Chapter 8
Working with VLANs

This chapter describes the management of VLANs on an Accelar 1000 Series routing switch using Accelar Device Manager including creating/editing/deleting VLANs and managing VLAN bridging. It includes the following sections:

For a description of the use of Accelar VLAN Manager to manage VLANs across multiple devices, refer to Chapter 4, "Using Accelar VLAN Manager."

For information about configuring IP routing on a VLAN, refer to Chapter 9, "IP Interfaces and Router Management."

Accelar 1000 Series VLANs

A VLAN is a collection of ports on one or more switches that define a broadcast domain. Accelar 1000 Series routing switches support three types of VLANs:

- Port-based VLANs
- Source IP-subnet-based VLANs
- Protocol-based VLANs

For further discussion of the types of VLANs, refer to the "Important Information about this Software Release" section in Chapter 1, "Networking Concepts."

When creating VLANs using Accelar Device Manager, keep in mind the following rules:

- The ports in a VLAN or Multi-Link Trunk must be a subset of a single spanning tree group.
- VLANs must have unique VLAN IDs and names.
- An access (nontagged) port can belong to one and only one protocol-based VLAN for a given protocol.
- An access (nontagged) port can belong to multiple IP-subnet-based VLANs.
- An access (nontagged) port can belong to one and only one port-based VLAN.
- A frame's membership in a source IP-subnet-based VLAN takes precedence over a protocol-based VLAN, which takes precedence over a port-based VLAN.
- The Default VLAN (VLAN ID 1) cannot be renamed or deleted, or it cannot have its type changed from port-based VLAN.

Managing VLANs

The main window for managing VLANs in Accelar Device Manager is the Edit VLAN window accessed by selecting VLAN -> VLANs from the main menu. The window is divided into Basic and Advanced display areas.

The Basic VLAN window pictured in Figure 8-1 displays all defined VLANs, their configurations, and their current status. For a description of the Basic VLAN Window fields, refer to Table 8-1.

Note: The Basic VLAN window contains the Bridging button and IP button, which access screens for managing the bridging and IP routing aspects of the VLAN, respectively. The options under the Bridging button are described in the "Managing VLAN Bridging" section on page 8-12 in this chapter. The options under the IP button are described in the following chapters.

The Advanced VLAN Window pictured in Figure 8-2 contains advanced options including the Action field, which may be useful in troubleshooting. For a description of the Advanced VLAN Window fields, refer to Table 8-2.

The Snoop Window is described in Chapter 13, "IP Multicasting."

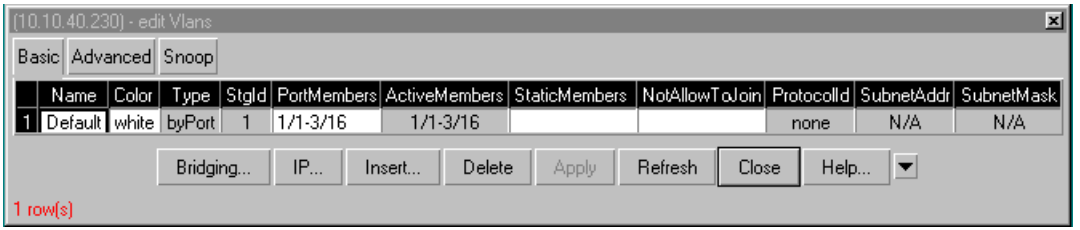

Figure 8-1. Basic VLAN Window

Table 8-1. Basic VLAN Window Fields

Field	Description
VlanId	The VLAN ID for the VLAN (unlabeled farthest left column).
Name	The name of the VLAN.
Color	The color is a proprietary color scheme used by Accelar VLAN Manager to associate a color with a VLAN. Color does not affect how frames are forwarded.
Type	Indicates the type of VLAN: ByPort or ByProtocolId.
StgId	The spanning tree group ID to which the VLAN belongs.
PortMembers	The slot/ports that are possible members of the VLAN.
ActiveMembers	The slot/ports that are active members of the VLAN. These include all static members and any potential member where the policy has been met.
StaticMembers	The slot/ports that are static (always) members of a protocol-based VLAN.
NotAllowToJoin	The slot/ports that are not allowed (never) to become members of a protocol-based VLAN.
ProtocolId	The protocol for protocol-based VLANs. This value is taken from the Assigned Numbers RFC. For port-based VLANs, none is the displayed value.
SubnetAddr	The source IP subnet address (IP subnet-based VLANs only).
SubnetMask	The source IP subnet mask (IP subnet-based VLANs only).

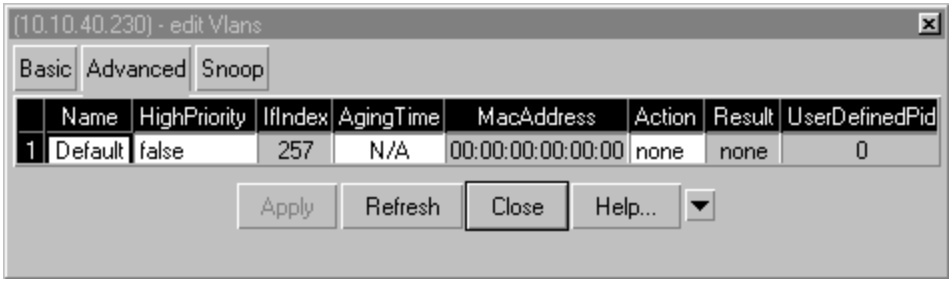

Figure 8-2. Advanced VLAN Window

Table 8-2. Advanced VLAN Window Fields

Field	Description
VlanId	The VLAN ID for the VLAN (unlabeled farthest left column).
Name	The name of the VLAN.
HighPriority	If true, frames in the VLAN will be forwarded through the switch fabric with high priority. If false, frames in the VLAN will be forwarded through the switch fabric with normal priority.
IfIndex	If routing is set to true for the VLAN, this value indicates the logical ifIndex that is assigned to the virtual router interface for the VLAN.
AgingTime	The timeout period in seconds for aging out the dynamic member ports of policy-based VLANs.
MacAddress	The MAC address assigned to the virtual router interface for this VLAN. *This field is relevant only when the VLAN is configured for routing*. This MAC address is used as the Source MAC in routed frames, ARP replies, or RIP and OSPF frames.
Action	One of the following VLAN-related actions: • flushMacFdb—flush MAC forwarding table for VLAN • flushArp—flush ARP table for VLAN • flushIp—flush IP route table for VLAN • flushDynMemb—flush Dynamic VLAN port members • all—flush all tables for VLAN
Result	Result code for Action.
UserDefinedPid	User-defined protocol ID if the user has selected and defined a protocol type.

Creating VLANs

Accelar Device Manager allows creating port-based, source IP subnet-based, or protocol-based VLANs off the Basic VLAN Window by clicking on Insert.

The window for creating a port-based VLAN opens (Figure 8-3).

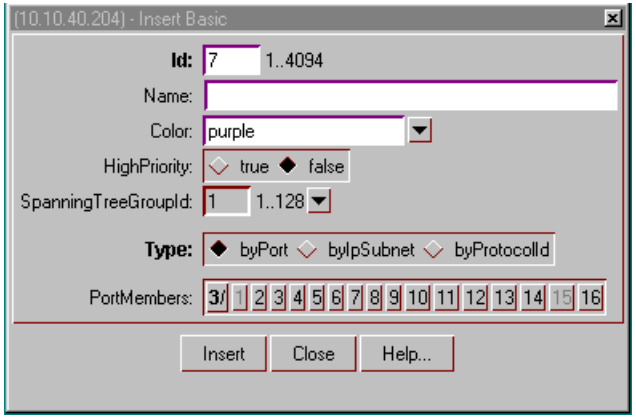

Figure 8-3. Create a Port-Based VLAN Window

To create a source IP subnet-based VLAN, changing the Type to byIpSubnet changes the screen layout as shown Figure 8-4 for creating an IP subnet-based VLAN.

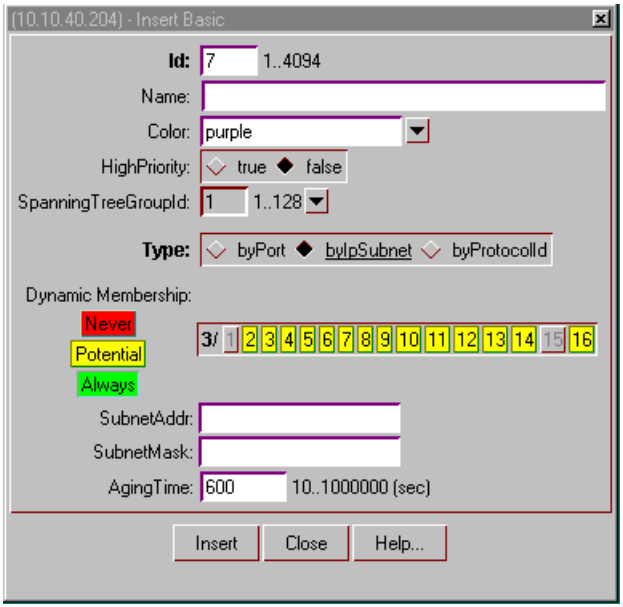

Figure 8-4. Create an IP Subnet-Based VLAN Window

To create a protocol-based VLAN, changing the Type to byProtocolId changes the screen layout to that pictured in Figure 8-5 for creating a protocol-based VLAN.

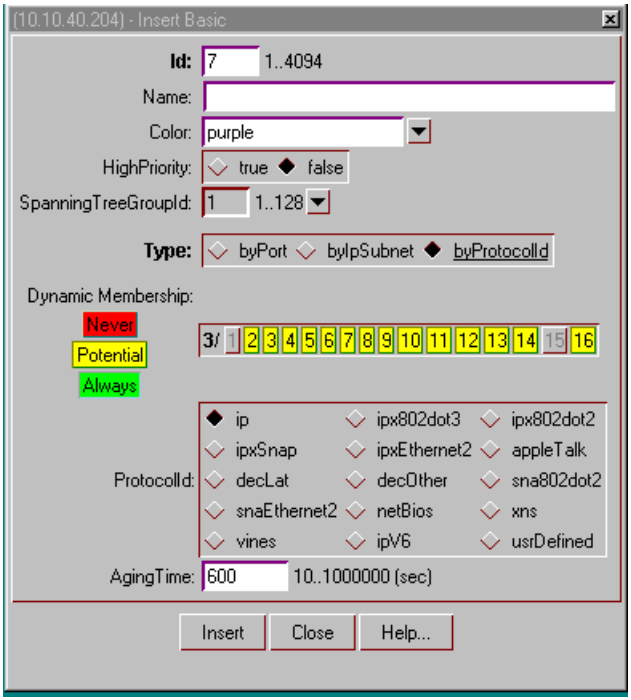

Figure 8-5. Create a Protocol-Based VLAN Window

To create all types of VLANs:

1. **Enter the VLAN ID.**
2. **Enter the VLAN name (optional).**

   If no name is entered, a default is created.
3. **Select the color (optional).**

   Accelar Device Manager will suggest a color, but it can be changed.
4. **Specify if traffic will be High Priority.**
5. **Select the Spanning Tree Group ID of the VLAN.**
6. **Select the Type of VLAN you are creating.**

7. **Specify the port membership by clicking on the port buttons.**
   - For port-based VLANs, you specify whether ports are Always members or Never members by selecting the member ports. These ports will display white, while the non-selected ports display as gray.
   - For source IP subnet- and protocol-based VLANs, you specify whether ports are:
     — Always members—static (green)
     — Never members—not allowed to join (red)
     — Potential members—dynamic (yellow)
8. **Source IP subnet-based VLANs and protocol-based VLANs also require the following:**
   a. **For source IP subnet-based VLANs, you must insert the source IP subnet address and IP subnet mask.**
   b. **For protocol-based VLANs, you must select the protocol.**

      If you are entering a UserDefined protocol, see the explanation under "User-Defined Protocols" on page 8-9.
   c. **For both source IP subnet-based and protocol-based VLANs, you should specify the Aging Time or use the default of 600 seconds (refer to Table 8-2 for a description of this parameter).**
9. **Click on the Insert button to create the VLAN.**

Note: In a protocol-based VLAN, a potential member becomes an active member of the VLAN when a frame of the specified protocol is received. In a source IP subnet-based VLAN, a potential member becomes an active member when a frame is received from the specified source IP address.

10. **Highlight the newly created VLAN and select IP->Insert to bring up the Insert IP Address window. Enter an IP address and click on Insert.**

User-Defined Protocols

You can create protocol-based VLANs with a user-defined protocol for integration into existing networks where nonstandard protocols are used. When the usrDefined button is selected, the Insert VLAN screen is displayed as shown in Figure 8-6.

In the UserDefinedPID field, enter the PID of the protocol in the format: 0x (protocol type in decimal value). The 16-bit PID assigned to a protocol-based VLAN specifies either an Ethertype, a DSAP/SSAP, or a SNAP PID, depending on whether the frame encapsulation is Ethernet 2, 802.2, or LLC-SNAP, respectively.

Refer to "User-Defined Protocols" and Table 1-2 on page 1-6 for more information on this topic, to see the actual values and how they are assigned.

The following PIDs are not valid:

- PID0x0000 through 0x05dc: overlap with the 802.3 frame length.
- PIDs of predefined protocols (for example, IP, IPX, AppleTalk).
- PID 0x8100: reserved by 802.1Q to identify tagged frames.
- PID0x9000: used by the diagnostic loopback frames.
- PID0x8808: used by 802.3x pause frames.
- PID0x4242: overlaps with the BPDU DSAP/SSAP.

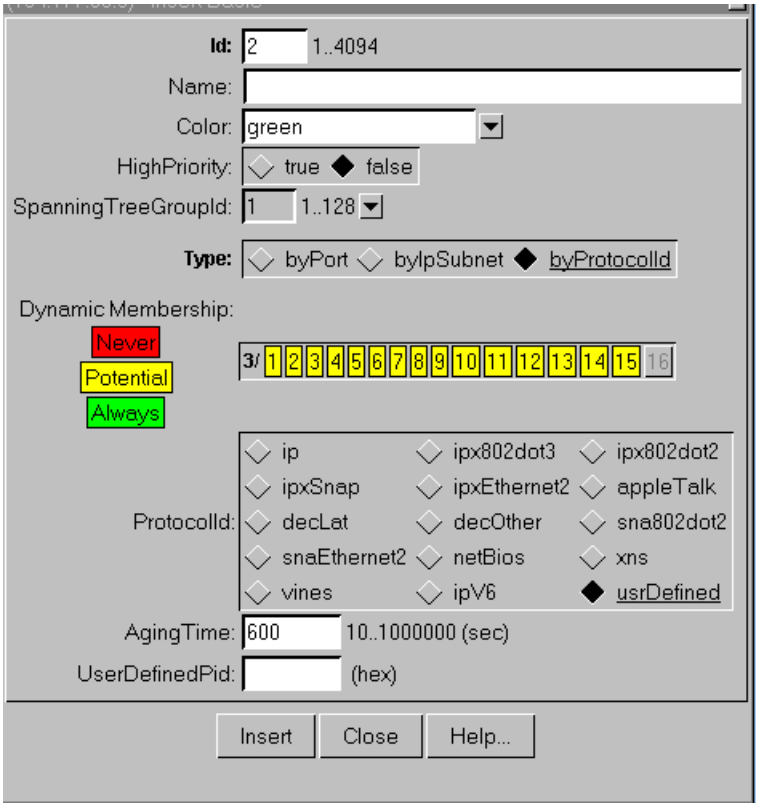

Figure 8-6. Entering User-Defined Protocol ID

Configuring Other VLAN Parameters

Some other VLAN parameters that should be considered are whether or not to discard frames on trunk and access ports and IGMP snooping parameters.

Accepting Tagged and Untagged Frames

You can select whether or not to discard tagged frames received on an access port and untagged frames received on a trunk port. The default is to discard the frames. You can also designate the port-based VLAN to which these frames are assigned by setting the trunk port's default VLAN ID (the default is VLAN 1).

To select to discard tagged frames received on a port, Click on the port and select Edit Port->VLAN.

Refer to "VLAN Window" on page 6-6 in Chapter 6, "Port Configuration and Graphing."

IGMP Snooping

IGMP Snooping allows the user to optimize the multicast data flow for a group within a VLAN only to the members of the group. This feature is set up through the VLAN->VLAN->Snoop window. For a description of this window along with more information about IGMP snooping and how to set it up, refer to Chapter 13, "IP Multicasting."

Modifying Existing VLANs

Existing VLANs are managed using the Basic VLAN Window (Figure 8-1) on page 8-3 and Advanced VLAN Window (Figure 8-2) on page 8-4 using the normal Device Manager GUI tools (refer to "Editing Objects" in Chapter 3, "Accelar Device Manager Basics").

Note: After a VLAN is created, the type of VLAN cannot be changed. The VLAN must be deleted and a new VLAN of the chosen type created.

Note: To edit the ports in a VLAN in the Basic VLAN Window, select a port member cell in the VLAN table and use the ellipses (...) icon to pull up a port selection tool with all ports in the STG of the VLAN available. Any changes made to the ports are made immediately.

Managing VLAN Bridging

Bridging occurs in layer 2 of the OSI model where only the MAC address in the packet header is considered when forwarding. With Accelar routing switches, all bridging is done within the context of a VLAN where each VLAN has its own bridging configuration and forwarding table.

To configure and monitor bridging:

1. **From the Accelar Device Manager menu bar, choose VLAN>VLANs>Basic>Bridging**.
2. **Select a VLAN; then click on Bridging at the bottom of the window.**

   The Edit Bridge window opens (Figure 8-7).Table 8-3 describes the fields.

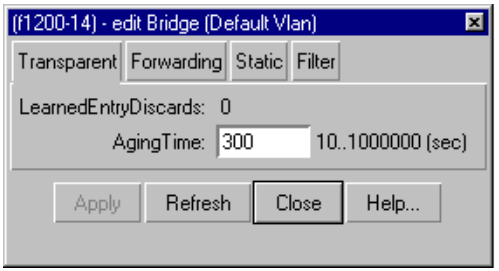

Figure 8-7. Edit Transparent Bridge Window

Table 8-3. Transparent Bridge Window Fields

Field	Description
LearnedEntryDiscards	The total number of Forwarding Database entries that have been or would have been learned but have been discarded due to a lack of space in the Forwarding Database. If this counter is increasing, it indicates that the Forwarding Database is regularly becoming full (a condition which has unpleasant performance effects on the subnetwork). If this counter has a significant value but is not presently increasing, it indicates that the problem has been occurring but is not persistent.
AgingTime	The timeout period in seconds for aging out dynamically learned forwarding information.The IEEE 802.1D-1990 standard recommends a default of 300 seconds. The actual aging time can be up to two times the AgingTime value.

VLAN Bridge Table

The VLAN Bridge Table window (Figure 8-8) is the forwarding database for the VLAN and contains information about unicast entries for which the bridge has forwarding and/or filtering information. This table is accessed by clicking on the Forwarding tab in the Edit Bridge window. This information is used by the transparent bridging function to determine how to forward a received frame. Refer to Table 8-4 for a description of the VLAN Bridge Table window fields.

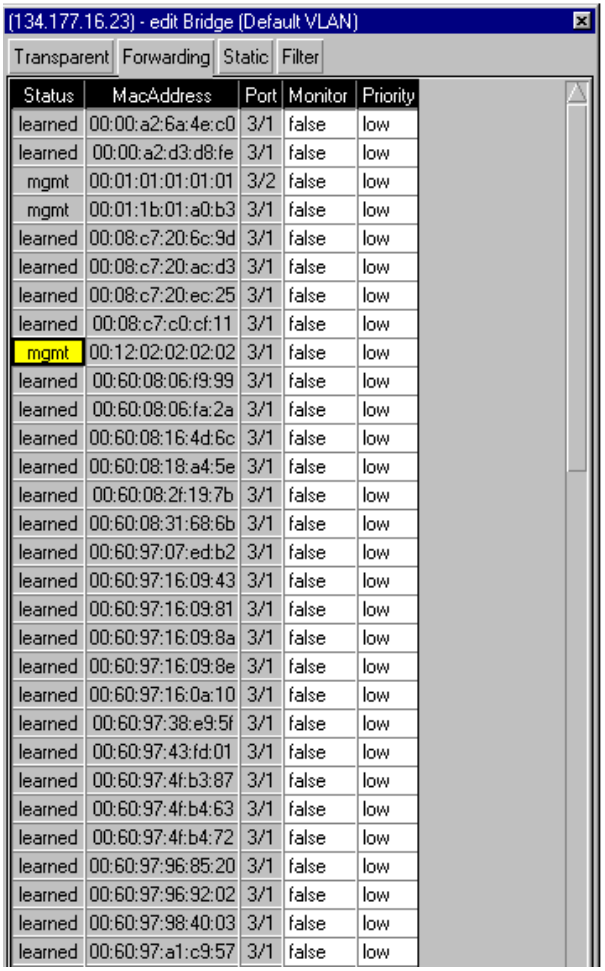

Figure 8-8. VLAN Bridge Table Window

Note: For troubleshooting purposes, it is sometimes necessary to manually flush the bridge forwarding database of learned MAC addresses.
The forwarding database can be flushed in two contexts:

- By Port—Delete all MAC addresses associated with a port for all VLANs under Edit Port -> Interface -> Action -> FlushMacFdb.
- By VLAN—Delete all MAC addresses associated with the VLAN under VLAN -> VLANs -> Advanced -> Action -> FlushMacFdb (Figure 8-9).

Table 8-4. VLAN Bridge Table Window Fields

Field	Description
Status	Values include: • self—one of the bridge's addresses. • learned—a learned entry that is being used. • mgmt—a static entry.
MacAddress	A unicast MAC address for which the bridge has forwarding and/or filtering information.
Port	Either a value of zero (0) or the port number of the port on which a frame having the specified MAC address has been seen. A value of 0 indicates a self-assigned MAC address.
Monitor	Select true or false to copy packets with MAC address in source or destination field. Used with port mirroring. For more information, refer to "Port Mirroring" on page 15-3.
Priority	Sets the priority of this entry as high or low relative to the other entries.

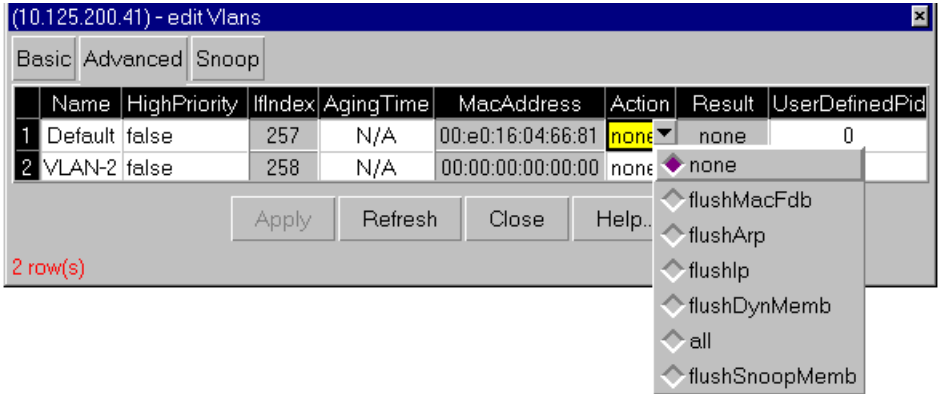

Figure 8-9. VLAN Flush Forwarding Database Window

VLAN Static Bridge Table

The VLAN Static Bridge Table window (Figure 8-10) contains static forwarding information configured into the bridge by (local or network) management specifying the set of ports to which frames received and containing specific destination addresses are allowed to be forwarded. Entries are valid for unicast and for group/broadcast addresses. Refer to Table 8-5 for a description of the VLAN Static Bridge Table window fields.

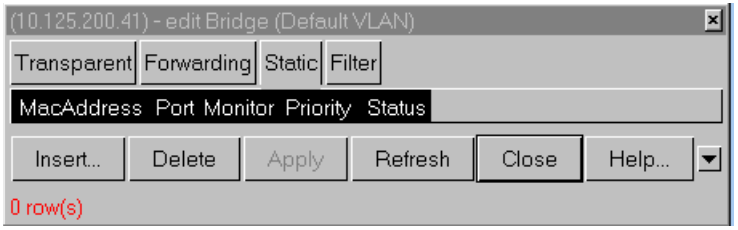

Figure 8-10. VLAN Static Bridge Table Window

Table 8-5. VLAN Static Bridge Table Window Fields

Field	Description
MacAddress	The destination MAC address in a frame to which this entry's forwarding information applies. This object can take the value of a unicast address.
Port	The port number of the port on which the frame will be received.
Monitor	Select true or false to copy packets with MAC address in source or destination field. Used with port mirroring. For more information, refer to "Port Mirroring" on page 15-3.
Priority	Sets the priority of this entry as high or low in relationship to the other entries.
Status	Indicates the status of this entry. Values can be one of the following: • permanent—in use and will remain so after the next bridge reset. This is the default value. • deleteOnReset—in use and will remain so until the next bridge reset. • deleteOnTimeout—currently in use and will remain so until it is aged. • other—in use but the conditions under which it will remain so are different from other values.

Bridging Filters

To perform MAC-layer bridging, the routing switch must know the destination MAC-layer address of each device on each attached network so it can forward packets to the appropriate destination. MAC-layer addresses are then stored in the bridging table, and you can filter packet traffic based on the destination MAC-layer address information.

The MAC filtering supported in the Accelar switches is the Bridge MIB filtering (RFC 1493). The number of MAC filters is limited to 100. You create a filter entry in much the same way as you create a static MAC entry, by entering a MAC address and the port on which it resides. In the MAC filter record, you also specify which ports are NOT ALLOWED to send traffic to that MAC on that port.

For example, if the filtered MAC address sends out an ARP request for a station on one of the NOT ALLOWED ports, the station will receive the ARP request and send a reply. The reply is what gets filtered by the routing switch in this instance, not the request from the filtered MAC.

To view a list of filters, select VLANs->VLANs->Bridging->Filters from Device Manager (Figure 8-11).

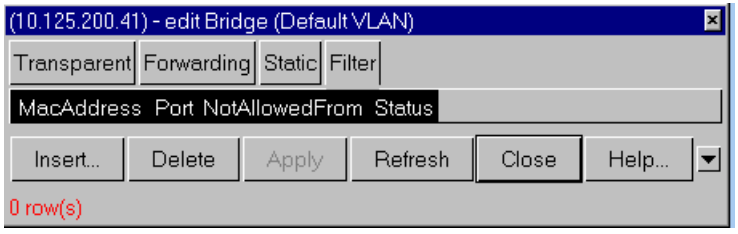

Figure 8-11. Bridge Filter Window

To filter traffic:

1. **Click on Insert at the bottom of the Bridge Filter window.**

   The Insert Filter window opens (Figure 8-12).

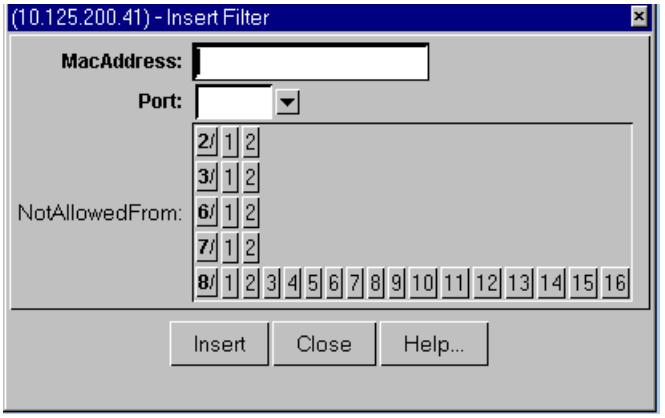

Figure 8-12. Insert Filter Window

2. **Enter a MAC address from the Bridge Filter window (Figure 8-11).**
3. **Select the port where this MAC address can be found.**
4. **Select the ports from which you do not want to receive packet traffic for this MAC address.**
5. **Click on Insert at the bottom of the window.**

Chapter 9
IP Interfaces and Router Management

This chapter describes basic IP router interfaces configuration and router management in Accelar Device Manager. It discusses the basic IP router interface configuration required before any routing protocols can be configured. Information about configuring RIP and OSPF are covered in Chapter 10, "Configuring and Managing RIP," and Chapter 11, "Configuring and Managing OSPF," respectively.

The router management features covered in this chapter apply regardless of which routing protocols are used and include router IP configuration, IP route table management, ARP configuration, ARP table management, BootP/DHCP relay configuration, and VRRP configuration.

Physical Versus Virtual Router Interfaces

There are two types of router interfaces: physical router interfaces (also called isolated router interfaces) and virtual router interfaces. These router interface types correspond to the two types of routing supported by an Accelar 1000 Series routing switch: routing on a physical port (also called routing on an isolated routing port) and routing on a virtual port that is associated with a VLAN.

Any port on the routing switch can be configured to be an isolated routing port. In this mode, the port only routes IP traffic and does not perform any bridging. The IP address is assigned to the port itself, and the router interface servicing the isolated routing port is called a physical router interface. Note that there is a one-to-one correspondence between the physical port and the router interface.

The other type of routing supported on an Accelar routing switch is the routing of IP traffic to and from a VLAN. Because a given port can belong to multiple VLANs (some of which are configured for routing on the switch and some of which are not), there is no longer a one-to-one correspondence between the physical port and the router interface. For VLAN routing, the router interface for the VLAN is called a virtual router interface because the IP address is assigned to an interface on the routing entity in the switch. This initial interface has a one-to-one correspondence with a VLAN on any given switch.

In an Accelar 1000 Series routing switch, the IP address of any physical or virtual router interface can be used for IP-based network management (SNMP, Telnet, and Web).

IP Interface Configuration

The steps required to configure IP for physical or virtual router interface can be broken down as follows:

1. **Verify that IP forwarding is enabled globally.**
2. **Assign an IP address and subnet mask to the interface.**
3. **Configure Address Resolution Protocol (ARP) for the interface.**
4. **Enable Dynamic Host Configuration Protocol (DHCP) relaying (optional).**
5. **Enable Virtual Router Redundancy Protocol (VRRP) (optional).**
6. **Enable Internet Group Management Protocol (IGMP) (optional).**
7. **Configure routing protocols (OSPF, RIP) for the interface (optional).**

The following sections give step-by-step instructions to complete the first five steps for isolated routing ports and virtual router ports. For information about configuring specific routing protocols on an IP interface, refer to the appropriate chapter (Chapter 13 for IGMP, Chapter 10 for RIP, and Chapter 11 for OSPF).

Configuring IP on an Isolated or Physical Routing Port

The following sections tell how to use the steps with Device Manager to configure IP on a physical or isolated routing port.

Assign an IP Address to the Port

To specify an IP address for an isolated routing port:

1. **Verify that routing is enabled by selecting Routing->IP from the menu and confirming that Forwarding is selected.**
2. **Select and Edit a port.**
3. **Select the IP Address tab and click on Insert.**

   The Insert IP Address window opens.
4. **Enter the IP address and mask.**
5. **Click on Insert.**

Note: You cannot edit the IP address, and you can assign only one IP address to any router interface (physical or virtual). Attempting to assign a second IP address returns an invalid IP address error.

Enable/Disable ARP on the Port

After the IP address is assigned, ARP can be configured. By default, ARP Response is enabled and Proxy ARP is disabled.

To enable or disable ARP on an isolated router port:

1. **Select and Edit a port.**
2. **Select the ARP tab, as illustrated in Figure 9-1.**
3. **In the DoResp field, click on disable or enable to select whether or not to respond to an ARP, and then click on Apply.**

   The default is enabled.
4. **In the DoProxy field, click on enable to enable Proxy ARP function (see "Using Proxy ARP" on page 9-14 for an explanation of the option).**

   The default is disabled.

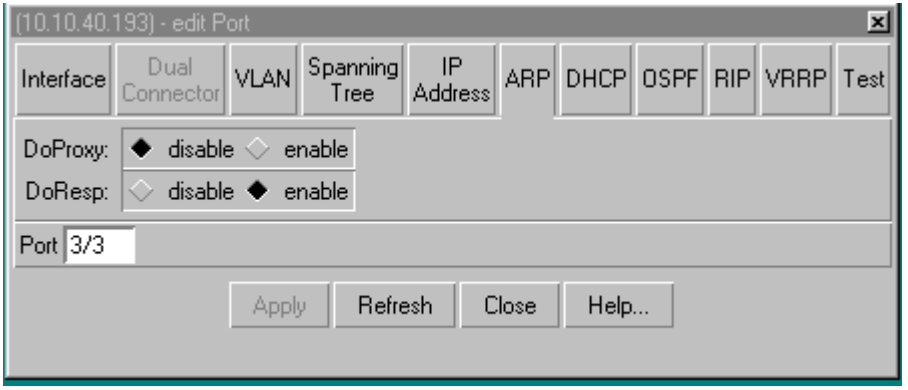

Figure 9-1. Edit Port ARP Configuration Window

Configuring IP on a Virtual Router Port

The following sections show how to configure IP on a virtual routing port.

Assign an IP Address to the VLAN

To specify an IP address for a virtual routing port:

1. **Verify that routing is enabled by selecting Routing->IP from the menu and confirming that Forwarding is selected.**
2. **From the Accelar Device Manager menu bar, choose VLAN->VLANs.**

   The Edit VLANs window opens (Figure 9-2).

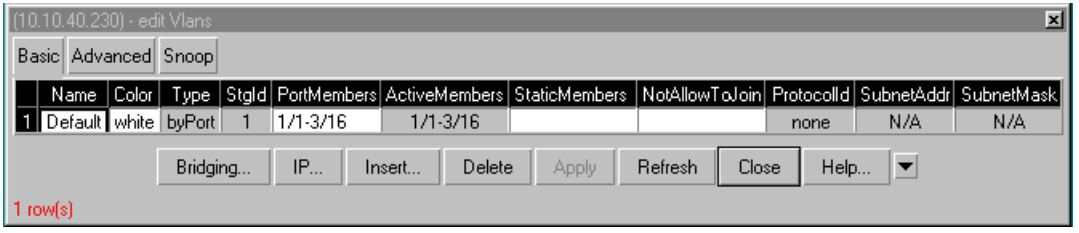

Figure 9-2. Edit VLANs Window

3. **Select the VLAN.**

4. **Click on IP at the bottom of the Edit VLANs window.**

   The Edit VLAN IP Address window opens (Figure 9-3).

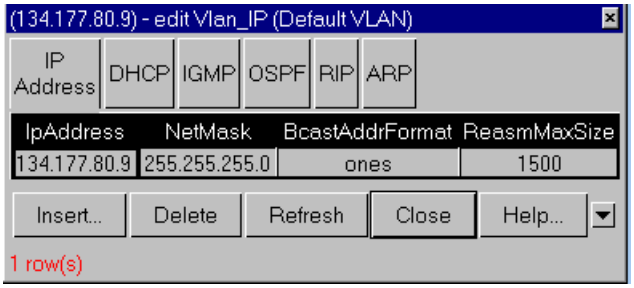

Figure 9-3. Edit VLAN IP Address Window

5. **Select the IP Address tab and click on Insert.**

   The Insert IP Address window opens.

6. **Enter the IP address and mask.**

7. **Click on Insert.**

Note: You can assign only one IP address to any router interface (physical or virtual). Attempting to assign a second IP address returns an invalid IP address error.

Enable/Disable ARP on the VLAN

After the IP address is assigned, ARP can be configured. By default, ARP Response is enabled and Proxy ARP is disabled.

To enable or disable ARP on a VLAN:

1. **From the Device Manager menu bar, choose VLAN->VLANs.**
2. **Click on the VLAN as shown in Figure 9-4.**

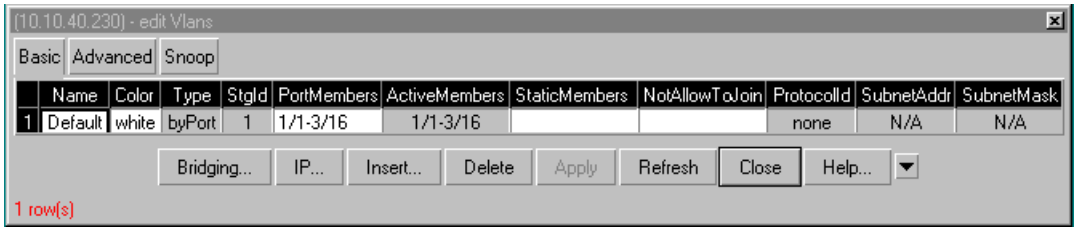

Figure 9-4. Selecting VLANs

3. **Select IP.**
4. **Click on ARP.**

   The VLAN ARP Configuration window opens (Figure 9-5).
5. **In the DoResp field, click on disable or enable to select whether or not to respond to an ARP, and then click on Apply.**

   The default is enabled.
6. **In the DoProxy field, click on enable to enable Proxy ARP function (see "Using Proxy ARP" on page 9-14 for an explanation of the option).**

   The default is disabled.

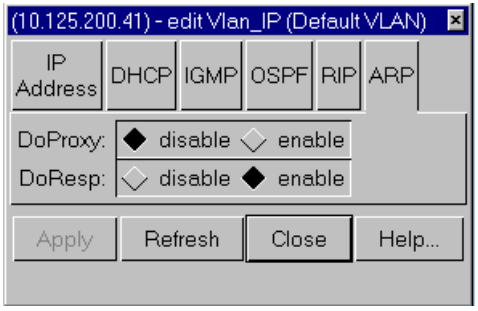

Figure 9-5. VLAN ARP Configuration Window

IP Router Management

In Accelar Device Manager, most of the windows related to managing the IP router are found under the Routing main menu selection.

The Routing->IP windows are the routing protocol independent windows and allow the network manager to configure the router's IP protocol stack and manage the routing tables. The Edit IP window includes the IP address, the IP route table, the IP flow table, and the ARP table.

The following sections describe the management facilities provided in the Routing -> IP windows except for high-priority IP Flow management, which is covered in Chapter 16, "Prioritization."

Router IP Configuration

The IP configuration window (Figure 9-6) contains parameters for configuring the router's IP protocol stack. The different options are described in Table 9-1.

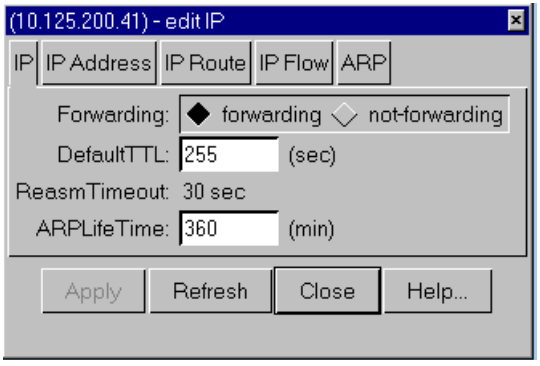

Figure 9-6. IP Configuration Window

Table 9-1. IP Configuration Window Fields

Field	Description
Forwarding	Sets the switch for IP forwarding (routing) or nonforwarding. The default is forwarding.
DefaultTTL	The default value inserted into the Time-To-Live field of the IP header of datagrams originated by the routing switch whenever a TTL value is not supplied by the transport layer protocol. Default is 255 seconds.
ReasmTimeout	The maximum number of seconds that received fragments are held while they are awaiting reassembly of this entity.This field cannot be changed by the user.
ARPLifeTime	The lifetime of an ARP entry within the system, global to the switch. Default is 360 minutes.

Router Interface Table

The Router Interface Table window (Figure 9-7) can be accessed from the Device Manager main menu under Routing->IP->IP Address. This window shows all the IP addresses defined in the routing switch and the associated router interfaces on which it is defined in one central location. The Interface column shows whether an IP address is configured for a physical router interface and a virtual router interface. For descriptions of the fields in the Router Interface Table window, refer to Table 9-2.

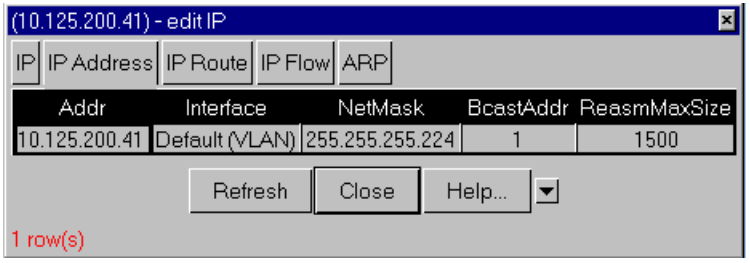

Figure 9-7. Router Interface Table Window

Note: This window shows all IP addresses defined in the box in a central location and is useful when trying to find IP address conflicts within the routing switch. None of the fields can be edited.

Table 9-2. Router Interface Window Fields

Field	Description
Addr	The IP address of the router interface.
Interface	The router interface. • Virtual router interfaces are identified by the name of the VLAN followed by the (VLAN) designation. • Physical interfaces are identified by the slot/port number of the isolated routing port.
NetMask	The subnet mask of the router interface.
BcastAddr	The IP broadcast address format used on this interface.
ReasmMaxSize	The size of the largest IP datagram that this entity can reassemble from incoming IP fragmented datagrams received on this interface (not editable).

IP Route Table

The IP Route Table window (Figure 9-8) is accessed from the Device Manager main menu under Routing->IP->IP Route. This window displays the contents of the system routing table and can be used to delete routes or to create static routes. From this window, you can delete any route, whether it is static or a dynamically learned route from RIP or OSPF. Therefore, you should exercise care when deleting entries from the route table.

The fields in the IP Route Table window are described in Table 9-3.

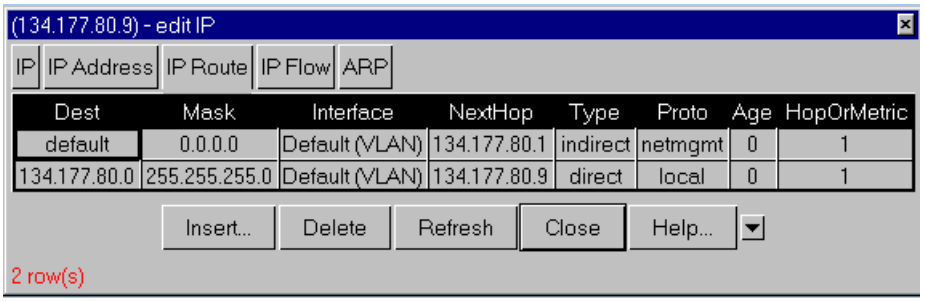

Figure 9-8. IP Route Table Window

Table 9-3. IP Route Table Window Fields

Field	Description
Dest	The destination IP network of this route. An entry with a value of 0.0.0.0 is considered a default route. Multiple routes to a single destination can appear in the table, but access to such multiple entries is dependent on the table access mechanisms defined by the network management protocol in use.
Mask	Indicate the network mask to be logically ANDed with the destination address before being compared to the value in the ipRouteDest field.
Interface	The router interface for this route. • Virtual router interfaces are identified by the VLAN number of the VLAN followed by the (VLAN) designation. • Physical interfaces are identified by the slot/port number of the isolated routing port.
NextHop	The IP address of the next hop of this route.
Type	The type of route: • direct • indirect Note that the values direct and indirect refer to the notion of direct and indirect routing in the IP architecture.
Proto	The routing mechanism through which this route was learned.
Age	The number of seconds since this route was last updated or otherwise determined to be correct.
HopOrMetric	The primary routing metric for this route. The semantics of this metric are specific to different routing protocols.

Creating Static Routes

Static routes are used to provide a mechanism to create routes to the destination IP address prefixes manually.

To create a static IP route:

1. **From the Accelar Device Manager menu bar, choose Routing->IP.**

   The Edit IP window opens.

2. **Select the IP Route tab and click Insert.**

   The Insert IP Route window opens, as shown in Figure 9-9.

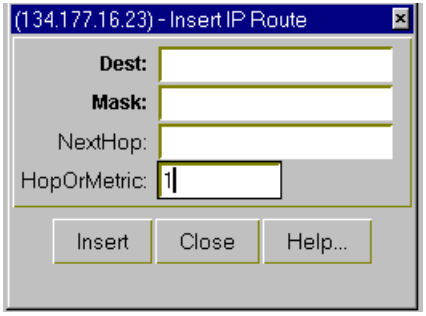

Figure 9-9. Insert IP Route Window

3. **Fill in the Dest and Mask fields with the IP route information.**
4. **Fill in the NextHop field (pointing to the router through which the specified route is accessible) and specify the HopOrMetric value. Click on Insert.**

   The route will now appear in the routing table.

Example: Creating a Static Default Route

The default route is used to specify a route to all networks for which there are no explicit routes in the Forwarding Information Base or the routing table. This route is by definition a route with the prefix length of zero [RFC1812]. The routing switches can be configured with the default route statically, or they can learn it via a dynamic routing protocol.

Note: To create a default static route, the destination address and subnet mask must be set to 0.0.0.0.

To create a static default route:

1. **From the Accelar Device Manager menu bar, choose Routing->IP.**

   The Edit IP window opens.

2. **Select the IP Route tab and click on Insert.**

   The Insert IP Route window opens, as shown in Figure 9-10.

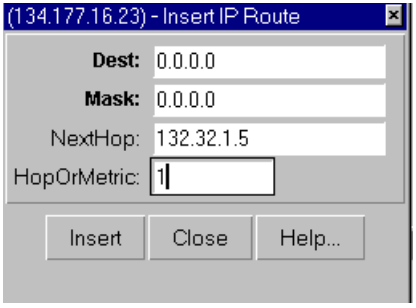

Figure 9-10. Insert IP Route Window

3. **Enter 0.0.0.0 in the Dest and the Mask fields.**

4. **Fill in the default NextHop router's IP address and the HopOrMetric value, and click on Insert.**

   The default route record is created in the routing table.

Router ARP Table

The router's ARP table window (Figure 9-11) can be accessed from the Device Manager main menu under Routing->IP->ARP. This window displays the known MAC address to IP address associations. Static ARP entries can be created and individual ARP entries deleted in this window. For descriptions of the fields in the ARP Table window, refer to Table 9-4.

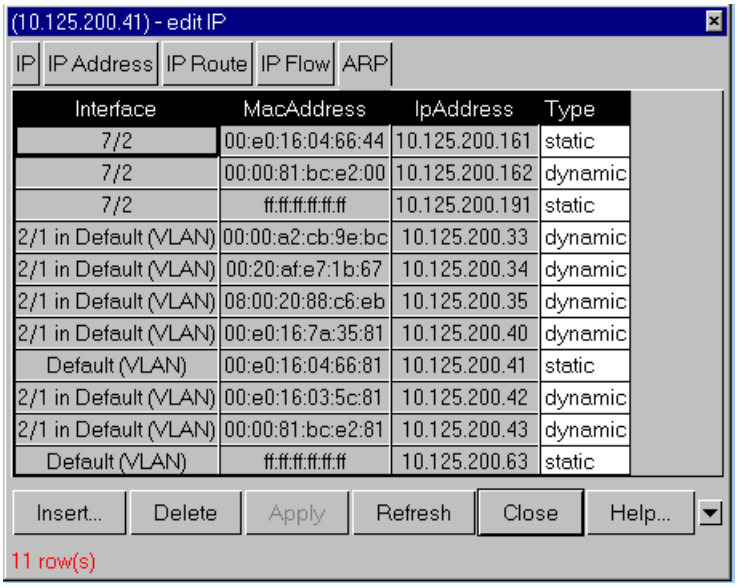

Figure 9-11. ARP Table Window

Table 9-4. ARP Table Window Fields

Field	Description
Interface	The router interface for this ARP entry: • Physical interfaces are identified by the slot/port number of the isolated routing port. • For virtual router interfaces, the physical slot/port and the name of the VLAN followed by the (VLAN) designation are specified.
MacAddress	The media-dependent physical address (that is, the Ethernet address).
IpAddress	The IP address corresponding to the media-dependent physical address.
Type	Type of ARP entry: • static—a statically configured ARP entry • dynamic—a learned ARP

Configuring Static ARP Entries

To configure static ARP entries:

1. **From the Accelar Device Manager menu bar, choose Routing->IP.**

   The Edit IP window opens.

2. **Select the ARP tab, and click on Insert.**

   The Insert ARP window opens, as shown in Figure 9-12.

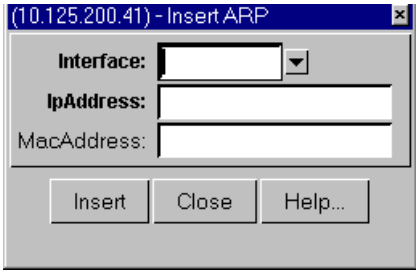

Figure 9-12. Insert ARP Window

3. **Specify the virtual or isolated router interface by selecting the down arrow in front of the Interface field.**

   This action specifies the interface connected to the station for which static ARP entry is being defined.

4. **Enter the IpAddress and the MacAddress fields, and click on Insert.**

   The static ARP entry appears in the ARP table, as shown in Figure 9-11.

Using Proxy ARP

Proxy ARP allows the Accelar routing switches to respond to an ARP request from a locally attached host or end station for a remote destination. It does so by sending an ARP response back to the local host with its own MAC address of the router interface for the subnet on which the ARP request was received. The reply is generated only if the switch has an active route to the destination network.

Figure 9-13 is an example of proxy ARP operation. Host B could send an ARP request for Host C. The Accelar routing switch would respond to the ARP request with Host C's IP address but with its own MAC address.

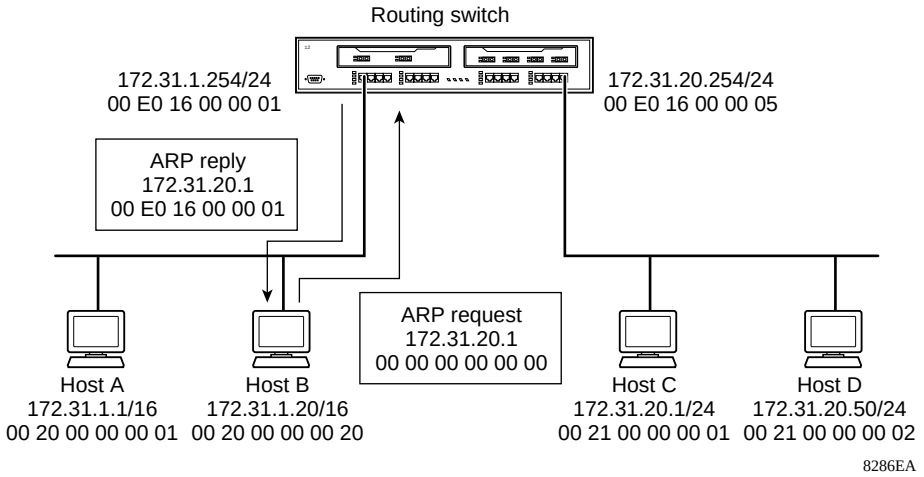

Figure 9-13. Proxy ARP Operation

To configure proxy ARP:

1. **From the Accelar Device Manager menu bar, choose VLAN->VLANs->Basic->IP->ARP.**

   The Edit VLANs window opens (Figure 9-14).

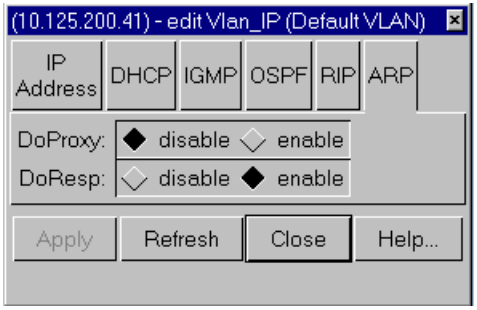

Figure 9-14. Enabling Proxy ARP

2. **Click on the DoProxy enable button, and then click on Apply.**

   Proxy ARP is enabled for the VLAN.

Flushing Router Tables

For administrative and/or troubleshooting purposes, it is sometimes necessary to flush the routing tables. Accelar Device Manager provides facilities for doing this in two contexts: by VLAN and by port.

In a VLAN context, all entries associated with the VLAN will be flushed. All the ARP entries and IP routes for a VLAN can be flushed under VLAN->VLANs->Advanced->Action. For more information, refer to Chapter 8, "Working with VLANs."

In a port context, all entries associated with the port will be flushed. The ARP entries and IP routes for a port can be flushed under Edit Port->Interface->Action. For more information, refer to Chapter 6, "Port Configuration and Graphing."

BootP/DHCP Relay

Dynamic Host Configuration Protocol (DHCP), an extension of the Bootstrap Protocol (BootP), is used to dynamically provide host configuration information to the workstations. To lower administrative overhead, network managers prefer to configure a small number of DHCP servers in a central location. Using few DHCP servers requires the routers connecting to the subnets or VLANs/bridge domains to support the BootP/DHCP relay function so that hosts can get the configuration information from servers several router hops away.

Differences Between DHCP and BootP

The following differences between DHCP and BootP are specified in RFC2131 and include functions that BootP does not address:

- DHCP defines mechanisms through which clients can be assigned a network address for a finite lease (allowing for reuse of IP addresses).
- DHCP provides the mechanism for clients to acquire all of the IP configuration parameters needed to operate.

DHCP uses the BootP message format defined in RFC 951. A packet is classified as DHCP if the first four octets in the options field are 99, 130, 83, 99 and the fifth octet is 53. The first four octets are referred to as the "Magic Cookie"; the fifth is the DHCP message type code. The remainder of the options field consists of a list of tagged parameters that are called "options" (RFC2131).

Summary of DHCP Relay Operation

BootP/DHCP clients (workstations) generally use UDP/IP broadcasts to determine their IP addresses and configuration information. If such a host is on a network or a subnet segment (or VLAN) that does not include a DHCP server, the UDP broadcasts are by default not forwarded to the server located on a different network segment or VLAN. The Accelar routing switches can be configured to overcome this issue by forwarding the broadcasts to the server through isolated or virtual router interfaces. The router interfaces can be configured to forward DHCP broadcasts to other locally connected network segments or directly to the server's IP address. DHCP must be enabled on a per-routable-interface basis.

In Figure 9-15, an end station is connected to subnet 1, corresponding to VLAN 1. The Accelar routing switch connects two subnets via the virtual routing function. When the end station generates a DHCP request as a limited UDP broadcast to the IP address of all 1s (that is, 255.255.255.255) with the DHCP relay function configured, the Accelar routing switch forwards DHCP requests to subnet 2 or to the host address of the DHCP server, depending on the configuration.

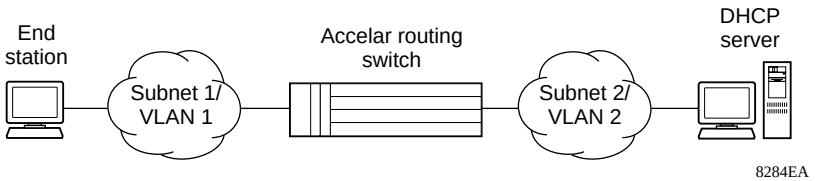

Figure 9-15. Example of DHCP Operation

To set up a forwarding path for BootP/DHCP packets received on an interface enabled for DHCP relaying:

1. **Choose Routing->DHCP.**

   The edit DHCP window (Figure 9-16) opens.

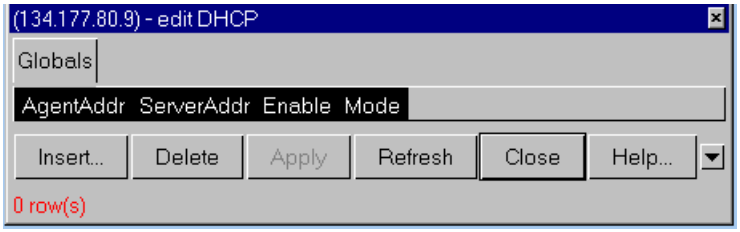

Figure 9-16. Edit DHCP Window

2. **Click on Insert.**

   The Insert Globals window (Figure 9-17) opens.

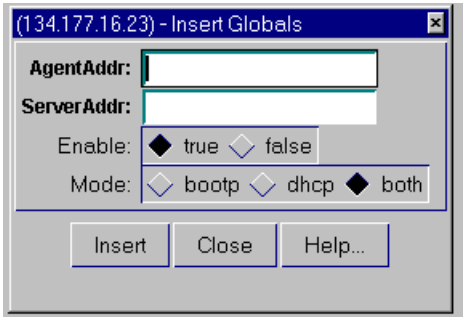

Figure 9-17. Insert Globals Window

3. **Type in the Agent Address.**

   This parameter specifies the IP address of the input interface on which the relaying of received BootP/DHCP packets must be enabled.

4. **Type in the Server Address.**

   This parameter is either the IP address of the BootP/DHCP server or the address of another local interface of the switch. If it is the address of the BootP/DHCP server, then the request is unicast to the server's address. If the address is one of the IP addresses of an interface on the switch, then the BootP/DHCP requests will be broadcast out of that local interface.

5. **Enable or disable BootP/DHCP relay.**

   The default is enabled.

6. **Select the type of messages to be relayed.**

   The default is to relay both BootP and DHCP messages.

Forwarding DHCP Packets

In the example shown in Figure 9-18, the Agent Address is: 10.10.1.2.

- To configure the Accelar routing switch to forward DHCP packets from the end station to the server, use 10.10.2.1. as the Server Address.

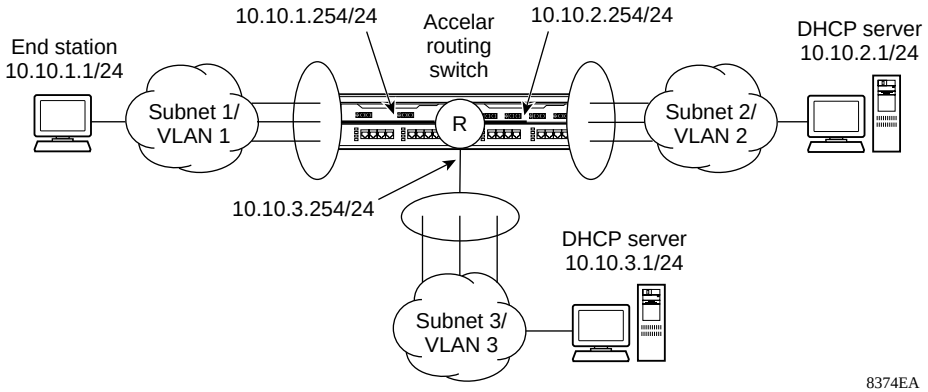

Figure 9-18. Forwarding DHCP Packets

All BootP broadcast packets, including DHCP packets that appear on the VLAN 1 router interface (10.10.1.2), will be forwarded to the DHCP server. In this case, the DHCP packets will be forwarded as unicast to the DHCP server's IP address.

- To forward BootP/DHCP packets as broadcast packets to VLAN 2, specify the IP address of the switch VLAN2 router interface (10.10.2.2.) as the Server Address.

Multiple BootP/DHCP Servers

Most enterprise networks use multiple BootP/DHCP servers for fault tolerance. The Accelar routing switches allow configuring to forward the BootP/DHCP requests to multiple servers. Up to 10 servers can be configured to receive copies of the forwarded/relayed BootP/DHCP messages.

If a DHCP client is connected to a routable interface, to configure DHCP requests to be sent to 10 different routable interfaces or 10 different server IP addresses, enable DHCP on the client (Agent Address) and then enable DHCP from the client to each of the interfaces or IP addresses (Server Addresses).

In the example shown in Figure 9-19, two DHCP servers are located on two different subnets. To configure the Accelar routing switch to forward the copies of the BootP/DHCP packets from the end station to both servers, specify the routing switch (10.10.1.254) as the Agent Address. Then enable DHCP to each of the DHCP servers by entering 10.10.2.1. and 10.10.3.1. as the Server Addresses.

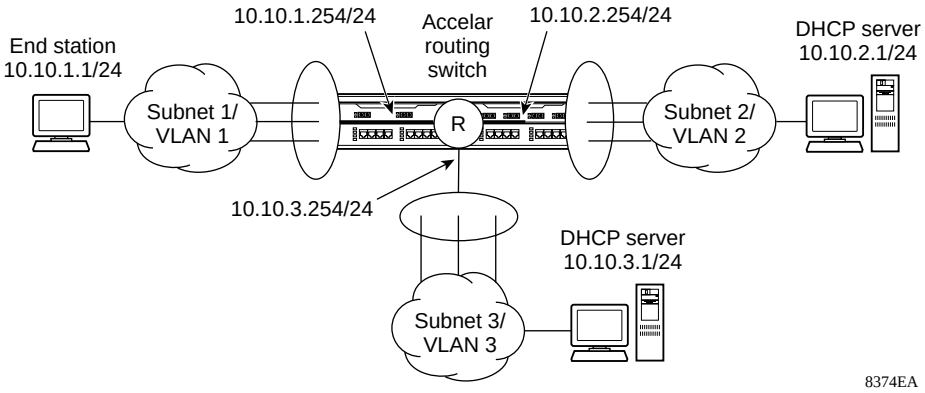

Figure 9-19. Configuring Multiple BootP/DHCP Servers

VRRP

End stations are often configured with a static default gateway IP address. Loss of the default gateway router can have catastrophic results. Virtual Router Redundancy Protocol (VRRP) is designed to eliminate this single point of failure routed environment by introducing the concept of a virtual IP address (transparent to users) shared between two or more routers connecting the common subnet to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides a dynamic default gateway redundancy in the event of a failure.

Four VRRP interfaces (isolated routing ports *and* VLANs) are allowed per Accelar switch and all VRIDs must be unique.

To set up VRRP parameters:

- On a port, select Edit->Port->VRRP.
- On a VLAN, select VLAN->VLANS->Basic->IP->VRRP.

The Port VRRP window (Figure 9-20) and the VLAN VRRP window have the same fields.

Click on Insert to view the Insert VRRP window (Figure 9-20). The window fields are described in Table 9-5.

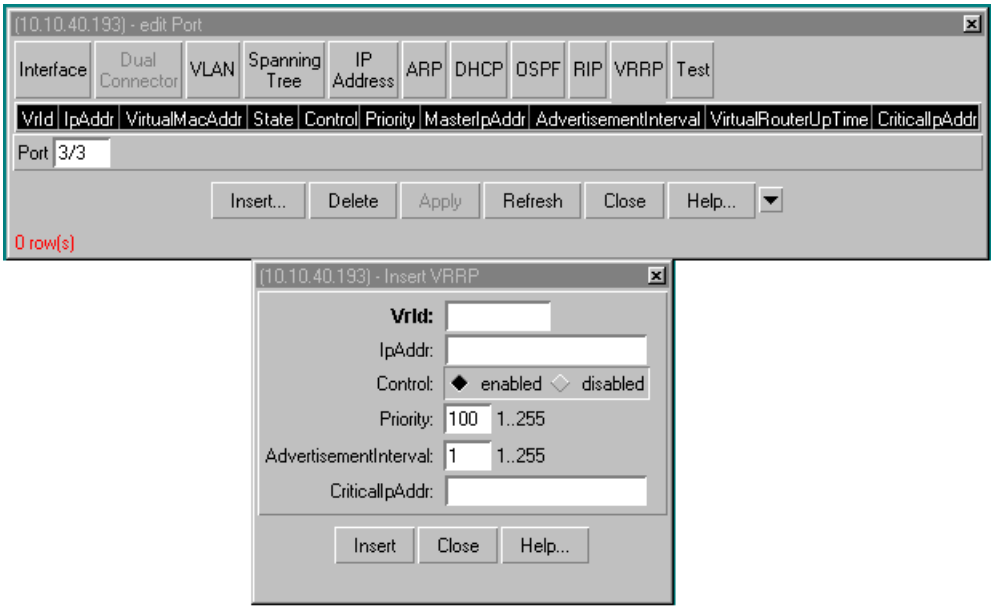

Figure 9-20. Edit Port VRRP and Insert VRRP Windows

Table 9-5. Port or VLAN VRRP Window Fields

Field	Description
Vrid	A number that uniquely identifies a virtual router on a given VRRP router. The virtual router acts as the default router for one or more assigned addresses. (1-255)
IpAddr	IP address of the virtual router interface.
VirtualMacAddr	MAC address of the virtual router interface.
State	The state of the virtual router interface: • initialize: waiting for a startup event • backup: monitoring availability and state of the master router • master: functioning as the forwarding router for the virtual router IP address(es)
Control	Whether VRRP is enabled or disabled for the port or VLAN.
Priority	Priority value to be used by this VRRP router. Set a value from 1 to 255, where 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
MasterIpAddr	The IP address of the physical interface of the master virtual router that has the responsibility of forwarding packets sent to the virtual IP address(es) associated with the virtual router.
Advertisement Interval	The time interval (in seconds) between sending advertisement messages. Set from 1 to 255 seconds with a default of 1 second. Only the master router sends advertisements.
VirtualRouter UpTime	The time interval (in hundredths of a second) since the virtual router was initialized.
CriticalIPAddr	An IP interface on the local router configured so that a change in its state would cause a role switch in the virtual router (for example, from master to backup) in case the interface went down.

Chapter 10
Configuring and Managing RIP

This chapter describes configuring and managing RIP on an Accelar 1000 Series routing switch using Accelar Device Manager.

For information about configuring OSPF, refer to Chapter 11, "Configuring and Managing OSPF."

There are three steps to configuring RIP on a router interface:

1. Configure RIP global parameters (this page).
2. Enable and configure RIP on the interface (page 10-3).
3. Configure the RIP version on the interface (page 10-5).

Note: The information in this chapter assumes the user has already created the router interface (either an isolated routing port or a virtual routing interface for a VLAN) and assigned an IP address. If an IP address has not been assigned, refer to Chapter 9, "IP Interfaces and Router Management," for information about creating router interfaces and assigning IP addresses.

Configure RIP Global Parameters

In the Accelar 1000 Series routing switch, the router has RIP global parameters that are used by all router interfaces using RIP. Both isolated routing ports and VLAN virtual routing interfaces use the same RIP global parameters.

The RIP global parameters in Accelar Device Manager are accessible off the main menu under Routing->RIP.

The RIP Globals window pictured in Figure 10-1 contains the two most important globals, including whether or not RIP is enabled for the routing switch (Operation: enable/disable) and the RIP update timer, which is the time between RIP updates on all interfaces.

Note: You can configure RIP on the interfaces with RIP globally disabled, thus having the flexibility to configure all interfaces before turning on RIP for the routing switch.

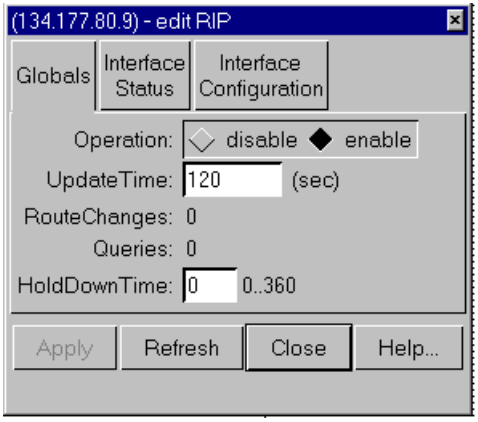

Figure 10-1. RIP Globals Window

Table 10-1 describes the RIP Globals window fields.

Table 10-1. RIP Globals Window Fields

Field	Description
Operation	Enable or disable the operation of RIP on all interfaces.
Update Time	The RIP Update Time refers to the time interval between RIP updates. It is a global parameter for the box; that is, it applies to all interfaces and cannot be set individually for each interface.
RouteChanges	The number of route changes made to the IP Route Database by RIP; does not include the refresh of a route's age.
Queries	The number of responses sent to RIP queries from other systems.
HoldDown Time	Sets the length of time that RIP will continue to advertise a network after determining it is unreachable. From 0 to 360 seconds.

Enable and Configure RIP on the Interface

Note: The screen shots in this section are for a virtual router interface for a VLAN. The screens for configuring an isolated routing port have the same parameters, and the parameters function the same.

The RIP enable and configuration parameters for an isolated router port are under Edit->Port->RIP. The RIP configuration parameters for a virtual router interface are part of a VLAN's routing parameters. They are found under VLAN->VLANs->Basic->IP->RIP (Figure 10-2).

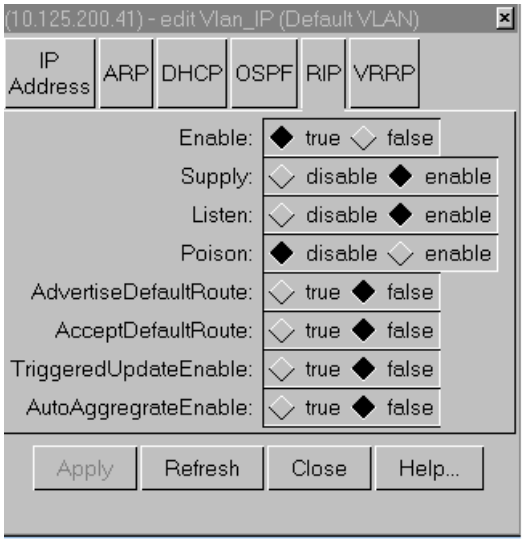

Figure 10-2. RIP Interface Parameters Window

The RIP Interface parameters are described in Table 10-2.

Table 10-2. RIP Interface Parameters Window Fields

Option	Description
Enable	The Enable field sets enabling RIP on the VLAN (or port) to true or false.
Supply	Enables or disables RIP route advertisements through the interface.
Listen	Enables or disables the learning of RIP advertised routes through this interface.
Poison	If disabled, split horizon is invoked, meaning that IP routes learned from an immediate neighbor are not advertised back to the neighbor from which the routes were learned. If enabled, the RIP update sent to a neighbor from which a route is learned is “poisoned” with a metric of 16. In this manner, the route entry is not passed along to the neighbor, because historically 16 is “infinity” in terms of hops on a network.
AdvertiseDefaultRoute	Set value to true if default route must be advertised out this interface. Default route will be advertised only if it exists in the routing table.
AcceptDefaultRoute	Set value to true if default route should be learned on this interface when advertised by another router connected to the interface.
TriggeredUpdateEnable	Sets whether to disable or enable automatic triggered updates for RIP.
AutoAggregateEnable	Enables automatic route aggregation. Only available when using RIPv2.

Configure the RIP Version

For interfaces configured to send (Supply) or receive (Listen to) RIP updates, the version of RIP to use can be configured under Routing->RIP->Interface Configuration (Figure 10-3).

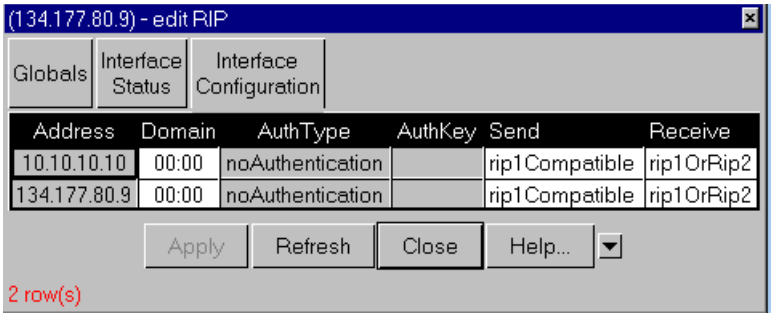

Figure 10-3. RIP Interface Configuration Window

Note: The AuthType and AuthKey parameters are not supported.

Table 10-3 describes the RIP Interface Configuration window fields.

Table 10-3. RIP Interface Configuration Window Fields

Field	Description
Address	The IP address of the router interface.
Domain	The value inserted into the Routing Domain field of all RIP packets sent on this interface.
AuthType	The type of authentication used on this interface.
AuthKey	The value to be used as the Authentication Key whenever the corresponding instance of rip2IfConfAuthType has a value other than noAuthentication.
Send	What the router sends on this interface (selected from a pull-down menu): • DoNotSend—no RIP updates sent on this interface • ripVersion1—RIP updates compliant with RFC 1058 • rip1Compatible—broadcast RIP-2 updates using RFC 1058 route subsumption rules • ripVersion2—multicasting RIP-2 updates
Receive	Indicates which versions of RIP updates are to be accepted: • rip1 • rip2 • rip1OrRip2 Note that rip2 and rip1OrRip2 imply reception of multicast packets.

RIP Interface Status

Statistics on RIP protocol are kept by the Accelar 1000 Series routing switch and are available under Routing->RIP->Interface Status (Figure 10-4). For a description of the RIP interface status statistics, refer to Table 10-4.

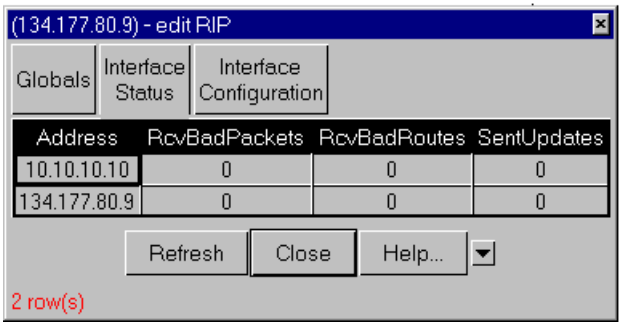

Figure 10-4. RIP Interface Status Window

Table 10-4. RIP Interface Status Window Fields

Field	Description
Address	The IP address of the router interface.
RcvBadPackets	The number of RIP response packets received by the RIP process that were subsequently discarded for any reason (Examples: a version 0 packet or an unknown command type).
RcvBadRoutes	The number of routes, in valid RIP packets, that were ignored for any reason (Examples: unknown address family or invalid metric).
SentUpdates	The number of triggered RIP updates actually sent on this interface. This field explicitly does *not* include full updates sent containing new information.

Chapter 11
Configuring and Managing OSPF

The Open Shortest Path First (OSPF) protocol is the primary TCP/IP routing protocol. Routers use OSPF to exchange network topology information among themselves, giving each router a map of the network. By searching their maps, routers know how to move packets through the network to their destinations.

This chapter contains information about the following OSPF routing topics:

- Classifications and descriptions of the different router types (this page)
- General examples of OSPF configurations on different network topologies (page 11-3)
- Creating virtual links (page 11-26)
- Specifying autonomous system border routers (ASBRs) (page 11-32)
- Creating stub areas (page 11-32)
- Changing metrics and specifying redistribution (page 11-33)
- Description of OSPF windows and fields (page 11-35)

Most of this chapter is task-oriented, showing you how to configure OSPF. However, Table 11-6 on page 11-35 lists and describes all the OSPF windows and fields.

Descriptions of Router Types

Routers deployed in an OSPF network can take on different roles depending on how they are configured. Table 11-1 provides a brief description of each possible router role. These descriptions are intended to assist you with terminology used in "OSPF Examples" starting on page 11-3.

Table 11-1. Router Classifications

Router Type	Description
AS Boundary Router (ASBR)	A router attached at the edge of an OSPF network is considered an AS Boundary Router (ASBR). An ASBR generally has one or more interfaces that run an Inter-Domain Routing Protocol (IDRP) such as BGP. In addition, any router distributing static routes or RIP routes into OSPF is considered an ASBR. The ASBR forwards routes learned from IDRP into the OSPF domain. In this way, routers inside the OSPF network learn about destinations outside their domain.
Area Border Router (ABR)	A router attached to two or more areas inside an OSPF network is considered an Area Border Router (ABR). ABRs play an important role in OSPF networks by limiting the amount of OSPF information that gets disseminated.
Internal Router (IR)	A router that only has interfaces within a single area inside an OSPF network is considered an Internal Router (IR). Unlike ABRs, IRs have topological information only about the area in which they are contained.
Designated Router (DR)	In a broadcast network, such as an Ethernet network that has more than one router locally attached, a single router is elected to be the Designated Router (DR) for that broadcast network. A DR assumes the responsibility of making sure all routers on the broadcast network are in synchronization with one another.
Backup Designated Router (BDR)	In a broadcast network, such as an Ethernet network, a Backup Designated Router (BDR) is elected in addition to the Designated Router (DR). The BDR assumes essentially the same responsibilities as the DR; if the DR fails, the BDR will assume the role of the DR in the broadcast network.
Other Router (OR)	In a broadcast network, such as an Ethernet network, any router not elected to be a Designated Router (DR) or Backup Designated Router (BDR) is considered to be an Other Router (OR).

This manual does not attempt to provide detailed information about how OSPF operates, but instead focuses on providing examples of how an OSPF network can be configured and provides three basic network examples. For detailed information about OSPF protocol concepts and terminology, refer to Chapter 7, “Customizing OSPF Services,” in *Configuring IP Services* (Bay Networks part number 117356-B).

OSPF Examples

The following sections contain examples of configuring OSPF in three basic types of network configurations. The examples are for the most common network configurations and are designed to provide procedural instructions for installing and operating OSPF on these common networks. The common network configurations of the examples are summarized in Table 11-2.

Table 11-2. Summary of Examples

Example	Network Configuration
Example 1	Configuring two routers for running OSPF on the same subnet.
Example 2	Configuring routers on different subnets in the same area.
Example 3	Configuring two routers on the same subnet in one area and two routers on a second subnet in a second area. The second switch is configured as the area border router for both networks.

Example 1: Configuring OSPF on One Subnet

The first example shows how to configure OSPF on two switches located on one subnet. This configuration is illustrated in Figure 11-1.

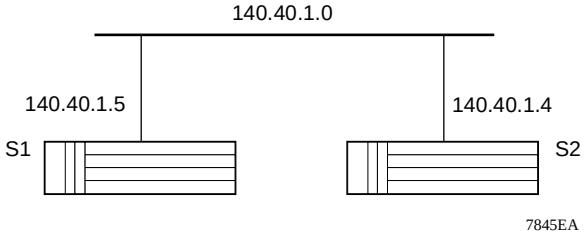

Figure 11-1. Example 1: OSPF on One Subnet

Table 11-3 identifies switches used in Example 1.

Table 11-3. Switch Identifiers for Example 1

Switch Number	Switch IP Address	Interface IP Address	Subnet
Switch 1	134.177.160.101	140.40.1.5	1
Switch 2	134.177.160.102	140.40.1.4	1

To enable OSPF on a switch:

1. **From the Accelar Device Manager menu bar, open the switch.**
2. **From the Accelar Device Manager menu bar, choose Routing->OSPF->General.**

   The OSPF General window, as shown in Figure 11-2, displays parameter values that apply globally to the router's OSPF configuration.

   Notice that the name or IP address of the device is always displayed in the upper left corner of the title bar.
3. **To activate OSPF, select enabled in the AdminStat field and click on Apply at the bottom of the window.**

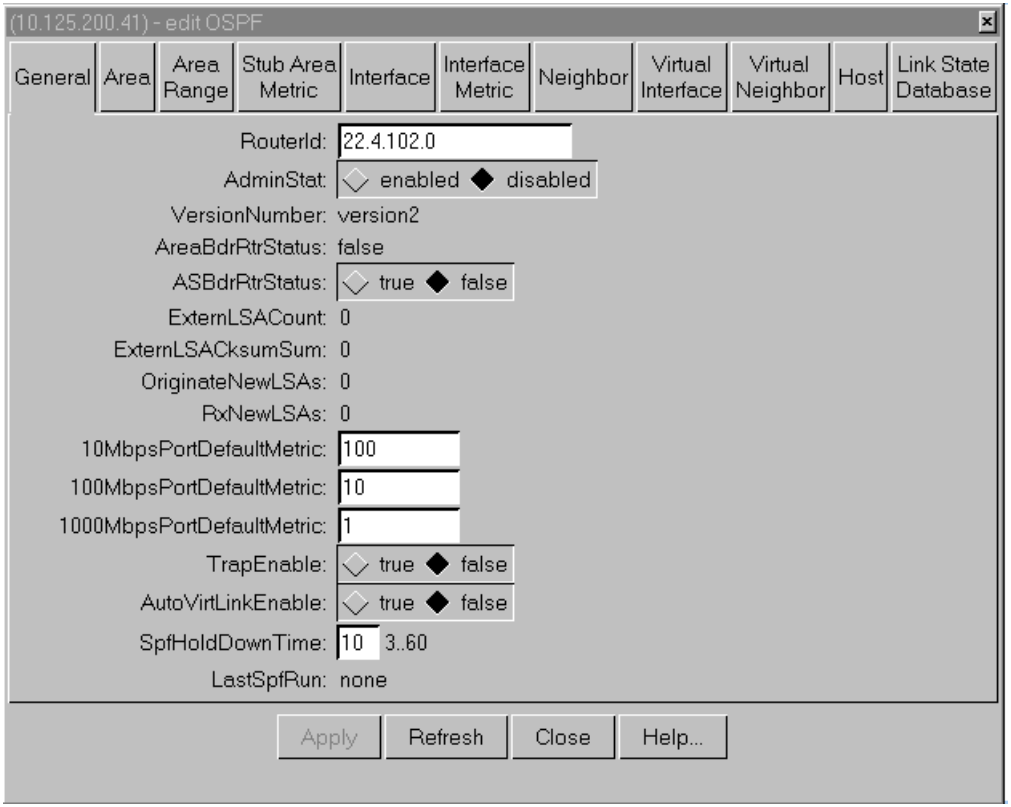

Figure 11-2. OSPF General Window

To assign the IP address:

1. **In the Edit Port window, click on the IP address tab.**
2. **Click on Insert at the bottom of the window.**

   The Insert IP Address dialog box opens, as shown in Figure 11-3.

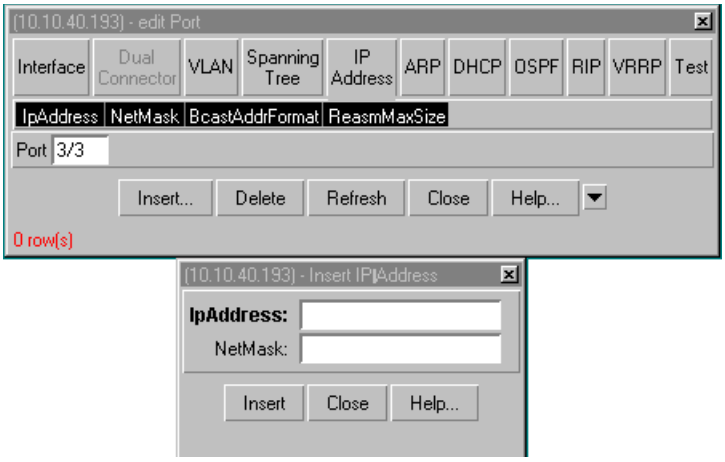

Figure 11-3. Insert IP Address Dialog Box

3. **In the IpAddress field, type the interface IP address and press [Tab].**

Note: Pressing [Tab] automatically enters the default net mask in the field below the IP address. This field also can be edited manually.

4. **Click on Insert at the bottom of the window.**
5. **To apply changes, click on Apply.**

To enable OSPF for an interface:

1. **On the Accelar Device Manager graphical representation of the switch, as shown in Figure 11-4, click on an interface that you want to enable for routing.**

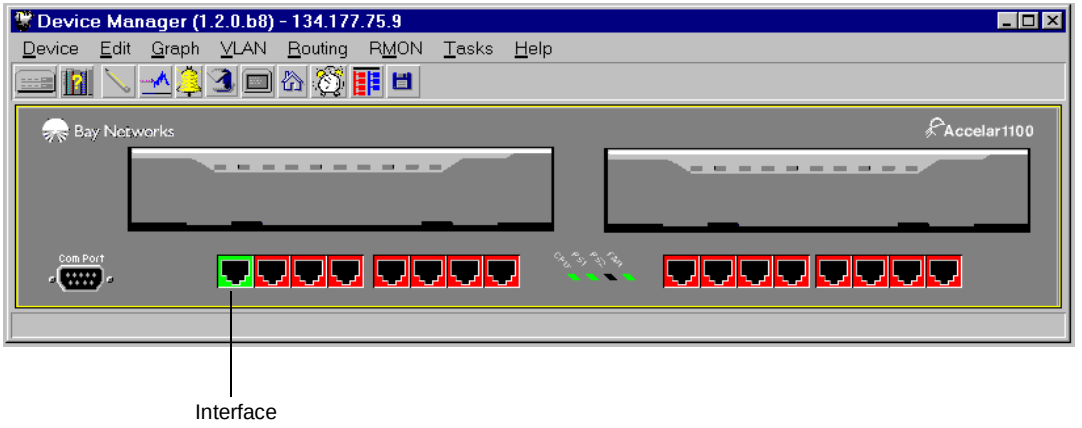

Figure 11-4. Accelar Device Manager Graphical Representation of a Switch

2. **From the Accelar Device Manager menu bar, choose Routing->IP.**

   The Edit IP window shown in Figure 11-5 opens. Verify that IP forwarding is forwarding for the switch. If not, select forwarding and click on Apply.

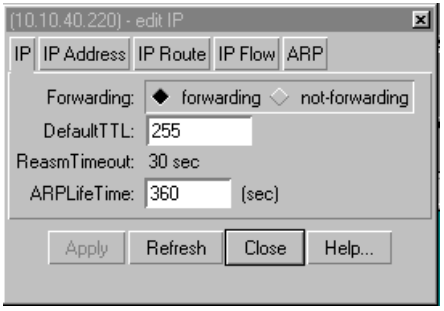

Figure 11-5. IP Forwarding Enable

3. **From the menu bar, select the Edit->Port->RIP and turn off RIP by selecting false in the Enable field, as shown in Figure 11-6.**

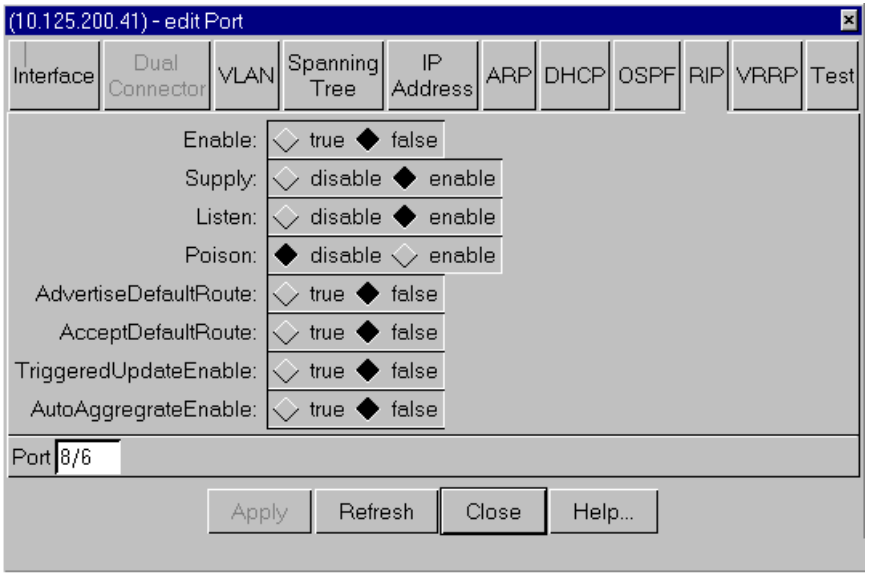

Figure 11-6. Edit Port RIP Window

4. **To apply changes, click on Apply.**
5. **To turn on OSPF, select Edit->Port->OSPF, and then select true in the Enable field shown in Figure 11-7.**

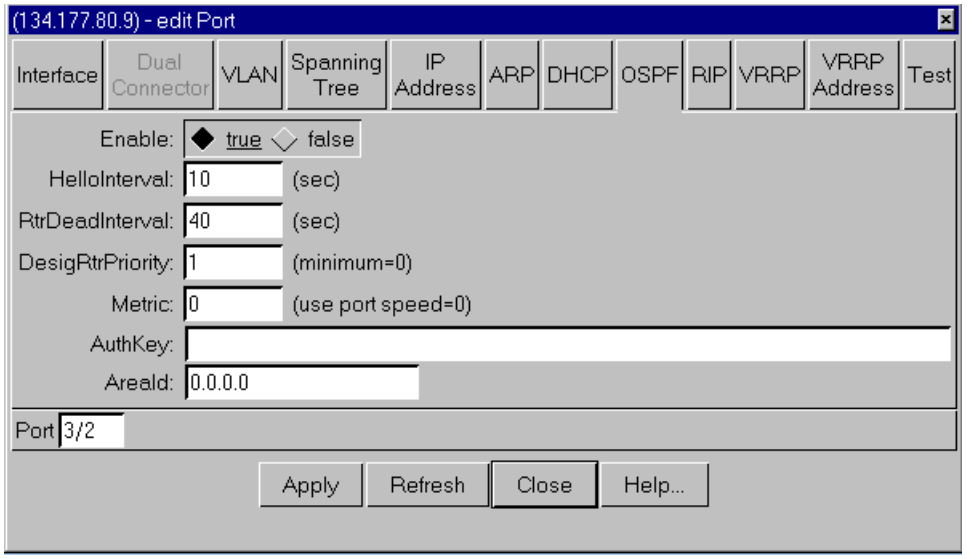

Figure 11-7. Edit Port OSPF Window

To configure a second interface:

1. **Enable OSPF on a switch (page 11-4).**
2. **Insert the IP address (page 11-5).**
3. **Enable OSPF for an interface (page 11-6).**

After you have configured a second interface, the two interfaces begin exchanging hello packets.

OSPF configuration for two interfaces on the same network is now complete. If you want to see how this network configuration appears under Accelar Device Manager, review the following sample screens that display information about IP routing and OSPF.

To view the network configuration under Accelar Device Manager:

1. **From the Accelar Device Manager menu bar, choose Routing->OSPF->Interface.**

   The OSPF Interface window opens, as shown in Figure 11-8. This window displays information about the OSPF interface configured on the router.

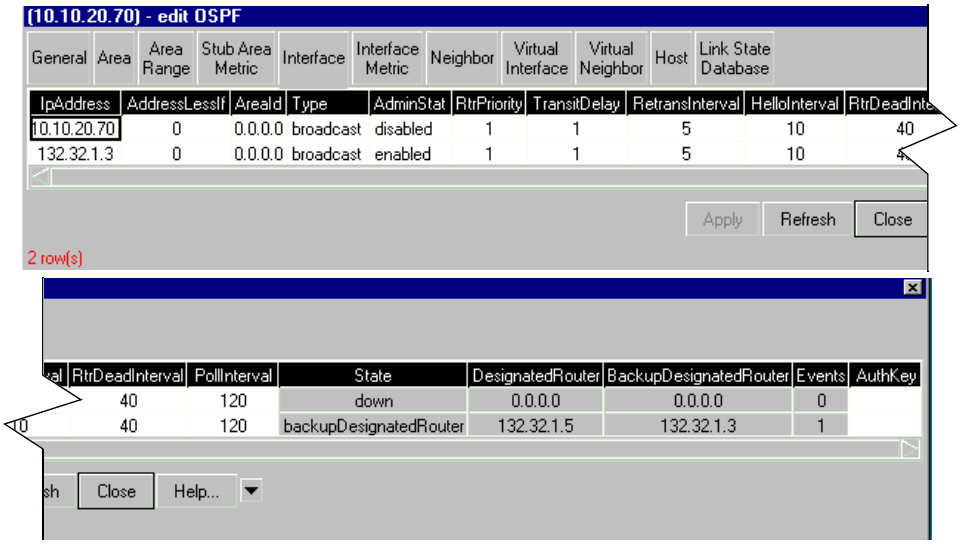

Figure 11-8. OSPF Interface Window

2. **Click on the Area tab.**

   The OSPF Area window opens (Figure 11-9).

   Notice that the backbone ID is always displayed as 0.0.0.0.

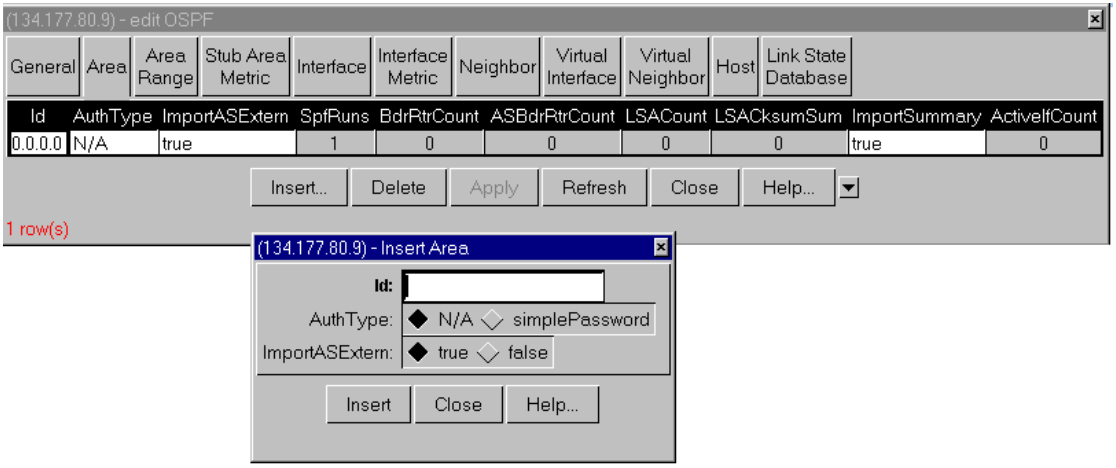

Figure 11-9. OSPF Area Tab and Insert Area Window

3. **Click on the Link State Database tab.**

   The Link State Database window opens, as shown in Figure 11-10. The Link State Database shows the advertisements in the area.

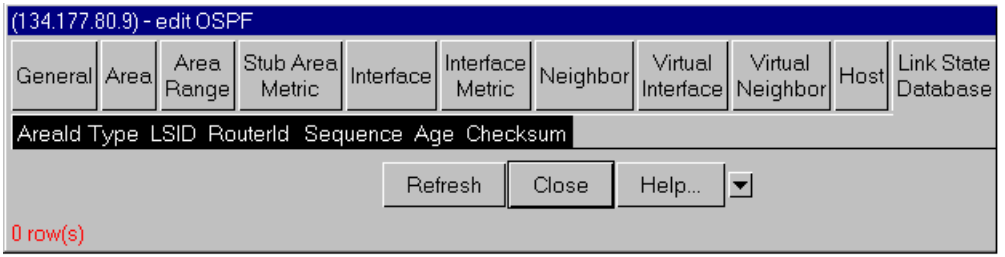

Figure 11-10. OSPF Link State Database Window

4. **Click on the Neighbor tab.**

   The Neighbor window opens (Figure 11-11), displaying the IP address and other information about the neighbor 140.40.1.4. Conversely, the neighbor to 140.40.1.4 would reciprocally display the same type of information for its neighbor.

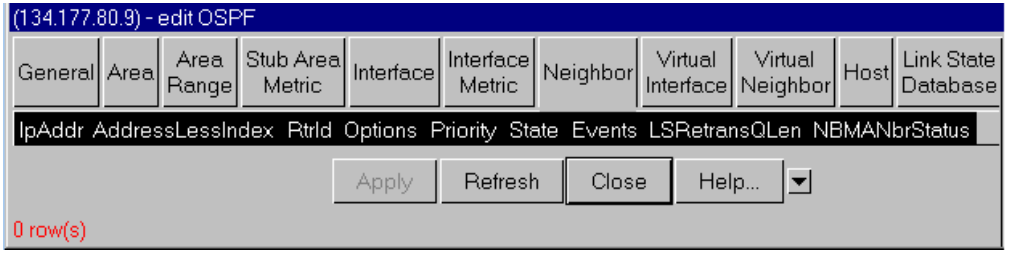

Figure 11-11. OSPF Neighbor Window

5. **Click on the graphing icon from the graphical representation of the chassis, and select OSPF.**

   The OSPF Statistics window opens, as illustrated in Figure 11-12. This window displays information about packet activity and errors.

   For a description of fields in the OSPF Statistics window, refer to page 11-35.

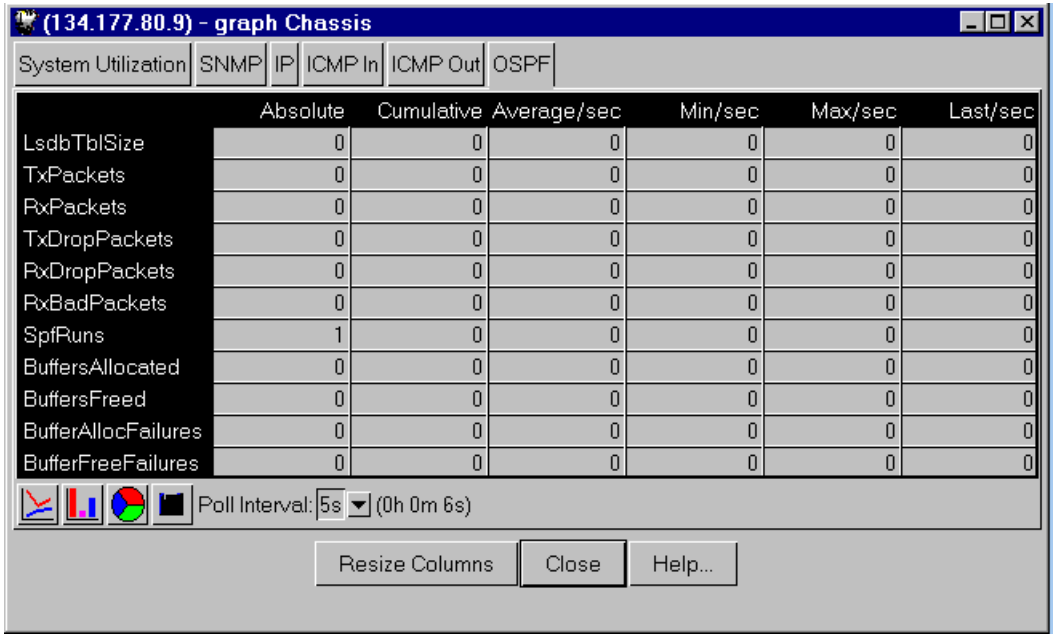

Figure 11-12. OSPF Statistics Window

Example 2: OSPF Among Multiple Networks

In this example, you will first configure OSPF on one interface on switch 1, on two interfaces on switch 2, and on one interface on switch 3. All switches are contained in one area but operate in two networks, as illustrated in Figure 11-13. You can review and verify the relationships of the configured switches, using the Interface, Neighbor, Link State Database, and IP Route windows under the Accelar Device Manager. Examples of these windows are found at the end of the procedure.

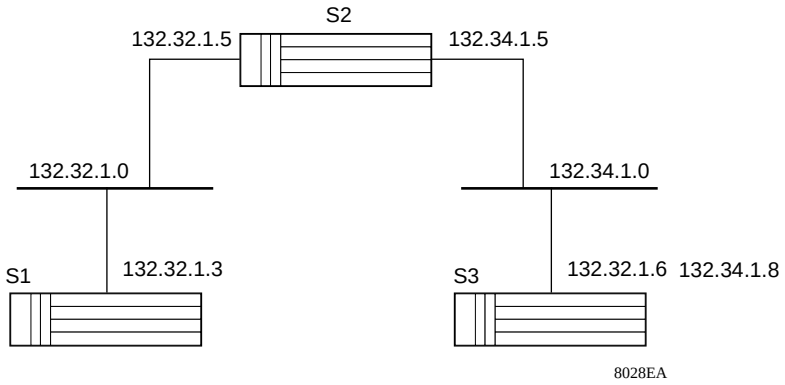

Figure 11-13. Example 2: Configuring OSPF Among Multiple Networks

Table 11-4 identifies switches used in Example 2.

Table 11-4. Switch Identifiers for Example 2

Switch Number	Switch IP Address	Interface IP Address	Mask Value	Networks in Area 0.0.0.0
Switch 1	134.177.160.101	132.32.1.3	255.255.255.0	N1 (backbone)
Switch 2	134.177.160.102	132.32.1.5 and 132.34.1.5	255.255.255.0	N1 and N2 (area border router)
Switch 3	134.177.160.97	132.34.1.8	255.255.255.0	N2

To configure OSPF on two switches in one network and on two switches in another network:

1. **Configure OSPF on one interface on switch 1, on two interfaces on switch 2, and on one interface on switch 3, using the following procedures from Example 1 and the identifiers in Table 11-4:**
   - Enable OSPF on each switch (page 11-4).
   - Insert the IP address for each switch (page 11-5).
   - Enable OSPF for an interface on each switch (page 11-6).
2. **Select a second interface on switch 2 for configuring with OSPF.**

   OSPF is already enabled on switch 2.

3. **Configure OSPF on the second interface on switch 2, using the following procedures from Example 1 and the identifiers in Table 11-4:**

   - Enable OSPF for an interface (page 11-6).
   - Insert the IP address (page 11-5).

Note: This second interface will be configured with OSPF to enable routing and establish an IP address related to a second network on a second interface.

All switches should now be configured for OSPF and should be exchanging hello packets.

You can now review the relationships among the three switches in the OSPF configuration.

To review and verify the relationships among the three switches:

1. **From switch 1, choose Routing->OSPF->Interface from the Accelar Device Manager menu bar.**

   The window displayed in Figure 11-14 opens, showing that switch 1 is the backup designated router.

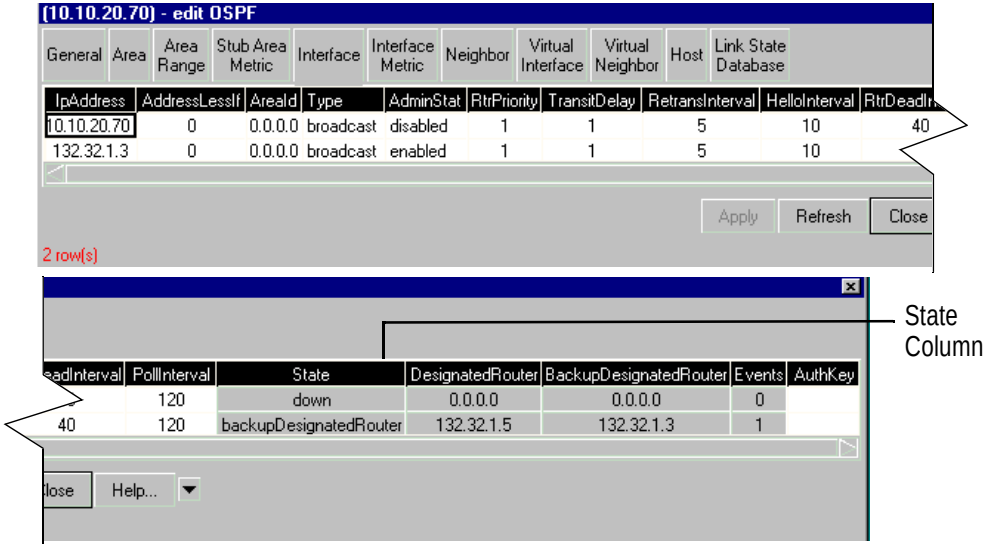

Figure 11-14. Switch 1 Interface Window

2. **From switch 2, choose Routing->OSPF->Interface from the Accelar Device Manager menu bar.**

   The window displayed in Figure 11-15 opens, showing that switch 2 has two interfaces and belongs to two separate networks. For the IP address 132.32.1.5, the switch is the designated router as indicated in the State column. In Figure 11-14, switch 1 is the backup designated router for the 132.32.1.0 network. However, notice in Figure 11-15 that switch 2 now has the role of backup designated router on a second network, 132.34.1.5.

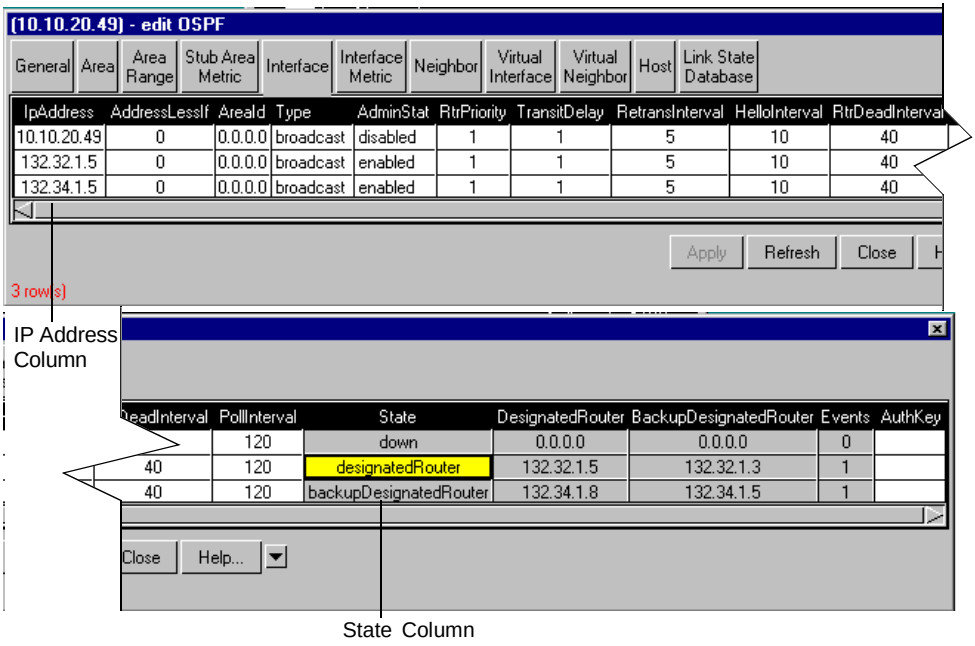

Figure 11-15. Switch 2 Interface Window

3. **From switch 1, choose Routing->OSPF->Neighbor from the Accelar Device Manager menu bar.**

   The Switch 1 Neighbor window is displayed, as shown in Figure 11-16, showing the IP address for the neighbor, switch 2. Likewise, the Switch 3 Neighbor window also will show the IP address for the neighbor, switch 2.

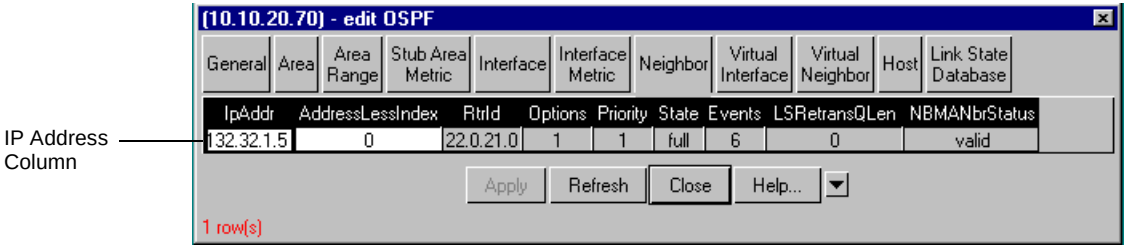

Figure 11-16. Switch 1 Neighbor Window

4. **From switch 2, choose Routing->OSPF->Neighbor from the Accelar Device Manager menu bar.**

   The Switch 2 Neighbor window displayed in Figure 11-17 opens, showing that both switch 1 and switch 3 are neighbors and that these switches are on separate networks.

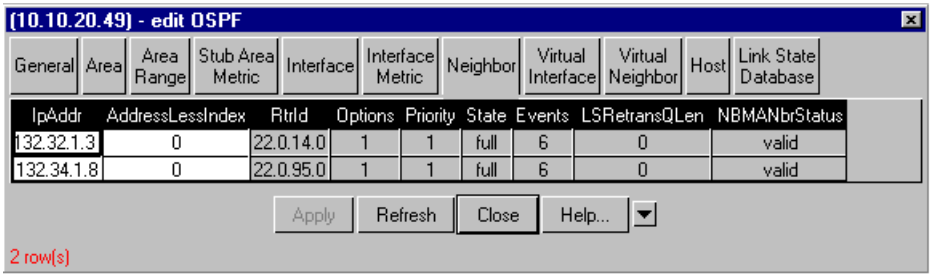

Figure 11-17. Switch 2 Neighbor Window

5. **From switch 1, choose Routing->OSPF->Link State Database from the Accelar Device Manager menu bar.**

   The Link State Database window displayed in Figure 11-18 opens, showing the OSPF roles the switches are assigned. If you check the Link State Database for all three switches, notice that they are identical.

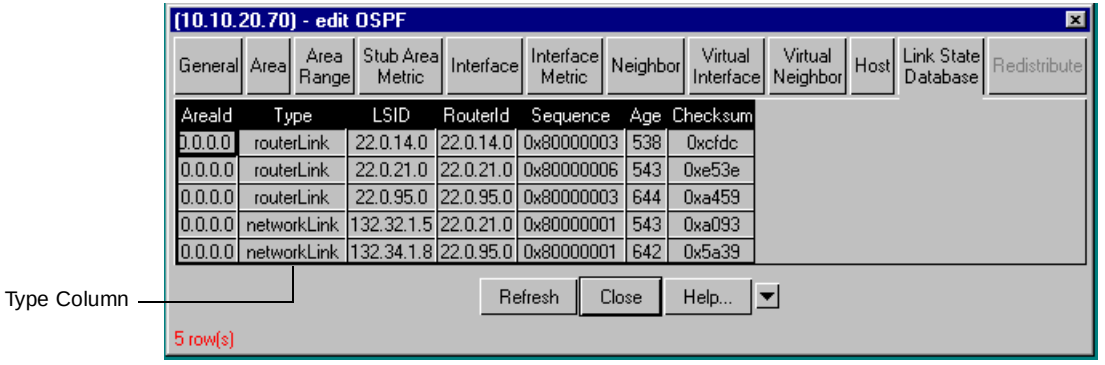

Figure 11-18. Switch 1 Link State Database Window

To see how the configuration of the three switches appears from a routing perspective:

1. **From switch 1, choose Routing->IP Route from the Accelar Device Manager menu bar.**

   The Switch 1 IP Route window opens, as shown in Figure 11-19. Notice in the Type column that anything on the 132.32.1.0 network is directly connected (switch 2), but anything on the 132.34.1.0 network is indirectly connected (switch 3) and needs to be forwarded by switch 2 to the other network.

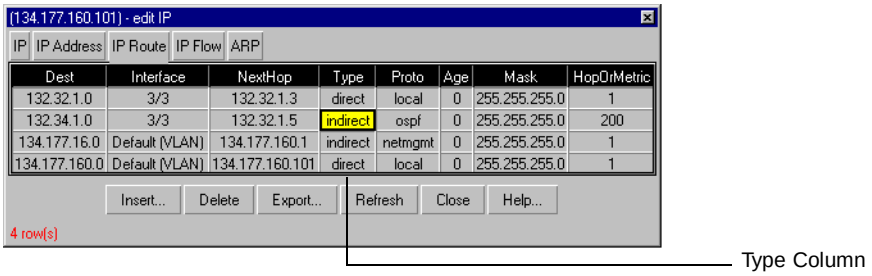

Figure 11-19. Switch 1 (vx3) IP Route Window

2. **From switch 3, choose Routing->IP->IP Route from the Accelar Device Manager menu bar.**

The Switch 3 IP Route window displayed in Figure 11-20 opens, showing the converse of the example illustrated in Figure 11-19. The IP route window for switch 3 shows that it is indirectly connected to 132.32.1.0 (switch 1), and that the next hop is 132.34.1.5 (switch 2). The relationship follows from switch 3 being directly connected to 132.34.1.0 (switch 2).

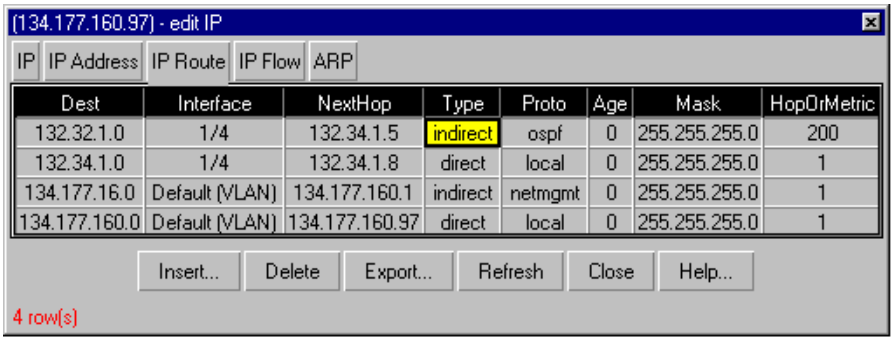

Figure 11-20. Switch 3 IP Route Window

Example 3: OSPF Among Multiple Areas

In this example, you will configure OSPF on two networks in separate areas.

- You will first configure OSPF on one network.
  - — Configure one interface on switch 1.
  - — Configure one interface on switch 2.
- Next you will configure OSPF on two additional interfaces in a second network.
  - — Configure a second interface on switch 2.
  - — Configure one interface on switch 3.

The second switch will become the area boundary router for both networks. The configuration described above is illustrated in Figure 11-21.

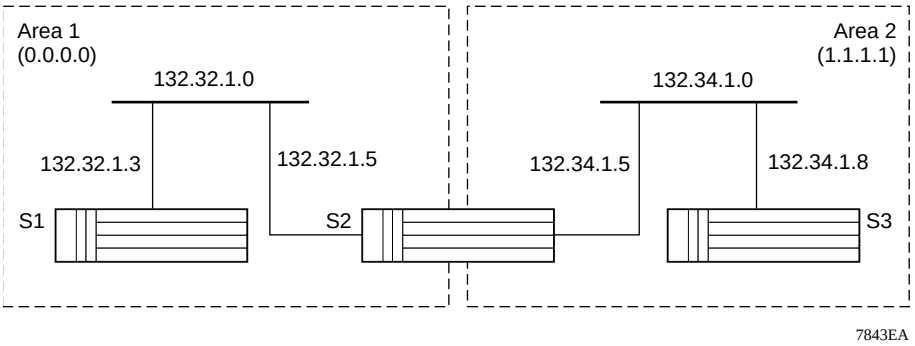

Figure 11-21. Example 3: Configuring OSPF Among Multiple Areas

Table 11-5 identifies switches used in Example 3.

Table 11-5. Switch Identifiers for Example 3

Switch Number	Switch Designation	Interface IP Address	Mask Value	Area
Switch 1	134.177.160.101	132.32.1.3	255.255.255.0	1
Switch 2	134.177.160.102	132.32.1.5 and 132.34.1.5	255.255.255.0 and 255.255.0.0	1 and 2
Switch 3	134.177.160.97	132.34.1.8	255.255.0.0	2

To configure OSPF on one interface on each of two switches:

1. **Configure OSPF on one port on switch 1, using the following procedures from Example 1 and the identifiers in Table 11-5:**
   - Enable OSPF on each switch (page 11-4).
   - Enable OSPF for a port on each switch (page 11-6).
   - Insert the IP address for each switch (page 11-5).

2. **Following the directions in step 1, configure OSPF on one port on switch 2.**

Note: Both routable ports belong to the same network. Therefore, by default, both ports are in the same area.

To create a new area:

1. **After configuring the two routers, select another port on switch 2.**

   Notice that OSPF is already enabled on the switch. However, follow the steps for enabling routing and establishing a second network on a second port on a switch.

2. **From switch 2, choose Routing->OSPF->Area from the Accelar Device Manager menu bar.**

3. **Click on Insert at the bottom of the Switch 2 Area window.**

   The Insert Area dialog box shown in Figure 11-22 opens.

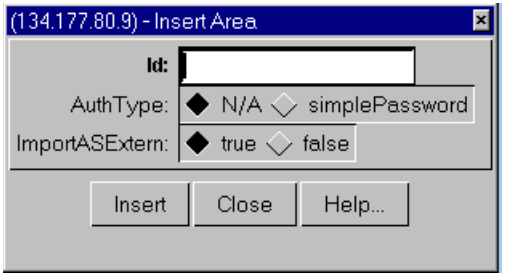

Figure 11-22. Switch 2 Insert Area Dialog Box

4. **In the Id field, type an area identification number and click on Insert at the bottom of the dialog box.**

   The identification number must be in dotted-decimal format. Any number is acceptable except 0.0.0.0, which is reserved for the backbone.

You have now created a new area, shown in Figure 11-23. Notice that the first field under the Id column lists the new area.

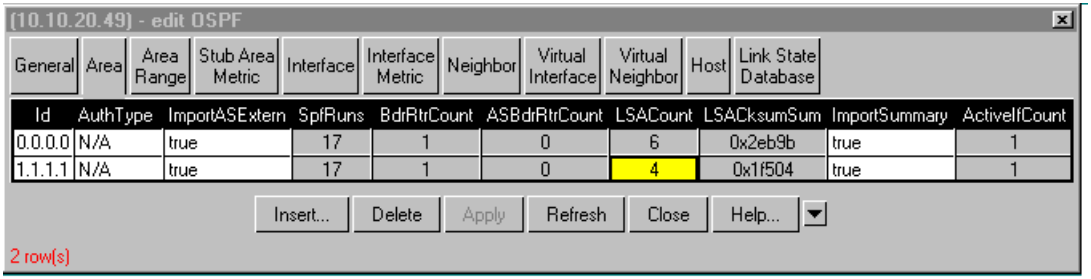

Figure 11-23. Switch 2 Area Window

To specify the range for the new area:

1. **Open the Area Range window.**
2. **Click on Insert at the bottom of the window.**

   The Insert Area Range dialog box shown in Figure 11-24 opens.

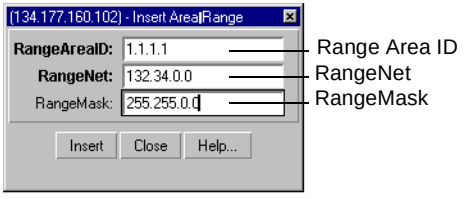

Figure 11-24. Switch 2 Insert Area Range Dialog Box

3. **In the RangeNet field, type the range of IP addresses to be associated with the designated area.**

   Refer to Table 11-5 for values in this example. In Figure 11-24, the RangeNet value is 132.34.0.0.

4. **In the RangeMask field, type the range mask value.**

   The RangeMask value summarizes the last 16 bits of the net mask. Refer to Table 11-5 for values in this example.

5. **Click on Insert at the bottom of the dialog box.**

   The values are inserted (Figure 11-25).

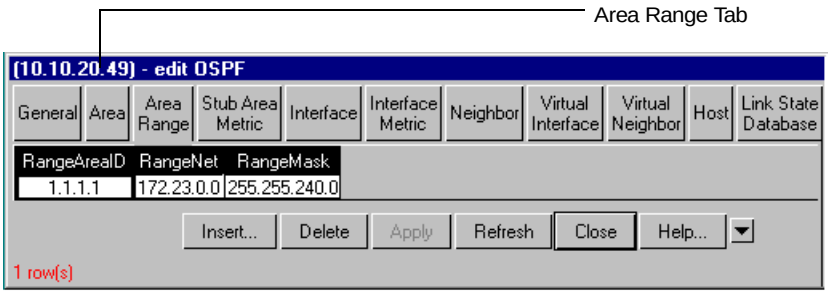

Figure 11-25. Switch 2 Area Range Window

Note that the value in the RangeArea ID field should match the ID number for the area created in step 4 on page 11-19.

To configure OSPF on one port on switch 3 and verify that all three switches are neighbors:

1. **Configure OSPF on one port on switch 3, using the following procedures from Example 1 and the identifiers in Table 11-5:**

   - Enable OSPF (page 11-4).
   - Enable OSPF for a port (page 11-6).
   - Insert the IP address (page 11-5).

   All three switches should now be configured for OSPF and should be exchanging hello packets.

   Now you can review the relationships among the three switches in the OSPF configuration.

2. **From switch 2, choose Routing->OSPF->General from the Accelar Device Manager menu bar.**

   The General window opens, as displayed in Figure 11-26, showing that switch 2 is the area border router. This status of switch 2 is confirmed by the word “true” in the AreaBdrRtrStatus field.

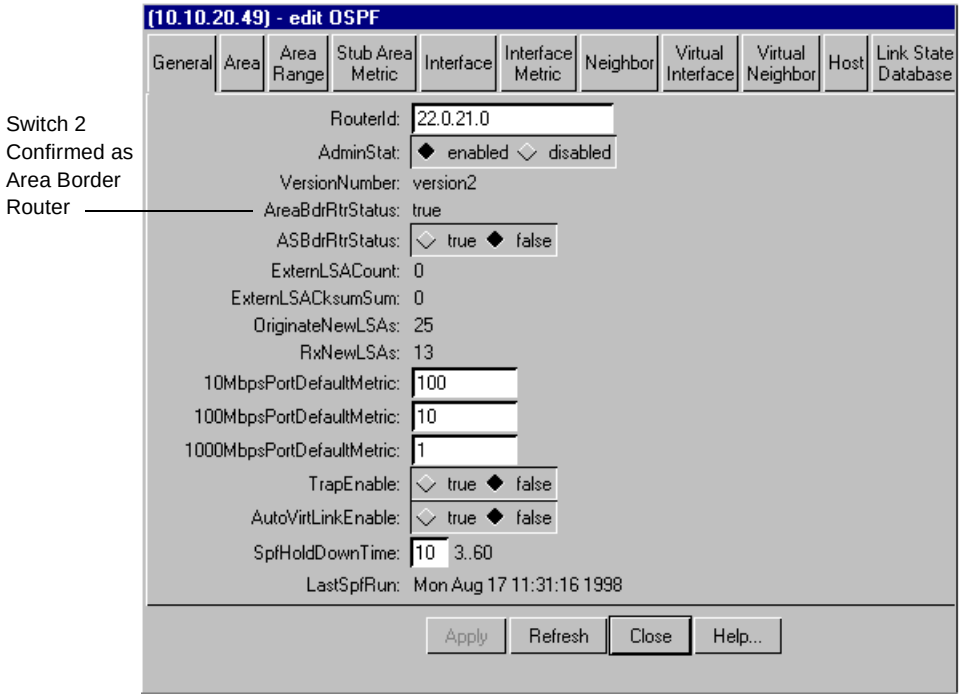

Figure 11-26. Switch 2 General Window

3. Click on the Neighbor tab.

The Neighbor window opens, as displayed in Figure 11-27, showing that switch 1 (network 132.32.1.0) and switch 3 (network 123.34.1.0) are neighbors to switch 2.

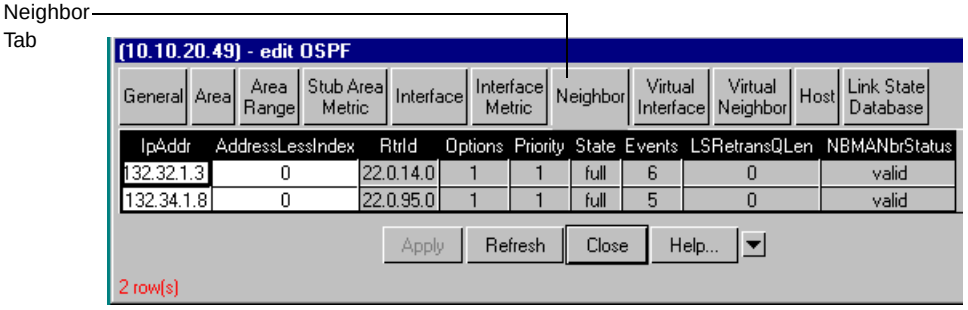

Figure 11-27. Switch 2 Neighbor Window

To compare the link state of the three switches:

1. **From switch 2, choose Routing->OSPF->Link State Database from the Accelar Device Manager menu bar.**

   The Link State Database window opens for switch 2, as displayed in Figure 11-28. The window shows information about the link state of each router in an area. In the AreaId column, area 1 is designated 0.0.0.0 and area 2 is 1.1.1.1.

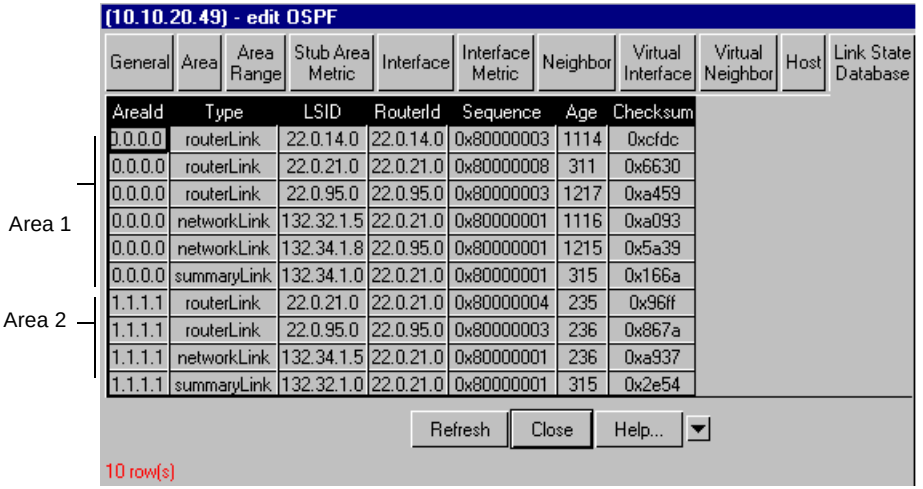

Figure 11-28. Switch 2 Link State Database Window

2. **From the Accelar Device Manager menu bar, choose Routing->IP->IP Route.**

   The Edit IP window opens (Figure 11-29), showing information about switch 1 IP routing. The Type column indicates that switch 1 is directly attached to the destination network 132.32.1.0 and indirectly attached to the destination network 132.34.0.0.

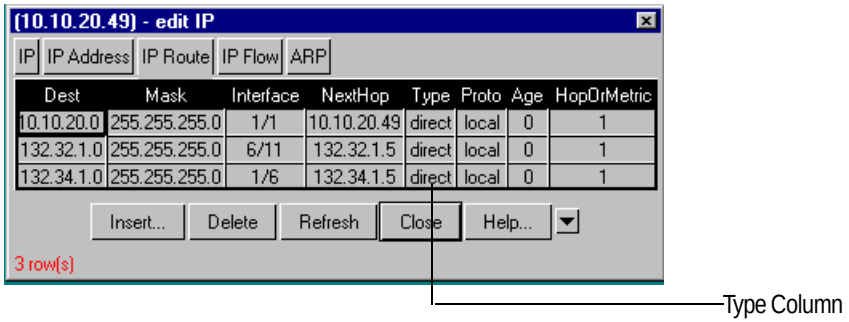

Figure 11-29. Switch 2 Edit IP Window

Compare Figure 11-29 for switch 1 with Figure 11-30 for switch 3. The Type column in Figure 11-30 contains information that indicates a direct attachment of switch 3 to area 2 (132.34.1.0) and an indirect attachment to area 1 (132.32.1.0), through the area border router.

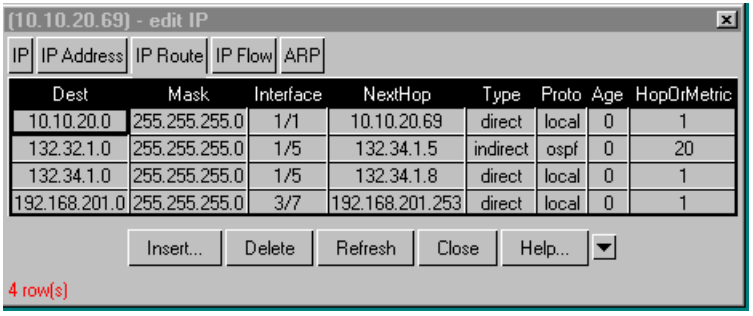

Figure 11-30. Switch 3 Edit IP Window

Next you can compare the link state database windows of switch 1 and switch 3.

3. **From switch 1, choose Routing->OSPF->Link State Database from the Accelar Device Manager menu bar.**

   The Link State Database window opens for switch 1 (Figure 11-31), showing information about the network link. The AreaId column indicates that switch 1 is the network link (0.0.0.0).

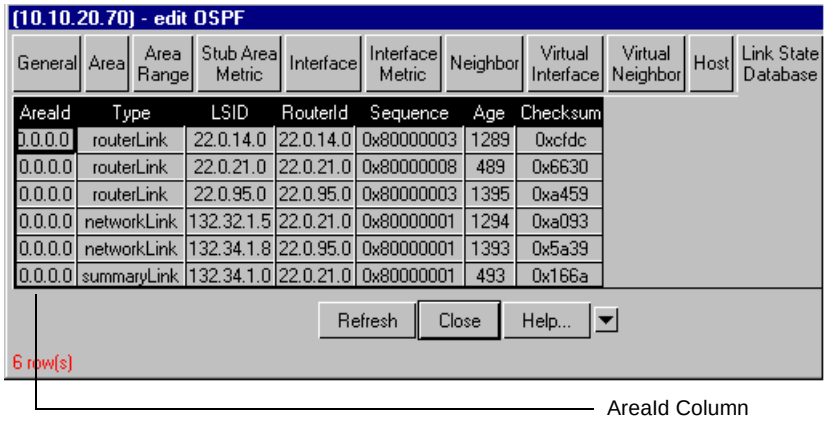

Figure 11-31. Switch 1 Link State Database Window

Figure 11-32 shows the switch 2 Link State Database window. Notice that the Type column indicates that switch 2 is the network link.

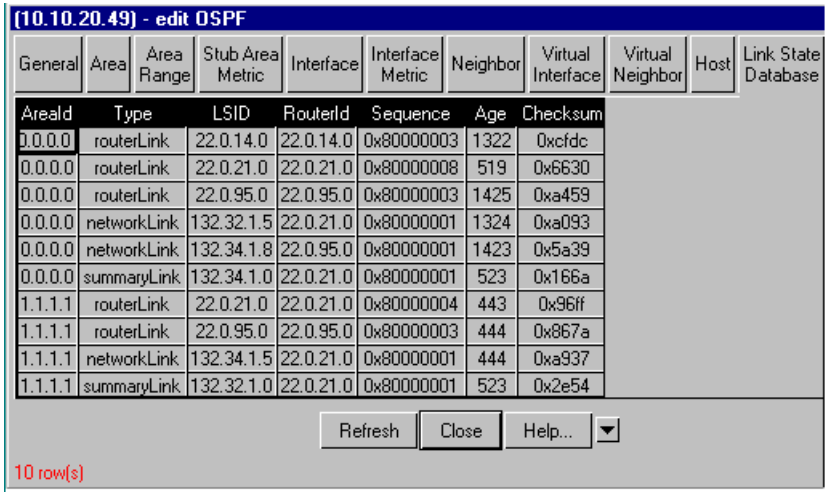

Figure 11-32. Switch 2 Link State Database Window

Figure 11-33 shows the switch 3 Link State Database window. Compare the information in the Type column to that shown in the previous figures.

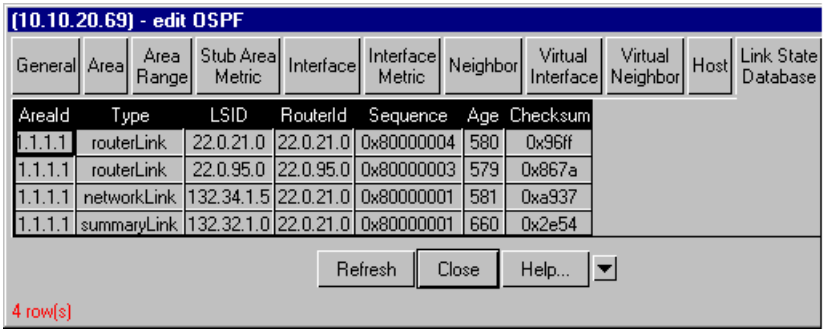

Figure 11-33. Switch 3 Link State Database Window

Creating a Virtual Link

When using OSPF, routing switches, which are ABRs, need to be connected directly to the backbone. If they are not directly connected, they need to have a virtual link. In the Accelar routing switches, you can specify that virtual links be automatically created, or you can manually configure a virtual link.

When automatic virtual linking is enabled, it acts like “insurance.” A virtual link will be created for vital traffic paths in your OSPF configuration if something goes amiss, such as when an interface cable providing connection to the backbone (either directly or indirectly) becomes disconnected from the switch. Specifying automatic virtual linking ensures that a link will be created via another routing switch. When you specify automatic virtual linking, it is always ready to create a virtual link. If automatic virtual linking uses more resources than you want to expend, creating a manual virtual link may be the better solution. This approach lets you conserve resources while having specific control of where virtual links are placed in your OSPF configuration.

Automatic Virtual Link

To specify that virtual links be automatically created:

1. **Choose Routing->OSPF->General to open the window shown in Figure 11-34.**

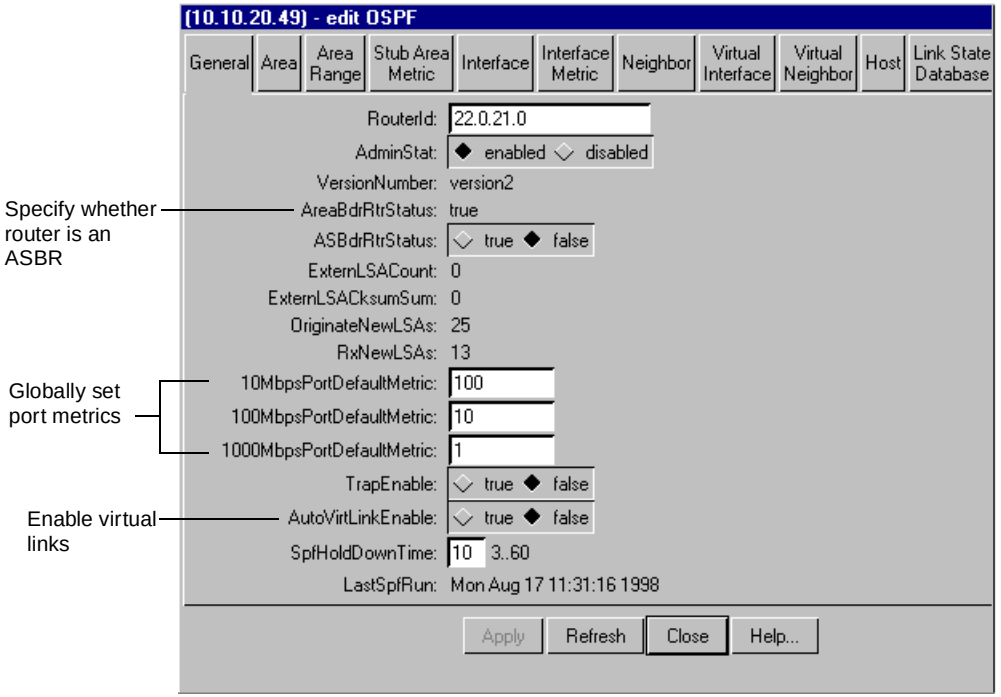

Figure 11-34. OSPF General Window

2. **Select true in the AutoVirtLinkEnable field.**

   By default, this feature is set to false, and virtual links are not automatically created.

3. **Click on Apply at the bottom of the window.**

Manual Virtual Link

The example that follows illustrates how to configure a virtual link between the ABR in area 2.2.2.2 and the ABR in area 0.0.0.0, as shown in Figure 11-35.

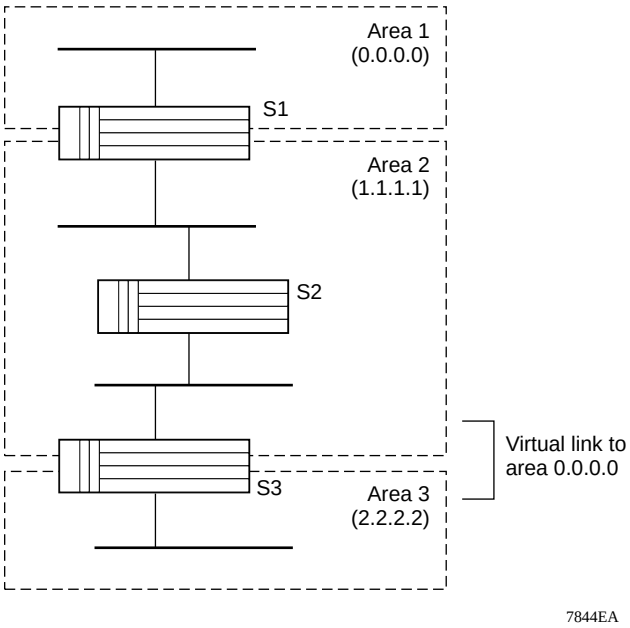

Figure 11-35. Area 3 Manual Virtual Link to Area 1 Via Area 2

The virtual link from area 2.2.2.2 is a necessary link, as shown in Figure 11-36 and Figure 11-37. Switch S3 needs to go through switch S2 to have a connection with switch S1; switch S1 has the connection to the backbone. The S1 interface window shows that the S1 switch recognizes the interface for area 1.1.1.1 but does not show any recognition of area 2.2.2.2. The S3 interface window shows that the S3 switch recognizes the interface for the 1.1.1.1 area and also shows that the interface for the 0.0.0.0 area is disabled.

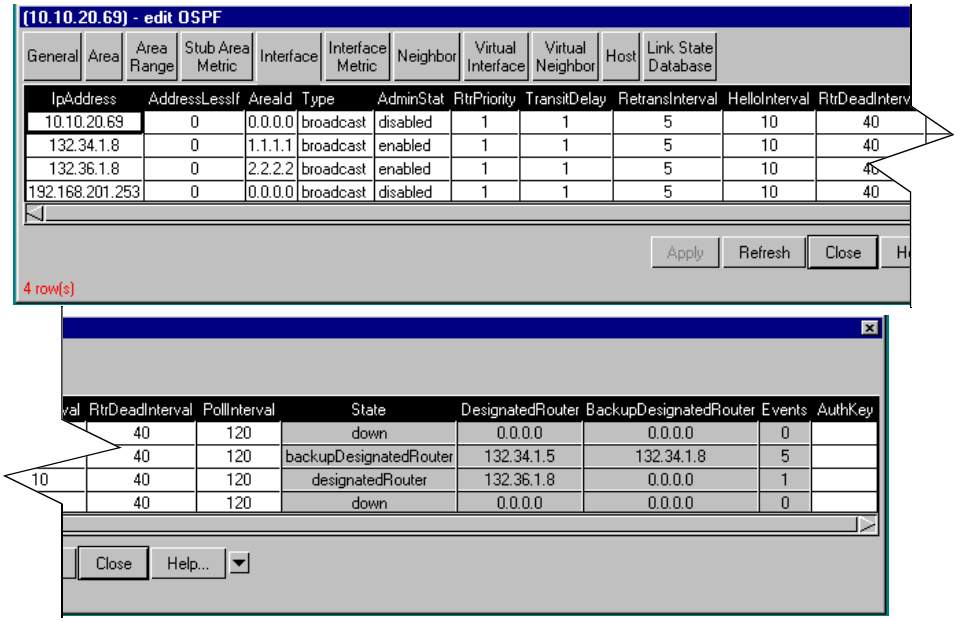

Figure 11-36. S3 Interface Window

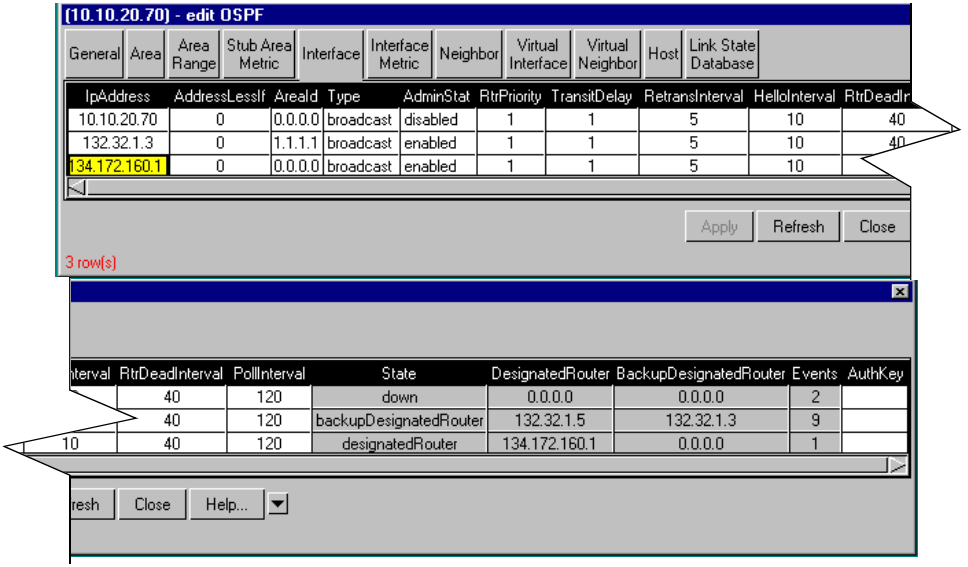

Figure 11-37. S1 Interface Window

This example assumes an OSPF configuration among three ABRs.

To manually configure a virtual link:

1. **Choose Routing->OSPF->Virtual Interface.**
2. **Click on Insert at the bottom of the window. The Insert Virtual Interface window is displayed, as shown in Figure 11-38.**
3. **In the Insert Virtual Interface window, specify the area ID of the transit area.**

   The transit area is the common area between two ABRs.

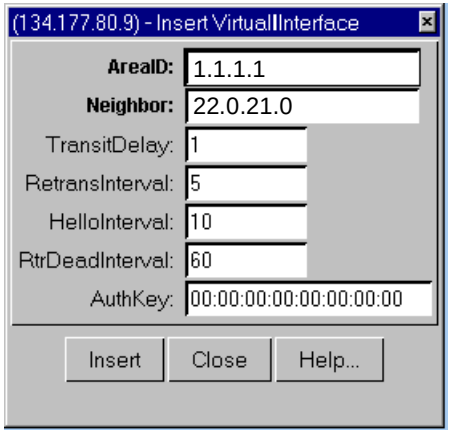

Figure 11-38. Insert Virtual Interface Window

4. **Specify the neighbor ID.**

   The neighbor ID is the IP router ID of the ABR that the other ABR needs to go through to get to the backbone.
5. **Click on Insert at the bottom of the window.**
6. **To verify that the virtual link is active, refresh the Virtual Interface window (Figure 11-39) and check the state column.**

   If the state displays “point to point,” the virtual link is active, as shown in Figure 11-39 and Figure 11-40. If the state column displays “down,” the virtual link is invalid.

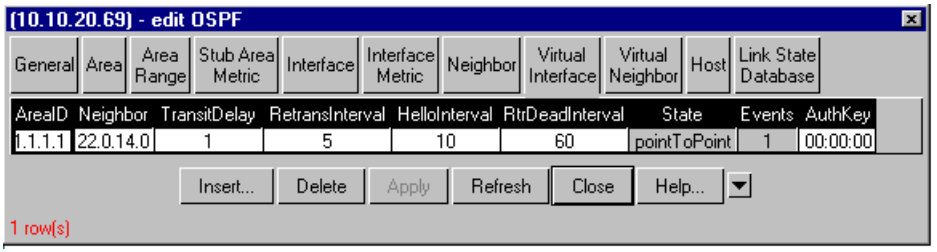

Figure 11-39. S3 Virtual Interface Window

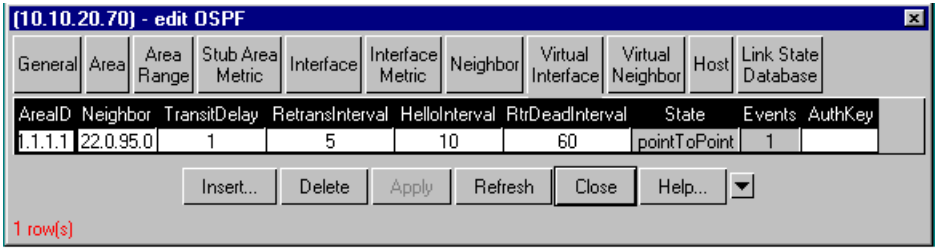

Figure 11-40. S1 Virtual Interface Window

In addition, you can check the virtual neighbor window (Figure 11-41) under the S3 ABR. The window reflects that the area 2.2.2.2 ABR now has a virtual neighbor going through area 1.1.1.1.

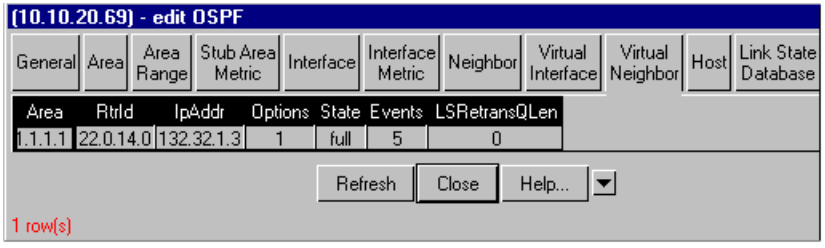

Figure 11-41. S3 Virtual Neighbor Window

Specifying ASBRs

ASBRs advertise non-OSPF routes into OSPF domains so that they can be passed along throughout the OSPF routing domain. A router can function as an ASBR if one or more of its interfaces is connected to a non-OSPF network (for example, RIP, BGP, or EGP).

To conserve resources, you may want to limit the number of ASBRs in your network or to specifically control which routers perform as ASBRs to control traffic flow.

To specify whether or not a router should be an ASBR:

1. **Choose Routing->OSPF->General.**
2. **From the ASBdrRtrStatus field, select true to designate the router as an ASBR or false to remove ASBR status from the router.**
3. **Click on Apply at the bottom of the window.**

Creating a Stub Area

A stub area does not receive advertisements for external routes, which reduces the size of the link state database. A stub area has only one area border router. Any packets destined outside the area are simply routed to that area border exit point, examined by the area border router, and forwarded to a destination.

To create a stub area:

1. **Choose Routing->OSPF->Area.**

Under the ImportASExtern field (Figure 11-42), select the area you want to change to a stub area, select false, and click on Apply.

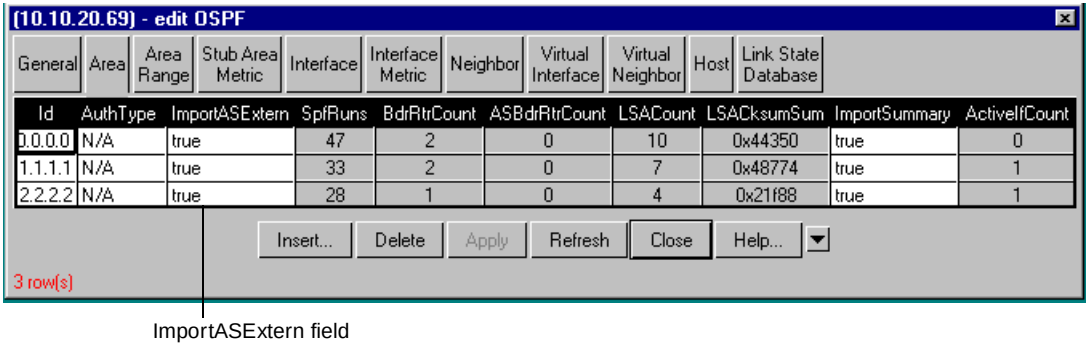

Figure 11-42. Area Window

Configuring Metric Speed

You can configure the metric speed globally or for specific ports and interfaces on your network. In addition, you can control redistribution options between non-OSPF interfaces and OSPF interfaces.

Global Default Metric Speed

To change the default metric speed on specific port types:

1. **Choose Routing->OSPF->General to open the window shown in Figure 11-34 on page 11-27.**
2. **Change the metric value in one or all of the following fields:**
   - 10MbpsPortDefaultMetric (default = 100)
   - 100MbpsPortDefaultMetric (default = 10)
   - 1000MbpsPortDefaultMetric (default = 1)

3. **Click on Apply at the bottom of the window.**

The default port metric speed will be changed on all port types for which you have specified a new metric speed.

Port-Specific Metric Speed

For finer control over metric speed, you can specify the metric speed when you enable OSPF on a port or when you edit a port.

To specify the metric speed on a specific port instead of a port type:

1. **From the main menu bar, choose Routing->OSPF->Interface Metric *or* choose Edit->Port->OSPF.**
2. **Specify a new metric speed in the metric field.**
3. **Click on Apply at the bottom of the window.**

Note: When you enable a port for OSPF routing, the default metric in the port window is “0.” A value of “0” (zero) means that the port will use the default metrics for port types that are specified on the OSPF general window.

Window and Field Reference

Table 11-6 lists the OSPF windows and describes their fields.

Table 11-6. OSPF Window and Field Descriptions

Window	Field	Description
General		Contains general information about the version of OSPF running and the configuration of the router.
	RouterID	The Router ID, which in OSPF has the same format as an IP Address but identifies the router independent of other routers in the OSPF domain.
	AdminStat	The administrative status of OSPF in the router. The value “enabled” denotes that the OSPF process is active on at least one interface; “disabled” disables it on all interfaces.
	VersionNumber	Current version number of OSPF.
	AreaBdrRtrStatus	A flag to note if this router is an area border router.
	ASBdrRtrStatus	Enabled as true indicates the router is configured as an Autonomous System border router.
	ExternalSACount	The number of external (LS type 5) link-state advertisements in the link-state database.
	ExternalSACksumSum	The 32-bit unsigned sum of the LS checksums of the external link-state advertisements contained in the link-state database. This sum can be used to determine if there has been a change in a router’s link state database and to compare the link-state databases of two routers.
	OriginateNewLSAs	The number of new link-state advertisements that have been originated. This number is incremented each time the router originates a new LSA.
	RxNewLSAs	The number of link-state advertisements received that are determined to be new instantiations. This number does not include newer instantiations of self-originated link-state advertisements.
	10MbpsPortDefaultMetric	Indicates the cost associated with 10 Mb/s interface (port).
	100MbpsPortDefaultMetric	Indicates the cost associated with 100 Mb/s interface (port).
	1000MbpsPortDefaultMetric	Indicates the cost associated with 1000 Mb/s interface (port).

Table 11-6. OSPF Window and Field Descriptions (continued)

Window	Field	Description
General	TrapEnable	Indicates whether or not traps relating to the Spanning Tree Protocol should be sent for this STG.
	AutoVirtLinkEnable	Enables or disables automatic creation of virtual links.
	rcIpConfOspfSpfHoldDownTime	Allows the user to change the OSPF hold down timer value (3 to 60 seconds).
	LastSpfRun	Used to indicate the time (SysUpTime) since the last SPF calculated by OSPF.
Area		Contains information describing the configured parameters and cumulative statistics of the router's attached areas.
	Id	A 32-bit integer uniquely identifying an area. Area ID 0.0.0.0 is used for the OSPF backbone.
	AuthType	The authentication type specified for an area. Additional authentication types may be assigned locally on a per area basis.
	ImportASExtern	The area's support for importing AS external link-state advertisements.
	SpfRuns	Used to indicate the number of SPF calculations performed by OSPF.
	BdrRtrCount	The total number of area border routers reachable within this area. The value, initially zero, is calculated in each SPF Pass.
	ASBdrRtCount	The total number of Autonomous System border routers reachable within this area. The value, initially zero, is calculated in each SPF Pass.
	LSACount	The total number of link-state advertisements in this area's link-state database, excluding AS External LSAs.
	LSACksumSum	The 32-bit unsigned sum of the link-state advertisements. This sum excludes external (LS type 5) link-state advertisements. The sum can be used to determine if there has been a change in a router's link state database and to compare the link-state database of two routers.
	ImportSummary	The area's support for importing Summary advertisements into a stub area. This field should be used only if ospfImportASExtern is set to FALSE.
	ActiveCount	

Table 11-6. OSPF Window and Field Descriptions (continued)

Window	Field	Description
Area Range		Contains a range of IP addresses specified by an IP address/IP network mask pair. For example, class B address range of x.x.x.x with a network mask of 255.255.0.0 includes all IP addresses from x.x.0.0 to x.x.255.255.
	RangeAreaID	The area the address range is located in.
	RangeNet	The IP address of the net or subnet indicated by the range.
	RangeMask	The subnet that pertains to the net or subnet.
Stub Area Metric		Contains the set of metrics that will be advertised by a default area border router into a stub area.
	AreaID	The 32-bit identifier for the stub area. On creation, it can be derived from the instance.
	TOS	The type of service associated with the metric. On creation, it can be derived from the instance.
	Metric	The metric value applied at the indicated type of service. By default, it equals the lowest metric value at the type of service among the interfaces to other areas.
	Status	This variable displays the status of the entry. Setting it to ‘invalid’ has the effect of rendering it inoperative.
Interface		Displays information describing the interfaces from the viewpoint of OSPF.
	IpAddress	The Internet Protocol address of the device used to represent a point of attachment in a TCP/IP internetwork.
	AddressLessIf	Used for the purpose of easing the instancing of addressed and addressless interfaces. This variable takes the value 0 on interfaces with IP addresses and the corresponding value of ifIndex for interfaces having no IP address.
	AreaId	A 32-bit integer uniquely identifying an area. Area ID 0.0.0.0 is used for the OSPF backbone.
	Type	Designates the type of routing designated: specifically routed, all paths explored, or spanning tree explored.
	AdminStat	The administrative status of OSPF in the router. The value “enabled” denotes that the OSPF process is active on at least one interface; “disabled” disables it on all interfaces.

Table 11-6. OSPF Window and Field Descriptions (continued)

Window	Field	Description
Interface	RtrPriority	The priority of this interface. Used in multiaccess networks, this field is used in the designated router election algorithm. The value 0 signifies that the router is not eligible to become the designated router on this particular network. In the event of a tie in this value, routers will use their router ID as a tie breaker.
	TransitDelay	The estimated number of seconds it takes to transmit a link-state update packet over this interface.
	RetransInterval	The number of seconds between link-state advertisement retransmissions. This value is also used for retransmission of database descriptions and link-state request packets.
	HelloInterval	The length of time, in seconds, between the hello packets that the router sends on the interface. This value must be the same for all routers attached to a common network.
	RtrDeadInterval	The number of seconds that a router's hello packets have not been seen before its neighbors declare the router down. This value should be some multiple of the hello interval. It must be the same for the virtual neighbor.
	PollInterval	The larger time interval, in seconds, between the Hello packets sent to an inactive nonbroadcast multiaccess neighbor.
	State	The OSPF interface state.
	DesignatedRouter	The IP address of the designated router.
	BackupDesignatedRouter	The IP address of the designated backup router.
	Events	The number of state changes or error events that have occurred through all interfaces.
	AuthKey	The area's authorization type as a simple password, and the key, which is shorter than 8 octets. The agent will left adjust and zero fill to 8 octets.
Interface Metric		Indicates the metrics associated with the peer layer interface.
	IpAddress	The Internet Protocol address of the device used to represent a point of attachment in a TCP/IP internetwork.

Table 11-6. OSPF Window and Field Descriptions (continued)

Window	Field	Description
	AddressLessIf	For the purpose of easing the instancing of addressed and addressless interfaces. This variable takes the value 0 on interfaces with IP addresses and the corresponding value of ifIndex for interfaces having no IP address.
	TOS	Type of service is a mapping to the IP type of service flags as defined in the IP forwarding table MIB.
	Metric	The metric advertised to other areas. The value indicates the distance from the OSPF router to any network in the range.
Neighbor		Displays information describing the interfaces from the viewpoint of OSPF.
	IpAddr	The device IP address.
	AddressLessIndex	On an interface having an IP address, zero. On addressless interfaces, the corresponding value of ifIndex in the Internet standard MIB. On row creation, this value can be derived from the instance.
	RtrId	The router ID of the neighboring router, which in OSPF has the same format as an IP address but identifies the router independent of its IP address.
	Options	A bit mask corresponding to the neighbor's options field.
	Priority	Assignment of preferential treatment to place the transmitted packets in queues and possible selection of the priority field in the data link header when the packet is forwarded.
	State	The OSPF Interface state.
	Events	The number of state changes or error events that have occurred between the OSPF router and the neighbor router.
	LSRetransQLen	The number of elapsed seconds between advertising retransmissions of the same packet to a neighbor.
	NBMANbrStatus	Status of the nonbroadcast multiaccess network.
Virtual Interface		Displays information describing a virtual interface from the viewpoint of OSPF.
	AreaID	An OSPF area identifier.
	Neighbor	Virtual neighbor interfaced on the same network.

Table 11-6. OSPF Window and Field Descriptions (continued)

Window	Field	Description
	TransitDelay	The estimated number of seconds it takes to transmit a link-state update packet over this interface.
	RetransInterval	The number of seconds between link-state advertisement retransmissions. This value is also used for retransmission of database descriptions and link-state request packets.
	HelloInterval	The length of time, in seconds, between the hello packets that the router sends on the interface. This value must be the same for all routers attached to a common network.
	RtrDeadInterval	The number of seconds that a router's hello packets have not been seen before its neighbors declare the router down. This value should be some multiple of the hello interval. It must be the same for the virtual neighbor.
	State	The OSPF interface state.
	Events	The number of state changes or error events that have occurred through the virtual interface.
	AuthKey	The area's authorization type as a simple password and the key, which is shorter than 8 octets. The agent will left adjust and zero fill to 8 octets.
Virtual Neighbor		Displays statistics for a virtual neighbor in a network.
	Area	The subnetwork in which the virtual neighbor resides.
	RtrId	A 32-bit integer (represented as a type IpAddress) uniquely identifying the neighboring router in the autonomous system.
	IpAddr	The IP address of the virtual neighboring router.
	Options	A bit mask corresponding to the neighbor's options field.
	State	The OSPF interface state.
	Events	The number of state changes or error events that have occurred between the OSPF router and the virtual neighbor router.
	LSRetransQLen	The number of elapsed seconds between advertising retransmissions of the same packet to a neighbor.

Table 11-6. OSPF Window and Field Descriptions (continued)

Window	Field	Description
Host		Indicates what hosts are directly attached to the router, and what metrics and types of service should be advertised for them.
	IpAddress	The IP address of the host used to represent a point of attachment in a TCP/IP internetwork.
	TOS	The type of service of the route being configured.
Host	Metric	The metric advertised to other areas. The value indicates the distance from the OSPF router to any network in the range.
Link State Database		Contains the link state advertisements from throughout the areas to which the device is attached.
	AreaId	A 32-bit integer uniquely identifying an area. Area ID 0.0.0.0 is used for the OSPF backbone.
	Type	The OSPF interface type. By way of a default, this field may be intuited from the corresponding value of ifType. Broadcast LANs, such as Ethernet and IEEE 802.5, take the value broadcast; X.25 and similar technologies take the value nbma; and links that are definitively point-to-point take the value pointToPoint.
	LSID	The Link State ID is an LS type-specific field containing either a router ID or an IP address. It identifies the piece of the routing domain that is being described by the advertisement.
	RouterID	A 32-bit integer uniquely identifying the router in the autonomous system.
	Sequence	The sequence number is a signed 32-bit integer that identifies old and duplicate link state advertisements.
	Age	The age in seconds of the link state advertisement.
	Checksum	This field is the checksum of the complete contents of the advertisement, excepting the age field. The age field is excepted so that an advertisement's age can be incremented without updating the checksum. The checksum used is the same that is used for ISO connectionless datagrams. It is commonly referred to as the Fletcher checksum.
Insert IP Address (pop-up)		Contains configured parameters for the router running OSPF.

Table 11-6. OSPF Window and Field Descriptions (continued)

Window	Field	Description
	IpAddress	The Internet Protocol address of the device used to represent a point of attachment in a TCP/IP internetwork.
	NetMask	The OSPF router subnet mask.
IP Route		Displays parameters and statistics of packets transmitted to routers in the network.
	Dest	The IP address of the destination router.
	InfIndex	Port position in chassis. For example, 1/1 is slot 1, port 1.
	NextHop	The IP address destination of the next packet to pass through the OSPF router.
	Type	Route type: direct or indirect.
	Proto	Routing protocol type used for the route.
	Age	The time in seconds that a packet was in transit.
	Mask	The mask value of the OSPF route advertisement.
	Hops	The number of packets passed through a router.
Insert Area (pop-up)		Contains configured parameters for areas attached to the network.
	Id	The unique identification number of an area in an autonomous system.
	AuthType	The authentication type specified for an area. Additional authentication types may be assigned locally on a per area basis.
	ImportASExtern	A value of true indicates the area is non-stub.
Insert Area Range (pop-up)		Displays information for the range of subnets attached to the network.
	RangeAreaID	Point of insertion for an OSPF area IP address.
	RangeNet	The IP address of the net or subnet indicated by the range.
	RangeMask	The subnet mask that pertains to the net or subnet.

Chapter 12
IP Policies

This chapter describes using Accelar Device Manager to configure IP policy features supported on an Accelar 1000 Series routing switch. Accept and announce policies are configured for the Accelar routing switch based on the selected protocol (OSPF or RIP).

A policy is made up of three parts: matching criteria, set parameters, and action. The matching criteria are used to decide whether or not a policy should be applied to a certain route. Once a policy is selected for a route, the set parameters are used to construct the route advertisement only if the action is “announce.”

Announce policies enable a user to selectively announce routes. Announce policies alter the routing information learned by the routers in a particular routing domain. OSPF announce policies are applied for non-OSPF routes in an Autonomous System Boundary Router (ASBR). Only an ASBR advertises the external route information into the OSPF domain. If no policies are configured or no matching policy exists for a given route, the default behavior is applied; that is, OSPF ignores the external route information.

OSPF Accept policies are applied whenever the OSPF engine computes the external routes due to a topology change or an external link-state advertisement (LSA). If there are no policies configured or no matching policy is found for a given route, the default behavior is applied; that is, the external route is included in the routing table.

RIP Announce policies are applied while sending a RIP update. The policy information is used to announce the route to other routers in the RIP routing domain. If no policies are configured or no matching policy exists for a given route, the default behavior is applied; that is, RIP-learned routes will be announced and all non-RIP routes will be ignored.

RIP Accept policies are applied whenever the router receives a RIP update. The policy is used to selectively accept routes from the RIP update. If no policies are configured or no matching policy exists for a given route, the default behavior is applied; that is, the route is included in the routing table.

Creating Policies

Creating policies involves specifying the following:

- match criteria—used to determine whether or not the policy will be applied to a route
- set parameters—used in controlling router advertisements
- action—used to determine whether to announce/accept or ignore the route that meets the matching criteria of the policy

A policy may be applicable for routes to a single network or a list of networks. Match criteria for such policies are specified in terms of network lists. A network list is created by grouping one or more network addresses together. Prior to use, the networks that are used in network lists should be entered in the network address table.

A policy may also be applicable for routes learned over specific interfaces or from specific gateways. Match criteria for such policies are specified in terms of interface/router lists. An interface/router list is created by grouping one or more IP addresses from the interface/router table. Prior to use, the IP addresses used in the interface/router lists should be entered in the interface/router table.

This chapter demonstrates the following steps leading to creation of policies:

1. Creating entries in the interface/router or network address tables
2. Creating interface/router or network lists
3. Including interface/router addresses or network addresses in the appropriate lists
4. Creating policies using interface/router lists or network lists as matching criteria

Figure 12-1 is an example of a routing table. The following procedures will include all the steps needed to:

- Create a RIP Announce policy to announce only routes for 12.128.0.0/ 255.255.255.128 and 12.120.0.128/255.255.255.128.
- Create a RIP Accept policy to accept only routes from interface 192.168.23.3.

The effect of the Accept policy will be shown in the resulting routing table in Figure 12-11 on page 12-12.

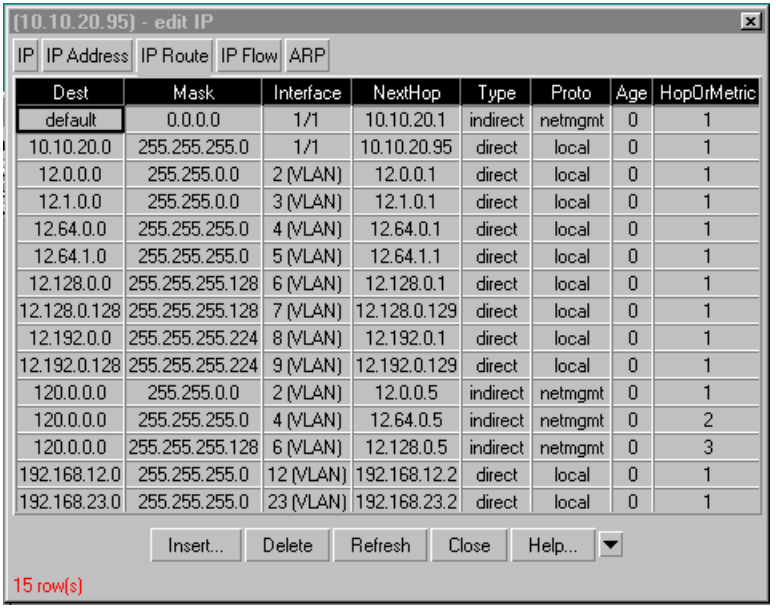

Dest	Mask	Interface	NextHop	Type	Proto	Age	HopOrMetric
default	0.0.0.0	1/1	10.10.20.1	indirect	netmgmt	0	1
10.10.20.0	255.255.255.0	1/1	10.10.20.95	direct	local	0	1
12.0.0.0	255.255.0.0	2 (VLAN)	12.0.0.1	direct	local	0	1
12.1.0.0	255.255.0.0	3 (VLAN)	12.1.0.1	direct	local	0	1
12.64.0.0	255.255.255.0	4 (VLAN)	12.64.0.1	direct	local	0	1
12.64.1.0	255.255.255.0	5 (VLAN)	12.64.1.1	direct	local	0	1
12.128.0.0	255.255.255.128	6 (VLAN)	12.128.0.1	direct	local	0	1
12.128.0.128	255.255.255.128	7 (VLAN)	12.128.0.129	direct	local	0	1
12.192.0.0	255.255.255.224	8 (VLAN)	12.192.0.1	direct	local	0	1
12.192.0.128	255.255.255.224	9 (VLAN)	12.192.0.129	direct	local	0	1
120.0.0.0	255.255.0.0	2 (VLAN)	12.0.0.5	indirect	netmgmt	0	1
120.0.0.0	255.255.255.0	4 (VLAN)	12.64.0.5	indirect	netmgmt	0	2
120.0.0.0	255.255.255.128	6 (VLAN)	12.128.0.5	indirect	netmgmt	0	3
192.168.12.0	255.255.255.0	12 (VLAN)	192.168.12.2	direct	local	0	1
192.168.23.0	255.255.255.0	23 (VLAN)	192.168.23.2	direct	local	0	1

Figure 12-1. Routing Table

Creating Interface/Router Lists

To create interface address lists:

1. **Select Routing->IP Policy.**

   The Edit Policy Interface/Router Addresses window opens (Figure 12-2).

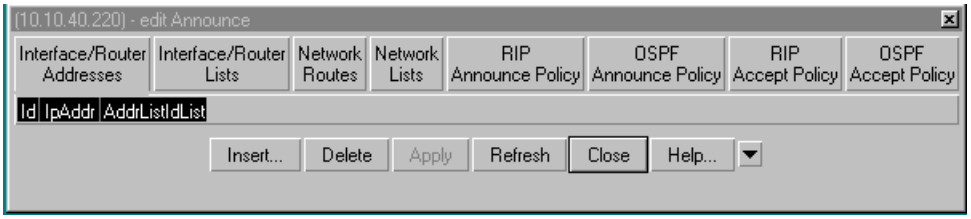

Figure 12-2. Interface/Router Addresses Window

2. **Click on Insert to insert an address.**

   The Insert Interface Address window opens (Figure 12-3).

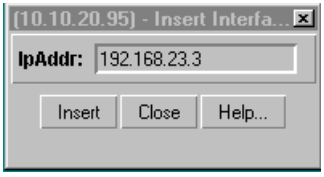

Figure 12-3. Insert Interface Address Window

3. **Enter an IP address and click on Insert to add the address.**

   The addresses in the Interface/Router Address table are used to create Interface/Router Lists.

4. **Select Routing->IP Policy->Interface/Router Lists (Figure 12-4).**

Figure 12-4. Edit Interface/Router Lists Window

5. **To create lists, click on Insert.**

   The Insert Interface/Routers window opens (Figure 12-5) with a list of the IP addresses that were entered in the Interface/Router Address table.

6. **Enter a list ID and name.**

7. **Select the IP addresses to be added to the list and click on Insert.**

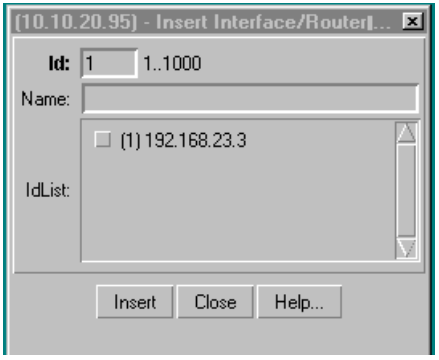

Figure 12-5. Insert Interface/Router List Window

Creating a RIP Accept Policy

The next step in this procedure is to create a RIP Accept policy that will accept RIP routes from gateway routers with IP addresses included in the interface router list just created.

1. **Select Routing->IP Policy->RIP Accept Policy, and click on Insert.**

   The Insert RIP Accept Policy window opens (Figure 12-6).

2. **Enter information in the fields as defined in Table 12-1.**

   You can select one or more matching criteria for the policy.

3. **Create RIP Accept policy 4001, and click on Insert.**

   The policy is now displayed in the RIP Accept Policy window (Figure 12-7).

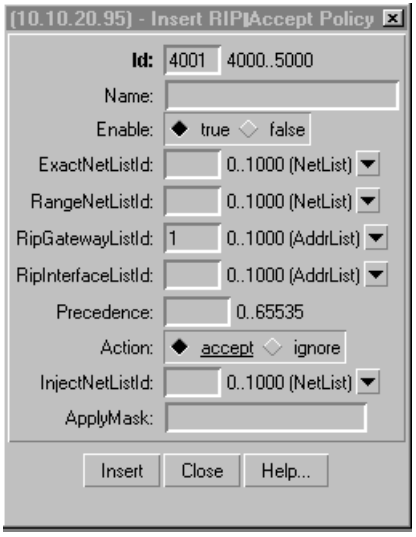

Figure 12-6. Insert RIP Accept Policy Window

Table 12-1. Insert RIP Accept Policy Window Fields

Field	Description
Id	The RIP Accept policy ID (4000-5000).
Name	The character string naming the Accept policy.
Enable	Set true to enable or false to disable the RIP Accept policy.
ExactNet	The exact network list ID (0 to 1000). For exact lists, the route and mask must both match. Empty means accept all.
RangeNet	The range network list ID (0 to 1000). For a range list, apply the mask and the result must match. Empty means accept all.
RipGateway	The RIP gateway address list ID (0 to 1000) learned from one of the listed gateways. Empty means accept all.
RipInterface	The RIP interface address list ID (0 to 1000) learned on the listed interfaces. Empty means accept all.
Precedence	If multiple policies match, the higher precedence is used (0-65535).
Action	To accept or ignore the route.
InjectNetListId	The inject network list ID (0 to 1000). After a match is found, all networks in this list will be included in the routing table.
ApplyMask	Applies the subnet mask to be used in the routing table.

Figure 12-7. RIP Accept Policy Window

Creating Network Lists

To create network lists:

1. **Select Routing->IP Policy->Network Routes.**

   The Network Routes window opens (Figure 12-8).

2. **Enter the IP addresses and subnet masks of network routes, and click on Insert.**

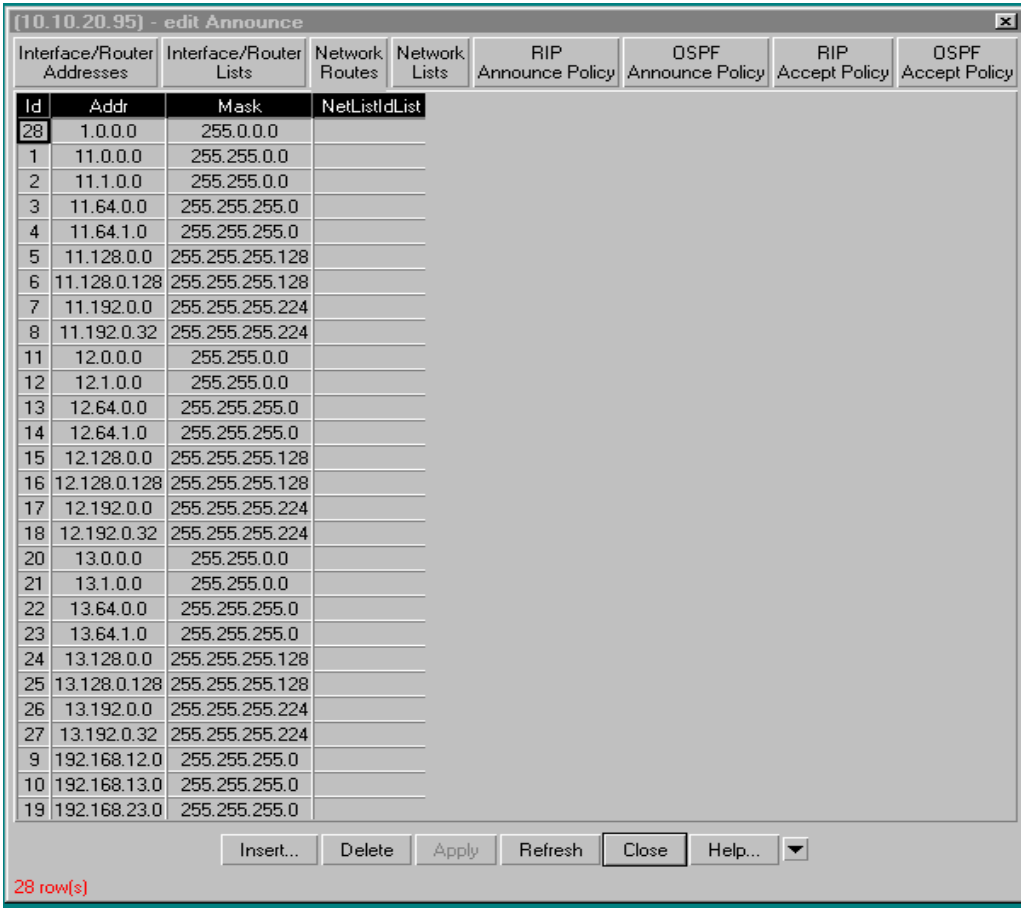

Figure 12-8. Network Routes Window

3. **To create network lists, select Routing->IP Policy->Network Lists. Then click on Insert.**

   The Insert Network Lists window is displayed (Figure 12-9).

4. **Enter a network list ID (1 to 1000) and name.**

5. **Select the network addresses to belong to the network list and click on Insert.**

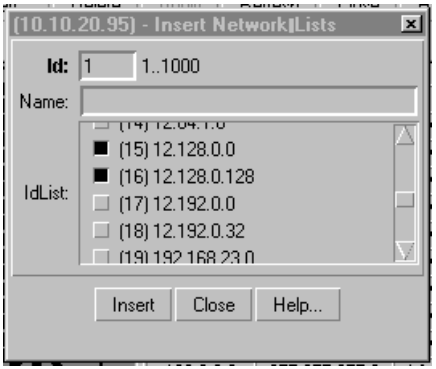

Figure 12-9. Insert Network Lists Window

Creating a RIP Announce Policy

To set up or edit a RIP Announce policy:

1. **Select Routing->IP Policy->RIP Announce Policy.**

   The Insert RIP Announce Policy window appears (Figure 12-10).

2. **Enter information in the fields as defined in Table 12-2.**

   In this case, create two policies and apply to the selected Net List:

   - The first policy matches all routes to ignore (not announce) the route.
   - The second policy announces routes in the network list ID 1.

More than one policy can be applied to the same route. When information is conflicting, precedences are used to determine which policy to apply. The policy with the higher value of precedence is used. If no precedence is entered, the policy with the higher ID number will be applied.

3. **Click on Insert to apply the policies.**

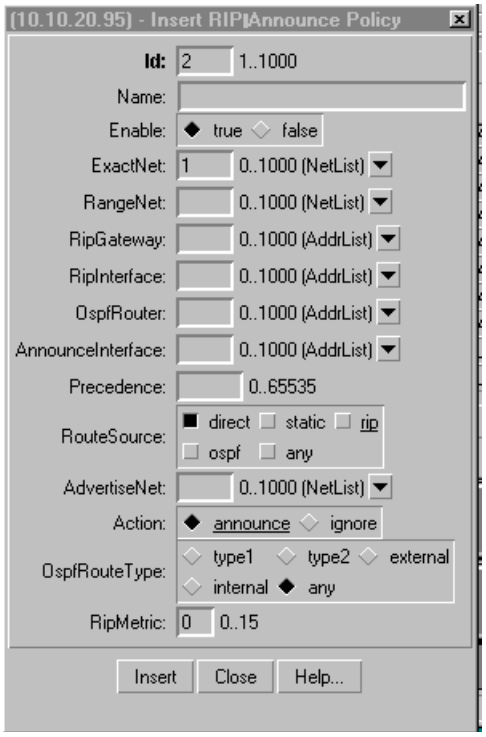

Figure 12-10. Insert RIP Announce Policy Window

Table 12-2. Insert RIP Announce Policy Window Fields

Field	Description
Id	The RIP Announce policy ID (1 to 1000).
Name	The character string naming the Announce policy.
Enable	Set true to enable or false to disable the RIP Announce policy.
ExactNet	The exact network list ID (0 to 1000). For exact lists, the route and mask must both match. Empty means accept all.
RangeNet	Å range of network lists ID (0 to 1000).
RipGateway	The RIP gateway address list ID (0 to 1000). Progagates only routes learned from specific RIP gateway.
RipInterface	The RIP interface address list ID (0 to 1000). Progagates only routes learned from specific RIP interfaces.
OspfRouter	The OSPF router address list ID (0 to 1000). Progagates only routes learned from specific OSPF gateways.
AnnounceInterface	The Announce Interface address list ID (0 to 1000).
Precedence	If multiple policies match, the higher precedence is used (0 to 65535).
RouteSource	Set to direct, static, or RIP.
AdvertiseNet	The advertise network list ID (0 to 1000).
Action	Announce or ignore.
OspfRouteType	Type1, Type2, External, Internal, or any.
RipMetric	The number of hops (0 to 15). Specifies the metric advertised with the networks defined in the Advertise net list.

Resulting Actions

The routing table in Figure 12-11 shows the result of the policies just applied. Only two routes will be announced:

- RIP Accept policy 4001 determines that only RIP routes from the gateway router in policy 4001 will be accepted. RIP routes are learned (accepted) only from IP address 192.168.23.3.
- RIP routes are announced only for networks 12.128.00 and 12.128.0.128.

(10.10.20.95) - edit IP

IP | IP Address | IP Route | IP Flow | ARP

Dest	Mask	Interface	NextHop	Type	Proto	Age	HopOrMetric
default	0.0.0.0	1/1	10.10.20.1	indirect	netmgmt	0	1
10.10.20.0	255.255.255.0	1/1	10.10.20.95	direct	local	0	1
11.0.0.0	255.255.0.0	23 (VLAN)	192.168.23.3	indirect	rip	8	2
11.1.0.0	255.255.0.0	23 (VLAN)	192.168.23.3	indirect	rip	8	2
11.64.0.0	255.255.255.0	23 (VLAN)	192.168.23.3	indirect	rip	8	2
11.64.1.0	255.255.255.0	23 (VLAN)	192.168.23.3	indirect	rip	8	2
11.128.0.0	255.255.255.128	23 (VLAN)	192.168.23.3	indirect	rip	8	2
11.128.0.128	255.255.255.128	23 (VLAN)	192.168.23.3	indirect	rip	8	2
11.192.0.0	255.255.255.224	23 (VLAN)	192.168.23.3	indirect	rip	8	2
11.192.0.128	255.255.255.224	23 (VLAN)	192.168.23.3	indirect	rip	8	2
12.0.0.0	255.255.0.0	2 (VLAN)	12.0.0.1	direct	local	0	1
12.1.0.0	255.255.0.0	3 (VLAN)	12.1.0.1	direct	local	0	1
12.64.0.0	255.255.255.0	4 (VLAN)	12.64.0.1	direct	local	0	1
12.64.1.0	255.255.255.0	5 (VLAN)	12.64.1.1	direct	local	0	1
12.128.0.0	255.255.255.128	6 (VLAN)	12.128.0.1	direct	local	0	1
12.128.0.128	255.255.255.128	7 (VLAN)	12.128.0.129	direct	local	0	1
12.192.0.0	255.255.255.224	8 (VLAN)	12.192.0.1	direct	local	0	1
12.192.0.128	255.255.255.224	9 (VLAN)	12.192.0.129	direct	local	0	1
110.0.0.0	255.255.0.0	23 (VLAN)	192.168.23.3	indirect	rip	8	2
110.0.0.0	255.255.255.0	23 (VLAN)	192.168.23.3	indirect	rip	8	3
110.0.0.0	255.255.255.128	23 (VLAN)	192.168.23.3	indirect	rip	8	4
120.0.0.0	255.255.0.0	2 (VLAN)	12.0.0.5	indirect	netmgmt	0	1
120.0.0.0	255.255.255.0	4 (VLAN)	12.64.0.5	indirect	netmgmt	0	2
120.0.0.0	255.255.255.128	6 (VLAN)	12.128.0.5	indirect	netmgmt	0	3
192.168.12.0	255.255.255.0	12 (VLAN)	192.168.12.2	direct	local	0	1
192.168.13.0	255.255.255.0	23 (VLAN)	192.168.23.3	indirect	rip	8	2
192.168.23.0	255.255.255.0	23 (VLAN)	192.168.23.2	direct	local	0	1

Insert... | Delete | Refresh | Close | Help...

27 row(s)

Figure 12-11. Routing Table with Policies Applied

OSPF Announce Policy

OSPF policies are configured in the same manner as RIP. After interface routes or network routes and lists are created, set up policies in the OSPF Policy windows.

To set up or edit an OSPF Announce policy:

1. **Select Routing->IP Policy->OSPF Announce Policy.**

   The Insert OSPF Announce Policy window opens (Figure 12-12).

2. **Enter information in the fields as defined in Table 12-3, and click on Insert to enter.**

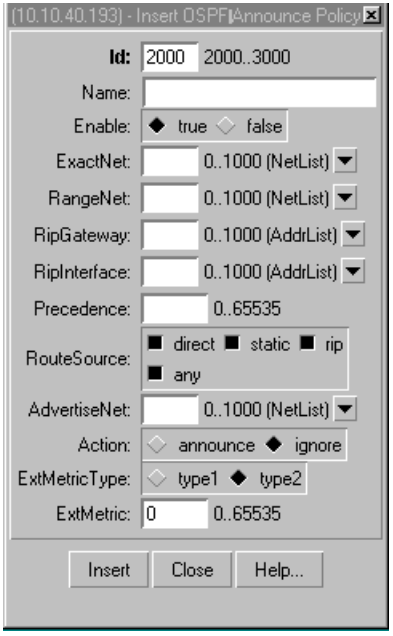

Figure 12-12. Insert OSPF Announce Policy Window

Table 12-3. Insert OSPF Announce Policy Window Fields

Field	Description
Id	OSPF Announce policy ID (2000 to 3000).
Name	The character string naming the Announce policy.
Enable	Set true to enable or false to disable the OSPF Announce policy.
ExactNet	The exact network list ID (0 to 1000). For exact lists, the route and mask must both match. Empty means accept all.
RangeNet	The network range list ID (0 to 1000).
RipGateway	The RIP gateway address list ID (0 to 1000). Progagates only routes learned from specific RIP gateway.
RipInterface	The RIP interface address list ID (0 to 1000). Progagates only routes learned from specific RIP interfaces.
Precedence	If multiple policies match, the higher precedence is used (0 to65535).
RouteSource	Set to direct, static, RIP, or any.
AdvertiseNet	The advertise network list ID (0 to 1000).
Action	Announce or ignore.
ExtMetricType	The external metric type: Type1 or Type 2.
ExtMetric	The external metric (0 to 65535).

OSPF Accept Policy

To set up or edit an OSPF Accept policy:

1. **Select Routing->IP Policy->OSPF Accept Policy.**

   The Insert OSPF Accept Policy window opens (Figure 12-13).

2. **Enter information in the fields as defined in Table 12-4, and click on Insert to enter.**

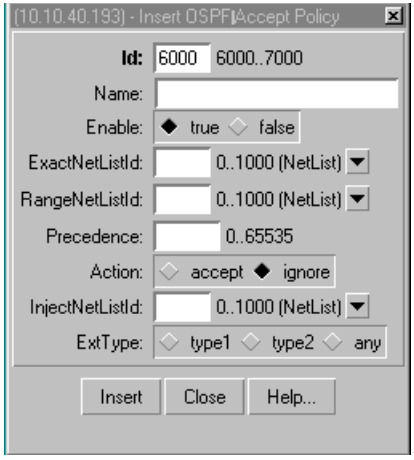

Figure 12-13. Insert OSPF Accept Policy Window

Table 12-4. Insert OSPF Accept Policy Window Fields

Field	Description
Id	OSPF Accept policy ID (6000 to 7000).
Name	The character string naming the Accept policy.
Enable	Set true to enable or false to disable the OSPF Accept policy.
ExactNet	The exact network list ID (0 to 1000). For exact lists, the route and mask must both match.Empty means accept all.
RangeNet	The range network list ID (0 to 1000). For a range list, apply the mask and the result must match. Empty means accept all.
Precedence	If multiple policies match, the higher precedence is used (0 to 65535).
Action	To accept or ignore the route.
InjectNetListId	The inject network list ID (0 to 1000). Once a match is found, all networks in this list will be included in the routing table.
ExtType	Type 1, Type 2, or both.