CHAPTER 2

IP Routing & Bridging

TCP/IP Routing: Ethernet Dialog Box

To access this dialog box (Figure 2-1), select Ethernet/TCP/IP Routing from the Device View.

Figure 2-1 TCP/IP Routing: Ethernet Configuration Dialog Box

If you need more information about the IP protocol, see the "IP 101" section on page A-1.

IP Routing/Bridging/Off

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

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

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

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

IP Address

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

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

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

Network IP Subnet Mask

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

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

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

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

Network IP Broadcast Address

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

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

Routing Protocol

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

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

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

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

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

- If **RIP 2** is selected with this pull-down menu, the router will send and/or accept RIP 2 packets over this interface, and will then periodically update its routing table with the information provided from these packets. On a large network, an up-to-date routing table will enhance network performance since the router will always be aware of the optimal path to use when sending packets.
- If **RIP 1** is selected with this pull-down menu, the router will send and/or accept RIP 1 packets, and will then periodically update its routing table with the information provided from these packets.
- If **None** is selected with this pull-down menu, the router will not be able to update its routing table and will always direct traffic for addresses it does not have a route for (addresses not on one of the networks connected to its interfaces) to the "gateway/port" defined in its IP Static Route dialog box. It will then be the responsibility of the default router to direct the packets to the correct address. For information on setting the default router see the discussion of the IP Static Route dialog box in the "IP Static Routing Dialog Box" section on page 2-18.

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

RIP Split Horizon

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

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

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

Output RIP - Input RIP

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

Directed Broadcast

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

Options

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

OSPF

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

TCP/IP Routing: WAN Configuration Dialog Box

To access this dialog box (Figure 2-2), select WAN/TCP/IP Routing from the Device View.

Figure 2-2 TCP/IP Routing: WAN Configuration Dialog Box

If you need more information about the IP protocol, see the "IP 101" section on page A-1.

IP Routing/Bridging/Off

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

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

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

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

Numbered Interface

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

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

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

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

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

IP Address

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

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

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

Network IP Subnet Mask

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

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

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

To have the VPN 5000 Manager calculate a default mask, make sure that the Subnet Mask field is empty, position the cursor in the IP Address field, then just tab through the Subnet Mask field.

Network IP Broadcast Address

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

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

Routing Protocol

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

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

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

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

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

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

- If **RIP 1** is selected with this pull-down menu, the router will send and/or accept RIP 1 packets, and will then periodically update its routing table with the information provided from these packets.
- If **None** is selected with this pull-down menu, the router will not be able to update its routing table and will always direct traffic for addresses it does not have a route for (addresses not on one of the networks connected to its interfaces) to the "default router" defined in its IP Static Route dialog box. It will then be the responsibility of the default router to direct the packets to the correct address. For information on setting the default router see the discussion of the IP Static Route dialog box in the "IP Static Routing Dialog Box" section on page 2-18.

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

Update Method

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

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

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

RIP Split Horizon

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

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

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

Output RIP - Input RIP

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

Directed Broadcast

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

Options

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

OSPF

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

TCP/IP Routing: VPN Configuration Dialog Box

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

Once you have created a VPN port, you may access the TCP/IP Routing: VPN Configuration dialog box (Figure 2-3) by clicking TCP/IP Routing under the VPN port's icon.

Figure 2-3 TCP/IP Routing: VPN Configuration Dialog Box

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

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

IP Routing/IP Bridging/IP Off

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

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

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

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

Numbered Interface

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

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

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

IP Address

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

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

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

Network IP Subnet Mask

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

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

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

To have the VPN 5000 Manager calculate a default mask, make sure that the Subnet Mask field is empty, position the cursor in the IP Address field, then just tab through the Subnet Mask field.

Network IP Broadcast Address

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

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

Routing Protocol

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

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

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

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

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

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

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

Update Method

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

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

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

- If **Periodic** is selected with this pull-down menu, the router will use the standard RIP protocol, which sends RIP packets over the link every 30 seconds.

RIP Split Horizon

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

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

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

Output RIP - Input RIP

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

Directed Broadcast

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

OSPF

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

TCP/IP Routing: Bridge Configuration Dialog Box

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

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

Logically, the IP Bridge Group is treated by the device as an interface (Bridge 0). The settings in the TCP/IP Routing: Bridge 0 Configuration dialog box (Figure 2-5) determine the IP parameters for all of the physical network interfaces which make up the IP Bridge Group. This is shown schematically in Figure 2-4.

Figure 2-4 Bridge Logical Diagram

If you need more information about bridging, see the "Bridging 101" section on page A-9.

To access the TCP/IP Routing :Bridge dialog box (Figure 2-5), select Bridge 0/TCP/IP Routing from the Device View.

Figure 2-5 TCP/IP Routing: Bridge 0 Configuration Dialog Box

IP Routing/Off

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

- If set to **IP Routing**, then IP packets received on a member interface of the IP Bridge Group which cannot simply be bridged to another member interface of the group are passed on for IP routing.

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

IP Address

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

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

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

Network IP Subnet Mask

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

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

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

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

Network IP Broadcast Address

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

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

Routing Protocol

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

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

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

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

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

- If **RIP 2** is selected with this pull-down menu, the router will send and/or accept RIP 2 packets via this Bridge Group's member interfaces, and will then periodically update its routing table with the information provided from these packets. On a large network, an up-to-date routing table will enhance network performance since the router will always be aware of the optimal path to use when sending packets.
- If **RIP 1** is selected with this pull-down menu, the router will send and/or accept RIP 1 packets, and will then periodically update its routing table with the information provided from these packets.
- If **None** is selected with this pull-down menu, the router will not be able to update its routing table and will always direct traffic for addresses it does not have a route for (addresses not on one of the networks connected to its interfaces) to the "default router" defined in its IP Static Route dialog box. It will then be the responsibility of the default router to direct the packets to the correct address. For information on setting the default router see the "IP Static Routing Dialog Box" section on page 2-18.

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

RIP Split Horizon

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

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

- If **Split Horizon** is selected with this pull-down menu, the router will apply the split horizon technique to routes being output over this Bridge Group's member interfaces.

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

Directed Broadcast

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

Options

The options button brings up the Bridge-TCP/IP Routing Options dialog box which provides access to Proxy ARP, UDP Relays and other configuration information. Refer to "TCP/IP Routing Options" section on page 2-24 for more information.

OSPF

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

IP Subinterface Dialog Box

Subinterfaces are added to the edit area of a device by right-clicking on any configuration item for the device, then choosing Sub interface/Add. This action opens the Add IP Subinterface dialog box (Figure 2-6). To delete a sub interface, right-click on the subinterface icon, then choose Subinterface/Delete. These functions are also available in the **Device** menu.

Figure 2-6 Add IP Subinterface Dialog Box

Once you have created a subinterface, you may access the IP Subinterface Configuration dialog box (Figure 2-7) by clicking on TCP/IP under the subinterface icon.

Figure 2-7 IP Subinterface Configuration Dialog Box

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

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

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

IP Connection Dialog Box

The IP Connection dialog box controls the IP settings for the IPsec-only port on a VPN 5000 concentrator with two or more Ethernet interfaces. This port will only handle IPsec traffic (i.e., authenticated and/or encrypted packets).

To access this dialog box (Figure 2-8), select Ethernet/TCP/IP Routing from the Device View.

Figure 2-8 IP Connection Dialog Box

IP On/IP Off

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

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

IP Address

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

This IP address must be on the same IP network as the IPsec Gateway, which is configured using the IPsec Gateway dialog box (under Global/IPsec Gateway).

Network IP Subnet Mask

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

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

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

Network IP Broadcast Address

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

IP Static Routing Dialog Box

To open the Static IP Routing Configuration dialog box (Figure 2-9), select Global/IP Static Routes. This dialog box displays static routes which have already been entered, but is not used to add or modify the entries.

Figure 2-9 IP Static Routes Dialog Box

To add or modify IP static route entries, you must access the Add Static Route dialog box by selecting the **Add...** or **Modify...** buttons in the Static IP Routing Configuration dialog box. The Add Static Route dialog box (Figure 2-10) allows you to set a default IP router and to assign multiple static routes.

Figure 2-10 Add Static Route Dialog Box

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

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

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

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

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

Destination

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

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

Mask

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

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

Gateway

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

- If **IP Address** is selected, enter the IP address of the gateway.
- If **Port** is selected, use the pull-down menu to select an interface on the device you are configuring.

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

Metric

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

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

Redistribute via

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

- If **None** is specified, the static route will not be redistributed. Only one routing protocol can be selected for redistributing each static route.
- If **RIP** is specified, the static route entry will be redistributed into the RIP routing protocol which means that other routers will be able to choose this device as a way to forward packets to the destination address, depending on the metric and what other routes are available.
- Routing information received via RIP from other routers will be redistributed out other interfaces where RIP processing is enabled. When routes are rebroadcast in this fashion, the metric for this route is increased by 1, which increases the cost of the route.
- If **OSPF1** or **OSPF2** is specified, the static route entry will be redistributed into the OSPF routing protocol. The 1 or 2 refer to the two types of external metrics which may be used in OSPF.
- A type 1 cost is the sum of both the external cost and the internal cost used to reach that router. The cost of a type 2 route is simply the external cost, regardless of the interior (i.e., within OSPF) cost to reach that router.
- If **BGP** is specified, the static route entry will be redistributed into the BGP routing protocol.

Ethernet IP Options or Bridge IP Options

To access this dialog box (Figure 2-11), select Ethernet/ or Bridge/TCP/IP Routing from the Device View, then click on the **Options** button.

This dialog box provides access to settings for IP Proxy ARP settings and the UDP Forwarding Agents dialog box described in the "UDP Forwarding Agents (Relays)" section on page 2-21.

Figure 2-11 Ethernet or Bridge TCP/IP Options Dialog Box

IP Proxy ARP

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

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

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

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

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

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

UDP Forwarding Agents (Relays)

The Relays button brings up a configuration dialog box (Figure 2-12) that can be used to turn on a relay agent in the device for UDP (User Datagram Protocol) broadcast packets.

Figure 2-12 UDP Forwarding Agents Dialog Box

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

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

Server IP Address

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

UDP Ports/Protocols

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

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

WAN IP Options

To access this dialog box (Figure 2-13), select Ethernet/ or Bridge/TCP/IP Routing from the Device View, then click on the **Options** button.

Figure 2-13 WAN IP Options Dialog Box

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

Optional Remote End-Node Address

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

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

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

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

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

Van Jacobson Header Compression

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

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

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

Send IP Address Configure Request

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

TCP/IP Routing Options

This dialog box (Figure 2-14) can be brought up selecting Options/TCP/IP Routing from the Device View. These parameters are not associated with a particular interface and are global to the device.

Figure 2-14 TCP/IP Routing Options Dialog Box

RIP V2 Password

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

IP Multiprotocol Precedence Dialog Box

This dialog box (Figure 2-15) sets the precedence order the router will follow for including routes in its routing table when multiple IP routing protocols are in use on the network. To access this dialog box, select Global/IP Multiprotocol Precedence from the Device View.

Figure 2-15 IP Multiprotocol Precedence Dialog Box

Protocol Precedence

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

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

An exception to the precedence rule is an OSPF external (i.e., type ASE) route. OSPF external routes will be overwritten by a RIP or static route, regardless of the precedence. This is because OSPF external routes originally come from another protocol, usually RIP or static. If the router is running both RIP

and OSPF, but another router on the network is redistributing RIP into OSPF, the RIP routes would be overwritten by OSPF external routes without this exception. In order to get the RIP routes via OSPF external routes, simply uncheck the **Input RIP** checkbox in the TCP/IP Routing dialog box, and it will then install the routes as OSPF externals.

IP Route Redistribution

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

Redistribution of static routes can be done using the IP Multiprotocol Precedence dialog box (Figure 2-16).

Figure 2-16 IP Route Redistribution Dialog Box

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

OSPF Route Aggregation

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

OSPF Route Aggregation is only used for importing static and RIP routes into OSPF.

RIP to OSPF

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

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

Default into OSPF

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

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

OSPF to RIP

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

BGP to OSPF

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

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

BGP to RIP

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

RIP to BGP

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

OSPF to BGP

This checkbox sets whether the router will redistribute OSPF routes into the BGP routing domain. BGP will provide its own hop count in its route advertisements.