Chapter 11
Configuring IP Multicast Protocols

This chapter covers how to configure the HP ProCurve 9304M, 9308M, and 6308M-SX routing switches for Protocol Independent Multicast (PIM) and Distance Vector Multicast Routing Protocol (DVMRP).

A summary of all CLI commands discussed in this chapter can also be found in “Command Line Interface Commands” on page B-1.

Overview of IP Multicasting

Multicast protocols allow a group or channel to be accessed over different networks by multiple stations (clients) for the receipt and transmit of multicast data.

Distribution of stock quotes, video transmissions such as news services and remote classrooms, and video conferencing are all examples of applications that use multicast routing.

The 9304M, 9308M, and 6308M-SX routing switches support two different multicast routing protocols—Distance Vector Multicast Routing Protocol (DVMRP) and Protocol-Independent Multicast (PIM) protocol along with the Internet Group Membership Protocol (IGMP).

PIM and DVMRP are broadcast and pruning multicast protocols that deliver IP multicast datagrams. The protocols employ reverse path lookup check and pruning to allow source-specific multicast delivery trees to reach all group members. DVMRP and PIM build a different multicast tree for each source and destination host group.

NOTE: Both DVMRP and PIM can concurrently operate on different ports of a routing switch.

Multicast Terms

The following are commonly used terms in discussing multicast-capable routers. These terms are used throughout this chapter:

Node: Refers to a router.

Root Node: The node that initiates the tree building process. It is also the router that sends the multicast packets down the multicast delivery tree.

Upstream: Represents the direction from which a router receives multicast data packets. An ***upstream router*** is a node that sends multicast packets.

Downstream: Represents the direction to which a router forwards multicast data packets. A ***downstream router*** is a node that receives multicast packets from upstream transmissions.

Group Presence: Means that a mulitcast group has been learned from one of the directly connected interfaces. Members of the multicast group are present on the router.

Intermediate nodes: Routers that are in the path between source routers and leaf routers.

Leaf nodes: Routers that do not have any downstream routers.

Multicast Tree: A unique tree is built for each source group (S,G) pair. A multicast tree is comprised of a root node and one or more nodes that are leaf or intermediate nodes.

Changing Global IP Multicast Parameters

IGMP allows routers to limit the multicast of IGMP packets to only those ports on the router that are identified as IP Multicast members.

The router actively sends out host queries to identify IP Multicast groups on the network, inserts the group information in an IGMP packet, and forwards the packet to IP Multicast neighbors.

The following parameters apply to PIM and DVMRP:

- IGMP query interval – Specifies how often the routing switch queries an interface for group membership. Possible values are 1 – 3600. The default is 60.
- IGMP group membership time – Specifies how many seconds an IP Multicast group can remain on a routing switch interface in the absence of a group report. Possible values are 1 – 7200. The default is 60.
- IGMP maximum response time – Specifies how many seconds the routing switch will wait for an IGMP response from an interface before concluding that the group member on that interface is down and removing the interface from the group. Possible values are 1 – 10. The default is 10.

To change these parameters, you must first enter the following CLI command at the global CLI level:

```
HP9300(config)# ip multicast-routing
```

Syntax: [no] ip multicast-routing

NOTE: You must enter the **ip multicast-routing** command before changing the global IP Multicast parameters. Otherwise, the changes do not take effect and the software uses the default values.

Modifying IGMP Query Interval Period

The IGMP query interval period defines how often a router will query an interface for group membership. Possible values are 1 – 3,600 seconds and the default value is 60 seconds.

USING THE CLI

To modify the default value for the IGMP query interval, enter the following:

```
HP9300(config)# ip igmp query 120
```

syntax: ip igmp query-interval <1-3600>

USING THE WEB MANAGEMENT INTERFACE

To modify the default value for the IGMP query interval:

1. Select <u>IGMP</u> from the DVMRP configuration sheet. The panel shown in Figure 11.1 will appear.
2. Enter a value from 1 – 3600.
3. Select the Apply button to assign the changes.

Modifying IGMP Membership Time

Group membership time defines how long a group will remain active on an interface in the absence of a group report. Possible values are from 1 – 7200 seconds and the default value is 140 seconds.

USING THE CLI

To define an IGMP membership time of 240 seconds, enter the following:

```
HP9300(config)# ip igmp group-membership-time 240
```

syntax: ip igmp group-membership-time <1-7200>

USING THE WEB MANAGEMENT INTERFACE

To modify the default value for the IGMP membership time, you would do the following:

1. Select IGMP from the DVMRP configuration sheet. The panel shown in Figure 11.1 will appear.
2. Enter a value from 1 – 7200.
3. Select the Apply button to assign the changes.

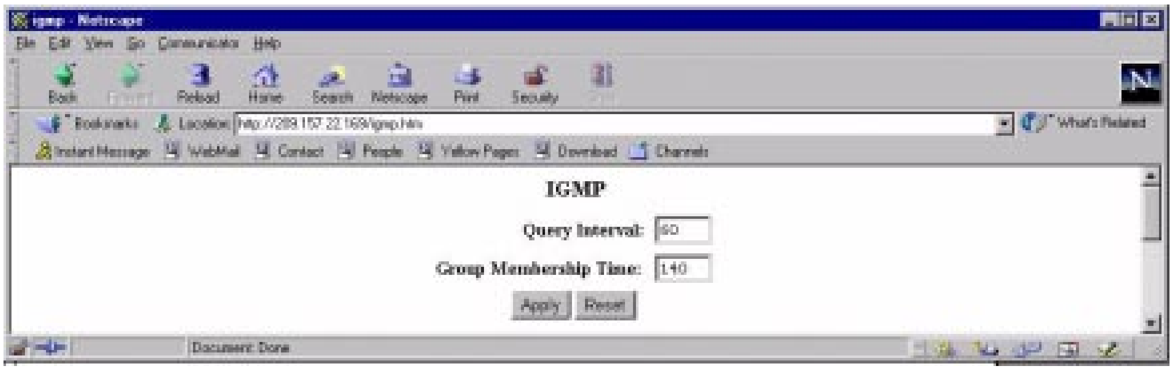

Figure 11.1 IGMP configuration panel noting default values

Modifying IGMP Maximum Response Time

Maximum response time defines how long the routing switch will wait for an IGMP response from an interface before concluding that the group member on that interface is down and removing the interface from the group. Possible values are 1 – 10. The default is 10.

USING THE CLI

To change the maximum response time to 5 seconds, enter the following:

```
HP9300(config)# ip igmp max-response-time 5
```

syntax: ip igmp max-response-time <1-10>

USING THE WEB MANAGEMENT INTERFACE

You cannot change this parameter using the Web management interface.

Disabling IGMP Queries on Individual Ports

By default, all ports are enabled to send and receive IGMP queries. To disable IGMP queries on a port, enter commands such as the following:

```
HP9300(config)# int e 1/5
HP9300(config-if-1/5) ip-multicast-disable
```

Syntax: [no] ip-multicast-disable

PIM Overview

PIM was introduced to simplify some of the complexity of the routing protocol at the cost of additional overhead tied with a greater replication of forwarded multicast packets. PIM is similar to DVMRP in that PIM builds source-routed multicast delivery trees and employs reverse path check when forwarding multicast packets.

There are two modes in which PIM operates: Dense and Sparse. The Dense Mode is suitable for densely populated multicast groups, primarily in the LAN environment. The Sparse Mode is suitable for sparsely populated multicast groups with the focus on WAN.

PIM primarily differs from DVMRP by using the IP routing table instead of maintaining its own, thereby being routing protocol independent.

Initiating PIM Multicasts on a Network

Once PIM is enabled on each router, a network user can begin a video conference multicast from the server on R1. When a multicast packet is received on a PIM-capable router interface, the interface checks its IP routing table to determine whether the interface that received the message provides the shortest path back to the source. If the interface does provide the shortest path back to the source, the multicast packet is then forwarded to all neighboring PIM routers. Otherwise, the multicast packet is discarded and a prune message is sent back upstream.

In Figure 11.2, the root node (R1) is forwarding multicast packets for group 229.225.0.1, which it receives from the server, to its downstream nodes, R2, R3, and R4. Router R4 is an intermediate router with R5 and R6 as its downstream routers. Because R5 and R6 have no downstream interfaces, they are leaf nodes. The receivers in this example are those workstations that are resident on routers R2, R3, and R6.

Pruning a Multicast Tree

As multicast packets reach these leaf routers, the routers check their IGMP databases for the group. If the group is not in a router’s IGMP database, the router discards the packet and sends a prune message to the upstream router. The router that discarded the packet also maintains the prune state for the source, group (S,G) pair. The branch is then pruned (removed) from the multicast tree. No further multicast packets for that specific (S,G) pair will be received from that upstream router until the prune state expires. You can configure the PIM Prune Timer (the length of time that a prune state is considered valid).

For example, in Figure 11.2 the sender with address 207.95.5.1 is sending multicast packets to the group 229.225.0.1. If a PIM router receives any groups other than that group, the router discards the group and sends a prune message to the upstream PIM router.

In Figure 11.3, Router R5 is a leaf node with no group members in its IGMP database. Therefore, the router must be pruned from the multicast tree. R5 sends a prune message upstream to its neighbor router R4 to remove itself from the multicast delivery tree and install a prune state, as seen in Figure 11.3. Router 5 will not receive any further multicast traffic until the prune age interval expires.

When a node on the multicast delivery tree has all of its downstream branches (downstream interfaces) in the prune state, a prune message is sent upstream. In the case of R4, if both R5 and R6 are in a prune state at the same time, R4 becomes a leaf node with no downstream interfaces and sends a prune message to R1. With R4 in a prune state, the resulting multicast delivery tree would consist only of leaf nodes R2 and R3.

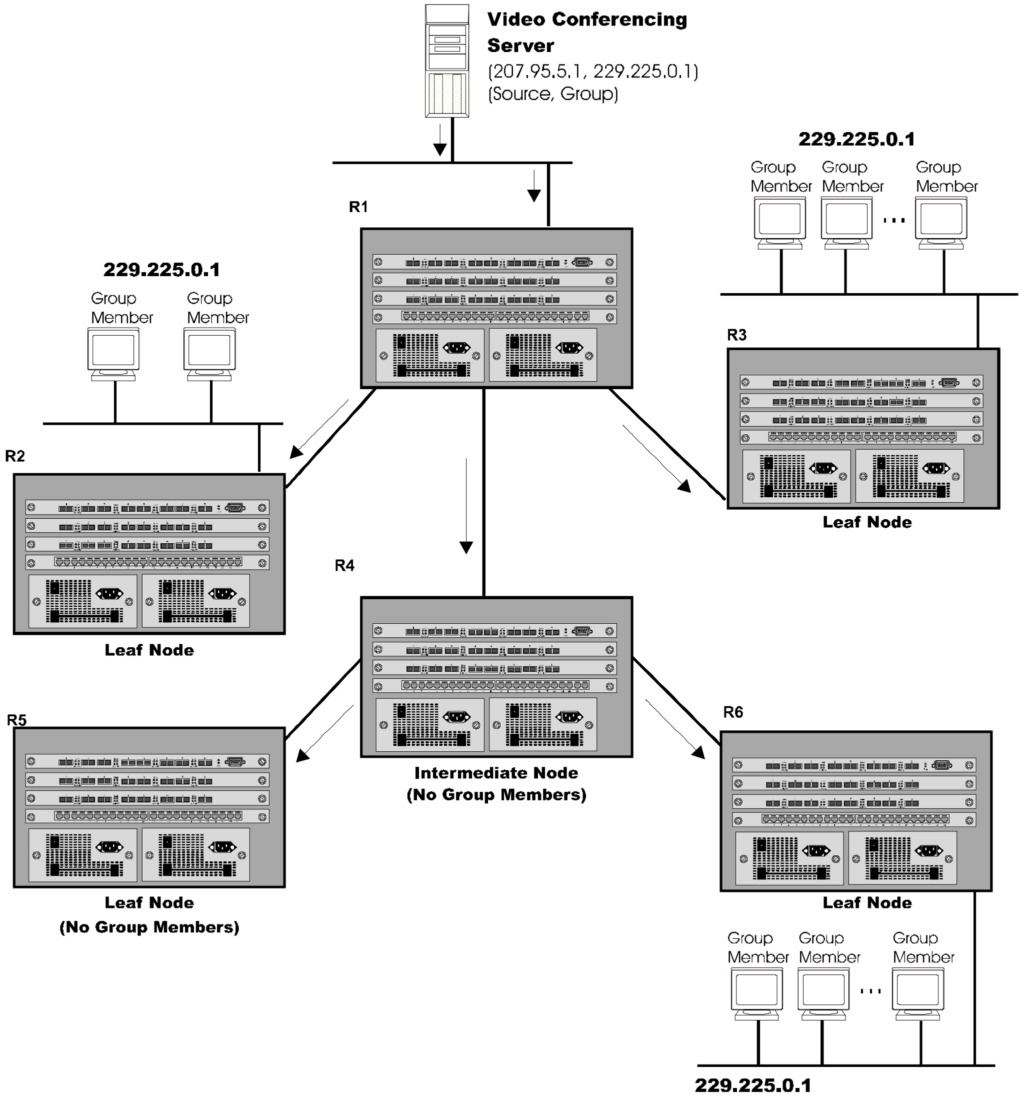

Figure 11.2 Transmission of multicast packets from the source to host group members

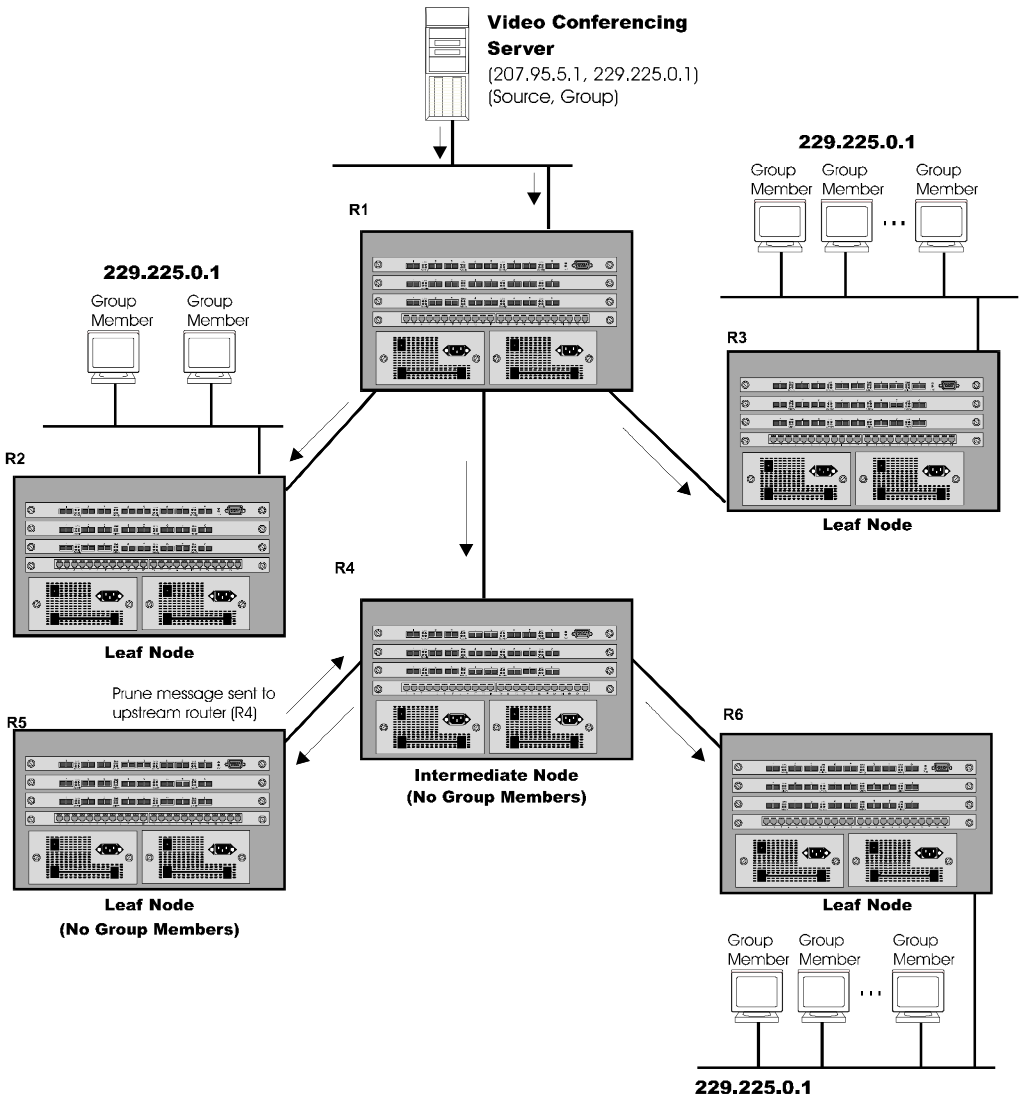

Figure 11.3 Pruning leaf nodes from a multicast tree

Grafts to a Multicast Tree

A PIM router restores pruned branches to a multicast tree by sending graft messages towards the upstream router. Graft messages start at the leaf node and travel up the tree, first sending the message to its neighbor upstream router.

In the example above, if a new 229.255.0.1 group member joins on router R6, which was previously pruned, a graft is sent upstream to R4. Since the forwarding state for this entry is in a prune state, R4 sends a graft to R1. Once R4 has joined the tree, R4 along with R6 once again receive multicast packets.

Prune and graft messages are continuously used to maintain the multicast delivery tree. No configuration is required on your part.

Configuring PIM

Enabling PIM on the Routing Switch and an Interface

By default, the PIM feature is disabled. You must enable PIM globally for the entire routing switch and individually for each interface on which you want the routing switch to perform PIM routing.

EXAMPLE: Suppose you want to initiate the use of desktop video for fellow users on a sprawling campus network. All destination workstations have the appropriate hardware and software but the routers that connect the various buildings need to be configured to support PIM multicasts from the designated video conference server as shown in Figure 11.2.

PIM is enabled on each of the routing switches shown in Figure 11.2, on which multicasts are expected. You can enable PIM on each routing switch independently or remotely from one of the routing switches with a Telnet connection. Follow the same steps for each routing switch. A reset of the routing switch is required when PIM is first enabled. Thereafter, all changes are dynamic.

USING THE CLI

EXAMPLE: To enable PIM on router1 and interface 3, enter the following:

```
HP9300(config)# router pim

HP9300(config)# int e 3

HP9300(config-if-3)# ip pim 207.95.5.1
```

USING THE WEB MANAGEMENT INTERFACE

1. Select the System link from the main menu. The System configuration panel will appear.
2. Enable PIM.
3. Select the Apply button to assign the change.
4. Select the PIM link from the main menu.
5. Select the PIM Virtual Interface link from the PIM configuration panel to configure an interface. The panel shown in Figure 11.5 will appear.
6. Select the type of interface that is being configured from the pulldown menu. You can select IP tunnel or sub-net.
7. Select the IP address of the interface being configured from the Local Address pulldown menu.

   NOTE: If you are configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the Remote Address field. IP tunneling must also be enabled and defined on the destination routing switch interface as well. Note that this field is left blank when configuring sub-net type interfaces.

8. Modify the time to live threshold (TTL) if necessary. TTL defines the minimum value required in a packet in order for the packet to be forwarded out the interface.

   NOTE: For example, if the TTL for an interface is set at 10, it means that only those packets with a TTL value of 10 or more will be forwarded. Likewise, if an interface is configured with a TTL Threshold value of 1, all packets received on that interface will be forwarded. Possible values are 1 – 64. The default value is 1.

9. Select the Add button to enable PIM on the interface.

Modifying PIM Global Parameters

PIM global parameters come with preset values. The defaults work well in most networks, but you can modify the following parameters if you need to:

- Neighbor timeout
- Hello timer
- Prune timer
- Graft retransmit timer
- Inactivity timer

Modifying Neighbor Timeout

Neighbor timeout is the interval after which a PIM router will consider a neighbor to be absent. Absence of PIM hello messages from a neighboring router indicates that a neighbor is not present.

The default value is 180 seconds.

USING THE CLI

To apply a PIM neighbor timeout value of 360 seconds to all ports on the routing switch operating with PIM, enter the following:

```
HP9300(config)# router pim
HP9300(config-pim-router)# pim-nbr-timeout 360
```

syntax: pim-nbr-timeout <60-8000>

USING THE WEB MANAGEMENT INTERFACE

1. Select the PIM link from the main menu. The PIM configuration panel shown in Figure 11.4 will appear.
2. Enter a value from 10 – 3600 into the Neighbor Router Timeout field.
3. Select the Apply button to assign the change.

Modifying Hello Timer

This parameter defines the time interval at which periodic hellos are sent out PIM interfaces. Routers use hello messages to inform neighboring routers of their presence. The default rate is 60 seconds.

USING THE CLI

To apply a PIM hello timer of 120 seconds to all ports on the routing switch operating with PIM, enter the following:

```
HP9300(config)# router pim
HP9300(config-pim-router)# pim-hello-timer 120
```

syntax: pim-hello-timer <10-3600>

USING THE WEB MANAGEMENT INTERFACE

1. Select the PIM link from the main menu. The PIM configuration panel shown in Figure 11.4 will appear.
2. Enter a value from 10 – 3600 in the Hello Time field.
3. Select the Apply button to assign the change.

Modifying Prune Timer

This parameter defines how long a PIM router will maintain a prune state for a forwarding entry.

The first received multicast interface is forwarded to all other PIM interfaces on the routing switch. If there is no presence of groups on that interface, the leaf node sends a prune message upstream and stores a prune state. This prune state travels up the tree and installs a prune state.

A prune state is maintained until the prune timer expires or a graft message is received for the forwarding entry. The default value is 180 seconds.

USING THE CLI

To set the PIM prune timer to 90, enter the following:

```
HP9300(config)# router pim
HP9300(config-pim-router)# pim-prune-timer 90
```

syntax: pim-prune-timer <10-3600>

USING THE WEB MANAGEMENT INTERFACE

1. Select the PIM link from the main menu. The PIM configuration panel shown in Figure 11.4 will appear.
2. Enter a value from 10 – 3600 into the Prune Time field.
3. Select the Apply button to assign the change.

Modifying Graft Retransmit Timer

The Graft Retransmit Timer defines the interval between the transmission of graft messages.

A graft message is sent by a router to cancel a prune state. When a router receives a graft message, the router responds with a Graft Ack (acknowledge) message. If this Graft Ack message is lost, the router that sent the graft message will resend it.

USING THE CLI

To change the graft retransmit timer from the default of 180 to 90 seconds, enter the following:

```
HP9300(config)# router pim
HP9300(config-pim-router)# pim-graft-retransmit-timer 90
```

syntax: pim-graft-retransmit-timer <10-3600>

USING THE WEB MANAGEMENT INTERFACE

1. Select the PIM link from the main menu. The PIM configuration panel shown in Figure 11.4 will appear.
2. Enter a value from 10 – 3600 into the Graft Retransmit Time field.
3. Select the Apply button to assign the change.

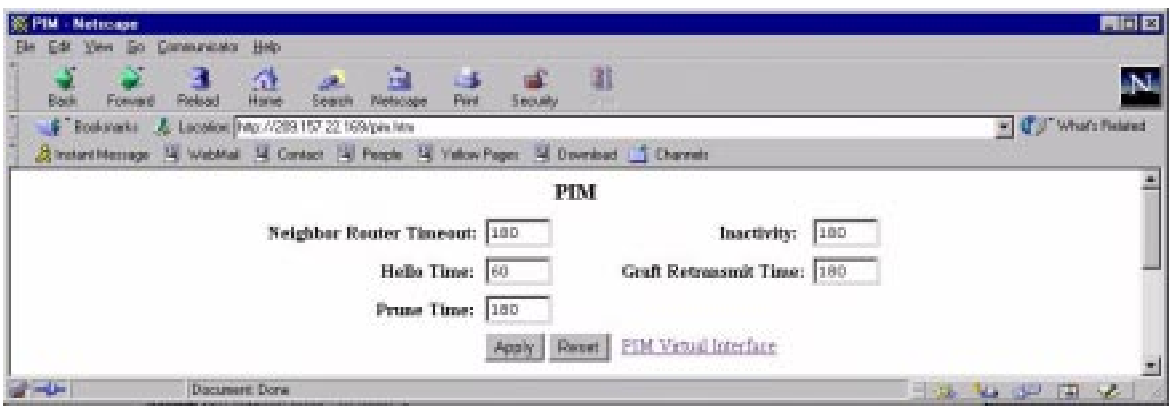

Figure 11.4 PIM configuration sheet

Modifying Inactivity Timer

The routing switch deletes a forwarding entry if the entry is not used to send multicast packets. The PIM inactivity timer defines how long a forwarding entry can remain unused before the routing switch deletes it.

USING THE CLI

To apply a PIM inactivity timer of 90 seconds to all PIM interfaces, enter the following:

```
HP9300(config)# router pim

HP9300(config-pim-router)# pim-inactivity-timer 90
```

syntax: pim-inactivity-timer <10-3600>

USING THE WEB MANAGEMENT INTERFACE

1. Select the PIM link from the main menu. The PIM configuration panel shown in Figure 11.4 will appear.
2. Enter a value from 10 – 3600 into the Inactivity field.
3. Select the Apply button to assign the change.

Modifying PIM Interface Parameters

TTL is the only interface parameter for PIM. TTL defines the minimum value required in a packet for it to be forwarded out of the interface.

For example, if the TTL for an interface is set at 10, it means that only those packets with a TTL value of 10 or more will be forwarded. Likewise, if an interface is configured with a TTL Threshold value of 1, all packets received on that interface will be forwarded. Possible TTL values are 1 to 64. The default TTL value is 1.

USING THE CLI

To configure a TTL of 45, enter the following:

```
HP9300(config-pim-router)# ip pim ttl 45
```

syntax: ip pim ttl <1-64>

USING THE WEB MANAGEMENT INTERFACE

To modify the PIM parameter (TTL) for an interface:

1. Select the PIM Virtual Interface link from the PIM configuration sheet. The panel show in Figure 11.5 will appear.
2. Select the type of interface that is being configured from the pull down menu. You can select IP tunnel or sub-net.
3. Select the IP address of the interface being configured from the Local Address pulldown menu.

NOTE: If you are configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the Remote Address field. IP tunneling must also be enabled and defined on the destination router interface as well. Note that this field is left blank when configuring sub-net type interfaces.

4. Enter a value from 1 – 64 in the Time To Live Threshold (TTL) field.
5. Select the Add button to enable PIM on the interface.

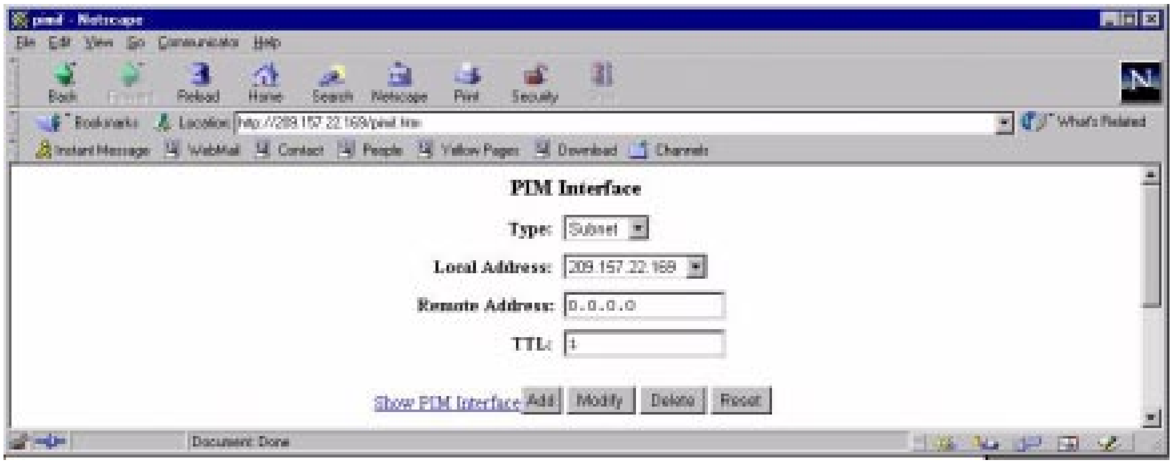

Figure 11.5 PIM interface panel

DVMRP Overview

The 9304M, 9308M, and 6308M-SX routing switches provide multicast routing with the ***Distance Vector Multicast Routing Protocol (DVMRP)*** routing protocol. DVMRP uses ***Internet Group Membership Protocol (IGMP)*** to manage the IP multicast groups.

DVMRP is a broadcast and pruning multicast protocol that delivers IP multicast datagrams to its intended receivers. The receiver registers the interested groups using IGMP. DVMRP builds a multicast delivery tree with the sender forming the root. Initially, multicast datagrams are delivered to all nodes on the tree. Those leaves that do not have any group members send ***prune messages*** to the upstream router, noting the absence of a group. The upstream router maintains a prune state for this group for the given sender. A prune state is aged out after a given configurable interval, allowing multicasts to resume.

DVMRP employs ***reverse path forwarding*** and ***pruning*** to keep source specific multicast delivery trees with the minimum number of branches required to reach all group members. DVMRP builds a multicast tree for each source and destination host group.

Initiating DVMRP Multicasts on a Network

Once DVMRP is enabled on each router, a network user can begin a video conference multicast from the server on R1. ***Multicast Delivery Trees*** are initially formed by source-originated multicast packets that are propagated to downstream interfaces as seen in Figure 11.6. When a multicast packet is received on a DVMRP-capable router interface, the interface checks its DVMRP routing table to determine whether the interface that received the message provides the shortest path back to the source. If the interface does provide the shortest path, the interface forwards the multicast packet to adjacent peer DVMRP routers, except for the router interface that originated the packet. Otherwise, the interface discards the multicast packet and sends a prune message back upstream. This process is known as ***reverse path forwarding***.

In Figure 11.6, the root node (R1) is forwarding multicast packets for group 229.225.0.2 that it receives from the server to its downstream nodes, R2, R3, and R4. Router R4 is an intermediate router with R5 and R6 as its downstream routers. Because R5 and R6 have no downstream interfaces, they are leaf nodes.

The receivers in this example are those workstations that are resident on routers R2, R3, and R6.

Pruning a Multicast Tree

After the multicast tree is constructed, ***pruning*** of the tree will occur after IP multicast packets begin to traverse the tree.

As multicast packets reach leaf networks (sub-nets with no downstream interfaces), the local IGMP database checks for the recently arrived IP multicast packet address. If the local database does not contain the address (the address has not been learned), the router prunes (removes) the address from the multicast tree and no longer receives multicasts until the prune age expires.

In Figure 11.7, Router 5 is a leaf node with no group members in its local database. Consequently, Router 5 sends a prune message to its upstream router. This router will not receive any further multicast traffic until the prune age interval expires.

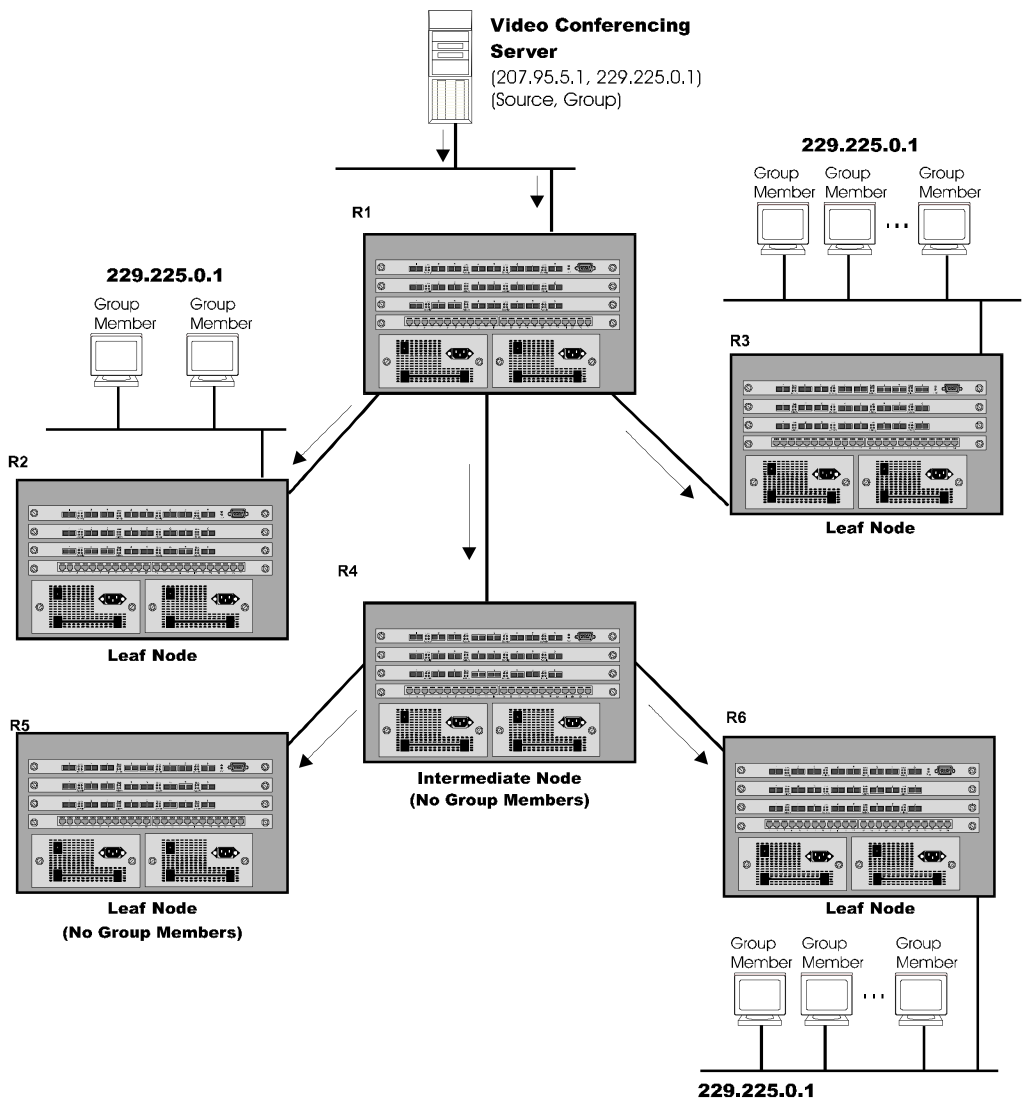

Figure 11.6 Downstream broadcast of IP multicast packets from source host

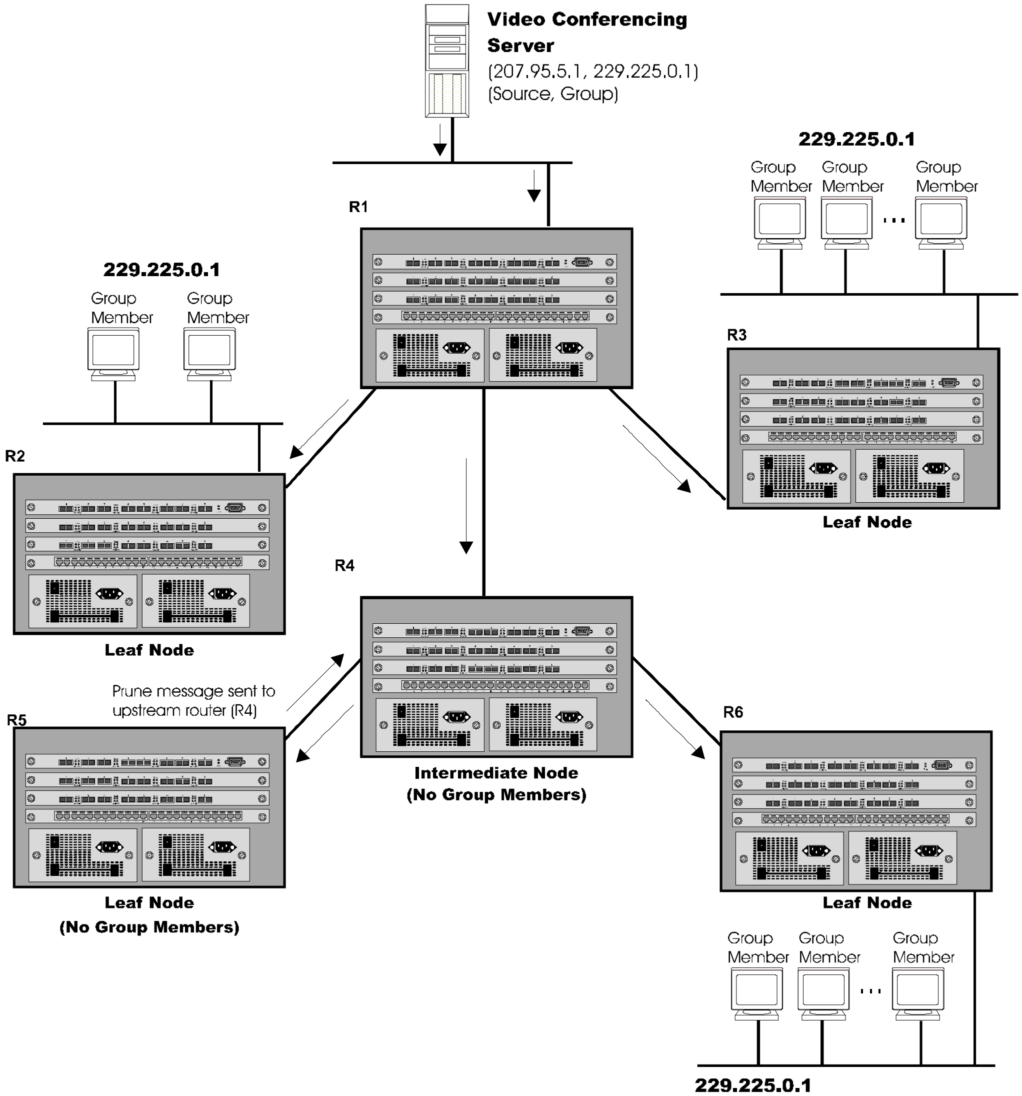

Figure 11.7 Pruning leaf nodes from a multicast tree

Grafts to a Multicast Tree

A DVMRP router restores pruned branches to a multicast tree by sending graft messages towards the upstream router. Graft messages start at the leaf node and travel up the tree, first sending the message to its neighbor upstream router.

In the example above, if a new 229.255.0.1 group member joins on router R6, which had been pruned previously, a graft will be sent upstream to R4. Since the forwarding state for this entry is in a prune state, R4 sends a graft to R1. Once R4 has joined the tree, it along with R6 will once again receive multicast packets.

You do not need to perform any configuration to maintain the multicast delivery tree. The prune and graft messages automatically maintain the tree.

Configuring DVMRP

Enabling DVMRP on the Routing Switch and Interface

Suppose you want to initiate the use of desktop video for fellow users on a sprawling campus network. All destination workstations have the appropriate hardware and software but the routers that connect the various buildings need to be configured to support DVMRP multicasts from the designated video conference server as seen in Figure 11.6.

DVMRP is enabled on each of the routers shown in Figure 11.6, on which multicasts are expected. You can enable DVMRP on each router independently or remotely from one of the routers by a Telnet connection. Follow the same steps for each router. A reset of the router is required when DVMRP is first enabled. Thereafter, all changes are dynamic.

NOTE: By default, the DVMRP feature is disabled. To enable DVMRP on router1, enable DVMRP at the global level and then on each interface that will support the protocol.

USING THE CLI

To enable DVMRP on Router 1 and interface 3, enter the following:

```
HP9300(config)# router dvmrp

HP9300(config)# int e 3

HP9300(config-if-3)# ip dvmrp
```

USING THE WEB MANAGEMENT INTERFACE

To enable DVMRP on Router 1 and interface 3, enter the following:

1. Select the System link from the main menu. The System configuration panel will appear.
2. Enable DVMRP.
3. Select the Apply button to assign the change.
4. Select the DVMRP link from the main menu to display the DVMRP configuration sheet, shown in Figure 11.8.
5. Select the DVMRP Virtual Interface link. The panel seen in Figure 11.9 will appear.
6. Select the type of interface to be configured (Sub-net or IP Tunnel). In this case, select Sub-net.
7. Enter the IP address of the interface in the Local Address field.
8. Modify the TTL and metric values is desired.
9. Enable local advertisement of routes and encapsulation if desired.
10. Select the Add button to assign the changes.
11. Repeat steps 5 – 9 for each interface you want to use for DVMRP.

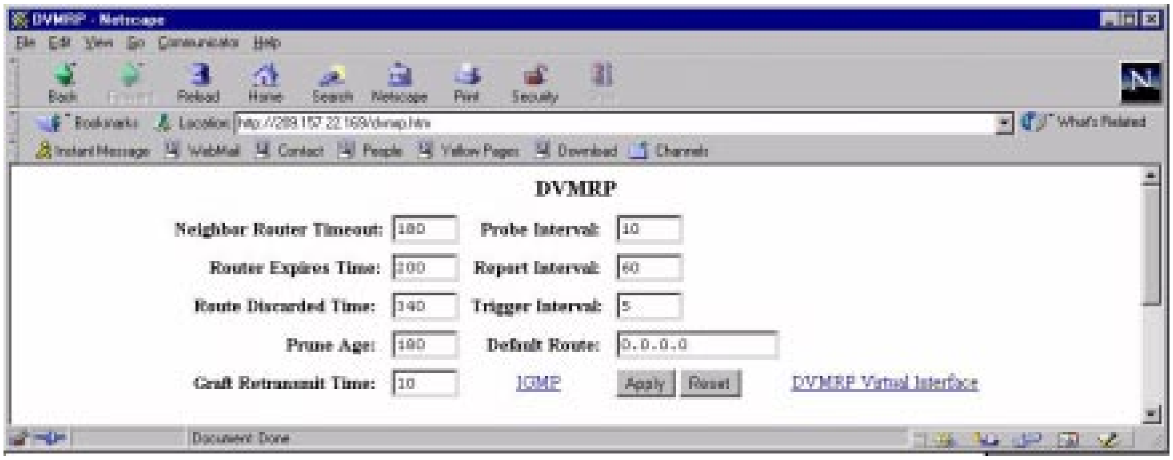

Figure 11.8 DVMRP configuration sheet

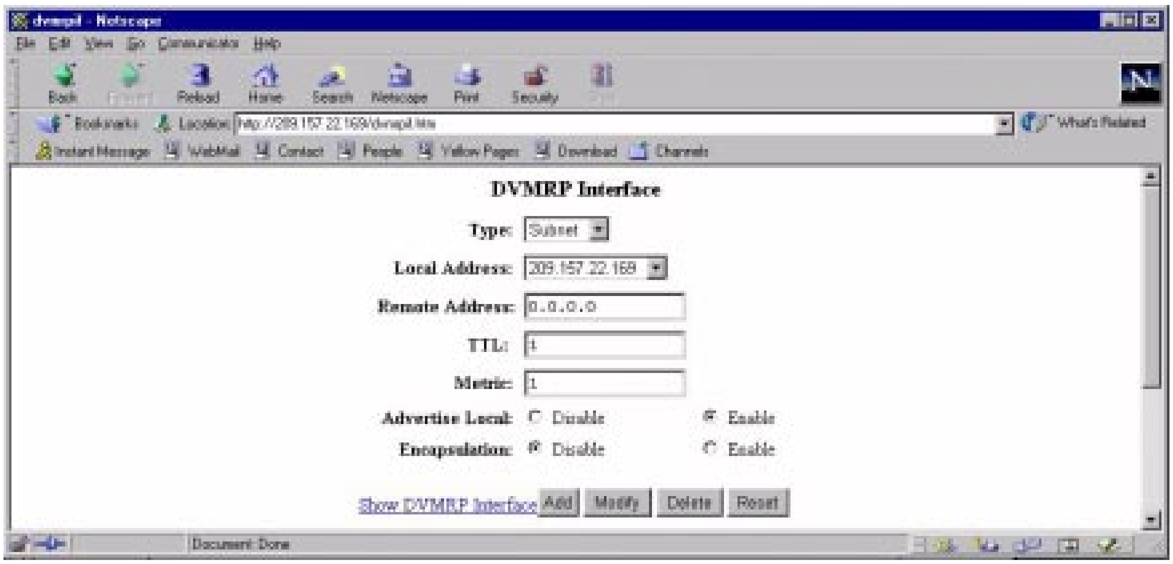

Figure 11.9 DVMRP virtual interface panel

Modifying DVMRP Global Parameters

DVMRP global parameters come with preset values. The defaults work well in most networks, but you can modify the following global parameters if you need to:

- Neighbor router timeout
- Route expire time
- Route discard time
- Prune age
- Graft retransmit time
- Probe interval
- Report interval
- Trigger interval
- Default route

Modifying Neighbor Router Timeout

The Neighbor Route Timeout specifies the period of time that a router will wait before it defines an attached DVMRP neighbor router as down. Possible values are 40 – 8000 seconds. The default value is 180 seconds.

USING THE CLI

To modify the neighbor timeout value to 100, enter the following:

```
HP9300(config-dvmrp-router)# nbr 100
```

syntax: nbr-timeout <40-8000>

USING THE WEB MANAGEMENT INTERFACE

1. Select DVMRP from the main menu bar to display the DVMRP configuration sheet as shown in Figure 11.8.
2. Enter a value from 40 – 8000 for neighbor router timeout.
3. Select the Apply button to assign the changes.

Modifying Route Expires Time

The Route Expire Time defines how long a route is considered valid in the absence of the next route update. Possible values are from 20 – 4000 seconds. The default value is 200 seconds.

USING THE CLI

To modify the route expire setting to 50, enter the following:

```
HP9300(config-dvmrp-router)# route-exp 50
```

syntax: route-expire <20-4000>

USING THE WEB MANAGEMENT INTERFACE

1. Select DVMRP from the main menu bar to display the DVMRP configuration sheet as shown in Figure 11.8.
2. Enter a value from 40 – 8000 in the Route Expire Time field.
3. Select the Apply button to assign the changes.

Modifying Route Discard Time

The Route Discard Time defines the period of time before a route is deleted. Possible values are from 40 – 8000 seconds. The default value is 340 seconds.

USING THE CLI

To modify the route discard setting to 150, enter the following:

```
HP9300(config-dvmrp-router)# route dis 150
```

syntax: route-discard <40-8000>

USING THE WEB MANAGEMENT INTERFACE

1. Select DVMRP from the main menu bar to reach the DVMRP configuration sheet as shown in Figure 11.8.
2. Enter a value from 40 – 8000 in the Route Discard Time field.
3. Select the Apply button to assign the changes.

Modifying Prune Age

The Prune Age defines how long a prune state will remain in effect for a source-routed multicast tree. After the prune age period expires, flooding will resume. Possible values are from 20 – 3600 seconds. The default value is 180 seconds.

USING THE CLI

To modify the prune age setting to 150, enter the following:

```
HP9300(config-dvmrp-router)# prune 25
```

syntax: prune-age <20-3600>

USING THE WEB MANAGEMENT INTERFACE

1. Select DVMRP from the main menu bar to display the DVMRP configuration sheet as shown in Figure 11.8.
2. Enter a value from 20 – 3600 in the Prune Age field.
3. Select the Apply button to assign the changes.

Modifying Graft Retransmit Time

The Graft Retransmit Time defines the initial period of time that a router sending a graft message will wait for a graft acknowledgement from an upstream router before re-transmitting that message.

Subsequent retransmissions are sent at an interval twice that of the preceding interval. Possible values are from 5 – 3600 seconds. The default value is 10 seconds.

USING THE CLI

To modify the setting for graft retransmit time to 120, enter the following:

```
HP9300(config-dvmrp-router)# graft 120
```

syntax: graft-retransmit-time <5-3600>

USING THE WEB MANAGEMENT INTERFACE

1. Select DVMRP from the main menu bar to display the DVMRP configuration sheet as shown in Figure 11.8.
2. Enter a value from 5 – 3600 in the Graft Retransmit Time field.
3. Select the Apply button to assign the changes.

Modifying Probe Interval

The Probe Interval defines how often neighbor probe messages are sent to the ALL-DVMRP-ROUTERS IP multicast group address. A router’s probe message lists those neighbor DVMRP routers from which it has received probes. Possible values are from 5 – 30 seconds. The default value is 10 seconds.

USING THE CLI

To modify the probe interval setting to 10, enter the following:

```
HP9300(config-dvmrp-router)# probe 10
```

syntax: probe-interval <5-30>

USING THE WEB MANAGEMENT INTERFACE

1. Select DVMRP from the main menu bar to display the DVMRP configuration sheet as shown in Figure 11.8.
2. Enter a value from 5 – 30 in the Probe Interval field.
3. Select the Apply button to assign the changes.

Modifying Report Interval

The Report Interval defines how often routers propagate their complete routing tables to other neighbor DVMRP routers. Possible values are from 10 – 2000 seconds. The default value is 60 seconds.

USING THE CLI

To support propagation of DVMRP routing information to the network every 90 seconds, enter the following:

```
HP9300(config-dvmrp-router)# report 90
```

syntax: report-interval <10-2000>

USING THE WEB MANAGEMENT INTERFACE

1. Select DVMRP from the main menu bar to display the DVMRP configuration sheet as shown in Figure 11.8.
2. Enter a value from 10 – 2000 in the Report Interval field.
3. Select the Apply button to assign the changes.

Modifying Trigger Interval

The Trigger Interval defines how often trigger updates, which reflect changes in the network topology, are sent. Example changes in a network topology include router up or down or changes in the metric. Possible values are from 5 – 30 seconds. The default value is 5 seconds.

USING THE CLI

To support the sending of trigger updates every 20 seconds, enter the following:

```
HP9300(config-dvmrp-router)# trig 20
```

Syntax: trigger interval <5-30>

USING THE WEB MANAGEMENT INTERFACE

1. Select DVMRP from the main menu bar to display the DVMRP configuration sheet as shown in Figure 11.8.
2. Enter a value from 5 – 30 in the Trigger Interval field.
3. Select the Apply button to assign the changes.

Modifying Default Route

This defines the default gateway for IP multicast routing.

USING THE CLI

To define the default gateway for DVMRP, enter the following:

```
HP9300(config-dvmrp-router)# default-gateway 192.35.4.1
```

syntax: default-gateway <ip address>

USING THE WEB MANAGEMENT INTERFACE

1. Select DVMRP from the main menu bar to display the DVMRP configuration sheet as shown in Figure 11.8.
2. Enter the IP address of the default gateway Default Route field.
3. Select the Apply button to assign the changes.

Modifying DVMRP Interface Parameters

DVMRP global parameters come with preset values. The defaults work well in most networks, but you can modify the following interface parameters if you need to:

- TTL
- Metric
- Advertising
- Encapsulation

The following paragraphs provide an overview and configuration details for DVMRP global parameters.

Modifying the TTL

The TTL defines the minimum value required in a packet in order for the packet to be forwarded out the interface. For example, if the TTL for an interface is set at 10 it means that only those packets with a TTL value of 10 or more are forwarded. Likewise, if an interface is configured with a TTL Threshold value of 1, all packets received on that interface are forwarded. Possible values are from 1 – 64. The default value is 1.

USING THE CLI

To set a TTL of 64, enter the following:

```
HP9300(config)# int e 1/4
HP9300(config-if-1/4)# ip dvmrp ttl 60
```

syntax: ttl-threshold <1-64>

USING THE WEB MANAGEMENT INTERFACE

1. Select the <u>DVMRP Virtual Interface</u> link from the main menu. The panel shown in Figure 11.9 will appear.
2. Select the type of interface that is being configured from the pulldown menu. You can select IP Tunnel or Sub-net.
3. Select the IP address of the interface being configured from the Local Address pulldown menu.

NOTE: If you are configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the Remote Address field. IP tunneling must also be enabled and defined on the destination router interface as well. Note that this field is left blank when configuring sub-net type interfaces.

---

4. Enter a value from 1 – 64 in the Time To Live Threshold (TTL) field.
5. Select the Add button to assign the change.

Modifying the Metric

The routing switch uses the metric when establishing reverse paths to some networks on directly attached interfaces. Possible values are from 1 – 31 hops. The default is 1.

USING THE CLI

To set a metric of 15 for a DVMRP interface, enter the following:

```
HP9300(config)# interface 3/5
HP9300(config-if-3/5)# ip dvmrp metric 15
```

syntax: ip dvmrp metric <1-31>

USING THE WEB MANAGEMENT INTERFACE

1. Select the <u>DVMRP Virtual Interface</u> link from the main menu. The panel shown in Figure 11.9 appears.
2. Select the type of interface that is being configured from the pulldown menu. You can select IP tunnel or sub-net.
3. Select the IP address of the interface being configured from the Local Address pulldown menu.

NOTE: If you are configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the Remote Address field. IP tunneling must also be enabled and defined on the destination router interface as well. Note that this field is left blank when configuring sub-net type interfaces.

---

4. Enter a value from 1 – 31 in the Metric field.
5. Select the Add button to assign the change.

Enabling Advertising

You can turn the advertisement of a local route on (enable) or off (disable) on the interface. By default, advertising is enabled.

USING THE CLI

To enable advertising on an interface, enter the following:

```
HP9300(config-if-1/4)# ip dvmrp advertise-local on
```

syntax: advertise-local <on | off>

USING THE WEB MANAGEMENT INTERFACE

1. Select the DVMRP Virtual Interface link from the main menu. The panel shown in Figure 11.9 appears.
2. Select the type of interface that is being configured from the pulldown menu. You can select IP Tunnel or Sub-net.
3. Select the IP address of the interface being configured from the Local Address pulldown menu.

NOTE: If you are configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the Remote Address field. IP tunneling must also be enabled and defined on the destination router interface as well. Note that this field is left blank when configuring sub-net type interfaces.

---

4. Enable the Advertise Local field.
5. Click the Apply button to assign the changes.

Enabling Encapsulation

Encapsulation varies based on the interface type. For type “IP tunnel”, DVMRP control messages such as probe and route report are encapsulated within the IP packet. For type “sub-net”, the IP data is encapsulated within an IP packet. Encapsulation is disabled by default.

USING THE CLI

To enable and define encapsulation type for DVMRP, enter the following:

```
HP9300(config)# int e 1/6
HP9300(config-if-1/6)# ip dvmrp encap ethernet-2
```

syntax: ip dvmrp encapsulation <ethernet-2 | snap>

USING THE WEB MANAGEMENT INTERFACE

1. Select the DVMRP Virtual Interface link from the main menu. The panel shown in Figure 11.9 appears.
2. Select the type of interface that is being configured from the pulldown menu. Select IP Tunnel or Sub-net.
3. Select the IP address of the interface being configured from the Local Address pulldown menu.

NOTE: If you are configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the Remote Address field. IP tunneling must also be enabled and defined on the destination router interface as well. Note that this field is left blank when configuring sub-net type interfaces.

---

4. Enable the Advertise Local field.
5. Click the Apply button to assign the changes.

Configuring an IP Tunnel

IP tunnels are used to send traffic through routers that do not support PIM or DVMRP multicasting. IP multicast datagrams are encapsulated within an IP packet and then sent to the remote address. Routers that are not configured for PIM or DVMRP route that packet as a normal IP packet. When the DVMRP or PIM router at the remote end of the tunnel receives the packet, the router strips off the IP encapsulation and forwards the packet as an IP Multicast packet.

NOTE: An IP tunnel must have a remote IP interface at each end. Also, for IP tunneling to work, the remote routers must be reachable by an IP routing protocol.

NOTE: Multiple tunnels configured on a routing switch cannot share the same remote address.

EXAMPLE: To configure an IP tunnel as seen in Figure 11.10, enter the IP tunnel destination address on an interface of the routing switch.

USING THE CLI

To configure an IP address on Router A, enter the following:

```
HP9300(config)# int e1/1
HP9300(config-if-1/1)# ip tunnel 192.3.45.6
```

NOTE: The IP tunnel address represents the configured IP tunnel address of the destination router. In the case of Router A, its destination router is Router B. Router A is the destination router of Router B.

For router B, enter the following:

```
HP9300(config-if-1/1)# ip tunnel 192.58.4.1
```

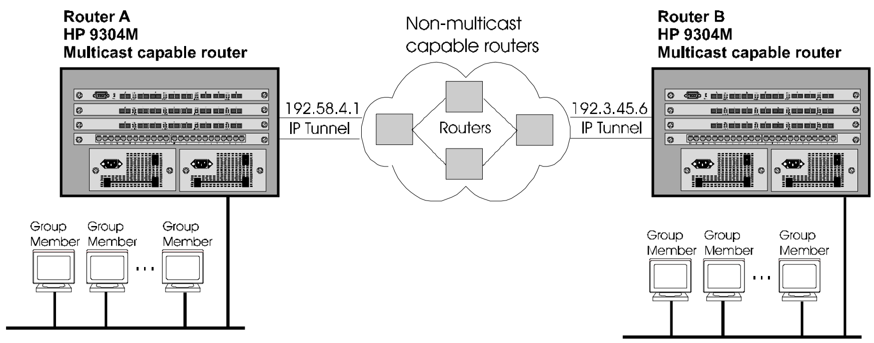

Figure 11.10 IP in IP tunneling on multicast packets in a unicast network

USING THE WEB MANAGEMENT INTERFACE

1. Select the DVMRP Virtual Interface link from the main menu. The panel shown in Figure 11.9 appears.

NOTE: IP tunnels can be configured for PIM and DVMRP. For example, if you user are configuring an IP tunnel for PIM, select the PIM virtual interface link instead in step 1 above.

---

2. Select IP tunnel as the type.
3. Select the IP address of the interface being configured from the Local Address pulldown menu.
4. Enter the IP address of the destination interface, the end point of the IP Tunnel, in the Remote Address field.
5. Modify the remaining parameters as required.

NOTE: For DVMRP, you can set a TTL value and enable advertisement of routes and encapsulation on the IP tunnel. For PIM, you can set a TTL value.

---

6. Select the Add button.
7. Repeat steps 1 – 6 for the interface on the remote end of the IP tunnel.