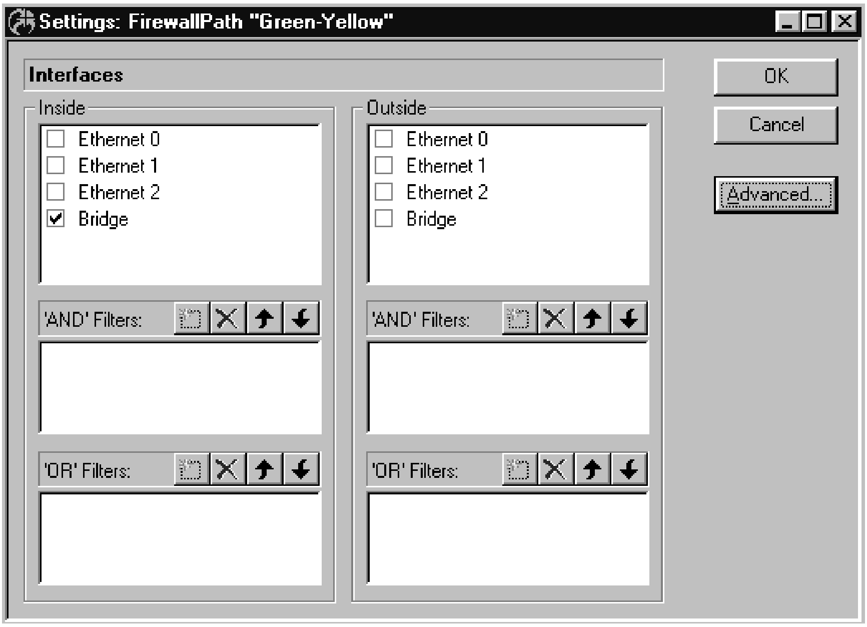

Settings: FirewallPath "Green-Yellow"

Interfaces

Inside

- ☐ Ethernet 0
- ☐ Ethernet 1
- ☐ Ethernet 2
- ☑ Bridge

'AND' Filters:

'OR' Filters:

Outside

- ☐ Ethernet 0
- ☐ Ethernet 1
- ☐ Ethernet 2
- ☐ Bridge

'AND' Filters:

'OR' Filters:

Settings: FirewallPath Dialog Box

New Button

Delete Button

Move Up Button

Move Down Button

To access this dialog box, select FirewallPath/Settings from the Device View.

Interfaces - Inside/Outside

These checkboxes control which interfaces will be specified as inside interfaces or outside interfaces for each path. Typically, **Inside** interfaces are secure while **Outside** interfaces are less secure.

If more than one interface is designated as an inside or outside interface on a particular path, those interfaces are considered to be open multiplexed and traffic will flow freely between them. For example, in the default configuration, both Ethernet 0 and the Bridge interface are inside interfaces on the Green-Red Path. Traffic between those two interfaces will not be subjected to firewall screening.

‘AND’ Filters

AND filters allow the device to accomplish packet filtering on packets that will be forwarded out the specified interface(s). AND filters are typically used to deny certain packets, so they are checked only for those protocols or ports which have been permitted by a Security Policy protocol setting, an Allow Ports/Protocol setting or an OR filter. Any packet not explicitly allowed by the rule set is dropped. Filters are created using the IP Filter Editor, described in the IP Filtering section of this manual. Up to four filter sets may be listed. The filters will be applied in the order listed.

Use the **New** button to add a named filter to the list or to select a named filter from a pull-down list.

Use the **Delete** button to remove a named filter from the list.

Use the **Move Up** and **Move Down** buttons to move the filters into the desired application order.

‘OR’ Filters

“OR” Filters allow the device to accomplish packet filtering on packets that will be forwarded out the specified interface(s). OR filters are typically used to permit certain packets, so they are checked only for those protocols or ports which have been denied by a Security Policy protocol setting or an Allow Ports/Protocol setting. Any packet not explicitly allowed by the rule set is dropped. Up to four filter sets may be listed. The filters will be applied in the order listed. Filters are created using the IP Filter Editor, described in the IP Filtering section of this manual.

Use the **New** button to add a named filter to the list or to select a named filter from a pull-down list.

Use the **Delete** button to remove a named filter from the list.

Use the **Move Up** and **Move Down** buttons to move the filters into the desired application order.

Advanced Settings: Firewall Path Dialog Box

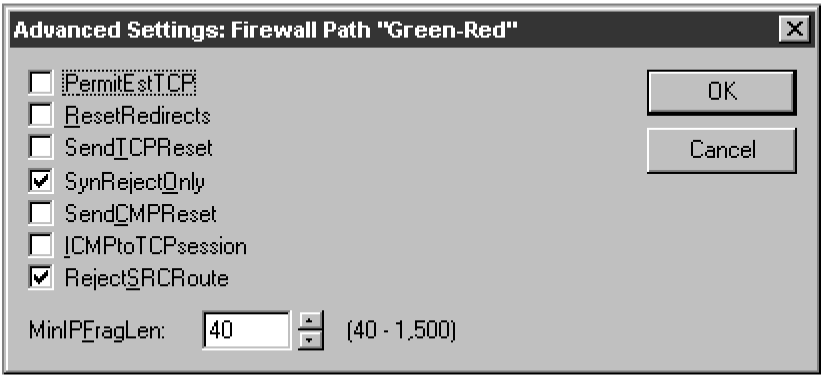
Advanced Settings: Firewall Path "Green-Red"

☐ PermitEstTCP
☐ ResetRedirects
☐ SendTCPReset
☑ SynRejectOnly
☐ SendCMPReset
☐ ICMPtoTCPsession
☑ RejectSRCRoute

MinIPFragLen: 40 (40 - 1,500)

Advanced Settings: Firewall Path Dialog Box

To access this dialog box, select FirewallPath/Settings from the Device View, then click on the **Advanced** button.

Advanced Options

These settings allow detailed control of how certain packet types and sessions will be handled on the path.

PermitEstTCP

This checkbox sets whether the path will permit TCP sessions for which the IntraGuard did not see the SYN flag. The SYN flag is included in the header of the first couple of TCP packets and indicates that a session is being established. When **checked**, this allows established connections to continue after rebooting the device, but it is also a less secure option. The default is unchecked.

ResetRedirects

This checkbox sets whether the device will terminate sessions on a firewall path where ICMP redirects have been sent. ICMP redirects are generated when a device cannot route a packet correctly on its own. The effect can be that three firewall path sessions will be created to route the packet correctly, two of which will not be needed after the first packet gets delivered. The default is unchecked.

SendTCPReset

This checkbox sets whether the device will send a TCP reset message to the client when a TCP session has been rejected. The default is unchecked.

SynRejectOnly

This checkbox sets whether the device will limit itself to sending TCP reset messages only when a TCP packet containing the SYN flag has been rejected. This can be useful when ICMP redirects are being sent, which could cause sessions to terminate prematurely. The default is checked.

SendICMPReset

This checkbox sets whether the device will send an ICMP message to the client when an IP or UDP packet has been rejected. The default is unchecked.

ICMPtoTCPsession

This checkbox sets whether the device will send an ICMP message to the client when a TCP packet has been rejected. This is in addition to sending a TCP reset message, if it has been enabled using the SendTCPReset checkbox. The default is unchecked.

RejectSRCRoute

This checkbox sets whether the device will reject source-routed IP packets. The default is checked.

MinIPFragLen

This field sets the minimum acceptable length of IP packets. Raising the minimum packet length can be useful in preventing "frag" attacks, which can take advantage of the use of partial header information in fragmented packets. The IntraGuard protects against overlapping fragmentation attacks, even when the MinIPFragLen is set to the minimum value of 40. Values may range between 40 and 1,500. The default is 40.

Security Policies: Firewall Path Dialog Box

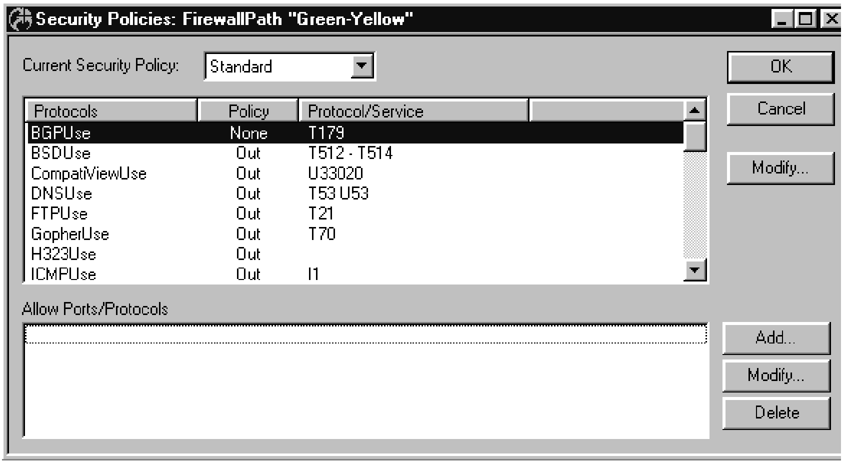

Security Policies: FirewallPath "Green-Yellow"

Current Security Policy: Standard

Protocols	Policy	Protocol/Service	
BGPUse	None	T179	
BSDUse	Out	T512 - T514	
CompatiViewUse	Out	U33020	
DNSUse	Out	T53 U53	
FTPUse	Out	T21	
GopherUse	Out	T70	
H323Use	Out		
ICMPUse	Out	I1	

Allow Ports/Protocols

Security Policies: Firewall Path Dialog Box

This dialog box can be accessed by selecting FirewallPath/Security Policies from the Device View. This dialog box displays the overall security policy for an IntraGuard Firewall path and the individual policy settings for each protocol. It can be used to change the overall security policy, but not the individual protocol policy settings. To change individual protocol settings, see the Security Policy Protocol Setting Dialog Box.

Current Security Policy

This pull-down menu sets the overall Security Policy for the path. There are five general policy sets, each of which has an associated list of protocol settings which define how the interfaces belonging to the path will handle those types of packets.

Definitions of the five sets of security policies follow:

- **Blocked** is the most secure policy set, which does not allow packets in or out along the path.
- **Strict** is a restrictive policy set. A small set of outgoing client sessions are permitted through the firewall and all incoming sessions are excluded.
- **Standard** is a moderately restrictive policy set. Almost all outgoing client sessions are permitted and almost all incoming server sessions are

excluded. The only exceptions to those rules are that the BPG and X Window protocols are excluded from going in or out along the path.

- **Lenient** is a less secure policy set. All outgoing client sessions are permitted and some incoming server sessions are permitted.
- **Open** is an insecure policy set. Everything is permitted through the firewall, thereby turning the firewall into a transparent bridge.

❖ **Note:** *Changing the Current Security Policy will override any individually made protocol settings.*

Security Policies at a Glance:

The following chart shows how each of the 31 protocols is treated by each of the five sets of security policies. The protocol BGPUse, for example, is assigned the security policy None by the Blocked policy set, but it is assigned the security policy Both by the Open policy set.

PROTOCOL	SECURITY POLICY				
	Blocked	**Strict**	**Standard**	**Lenient**	**Open**
BGPUse	None	None	None	Both	Both
BSDUse	None	None	Out	Out	Both
CompatiViewUse	None	Out	Out	Both	Both
DNSUse	None	Out	Out	Both	Both
FTPUse	None	Out	Out	Both	Both
H323Use	None	None	Out	Out	Both
ICMPUse	None	None	Out	Out	Both
IPSecUse	None	Out	Out	Both	Both
IRCUse	None	None	Out	Out	Both
LPRUse	None	None	Out	Out	Both
MailUse	None	Out	Out	Both	Both
NFSUse	None	None	Out	Out	Both
NetBIOSUse	None	None	Out	Out	Both
NewsUse	None	None	Out	Out	Both
NonIPUse	None	None	Out	Out	Both
OSPFUse	None	None	Out	Out	Both
POPUse	None	None	Out	Out	Both
RIPUse	None	None	Out	Out	Both
RealAudioUse	None	None	Out	Out	Both
SunRPCUse	None	None	Out	Out	Both
TelnetUse	None	Out	Out	Out	Both
TFTPUse	None	Out	Out	Out	Both
TunnelUse	None	None	Out	Out	Both
WebUse	None	Out	Out	Both	Both
XWinUse	None	None	None	In	Both
ISAKMPUse	None	Out	Out	Both	Both
GopherUse	None	Out	Out	Out	Both
NTPUse	None	None	Out	Both	Both
OtherTCPUse	None	None	Out	Out	Both
OtherUDPUse	None	None	Out	Both	Both
OtherUse	None	None	Out	Both	Both

Security Policy Protocol Setting Dialog Box

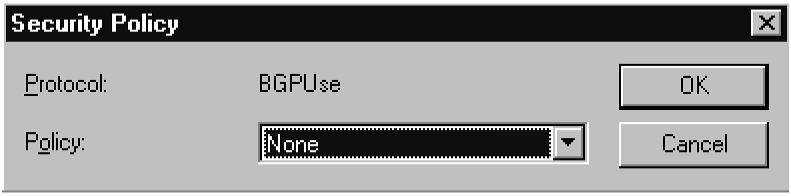

Protocol: BGPUse

Policy: None

Security Policy Protocol Setting Dialog Box

To change the individual protocol settings, select a protocol in the Security Policies: Firewall Path Dialog Box and then click the **Modify...** button. The Security Policy Dialog Box will appear in the Main Window.

❖ **Note:** *Changing the Current Security Policy will override any individually made protocol settings.*

Policy

This pull-down menu allows you to set how the selected protocol's packets will be handled on the path.

- **In** means that a protocol will be allowed through to the inside interface(s) of a path.
- **Out** means that a protocol will be allowed through to the outside interface(s) of a path.
- **None** means that a protocol will be allowed neither in nor out.
- **Both** means that a protocol will be allowed both in and out.

Protocols

- **BGPUse** defines how BGP (Border Gateway Protocol) packets will be handled on the path. BGP is the routing protocol between Internet backbone routers.
- **BSDUse** defines how BSD packets will be handled on the path. BSD is the UC Berkeley remote execution and terminal session protocol. RSH, RCP, RLogin, and RExec are the protocols supported.
- **CompatiViewUse** defines how CompatiView packets will be handled on the path. CompatiView is Compatible System's GUI manager. This option also defines handling for earlier versions of STAMP, Compatible System's tunnel authentication protocol.

- **DNSUse** defines how DNS (Domain Name Service) packets will be handled on the path. DNS is the protocol which translates IP addresses into hostnames and hostnames into IP addresses.
- **FTPUse** defines how FTP (File Transfer Protocol) packets will be handled on the path. Dynamic sessions are created for file transfers using the PASV and PORT commands.
- **H323Use** defines how H323 packets will be handled on the path. H323 is a video and audio conferencing protocol.
- **IPSecUse** defines how IPSec (Internet Protocol Security) packets will be handled on the path. Both encrypted (ESP) and authenticated (AH) packets are supported.
- **IRCUse** defines how IRC (Internet Relay Chat Protocol) packets will be handled on the path.
- **LPRUse** defines how LPR packets will be handled on the path. LPR is a network printing protocol.
- **MailUse** defines how SMTP (Simple Mail Transfer Protocol) packets will be handled on the path. This protocol is used to send mail between servers.
- **NFSUse** defines how NFS (Network File Sharing Protocol) packets will be handled on the path. To permit NFS In, it may be necessary to set SunRPCUse to In as well.
- **NetBIOSUse** defines how NetBIOS packets will be handled on the path. NetBIOS is Microsoft’s file sharing protocol.
- **NewsUse** defines how NNTP (Network News Transfer Protocol) packets will be handled on the path.
- **NonIPUse** defines how non-IP packets will be handled on the path. This would include other protocols such as AppleTalk and IPX.
- **OSPFUse** defines how OSPF (Open Shortest Path First) packets will be handled on the path. OSPF is a link state routing protocol.
- **POPUse** defines how POP packets will be handled on the path. POP is a mail client protocol. This protocol allows users to receive mail.
- **RIPUse** defines how RIP (Routing Information Protocol) packets will be handled on the path.
- **RealAudioUse** defines how Internet Real Audio Protocol packets will be handled on the path. Real Audio is an audio and video conferencing protocol.

- **SunRPCUse** defines how SunRPC (Sun’s Remote Procedure Call Protocol) packets will be handled on the path. The SunRPC Protocol is used by NFS and other UNIX utilities to get the server’s port address.
- **TelnetUse** defines how Telnet packets will be handled on the path. Telnet is a virtual terminal protocol.
- **TFTPUse** defines how TFTP (Trivial File Transfer Protocol) packets will be handled on the path.
- **TunnelUse** defines how GRE (General Router Encapsulation) packets will be handled on the path. GRE packets are IP-encapsulated tunneled packets. This option does not work with non-STEP tunnels (e.g. STAMP tunnels), which are enabled using the CompatiViewUse protocol.
- **WebUse** defines how HTTP (Hypertext Transfer Protocol) packets will be handled on the path. HTTP is the World Wide Web protocol. This option affects only HTTP packets; Telnet and FTP must be enabled individually to allow users to reach FTP sites or Telnet via the web. See the TelnetUse and FTPUse protocols.
- **XWinUse** defines how X Windows packets will be handled on the path. X Windows is the UNIX GUI.
- **GopherUse** defines how Gopher packets will be handled on the path. Gopher is a file transfer and browsing protocol.
- **ISAKMPUse** defines how ISAKMP (Internet Security Association Key Management Protocol) packets will be handled on the path. ISAKMP is the VPN (Virtual Private Network) key management protocol used by Compatible’s VPN products.
- **NTPUse** defines how NTP (Network Time Protocol) packets will be handled on the path.
- **OtherTCPUse** defines how all other TCP-based protocols will be handled on the path.
- **OtherUDPUse** defines how all other UDP-based protocols will be handled on the path.
- **OtherUse** defines how IP packets which are not included in the other pushbutton options will be handled on the path.

Allow Ports/Protocols Dialog Box

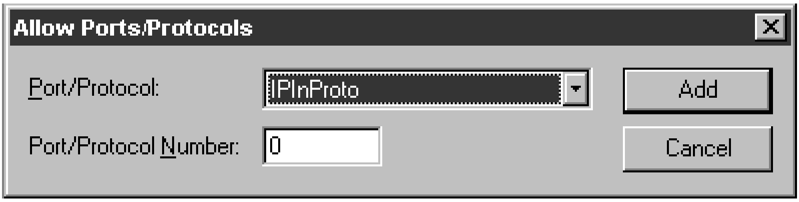

Port/Protocol: IPInProto

Port/Protocol Number: 0

Security Policy Protocol Setting Dialog Box

To access the Allow Ports/Protocols Dialog Box, select the **Add...** button to the right of the Allow Ports/Protocols list in the Security Policies: Firewall Path Dialog Box.

This dialog box allows you to specify a handling method for any numbered port or named protocol which isn’t already an explicit Security Policy option. All Security Policy protocol settings take precedence over the Allow Ports/Protocols options. For example, if the **OtherTCPUse** option is set to In in the Security Policy settings, then it would be unnecessary to specify any particular TCP port using the **TCPInPort** option below.

Port/Protocol

- The **TCPInPort** option specifies that a TCP port number will be allowed in along the path.
- The **TCPOutPort** option specifies that a TCP port number will be allowed out along the path.
- The **UDPInPort** option specifies that a UDP port number will be allowed in along the path.
- The **UDPOutPort** option specifies that a UDP port number will be allowed out along the path.
- The **IPInProto** option specifies that an IP protocol will be allowed in along the path.
- The **IPOutProto** option specifies that an IP protocol will be allowed out along the path.

Port/Protocol Number

The port or protocol number must be specified as a decimal number between 0 and 65,535. RFC 1700 "Assigned Numbers" contains a listing of all currently assigned IP protocol numbers.

Firewall Logging Dialog Box

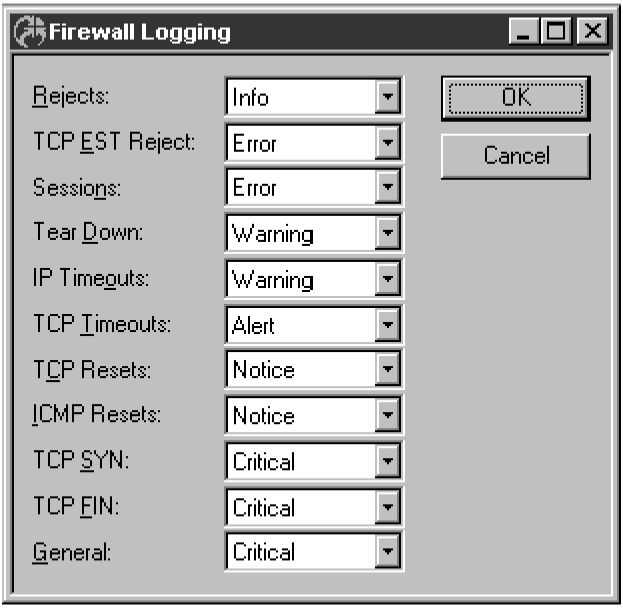

Rejects: Info

TCP EST Reject: Error

Sessions: Error

Tear Down: Warning

IP Timeouts: Warning

TCP Timeouts: Alert

TCP Resets: Notice

ICMP Resets: Notice

TCP SYN: Critical

TCP FIN: Critical

General: Critical

Firewall Logging Dialog Box

To access this dialog box, select Global/Firewall Logging from the Device View.

The logging settings define the level at which specific events are logged. The nine logging levels are listed below in descending order of importance.

- Off
- Emergency
- Alert
- Critical
- Error
- Warning
- Notice
- Info
- Debug

The IntraGuard “tags” the log messages associated with each type of event with the specified log level. The **Off** setting will disable log messages for the event.

The event log messages will appear in the log buffer (or wherever log messages are being sent), only if the global log level is at the same level or a lower level of importance. This allows you to closely monitor certain events while excluding events you do not wish to closely monitor from the log.

Logging parameters for the device, including the global log level, are set in the Logging Configuration Dialog Box, which can be accessed by selecting Logging from the Device View.

Using the default configuration as an example, if you wish to see log messages for TCP Resets, which have a default setting of Notice, you would need to set the **Log Level** in the Logging Configuration Dialog Box to Notice, Info or Debug. Any other setting would mean that TCP Resets would not appear in the log.

Rejects

Rejects messages are created by the firewall whenever an IP packet is rejected for any reason. The default is Info.

TCP EST Reject

TCP EST Reject messages are created by the firewall whenever an established TCP session is rejected. These messages are also created when a TCP session for which the firewall has not seen the SYN flag is established. The default is Error.

Sessions

Sessions messages are created by the firewall whenever an IP session is established. The default is Error.

TearDown

TearDown messages are created by the firewall whenever an IP session is torn down. The default is Warning.

IP Timeouts

IP Timeouts messages are created by the firewall whenever a non-TCP session (i.e. IP or UDP session) is timed out. The default is Warning.

TCP Timeouts

TCP Timeouts messages are created by the firewall whenever a TCP session is timed out due to inactivity. The default is Alert.

TCP Resets

TCP Resets messages are created by the firewall whenever a TCP session is reset. The default is Notice.

ICMP Resets

ICMPResets messages are created by the firewall whenever a non-TCP session (i.e. UDP or ICMP session) is reset. The default is Notice.

TCP SYN

TCP SYN messages are created by the firewall whenever a TCP connection cannot be completed because it was timed out. The default is Critical.

TCP FIN

TCP FIN messages are created by the firewall whenever a TCP connection cannot be properly torn down and is instead timed out. The default is Critical.

Redirects

Redirects messages are created by devices on the network when they receive a misdirected packet. These messages sometimes indicate route instability or the presence of an incorrectly configured IP host, but they do not necessarily indicate a problem on the network. The default is Critical.

General

General messages are created when errors occur within the IntraGuard. This might include running out of memory or internal state errors, and should be infrequent. The default is Critical.

Firewall Settings Dialog Box

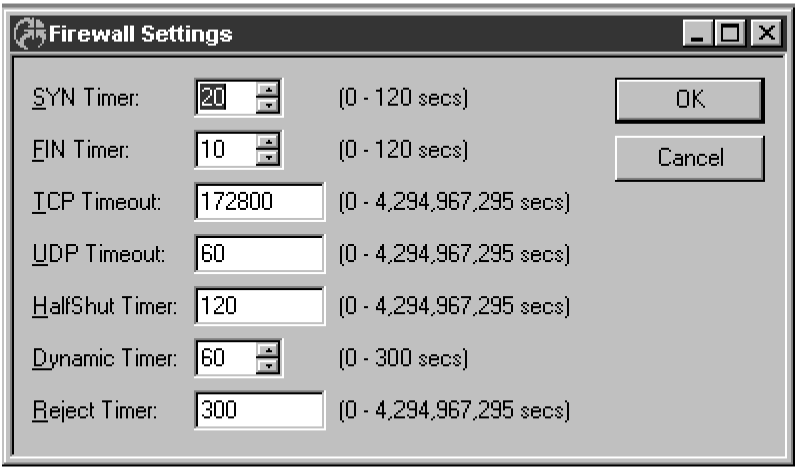

SYN Timer: 20 (0 - 120 secs)
FIN Timer: 10 (0 - 120 secs)
TCP Timeout: 172800 (0 - 4,294,967,295 secs)
UDP Timeout: 60 (0 - 4,294,967,295 secs)
HalfShut Timer: 120 (0 - 4,294,967,295 secs)
Dynamic Timer: 60 (0 - 300 secs)
Reject Timer: 300 (0 - 4,294,967,295 secs)

Firewall Settings Dialog Box

To access this dialog box, select Global/Firewall Settings from the Device View. The dialog box Firewall Settings appears on the Main Screen.

This dialog box is used to set global timers for the firewall.

SYN Timer

This field sets the number of seconds the firewall will wait without receiving a response to a SYN TCP packet before clearing a TCP session. The SYN flag is included in the header of the first couple of TCP packets and indicate that a session is being established. If the SYN Timer is set too low, half-open sessions may accumulate. If the SYN Timer is set too high, there may not be enough time to complete the handshake and establish a session. Values may range from 0 to 120. The default is 20 seconds.

FIN Timer

This field sets the number of seconds the firewall will wait without receiving a response to a FIN TCP packet before clearing a TCP session. TCP specifies that for a session to be fully closed down, both ends of the connection must send out a FIN packet. If the FIN Timer is too high, half-shut sessions may accumulate. If the FIN Timer is too low, sessions may be shut down too quickly. Values may range from 0 to 120. The default is 10 seconds.

TCPTimeout

This field sets the number of seconds the firewall will wait before shutting down an inactive TCP session. Values may range from 0 to 0xFFFFFFFF. The default is 172,800 seconds (48 hours).

UDPTimeout

This field sets the number of seconds the firewall will wait before shutting down an inactive non-TCP session. Values may range from 0 to 0xFFFFFFFF. The default is 60 seconds.

HalfShutTimer

This field sets the number of seconds the firewall will wait to close down a half-shut, inactive TCP session. TCP specifies that for a session to be fully closed down, both ends of the connection must send out a FIN packet. If the firewall has not received a FIN packet from the other end and there has been no activity during the specified length of time, the firewall will clear the session. Values may range from 0 to 0xFFFFFFFF. The default is 120 seconds. Setting a value of 0 will disable the timer.

DynamicTimer

This field sets the number of seconds the firewall will wait before shutting down an inactive dynamic session. Dynamic sessions are created by the firewall to allow TCP sessions or non-TCP packets to come through the firewall. The firewall does this by monitoring packet headers and data, and then opening permitted sessions only when necessary. Values may range from 0 to 300. The default is 60 seconds.

RejectTimer

This field sets the number of seconds the firewall will keep track of rejected packets after the packet flow has ended. The firewall tallies the different types of rejected packets and summarizes the information in a display using the **show firewall rejects** command (see **firewall(show)** in the ***Text-Based Configuration and Command Line Reference Guide***). Values may range from 0 to 0xFFFFFFFF. The default is 300 seconds. If the Reject Timer is set to 0, the firewall will log every rejected packet individually, without summarizing them in a tally.

Chapter 9 - Bridging

Global Bridging Configuration Dialog Box

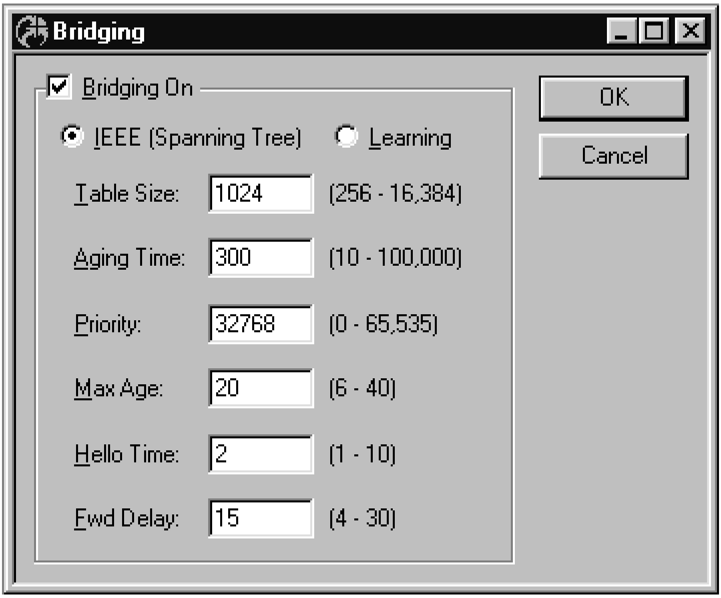

☑ Bridging On

● IEEE (Spanning Tree) ○ Learning

Table Size: 1024 (256 - 16,384)

Aging Time: 300 (10 - 100,000)

Priority: 32768 (0 - 65,535)

Max Age: 20 (6 - 40)

Hello Time: 2 (1 - 10)

Fwd Delay: 15 (4 - 30)

Global Bridging Configuration Dialog Box

❖ **Note:** *If you need more information about bridging, see “Bridging 101” in the Appendices to this manual.*

Bridging operates on physical network addresses (such as Ethernet addresses), rather than logical addresses (such as IP or IPX addresses). From the standpoint of routing, router interfaces which are set to bridge between themselves appear as a single logical entity.

Thus, a “Bridge Group” is made up of all of the physical network interfaces in a router which have been set to bridge the same protocol. Whether an interface bridges a protocol is set in the protocol configuration dialog box for each individual physical interface. See, for instance, the IP Routing/Bridging/Off radio buttons in the TCP/IP Routing: Ethernet Configuration Dialog Box.

Each Bridge Group can have routing parameters set for it. All of the interfaces in the group share these parameters.

To access the Main Bridging Configuration dialog box, select Global/Bridging from the Device View.

> Bridge On

This checkbox sets a global parameter which determines whether this router will perform bridging or not. Whether an individual interface actually participates in bridging is determined by settings for that interface.

- If **checked**, the router will bridge packets between interfaces which have the Bridging On checkbox set in their Interface Bridging Configuration dialog boxes. To access that dialog box select Interface/Bridging from the Device View. This can be done for any type of interface except IP subinterfaces.
- If **unchecked**, no packets will be bridged by the router.

Learning/IEEE (Spanning Tree)

This set of radio buttons determines which type of bridging will be performed.

- If **IEEE (Spanning Tree)** is selected, the bridge will act as an IEEE spanning tree bridge. It will send spanning tree BPDU (Bridge Protocol Data Unit) packets, and it will not forward any BPDU packets it receives. This is the default setting.
- If **Learning** is selected, the bridge will act as a simple learning bridge. This means care must be taken not to introduce loops in your network architecture. The bridge will not send spanning tree BPDU packets, and it will forward any BPDU packets it receives.

Table Size

This field sets the maximum number of entries allowed in the bridge table. The bridge table tracks the correspondence between physical interfaces and known addresses. The bridge will only allocate the amount of memory actually needed for the table. If the maximum size is too small, network traffic on all bridged segments will increase since the bridge will not know the specific segment for an address.

For the bridge to operate most efficiently, the table must be as large as the number of network nodes on all network segments which are physically connected to the router plus all network segments connected to the router through other bridges. Values may range from 256 to 16,384. The default value is 1024.

❖ **Note:** *Nodes on segments connected through routers which are not doing bridging do not need to be counted. This is because a router hides the physical addresses of the nodes behind it.*

Aging Time

This is the number of seconds since a node’s last transmission before its address will be removed from the bridge table. Values may range from 10 to 100,000. The default is 300.

Priority (Spanning Tree)

The spanning tree algorithm uses this value to help determine the Root Bridge for a network. The priority is combined with the bridge’s node address to create an eight byte bridge ID. The bridge which has the numerically lowest bridge ID on a network will become the Root Bridge for that network. There will only be one Root Bridge per network.

Values may range from 0 to 65,535. The default value is 32,768.

❖ **Note:** *Setting a bridge’s priority to 0 should make it the Root Bridge on the network.*

Max Age (Spanning Tree)

This parameter determines the maximum amount of time before the information from the last BPDU packet received is considered stale and the spanning tree is recalculated. Values may range from 6 to 40. The default value is 20 seconds.

❖ **Note:** *All bridges on a network use the **Max Age** value configured into the Root Bridge.*

Hello Time (Spanning Tree)

This parameter determines the amount of time between spanning tree BPDU packets sent by the bridge. Values may range from 1 to 10. The default is 2 seconds.

❖ **Note:** *All bridges on a network use the **Hello Time** value configured into the Root Bridge.*

Fwd Delay (Spanning Tree)

This parameter determines the amount of time between interface state transitions in the bridge. After startup the bridge spends this amount of time determining whether an interface should participate in the network’s spanning tree or be blocked. Once that decision has been made, the same amount of time is also spent learning addresses on the interface before forwarding is enabled.

Values may range from 4 to 30. The default value is 15 seconds.

❖ **Note:** *All bridges on a network use the **Fwd Delay** value configured into the Root Bridge.*

Bridging: Ethernet Configuration Dialog Box

Bridging: WAN Configuration Dialog Box

Bridging: VPN Configuration Dialog Box

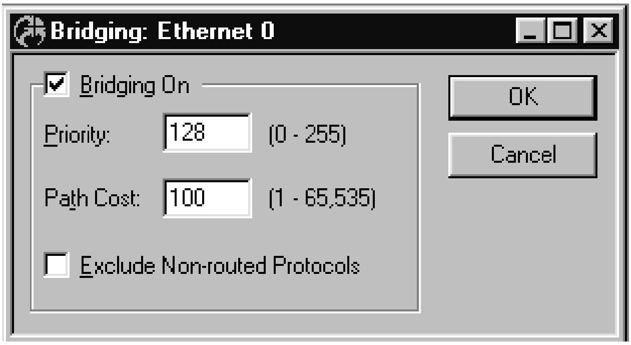

Bridging: Ethernet 0

☑ Bridging On

Priority: 128 (0 - 255)

Path Cost: 100 (1 - 65,535)

☐ Exclude Non-routed Protocols

Interface Configuration Dialog Box

❖ **Note:** *If you need more information about bridging, see “Bridging 101” in the Appendices to this manual.*

To access this dialog box, select Interface/Bridging from the Device View. This can be done for any type of interface except IP subinterfaces.

Bridging operates on physical network addresses (such as Ethernet addresses), rather than logical addresses (such as IP or IPX addresses). From the standpoint of routing, router interfaces which are set to bridge between themselves appear as a single logical entity.

Thus, a “Bridge Group” is made up of all of the physical network interfaces in a router which have been set to bridge the same protocol. Whether an interface bridges a protocol is set in the protocol configuration dialog box for each individual physical interface. See, for instance, the IP Routing/Bridging/Off radio buttons in the TCP/IP Routing: Ethernet Configuration Dialog Box.

Each Bridge Group can have routing parameters set for it. All of the interfaces in the group share these parameters. To access the dialog boxes which set these parameters, select Protocol/Bridging for the desired protocol (e.g. IP, IPX, etc.).

❖ **Note:** *WAN bridging is not recommended for ports set to On Demand PPP Link operation. Bridging requires that any broadcast traffic received on one Bridge Group port be resent on all other Bridge Group ports. The net effect is to keep on-demand links up all the time.*

❖ **Note:** *This CompatiView dialog box is only used to set the per interface values for bridging parameters. The majority of bridging parameters are set in the Main Bridging Configuration Dialog Box. To access this dialog box, select Global/Bridging from the Device View.*

> Bridging On

This checkbox determines whether this interface will perform bridging or not. In order for bridging to occur, the global Bridge On checkbox must also be set in the Main Bridging Configuration Dialog Box.

- If **checked**, the interface will participate in bridging. All non-routable protocols will be bridged. Routable protocols may or may not be bridged, depending on the setting in the individual protocol dialog box for this interface.
- If **unchecked**, no packets will be bridged by this interface.

Priority

This parameter determines the precedence given to this interface by the bridge. By default the lowest numbered interface (i.e. interface 0) will have precedence. Values may range from 0 to 255. The default is 128.

Path Cost (Spanning Tree)

This parameter sets the “cost” of using the interface, which in turn sets the calculated “distance” from the spanning tree’s Root Bridge. Distances are used when calculating spanning tree topology. This value may be used to artificially change the topology of a spanning tree network. Values may range from 1 to 65,535.

The default value for Ethernet ports is 100, as recommended by the IEEE. The default value for WAN ports is 5000.

Exclude Non-Routed Protocols

This checkbox determines whether this interface will bridge protocols which the router does not route. Examples are NetBEUI and DEC LAT.

- If **checked**, the interface will not bridge protocols that the router does not route.
- If **unchecked**, protocols which the router does not route will be bridged to all other interfaces which also have bridging turned on (and do not have this checkbox checked).

Chapter 10 - WAN Link Protocols

Link Configuration: WAN Dialog Box

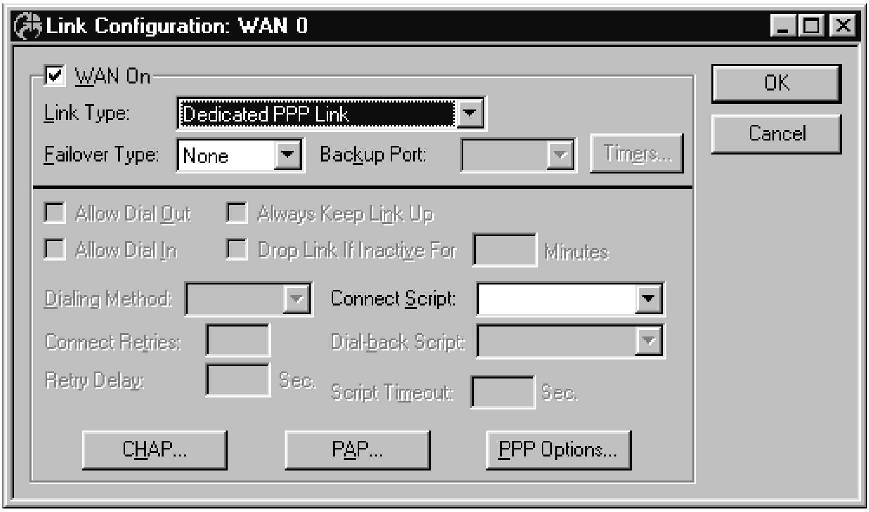

Link Configuration: WAN 0

☑ WAN On

Link Type: Dedicated PPP Link

Failover Type: None

Link Configuration: WAN Dialog Box

To access this dialog box, select WAN/Link Configuration from the Device View.

> **WAN On**

This checkbox controls how wide area network traffic is handled for this interface.

- If **checked**, then the interface will be active, link information can be configured with this dialog box, and network protocol configurations (TCP/IP Routing, IPX Routing, etc.) for the interface will take effect.
- If **unchecked**, then the interface will be inactive, no link information can be configured into this dialog box, and protocol configurations for the interface will not be in effect.

> Link Type

This pull-down menu determines how the router will maintain the WAN link, and sets the low-level communications protocol which will be used on the line connected to this interface.

- If **On Demand PPP Link** is selected, the router will treat the line connected to this interface as an intermittent "on-demand" connection which may require dialing commands to be issued. The router will use the Point-to-Point protocol to establish communications with the system at the other end of the line. Whether a connection can be initiated by this router, another router (or remote node client), or both, can be set using the **Allow Dial Out** and **Allow Dial In** checkboxes (as explained later in this chapter).
- If **Dedicated PPP Link** is selected, the router will treat the line connected to this interface as a connection which is always available regardless of traffic activity. The router will use the Point-to-Point protocol to establish communications with the system at the other end of the line.
- If **Frame Relay Link** is selected, the router will treat the line connected to this interface as a connection which is available regardless of traffic activity. The router will use the Frame Relay protocol to establish communications with the system (typically a Frame Relay switch) at the other end of the line.

❖ **Note:** *For **On Demand PPP Link** operation over RS-232C DIN-8 interfaces, Compatible Systems routers require that your communications device (modem, CSU/DSU, TA, etc.) be set to raise the DCD (data carrier detect) and/or DSR (data set ready) line when a connection is established, and drop it when the connection is terminated.*

❖ **Note:** *If an interface is set to **On Demand PPP Link**, there are certain maintenance packets for each protocol (IP, IPX, etc.) which will not cause an inactive connection to be dialed. This is a security measure that keeps intruders out and allows on-demand links to be useful.*

❖ **Note:** *The push buttons at the bottom of this dialog box will change depending on the choice you make for this pulldown.*

Failover Type

WAN ports can be set to divert their traffic to a secondary port (known as "failing over") if a line problem is detected. This pull-down menu determines the failover mode for this port.

Ports set for PPP operation will fail over if the PPP echo protocol determines that the line is down. Ports set for Frame Relay operation will fail over if the router stops receiving Frame Relay switch maintenance packets, or if all user PVCs go down.

- If **None** is selected, failover mode on the port will not be used.
- If **Primary** is selected, the router will monitor the status of the line connected to the port. If problems are detected on the line, traffic to this port will be diverted to the port selected with the **Backup Port** pull-down menu as described below.

❖ **Note:** *This pull-down menu will be disabled and will show “Backup” on a port which has been selected as a backup for a Primary port.*

Backup Port

When the port has been set as a Primary failover port using the **Failover Type** pulldown menu, this pulldown allows a backup port to be set. If the line on the primary port goes down, traffic will be diverted to the designated backup port.

Once a port has been selected as a backup port for one primary it cannot be used as a backup for another.

❖ **Note:** *This pull-down menu will be disabled, renamed to “Primary Port,” and will show the Primary port’s name on a port which has been selected to be a backup.*

Timers

This button brings up the **Failover Timers** screen, which controls the amount of time before traffic is diverted from the Primary to a backup port when a Primary’s line goes down, and the amount of time before traffic is diverted back to the Primary port when its line comes back up. The screen is described later in this chapter.

Allow Dial Out

This checkbox tells the router whether traffic forwarded from other interfaces on this router will cause an on-demand connection to be established on this interface. This checkbox can only be set if the Link Type is **On Demand PPP Link**.

- If **checked**, then incoming packets from another interface on this router whose destination is via this port will initiate a dialing sequence if the link is not already connected. If the link is already connected, the packets will simply be forwarded.
- If **unchecked**, then incoming packets from another interface on this router will be dropped if the link is not already connected.

Allow Dial In

This checkbox tells the router whether it should accept incoming on-demand PPP connections from other routers (or end-node clients). This checkbox can only be set if the Link Type is **On Demand PPP Link**.

- If **checked**, then incoming PPP connections will be accepted.
- If **unchecked**, then incoming PPP connections will be rejected.

Always Keep Link Up

This checkbox tells the router whether it should always initiate a dialing sequence if there is no connection established for this interface. This checkbox can only be set if the **Allow Dial Out** checkbox is checked and the **Drop Link If Inactive For** checkbox is unchecked.

- If **checked**, then whenever the connection for this interface is down, a dialing sequence will be initiated.
- If **unchecked**, then a dialing sequence will only be initiated when there is network traffic which needs to be forwarded out this interface.

Drop Link If Inactive For

This checkbox and edit box tell the router how long it should wait once all traffic has been forwarded across the connection before dropping the link. If additional traffic is forwarded from another interface on the router before the link has been dropped, the timer will be reset.

- If **checked**, then the link will be dropped after the specified number of minutes have passed with no packets being forwarded out this interface. The maximum value is 65535 minutes. The default value is 10 minutes.
- If **unchecked**, then the link will only be dropped when the router (or remote end-node client) drops its end of the connection.

❖ **Note:** *There are certain maintenance packets for each protocol (IP, IPX, etc.) which will not cause the inactivity timer to be reset. This is a security measure that keeps intruders out and allows on-demand links to be useful.*

Dialing Method

This pull-down menu lets you pick the dialing method which will be used for on-demand dialing on this interface. Which dialing method is used depends on the type of equipment being dialed. In general, asynchronous devices, such as modems, use **AT** style dialing. Synchronous devices, such as dialed CSU/DSU’s and ISDN terminal adapters, generally use **V.25bis** style dialing.

- The **AT** dialing specification is the industry standard for dialing modems. If you select this option, make sure you enter AT-style commands in the chat scripts you select as the **Dial-Out Script** and/or **Dial-back Script**.

- If you select **V.25bis** dialing, make sure you enter V.25bis-style commands in the chat scripts you select as the **Dial-Out Script** and/or **Dial-back Script**.

❖ **Note:** *Please check the manual for the communications device you are using to determine the best available dialing method for this interface.*

❖ **Note:** *Compatible Systems routers support "chat scripts" which let you provide a sequence of commands (using chat "send" statements), and anticipated responses (using chat "expect" statements) to devices which need to be dialed.*

Dial-Out / Connect Script

This pull-down menu selects the main chat script the router will run when attempting to initiate a connection.

You may choose any of the chat scripts which have been configured into the router. For more information on creating chat scripts, see the section on the Chat Script Editor Dialog Box later in this manual.

- If you selected the **On Demand PPP Link** pulldown discussed earlier, this pulldown will be labeled **Dial-Out**, and you must select a chat script here. The chat script you select will be executed whenever dialing is initiated.
- If you selected the **Dedicated PPP Link** or **Frame Relay Link** pulldown discussed earlier, this pulldown will be labeled **Connect**, and you may optionally select a chat script here. This script will be run when the router starts up and again whenever PPP communications are lost for some reason, and can be used to provide a set of required connect responses to a device (such as a terminal server) at the other end of the dedicated line.

Dial-back Script

This pull-down menu provides a way to select a chat script which will provide global dial-back security on incoming connections to this interface. This option can only be used if you have checked both the **Allow Dial Out** and **Allow Dial In** boxes discussed above.

You may use this menu to choose any of the chat scripts which have been configured into the router. For more information on creating chat scripts, see the section on the Chat Script Editor Dialog Box later in this manual.

- If you select a **chat script** here, the router will accept a PPP dial-in connection and then automatically drop the link and initiate dialing using the chat script you have selected.

- If you select **None** here, the router will not initiate a global dial-back on all incoming connections to this interface.

❖ **Note:** *You may still enforce dial-back security on selected connections by correctly setting the parameters in the User Authentication Database Dialog Box discussed later in this chapter.*

Dialing Retries / Connect Retries

Use this parameter to set the number of dialing retry attempts the router will make following an unsuccessful connection effort.

- If you selected the **On Demand PPP Link** pulldown discussed earlier, this field will be labeled **Dialing Retries**. This option can only be used if **Allow Dial Out** has been checked and a **Dial-Out Script** has been set.
- If you selected the **Dedicated PPP Link** or **Frame Relay Link** pulldown discussed earlier, this field will be labeled **Connect Retries**. On those types of links, this option can only be used if a **Connect Script** has been set.

Values may range between 1 and 255.

Retry Delay Setting

This parameter sets the amount of time in seconds the router will wait between dialing attempts.

- For an **On Demand PPP Link**, this option can only be used if **Allow Dial Out** has been checked and a **Dial-Out Script** has been set.
- For a **Dedicated PPP Link** or a **Frame Relay Link**, this option can only be used if a **Connect Script** has been set.

Values may range between 1 and 255.

Script Timeout

This is the amount of time in seconds the router will wait for input when it encounters an “Expect” statement in one of your chat scripts.

For more information on Expect statements and chat scripts in general, see the section on the Chat Script Editor Dialog Box later in this manual.

Failover Timers Configuration Dialog Box

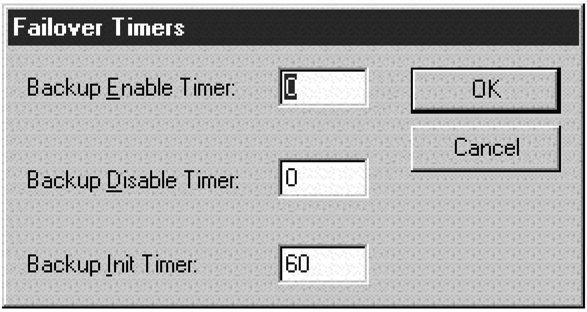

Backup Enable Timer: 0

Backup Disable Timer: 0

Backup Init Timer: 60

Failover Timers Configuration Dialog Box

You can access the Failover Timers Configuration Dialog Box by selecting **Primary** in the **Failover Type** pulldown in the Link Configuration: WAN Dialog Box (under WAN/Link Configuration), and then selecting the **Timers** button.

> **Backup Enable Timer**

This is the number of seconds from the time the Primary port's line is detected as being down until traffic is diverted to the Backup port. This is also known as the "failover time."

> **Backup Disable Timer**

This is the number of seconds from the time the Primary port's line is detected as having come back up until traffic is restored to the Primary port. This is also known as the "failback time" and is used to keep the router from switching out of failover mode too soon if the Primary link has an intermittent connection.

Backup Init Timer

This is the number of seconds after router startup before failover operation will go into effect. This timer allows PPP or Frame Relay communications time to stabilize before Primary port line status is checked.

Frame Relay Configuration Dialog Box

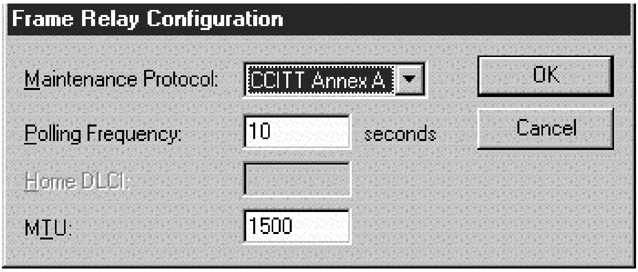

Frame Relay Configuration

Maintenance Protocol: CCITT Annex A

Polling Frequency: 10 seconds

Home DLCI:

MTU: 1500

Frame Relay Configuration Dialog Box

❖ **Note:** *If you need more information about the Frame Relay protocol, see “Frame Relay 101" in the Appendices to this manual.*

You can access the Frame Relay Configuration Dialog Box by selecting **Frame Relay Link** from the **Link Type** pulldown in the Link Configuration: WAN Dialog Box (under WAN/Link Configuration), and then clicking on the **Frame Relay** button at the bottom of the dialog box.

> Maintenance Protocol

This checkbox controls which Frame Relay maintenance protocol is used on this WAN interface. The maintenance protocol is used to send link status and virtual circuit information between Frame Relay switches and other devices (such as routers) that communicate with them.

- **ANSI Annex D** is the most commonly used standard in the United States.
- **CCITT Annex A** is a European standard.
- **LMI** was developed by a vendor consortium and is also known as the “consortium” management interface specification. It is still used by some carriers in the United States.
- **Static** allows the emulation of a Frame Relay network over WAN broadcast media. Examples include satellite ground stations and multipoint packet radio installations. Do not use this setting for normal Frame Relay switch communications.

❖ **Note:** *Your Frame Relay carrier may or may not give you a choice of management protocols. If you are given a choice, we suggest Annex D since it is the most widely used.*

> Polling Frequency

The router is required to periodically poll the Frame Relay switch at the other end of the communications link in order to determine whether the link is active. This field determines how often the router polls the switch, using the **Maintenance Protocol** you have selected.

If any three out of four polls go unanswered by the switch, the router will assume the Frame Relay link is down. Every sixth poll, the router requests a full status packet from the switch in order to update its table of active permanent virtual circuits (PVCs).

This value is in seconds. The allowable range for the value is 5 to 30. The default is 10.

Home DLCI

When Static maintenance is used on a WAN broadcast medium, this edit box can be filled in to provide a statically assigned DLCI (Data Link Control Identifier) number for this interface.

This number can be configured into other routers’ DLCI Mapping Dialog Boxes so that they can communicate with this router. In order to reject packets that were sent out its own interface, this router will ignore any packets with a sending DLCI number that matches this number.

MTU

This is the Maximum Transmission Unit in bytes for the interface. This setting may need to be adjusted in order to communicate with switches or routers from other vendors which do not support full size frame packets. The allowable range for the value is 262 to 1700. The default for this value is 1500.

❖ **Note:** *Adjusting the MTU to a smaller size will cause fragmentation of Frame Relay packets, which will impact performance. This setting should be left at the default unless it must be changed for compatibility reasons.*

DLCI Database Dialog Box

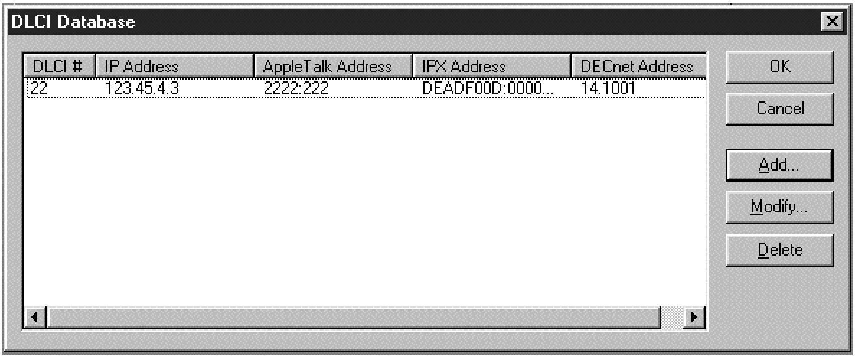

DLCI #	IP Address	AppleTalk Address	IPX Address	DECnet Address
22	123.45.4.3	2222:222	DEADF00D:0000...	14.1001

DLCI Database Configuration Dialog Box

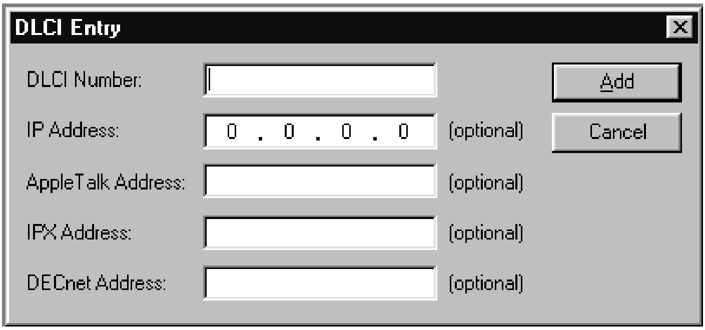

DLCI Number:		
IP Address:	0 . 0 . 0 . 0	(optional)
AppleTalk Address:		(optional)
IPX Address:		(optional)
DECnet Address:		(optional)

DLCI Entry Dialog Box

❖ **Note:** *If you need more information about the Frame Relay protocol, see “Frame Relay 101" in the Appendices to this manual.*

You can access the Frame Relay DLCI Database Dialog Box by selecting **Frame Relay Link** from the **Link Type** pulldown in the Link Configuration: WAN Dialog Box (under WAN/Link Configuration), and then clicking on the **DLCI** button at the bottom of the dialog box. This window displays all DLCI mapping entries, but is not used to add or modify the entries. To add or modify the entries, you must access the DLCI Entry Dialog Box by selecting the **Add...** or **Modify...** buttons in the Frame Relay DLCI Database Dialog Box.

The Data Link Connection Identifier (DLCI) is a number which uniquely identifies one end of a Permanent Virtual Circuit (PVC) to your Frame Relay carrier's Frame Relay switch. The DLCIs are not interchangeable between the two ends of a PVC, since they only identify one end of the PVC. Unless you use the correct DLCI numbers at each end of your PVC, two-way communications cannot take place.

This database lets you create static mappings between the Frame Relay PVCs on this interface (identified by their DLCI number) and the protocol (e.g. IP, IPX, etc.) addresses of the router interfaces at the far ends of the PVCs.

❖ **Note:** *If the router at the far end of a PVC is another Compatible Systems router, you will generally not need any entries in the DLCI database for it. Compatible Systems routers use the IARP (Inverse Address Resolution Protocol) to dynamically create the same type of mappings that are manually entered in the DLCI database.*

❖ **Note:** *A router will not use IARP to attempt to discover addresses for a particular protocol on a PVC if there is already a DLCI database entry for the PVC for that protocol. Therefore, if you wish to use IARP to dynamically discover the addresses at the far end of a PVC, do not make any entries for its DLCI number in the DLCI database.*

❖ **Note:** *Frame Relay DLCIs must be statically mapped using the DLCI mapping database when IP subinterfaces are in use, because IARP can only resolve a physical port, not a logical subinterface on that port.*

> DLCI

This is the decimal number between 16 and 991 which uniquely identifies this end of a PVC. A DLCI number will be provided to you by your Frame Relay carrier for each end of each PVC.

IP Address

This is the IP address of the router interface at the other end of the PVC. It should be entered in standard IP dotted-decimal notation (e.g. 198.041.9.1).

AppleTalk Address

This is the AppleTalk address of the interface of the router WAN interface at the other end of the PVC. It should be entered in decimal as a "network:node" pair (e.g. 24:1).

The AppleTalk network number must be between 1 and 65,279. The node address must be between 1 and 254.

IPX Address

This is the IPX address of the interface of the router WAN interface at the other end of the PVC. It should be entered in hexadecimal as a “network:node” pair (e.g. 12F0A:00A510123456).

The IPX network number must be between 1 and FFFFFFFE. The IPX node address must be 12 hexadecimal digits.

❖ **Note:** *The IPX node address at the other end is generally a “borrowed” Ethernet address from one of the other router’s Ethernet interfaces. There is no addressing conflict because the actual Ethernet interface is on a network with a different IPX network number.*

DECnet Address

This is the DECnet address of the router at the other end of the PVC. The address consists of a decimal “area.node” pair (e.g. 14.1001).

The area value must be within the range of 1 to 63. The node value must be within the range of 1 to 1023.

❖ **Note:** *A period is traditionally used as the separator for DECnet area:node pairs. Other protocols use a colon.*

CHAP Configuration Dialog Box

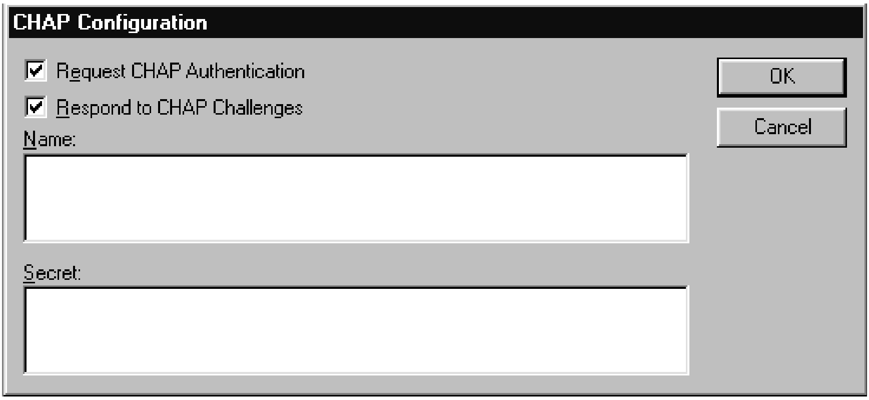

☑ Request CHAP Authentication

☑ Respond to CHAP Challenges

Name:

Secret:

CHAP Configuration Dialog Box

You can access the CHAP (Challenge Handshake Authentication Protocol) Configuration Dialog Box by selecting **On Demand PPP Link** or **Dedicated**

PPP Link from the **Link Type** pulldown in the Link Configuration: WAN Dialog Box (under WAN/Link Configuration), and then clicking on the **CHAP** button at the bottom of the dialog box.

CHAP is a security protocol that allows devices using PPP to authenticate their identities to each other through the use of a message digest (MD5) calculation. Either or both ends of a link can request that the opposite end of the link authenticate itself. CHAP requests do not depend on knowing which device initiated a call, so a calling device can request and/or provide authentication, as can a device that receives a call.

CHAP authentications can be performed at any time after a communications link is connected. A CHAP authentication sequence begins with a “challenge” from one end of the link. The challenge includes the name of the challenging router.

The response to the challenge includes the name of the responding router. This name will be looked up in the challenging router’s database or on a configured RADIUS server. The name, along with a “secret” value that is stored in the database or RADIUS server and is shared by both ends, will be processed by the challenging end using the MD5 algorithm.

If the result of an identical MD5 calculation performed by the challenging end is not the same, the challenging end drops the link.

To access the User Authentication Database Configuration Dialog Box, select Global/User Authentication Database in the Device View. To access the RADIUS Configuration Dialog Box, select Global/System Configuration in the Device View and click on the RADIUS button.

❖ **Note:** *Because the secret is never passed across the link, even in encrypted form, CHAP is considered to be significantly more secure than PAP.*

Request CHAP Authentication

This checkbox controls whether this router will send a CHAP challenge to the other end before allowing PPP negotiation to complete. Each challenge will include this router’s **Name** (as described below), along with a random value selected by this router.

- If **checked** this router will send a CHAP challenge to the device at the other end of the link.
- If **unchecked** this router will not send a CHAP challenge to the device at the other end of the link.

Respond to CHAP Challenges

This checkbox controls whether this router will respond to CHAP challenges from the other end.

- If **checked** this router will use the values in the **Name** and **Secret** fields to respond to a CHAP challenge from the other end.
- If **unchecked** this router will not respond to CHAP challenges.

Name

This is the name that the router will include in any CHAP challenges it makes, and in any CHAP responses it provides. A name is required if either **Request CHAP Authentication** or **Respond to CHAP Challenges** is checked. The name can be from 1 to 255 characters in length.

Secret

This is the shared information that is used to calculate expected CHAP responses to challenges issued by this router. A secret is required if **Respond to CHAP Challenges** is checked. The secret can be from 1 to 255 characters in length.

❖ **Note:** *CHAP functionality was changed in version 3.04 and higher of Compatible Systems' router software in order to allow for effective use of RADIUS servers. CHAP in versions 3.04 and later are not downward compatible with earlier versions.*

PAP Configuration Dialog Box

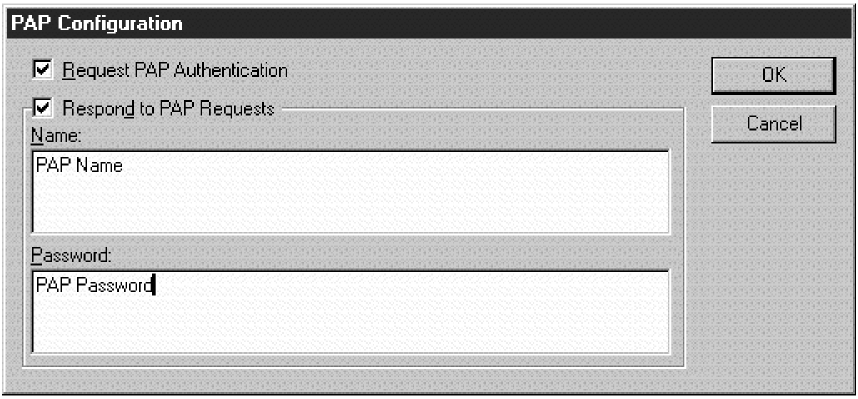

☑ Request PAP Authentication

☑ Respond to PAP Requests

Name:
PAP Name

Password:
PAP Password

PAP Configuration Dialog Box

You can access the PAP (Password Authentication Protocol) Configuration Dialog Box by selecting **On Demand PPP Link** or **Dedicated PPP Link** from the **Link Type** pulldown in the Link Configuration: WAN Dialog Box (under WAN/Link Configuration), and then clicking on the **PAP** button at the bottom of the dialog box.

PAP is a security protocol that allows devices using PPP to authenticate their identities to each other through the use of passwords. Either or both ends of a link can request that the opposite end of the link authenticate itself. PAP requests do not depend on knowing which device initiated a call, so a calling device can request and/or provide authentication, as can a device that receives a call.

PAP authentications are only performed after a communications link is connected, but before PPP has completely negotiated the communications parameters which will be used on the link. A PAP authentication sequence begins with a “PAP request” from one end of the link. The other end must respond with a valid name and password. If it does not, the requesting end drops the link.

❖ **Note:** *Because PAP passes the name and password values back across the link in “cleartext,” it is considered to be less secure than CHAP.*

Request PAP Authentication

This checkbox controls whether this router will request a PAP name and password from the other end before allowing PPP negotiation to complete.

All name/password combinations received are checked against the entries in the User Authentication Database, or in a configured RADIUS server.

To access the User Authentication Database Configuration Dialog Box, select Global/User Authentication Database in the Device View. To access the RADIUS Configuration Dialog Box, select Global/System Configuration in the Device View and click on the RADIUS button.

- If **checked** this router will request a PAP name and password from the device at the other end of the link. The name and password will be checked against all entries in the User Authentication Database or configured RADIUS server.
- If **unchecked** this router will not request a PAP name and password from the device at the other end of the link.

Respond to PAP Requests

This checkbox controls whether this router will supply a PAP name and password to the other end if they are requested.

- If **checked** this router will provide the name and password entered into the **PAP Name** and **PAP Password** edit areas on this screen when PAP information is requested by the device at the other end of the link.

- If **unchecked** this router will not provide any PAP information if it is requested by the device at the other end of the link.

Name

This is the name that the router will provide to the device at the other end if PAP name/password information is requested and the **Provide PAP Information** checkbox is checked. The name can be from 1 to 255 characters in length.

Password

This is the password that the router will provide to the device at the other end if PAP name/password information is requested and the **Provide PAP Information** checkbox is checked. The password can be from 1 to 255 characters in length.

SMDS Dialog Box

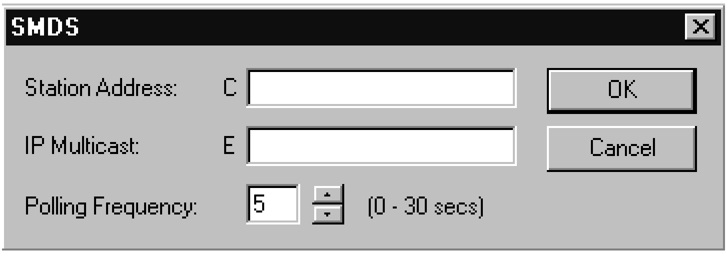

Polling Frequency: 5 (0 - 30 secs)

SMDS Dialog Box

You can access the SMDS Dialog Box by selecting **SMDS** from the **Link Type** pull-down in the Link Configuration: WAN Dialog Box (under WAN/Link Configuration), and then clicking on the **SMDS** button at the bottom of the dialog box.

Station Address

This is the SMDS physical station address. The address is assigned by the service provider and follows the E.164 format (i.e., 64-bit/15-digit addressing). The station address must start with the letter C and be followed by at least 10 digits.The missing digits will be filled in with F. The address should be entered exactly as it is assigned by the service provider.

IP Multicast

This is the IP multicast address. This address is the SMDS group address assigned by the service provider and follows the E.164 format. The multicast address must start with the letter E and be followed by at least 10 digits. The missing digits will be filled in with F. The address should be entered exactly as it is assigned by the service provider.

Polling Frequency

This number specifies the interval that the router uses to poll the SMDS switch. The interval is specified in seconds and must be between 0 and 30.

If the switch does not respond to the polling, the router will eventually declare the SMDS link down and start dropping packets designated for that interface. A value of 0 will disable the polling mechanism. Disabling the polling mechanism will automatically declare the SMDS link up.

❖ **Note:** *The keepalive mechanism is also referred to as "heartbeat exchange" in SMDS literature.*

PPP Options Dialog Box

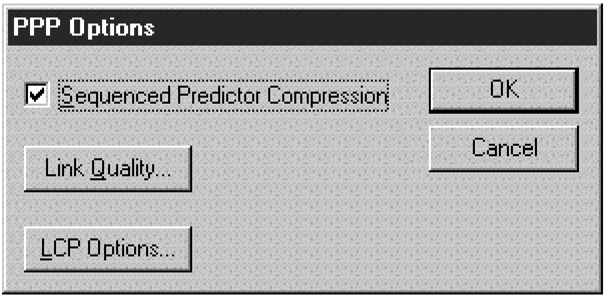

☑ Sequenced Predictor Compression

PPP Options Dialog Box

You can access the PPP Options Dialog Box by selecting **On Demand PPP Link** or **Dedicated PPP Link** from the **Link Type** pulldown in the Link Configuration: WAN Dialog Box (under WAN/Link Configuration), and then clicking on the **PPP Options** button at the bottom of the dialog box.

Sequenced Predictor Compression

Packet data can be compressed to provide better throughput across slower WAN links. Sequenced Predictor is a compression algorithm used in some Compatible Systems routers.

If **checked** this router will compress packet data being sent on this interface using the Sequenced Predictor algorithm.

❖ **Note:** *A general rule of thumb for Compatible Systems routers would be to use Sequenced Predictor on uncompressed links at up to 128K rates, but to turn it off at higher speeds or if other means of compression (such as the V.42 compression built into modems) are in use. A few simple file copy transfer tests over your particular WAN setup will yield a more exact answer.*

PPP Link Quality Configuration Dialog Box

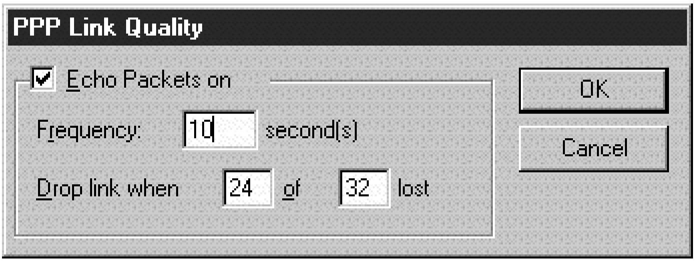

PPP Link Quality

☑ Echo Packets on

Frequency: 10 second(s)

Drop link when 24 of 32 lost

PPP Link Quality Configuration Dialog Box

You can access the PPP Link Quality Configuration Dialog Box by selecting **On Demand PPP Link** or **Dedicated PPP Link** from the **Link Type** pull-down in the Link Configuration: WAN Dialog Box (under WAN/Link Configuration), and then clicking on the **PPP Options** button at the bottom of the dialog box, and then clicking on the **Link Quality** button at the bottom of the PPP Options Dialog Box.

This dialog box is used to set parameters which allow a router using PPP to monitor the quality of an on-demand WAN link. If poor link quality is detected, the line can be dropped and redialed to improve performance.

> Echo Packets On

This checkbox controls whether this router will use an echo protocol to monitor the quality of the line.

The number of echo packets sent, and the number of responses, are counted. If the conditions set in the **Drop Link When...** (discussed below) fields are met, the link is dropped.

- If **checked**, echo packets will be regularly sent, and line quality monitored.

Frequency in Seconds (Echo Packets)

This parameter determines how often an echo packet will be sent to the other end. The value must be in the range of 1 to 32.

Drop Link When

These parameters set the size of the echo sequence that will be tracked, and the number of packets that must be lost out of a sequence before the link will be dropped. The values must be in the range of 1 to 32.

LCP Options Configuration Dialog Box

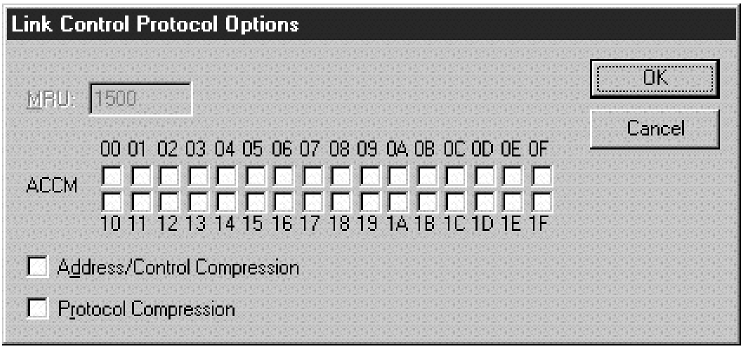

Link Control Protocol Options

MRU: 1500

ACCM

00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
□	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
10	11	12	13	14	15	16	17	18	19	1A	1B	1C	1D	1E	1F

□ Address/Control Compression

□ Protocol Compression

LCP Options Configuration Dialog Box

You can access the LCP Options Configuration Dialog Box by selecting **On Demand PPP Link** or **Dedicated PPP Link** from the **Link Type** pulldown in the Link Configuration: WAN Dialog Box (under WAN/Link Configuration), and then clicking on the **PPP Options** button at the bottom of the dialog box, and then clicking on the **LCP Options** button at the bottom of the PPP Options Dialog Box.

This dialog box is used to set parameters relating to PPP's internal operation. You will probably never need to change the settings in this dialog box.

MRU

This is the Maximum Receive Unit size in bytes for PPP packets. The default value is 1500 bytes.

ACCM

The Asynchronous Character Control Map allows you to set characters which must be “escaped” for your particular communications link. For the vast majority of communications links, the default (no characters escaped) is correct.

❖ **Note:** *If you set Flow Control to XOn/XOff in the Interface Configuration dialog box (under WAN/Physical Configuration) for this WAN interface, the characters for XOn and XOff will automatically be escaped by the router.*

Address/Control Compression

This checkbox controls whether this router will use the method defined in the PPP specification for compression of the PPP address and control fields. The default is **checked**.

Protocol Compression

This checkbox controls whether this router will use the method defined in the PPP specification for compression of the PPP protocol fields. The default is **checked**.

Multilink PPP Dialog Box

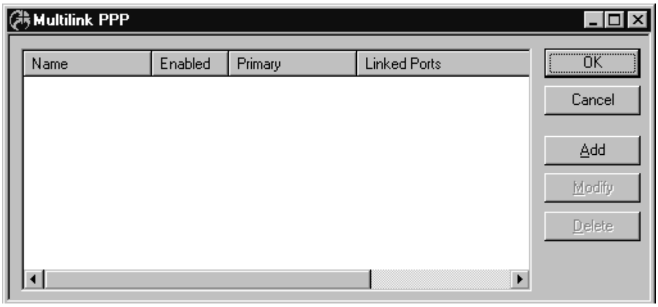

Multilink PPP Dialog Box

This dialog box is used to configure Multilink PPP (MPPP) parameters for multiple WAN interfaces. MPPP allows multiple physical links to be combined into a “bundle” which provides a virtual link with greater bandwidth than a single link

To access this dialog box, select Global/Multilink PPP from the Device View. This dialog box defines a list of MPPP bundles and the physical WAN ports that are included in each bundle. To add or modify this list, click on the appropriate button to open the MPPP Bundle Dialog Box.

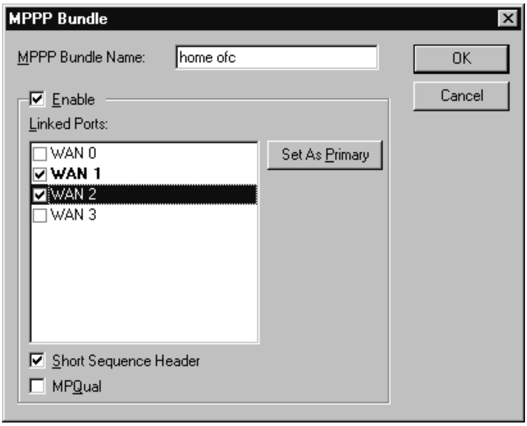

MPPP Bundle Name: home ofc

☑ Enable

Linked Ports:

□ WAN 0
☑ WAN 1
☑ WAN 2
□ WAN 3

☑ Short Sequence Header

□ MPQual

MPPP Bundle Dialog Box

> MPPP Bundle Name

This edit box allows you to specify a name for the multilink virtual port.

> Enable

This checkbox is used to specify whether multilink bundling will function on this router.

> Linked Ports

Check each of the physical WAN ports that you wish to include in the bundle. You must select at least two ports.

> Set as Primary

Select which interface in the bundle should be used by the router to configure the network protocol for the multilink, and click on the **Set as Primary** button.

Short Sequence Header

This checkbox allows the router to use an abbreviated sequence number in its multilink headers.

❖ **Note:** *While the shorter header can enhance performance slightly, routers from other vendors may not be compatible with this feature.*

MPQual

This checkbox allows the router to use echo packets on each of the physical ports in the bundle to determine whether individual links are up. If one link in a bundle goes down, the router can divert data away from that port.

❖ **Note:** *If the primary port goes down, the entire link will go down, even if* ***MPQual*** *is enabled. If left unchecked, any individual link in the bundle can bring down the entire multilink. (Parameters for echo packets are configured in the PPP Options/PPP Link Quality dialog box.)*

WAN Chat Script Editor Dialog Box

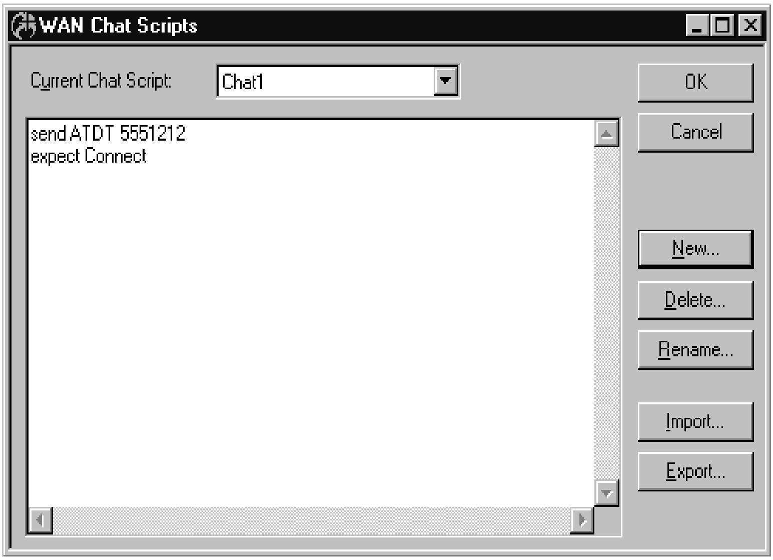

Current Chat Script: Chat1

```
send ATDT 5551212
expect Connect
```

WAN Chat Script Editor Dialog Box

You can access the Chat Script Editor Dialog Box by selecting Global/WAN Chat Scripts in the Device View.

Compatible Systems routers support standard communications chat scripts that let you specify dialing and/or connect sequences between this router and remote routers or terminal servers.

All of the chat scripts stored in a router are available to any of the router's WAN interfaces. To select the scripts which will be used on a specific interface, use the **Dial-out Script / Connect Script** and **Dial-back Script** pull-down menus in CompatiView's Link Configuration: WAN Dialog Box. You can access this dialog box by selecting WAN/Link Configuration from the Device View.

These scripts may also be used for user-specific dial-back scripts in the User Authentication Dialog Box, and can be selected from there. Access this dialog box by selecting Global/User Authentication Database in the Device View.

Chat Script Editor Dialog Box Buttons & Controls

- The **Current Chat Script** pull-down menu lets you select a script for editing.
- The **New** button brings up a dialog box which asks you to name the new script, then creates a blank chat script and selects it in the **Current Chat Script** pull-down menu. Names can be up to 16 characters long.
- The **Rename** button lets you change the name of the chat script you are currently editing.
- The **Delete** button deletes the chat script which is currently selected in the **Current Chat Script** pull-down menu.
- The **Import** button lets you bring a previously exported chat script in from a disk on your computer. The imported information will be appended to the script which is currently selected in the **Current Chat Script** pull-down menu.
- The **Export** button lets you save the chat script which is currently selected in the **Current Chat Script** pull-down menu to a disk file on your computer.

Chat Script Rules and Syntax

Every line in a chat script must start with either **send** or **expect** in order to be a valid chat script line.

- Lines which begin with **send** will cause all other characters (except escaped control characters, as described below) on the line to be output through the WAN interface which is running the script.
- Lines which begin with **expect** will cause the router to wait for matching input characters from the WAN interface which is running the script. The router is case-sensitive when examining returned data.

❖ **Note:** *The amount of time the router will wait is determined by the* ***Script Timeout*** *parameter in the Link Configuration: WAN Dialog Box.*

All control characters are preceded by a backslash character (\) which tells the router that what follows is an escaped character and should not be literally sent on the WAN interface.

- **\r** insert a carriage return
- **\c** don’t add a carriage return to end of line – valid at end of line only
- **\x** insert a hex digit (range 0 to FF)
- **\p** pause for 0.3 seconds
- **\b** send a break character
- \ <*space*> follow the backslash with a space to insert a space; space characters between **send** or **expect** commands and the first character of a line are normally stripped
- **\t** insert a tab
- **\n** insert a new line
- **\q** set “quiet mode” – do not log output until another **\q** encountered
- \\ insert a backslash

A Note About the AT Command Set

Most asynchronous devices (e.g. modems and some terminal adapters) expect AT commands from the router in order to dial or perform other functions. Different modems support different subsets of AT commands. To be certain that the AT commands you are using are correct for your modem, you must refer to the manual that came with your modem.

Every AT command is preceded by “AT,” which tells the modem that the string is destined for it. Listed below are the most common (and commonly supported) AT commands:

- **ATDT** -- Originate a call by dialing the number sequence which follows this command using tones (note: use a comma in the sequence for a delay)

❖ **Note:** *An asynchronous terminal adapter does not use tones to dial ISDN phone numbers. Use* **ATD** *to dial ISDN phone numbers.*

- **ATH0** -- Hang up (note: the final character is a zero)
- **ATM0** -- Set speaker off (note: the final character is a zero)
- **ATM1** -- Set speaker on until connect

Modems typically provide a response message depending on the success of an attempted call:

- **CONNECT** -- The other end has successfully answered. Note that some modems require a switch to be set correctly to receive text responses (as opposed to result codes).

❖ **Note:** *Compatible Systems routers automatically send standard modem setup parameters when a port's **Dialing Method** is set for **AT** dialing. These setup parameters are adequate for virtually all dial-up applications. In almost all cases, your modem should work right out of the box.*

A Note About the V.25bis Command Set

Different CSU/DSU's and Terminal Adapters support different subsets of the V.25bis commands. To be certain that the V.25bis commands you are using are correct for your communications device, you should refer to the manual that came with the device.

The V.25bis commands use hardware signaling to denote whether the information they are sending is destined for the communications device or the data link itself. Listed below are the most common (and commonly supported) V.25bis commands:

- **CRN** -- Originate a call by dialing the number sequence which follows this command

❖ **Note:** *To include a pound sign (#) as part of the number sequence, it must be enclosed in double quotes (" ").*

- **CIC** -- Connect an incoming call

Communications devices provide several responses depending on the outcome of an attempted call:

- **CNX** -- The other end has successfully answered
- **INC** -- An incoming call has been detected
- **VAL** -- The command received is valid
- **INV** -- The command received is invalid or is not supported (may be followed by an error code)
- **CFI** -- Call Failure Indicator; the call could not be completed

❖ **Note:** *If your router is connected to a device synchronously, make sure to configure it to accept V.25bis commands in bit-synchronous format (i.e. within HDLC packets). This is the format Compatible Systems routers use to send V.25bis commands.*

Chat Script Examples

There are as many variations of chat scripts as there are specific installation requirements. However, all chat scripts generally follow the same format, which is a series of **send** and **expect** statements.

- To connect to another router using a modem. This script dials through a PBX which requires a 9 to be dialed followed by a delay in order to access an outside line:

```
send atdt 9,13035559000
expect CONNECT
```

- To connect to another router via an ISDN line, using V.25bis dialing:

```
send CRN 5554000
expect CNX
```

- To connect to an Internet Service Provider using a modem:

```
send atdt 5551000
expect CONNECT
expect login:
send myname
expect ssword:
send im4CSCru2
expect connecting
```

❖ **Note:** *As demonstrated in this script, it may be convenient to only put part of the expected response in an* ***expect*** *statement. This can make it easier to get an exact match when the actual expected string is long (e.g. Please login:, Please enter your password:, etc.).*

User Authentication Database Dialog Box

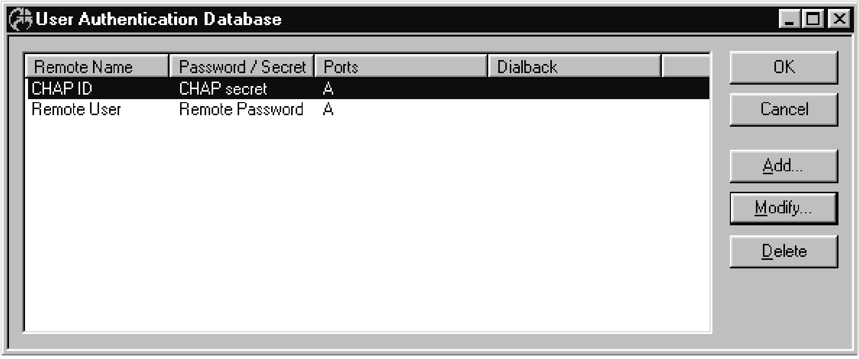

Remote Name	Password / Secret	Ports	Dialback
CHAP ID	CHAP secret	A	
Remote User	Remote Password	A	

User Authentication Database Configuration Dialog Box

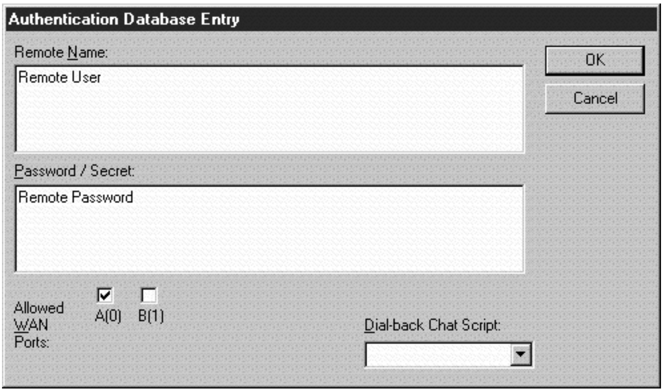

Remote Name: Remote User

Password / Secret: Remote Password

Allowed WAN Ports: ☑ A(0) □ B(1)

Dial-back Chat Script:

Authentication Database Entry Dialog Box

You can access the User Authentication Database Configuration Dialog Box by selecting Global/User Authentication Database in the Device View. This dialog box displays all database entries, but is not used to add or modify the entries.

To add or modify database entries, you must access the Authentication Database Entry Dialog Box by selecting the **Add...** or **Modify...** buttons in the User Authentication Database Configuration Dialog Box.

This database is global to the router. If you have configured a RADIUS server, entries in this database will take precedence over RADIUS entries.

> Remote Name

This is the name of the remote device.

- For PAP entries, this is the name of the device we are requesting a password from, when the **Request PAP Authentication** checkbox is set in this router's PAP Configuration Dialog Box.
- For CHAP entries, this is the name of the device we will send a challenge to, when the **Request CHAP Authentication** checkbox is set in this router's CHAP Configuration Dialog Box.

❖ **Note:** *If there is a Compatible Systems router at the far end, these names correspond to the names entered in the CHAP Configuration Dialog Box and/or PAP Configuration Dialog Box **Name** fields.*

> Password/Secret

This is the password or secret string for the remote device.

- For PAP entries, this is the password value which must be returned from the remote device before we will grant it access to this router.
- For CHAP entries, this is the secret value which is shared with the remote device which will be challenged by this router. This value, along with the random value in the challenge, will be used to determine whether a response is valid.

❖ **Note:** *If there is a Compatible Systems router at the far end, these strings correspond to the **Password** entered in the PAP Configuration Dialog Box and/or the **Secret** entered in the CHAP Configuration Dialog Box.*

> Interfaces

This is the list of interfaces on which we will accept the entered **Name** and **Password** as valid. The entry will be invalid on interfaces not selected here.

Dial-back Chat

If a chat script is selected in this pulldown, then upon successful negotiation of PAP or CHAP, the link will be dropped and the selected chat script will be executed.

Chapter 11 - TCP/IP Filtering

Main TCP/IP Filtering Dialog Box

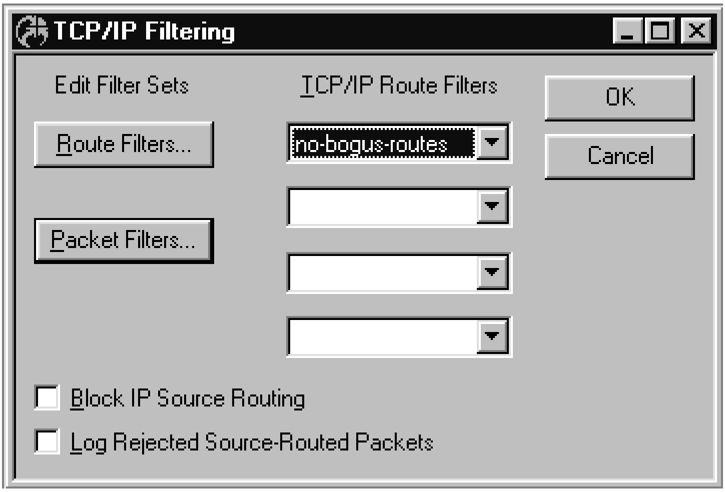

TCP/IP Route Filters

no-bogus-routes

□ Block IP Source Routing

□ Log Rejected Source-Routed Packets

Main TCP/IP Filtering Configuration Dialog Box

To access this dialog box, select Global/Filtering/TCP/IP Filtering from the Device View.

Route Filters Button

This button brings up a filter editor screen for creating route filters. The screen is described later in this chapter.

Packet Filters Button

This button brings up a filter editor screen for creating packet filters. The screen is described later in this chapter.

TCP/IP Route Filters

This set of pull-downs allows you to select previously defined sets of internetworking device filter rules. These rules are global for the device and are not associated with any interface. Up to four sets of rules can be selected.

Block IP Source Routing

This check box sets a filter in the device which drops any received packet which has the “source route” option set.

Log Rejected Source-Routed Packets

This checkbox tells the device to add a log entry (if logging is turned on) whenever the **Block IP Source Routing** checkbox is set and a packet is received with the source route option set. See the section on the Logging Configuration Dialog Box of this manual for more information.

TCP/IP Filter Editor Window

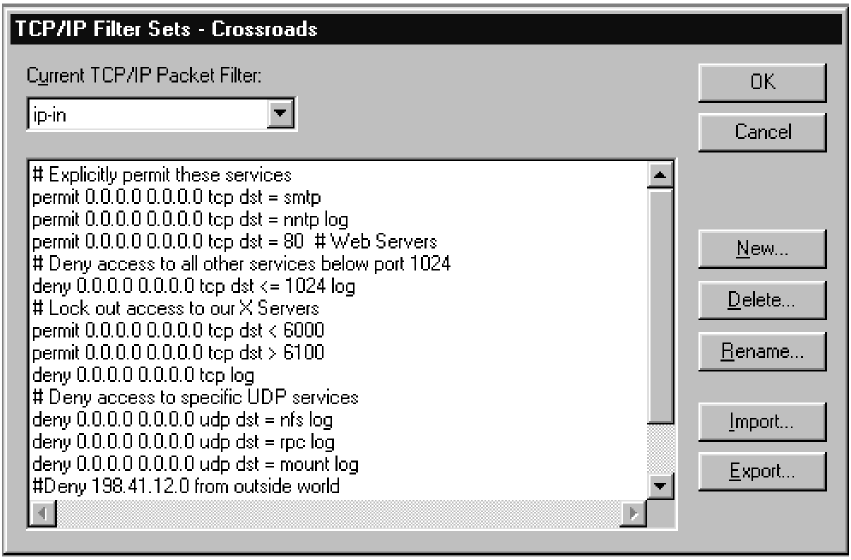

TCP/IP Filter Sets - Crossroads

Current TCP/IP Packet Filter: ip-in

```
# Explicitly permit these services
permit 0.0.0.0 0.0.0.0 tcp dst = smtp
permit 0.0.0.0 0.0.0.0 tcp dst = nntp log
permit 0.0.0.0 0.0.0.0 tcp dst = 80  # Web Servers
# Deny access to all other services below port 1024
deny 0.0.0.0 0.0.0.0 tcp dst <= 1024 log
# Lock out access to our X Servers
permit 0.0.0.0 0.0.0.0 tcp dst < 6000
permit 0.0.0.0 0.0.0.0 tcp dst > 6100
deny 0.0.0.0 0.0.0.0 tcp log
# Deny access to specific UDP services
deny 0.0.0.0 0.0.0.0 udp dst = nfs log
deny 0.0.0.0 0.0.0.0 udp dst = rpc log
deny 0.0.0.0 0.0.0.0 udp dst = mount log
#Deny 198.41.12.0 from outside world
```

TCP/IP Filter Editor Window

The editor window shown above is used in CompatiView for editing all TCP/IP filter sets, including those for TCP/IP Route and Packet filters. The editor window type can be identified by the text at the top of the window, and will only allow you to create or select the type of filter set for which it was selected.

Filter Editor Dialog Box Buttons and Controls

- The **Current Filter** pull-down menu lets you select a filter set for editing.
- The **New** button lets you create a new set of filter rules. A dialog box will pop up to ask you to name the filter set. The name must be 16 characters or less.
- The **Delete** button lets you delete the selected set of filter rules.
- The **Rename** button lets you rename the selected set of filter rules.
- The **Import** button lets you import a previously exported set of filter rules, or a text file in which you have stored filter rules. A file dialog box will pop up to ask you to locate an import file.
- The **Export** button lets you export a set of filter rules to disk. A dialog will pop up to ask you to name the export file.

TCP/IP Route Filter Rules

To access an editor window for TCP/IP route filters, open the Main TCP/IP Filtering Dialog Box (under Global/Filtering/TCP/IP Filtering) and then select the **Route Filters** button.

Route filtering rules are applied globally in the device and are not associated with any interface. However, they can be restricted to an interface using the “from” or “to” modifiers in the rule.

A device does not reorder rule sets as they have been specified before they are applied. They are applied in the order they were written. When multiple filter sets are selected with CompatiView, the filter sets will be concatenated in the device from first to last (top to bottom on screen).

Any IP network not explicitly allowed by the rules will not be included in the routing table on input or in the routing update on output. To allow all other network numbers not filtered, the last rule must be:

```
permit 0.0.0.0
```

Because direct and static routes are configured in the device and not received via an interface, they are always installed and cannot be filtered.

Rules that have been specified using CompatiView may be edited or examined through the command line interface, and vice-versa. When the rules are downloaded into the device from CompatiView, they will be encrypted.

Rule sets that have been created with the TCP/IP Route Filter Editor Window must be applied using the pull-down menus in the Main TCP/IP Filtering Dialog Box.

Basic IP Route Filter Rules and Syntax

At a minimum, every non-comment line in a filter set must include an action, and an IP address. Together these components specify a filter rule that the device will follow when sending and/or receiving IP routing packets.

Every line in a route filter set must begin with the actions **permit** or **deny**, or the comment indicator **#**.

- Lines which begin with **permit** specify that information from routing packets meeting the conditions should be included in the IP routing table.
- Lines which begin with **deny** specify that information from packets meeting the conditions should not be included in the IP routing table.
- Lines which begin with **#** specify that the text on the line is a comment and should be ignored.

Every line which begins with permit or deny must be followed by an IP address. This IP address can be specified in a number of different ways.

- Addresses can be specified in **dotted-decimal notation**. If the rightmost components are 0, they are treated as wildcards. For example, 128.138.12.0 matches all hosts on the 128.138.12 subnet. An address with all zeros matches anything.
- A **factorized format** can also be used where a set of components are substituted into an address. These addresses take the form of #.#.#.{#,#,...}. For example, 192.12.9.{1,2,15} matches the hosts 192.12.9.1, 192.12.9.2, and 192.12.9.15. The factor set must be at the end of the address, but addresses of the form #.{#,#,...}, #.#.{#,#,...}, etc., are allowed. Any components past the factor set's position are implicitly assumed to be 0.
- IP addresses may also be specified as a **hexadecimal number** (for example, 0x82cc0801 matches the host address 130.204.8.1).

Any address may have an optional **/bits** field at its end. This denotes the number of bits, starting with the most significant, that will be considered by the device when it compares the address in a routing packet to the filter rule. For example, an address specified in the rules as 192.15.32.0/19 would match all host addresses from 192.15.32.1 to 192.15.63.255.

Any part of an address which is past the number of significant bits specified is ignored and assumed to be zero.

IP Route Filter Rule Options

A direction can optionally be specified with **in**, **out** or **both**. If no direction is specified, **both** is assumed.

- Filter rules specifying **in** are only applied to routing packets coming into the device.
- Filter rules specifying **out** are only applied to routing packets being sent from the device.
- Filter rules specifying **both** are applied to routing packets in both directions.

IP Route Filter Rule Modifiers

Filter rules can be modified with the following parameters. When used, the modifiers must be put in a filter rule in the order shown below. By default, a filter rule is applied to all routing data.

- **via** <protocol(s)> This modifier specifies that the filtering rule should only be applied to routing data being received or transmitted by the designated routing protocol. Allowed values are **icmp**, **rip**, and **ripv2**. Multiple protocols may be listed, each separated by white space. The **icmp** keyword implies redirected routes.
- **origin** <protocol(s)> This modifier limits output rules to routes originating from the designated protocol. Allowed values are **icmp**, **rip**, **ripv2**, **static**, and **direct**. Multiple protocols may be listed, each separated by white space.
- **metricin** <increment value> This modifier tells the device to increment the metric on incoming routes which match the filter rule. The metric is the number of routers on a route. By increasing or decreasing the metric, a particular route can be made more or less attractive. The value to increment by can be from 1 to 15.
- **metricout** <increment value> This modifier tells the device to increment the metric on outgoing routes which match the filter rule. The metric is the number of routers on a route. By increasing or decreasing the metric, a particular route can be made more or less attractive. The value to increment by can be from 1 to 15.
- **from** <IP address> *or* **from** <interface> This modifier tells the device to apply the rule only to routes coming from a specified IP address (where the address is in the same format as discussed above), or interface (e.g. Ethernet 0, WAN 1, etc.).
- **to** <IP address> *or* **to** <interface> This modifier tells the device to apply the rule only to routes being sent to a specified IP address (where the address is in the same format as discussed above), or interface (e.g. Ethernet 0, WAN 1, etc.).

IP Route Filter Rule Notification

Filter rule matches can optionally cause a log message to be sent. By default, no logging of matches is performed. See the section on the Logging Configuration Dialog Box of this manual for more information.

- **log** The log option causes the device to log data about the packet to syslog when the condition of the rule is met.

IP Route Filter Rule Examples

The following example specifies a rule to allow routes to be input only from RIP and only from 198.41.11.1.

```
permit 0.0.0.0 in via rip from 198.41.11.1
```

The rule below specifies that routing information should only be sent which originates from RIP, directly connected routes, and static routes.

```
permit 0.0.0.0 out origin rip direct static
```

TCP/IP Packet Filter Rules

❖ **Note:** *Due to the nature of the IP protocol, IP packet filtering can be quite complicated. If you are attempting to design and implement a comprehensive set of filters, or an Internet Firewall, there are a number of references you should consult. Two good starting points are: Building Internet Firewalls, by D. Brent Chapman and Elizabeth D. Zwicky, O'Reilly & Associates, 1995, and Firewalls and Internet Security, by William R. Cheswick and Steven M. Bellovin, Addison-Wesley Publishing Company, 1994.*

To access a filter editor window for TCP/IP packet filters, open the Main TCP/IP Filtering Dialog Box (under Global/Filtering/TCP/IP Filtering) and then select the **Packet Filters** button.

Packet filtering rules are selected for individual device interfaces. Whether they are used as input filters, output filters, or both, depends on which pulldown is used to select them in the TCP/IP Filtering Dialog Box for a particular interface.

A device does not reorder rule sets as they have been specified before they are applied. They are applied in the order they were written. When multiple filter sets are selected with CompatiView, the filter sets will be concatenated in the device from first to last (top to bottom on screen).

Any IP packet not explicitly allowed by the rules will be filtered. To allow all other packets not filtered, the last rule must be:

```
permit 0.0.0.0 0.0.0.0 ip
```

Rules that have been specified using CompatiView may be edited or examined through the command line interface, and vice-versa. When the rules are downloaded into the device from CompatiView, they will be encrypted.

Basic IP Packet Filter Rules and Syntax

At a minimum, every non-comment line in a filter set must include an action, a source IP address, and a destination IP address. Together these components specify the action to be taken when a packet meets the condition of the rule.

Every line in a packet filter set must begin with the actions **permit** or **deny**, or the comment indicator #.

- Lines which begin with **permit** specify that packets meeting the conditions should be passed through the filter.
- Lines which begin with **deny** specify that packets meeting the conditions should be dropped by the filter.
- Lines which begin with # specify that the text on the line is a comment and should be ignored.

Every line which begins with permit or deny must be followed by a source and destination IP address. These IP addresses can be specified in a number of different ways.

- Addresses can be specified in **dotted-decimal notation**. If the rightmost components are 0, they are treated as wildcards. For example, 128.138.12.0 matches all hosts on the 128.138.12 subnet. An address with all zeros (0.0.0.0) matches anything.
- A factorized format can also be used where a set of components are substituted into an address. These addresses take the form of #.#.#.{#,#,...}. For example, 192.12.9.{1,2,15} matches the hosts 192.12.9.1, 192.12.9.2, and 192.12.9.15. The factor set must be at the end of the address, but addresses of the form #.{#,#,...}, #.#.{#,#,...}, etc., are allowed. Any components past the factor set’s position are implicitly assumed to be 0.

❖ **Note:** *When the factorized format is used, one line is substituted for many. However, when the device reads the filters and installs them, it expands each address into a separate rule. In the example given, three rules would be created. This can make the number of rules to process greater, which can affect performance.*

- IP addresses may also be specified as a hexadecimal number (for example, 0x82cc0801 matches the host address 130.204.8.1).

Any address may have an optional **/bits** field at its end. This denotes the number of bits, starting with the most significant, that will be considered by the device when it compares the address in a packet to the filter rule. For example, an address specified in the rules as 192.15.32.0/19 would match all host addresses from 192.15.32.1 to 192.15.63.255.

Any part of an address which is past the number of significant bits specified is ignored and assumed to be zero.

IP Packet Filter Rule Operators and Port Names

Filter rules can accept certain modifiers, which are described in the next subsection of this manual. All of these modifiers use a set of expression operators to allow information in a packet to be compared to the modifier’s parameters. These operators are discussed below:

- **eq**, **==**, or **=** These are allowable ways of writing an "equality" operator which will match a packet if its port number is equal to the port specified in the modifier.
- **lt** or **<** These are allowable ways of writing a "less than" operator which will match a packet if its port number is less than the port specified in the modifier.
- **lteq**, **le**, **<=**, or **=<** These are allowable ways of writing a "less than or equal to" operator which will match a packet if its port number is less than or equal to the port specified in the modifier.
- **gt** or **>** These are allowable ways of writing a "greater than" operator which will match a packet if its port number is greater than the port specified in the modifier.
- **gteq**, **ge**, **>=**, or **=>** These are allowable ways of writing a "greater than or equal to" operator which will match a packet if its port number is greater than or equal to the port specified in the modifier.
- **ne**, **<>**, or **!=** These are allowable ways of writing an "inequality" operator which will match a packet if its port number is not equal to the port specified in the modifier.

❖ **Note:** *In rules where expressions are used, the syntax checker requires a space before and a space after the expression operator(s).*

All of the modifiers also require a port number between 0 and 65535. Port numbers can also be specified using the names in the following list of services with known ports:

TCP Ports:		
systat (11)	netstat (13)	ftp-data (20)
ftp (21)	telnet (23)	smtp, mail (25)
whois (43)	gopher (70)	rje (77)
pop-2 (109)	pop-3 (110)	auth (113)
nntp, usenet (119)	netbios-ssn (139)	news (144)
rexec (512)	rlogin (513)	rshell (514)
printer, lpd (515)	uucp (540)	listen, rfs (1025)
x, xwin (6000)	irc (6667)	www, http (80)
UDP Ports:		
name (42)	bootps (67)	bootpc (68)
tftp (69)	snmp (161)	snmp-trap (162)
biff, comsat (512)	rwho (513)	syslog (514)
talk (517)	ntalk (518)	route, rip (520)
timed (525)	mount (635)	pcnfs (640)
nfs (2049)		
Common UDP and TCP Ports:		
echo (7)	discard (9)	daytime (13)
chargen (19)	time (37)	dns, domain (53)
sunrpc, rpc, portmapper (111)	ntp (123)	netbios-ns (137)
netbios-dgm (138)		

ICMP TYPES:		
echo-reply (0)	dest-unrch (3)	src-quench (4)
redirect (5)	echo, ping (8)	time-exceed (11)
param-prob (12)	time (13)	time-reply (14)
info (15)	info-reply (16)	mask (17)
mask-reply (18)		

❖ **Note:** *RFC 1700 "Assigned Numbers" contains a listing of all currently assigned IP protocol keywords and numbers.*

IP Packet Filter Rule Modifiers

These modifiers act to restrict the type of packets which will match a filter rule.

- **IP** This option specifies that all packets from the source and destination IP address and mask will match this rule. If no particular IP protocol packet type (**TCP**, **UDP**, **ICMP**, **GRE**, **AH**, **ESP** or **OSPF**) is specified, **IP** is assumed.

  The IP protocols, other than IP itself, may be specified as a decimal number or as a keyword. The supported keywords are followed by their protocol numbers for your reference.

TCP (6)	UDP (17)
ICMP (1)	GRE (47)
AH (51)	OSPF (89)
ESP (50)	

- **TCP**
  or **TCP src** <expression> <port>
  or **TCP dst** <expression> <port>
  or **TCP est**
  or **TCP src** <expression> <port> **est**
  or **TCP dst** <expression> <port> **est**

  This modifier allows filtering on TCP (Transmission Control Protocol) packets. A source or destination port may be filtered by including the **src** or **dst** specifiers, followed by a logical expression and a port (as described in the subsection above).

The **est** keyword allows a rule to be established in which an external connection to a particular port is not allowed, but two way traffic established by an internal machine will pass through the device.

The device performs this operation by examining the flags in the TCP header. When a session is being established, the first packet only contains the "SYN" flag while subsequent packets contain the "ACK" flag. A permit packet filter rule using the **est** keyword will not match a packet with only the "SYN" flag and the packet will be dropped. Unless another rule allows it through, the "SYN" packet doesn't reach its destination, no reply will be returned to the sender, and a connection will never be established.

Examples using the **est** keyword are shown later in this chapter.

- **UDP**
  or **UDP src** <expression> <port>
  or **UDP dst** <expression> <port>
  This modifier allows filtering on UDP (User Datagram Protocol) packets. A source or destination port may be filtered by including the optional **src** and **dst** specifiers, followed by a logical expression and a port (as described in the subsection above).

❖ **Note:** *CompatiView uses UDP port 33020. Care should be taken not to deny this port if CompatiView configuration is desired.*

- **ICMP**
  or **ICMP type** <expression> <port>
  This modifier allows filtering on ICMP (Internet Control Message Protocol) packets. The ICMP type may be filtered by using the type specifier and the list of types from the subsection above.

- **GRE**
  This modifier allows filtering on GRE (Generic Routing Encapsulation) packets. GRE provides a simple, general purpose mechanism to encapsulate network protocols into IP for the purpose of tunneling across the Internet.

❖ **Note:** *If VPN tunneling without authentication is enabled on an interface to which an IP filter is applied, then the filter must specifically* ***permit*** *GRE packets.*

- **AH**
  This modifier allows filtering on AH (Authentication Header) packets. AH is used for authentication of tunneled packets across the Internet.

❖ **Note:** *If VPN tunneling with authentication is enabled on an interface to which an IP filter is applied, then the filter must specifically **permit** AH packets.*

- **ESP**
  This modifier allows filtering on ESP (Encapsulating Security Payload) packets. ESP is used for encryption of tunneled packets across the Internet.

❖ **Note:** *If VPN tunneling with encryption only (i.e. no authentication) is enabled on an interface to which an IP filter is applied, then the filter must specifically **permit** ESP packets.*

- **OSPF**
  This modifier allows filtering on OSPF (Open Shortest Path First) packets. OSPF IP packets carry OSPF routing data.
- **proto** <operator> <protocol number>
  This modifier allows general filtering of IP protocol numbers that don't have established keywords as specified above. The rule also allows an expression to be specified which allows filtering on ranges of protocol numbers (i.e. proto > 51).

IP Packet Filter Rule Notification

There are two notification actions which the device can take when a packet matches a particular rule. By default, no logging or notification of matches is performed.

- **log** The log option causes the device to log data about the packet to syslog when the condition of the rule is met. See the section on the Logging Configuration Dialog Box of this manual for more information.
- **icmp** The icmp option is valid only on a **deny** rule and directs the device to return an ICMP notification to the source of the matching packet.

Simple IP Packet Filter Rule Examples

This rule allows TCP packets with a source port greater than or equal to 1024 and a destination port of 25 (SMTP mail):

```
permit 0.0.0.0 0.0.0.0 TCP src >= 1024 dst = 25
```

A rule to allow UDP packets with a source port greater than 910 and a destination port of 53 (Domain Name Service) would look like:

```
permit 0.0.0.0 0.0.0.0 UDP src > 910 dst = 53
```

A rule to deny ICMP echo request (pings) would look like:

```
deny 0.0.0.0 0.0.0.0 ICMP type = 8
```

This rule would drop all packets with the source host address 192.15.1.10:

```
deny 192.15.1.10 0.0.0.0
```

A rule to drop all packets with a source network address of 192.15.1.0. All packets from hosts on that network would be denied:

```
deny 192.15.1.0/24 0.0.0.0
```

IP Packet Filter Rule Set Examples

The rule set below allows only inbound and outbound mail from 192.15.14.1.

The input-filter:

```
permit 0.0.0.0 192.15.14.1 TCP src >= 1024 dst = 25
permit 0.0.0.0 192.15.14.1 TCP src = 25 dst >= 1024
```

The output-filter:

```
permit 192.15.14.1 0.0.0.0 TCP src = 25 dst >= 1024
permit 192.15.14.1 0.0.0.0 TCP src >= 1024 dst = 25
```

These sets of rules are intended to filter out all traffic and only allow incoming and outgoing mail to a server inside a net with an IP address of 192.15.14.1. However they aren’t enough to prevent access from someone outside using source port 25. This is because a connection to destination ports greater than 1024 can be initiated according to the second rule in the input filter. To prevent this from happening, add the **est** keyword to the second rule in the input filter:

```
permit 0.0.0.0 192.15.14.1 TCP src = 25 dst >= 1024 est
```

The **est** keyword in this rule tells the device to only accept TCP packets on the input to this interface when the connection has already been established. A TCP packet which is attempting to initiate a connection will have only the “SYN” flag set. If someone tries to establish a connection from the outside using source port 25, the rule won’t match (no permit will occur). The connection can’t be established since the packet will be dropped by the default rule.

TCP/IP Packet Filtering: Ethernet Dialog Box

TCP/IP Packet Filtering: WAN Dialog Box

TCP/IP Packet Filtering: VPN Dialog Box

TCP/IP Packet Filtering: Bridge Dialog Box

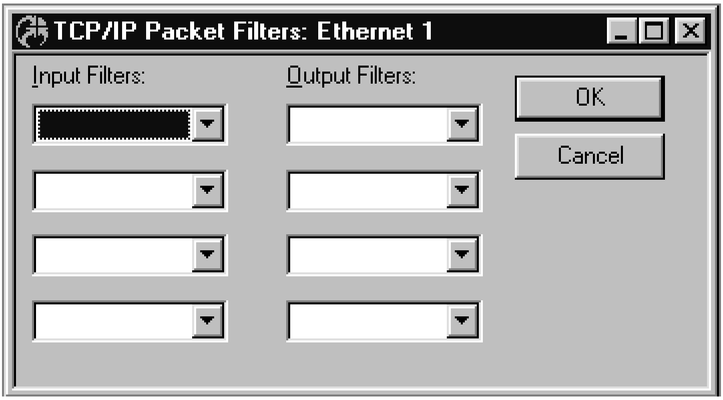

Interface TCP/IP Packet Filtering Configuration Dialog Box

To access this dialog box, select Interface/Filtering/TCP/IP Filtering from the Device View. This can be done for any type of interface except IP subinterfaces.

Input Filters

This set of pulldowns allows you to select previously defined sets of packet filter rules. These rules will be applied to packets arriving on this interface. Up to four sets of rules can be selected.

Output Filters

This set of pulldowns allows you to select previously defined sets of packet filter rules. These rules will be applied to packets which are to be sent on this interface. Up to four sets of rules can be selected.

Chapter 12 - IPX Filtering

Main IPX Filtering Dialog Box

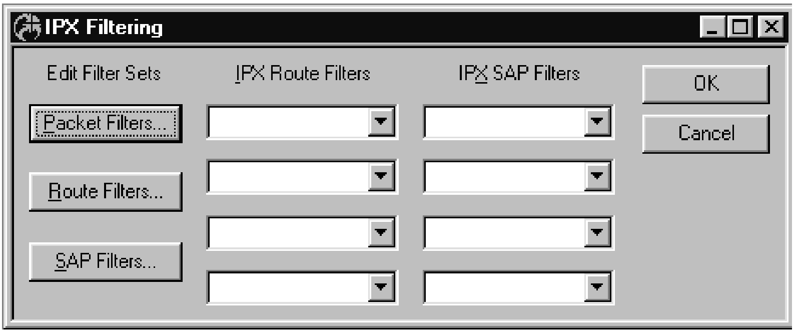

Main IPX Filtering Configuration Dialog Box

To access this dialog box, select Global/Filtering/IPX Filtering from the Device View.

IPX Route Filters

This set of pulldowns allows you to select previously defined sets of internetworking device filter rules that operate on the IPX Routing Information Protocol (RIP). These rules are global for the device and are not associated with any interface. Up to four sets of rules can be selected.

IPX SAP Filters

This set of pull-downs allows you to select previously defined sets of internetworking device filter rules that operate on the IPX Service Advertising Protocol (SAP). These rules are global for the device and are not associated with any interface. Up to four sets of rules can be selected.

IPX Filter Editor Window

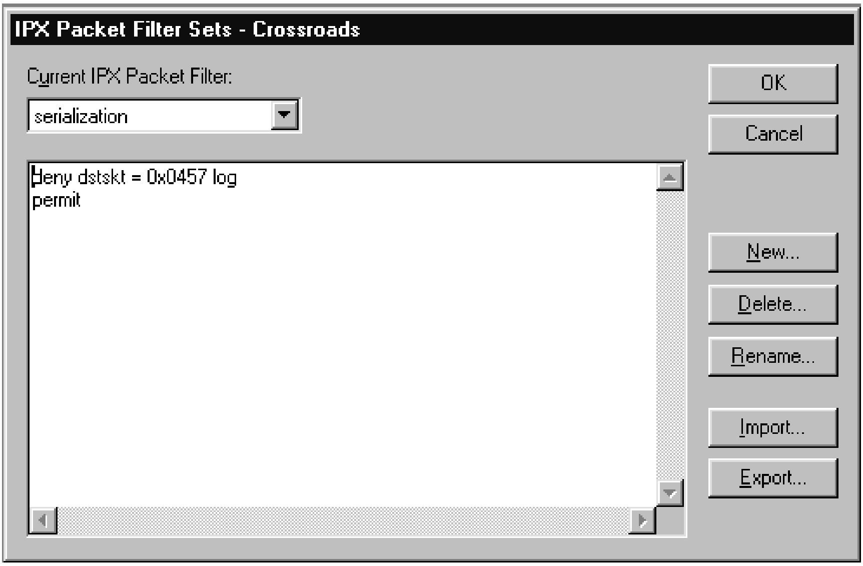

Current IPX Packet Filter: serialization

```
deny dstskt = 0x0457 log
permit
```

IPX Filter Editor Window

The editor window shown above is used in CompatiView for editing all IPX filter sets, including those for IPX Route, SAP, and Packet filters. The editor window type can be identified by the text at the top of the window, and will only allow you to create or select the type of filter set for which it was selected.

Filter Editor Window Buttons and Controls

- The **Current Filter** pull-down menu lets you select a filter set for editing.
- The **New** button lets you create a new set of filter rules. A dialog will pop up to ask you to name the filter set. The name must be 16 characters or less.
- The **Delete** button lets you delete the selected set of filter rules.
- The **Rename** button lets you rename the selected set of filter rules.

- The **Import** button lets you import a previously exported set of filter rules, or a text file in which you have stored filter rules. A file dialog will pop up to ask you to locate an import file.
- The **Export** button lets you export a set of filter rules to disk. A dialog will pop up to ask you to name the export file.

IPX Packet Filter Rules

To access an editor window for IPX Packet filters, open the Main IPX Filtering Dialog Box (under Global/Filtering/IPX Filtering) and then select the **Packet Filters** button.

Packet filtering rules are applied on a per interface basis. Whether they are used as input filters, output filters, or both, depends on which pulldown is used to select them in the IPX Filtering Dialog Box for a particular interface.

A device does not reorder rule sets as they have been specified before they are applied. They are applied in the order they were written. When multiple filter sets are selected with CompatiView, the filter sets will be concatenated in the device from first to last (top to bottom on screen).

Any IPX packet not explicitly allowed by the rules will not be passed through the filter. To allow all other packets not filtered, the last rule must be:

```
permit
```

Rules that have been specified using CompatiView may be edited or examined through the command line interface, and vice-versa. When the rules are downloaded into the device from CompatiView, they will be encrypted.

Basic IPX Packet Filter Rules and Syntax

At a minimum, every non-comment line in a filter set must include an action. However, an action alone will not create a useful filter rule (except for setting a default rule as noted above).

Every line in a packet filter set must begin with the actions **permit** or **deny**, or the comment indicator #.

- Lines which begin with **permit** specify that a packet meeting the conditions should be passed by the filter.
- Lines which begin with **deny** specify that a packet meeting the conditions should be dropped by the filter.
- Lines which begin with # specify that the text on the line is a comment and should be ignored.

IPX Packet Filter Options

The basic action specified in the rule will almost always be accompanied with an option. IPX Packet filter options use some or all of a set of operators to determine whether the filter rule matches information in a packet or not. These operators are discussed below:

- **eq**, **==**, or **=** These are allowable ways of writing an “equality” operator which will match if the value in the packet is equal to the value specified in the option expression.
- **lt** or **<** These are allowable ways of writing a “less than” operator which will match a packet if its value is less than the value specified in the option expression.
- **lteq**, **le**, **<=**, or **=<** These are allowable ways of writing a “less than or equal to” operator which will match a packet if its value is less than or equal to the value specified in the option expression.
- **gt** or **>** These are allowable ways of writing a “greater than” operator which will match a packet if its value is greater than the value specified in the option expression.
- **gteq**, **ge**, **>=**, or **=>** These are allowable ways of writing a “greater than or equal to” operator which will match a packet if its value is greater than or equal to the value specified in the option expression.
- **ne**, **<>**, or **!=** These are allowable ways of writing an “inequality” operator which will match if the value in the packet is not equal to the value specified in the option expression.

❖ **Note:** *In rules where expressions are used, the syntax checker requires a space before and a space after the expression operator(s).*

The options available for IPX Packet filter rules allow rules to be more narrowly specified to exclude all but certain **types** of packets, packets with a given source network number (**srcnet**), packets with a specified destination network numbers (**dstnet**), packets with a particular source socket number (**srcskt**), packets with a selected destination socket number (**dstskt**), packets with a chosen source node address (**srcnode**), and/or packets with a stated destination node address (**dstnode**).

- **type** <operator> <IPX packet type> This option allows filtering using the IPX packet **type** contained in the packet. The IPX packet type value must be a hex number. The keyword **all** may be used to specify all network number values.

  For some versions of NetWare, the packet type field is not a reliable indicator of the type of packet encapsulated by the IPX header. Generally, the source and destination sockets should be used to implicitly filter the packet type. NetBIOS propagate packets (type 14h) are an exception to this rule.

- **srcnet** <operator> <network number> This option allows filtering of the source network number contained in the packet. The number is specified in hex. The keyword **all** may be used to specify all network number values.
- **dstnet** <operator> <network number> This option allows filtering of the destination network number contained in the packet. The number is specified in hex. The keyword **all** may be used to specify all network number values.
- **srcskt** <operator> <socket number> This rule allows filtering of the source socket contained in the packet. The number is specified in hex.

  The following keywords may be used for well known socket values: **NCP**(0451h), **SAP**(0452h), **RIP**(0453h), or **DIAG**(0456h). The keyword **all** may be used to specify all socket numbers.
- **dstskt** <operator> <socket number> This rule allows filtering of the destination socket contained in the packet. The number is specified in hex. The keywords listed above for **srcskt** may also be used. The keyword **all** may be used to specify all socket numbers.
- **srcnode** <operator> <node address> This rule allows filtering of the source node address contained in the packet. The operator in this option can only be “equality” or “inequality.”

  The node address parameter is the IPX server node number specified as an Ethernet address. An Ethernet address is specified as six hexadecimal octets separated by dots or colons (e.g. 0.0.A5.0.0.1 or 0:0:A5:0:0:1). The keyword **all** may be used to specify all node values.
- **dstnode** <operator> <node address> This rule allows filtering of the destination node address contained in the packet. The operator in this option can only be “equality” or “inequality.” The address parameter should be entered as shown above for **srcnode**. The keyword **all** may also be used.

IPX Packet Filter Rule Notification

Filter rule matches can optionally cause a log message to be sent. By default, no logging of matches is performed. See the section on the Logging Configuration Dialog Box of this manual for more information.

- **log** The log keyword causes the device to send information about the packet to syslog when the condition of the rule is met.

IPX Packet Filter Rule Examples

Drop all packets where the source network number is greater than or equal to 1000 and permit all other packets:

```
deny srcnet >= 1000
permit type = ALL
```

Drop all packets from a specific IPX network and node and permit all other packets:

```
deny srcnet = FAB4 srcnode = 0.0.A5.0.0.1
permit
```

Drop all packets where the source socket is a diagnostic packet, log the denial and permit all other packets through:

```
deny srcskt = DIAG log
permit
```

IPX Route Filter Rules

To access an editor window for IPX Route filters, open the Main IPX Filtering Dialog Box (under Global/Filtering/IPX Filtering) and then select the **Route Filters** button.

Route filtering rules are applied globally in the device and are not associated with any interface. However, they can be restricted to an interface using the “from” or “to” modifiers in the rule.

A device does not reorder rule sets as they have been specified before they are applied. They are applied in the order they were written. When multiple filter sets are selected with CompatiView, the filter sets will be concatenated in the device from first to last (top to bottom on screen).

Any IPX network not explicitly allowed by the rules will not be included in the routing table on input or in the routing update on output. To allow all other network numbers not filtered, the last rule must be:

```
permit network = ALL
```

Rules that have been specified using CompatiView may be edited or examined through the command line interface, and vice-versa. When the rules are downloaded into the device from CompatiView, they will be encrypted.

Rule sets that have been created with the IPX Route Filter Editor Window must be selected using the pull-downs in the Main IPX Filtering Dialog Box.

Basic IPX Route Filter Rules and Syntax

At a minimum, every non-comment line in a filter set must include an action and a network expression. Together these components specify a filter rule that the device will follow when sending and/or receiving IPX RIP packets.

Every line in an IPX Route filter set must begin with the actions **permit** or **deny**, or the comment indicator **#**.

- Lines which begin with **permit** specify that information meeting the conditions should be included in the IPX routing table.
- Lines which begin with **deny** specify that information meeting the conditions should not be included in the IPX routing table.
- Lines which begin with **#** specify that the text on the line is a comment and should be ignored.

The network expression uses a set of operators to specify the conditions under which the rule will be satisfied. These operators are discussed below:

- **eq**, **==**, or **=** These are acceptable ways of writing an "equality" operator which will match if the value in the routing information is equal to the value specified in the network expression.
- **lt** or **<** These are acceptable ways of writing a "less than" operator which will match if the value in the routing information is less than the value specified in the network expression.
- **lteq**, **le**, **<=**, or **=<** These are acceptable ways of writing a "less than or equal to" operator which will match if the value in the routing information is less than or equal to the value specified in the network expression.
- **gt** or **>** These are acceptable ways of writing a "greater than" operator which will match if the value in the routing information is greater than the value specified in the network expression.
- **gteq**, **ge**, **>=**, or **=>** These are acceptable ways of writing a "greater than or equal to" operator which will match if the value in the routing information is greater than or equal to the value specified in the network expression.
- **ne**, **<>**, or **!=** These are acceptable ways of writing an "inequality" operator which will match if the value in the routing information is not equal to the value specified in the network expression.

The keyword **all** may be used to specify all network number values in the network expression.

❖ **Note:** *In rules where expressions are used, the syntax checker requires a space before and a space after the expression operator(s).*

IPX Route Filter Rule Options

Filter rules can optionally include the following parameter. When used, the options must be inserted after the required part of the rule, but before any modifiers.

The direction is specified with **in**, **out**, or **both**. If no direction is specified, **both** is assumed.

- Filter rules specifying **in** are only applied to routing information coming into the device.
- Filter rules specifying **out** are only applied to routing information being sent from the device.
- Filter rules specifying **both** are applied to routing information in both directions.

IPX Route Filter Rule Modifiers

The source address, destination address, source interface or destination interface can be specified using the **from** and **to** modifiers. These keywords modify the global nature of a RIP filter rule.

- **from** <IPX address> *or* **from** <interface> This modifier tells the device to apply the rule only to routes coming from a specified IPX address, or interface (e.g. Ethernet 0, WAN 1, etc.).

  The IPX address parameter is specified as a hexadecimal network number and node number separated by a dash (e.g. A011-0:0:A5:0:0:1 indicates a node with the hexadecimal network number of A011 and a node address of 0:0:A5:0:0:1).

- **to** <IPX address> *or* **to** <interface> This modifier tells the device to apply the rule only to routes being sent to a specified IP address (where the address is in the same format as discussed above), or interface (e.g. Ethernet 0, WAN 1, etc.).

Filter rules can also optionally be set to modify some RIP information as it is handled by the device.

- **metricin** <increment value> This modifier tells the device to increment the metric on incoming routes which match the filter rule. The metric is the number of routers on a route. By increasing or decreasing the metric, a particular route can be made more or less attractive. The value to increment by can be from 1 to 15.

- **metricout** <increment value> This modifier tells the device to increment the metric on outgoing routes which match the filter rule. By increasing or decreasing the metric, a particular route can be made more or less attractive. The value to increment by can be from 1 to 15.

IPX Route Filter Rule Notification

Filter rule matches can optionally cause a log message to be sent. By default, no logging of matches is performed. See the section on the Logging Configuration Dialog Box of this manual for more information.

- **log** The log keyword causes the device to send information about the packet to syslog when the condition of the rule is met.

IPX Route Filter Rule Examples

The following example specifies a rule to allow routes to be input from any IPX network except network number 7.

```
permit network != 7
```

The rule below specifies that routing information should only be accepted from the Ethernet 0 interface.

```
permit network = ALL from ethernet 0
```

IPX SAP Filter Rules

To access a dialog box for IPX SAP filters, open the Main IPX Filtering Dialog Box (under Global/Filtering/IPX Filtering) and then select the **SAP Filters** button.

SAP filtering rules are applied globally in the device and are not associated with any interface. However, they can be restricted to an interface using the “from” or “to” modifiers in the rule.

A device does not reorder rule sets as they have been specified before they are applied. They are applied in the order they were written. When multiple filter sets are selected with CompatiView, the filter sets will be concatenated in the device from first to last (top to bottom on screen).

Any server not explicitly allowed by the rules will not be included in the SAP table on input or in the SAP update on output. To allow all other servers not filtered, the last rule must be:

```
permit
```

Rules that have been specified using CompatiView may be edited or examined through the command line interface, and vice-versa. When the rules are downloaded into the device from CompatiView, they will be encrypted.

Rule sets that have been created with the IPX SAP Filter Editor Window must be applied using the pull-down menus in the Main IPX Filtering Dialog Box.

Basic IPX SAP Filter Rules and Syntax

At a minimum, every non-comment line in a filter set must include an action. However, an action alone will not create a useful filter rule (except for setting a default rule as noted above).

Every line in a SAP filter set must begin with the actions **permit** or **deny**, or the comment indicator #.

- Lines which begin with **permit** specify that server information meeting the conditions should be inserted into the device's SAP table.
- Lines which begin with **deny** specify that server information meeting the conditions should not be included in the device's SAP table.
- Lines which begin with # specify that the text on the line is a comment and should be ignored.

IPX SAP Filter Options

The basic action specified in the rule will almost always be accompanied with an option. IPX SAP options use some or all of a set of operators to determine whether the filter rule matches information in a SAP packet or not. These operators are discussed below:

- **eq**, **==**, or **=** These are allowable ways of writing an "equality" operator which will match if the value in the server information is equal to the value specified in the option expression.
- **lt** or **<** These are allowable ways of writing a "less than" operator which will match server information if its value is less than the value specified in the option expression.
- **lteq**, **le**, **<=**, or **=<** These are allowable ways of writing a "less than or equal to" operator which will match server information if its value is less than or equal to the value specified in the option expression.
- **gt** or **>** These are allowable ways of writing a "greater than" operator which will match server information if its value is greater than the value specified in the option expression.

- **gteq**, **ge**, **>=**, or **=>** These are allowable ways of writing a “greater than or equal to” operator which will match server information if its value is greater than or equal to the value specified in the option expression.
- **ne**, **<>**, or **!=** These are allowable ways of writing an “inequality” operator which will match if the value in the server information is not equal to the value specified in the option expression.

❖ **Note:** *In rules where expressions are used, the syntax checker requires a space before and a space after the expression operator(s).*

The options available for IPX SAP filter rules allow rules to be more narrowly specified to exclude all but certain **types** of servers, an individual **service**, servers on certain **networks**, servers with a certain **node** address, and/or servers using a certain IPX **socket** address.

- **type** <operator> <server type> This option allows filtering using the server **type** contained in the SAP update tuple. The server type value must be a hex number. The keyword **all** may be used to specify all types.
- **service** <operator> <server name> This option allows filtering using the **service** name contained in the SAP update tuple. The operator in this option can only be “equality” or “inequality.” The name must be 48 characters or less, and enclosed in quotation marks ("").
- **network** <operator> <network number> This option allows filtering of the server **network** number contained in the SAP table. The number is specified in hex. The keyword **all** may be used to specify all network number values.
- **node** <operator> <node address> This rule allows filtering of the server **node** address contained in the SAP table. The operator in this option can only be “equality” or “inequality.”

  The node address parameter is the IPX server node number specified as an Ethernet address. An Ethernet address is specified as six hexadecimal octets separated by dots or colons (e.g. 0.0.A5.0.0.1 or 0:0:A5:0:0:1). The keyword **all** may be used to specify all node values.
- **socket** <operator> <socket number> This rule allows filtering of the server **socket** contained in the SAP table. The number is specified in hex. The keyword **all** may be used to specify all socket numbers.

A final option is the ability to specify a direction using **in**, **out**, or **both**. If no direction is specified, **both** is assumed.

- Filter rules specifying **in** are only applied to server information coming into the device.

- Filter rules specifying **out** are only applied to server information being sent from the device.
- Filter rules specifying **both** are applied to server information in both directions.

IPX SAP Filter Rule Modifiers

The source address, destination address, source interface or destination interface can be specified using the **from** and **to** options. These keywords modify the global nature of a SAP filter rule.

- **from** <IPX address> *or* **from** <interface> This modifier tells the device to apply the rule only to server information coming from a specified IPX address, or interface (e.g. Ethernet 0, WAN 1, etc.).

  The IPX address parameter is specified as a hexadecimal network number and node number separated by a dash (e.g. A011-0:0:A5:0:0:1 indicates a node with the hexadecimal network number of A011 and a node address of 0:0:A5:0:0:1).
- **to** <IPX address> *or* **to** <interface> This modifier tells the device to apply the rule only to server information being sent to a specified IPX address (where the address is in the same format as discussed above), or interface (e.g. Ethernet 0, WAN 1, etc.).

Filter rules can also optionally be set to modify some SAP information as it is handled by the device.

- **metricin** <increment value> This modifier tells the device to increment the metric on incoming servers which match the filter rule. The value to increment by can be from 1 to 15.
- **metricout** <increment value> This modifier tells the device to increment the metric on outgoing servers which match the filter rule. The value to increment by can be from 1 to 15.

IPX SAP Filter Rule Notification

Filter rule matches can optionally cause a log message to be sent. By default, no logging of matches is performed. See the section on the Logging Configuration Dialog Box of this manual for more information.

- **log** The log keyword causes the device to send information about the packet to syslog when the condition of the rule is met.

IPX SAP Filter Rule Examples

The following example specifies a rule set to ignore any server named “Test Server.” The permit line states that all other servers should be entered into the device’s SAP table.

```
deny server = “Test Server”
permit
```

The rule below specifies that only servers from network 7 should be entered into the device’s SAP table. All other SAP types will be dropped.

```
permit network = 7
```

IPX Packet Filtering: Ethernet Dialog Box

IPX Packet Filtering: WAN Dialog Box

IPX Packet Filtering: VPN Dialog Box

IPX Packet Filtering: Bridge Dialog Box

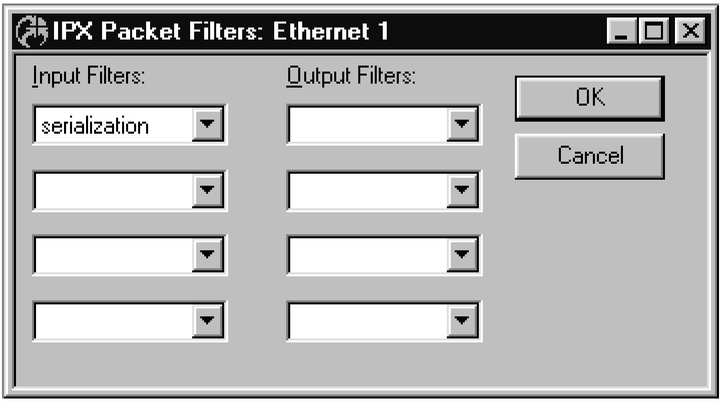
IPX Packet Filters: Ethernet 1

Input Filters: serialization

Interface IPX Packet Filtering Configuration Dialog Box

To access this dialog box, select Interface/Filtering/IPX Filtering from the Device View.

Input Filters

This set of pulldown menus allows you to select previously defined sets of packet filter rules. These rules will be applied to packets arriving on this interface. Up to four sets of rules can be selected.

Output Filters

This set of pull-downs allows you to select previously defined sets of packet filter rules. These rules will be applied to packets which are to be sent on this interface. Up to four sets of rules can be selected.

Chapter 13 - AppleTalk Filtering

Main AppleTalk Filtering Editor Window

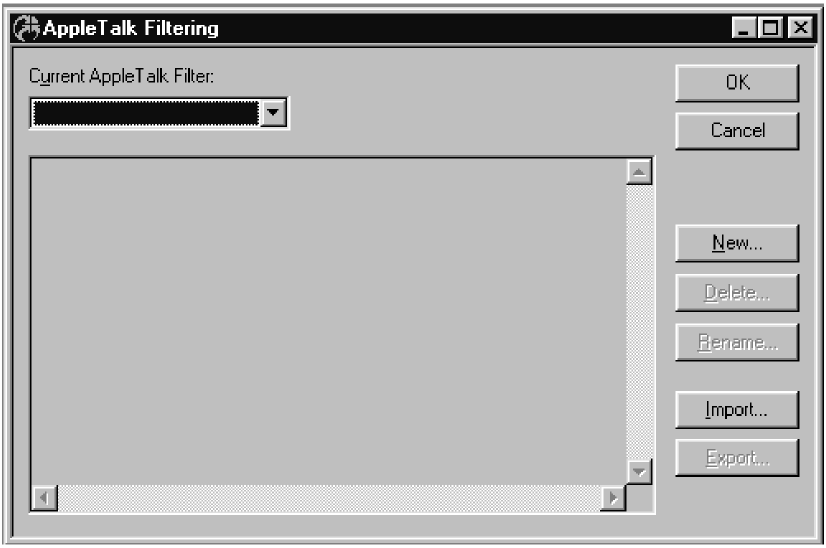

Main AppleTalk Filter Editor Window

To access this editor window, select Global/Filtering/AppleTalk Filtering from the Device View.

The editor window shown above is used in CompatiView for editing all AppleTalk filter sets, including those for AppleTalk Route, Zone List, and Packet filters.

Filter Editor Dialog Box Buttons and Controls

- The **Current Filter** pull-down menu lets you select a filter set for editing.
- The **New** button lets you create a new set of filter rules. A dialog will pop up to ask you to name the filter set. The name must be 16 characters or less.
- The **Delete** button lets you delete the selected set of filter rules.

- The **Rename** button lets you rename the selected set of filter rules.
- The **Import** button lets you import a previously exported set of filter rules, or a text file in which you have stored filter rules. A file dialog will pop up to ask you to locate an import file.
- The **Export** button lets you export a set of filter rules to disk. A dialog will pop up to ask you to name the export file.

AppleTalk Packet Filter Rules

The AppleTalk filter editor window allows a set of AppleTalk filtering rules to be defined, edited and identified with a specific name.

Once a set of rules is defined and named, those rules may be linked to several different AppleTalk filter interpreters to accomplish different types of filtering.

Each interpreter understands and uses a subset of the complete AppleTalk rules. The interpreters available are: general packet filtering, get-zone-list filtering and route (RTMP) filtering. Each is described below.

The interpreters will not reorder the rules as they are specified. They will be applied sequentially from the first rule through the last. Any filtered information not specifically allowed by the set of rules will be dropped silently. If that information is to be allowed, a final permit rule must be specified:

```
permit
```

There is an interaction between the packet filtering interpreter and the other interpreters. The packet filter interpreter will be applied to incoming packets before the other interpreters, and it will be applied to outgoing packets after the other interpreters. For example, a received get-zone-list request may be filtered by an input packet filter before it arrives at the get-zone-list interpreter and the reply may also be filtered again by an outgoing packet filter.

Rules that have been specified using CompatiView may be edited or examined through the command line interface. Likewise, rules defined through the command line interface may be edited through CompatiView. When the rules are downloaded into the device from CompatiView, they will be encrypted.

General Packet Filtering

This interpreter allows packets being forwarded by the device to be filtered on the input and output side of an interface. The only rules used in this interpreter are the **type**, **srcnet**, **dstnet**, **srcnode**, **dstnode**, **srcskt** and **dstskt** for all packets. For NBP request and reply packets the **NBPName**, **NBPType** and **NBPZone** rules are also used. All other rules are ignored.

Get Zone List

The get-zone-list interpreter allows the filtering of outgoing get-zone-list replies on an interface. These replies contain the zone list displayed by the Chooser on a Macintosh when it is opened. Thus, the get-zone-list interpreter allows control of the zones that are seen on a Macintosh behind a device. The only rules used in this interpreter are the **network**, **net-range** and **zone** rules. All other rules are ignored.

Routing Filters (RTMP)

The RTMP interpreter allows network numbers in input and output AppleTalk RTMP routing packets to be filtered on an interface. The only rules used in this interpreter are the **network** and **net-range** rules. All other rules are ignored.

Basic AppleTalk Filter Rules and Syntax

At a minimum, every non-comment line in a filter set must include an action. However, an action alone will not create a useful filter rule (except for setting a default rule as noted above).

Every line in a packet filter set must begin with the actions **permit**, or **deny**, or the comment indicator #.

- Lines which begin with **permit** specify that a packet meeting the conditions should be passed by the filter.
- Lines which begin with **deny** specify that a packet meeting the conditions should be dropped by the filter.
- Lines which begin with # specify that the text on the line is a comment and should be ignored.

AppleTalk Filter Options

The basic action specified in the rule will almost always be accompanied with an option. AppleTalk filter options use some or all of a set of operators to determine whether the filter rule matches the information being examined or not. These operators are discussed below:

- **eq**, **==**, or **=** These are allowable ways of writing an “equality” operator which will match if the value in the packet/information is equal to the value specified in the option expression.
- **lt** or **<** These are allowable ways of writing a “less than” operator which will match the packet/information if its value is less than the value specified in the option expression.

- **lteq**, **le**, **<=**, or **=<** These are allowable ways of writing a “less than or equal to” operator which will match the packet/information if its value is less than or equal to the value specified in the option expression.
- **gt** or **>** These are allowable ways of writing a “greater than” operator which will match the packet/information if its value is greater than the value specified in the option expression.
- **gteq**, **ge**, **>=**, or **=>** These are allowable ways of writing a “greater than or equal to” operator which will match the packet/information if its value is greater than or equal to the value specified in the option expression.
- **ne**, **<>**, or **!=** These are allowable ways of writing an “inequality” operator which will match if the value in the packet/information is not equal to the value specified in the option expression.

The options available for AppleTalk filter rules allow rules to be more narrowly specified to exclude packets or other information based on a number of additional factors.

- **type** <operator> <AppleTalk packet type> This option allows filtering of the packet type from the AppleTalk DDP header. The value must be between 1 and 255. The numbers of some well-known packet types are listed below.

  RTMP (1)NBP (2)ATP (3)

  ECHO (4)RTMP Request (5) ZIP (6)

  ADSP (7)SNMP (8)IP-in-AppleTalk (22)

  DECnet-in-AppleTalk (68)

- **srcnet** <operator> <network number> This option allows filtering of packets by the source network from the AppleTalk DDP header. The value must be between 1 and 65279. The keyword **all** may be used to specify all network numbers.
- **dstnet** <operator> <network number> This option allows filtering of packets by the destination network from the AppleTalk DDP header. The value must be between 1 and 65279. The keyword **all** may be used to specify all network numbers.
- **srcnode** <operator> <node address> This option allows filtering of packets by the source node from the AppleTalk DDP header. The node value must be between 1 and 253.
- **dstnode** <operator> <node address> This option allows filtering of packets by the destination node from the AppleTalk DDP header. The node value must be between 1 and 253.

- **srcskt** <operator> <socket number> This option allows filtering of packets by the source socket from the AppleTalk DDP header. The value must be between 1 and 255.

- **dstskt** <operator> <socket number> This option allows filtering of packets by the destination socket from the AppleTalk DDP header. The value must be between 1 and 255.

- **network** <operator> <network number> This option allows by the network number in Get Zone List and RTMP packets. The value must be between 1 and 65279. The keyword **all** may be used to specify all network numbers.

- **net-range** <operator> <network range> This option allows filtering of Get Zone List and RTMP packets using a network range. Two network numbers separated by a space make up the network range. Each number must be between 1 and 65279, and the first number must be equal to or smaller than the second. The operator in this option can only be "equality" or "inequality."

- **zone** <operator> <zone name> This option allows filtering of the zone name in Get Zone List and RTMP packets. The zone name must be enclosed in quotes (e.g. "My Zone"), no greater than 32 characters long, and cannot contain the Ý symbol or *. The operator in this option can only be "equality" or "inequality."

- **NBPName** <operator> <NBP name> This option allows filtering of the NBP name in an NBP request or reply packet. The NBP name must be between 1 and 32 characters long and enclosed in quotes (e.g. "Laser-Writer"). The name may contain Ý. The operator in this option can only be "equality" or "inequality."

- **NBPType** <operator> <NBP type> This option allows filtering of the NBP type in an NBP request or reply packet. The NBP name must be between 1 and 32 characters long and enclosed in quotes (e.g. "AFP Server"). The name may contain Ý. The operator in this option can only be "equality" or "inequality."

- **NBPZone** <operator> <zone name> This option allows filtering of the NBP zone name in an NBP request or reply packet. The NBP name must be between 1 and 32 characters long and enclosed in quotes (e.g. "Administration Zone"). The name may contain Ý. The operator in this option can only be "equality" or "inequality."

- **log** The log option causes the device to log data about the packet to syslog when the condition of the rule is met.

Simple AppleTalk Packet Filter Rule Examples

The following is an AppleTalk packet filter which denies echo packets (type 4) from network 55, and permits everything else.

```
deny srcnet = 55 type = 4

permit
```

The following is an AppleTalk packet filter which denies NBP lookups for the printer named "Engineering Printer," permits NBP lookups for the printer named "HP Printer" by the NBP zone "Sales," and permits everything else.

```
deny NBPName = "Engineering Printer"

permit NBPName = "HP Printer" NBPZone = "Sales"

permit
```

AppleTalk Get Zone List Filter Rule Set Examples

AppleTalk Get Zone List filter rules filter what is seen in the Chooser of Macintoshes attached to the network to which the rules are assigned. The example would: deny all zone names from networks 1-10; permit the zone name "Engineering;" deny the zone name "Sales;" permit all networks not equal to 100; and permit everything else.

```
deny net-range = 1 10

permit zone = "Engineering"

deny zone = "Sales"

permit network != 100

permit
```

AppleTalk RTMP Filter Rule Set Examples

AppleTalk RTMP filter rules can be used to limit the network numbers that are allowed into the routing table or to be advertised from the device. The example performs the following actions: deny networks with a number of 100; permit networks between 200 and 300; deny networks numbered greater than 301; and permit everything else.

```
deny network = 100

permit net-range = 200 300

deny network > 301

permit
```

AppleTalk Filtering: Ethernet Dialog Box

AppleTalk Filtering: WAN Dialog Box

AppleTalk Filtering: VPN Dialog Box

AppleTalk Filtering: Bridge Dialog Box

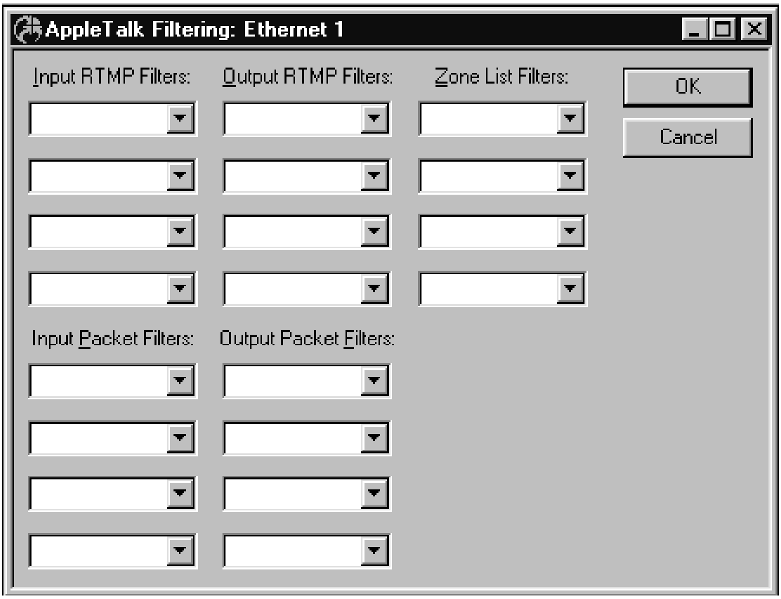

Interface AppleTalk Filtering Configuration Dialog Box

To access this dialog box, select Interface/Filtering/AppleTalk Filtering from the Device View.

Input RTMP Filters

This set of pull-downs allows you to select previously defined sets of routing (RTMP) filter rules. These rules will be applied to information arriving on this interface. Up to four sets of rules can be selected.

Output RTMP Filters

This set of pulldowns allows you to select previously defined sets of routing (RTMP) filter rules. These rules will be applied to information which is to be sent on this interface. Up to four sets of rules can be selected.

Zone List Filters

This set of pulldowns allows you to select previously defined sets of get-zone-list filter rules. These rules will be applied to replies to AppleTalk get-zone-list requests which are received on this interface. Up to four sets of rules can be selected.

Input Packet Filters

This set of pulldowns allows you to select previously defined sets of packet filter rules. These rules will be applied to packets arriving on this interface. Up to four sets of rules can be selected.

Output Packet Filters

This set of pulldowns allows you to select previously defined sets of packet filter rules. These rules will be applied to packets which are to be sent on this interface. Up to four sets of rules can be selected.

Chapter 14 - General

Physical RS-232 Configuration: WAN Dialog Box

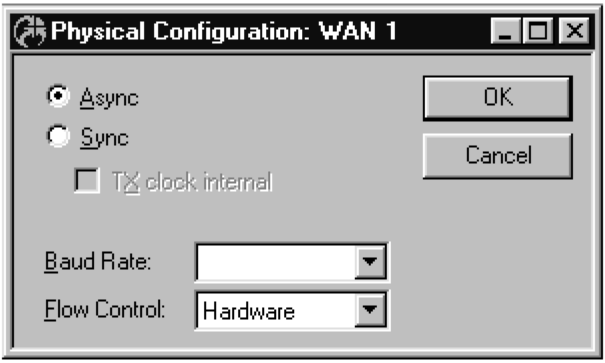

Physical Configuration: WAN 1

● Async

○ Sync

□ TX clock internal

Baud Rate:

Flow Control: Hardware

Physical RS-232 Configuration: WAN Dialog Box

To access this dialog box, select WAN/Physical Configuration from the Device View.

> **Async/Sync**

This set of radio buttons determines whether this interface will use the asynchronous or synchronous mode of communication.

- If **Async** is selected, the interface will communicate asynchronously (using start and stop bits) with the device it is connected to. This mode of communication is typically used by modems.
- If **Sync** is selected, the interface will communicate synchronously (using a separate clock) with the device it is connected to. This mode of communication is typically used by CSU/DSU’s and ISDN Terminal Adapters.

❖ **Note:** *Interfaces set for asynchronous operation do not use parity, and use one stop bit.*

❖ **Note:** *Some high-speed WAN interfaces (i.e. V.35) may only support synchronous communications. This set of radio buttons will not appear in the CompatiView Physical Configuration: WAN Dialog Box for these interfaces.*

Tx Clock Internal (Sync Only)

This parameter determines whether the interface will source a clock signal or expect to receive an external clock.

- If **checked**, the interface will expect to source a clock, and will ignore an external clock signal.
- If **unchecked**, the interface will expect to receive an external clock. This is the default setting.

❖ **Note:** *In addition to this setting, some WAN interfaces require internal hardware jumpers to be changed in order to source a clock signal. Check the Installation Reference Guide for the device.*

Baud Rate

This pull-down menu determines the speed of the port's internal clock. In **Async** operation, this value must match the baud rate of the external communications device. In **Sync** operation this value is ignored unless **Tx Clock Internal** is checked, in which case it determines the speed of the clock which is sourced.

Flow Control (Async Only)

This setting determines the type of flow control used on interfaces set for Async operation.

- If **none** is selected, the interface will not pace the rate at which it sends characters.
- If **hardware** is selected, the interface will use the state of the CTS (Clear To Send) line to determine whether characters may be sent. This is the default setting.
- If **XOn/XOff** is selected, the interface will trap XOn and XOff characters to determine whether characters may be sent. This method is also known as "software" flow control.

❖ **Note:** *You may also need to configure your communications device (through switch settings or internal registers) to run with software flow control. Software flow control is* ***not*** *recommended at speeds above 9600 Baud.*

Physical T1 Configuration: WAN Dialog Box

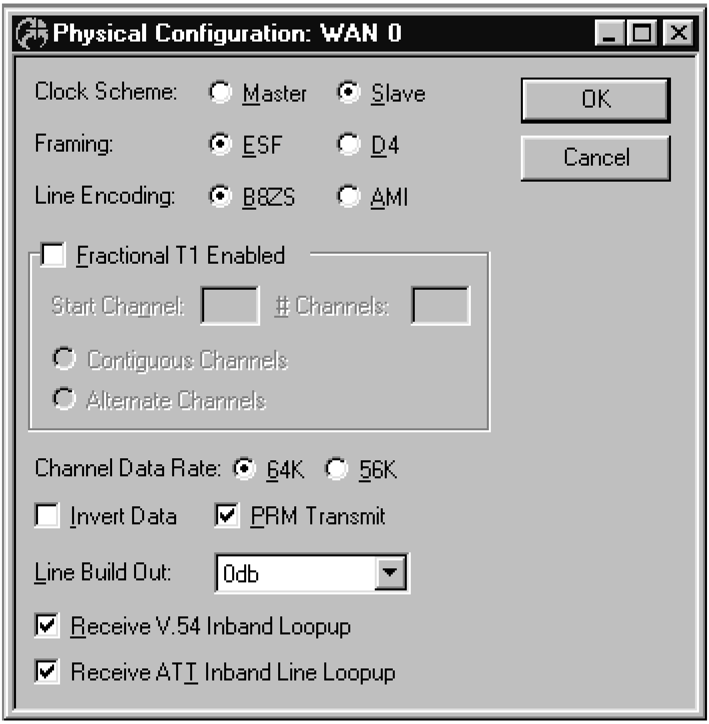

Physical Configuration: WAN 0

Clock Scheme: ○ Master ● Slave

Framing: ● ESF ○ D4

Line Encoding: ● B8ZS ○ AMI

□ Fractional T1 Enabled

Start Channel: # Channels:

○ Contiguous Channels

○ Alternate Channels

Channel Data Rate: ● 64K ○ 56K

□ Invert Data ☑ PRM Transmit

Line Build Out: 0db

☑ Receive V.54 Inband Loopup

☑ Receive ATT Inband Line Loopup

Physical T1: WAN Configuration Dialog Box

To access this dialog box, select WAN/Physical Configuration from the Device View.

Since many of the settings for a T1 line are dependent upon the service provided by your ISP or telco, you may need to contact them to find out the appropriate specifications. Unless otherwise noted, both ends of a T1 WAN connection should have the same physical configuration settings.

> Clock Scheme

This set of radio buttons determines whether this interface will source clock onto the T1 line (dry line operation) or accept clock from an external source on a T1 line.

- If **Master** is selected, the interface will source clock onto the line.
- If **Slave** is selected, the interface will sync to the clock received on the line. The default setting is Slave.

❖ **Note:** *Units connected to telco lines should always be set for slave mode. Units driving a dry line should have one end set to master and the other set to slave.*

> Framing

This parameter determines the type of T1 framing to be used on the interface. Both ends of a WAN connection must be configured with the same framing format

- If **ESF** is selected, the interface will expect extended superframe framing on the line. ESF is the preferred format because it offers a Facility Data Link which can provide performance monitoring, error checking and other features. ESF is the default.
- If **D4** is selected, the interface will expect the older D4 superframe framing. D4 may be the only framing format available in some areas.

> Line Encoding

This parameter determines the type of T1 encoding to be used on the interface.

- If **B8ZS** is selected, the interface will expect this type of encoding on the line. With B8ZS, the **Channel Data Rate** should be set to 64 Kbps. This is the default setting.
- If **AMI** is selected, the interface will expect this type of encoding. With AMI, the **Channel Data Rate** should be set to 56 Kbps if the line is Full T1 or if **Contiguous** channels are being used on a Fractional T1 line. If **Alternate** channels are being used on a Fractional T1 line, then the **Channel Data Rate** can be set to 64 Kbps.

Fractional T1 Enabled

This checkbox enables fractional T1 operation, where the device's built-in CSU/DSU will not use all of the T1 channels in the T1 data stream.

Start Channel & Number of Channels

This pair of edit boxes determines which T1 channel the internal CSU/DSU will use as the beginning of its T1 fraction, and how many channels it will occupy.

Contiguous Channels or Alternate Channels

This set of radio buttons determine whether the T1 fraction will occupy every channel starting with the requested channel, or every other channel. If more than 12 channels will be used, the Contiguous Channels radio button must be selected.

- If **Contiguous** is selected, the T1 fraction will occupy every channel beginning at the Start Channel. This is the default setting.
- If **Alternate** is selected, the T1 fraction will occupy every other channel beginning at the Start Channel.

> Channel Data Rate

These two radio buttons determine whether the internal CSU/DSU will use 64Kbps channels or 56Kbps channels. The default is 64Kbps.

Invert Data

This checkbox instructs the CSU/DSU to invert the data it transmits and to expect to receive inverted data on the line. This is sometimes done to insure ones density on the line. The default setting is unchecked.

PRM Transmit

This checkbox enables sending and receiving of Performance Report Messages (PRM) on the Facility Data Link (FDL). This is only possible when ESF framing has been selected. The default is enabled.

Line Build Out

This pulldown sets the expected signal range for the CSU/DSU’s receiver. 0db is the default and will be satisfactory in virtually all telco applications. Other settings may be necessary for dry line applications.

Receive V.54 Inband Loopup

This checkbox instructs the CSU/DSU to respond to V.54 style loopup commands received over the line. The default setting is on.

Receive ATT Inband Line Loopup

This checkbox instructs the CSU/DSU to respond to ATT 64211 style loopup commands received over the line. The default setting is on.

Physical V.35 Configuration: WAN Dialog Box

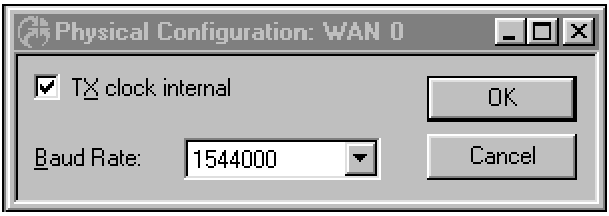

Physical Configuration: WAN 0

☑ TX clock internal

Baud Rate: 1544000

Physical V.35 Configuration: WAN Dialog Box

To access this dialog box, select WAN/Physical Configuration from the Device View.

Tx Clock Internal

This parameter determines whether the interface will source a clock signal or expect to receive an external clock.

- If **checked**, the interface will expect to source a clock, and will ignore an external clock signal.
- If **unchecked**, the interface will expect to receive an external clock. This is the default setting.

Baud Rate

This pull-down menu determines the speed of the port’s internal clock when **Tx Clock Internal** is checked.

Physical DS3 Configuration: WAN Dialog Box

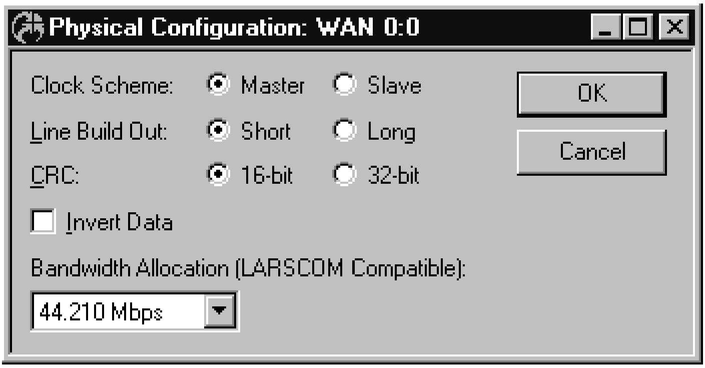
Physical Configuration: WAN 0:0

Clock Scheme: ● Master ○ Slave

Line Build Out: ● Short ○ Long

CRC: ● 16-bit ○ 32-bit

□ Invert Data

Bandwidth Allocation (LARSCOM Compatible): 44.210 Mbps

Physical DS3 Configuration: WAN Dialog Box

To access this dialog box, select WAN/Physical Configuration from the Device View.

Clock Scheme

These radio buttons set whether the DSU will use its own internal clock or obtain the clock from the network to use for the DSU’s DS3 transmit signal towards the network.

- **Master** means an internal clock will be used.
- **Slave** means the clock derived from the DS3 receive signal will be used. This is the default setting

Line Build Out

These radio buttons should be set based on the distance between the device and the DS3 terminal located in your building.

- **Short** should be used for cable lengths from 0 - 100 feet.
- **Long** should be used for cable lengths from 101 - 900 feet.

CRC

These radio buttons control whether the DSU will use a 16-bit or 32-bit frame check sequence. Both ends of a DS3 connection must use the same CRC (Cyclical Redundancy Check) setting. The default is **16 bit**.

Invert Data

This checkbox determines whether data will be inverted. Data inversion can be used to meet pulse density requirements. Always leave this unchecked unless otherwise instructed by your ISP.

- If **checked**, data will be inverted. If a DSU at one end of a DS3 line inverts its data, then the DSU at the other end must do the same.
- If **unchecked**, data will not be inverted. This is the default setting.

Bandwidth Allocation

This pull-down menu allows you to select the data rate for the CSU/DSU. This can be used to set the throughput to match the bandwidth provided by your NSP (Network Service Provider). The values are specified in megabits per second, using an underscore (_) as the decimal point (e.g., 3_158 is 3.158 Mbps). Both ends of the DS3 connection must have the same rate specified. Unless the remote end is a Larscom CSU/DSU (or equivalent) or another Compatible Systems DS3 interface, the default setting of **44_210** must be used.

System Configuration Dialog Box

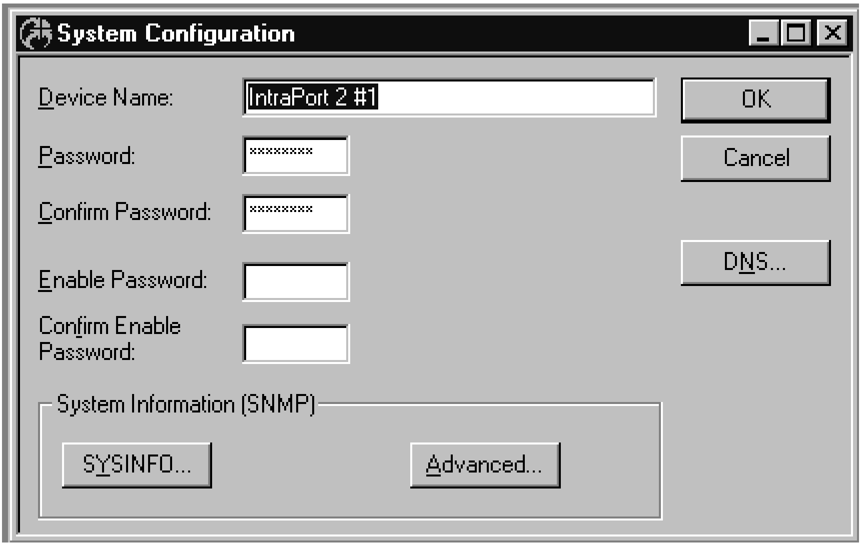

System Configuration

Device Name: IntraPort 2 #1

Password: *********

Confirm Password: *********

Enable Password:

Confirm Enable Password:

System Configuration Dialog Box

To access this dialog box, select Global/System Configuration from the Device View.

> Device Name

This is the name which is used to advertise this device on both AppleTalk and IPX networks. Thus, it is the name CompatiView displays in the Open - Device screen (accessed from the File menu).

> Password

This is the main password used to access the device from CompatiView and from the command line (either Telnet or auxiliary port operation). This login level will allow a user to display tables and statistics, but does not permit a user to view or make any changes to the configuration.

❖ **Note:** *If you lose the password for a Compatible Systems device, you can enable the default* ***letmein*** *password for five minutes by setting the switch marked "Test" or "Mode" on the back of the device to 9 and restarting the device. Make sure you set the switch back to 0 after you have set a new password into the device.*

> Confirm Password

This box is used to confirm the entered **Password**.

> Enable Password

This password will enable supervisor mode for viewing or making changes to a device's configuration. If no **Enable Password** is created, then the **Password** will be used.

> Confirm Enable Password

This box is used to confirm the entered **Enable Password**.

SNMP Configuration

SNMP System Info Configuration Dialog Box

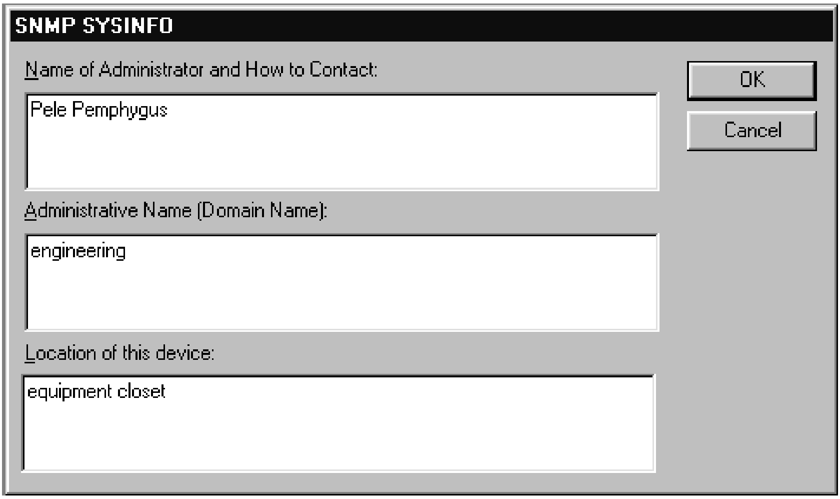

Name of Administrator and How to Contact: Pele Pemphygus

Administrative Name (Domain Name): engineering

Location of this device: equipment closet

SNMP System Info Configuration Dialog Box

To access this dialog box, select Global/System Configuration from the Device View, and then select the **SYSINFO** button.

The information in this dialog box is returned by the device in response to SNMP (Simple Network Management Protocol) queries from SNMP consoles for the SNMP MIB-II System Group, as specified in RFC 1213. Each of the entries may be up to 255 characters.

> Name of Administrator and How to Contact

This is the name of the contact person for this device, together with information on how to contact the administrator.

> Administrative Name

This is the administratively assigned name for this device. By convention, this is the fully qualified domain name for the device.

> Location of this device

This is the physical location of the device (e.g. telephone closet, 3rd floor).

Advanced SNMP Configuration Dialog Box

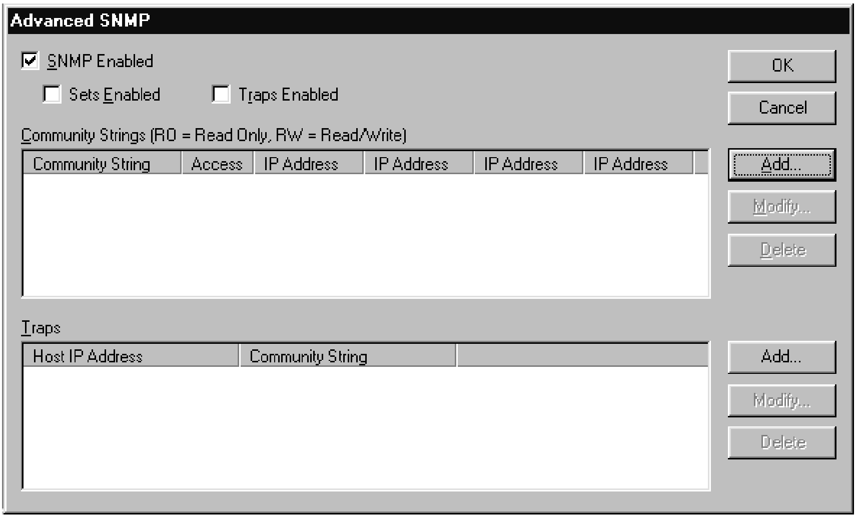

Advanced SNMP

☑ SNMP Enabled

□ Sets Enabled □ Traps Enabled

Advanced SNMP Configuration Dialog Box

To access this dialog box, select Global/System Configuration from the Device View, and then select the **ADVANCED** button.

This dialog box displays Community Strings and Traps, but is not used to add or modify the entries.

To add or modify entries, you must access the Community Strings and/or Traps Dialog Boxes by selecting the **Add...** or **Modify...** buttons in the Advanced SNMP Configuration Dialog Box.

> SNMP Enabled

This checkbox controls whether Advanced SNMP management of the device can be done.

> Sets Enabled

This checkbox controls whether SNMP Sets can be done to the device.

> Traps Enabled

This setting controls whether SNMP Traps will be done by the device when trap conditions are encountered.

Compatible Systems devices support the following SNMP Traps (as outlined in RFC 1157):

- **coldStart** - this will be generated when a restart to save a configuration or software download is accomplished.
- **warmStart** - this will be generated when a restart event is received.
- **linkDown** - this will be generated from a WAN interface when a link is dropped due to abnormal conditions, such as lost carrier, lost PVC, etc.
- **linkUp** - this will be generated from a WAN interface when a link which was lost due to abnormal conditions comes back up.
- **authenticationFailure** - this will be generated when a protocol message is not properly authenticated.

SNMP Community Strings Configuration Dialog Box

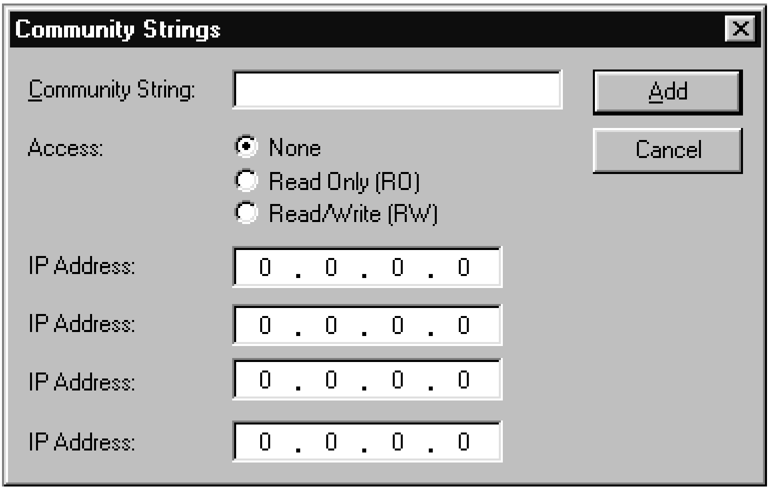

Community String:

Access: ● None ○ Read Only (RO) ○ Read/Write (RW)

IP Address: 0 . 0 . 0 . 0

IP Address: 0 . 0 . 0 . 0

IP Address: 0 . 0 . 0 . 0

IP Address: 0 . 0 . 0 . 0

SNMP Community Strings Configuration Dialog Box

Community String

This is the string associated with the administrator(s) who have access to the SNMP console. It is included in every message and is used, along with the IP address(es) configured below, for access authentication.

Access

This set of radio buttons controls the type of access the administrator(s) within the Community String will have to this device.

- **None** - no access.
- **Read Only (RO)** - receives information such as Traps, but can not do Sets.
- **Read/Write (RW)** - can perform Sets to, and receive Traps from, this device.

IP Address

This is the IP address, or addresses, of the SNMP console. The address is used, along with the Community String, for access authentication. Up to four IP addresses may be entered.

They should be entered in standard IP dotted-decimal notation (e.g. 198.41.9.1). An address with all zeros (0.0.0.0) can be used as a wildcard to allow the Community String access from any console.

SNMP Traps Configuration Dialog Box

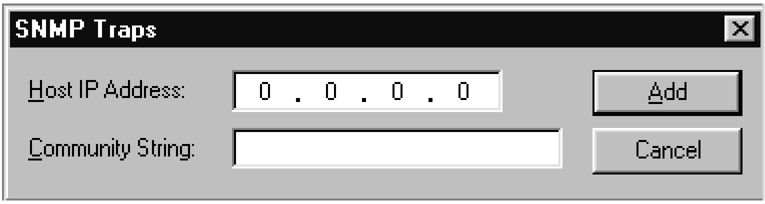

Host IP Address: 0 . 0 . 0 . 0

Community String:

SNMP Traps Dialog Box

Host IP Address

This is the IP address of the SNMP console to which the device will transmit a Trap message whenever one is generated.

It should be entered in standard IP dotted-decimal notation (e.g. 198.41.9.1).

Community String

This is the Community String on the SNMP console to which the Trap message will be sent.

Domain Name Server (DNS) Dialog Box

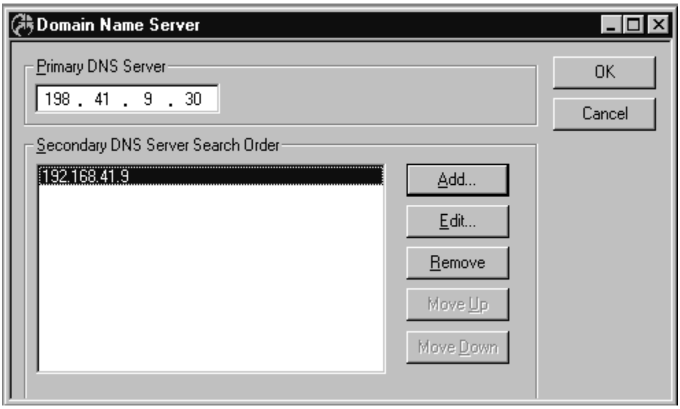

Primary DNS Server: 198 . 41 . 9 . 30

Secondary DNS Server Search Order: 192.168.41.9

Domain Name Server Dialog Box

To access this dialog box, select Global/Domain Name Server from the device view.

DNS allows the device to report DNS names instead of raw IP addresses when using the **Traceroute** command, and also allows the **Ping** command to be optionally issued with a DNS name.

❖ **Note:** *The Traceroute and Ping commands themselves are not supported from CompatiView. To access these commands, use the command line interface via Telnet or the Console port.*

> Primary DNS Server

This is the IP address of the DNS server which should be queried first for the identity of a name or an IP address.

This address should be entered directly into the edit box as four decimal numbers separated by periods – for example 198.238.9.1

Secondary DNS Server Search Order

These are the IP addresses of other DNS servers which should be queried for the identity of a name or an IP address.

Use the **Move Up** and **Move Down** buttons to manipulate the addresses in this list.

To add or modify this list, click on the appropriate button to access the Add TCP/IP DNS Server Dialog Box.

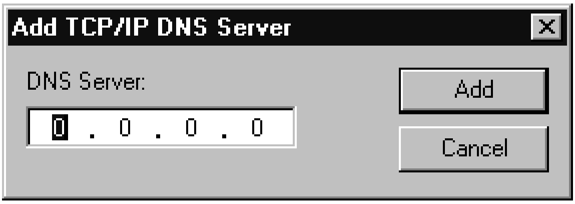

DNS Server: 0 . 0 . 0 . 0

Add TCP/IP DNS Server

Enter the IP address of other DNS servers which should be queried for the identity of a name or an IP address.

Time Server Dialog Box

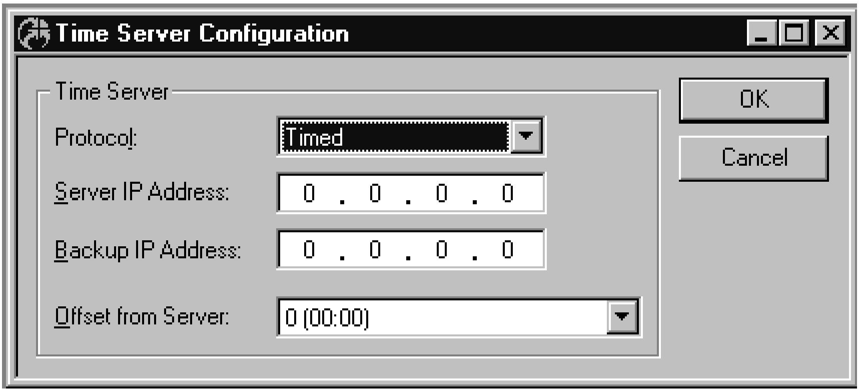

Protocol: Timed

Server IP Address: 0 . 0 . 0 . 0

Backup IP Address: 0 . 0 . 0 . 0

Offset from Server: 0 (00:00)

Time Server Configuration Dialog Box

To access this dialog box, select Global/Time Server Configuration from the Device View.

This dialog box is used to enable the setting of the device’s internal clock from a network time server. The device’s time server will connect to most UNIX systems running "inetd" using either the time server port (UDP 37) or NTP port (UDP 123).   Automatic daylight savings adjustment is not supported.

Protocol

This pulldown identifies the type of time server protocol to use. In most cases, the time server being used will dictate the protocol type. UNIX servers generally use **Timed**. Windows servers generally use **SNTP** (Simple Network Time Protocol). The default is Timed.

Server IP Address

This field is used to specify the IP address of the primary time server. It is recommended that you use a time server which is local to your network.

Backup IP Address

This field is used to specify the IP address of the backup time server. All time requests go to the primary server first. If there is no response the backup will be used. This address is optional.

Offset from Server

Most time servers return GMT. You can use this option to set the device to local time. Accepted values range from -720 to 720 minutes.

RADIUS Configuration Dialog Box

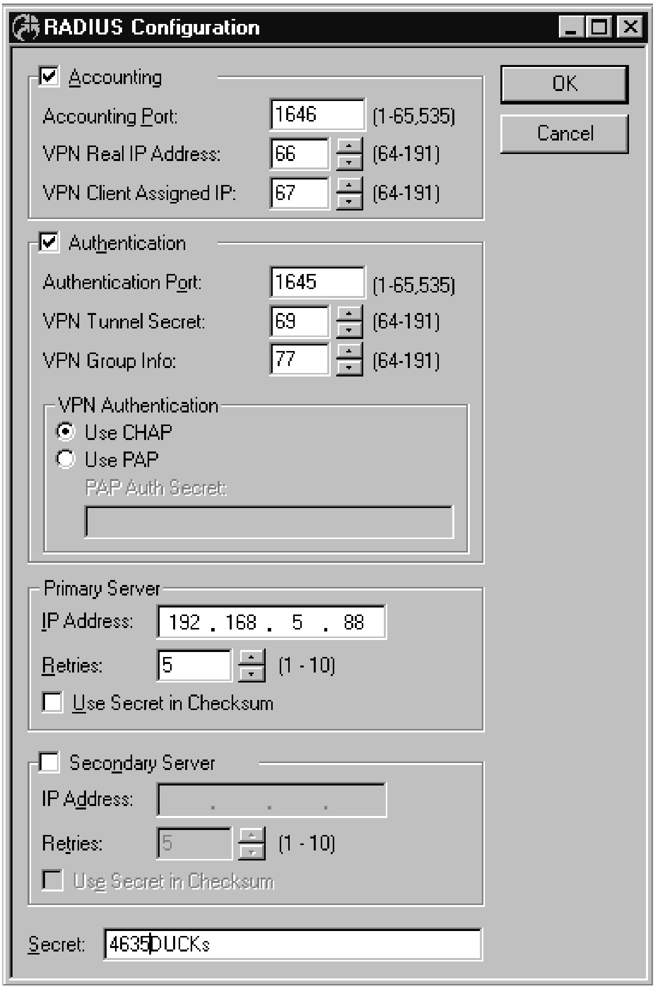

RADIUS Configuration

☑ Accounting

Accounting Port: 1646 (1-65,535)

VPN Real IP Address: 66 (64-191)

VPN Client Assigned IP: 67 (64-191)

☑ Authentication

Authentication Port: 1645 (1-65,535)

VPN Tunnel Secret: 69 (64-191)

VPN Group Info: 77 (64-191)

VPN Authentication

● Use CHAP

○ Use PAP

PAP Auth Secret:

Primary Server

IP Address: 192 . 168 . 5 . 88

Retries: 5 (1 - 10)

□ Use Secret in Checksum

□ Secondary Server

IP Address: . . .

Retries: 5 (1 - 10)

□ Use Secret in Checksum

Secret: 4635DUCKs

RADIUS Configuration Dialog Box

To access this dialog box, select Global/RADIUS from the Device View.

RADIUS (Remote Authentication Dial In User Service) can be used for remote access authentication using PAP or CHAP and for remote access accounting. The device acts as a client and exchanges packets with a RADIUS server running on an external host computer.

The device can be configured with a primary and a secondary server. If the device is unable to reach the primary server, it will attempt to use the secondary server if one has been configured.

❖ **Note:** *RADIUS servers are available in the public domain, and can also be purchased from a variety of commercial suppliers.*

Accounting

This setting determines whether the device will attempt to exchange user accounting information with a RADIUS server.

- If **checked**, each time a user logs into the device, a record of their login is sent to the RADIUS server where it is catalogued.

Accounting Port

This edit box allows you to set the UDP port on the RADIUS server(s) on which accounting information will be exchanged. The default is 1646.

VPN Real IP Address

This value sets the attribute number for the reporting of the actual IP address of an IntraPort user. This attribute number must also be set up in the RADIUS server's dictionary file. If this number has been set both here and in the RADIUS server's dictionary file, then the actual IP address of a user will be reported by the VPN Client software and will be recorded by the RADIUS server. The value may range between 64 and 191. The default is 66.

VPN Client Assigned IP

This value sets the attribute number for the reporting of the IP address which the IntraPort server assigns to an IntraPort user. This attribute number must also be set up in the RADIUS server's dictionary file. If this number has been set both here and in the RADIUS server's dictionary file, then the assigned IP address will be reported by the VPN Client software and will be recorded by the RADIUS server. The value may range between 64 and 191. The default is 67.

Authentication

This setting determines whether the device will exchange user authentication information with a RADIUS server.

- If **checked**, user authentication information will be exchanged.

Authentication Port

This edit box allows you to set the UDP port on the RADIUS server(s) on which authentication information will be exchanged. The default is 1645.

VPN Tunnel Secret

This value sets the attribute number for the VPN tunnel secret. The tunnel secret is a shared secret between the VPN Client and the RADIUS server which is used for authentication of tunnel connections. This attribute number must also be set up in the RADIUS server’s dictionary file. The value may range between 64 and 191. The default is 69.

VPN Group Info

This value sets the attribute number for the VPN group configuration. The group configuration defines tunneling profiles for a group of one or more VPN Client users. This attribute number must also be set up in the RADIUS server’s dictionary file. The value may range between 64 and 191. The default is 77.

VPN Authentication

This set of radio buttons sets the authentication protocol to be used for validation of remote VPN Client users to the RADIUS server.

- If **Use CHAP** is selected, CHAP will be used to validate remote VPN Client users to the RADIUS server.
- If **Use PAP** is selected, PAP will be used to validate remote VPN Client users to the RADIUS server. This should only be used for an older RADIUS server which does not support CHAP authentication.

PAP Authentication Secret

This is the secret used to authenticate and encrypt packets before they are passed on to the RADIUS server. The PAP authentication secret can be a string from 1 to 255 ASCII characters in length.

> Primary Server IP Address

The device will attempt to contact this RADIUS server first when it needs to exchange RADIUS information. The address should be entered in dotted-decimal notation (e.g. 198.238.41.7).

> Primary Server Retries

The device will try to resend a packet if the primary RADIUS server doesn’t acknowledge it within a timeout period. The timeout period for packets 1 through 10 is (in seconds): 1, 1, 2, 2, 3, 3, 4, 4, 5, 5.

If the retry limit is reached and a secondary server is configured, the device will attempt to communicate with the secondary server.

Possible values range between 1 and 10 with a default of 5.

Use Secret in Checksum

Some RADIUS servers calculate packet validation checksums using both the secret value and the packet data. Earlier RADIUS servers typically do not. Check the documentation for your RADIUS server to determine whether this parameter should be set.

- If **checked**, packet checksums will be calculated using both the data and the checksum.

Secondary Server

The device may be configured to use a secondary server if the primary server cannot be contacted.

- If **checked**, a secondary server can be configured and will be used in the event the primary server cannot be contacted.

Secondary Server IP Address

The device will attempt to contact this RADIUS server if the primary server does not respond after the configured number of primary server retries. The address should be entered in dotted-decimal notation (i.e. 198.238.41.7).

Secondary Server Retries

The device will try to resend a packet if the secondary RADIUS server doesn’t acknowledge it within a timeout period. The timeout period for packets 1 through 10 is (in seconds): 1, 1, 2, 2, 3, 3, 4, 4, 5, 5.

Possible values range between 1 and 10 with a default of 5.

Use Secret in Checksum

Some RADIUS servers calculate packet validation checksums using both the secret value and the packet data. Earlier RADIUS servers typically do not. Check the documentation for your RADIUS server to determine whether this parameter should be set.

- If **checked**, packet checksums will be calculated using both the data and the checksum.

Secret

The secret is the shared secret used by the device and RADIUS server to validate packets exchanged between them. This secret must match the client secret configured in the RADIUS server. It can be from 1 to 31 ASCII characters in length.

SecurID Configuration Dialog Box

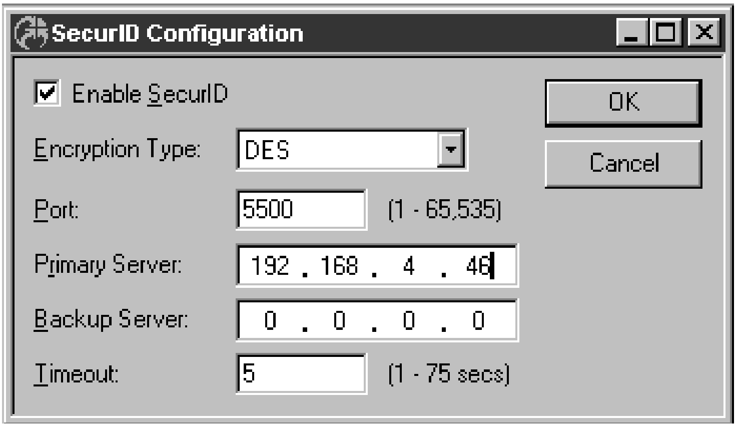

SecurID Configuration

☑ Enable SecurID

Encryption Type: DES

Port: 5500 (1 - 65,535)

Primary Server: 192 . 168 . 4 . 46

Backup Server: 0 . 0 . 0 . 0

Timeout: 5 (1 - 75 secs)

SecurID Configuration Dialog Box

To access this dialog box, select Global/SecurID from the Device View.

All IntraPort VPN Access Servers and the VPN Client software are SecurID-ready. SecurID is Security Dynamic’s proprietary system which requires ACE/Server software and SecurID tokens to perform dynamic two-factor authentication.

> Enable SecurID

This checkbox determines whether SecurID authentication will be performed by the device.

Port

This edit box allows you to set which UDP port on the ACE/Server will be used to exchange information. The default is 5500. The value may range between 1 and 65,535.

Encryption Type

This edit box allows you to select the encryption algorithm for data exchanged between the IntraPort and the ACE/Server. **DES** specifies that the DES algorithm will be used to scramble the data in both directions. **SDI** specifies that Security Dynamic’s propriety algorithm will be used. The default is **DES**.

> Primary Server

The device will attempt to contact this SecurID server first when attempting to authenticate a user. The address should be entered in dotted-decimal notation (i.e. 198.238.41.7).

If the timeout period is reached and a secondary server is configured, the device will attempt to communicate with the backup server.

Backup Server

The device will attempt to contact this SecurID server if the primary server does not respond after the configured timeout period. The address should be entered in dotted-decimal notation (i.e. 198.238.41.7).

Timeout

This edit box allows you to set which UDP port on the ACE/Server will be used to exchange information. the number of seconds the device will wait before trying the backup ACE/Server. The default is 5. The value may range between 1 and 75.

NAT Configuration Dialog Box

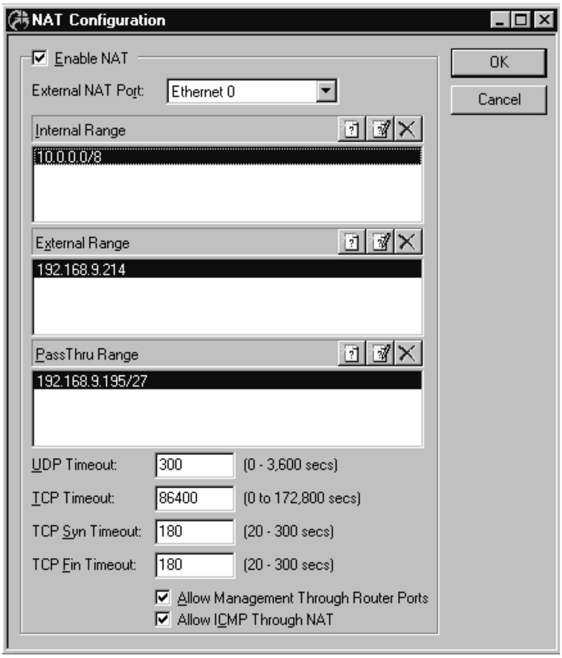

☑ Enable NAT

External NAT Port: Ethernet 0

Internal Range
10.0.0.0/8

External Range
192.168.9.214

PassThru Range
192.168.9.195/27

UDP Timeout: 300 (0 - 3,600 secs)

TCP Timeout: 86400 (0 to 172,800 secs)

TCP Syn Timeout: 180 (20 - 300 secs)

TCP Fin Timeout: 180 (20 - 300 secs)

☑ Allow Management Through Router Ports

☑ Allow ICMP Through NAT

NAT Configuration Dialog Box

New Button

Modify Button

Delete Button

To access this dialog box, select Global/NAT Configuration from the Device View.

NAT allows internal networks which use private IP addresses to be translated into a valid external “global” IP address (or addresses). (See RFC 1918 "Address Allocation for Private Internets" for more information about private IP addresses.) This can allow a private network to provide Internet access through a single "official" IP address. It can also function as a minimal firewall by limiting access to the internal network from external networks while allowing the internal network easy access to the Internet.

> Internal Range

This is the address range of the internal NAT network. This range will be translated into the range of IP addresses defined by the External Range. Any interface or subinterface on the device which is part of the same network as the Internal Range is considered to be an internal NAT port.

This window displays a list of all entered Internal Range but is not used to add or modify the entries. To add or modify the entries, you must access the NAT Map Dialog Box by selecting the **New** or **Modify** buttons above the Internal Range window.

External Range

This is the address range of the external NAT network. This range will be translated into the range of IP addresses defined by the Internal Range. The address or range of addresses specified must be a valid, globally recognized Internet address (or addresses) which can be routed on the network.

If only a single Internet IP address is available, then the **External Range** must be the same as the IP address on the IP port communicating with the Internet. In this case, care must be taken not to create a one-to-one translation pair using this IP address in the NAT Mapping Dialog Box (under Global/Nat Mapping). If a range of addresses is specified, the NAT software makes the decision about which Internet address is assigned to outgoing packets.

This window displays a list of all entered External Range but is not used to add or modify the entries. To add or modify the entries, you must access the NAT Map Dialog Box by selecting the **New** or **Modify** buttons above the External Range window.

PassThru Range Addresses

This is the address range which may pass through the external NAT port without being translated. This is used when the NAT router has an IP interface (or interfaces), in addition to the NAT internal port and NAT external port, which is connected to part of the local network which is configured with global IP addresses.

❖ **Note:** *If an IP address or range of addresses is included in both the* **External Range** *and* **PassThru Range**, *NAT will treat the IP address(es) as being members of the* **External Range** *only.*

This window displays a list of all entered PassThru Range but is not used to add or modify the entries. To add or modify the entries, you must access the NAT Map Dialog Box by selecting the **New** or **Modify** buttons next to the PassThru Range window.

TCP Timeout

This edit box allows you to set the amount of time to lapse without any IP Network Address Translations using this NAT session before the router removes an active NAT session for TCP. The value may range from 0 to 172,800 seconds (48 hours). A value of zero will cause TCP NAT sessions to never be removed due to inactivity. Extending the amount of time will cause more router memory to be used by the NAT translation session database. The default is 86,400 seconds (24 hours).

UDP Timeout

This edit box allows you to set the amount of time to lapse without any IP Network Address Translations using this NAT session before the router removes an active non-TCP NAT session. The value may range from 0 to 3600 seconds (1 hour). A value of zero will cause non-TCP NAT sessions to never be removed due to inactivity. Extending the amount of time will cause more router memory to be used by the NAT translation session database. The default is 300 seconds (5 minutes).

TCP Syn Timeout

This edit box allows you to set the amount of time to lapse without a response to a SYN TCP packet before the router removes an active NAT session for TCP. The value may range from 20 to 300 seconds. The default is 180 seconds (3 minutes).

TCP Fin Timeout

This edit box allows you to set the amount of time to lapse without a response to a FIN TCP packet before the router removes an active NAT session for TCP. The value may range from 20 to 300 seconds. The default is 180 seconds (3 minutes).

Allow Management Through Router Ports

This checkbox allows communication with the router through the IP addresses of the router’s ports. This allows the user to communicate with the router (e.g., establish a telnet session with the router). The default is **checked**.

Allow ICMP Through NAT

This checkbox allows external workstations/routers to ping workstations/routers in the internal NAT network if a one-to-one translation pair allowing such a translation has been set using the NAT Mapping Dialog Box. The default is **checked**. The workstation/router on the internal NAT network will not be allowed to respond to a ping if this parameter is **unchecked**.

NAT Range Dialog Box

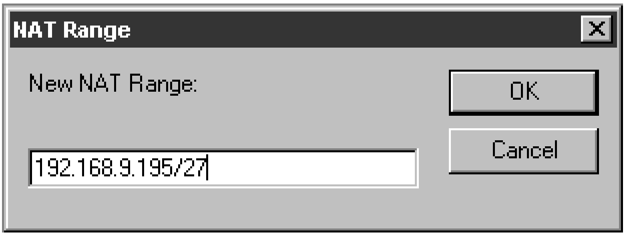

192.168.9.195/27

NAT Range Dialog Box

You can access the NAT Range Dialog Box by selecting one of the **New** or **Modify** buttons in the NAT Configuration Dialog Box (under Global/NAT Configuration). This dialog box allows you to enter a NAT address range. It can be a single IP address or a range of addresses.

The address range may be specified in several different ways:

- Addresses can be specified in normal dotted-decimal notation. If the rightmost components are 0, they are treated as wild cards (e.g., 128.138.12.0 matches all hosts on the 128.138.12 subnet).
- An inclusive range of addresses can be specified using a "dash notation" in the form of #.#.#.{# -#}. For example, 10.5.3.{1-30} would be parsed as the IP addresses 10.5.3.1, 10.5.3.2,..... 10.5.3.29, and 10.5.3.30 (and every IP address in between). Each of these parsed addresses would have a mask of /32 or 255.255.255.255
- Addresses may also be specified as a hexadecimal number (e.g., 0x82cc0801 matches the host address 130.204.8.1).
- A bit field can also be used to indicate a range of addresses by denoting the top or most significant bits which define the range. For example, an address specified as 192.15.32.0/19 would indicate a range from 192.15.32.1 to 192.15.63.255.

NAT Mapping Dialog Box

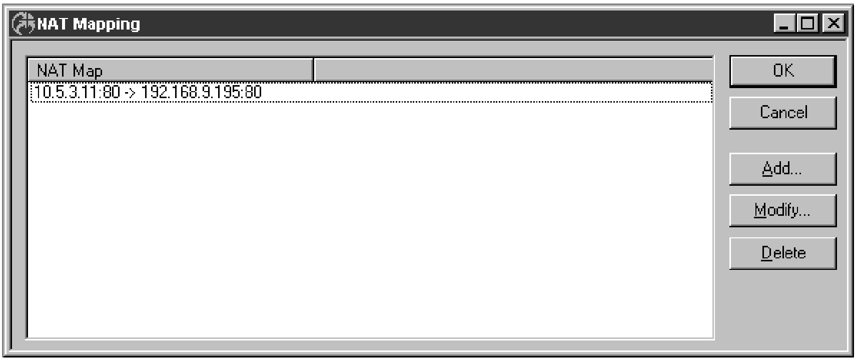

NAT Map

10.5.3.11:80 -> 192.168.9.195:80

NAT Mapping Dialog Box

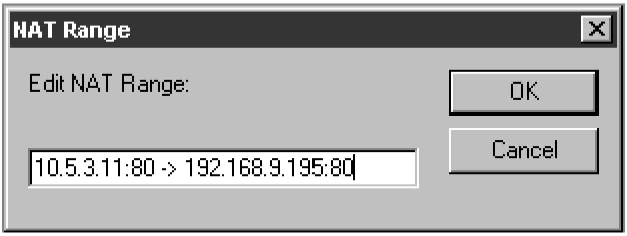

Edit NAT Range:

10.5.3.11:80 -> 192.168.9.195:80

NAT Range Dialog Box

You can access the NAT Mapping Dialog Box by selecting Global/NAT Configuration from the Device List. This window displays a list of all entered one-to-one NAT Mapping translation pairs but is not used to add or modify the entries. To add or modify the entries, you must access the NAT Range Dialog Box by selecting the **Add...** or **Modify...** buttons.

These one-to-one translation pairs allow the user to provide access from the internal or external network to selected parts of the NAT internal network, such as a web server.

Each translation pair must be entered using the following syntax:

<internal IP address> [/*<bits>* | :*<port>*] [-> | =] *<external IP address>* [/*<bits>* | :*<port>*]

<*internal IP address*>

This is the IP address on the internal network to be mapped to the external IP address. It must be entered first, followed by " -> " or " = " and the external IP address. The *internal IP address* must be within the range (or ranges) of IP addresses defined by the **Internal Range Addresses**. IP addresses must be specified in normal dotted-decimal notation. If the rightmost components are 0, they are treated as wild cards (e.g., 128.138.12.0 includes all devices on the 128.138.12 subnet).

<*external IP address*>

This is the IP address on the external network to be mapped to the internal IP address. The *external IP address* must be within the range of IP addresses defined by the **External Range Addresses**.

❖ **Note:** *If only a single external IP address is available for the NAT router, do* ***not*** *map that IP address to an internal IP address because you will no longer be able to communicate with the router. Mapping single ports of the single external IP address to internal IP address:port combinations (e.g., creating access to a web server in the internal NAT network) is acceptable, however.*

:<*port*>

The :*port* option allows an individual socket (IP address and port combination) to be mapped as part of a translation pair.

❖ **Note:** *An IP address:port combination cannot be paired with an IP address range (even if that range is a single IP address). It can only be paired with another IP address:port combination.*

/<*bits*>

The */bits* option allows a range of IP addresses to be mapped as part of a translation pair. The *bits* field denotes the top or most significant bits which define the range. For example, an address specified as 192.15.32.0/19 would indicate a range from 192.15.32.1 to 192.15.63.255.

NAT Mapping Translation Pair Examples

The following example shows one IP address being translated into another.

```
[ NAT Mapping ]

10.5.3.20 -> 198.41.9.194
```

The following example shows individual sockets (IP address and port combination) being mapped as a translation pair.

```
[ NAT Mapping ]

10.5.3.10:80 -> 198.41.9.195:80
```

The following example shows a range of IP addresses being mapped as a translation pair.

```
[ NAT Mapping ]

10.5.3.0/29 -> 198.41.9.200/29
```

Logging Configuration Dialog Box

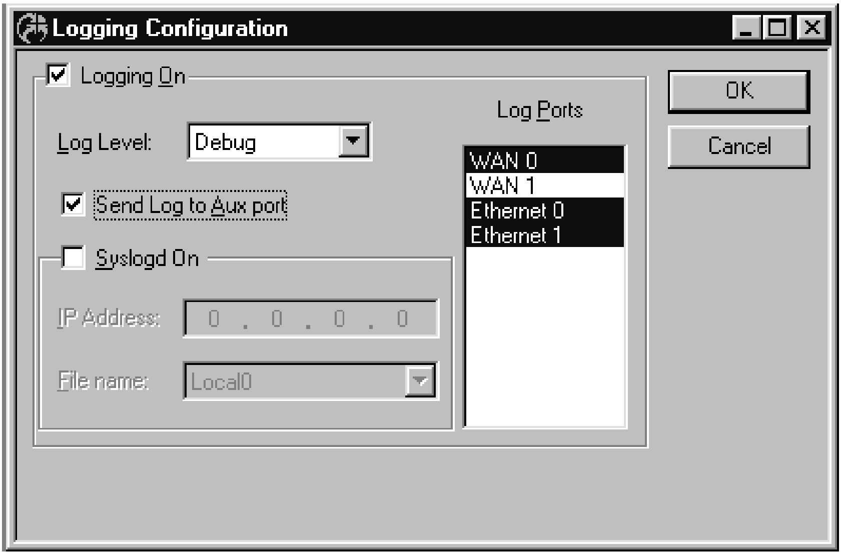

Logging Configuration

☑ Logging On

Log Level: Debug

☑ Send Log to Aux port

□ Syslogd On

IP Address: 0 . 0 . 0 . 0

File name: Local0

Log Ports: WAN 0, WAN 1, Ethernet 0, Ethernet 1

Logging Configuration Dialog Box

To access this dialog box, select Logging from the Device View.

> Logging On

This setting determines whether the internetworking device will output logging information via any of the possible output methods (as discussed below). Logging is on by default.

Log Level

This pull-down menu selects the detail of the logging information provided.

- The **Notice** setting provides information that may be useful on a day-to-day basis by an administrator but generally does not require any response. Examples include login/logout, serial line resets, and LAN-to-LAN connections. This is the default setting and is suitable for most conditions.

- The **Emergency** level means that you will receive logging information only when the system is unusable. These log messages will help indicate the source of the problem.
- The **Alert** level reports only alert and emergency messages. An alert message requires immediate attention.
- The **Critical** level outputs critical, alert, and emergency messages. A critical condition requires imminent action.
- **Error** messages include exception cases pertaining to violations of protocols or other operational rules. Such violations may include illegal packets and improper command syntax.
- If **Warning** messages are repeated, they require a response. Examples of warning-level messages include network number conflicts and resource allocation problems.
- The **Info** option reports routine information, such as WAN network connect and disconnect messages.
- The **Debug** option logs every action of the device and should not be used on a day-to-day basis since it generates a large number of log messages.

Send Log to Aux Port

This checkbox determines whether the auxiliary port will receive logging messages.

Syslogd On

This checkbox determines whether the logging messages will be sent to a UNIX host system running the syslog daemon.

IP Address (Syslogd On Only)

This is the IP address of the UNIX system which is running the syslog daemon, in dotted-decimal notation (i.e. 198.238.41.7).

File (Syslogd On Only)

This pull-down menu determines which syslogd file the device’s logging messages will be written into.

Log Ports

This list shows the ports for which logging information will be generated. If an interface is highlighted, logging information will be generated for that interface. To select or deselect more than one interface, press Control while clicking on the interface.

LDAP Configuration

This section configures LDAP (Lightweight Directory Access Protocol) parameters into a device. LDAP can be used to serve configurations to a Compatible Systems device. LDAP configuration server settings are set in the LDAP Server Dialog Box.

LDAP can also be used for VPN user authentication. LDAP user authentication is configured with the LDAP Authentication Dialog Box.

LDAP Server Dialog Box

Each LDAP configuration specifies an LDAP server and information about the configuration to be served. The configuration can be a full device configuration, or just a portion of one. When new configurations are added, the device's configuration is rebuilt to include the one that was just added.

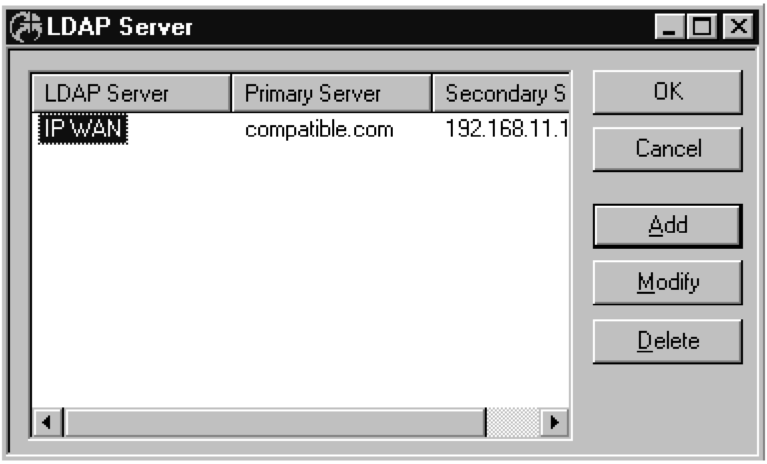

LDAP Server	Primary Server	Secondary S
IP WAN	compatible.com	192.168.11.1

LDAP Server Dialog Box

To access this dialog box, select Global/LDAP Server from the Device View.

To add to or modify this list of LDAP configuration servers, click on the appropriate button to open the Add LDAP Server Dialog Box.

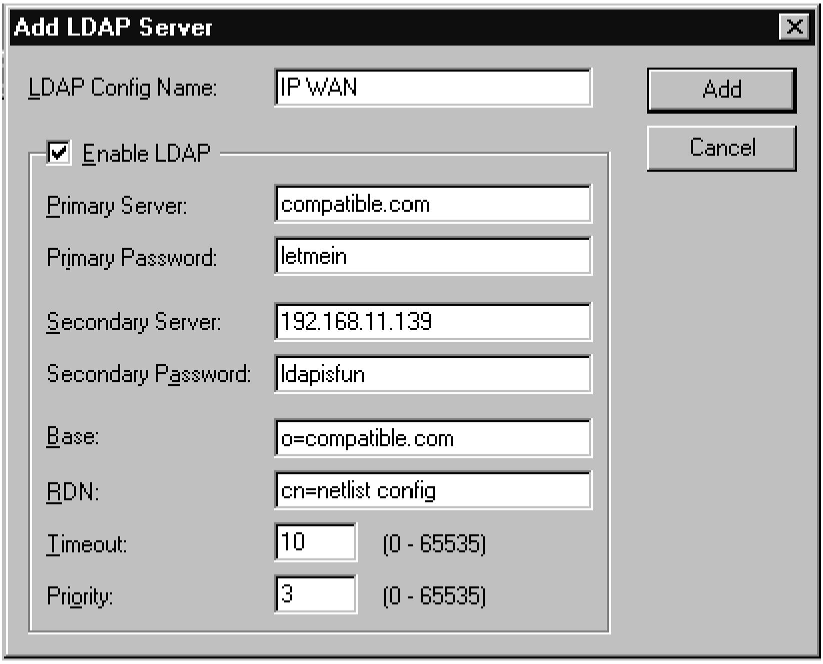

Add LDAP Server

LDAP Config Name: IP WAN

☑ Enable LDAP

Primary Server: compatible.com

Primary Password: letmein

Secondary Server: 192.168.11.139

Secondary Password: ldapisfun

Base: o=compatible.com

RDN: cn=netlist config

Timeout: 10 (0 - 65535)

Priority: 3 (0 - 65535)

Add LDAP Server Dialog Box

LDAP Config Name

This specifies a name which uniquely defines this LDAP configuration. It can be up to 16 characters long.

Enable LDAP

This checkbox enables this entire section. If checked, the settings from this section will be used to get a configuration from an LDAP server. If left unchecked, no settings from this section will be used.

Primary Server

This sets the IP address (e.g., 192.168.9.99) or fully qualified domain name (e.g., monkeywrench.com) of the primary LDAP server which contains the configuration.

Primary Password

This string is used to authenticate the device to the primary LDAP server. If this is not set, then the device will attempt an anonymous bid to the server. The Primary Password may be up to 32 characters long.

Secondary Server

This sets the IP address (e.g., 192.168.9.99) or fully qualified domain name (e.g., monkeywrench.com) of the secondary LDAP server which contains the configuration.

Secondary Password

This string is used to authenticate the device to the secondary LDAP server. If this is not set, then the device will attempt an anonymous bid to the server. The Secondary Password may be up to 32 characters long.

Base

This specifies the portion of the LDAP tree where the configuration is located.

RDN

This string specifies the relative distinguished name used in the LDAP server to identify the entry which contains the configuration.

Timeout

This value is the number of seconds the device will wait for a response from the LDAP server.

Priority

This value specifies which configurations take precedence. When new configurations are added, the device’s configuration is rebuilt to include the one that was just added.

If a new configuration contains a section which contains a higher priority than one already in place, The new configuration (or configuration portion) is added above the one that is already there. This enables higher priority sections to take precedence.

LDAP Authentication Dialog Box

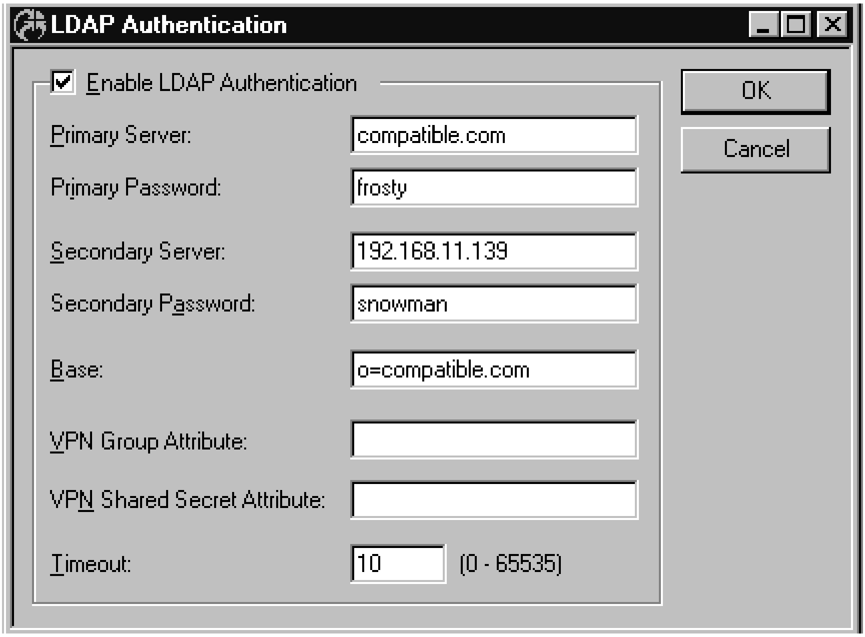

LDAP Authentication

☑ Enable LDAP Authentication

Primary Server: compatible.com

Primary Password: frosty

Secondary Server: 192.168.11.139

Secondary Password: snowman

Base: o=compatible.com

VPN Group Attribute:

VPN Shared Secret Attribute:

Timeout: 10 (0 - 65535)

LDAP Authentication Dialog Box

LDAP authentication is done only if the user cannot be found in the VPN User Authentication Database first. The device acts as a client and exchanges packets with an LDAP server.

❖ **Note:** *For more information on VPN user authentication, refer to* ***Chapter 7 - VPN Client Tunnels****.*

To access this dialog box, select Global/LDAP Authentication from the Device View.

Enable LDAP Authentication

This checkbox enables this entire section. If checked, the settings from this section will be used to get a VPN user authentication from an LDAP server. If left unchecked, no settings from this section will be used.

Primary Server

This sets the IP address (e.g., 192.168.9.99) or fully qualified domain name (e.g., monkeywrench.com) of the primary LDAP server which contains the authentication information.

Primary Password

This string is used to authenticate the device to the primary LDAP server. If this is not set, then the device will attempt an anonymous bid to the server. The Primary Password may be up to 32 characters long.

Secondary Server

This sets the IP address (e.g., 192.168.9.99) or fully qualified domain name (e.g., monkeywrench.com) of the secondary LDAP server which contains the authentication information.

Secondary Password

This string is used to authenticate the device to the secondary LDAP server. If this is not set, then the device will attempt an anonymous bid to the server. The Secondary Password may be up to 32 characters long.

Base

This specifies the portion of the LDAP tree where the authentication information is located.

VPN Group Attribute

This value specifies the attribute name given to the VPN group attribute which has been defined in the LDAP server. There are no standard attributes defined by LDAP for this attribute, so you must specify one. If this field is left blank, the device will assume the attribute name to be “vpngroupattr”.

VPN Shared Secret Attribute

This value specifies the name given to the VPN shared secret attribute which has been defined in the LDAP server. There are no standard attributes defined by LDAP for this attribute, so you must specify one. If this field is left blank, the device will assume the attribute name to be “sharedsecret”.

Timeout

This value is the number of seconds the device will wait for a response from the LDAP server.

Chapter 15 - OSPF

This chapter provides instructions for configuring a network utilizing the OSPF (Open Shortest Path First) Protocol. OSPF uses a link-state algorithm in order to build and calculate the shortest path to all known destinations. Each router in an OSPF area contains an identical link-state database, which is a list of each router's usable interfaces and reachable neighbors.

Unlike RIP updates, OSPF link-state database updates are only sent when routing changes occur, instead of periodically, and the link-state database is updated instantly, rather than gradually, as stale information is timed out.

OSPF routing decisions are based on "cost" which is an indication of the overhead required to send packets across a certain interface. The cost of an interface is calculated based on link bandwidth rather than the number of hops to the destination. The cost can also be configured to specify preferred paths.

OSPF Dialog Box

To access this dialog box, select the interface you wish to configure (Ethernet, WAN, VPN) from the device view, select TCP/IP Routing, then click on the OSPF option button. **OSPF must be configured for each interface on which you want to run OSPF.**

❖ **Note:** *The router will automatically be in the Backbone area, area 0, on each interface. It is also necessary to add configuration information for an interface to be assigned to a non-backbone area.*

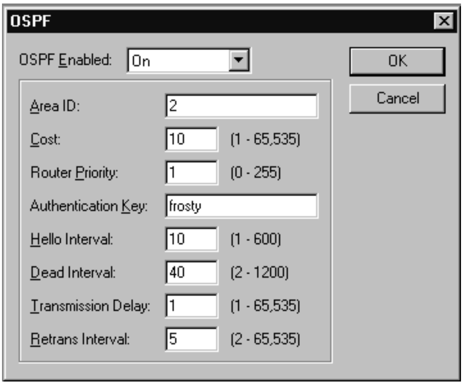

OSPF Enabled: On

Area ID: 2

Cost: 10 (1 - 65,535)

Router Priority: 1 (0 - 255)

Authentication Key: frosty

Hello Interval: 10 (1 - 600)

Dead Interval: 40 (2 - 1200)

Transmission Delay: 1 (1 - 65,535)

Retrans Interval: 5 (2 - 65,535)

OSPF Dialog Box

OSPF Enabled

This pull-down menu sets how the interface will function on a network utilizing OSPF.

- If **On** is specified, the interface will serve as an active interface on an OSPF network. The router will establish adjacencies with other routers. Adjacent routers exchange database information with the Designated Router, which then floods the information to all other routers in their area.
- If **Passive** is specified, the interface will not send out Hello packets and thus will not establish any adjacencies with other routers on that network, even if they are running OSPF.

  A Passive interface will have its network advertise to other OSPF networks. This can be used to have a non-OSPF interface's network advertised into OSPF. A Passive interface must also be associated with an OSPF Area.
- If **Off** is specified, the interface's network is not advertised to the router's other interfaces.

Area ID

This sets the area to which this interface belongs. An area is a generalization of an IP subnetted network. It can be specified as a number between 0 and 0xFFFFFFFF, or as an IP address in dotted-decimal notation. Area 0 is the backbone area and is the default setting.

Cost

This number specifies the priority of one particular path over another path. An OSPF router will choose the gateway with the lowest cost to enter into its routing table. To give preference to a path, set a lower cost on that interface. The value can be a number between 1 and 65,535.

Router Priority

This number sets the router priority and is only used on multi-access networks such as LANs. This establishes whether the router is eligible to become the Designated Router for the LAN. The value for OSPF Router Priority can be a number between 0 and 255.

❖ **Note:** *At least one router on a LAN must have a priority greater than 0 in order for OSPF to work, since there must be a Designated Router.*

The Designated Router is the single router within an area which broadcasts the Link State Advertisement for the area. A priority of 0 means that the router is not eligible. The router with the highest priority becomes the Designated Router.

❖ **Note:** *If a router with a lower priority is the Designated Router and a new router with a higher priority comes online, the Designated Router will not change.*

Authentication Key

This string sets the OSPF packet Authentication Key. The string can be between one and eight alphanumeric characters.

In order to use authentication, the OSPF Authentication Type for this interface should be set to **Simple**. (This is set in the OSPF Area Dialog Box which is discussed in the following section.)

Hello Interval

This value sets the interval, in seconds, that the router sends out OSPF keepalive packets which let other routers know the router is up.The value must be greater than one. The default settings of 10 seconds for a LAN and 30 seconds for a point-to-point connection are recommended for most applications.

Dead Interval

This value sets the length of time, in seconds, that OSPF neighbors will wait without receiving an OSPF keepalive packet from a neighbor before assuming the router is down. **This value must be at least twice the value of Hello Interval.**

❖ **Note:** *Each connected router must have the same values for* **Hello Interval** *and* **Dead Interval** *or the routers will not be able to communicate. If you change the defaults on one router, you must change them on all attached routers within an area.*

Transmission Delay

This value sets the amount of time added to the age of OSPF Link State Update packets before transmission. It is the estimated number of seconds to transmit a packet over the interface.

Retrans Interval

This value sets the interval, in seconds, between retransmission of Link State Update packets. The value must be between 2 and 65,535.

OSPF Area

This dialog box defines a list of OSPF Area names. An area is a generalization of an IP subnetted network within an Autonomous System (AS). An AS is a collection of networks under a common administration sharing a common routing strategy.

All routers within an area have the same link-state database. An interface can only belong to one area, although different interfaces on a router can belong to different areas, making the router an Area Border Router. Area Border Routers disseminate routing information or routing changes between areas.

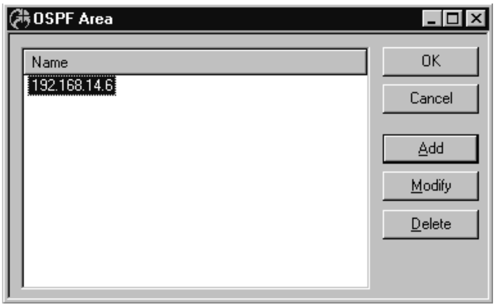
OSPF Area

Name

192.168.14.6

OSPF Area Dialog Box

You can access the OSPF Area Dialog Box by selecting OSPF/OSPF Area from the device view.

To add or modify the entries, select the appropriate button to access the Add OSPF Area Dialog Box.

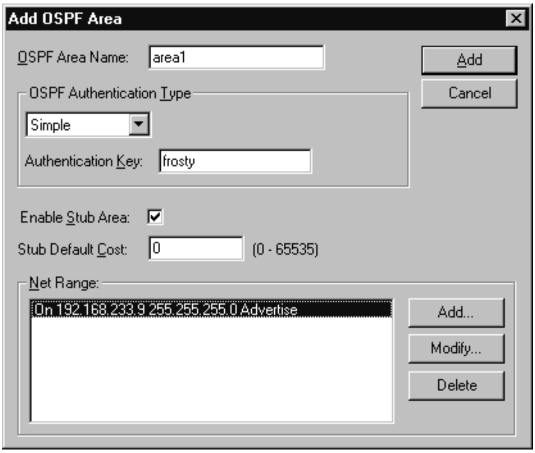
Add OSPF Area

OSPF Area Name: area1

OSPF Authentication Type: Simple

Authentication Key: frosty

Enable Stub Area: ☑

Stub Default Cost: 0 (0 - 65535)

Net Range:

On 192.168.233.9 255.255.255.0 Advertise

Add OSPF Area Dialog Box

OSPF Area Name

The OSPF Area Name is an integer or IP address. If more than one area is configured within an AS, then one of these areas has to be area 0, which is the backbone. The backbone has to be physically connected to all other areas. It is recommended that you start with area 0 and then expand into other areas.

❖ **Note:** *The only exception to starting with area 0 is with virtual links, which are explained in the following section.*

OSPF Authentication Type

This pull-down box specifies whether the router will perform authentication of Link State Advertisements received from other routers.

- If **Simple** is selected, you must specify an authentication key for any interface which is associated with this area.
- If **None** is selected, no authentication will be done on Link State Advertisements.

Authentication Key

This sets the OSPF packet authentication key. The authentication key must match for each router connected to the interface and belonging to the area. The authentication key must be between one and eight alphanumeric characters.

Enable Stub Area

This checkbox sets whether this area will function as a stub area. A stub area is an area which cannot receive external advertisements, which means RIP or static routes will not be redistributed into this area.

If routing from a stub area to external routes (i.e., non-OSPF routes) is needed, a default route must be set. A stub area may not be a transit area for a virtual link.

❖ **Note:** *The backbone area (area 0) cannot be designated as a stub area.*

Stub Default Cost

This value sets the cost of the default route which will be used by routers within the stub area to route to external destinations. The value must be a number between 0 and 65,535.

Net Range

The Net Range can be used to consolidate routing information at area boundaries, or to hide routing information from routers outside the area. Net Ranges only apply to inter-area networks. If all the routers are in one area, any defined Net Ranges will not be used by the router. You may specify several different Net Ranges.

To add or modify a Net Range, click the appropriate button to open the Net Range Dialog Box.

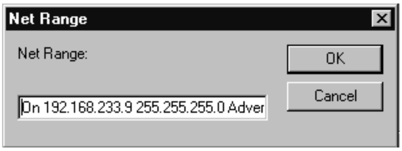

On 192.168.233.9 255.255.255.0 Adver

Net Range Dialog Box

The Net Range string has the following syntax:

{On|Off <IP Address> < IP Subnet Mask> Advertise|DoNotAdvertise}

- **On** specifies that a Net Range will be used, **Off** indicates that a Net Range is not being used.
- **IP Address** is the IP address of the Net Range
- **IP Subnet Mask** is the subnet mask of the Net Range.
- **Advertise** specifies that the Net Range will be advertised to other areas. **DoNotAdvertise** specifies that the network in the Net Range will not be advertised to other areas. This parameter is optional.

❖ **Note:** *DoNotAdvertise only applies to OSPF routes and not to routes learned from external protocols using IP route redistribution. External routes must be excluded by using route filtering.*

OSPF Virtual Link

This dialog box displays a list of all IP addresses being used for virtual neighbors. Configuring a virtual link is the only way to allow an area which is not contiguous to the backbone area (area 0) to operate.

❖ **Note:** *The virtual link must be configured in both routers which are providing the tunnel to the backbone. These two routers do not need to be physically connected, but they must share a common area called the **transit area**. (The transit area is designated in the Add OSPF Virtual Link Dialog Box.)*

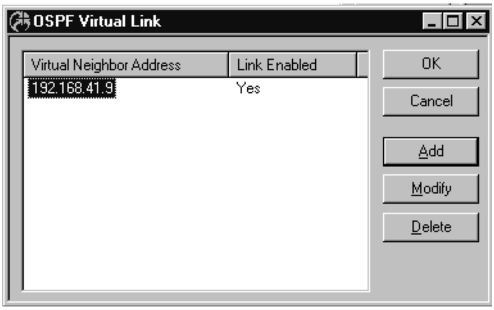

Virtual Neighbor Address	Link Enabled
192.168.41.9	Yes

OSPF Virtual Link Dialog Box

To access the OSPF Virtual Link Dialog Box, select OSPF/OSPF Virtual Link from the device view.

To add or modify the entries, select the appropriate button to access the Add OSPF Virtual Link Dialog Box.

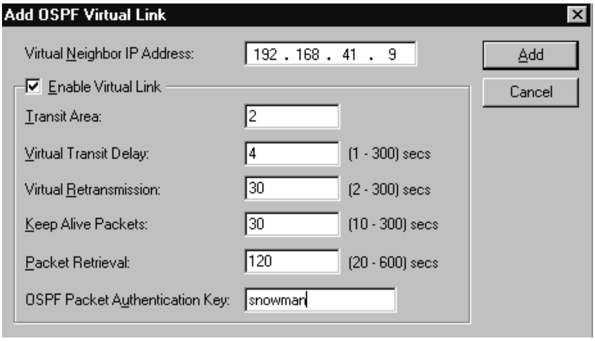

Virtual Neighbor IP Address: 192 . 168 . 41 . 9

☑ Enable Virtual Link

Transit Area:	2	
Virtual Transit Delay:	4	(1 - 300) secs
Virtual Retransmission:	30	(2 - 300) secs
Keep Alive Packets:	30	(10 - 300) secs
Packet Retrieval:	120	(20 - 600) secs
OSPF Packet Authentication Key:	snowman	

Add OSPF Virtual Link Dialog Box

This dialog box defines configuration parameters for an OSPF Virtual Link.

Virtual Neighbor IP Address

The virtual neighbor IP address is the largest IP address associated with the router used for the virtual link.

Enable Virtual Link

This checkbox will specify whether an OSPF virtual link will operate. When checked, it will activate the virtual link.

Transit Area

The transit area is the number assigned to the tunnel between the two routers of the virtual link. Each router must have at least one interface attached to the transit area. The transit area can be specified as a number between 0 and 0xFFFFFFFF, or as an IP address.

Virtual Transit Delay

The virtual transit delay sets the amount of time added to the age of Link State Update packets before transmission. It is the estimated number of seconds to transmit a packet over the virtual link. The numeric value can be between 2 and 65,535 seconds.

Virtual Retransmission

The virtual retransmission value sets the interval, in seconds, between retransmission of the Link State Update packets across the virtual link. The value can be between 2 and 65,535 seconds.

Keep Alive Packets

The keepalive value sets the interval, in seconds, that the router sends out “keepalive” packets across the virtual link to let the other end of the link know the router is up. The value must be greater than 10 seconds.

Packet Retrieval

The packet retrieval value sets the length of time, in seconds, that this router will wait without receiving a “keepalive” packet from the other end of the virtual link before assuming it’s down. **The packet retrieval value must be at least twice that of the keepalive packet value.** The default value is 4 times the keepalive packet value.

❖ **Note:** *The KeepalivePacket and Packet Retrieval values for each end of the virtual link must match or the virtual link will not function.*

OSPF Packet Authentication Key

This string sets the OSPF Authentication key for the virtual link. The string may be between one and eight alphanumeric characters.

❖ **Note:** *The authentication key must be the same for both ends of the virtual link.*

Chapter 16 - BGP

This chapter explains how to modify parameters that affect the way Border Gateway Protocol (BGP) operates. These parameters are global to the device and are not associated with a particular interface.

BGP is an exterior gateway protocol that runs on the Internet backbone and allows Autonomous Systems to exchange routing information with each other. BGP Autonomous Systems are separately administered sites which run other routing protocols such as RIP or OSPF internal to the site. Internet Service Providers are transit AS’s, which means that Internet traffic passes through their Autonomous System. Companies such as Compatible Systems are also separate AS’s, but are only a termination and origination point for Internet traffic.

BGP routers communicate via the TCP protocol. BGP routers that have established a BGP session are called BGP peers.

BGP General Dialog Box

To access this dialog box, select BGP/BGP General from the device view.

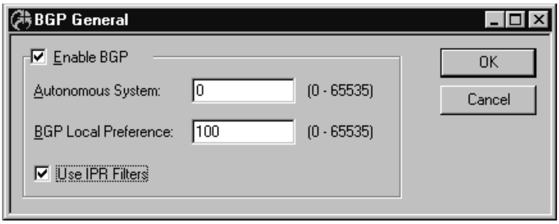

☑ Enable BGP

Autonomous System: 0 (0 - 65535)

BGP Local Preference: 100 (0 - 65535)

☑ Use IPR Filters

BGP General Dialog Box

Enable BGP

This checkbox turns BGP on globally for this device.

❖ **Note:** *You must also configure BGP peers to enable BGP. If no peers have been configured, BGP will not operate on the router, even if the Enable BGP box has been checked. BGP peers are configured later in this chapter.*

Autonomous System

This number specifies the Autonomous System (AS) to which this router belongs. An AS must be assigned a unique 16-bit number by the American Registry for Internet Numbers (ARIN). If an installation has only one ISP, the AS will be provided.

For multi-homed installation where more than one ISP is used, an “official” AS number is required.

BGP Local Preference

The local preference number sets the local preference of this router. The local preference is exchanged among routers in the same AS and is an indication about which path is preferred to exit the AS. A path with a higher local preference is more preferred.

Use IPR Filters

This checkbox sets whether the router will use IP route filters instead of BGP route maps. BGP uses BGP route maps to filter routes and set attributes. If no BGP route maps have been configured, the router will automatically use any configured IP route filters. IP Route Filters are set in **Chapter 11 - TCP/IP Filtering**.

BGP Aggregates Dialog Box

This dialog box defines a list of networks which are to be aggregated before being advertised to external peers. The router’s IP routing table must contain the networks which are a subset of the aggregate in order for the aggregate to be advertised.

❖ **Note:** *Only the aggregate, and not the individual routes, will be advertised to external peers. Internal peers will receive the individual routes if they originated outside the Autonomous System. Internal peers do not exchange internal routes via BGP.*

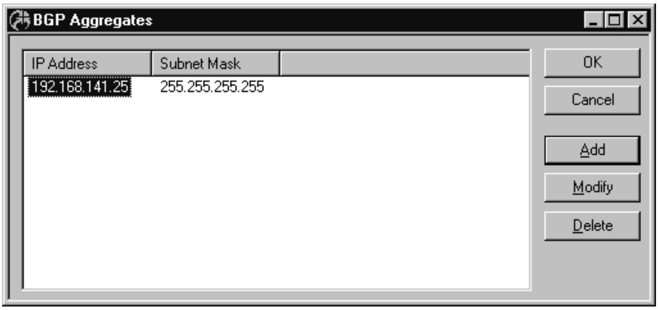

IP Address	Subnet Mask
192.168.141.25	255.255.255.255

BGP Aggregates Dialog Box

To access this dialog box, select BGP/BGP Aggregates from the device view.

To add or modify a BGP aggregate network on the list, click on the appropriate button to open the Add BGP Aggregate Dialog Box.

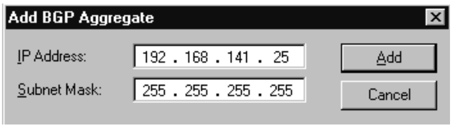

IP Address: 192 . 168 . 141 . 25

Subnet Mask: 255 . 255 . 255 . 255

Add BGP Aggregate Dialog Box

IP Address

The IP address specifies the IP address of the network to be aggregated and is entered in the standard dotted-decimal form.

Subnet Mask

The mask field is the subnet mask field of the aggregate network. If a mask is not provided, an all 255's mask will be assumed.

BGP Peer Configs Dialog Box

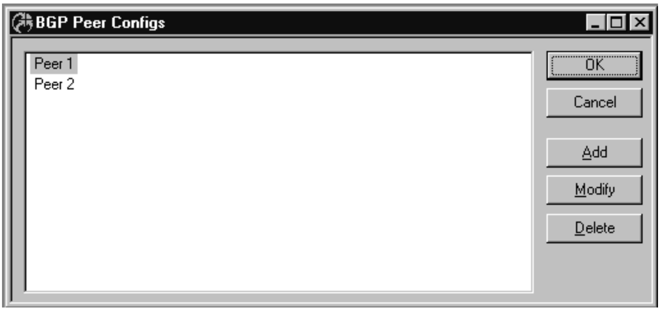

Peer 1

Peer 2

BGP Peer Configs Dialog Box

This dialog box defines a list of BGP peer configurations for a single BGP peer or for a group of BGP peers of this router. Any two routers that have opened a TCP connection to each other for the purpose of exchanging BGP routing information are known as peers.

To access this Dialog Box, select BGP/BGP Peer Configs from the device view.

A peer configuration should only be used for more than one peer if all the same parameters are desired. To add or modify these entries, click on the appropriate button to open the Add BGP Peer Config Dialog Box.

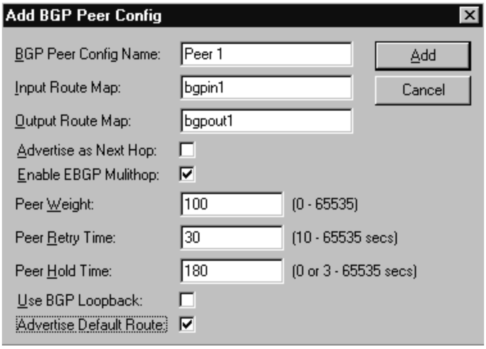

BGP Peer Config Name: Peer 1

Input Route Map: bgpin1

Output Route Map: bgpout1

Advertise as Next Hop: □

Enable EBGP Mulithop: ☑

Peer Weight: 100 (0 - 65535)

Peer Retry Time: 30 (10 - 65535 secs)

Peer Hold Time: 180 (0 or 3 - 65535 secs)

Use BGP Loopback: □

Advertise Default Route: ☑

Add BGP Peer Config Dialog Box

This dialog box defines configuration parameters for a BGP Peer.

BGP Peer Config Name

The Peer Config Name specifies the name of the BGP peer configuration that you wish to add or modify. Names can be up to 16 characters long.

Input Route Map

This specifies a named BGP route map or IP route filter to be used for the input route for this peer configuration.

Output Route Map

This specifies a named BGP route map or IP route filter to be used for the output route for this peer configuration.

❖ **Note:** *No input or output routes will be accepted by the router unless a BGP route map or IP route filter has been defined. BGP route maps are configured later in this chapter. Refer to* ***Chapter 11 - TCP/IP Route Filtering*** *for more information on IP route filters.*

Advertise as Next Hop

This checkbox sets whether the router will advertise itself as the next hop to the routes it advertises to this peer.

Enable EBGP Multihop

BGP usually requires external peers to be directly connected. This checkbox allows routers which are not directly connected to be peers. If checked, the router must also have a route to the external peer that is not directly connected in order to establish a connection.

Peer Weight

The peer weight value assigns an internal rating to the peer. Peers with a higher weight are preferred when multiple routes exist to the same destination. The number must be within the range of 0 to 65,535.

Peer Retry Time

The amount of time, in seconds, between retries to establish a connection to configured peers which have gone down for some reason. The value must be at least 10 seconds.

Peer Hold Time

The interval, in seconds, the router will wait for an update or keepalive packet from the peer before declaring the peer down. The hold time is actually negotiated between peers, which will use the smaller of the two hold times proposed. The value must be either 0, or at least 3 seconds. If zero is selected, keepalive packets will not be sent.

Use BGP Loopback

This checkbox allows the router's loopback address to be used as the IP source in TCP packets to that peer rather than a specific IP address of one of its interfaces.

❖ **Note:** *If this box is checked, you must specify an IP address in the IP Loopback Dialog Box.*

Advertise Default Route

This checkbox sets whether the default route to this peer will be advertised to other peers.

IP Loopback Dialog Box

This dialog box allows an IP Loopback address to be specified for the router. This parameter only needs to be set if you are using the BGP protocol and the **Use BGP Loopback** option is specified.

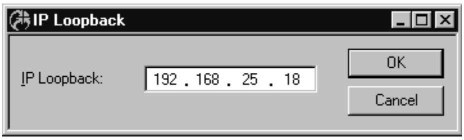

IP Loopback: 192 . 168 . 25 . 18

IP Loopback Dialog Box

To access this dialog box, select Global/IP Loopback from the device view.

IP Loopback

This specifies the IP address of the Loopback interface on the router. This can be used to provide a separate IP address for the router which is not tied to one of its IP interfaces.

BGP Peers Dialog Box

This dialog box defines a list of configured peers for this router. Routers that exchange BGP information are called BGP peers. A router may have both external peers in other Autonomous Systems, and internal peers within its own AS.

Routers establish BGP sessions using the TCP protocol. Upon startup of a new BGP session, BGP peers will exchange their full routing tables, and then only incremental updates are sent as the routing table changes.

❖ **Note:** *The router will not establish a BGP connection with any router not on this list. If there is no BGP Peers list, BGP will not be enabled even if the BGP Enabled box has been checked in the BGP General Dialog Box.*

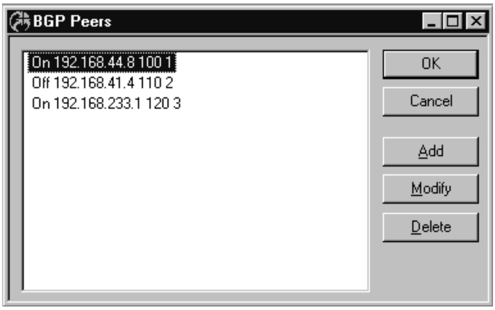

```
On 192.168.44.8 100 1
Off 192.168.41.4 110 2
On 192.168.233.1 120 3
```

BGP Peers Dialog Box

To access this dialog box, select BGP/BGP Peers from the device view. To add or modify this list, click on the appropriate button to open the Add BGP Peer Dialog Box.

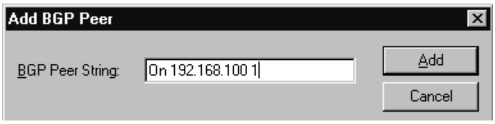

BGP Peer String: On 192.168.100 1

Add BGP Peer Dialog Box

The BGP Peer String specifies a BGP peer for this router. The BGP Peer String has the following syntax:

On|Off <IP Address> < AS Number> <Peer Config ID>

- **On/Off**

  This parameter determines whether the router will try to establish a BGP session with the peer at start-up.

  If this parameter is set to **Off**, the peer will not be contacted at start-up. The router can still establish a BGP session with this peer when the BGP Enable box is checked (in the BGP General Dialog Box). However, the next time the router is booted, the peer will come up in the **Off** state.

- **IP Address**

  This specifies the IP address of the interface which will be a BGP peer for this router. The router will contact the peer using this IP address. The router must have the network of the supplied IP address in its routing table in order for the session to be established.

  External peers should be directly connected to the router (usually over a WAN link). Internal peers do not need to be directly connected.

- **AS Number**

  This specifies the number of the Autonomous System (AS) of the BGP peer. The router determines if a peer is internal or external based on the AS number of the peer, since internal peers have the same AS number as the router itself.

- **Peer Config ID**

  This optional parameter specifies the number of the BGP Peer Configuration to which this peer will belong. A BGP Peer Configuration may be used for more than one peer only if all the same parameters are desired. BGP Peer Configuration is done using the BGP Peer Config Dialog Box.

BGP Route Maps Editor Dialog Box

Route maps help the administrator influence the route selection process. BGP uses weight, preference, and multi-exit discriminator (MED), among other things, to determine the optimal route. BGP uses the following criteria, in the order presented to select its best route for a destination.

- The most preferred path is the path with the largest weight.
- If the weights are the same, the protocol selects the path with the largest local preference.

- If the preferences are the same, the protocol selects the path that has the shortest AS path length.
- If all paths have the same AS path length, the protocol selects the path with the lowest MED.
- If the paths have the same MED, the protocol selects the path from the BGP peer with the lowest router ID.

❖ **Note:** *IP route filters may be used with BGP instead of BGP route maps. However, the matching conditions are more limited, and various parameters such as community, local preference, and weight cannot be set with IP route filters.*

The configuration of BGP route maps includes several elements which are optional. There can be multiple route names, each uniquely identified by the selection of these elements.

BGP route maps are used only by the BGP protocol to filter routes and set certain attributes. Route maps are not associated with a particular interface. They are applied to the device in the BGP Peer Config Dialog Box.

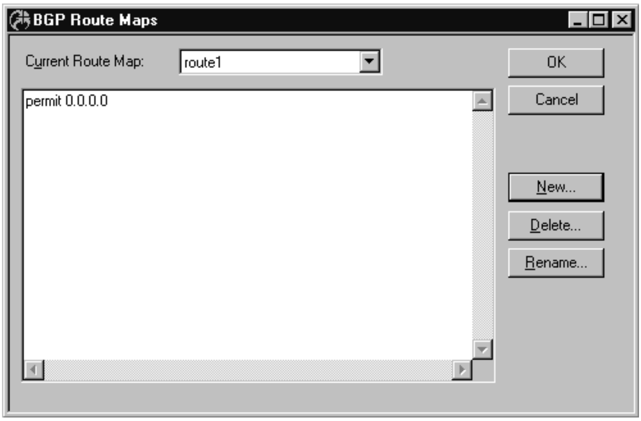

route1

permit 0.0.0.0

BGP Route Maps Dialog Box

To access the BGP Route Maps Dialog Box, select BGP/BGP Route Maps from the device view.

BGP Route Map Editor Dialog Box Buttons & Controls

- The **Current Route Map** pull-down menu lets you select a route map for editing.

- The **New** button brings up a dialog box which asks you to name the new route map, then creates a blank editor window and selects the new name in the **Current Route Map** pull-down menu. Names can be up to 16 characters long.

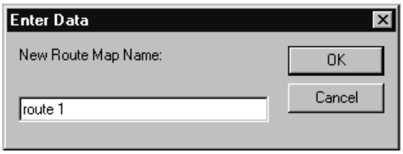

route 1

Enter Data Dialog Box

- The **Rename** button lets you change the name of the route map you are currently editing.
- The **Delete** button deletes the route map which is currently selected in the **Current Route Map** pull-down menu.

BGP Route Mapping Rules

At a minimum, every line in a route map must include an **Action**, a **Route**, and a **Direction**. Together, these components, along with optional input and output modifiers, specify a rule that the router will follow when a route meets the condition of the rule.

❖ **Note:** *No input routes will be accepted by the router unless a BGP route map or IP route filter has been defined. To allow all other network numbers not filtered, include the following rule:*

```
permit 0.0.0.0
```

Action

These parameters specify the action to be taken when a route meets the condition of the rule. Select **permit** or **deny**.

Route

The Route consists of an IP address with the optional **/bits** at the end of the IP address. IP addresses can be specified in the following three ways.

- Normal dotted-decimal notation. Example 192.168.12.4
- Factorized addresses in the form of #.#.#.# {#,#,...} Example 192.168.12.{1,4,8} This means IP addresses 192.168.12.1, 192.168.12.4, and 192.168.12.8

- Hexadecimal numbers. Example 0x82cc0801 (This matches the host address 130.204.8.1)

The optional **/bits** at the end of the IP address is a bit field denoting the number of bits that are significant when doing the comparison against the addresses from the IP packet. It denotes the top or most significant bits to use.

Direction

These parameters allow users to specify the direction for which the rule is applied. Select **in** or **out**.

Options/Output Modifiers

- **ipaddr** (IP address/<bits>) **toas** (AS number)

  This modifier limits output rules to routes going to the designated IP address or Autonomous System (AS) number. IP address may be specified in any of the ways described above. The AS number is specified as an integer.

- **origin** (protocol(s))

  This modifier limits output rules to routes originating from the designated protocol. Multiple protocols may be listed. Possible values are **icmp**, **rip**, **ripv2**, **static**, **OSPF**, **BGP**, and **direct**.

- **setnhop** (IP address)

  This modifier allows the next hop to be set on the outgoing route. The hop is specified as an IP address in standard dotted-decimal notation.

- **setmed** (MED number)

  This modifier allows the multi-exit discriminator (MED) to be set on the outgoing route. This is a metric which is used only when there are multiple paths to an AS. The MED is used to set a preference to a particular path to the AS, and is specified as an integer.

- **setasp** (AS number)

  This modifier allows the specified AS list to be prepended to the outgoing AS path attribute. Up to 6 AS numbers may be entered. The AS number is specified as an integer.

- **setcomm** (community number)

  This modifier allows a community list to be set on the outgoing route. A community is a group of destinations to which routing decisions can be applied. The community number can be specified with up to 6 community numbers, specified as integers, or can be listed as one of the three special communities, **noexport**, **noadv**, or **noexpsub**.

Special Communities

1. **noexport** - specifies that this route will not be advertised outside a BGP confederation boundary. A BGP confederation is a collection of several AS's that are advertised as a single AS to all BGP peers which are not members of the confederation.
2. **noadv** - specifies that this route will not be advertised to any BGP peers (including external peers).
3. **noexpsub** - specifies that this route will not be advertised to external peers. This means that this route can be advertised to internal peers only and will not be advertised outside its AS.

- **addcomm** (community number)

  This modifier allows a community list to be prepended on the outgoing route. The parameters can be up to 6 community numbers and are specified as integers.

Options/Input Modifiers

- **ipaddr**(IP address/optional bits) **hasas**(AS number) **srcas**(AS number) **nhop**(IP address) **comm**(community number)

  This modifier (with the exception of **hasas**), limits input rules to routes originating from the designated IP address, AS number, next hop, or community. A BGP route contains information concerning each AS it has traversed. The **hasas** parameter specifies that the rule will be applied if the AS path contains the specified AS number anywhere in the AS path.

  The IP address may be specified in any of the ways described previously. The AS number and community number are both specified as an integer.

- **setpref** (preference)

  This modifier allows the preference to be set on incoming routes from the given IP address, AS number, community, or next hop. The preference is specified as an integer.

- **setwt** (weight)

  This modifier allows the weight to be set on incoming routes from the given IP address, AS number, community, or next hop. The weight is specified as an integer.

BGP Networks

This dialog box defines a list of routes which will be advertised as originating inside the Autonomous System this router belongs to. These may be directly connected routes, static routes, RIP routes, or OSPF routes.

The route must be contained in the router's IP routing table or it will not be advertised. To advertise local networks which are not in the router's own IP routing table, they must be added as static routes.

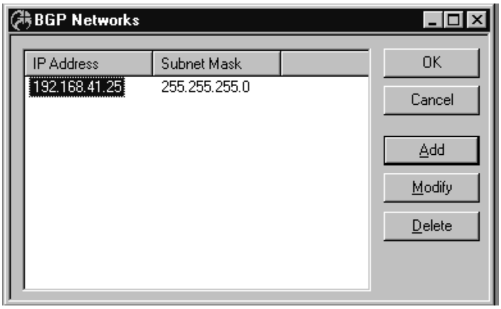

IP Address	Subnet Mask
192.168.41.25	255.255.255.0

BGP Networks Dialog Box

To access this dialog box, select BGP/BGP Networks from the device view.

❖ **Note:** *The only way to get directly connected routes advertised into BGP is to include them in this list. Static, RIP, and OSPF routes can also be imported into BGP by using IP Route Redistribution. Refer to* ***Chapter 2 - IP Routing & Bridging*** *for more information on IP Route Redistribution.*

To add or modify this list, select the appropriate button to open the Add BGP Network Dialog Box.

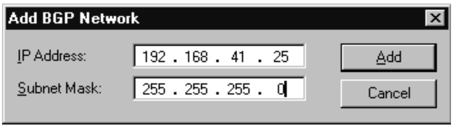

IP Address: 192 . 168 . 41 . 25
Subnet Mask: 255 . 255 . 255 . 0

BGP Network Dialog Box

IP Address

The IP address specifies a route to be advertised as originating inside the local Autonomous System to which the router belongs. The IP address is entered in standard dotted-decimal notation.

Subnet Mask

The optional mask parameter tells the router how many bits of the IP routing table entry to match against the IP address listed in the BGP Network. If a mask is not provided, an all 255’s mask will be assumed.

❖ **Note:** *This is not necessarily the actual mask of the network you wish to advertise because subnet masks more specific than Class C are automatically truncated. This truncation is not the same as aggregation, and only applies to internal networks, and only to masks more specific than Class C.*

Appendices

IP 101

❖ **Note:** *This is a very brief introduction to IP networking. For more in-depth information, there are a number of excellent references. In particular, Douglas Comer's Internetworking with TCP/IP (Prentice Hall) is one of the standard references and provides a wealth of information on the subject.*

IP Addresses

Each device on an IP network requires 3 different pieces of information in order to correctly communicate with other devices on the network: an IP address, a subnet mask, and a broadcast address. You will usually see each of these numbers written as four "octets" (e.g. 198.41.12.151, 255.255.255.0, and 198.41.12.255).

Every IP address is really made up of two pieces: a "network" portion, which tells routers what group of devices a packet should go to (e.g., any, a campus, etc.) and a "host" portion which tells routers what specific device among that group the packet should go to.

By examining the destination address in an IP packet that must be forwarded, and by using information that has either been statically configured or dynamically gathered from other routers, any router can determine the optimal path for forwarding packets from one group to another.

Each group of devices on an IP internet needs to have a unique network portion, and each device within that group also needs a unique host portion. In the case of the Internet, this uniqueness is made possible by indirectly getting all network portion assignments through a central clearinghouse called the Network Information Center or "NIC." The NIC assigns blocks of addresses to Internet Service Providers (ISPs), who then assign these addresses to their customers.

If your network is, or will be, connected to the Internet, you will need to get a unique network address from your ISP or network administrator.

How much of any given address is the network part and how much is the host part is determined by the "class" of the network. In each case, the part of the address not used for the network portion is left as the host portion.

CLASS	NETWORK PORTION	HOSTS ALLOWED
A	from 1. to 127.	about 16 million
B	from 128.0 to 191.255	65,536
C	from 192.0.0 to 223.255.255	255

Chart 1: IP Address Classes

You can always tell what class an address is by looking at the first octet and comparing it to the chart above. For instance, the address at the top of this appendix has 198 as the first octet, so it is Class C.

Subnet Masks

A subnet mask tells a router how much of an address it should treat as the network portion. The masks for traditional Class A, B and C networks are shown below.

CLASS	SUBNET MASK
A	255.0.0.0
B	255.255.0.0
C	255.255.255.0

Chart 2: Standard IP Subnets

Comparing the masks above to the first chart, you can see that the 255's in a mask identify the network portion of the address.

Just as the masks above specify what portion of the global IP address range a network is using, a subnet mask can also be used to subdivide a Class A, B or C network range into multiple groups of hosts, or "subnets."

This is done by telling the router that more than the traditional number of bits in the mask are to be treated as the network portion of the address. The chart below shows all of the possible Class C subnet masks, and how many hosts are then allowed on each subnet.

SUBNET MASK	HOST RANGES
255.255.255.0	1-254 (traditional C)
255.255.255.128	1-126, 129-254
255.255.255.192	1-62,65-126,129-190,193-254
255.255.255.224	1-30,33-62,65-94,97-126, 129-158,161-190,193-222, 225-254
255.255.255.240	1-14,17-30,33-46,49-62, 65-78,81-94,97-110, 113-126,129-142,145-158, 161-174,177-190,193-206, 209-222,225-238,241-254
255.255.255.248	1-6,etc.

Chart 3: Subnetted Class C Host Ranges

❖ **Note:** *The lowest calculated address in each range (0 in the traditional C range) is not shown, cannot be used, and is skipped in the chart. The highest address in each range (255 in the traditional C range) is also not shown, and is the broadcast address for the subnet.*

With each mask above, the 1's in the binary value represent the network portion, and the 0's represent the host portion (128 is 10000000, 192 is 11000000, etc.). As you use more bits to represent the network portion, fewer bits are left to use as host addresses.

The same idea can be extended to Class A and Class B networks.

Broadcast Addresses

The broadcast address is the address to which devices send packets meant for all other devices. All devices "listen" for broadcasts in addition to their own address. Address Resolution Protocol (ARP) packets and routing information are examples of packets sent to the broadcast address. Most often, the broadcast address is the last address in the network (or subnet) with the host portion being all 1's binary (some networks use 0.0.0.0 or 255.255.255.255, however.). Below are some examples of broadcast addresses.

CLASS	NETWORK	SUBNET MASK	BROADCAST
A	45.0.0.0	255.0.0.0	45.255.255.255
B	128.138.0.0	255.255.0.0	128.138.255.255
C	198.41.9.0	255.255.255.0	198.41.9.255
A*	45.21.16.0	255.255.252.0	45.21.19.255
C*	198.41.9.64	255.255.255.224	198.41.9.95

Chart 4: Broadcast Address Examples

The first three entries are traditional Class A, B and C network addresses and use traditional masks. The last two are less traditional, "real world" examples. Note in line 4 the change in the third octet between network address and broadcast address. Line 5 shows what happens when a Class C network has been subnetted.

Assigning an IP address

Use the network portion you were given by your administrator or ISP. Assign the router interface a unique (i.e. unused) host portion. For example, if your ISP tells you your network portion is 198.41.9, you could assign an interface to 198.41.9.1. If you have a router with more than one interface, the network (+ subnet) portions of each port's IP address must be different.

Assigning a Subnet Mask

If you are using traditional Class A, B or C networks, CompatiView will automatically calculate the value for you. If you wish to compute it yourself, use the values in Chart 2.

If you are subnetting, use Chart 3 as a guide for Class C, or follow the same scheme for Class A or B. Note that the IP address for a subnetted interface (including the router interface on that subnet) must be in the correct subnet range, as shown in Chart 3.

Assigning a Broadcast Address

CompatiView will automatically compute the broadcast address for you. If you wish to compute it yourself, use the examples in Chart 4 above as a guide. You can then use CompatiView to check your results.

Static Routes & Routing Protocols

In addition to the three required values, you must also decide whether to use an IP routing protocol. Routing protocols are how routers tell each other about networks they are responsible for. Virtually all routers support the IP Routing Information Protocol (RIP).

There are also a variety of other routing protocols which have been developed, some proprietary and some open. A router which is using one of these other protocols can always accept routes using RIP and then supply information about them using the other protocol.

If you choose not to use RIP, or other routers on your network are not broadcasting routing information, you may need to set a default router or define some static routes.

The default router is the place where your router will send any packets addressed to IP networks that it does not know about. With RIP turned off, it will only know about statically configured routes. For very simple IP connections, such as a small network being connected out to the Internet through an ISP, a default route is probably the only routing information needed by your router.

A default router provides a generic location for packets to be sent to, while static routes are more detailed definitions where you specify the route for certain networks, and a “metric” which defines how attractive the route should be considered.

When specifying default routes, you must provide a mask value (as discussed earlier) which tells the router how much of the address you are entering the route for should be considered as the network portion.

IPX 101

❖ **Note:** *This is a very brief introduction to IPX networking. For more in-depth information, there are a number of excellent references. In particular, Rick Sant’Angelo’s NetWare Unleashed (SAMS Publishing) provides a good overview of IPX routing along with tips on getting IPX drivers correctly loaded on client machines.*

IPX Routing Basics

All routable protocols work by dividing the physical devices on a network into logical groups. A logical group will typically consist of all of the machines on a physical network segment (such as an Ethernet segment).

Each group of devices is assigned a unique “network number” which represents that particular group to all of the routers on the network. Packets which are sent between members of the same group are simply sent directly from one member to another.

Packets which must go between devices belonging to two different groups travel through routers, which forward them along an optimal path.

By examining the destination network number in a packet that must be forwarded, and by using information that routers automatically pass between themselves in IPX Routing Information Protocol (RIP) packets, any router can determine the optimal path for forwarding packets from one group to another.

This scheme relies on the fact that each segment is assigned a unique network number. If not, the routers have no way of knowing which of the physical segments with that number should actually receive a packet.

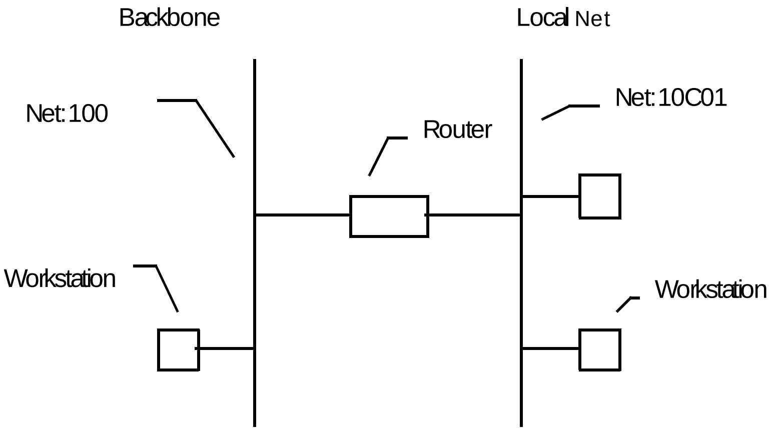

IPX Routing Example

Among routable network protocols, IPX is relatively simple. Each physical network segment is assigned a network number by the routers on the segment. The network number can be in the range of 1 to FFFFFFFE (that’s 8 hexadecimal digits). In the diagram above, 100 and 10C01 are the network numbers for the two segments shown.

Establishing the network number for an IPX network segment is referred to as “seeding” the network. You should generally only have one seed router per network cable segment. It may sometimes be desirable for redundancy to have several seed routers on a segment. This is acceptable as long as all seed routers on the segment are seeding the same network number.

Service Advertising Protocol

Routers participate in allowing end nodes to access IPX services (such as file servers, print servers, communications servers, etc.) by keeping a list of all of the services on an IPX internetwork. This list is maintained by examining the Service Advertising Protocol (SAP) packets which are sent by servers and other routers on the local segment, and by rebroadcasting this information out of their other interfaces.

A “split-horizon” technique is used so that routers do not duplicate information which is already known on the segment being broadcast to.

Client Machine Addressing

Unlike TCP/IP, IPX workstations do not have fixed network/node addresses that need to be configured. Instead, a workstation gets its network number from the router(s) on the segment it is connected to, and uses its Ethernet address for its node number.

This means that an IPX workstation may have as much as 18 hexadecimal digits of network/node address. Fortunately for workstation users, the NetWare client software does the work of discovering the network number and setting the address. Users only need to install Novell drivers to be able to use the IPX protocols over their network.

Routers which support IPX can use any of four “frame types” to send IPX packets. Each frame type organizes the IPX information in a network packet (i.e. frame) in a slightly different fashion. Although the basic information may be the same, clients or servers using different frame types cannot communicate with each other without an intermediate translation occurring between frame types. This translation is called “transitional routing,” and is one of the functions that can be performed by routers.

The four IPX frame types are:

- Ethernet_Type_II
- Ethernet_802.3 (Raw)
- Ethernet_802.2
- Ethernet_SNAP

Older versions of NetWare defaulted to the 802.3 Raw frame type, whereas NetWare 4.0 uses the 802.2 frame type.

For this reason, the default configuration for Compatible Systems routers which support IPX has both 802.3 Raw and 802.2 set to autoseed (they will come up regardless of whether there is a server on line or not) and the other two frame types set to non-seed (they won’t come up unless they “hear” another router using this frame type.

This autoseeding default router configuration simplifies administration of the router since IPX can be routed right out of the box without any configuration. To determine a network number to use for autoseeding, a router listens to the network for several RIP periods, and then examines its routing table (which is filled in with information from RIP packets), and picks an unused number.

AppleTalk 101

❖ **Note:** *This is a very brief introduction to AppleTalk networking. For more in-depth information, the definitive reference is Gursharan Sidhu et al’s Inside AppleTalk (Addison-Wesley Publishing). This book provides an in-depth look at the AppleTalk protocol suite and AppleTalk routing.*

AppleTalk Routing Basics

All routable protocols work by dividing the physical devices on a network into logical groups. A logical group will typically consist of all of the machines on a physical network segment (such as an Ethernet segment).

Each group of devices is assigned a unique “network number” (or a range of network numbers) which represents that particular group to all of the routers on the network. Packets which are sent between members of the same group are simply sent directly from one member to another.

Packets which must go between devices belonging to two different groups travel through routers, which forward them along an optimal path.

By examining the destination network number in a packet that must be forwarded, and by using information that routers automatically pass between themselves in AppleTalk Routing Table Maintenance Protocol (RTMP) packets, any router can determine the optimal path for forwarding packets from one group to another.

This scheme relies on the fact that each segment is assigned a unique network number/range. If not, the routers have no way of knowing which of the physical segments with that number should actually receive a packet.

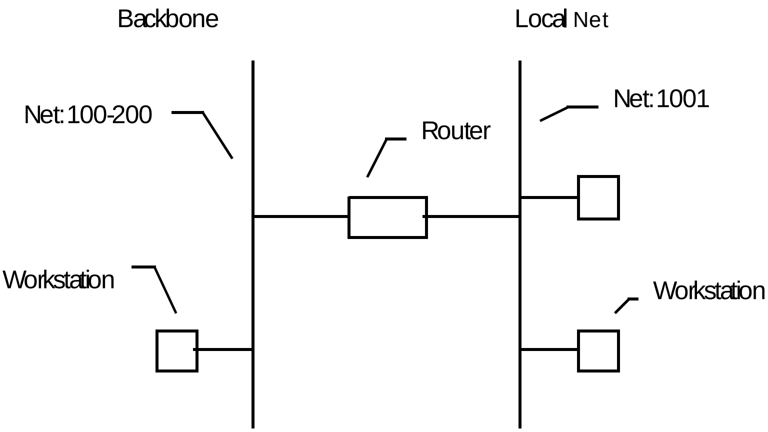

AppleTalk Routing Example

Each AppleTalk physical network segment is assigned a network number/range by the routers on the segment. The network number (or range of numbers) can be between 1 and 65,279. In the diagram above, 100-200 is the network range for the backbone, and 1001 is the network number for the local net segment.

Non-extended and Extended AppleTalk Networks

The original AppleTalk specification, which is now referred to as AppleTalk Phase 1, used only a network number, not a network range. A network number was a sixteen bit value, which allowed numbers between 1 and 65,534 to be used. The address of an individual device on the segment consisted of the network number, along with an 8 bit node address value. This scheme meant there could be a maximum of 254 devices per network segment. While this was more than adequate for LocalTalk networks, it was a major constraint on Ethernet networks.

AppleTalk Phase 2 introduced the concept of extended networks. While the node address remained an 8 bit number, network segments could now be identified by a range of network numbers between 1 and 65,279. Each number in the range allows 253 node addresses. These Phase 2 extended ranges should be used for all new AppleTalk installations.

"Seeding" a Network Segment

Establishing the network number/range for an AppleTalk network segment is referred to as "seeding" the network. You should generally only have one seed router per network cable segment. It may sometimes be desirable for redundancy to have several seed routers on a segment. This is acceptable as long as all seed routers on the segment are seeding the same network number/range.

Unlike TCP/IP, AppleTalk workstations do not have fixed network/node addresses that need to be configured. Instead, a workstation gets a network number from the router(s) on the segment it is connected to, and picks an unused address for its node number through a process called probing.

Probing

When a device comes up on a non-extended AppleTalk network, it will set its network number to the number seeded on the network, and then try to claim a node address. It does this by broadcasting a packet to all other nodes on its segment asking whether the node address is already in use. If another node on the segment responds, the original node will randomly select another node address value and try again.

When a device comes up on an extended AppleTalk network, it will set its network number randomly to one of the numbers in the range seeded on the network, and then try to claim a node address. It does this by sending out a packet to all other nodes on its segment asking whether the node address is already in use. If another node using the same number in the network range responds, the original node will randomly select another network number and node address value and try again.

Zones

While network numbers/ranges logically group devices together according to which network segment they are connected to, AppleTalk zones provide a way of creating groupings which can correspond to any concept a network administrator cares to use. This could be the department the devices are used in, the physical location of the devices, or some other method of categorization.

Zones are configured into a router by an administrator, and are logically tied to a segment and its network number/range by the router. However, the same zone names can be used on different segments. This gives an administrator the opportunity to make zone names represent groups of devices which are on more than one segment. A non-extended network can only have one zone (which will also be the “default zone” for the segment). An extended network can have from 1 to 255 zones, one of which will be the default zone.

Once a device has successfully claimed an address, it contacts a router on its segment and asks for a list of zones for the segment. Unless it has been configured to pick one of the other zones, it will use the “default zone” which is returned by the router.

When a device on the network attempts to discover services (such as servers or printers) using a Chooser program, an NBP (Name Binding Protocol) lookup packet is sent to a router on the same segment, which then performs a lookup in its tables to determine the network number(s)/range(s) for a particular zone. These tables are maintained using the ZIP (Zone Information Protocol).

The lookup is then forwarded to the appropriate segment(s). Devices whose services match the information in the lookup will respond to it, and the response will be forwarded back to the original machine.

Router Autoconfiguration

An autoseeding default router configuration simplifies administration of routers since AppleTalk can be routed right out of the box without any additional configuration.

To determine a network number to use for autoseeding, a router listens to the network for several RTMP periods, and then examines its routing table (which is filled in with information from RTMP packets), and picks an unused number for each interface. Only Phase 2 extended networking is turned on in the default configuration, with network ranges of 1.

A default zone name is created for each interface that incorporates the router’s Ethernet address, which is guaranteed to be unique.

Bridging 101

❖ **Note:** *This is a very brief introduction to the concept of bridging networks. For more in-depth information, there are a variety of references, including the IEEE 802.3d spanning tree specification. A good general purpose reference is Radia Perlman’s Interconnections (Addison-Wesley Publishing).*

Bridging Basics

Bridges are used to limit the amount of traffic appearing on network segments other than the destination segment. They do not provide for the logical grouping of network devices, which makes them considerably less flexible than routers from the standpoint of network management.

In contrast to routers, bridges operate on the “physical” network layer. While protocols such as IP or IPX are concerned with their own addressing schemes and routing tables (see IP 101 or IPX 101), bridging is only concerned with physical (i.e. Ethernet) addresses, and which bridge interface they are located on.

This simplicity is both the strength of bridging, and also a weakness. Because bridges maintain very little information about network topology, they are easier to configure than routers. But for this same reason, they do not limit traffic on network segments as well as routers do, and they are more prone to propagating network problems from one segment to another.

❖ **Note:** *“Ethernet switches” are actually just a new name for multiport bridges.*

Transparent Bridging

The simplest kind of bridge is called a transparent bridge. It operates according to the following rules:

1. Examine all packets on all active network interfaces for their source address.

2. Maintain a table that tracks which interface a source address has appeared on.

3. Look up the destination addresses in this table for all packets, and if a packet’s matching interface is different than the interface it was received on, forward the packet to the matching interface.

4. If no match is found, or if the destination address is the broadcast address, forward the packet out all active interfaces.

This scheme is acceptable on very simple network topologies. It will not work correctly if there are multiple paths to the same destination. In this case, packets will be forwarded in a “bridging loop” which will quickly use up the available network bandwidth on the segments in question.

Spanning Tree Bridging

To avoid bridging loops, an algorithm was developed which lets bridges shut off interfaces which provide duplicate paths to the same destination. This “spanning tree” algorithm was ratified as an IEEE standard (802.3d), and is supported by most bridge/switch vendors.

The algorithm relies on the use of Bridge Protocol Data Units (BPDU packets), which provide information to all bridges about the “distance” in hops to each bridge interface from a “root bridge.” The root bridge is selected using settings entered into each bridge (with the Ethernet address acting as a tie-breaker).

Using BPDU information, a bridge can determine whether one of its interfaces provides an optimal path to the root bridge. If it does not, the interface is shut down. If the path distance is optimal but is the same as another bridge’s path, a simple protocol allows one of the interfaces to be shut down.

In all other respects, spanning tree bridges operate in the same fashion as simple learning bridges.

A Simple Bridging Example

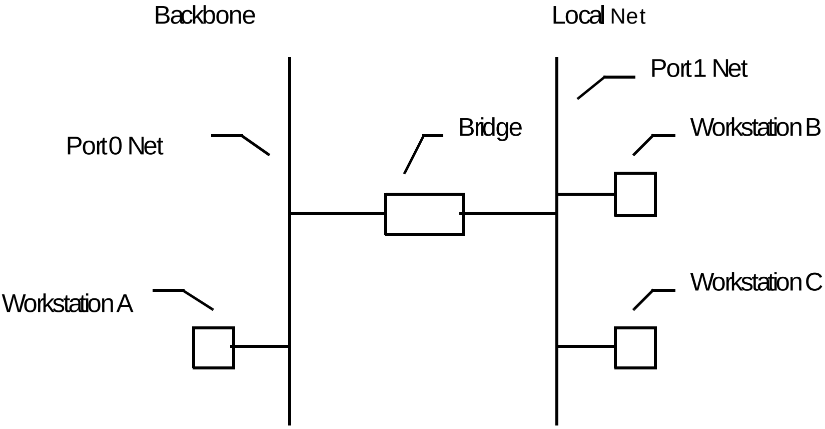

Bridging Example

In the diagram above, the bridge develops a table by listening to both the Port 0 net and the Port 1 net. Through the listening process, it associates Workstation A with Port 0 and Workstations B and C with Port 1. A simplified bridging table is shown below:

```
A -> Port 0
B -> Port 1
C -> Port 1
```

When a packet with A as a destination address arrives at Port 0, the packet is dropped (A is on the same interface). When a packet with a destination address of C arrives at Port 0, the packet is forwarded to Port 1. When a packet with a destination address which isn't in the table (or a broadcast address) arrives at Port 1, it is forwarded to Port 0.

Multiport Bridges/Switches and "Bridge Groups"

When a router has multiple interfaces, and also supports bridging/switching, some new concepts are required to understand the organization of the available functions.

The following diagram shows a four interface router which also supports bridging. Two of the router's interfaces (Port 0 and Port 1) are set to bridge IP, and two interfaces (Port 1 and Port 2) are also set to bridge IPX.

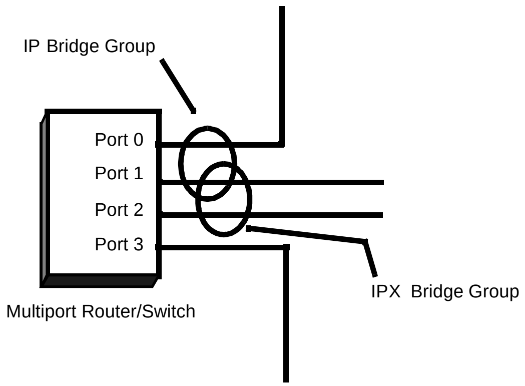

Bridge Groups on a Multiport Router

The diagram illustrates two Bridge Groups. The IP Bridge Group consists of Port 0 and Port 1. The network segments connected to these two interfaces appear as a single logical segment for IP routing purposes. That is, they will share a single IP network number, subnet mask, and broadcast address. IP communications between these two segments will be bridged, not routed.

The IPX Bridge Group consists of Port 1 and Port 2. The network segments connected to these two interfaces appear as a single logical segment for IPX routing purposes. That is, they will share a single IPX network number. IPX communications between these two segments will be bridged, not routed.

In this example, the segment connected to Port 3 has its own IP network number, subnet mask, and broadcast address. It also has its own IPX network number. Thus all IP and IPX communications between this segment and the two Bridge Groups is routed.

Finally, assuming that non-routable protocols have not been excluded, Ports 0, 1 and 2 all appear as a single physical segment to NetBEUI and DEC LAT.

Frame Relay 101

Frame Relay is a streamlined subset of the X.25 packet switching protocol which has been used by many corporations for wide area communications for a number of years. By removing a number of the X.25 protocol's seldom-used functions and their associated overhead, the Frame Relay protocol allows communications at up to T1 speeds (about 1.5 megabits per second).

The generic advantage provided by Frame Relay is its ability to combine multiple streams of "bursty" data (such as LAN protocol traffic) all of which have relatively low average usage rates, into a single channel with a relatively higher average usage rate. This "statistical multiplexing" effect allows your Frame Relay carrier to provide high bandwidth wide area connectivity to you at a price which is often significantly lower than standard leased line rates.

Virtual Circuits

Like X.25, Frame Relay is a connection oriented service requiring circuits to be configured by your carrier to establish a physical link between two or more locations. Multiple virtual circuits (which appear as virtual point-to-point links) can be run through the same physical connection.

There are two types of virtual circuits supported in Frame Relay: Permanent Virtual Circuits (PVC) and Switched Virtual Circuits (SVC).

PVCs are like dedicated point-to-point private lines. Since the physical connection is always there in the form of a leased line, call setup and tear down is done by a carrier via a network management system.

SVCs are analogous to X.25 connections, which require call setup and tear down.

❖ **Note:** *SVCs are generally not yet available from Frame Relay carriers. Virtually all Frame Relay communications is presently being done using PVCs.*

Addressing

A number called the Data Link Connection Identifier (DLCI) identifies each virtual circuit within a shared physical channel.

Frame relay packets are exchanged between nodes by mapping packets containing the source node's DLCI address to the destination DLCI address at the switch. Each switch contains a table identifying the various DLCIs with their associated user lines and interface trunks. However, the switch has more or less work depending on if the DLCI has global or local significance.

Local & Global DLCIs

Local DLCI addressing means that DLCI numbers are only significant at one end of a Frame Relay virtual circuit (VC). In other words, the same VC will be identified by different DLCIs at each end. To accomplish this, a mapping occurs across a VC. Frame Relay switches are required to translate the “source” DLCI in a packet to the “destination” DLCI when forwarding the packet.

Global DLCI addressing is a Local Management Interface (LMI) extension that allows a DLCI number to have universal significance. A global DLCI identifies the same VC at both ends. Global addressing simplifies address administration but allows for only 1024 DLCIs in the entire network. The switch is not required to translate the DLCI in a packet as it does with local DLCIs.

❖ **Note:** *The majority of Frame Relay connections use **Local** DLCI addressing, where a DLCI number is only significant at one end of the PVC.*

Local Management Interface (LMI)

The local management interface specifies communication between different Frame Relay devices (i.e. frame relay switches, routers, access devices, etc.). Over the years, three different local management interface specifications have been developed for Frame Relay: “consortium” LMI (an early cooperative effort by a group of frame relay vendors), CCITT Annex A, and ANSI Annex D. The CCITT and ANSI specifications are formal outgrowths of the consortium LMI specification.

The Annex D specification is the most widely used in the United States, although consortium LMI is still in use by some carriers. The Annex A specification is primarily a European specification.

Encapsulation and Fragmentation

RFC 1490 describes an encapsulation method for carrying packets across a Frame Relay network. All protocol packets are encapsulated within a Q.922 Annex A frame (a CCITT specification for data frames). Additionally, the frames must contain information necessary to identify the protocol being carried, allowing the receiver to properly process the incoming packet.

RFC 1490 also specifies a simple fragmentation procedure for carrying large frames over a frame relay network with a smaller maximum frame size.

Network/Protocol Addressing and Virtual Interfaces

Routing between LANs across a Frame Relay network is similar to routing across a point-to-point connection. A PVC on one router is directly connected to a PVC on another router. The difference is that multiple PVCs can be supported on the same physical interface of a router.

Network/protocol addresses are associated with each PVC using one of two methods: static mapping, or the Inverse Address Resolution Protocol (IARP).

IARP is outlined in RFC 1293. IARP allows dynamic mapping of protocol addresses to a DLCI. It can be used for IP, IPX and AppleTalk. It is more flexible and easier to configure than static configuration.

IARP is used when a router discovers a new PVC with its corresponding DLCI on a physical interface. The PVC is discovered by communicating with the Frame Relay switch using the LMI protocol. This may be done when the router is coming up or when a PVC has come back up after going down for some reason.

A Frame Relay Example

The following diagram shows three remote office 56K Frame Relay connections feeding one central office T1 connection. Note that the PVCs are shown as virtual point-to-point links which run through the physical connections and the Frame Relay cloud.

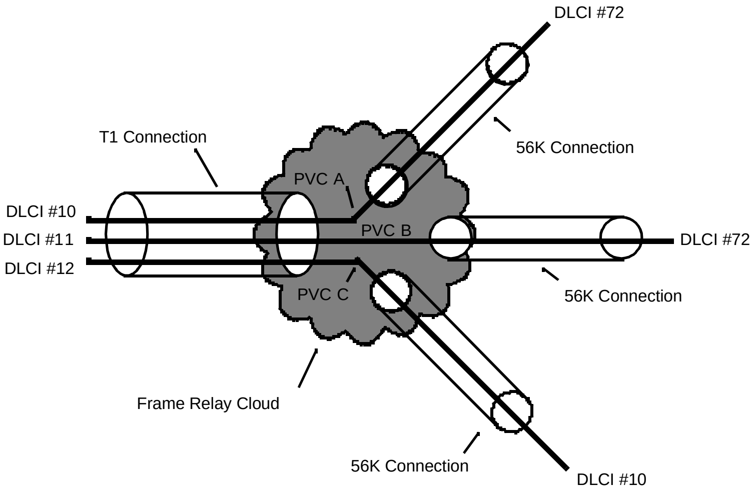

A Frame Relay Example

The DLCI numbers in the diagram are only locally significant. That is, DLCI numbers can only be guaranteed to not be duplicated locally, and a DLCI at one location has no significance at another location.