Chapter 10
Configuring IP Multicast Protocols

This chapter covers how to configure Protocol Independent Multicast (PIM) and Distance Vector Multicast Routing Protocol (DVMRP), two IP multicast protocols supported on the HP 9304M and 9308M routing switches.

A summary of all CLI commands discussed in this chapter can also be found in **Appendix B**.

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 multicast routing.

HP 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. They 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 an HP 9304M or 9308M routing switch.

Multicast Terms

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

Node: Refers to a router.

Root Node: This is the node which 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 defined as 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: Those routers which have downstream interfaces in the path of the end points, the source and leaf routers.

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

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

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. It is similar to DVMRP in that it 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, the user can begin a video conference multicast from the server on R1. When a multicast packet is received on a PIM-capable router interface, it will check its IP routing table to determine if the interface that received the message provides the shortest path back to the source. If it does represent the shortest path, 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 10.1**, 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

As multicast packets reach these leaf routers, the routers check their IGMP database for the group. If the group is not in the IGMP database, the packet is discarded and a prune message is sent to the upstream router and a prune state is maintained for that source, group (S,G) pair. The branch will then be removed or pruned 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. The length of time that a prune state is considered valid is a configurable parameter (PIM Prune Timer) on the HP routing switch.

For example in **Figure 10.1** the sender with address 207.95.5.1 is sending multicast packets to the group 229.225.0.1. Should any groups other than that group be received at the PIM routers, it will be discarded and a prune message sent upstream.

In **Figure 10.2**, Router R5 is determined to be a leaf node with no group members in its IGMP database and it must be pruned from the multicast tree. Therefore, 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 10.2**. 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, should both R5 and R6 be in a prune state at the same time, then R4 would in effect become a leaf node with no downstream interfaces and would send a prune message to R1. With R4 in a prune state, the resulting multicast delivery tree would be just leaf nodes R2 and R3.

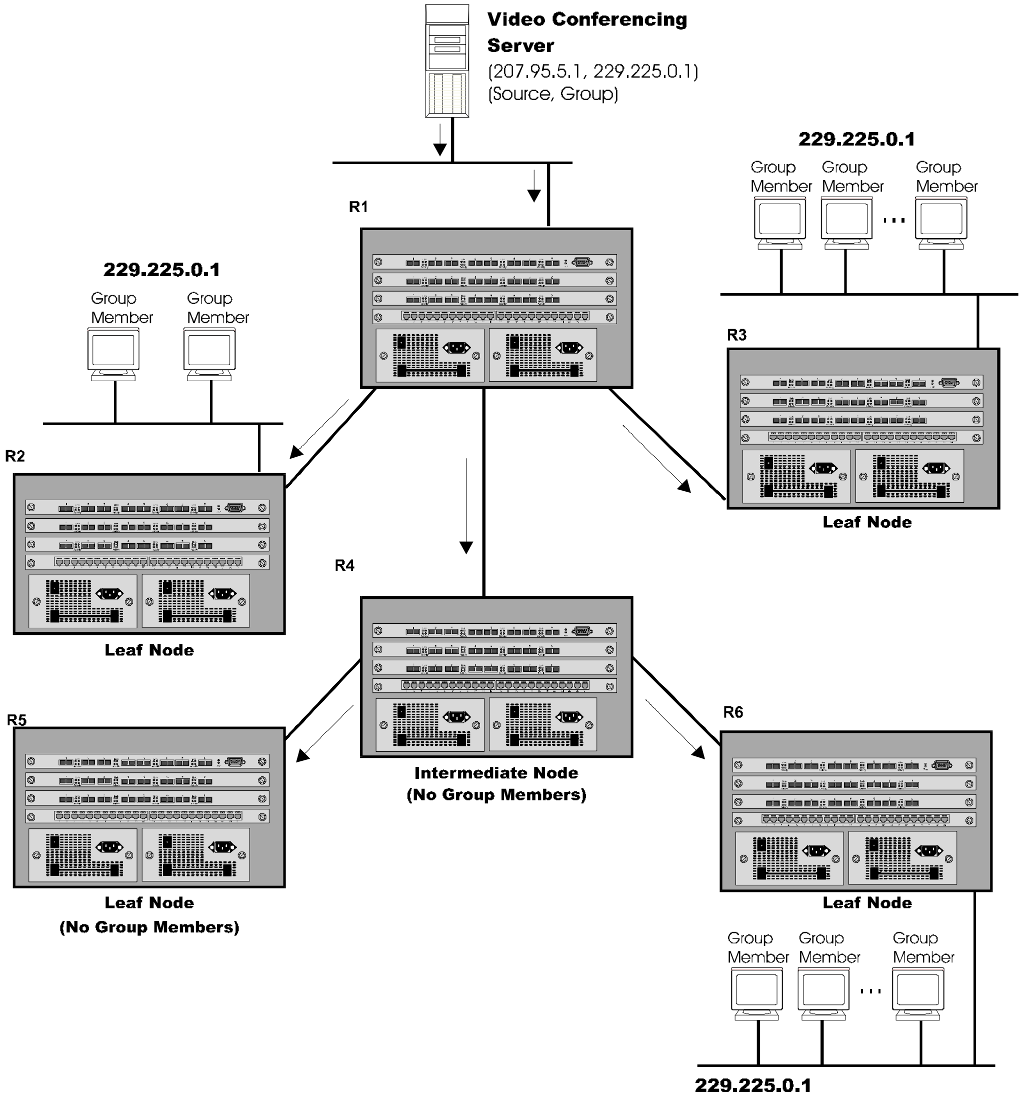

Figure 10.1 Transmission of multicast packets from the source to host group members

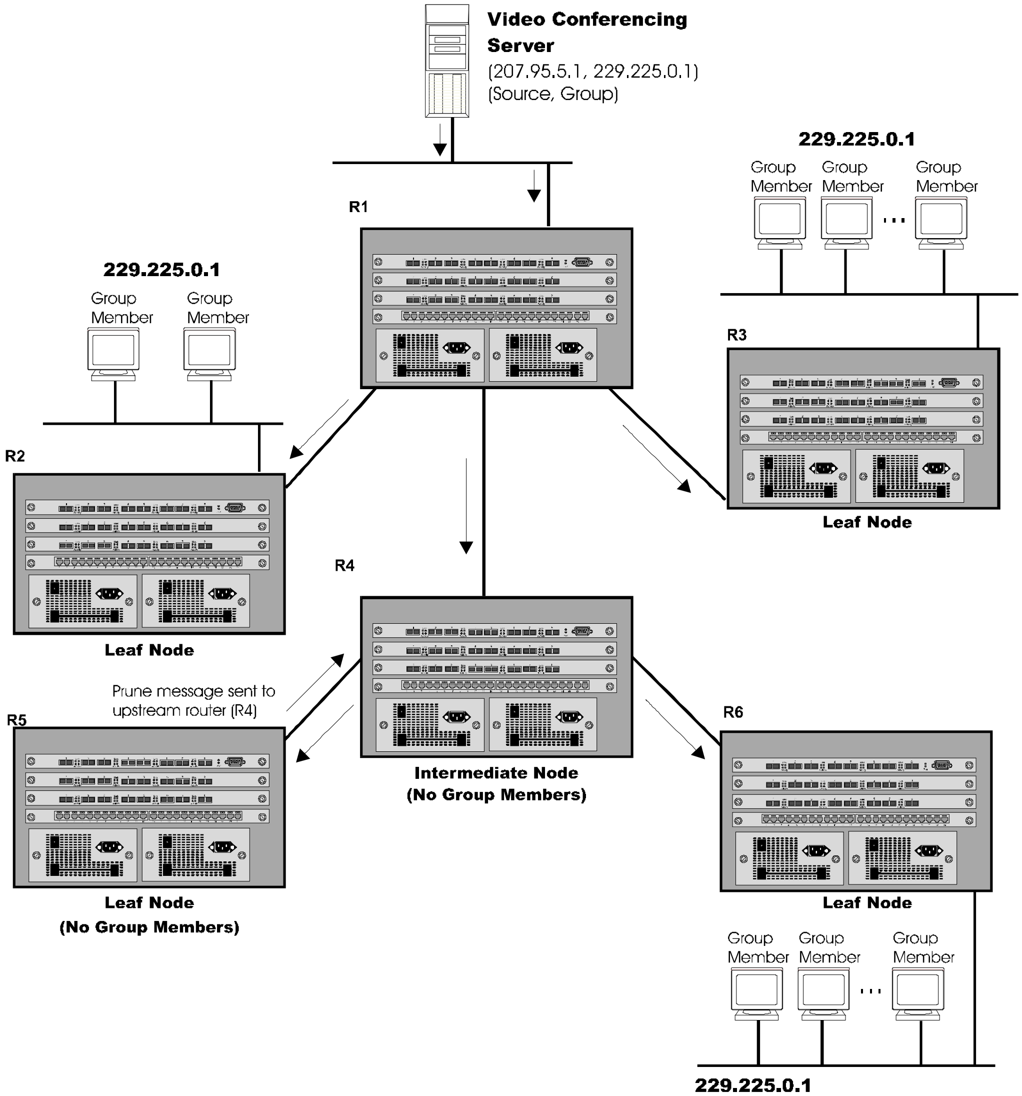

Figure 10.2 Pruning leaf nodes from a multicast tree

Grafts to a Multicast Tree

Pruned branches are restored 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 will be sent upstream to R4. Since the forwarding state for this entry is in a prune state, R4 will send a graft to R1. Once R4 has joined the tree, it along with R6 will once again receive multicast packets.

Prune and graft messages are continuously used to maintain the multicast delivery tree. No configuration is required of the user.

Configuring PIM

Enabling PIM on the Router and an Interface

PIM must be enabled at both the global and interface levels. PIM should be enabled on each interface that PIM routing is desired.

EXAMPLE: A user wants to initiate the use of desktop video for fellow users on a sprawling campus. All destination workstations have the appropriate hardware and software but the HP routing switch that connect the various buildings need to be configured to support PIM multicasts from the designated video conference server as seen in **Figure 10.1**.

PIM is enabled on each of the HP routing switches shown in **Figure 10.1**, on which multicasts are expected. The user can enable PIM on each router independently or remotely from one of the routers with a Telnet connection. The user would follow the same steps for each router. A reset of the router is required when PIM is first enabled. Thereafter, all changes will be dynamic. By default, the PIM feature is disabled.

USING THE CLI

EXAMPLE: To enable PIM on router1 and interface 3 (slot 2), the user would enter the following:

```
Router1 (config)# router pim

Router1 (config)# int e 2/3

Router1 (config-if-2/3)# 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 virtual interface link from the PIM configuration panel to configure an interface. The panel shown in **Figure 10.4** will appear.
5. Select the type of interface that is being configured from the pull down menu—**IP tunnel** or **sub-net**.
6. Select the IP address of the interface being configured from the **local address** pull down menu.

NOTE: If configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the **remote address** field of the screen. 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.

7. 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 to 64. The default value is 1.

8. Select the **add** button to enable PIM on the interface.

Modifying PIM Global Parameters

PIM global parameters come with preset values. The user does not need to modify the values unless network requirements deem the changes.

The following parameters can be modified for PIM:

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

Modifying Neighbor Timeout

Neighbor timeout is the interval after which a PIM capable router will consider a neighbor to not be present. Absence of PIM hello messages from a neighboring router is what will indicate 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 router 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 10.3** will appear.
2. Enter a value between 10 and 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 router operating with PIM, the user would 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 10.3** will appear.
2. Enter a value between 10 and 3600 into the **hello time** field.
3. Select the **apply** button to assign the change.

Modifying Prune Timer

This parameter is used to define how long an HP 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 router. If there is no presence of groups on that interface, the leaf node will send a prune message upstream and store a prune state. This prune state will travel up the tree and install 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 define a PIM prune timer of 90, the user would 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 10.3** will appear.
2. Enter a value between 10 and 3600 into the **prune time** field.
3. Select the **apply** button to assign the change.

Modifying 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 it will respond with a Graft Ack message. If this Graft Ack message is lost, the router that sent the graft message, send it again. The interval between the transmission of the first and subsequent graft message is what is configurable with the PIM graft retransmit timer.

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 10.3** will appear.
2. Enter a value between 10 and 3600 into the **graft retransmit time** field.
3. Select the **apply** button to assign the change.

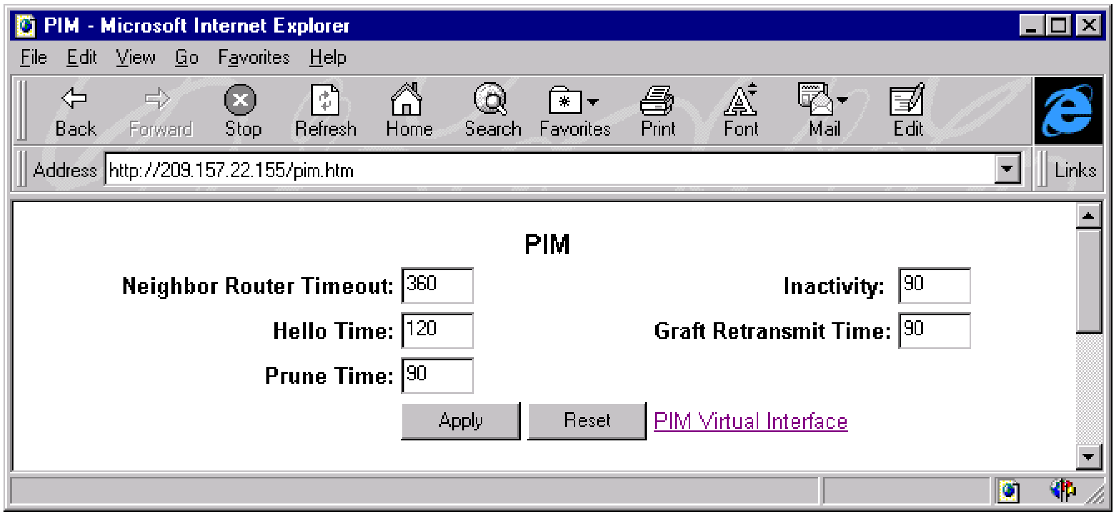

Figure 10.3 PIM configuration sheet

Modifying Inactivity Timer

A forwarding entry is deleted if it is not used to send multicast packets. The PIM inactivity timer defines the time interval after which an inactive forwarding entry is deleted.

USING THE CLI

To apply a PIM inactivity timer of 90 seconds to all PIM interfaces, the user would 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 10.3** will appear.
2. Enter a value between 10 and 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 values are 1 to 64. The default value is 1.

USING THE CLI

To configure a TTL of 45, the user would 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, the user would do the following:

1. Select the PIM virtual interface link from the PIM configuration sheet. The panel show in **Figure 10.4** will appear.
2. Select the type of interface that is being configured from the pull down menu—**IP tunnel** or **sub-net**.
3. Select the IP address of the interface being configured from the **local address** pull down menu.

NOTE: If configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the **remote address** field of the screen. 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 between 1 and 64 in the **time to live threshold (TTL)** field.
5. Select the **add** button to enable PIM on the interface.

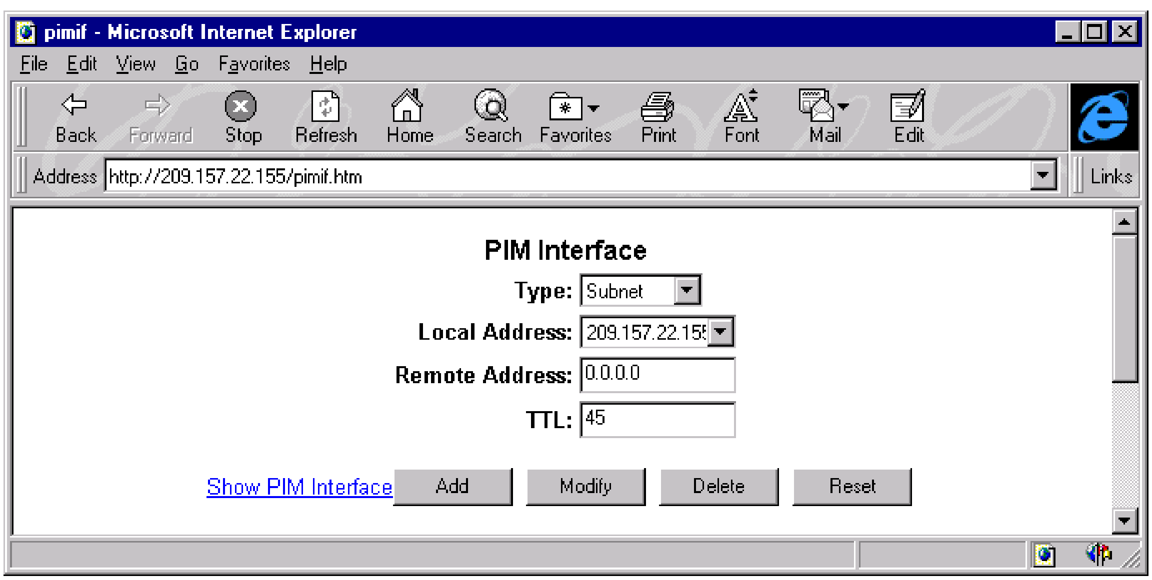

Figure 10.4 PIM interface panel

DVMRP Overview

HP routing switches provide multicast routing with the **Distance Vector Multicast Routing Protocol (DVMRP)** routing protocol and **Internet Group Membership Protocol (IGMP)** providing management of 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, the 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 10.5**. When a multicast packet is received on a DVMRP-capable router interface, it will check its DVMRP routing table to determine if the interface that received the message provides the shortest path back to the source. If it does represent the shortest path, the multicast packet is then forwarded to adjacent peer DVMRP routing switches, except for the router interface that originated the packet. Otherwise, the multicast packet is discarded and a prune message is sent back upstream. This process is known as **reverse path forwarding**.

In **Figure 10.5**, 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 routing switches. Because R5 and R6 have no downstream interfaces, they are leaf nodes.

The receivers in this example are those workstations that are resident on routing switches 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 (i.e. those sub-nets with no downstream interfaces), the local IGMP database checks for the recently arrived IP multicast packet address. If the address is not learned in the leaf router, then it is removed or **pruned** from the multicast tree and it will no longer receive multicasts until the prune age expires.

In **Figure 10.6**, Router 5 is determined to be 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.

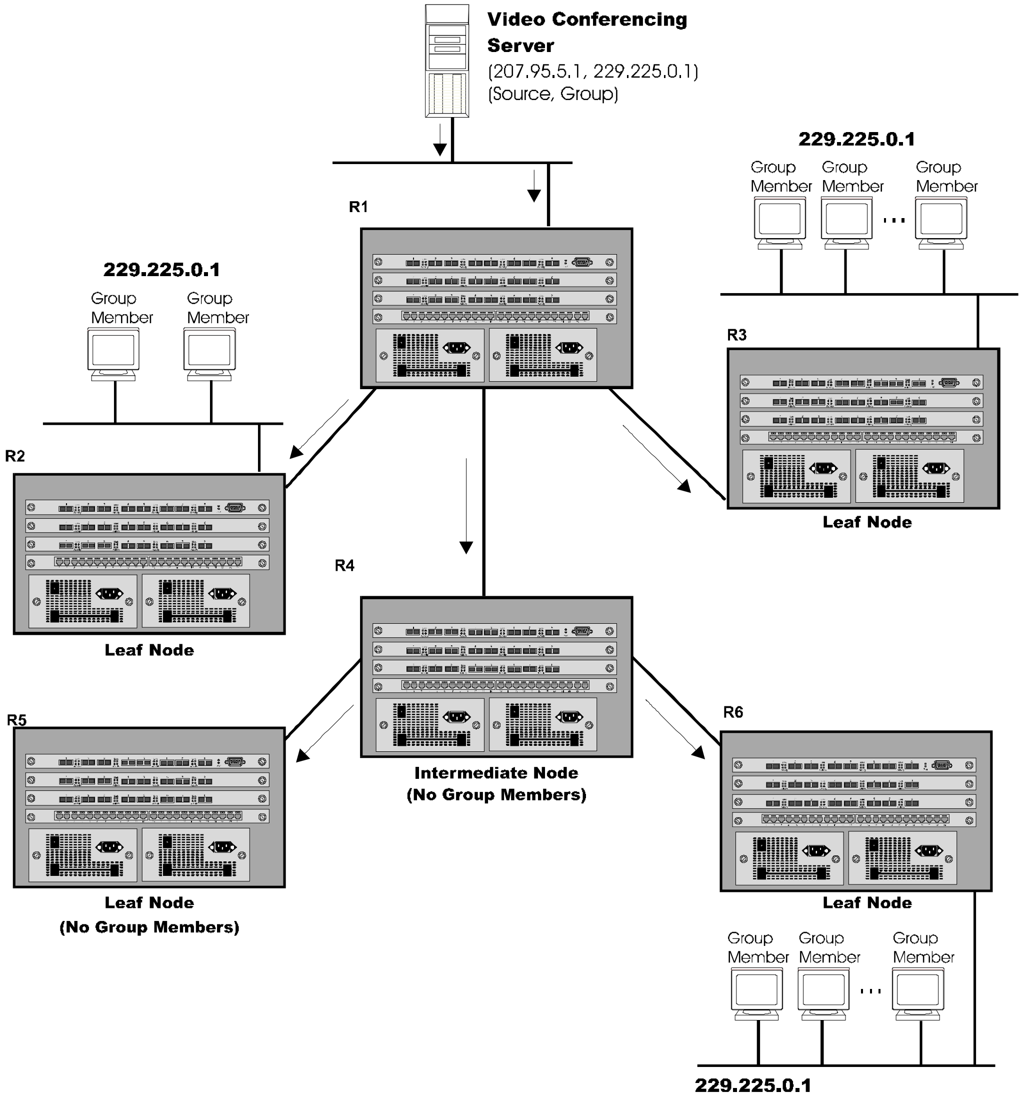

Figure 10.5 Downstream broadcast of IP multicast packets from source host

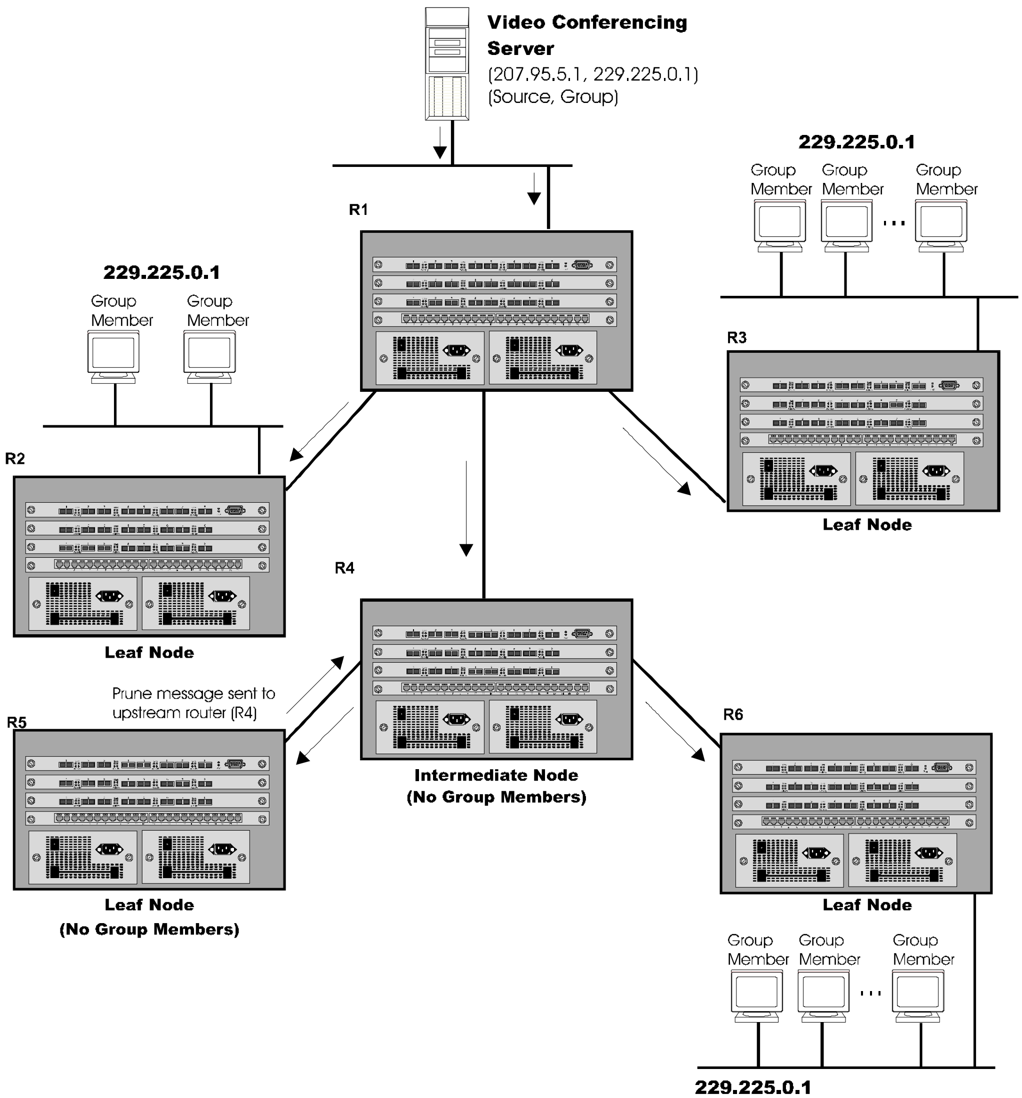

Figure 10.6 Pruning leaf nodes from a multicast tree

Grafts to a Multicast Tree

Pruned branches are restored 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 that had been previously pruned, a graft will be sent upstream to R4. Since the forwarding state for this entry is in a prune state, R4 will send a graft to R1. Once R4 has joined the tree, it along with R6 will once again receive multicast packets.

Prune and graft messages are continuously used to maintain the multicast delivery tree. No configuration is required of the user.

Configuring DVMRP

Enabling DVMRP on the Router and Interface

A user wants to initiate the use of desktop video for fellow users on a sprawling campus. All destination workstations have the appropriate hardware and software but the routing switches that connect the various buildings need to be configured to support DVMRP multicasts from the designated video conference server as seen in **Figure 10.5**

DVMRP is enabled on each of the HP routing switches shown in **Figure 10.5**, on which multicasts are expected. The user can enable DVMRP on each router independently or remotely from another HP 9304M or 9308M using a Telnet connection. The user would follow the same steps for each router. A reset of the router is required when DVMRP is first enabled. Thereafter all changes will by dynamic.

By default, the DVMRP feature is disabled.

NOTE: To enable DVMRP on router1, the user would enable DVMRP at the global level and then on each interface it is to be supported.

USING THE CLI

To enable DVMRP on Router 1 and interface 3 (slot 2), the user would enter the following:

```
Router1 (config)# router dvmrp

Router1 (config)# int e 2/3

Router1 (config-if-2/3)# ip dvmrp
```

USING THE WEB MANAGEMENT INTERFACE

To enable DVMRP on Router 1 and interface 3 (slot 2), the user would 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 virtual interface link. The panel seen in **Figure 10.8** will appear.
5. Select the **type** of interface to be configured, options are sub-net or IP tunnel. In this case, sub-net would be chosen.
6. Enter the IP address of the interface in the **local address** field.
7. Modify the **TTL** and **metric** values is desired.
8. Enable **local advertisement** of routes and **encapsulation** if desired.
9. Select the **add** button to assign the changes.
10. Repeat steps 5 through 9 for each interface that DVMRP is to be assigned.

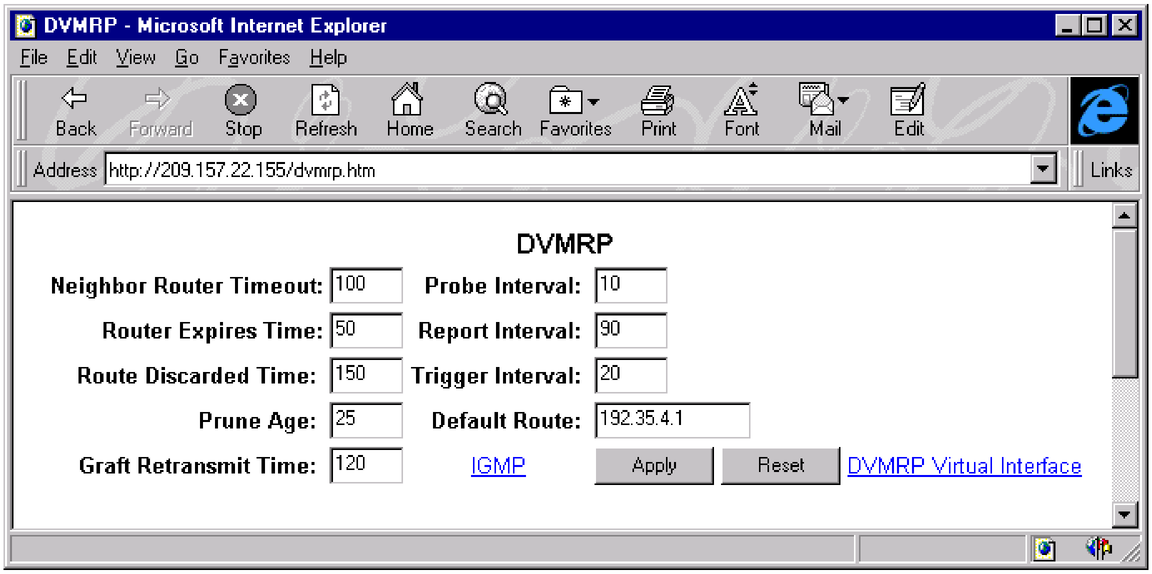

Figure 10.7 DVMRP configuration sheet

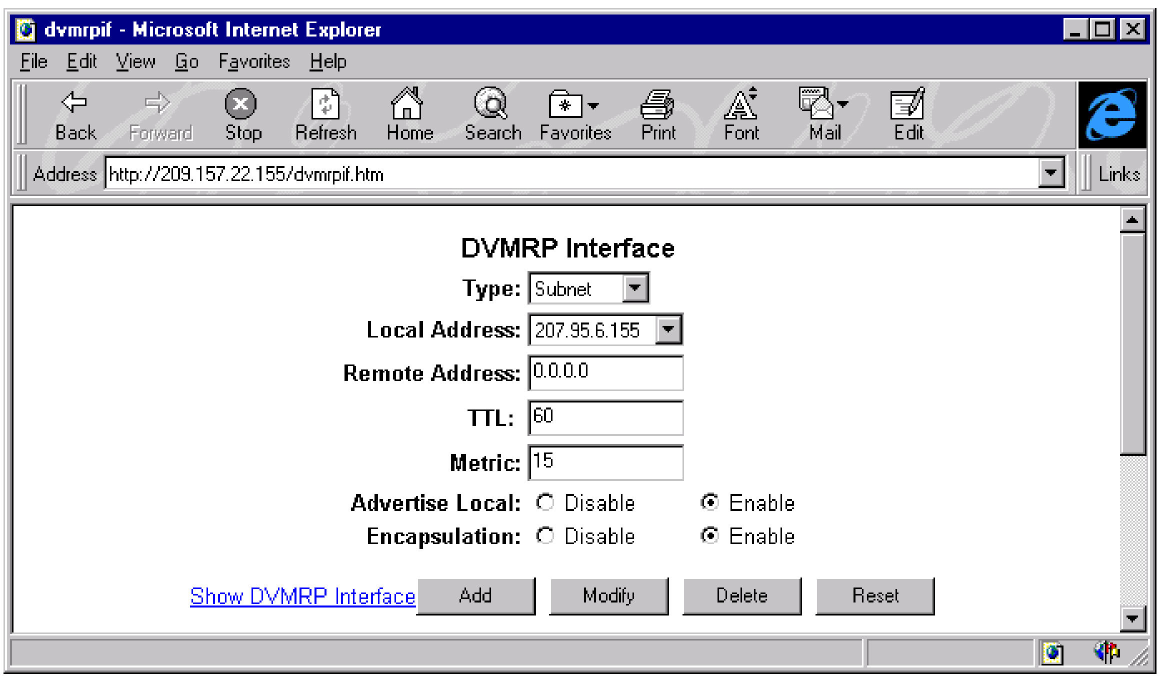

Figure 10.8 DVMRP virtual interface panel

Modifying DVMRP Global Parameters

All DVMRP parameters come with default settings. There is no need to modify these parameters unless network needs require it.

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

Modifying Neighbor Router Timeout

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

USING THE CLI

To modify the neighbor timeout value to 100, the user would 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 reach the DVMRP configuration sheet as seen in **Figure 10.7**.
2. Enter a value between 40 and 8000 for **neighbor router timeout**.
3. Select the **apply** button to assign the changes.

Modifying Route Expires Time

This parameter defines how long a route is considered valid in the absence of the next route update. Possible values are 20 to 4000 seconds. The default value is 200 seconds.

USING THE CLI

To modify the route expire setting to 50, the user would 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 reach the DVMRP configuration sheet as seen in **Figure 10.7**.
2. Enter a value between 40 and 8000 in the **route expire time** field.
3. Select the **apply** button to assign the changes.

Modifying Route Discard Time

Defines the period of time before a route is deleted. Possible values are 40 to 8000 seconds. The default value is 340 seconds.

USING THE CLI

To modify the route discard setting to 150, the user would 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 seen in **Figure 10.7**.
2. Enter a value between 40 and 8000 in the **route discard time** field.
3. Select the **apply** button to assign the changes.

Modifying Prune Age

This 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 20 to 3600 seconds. The default value is 180 seconds.

USING THE CLI

To modify the prune age setting to 150, the user would 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 reach the DVMRP configuration sheet as seen in **Figure 10.7**.
2. Enter a value between 20 and 3600 in the **prune age** field.
3. Select the **apply** button to assign the changes.

Modifying Graft Retransmit Time

This defines the initial period of time that a router sending a graft message will wait for a graft acknowledgment from an upstream router before re-transmitting that message.

Subsequent re-transmissions will be sent at an interval twice that of the preceding interval. Possible values are 5 to 3600 seconds. The default value is 10 seconds.

USING THE CLI

To modify the setting for graft retransmit time to 120, the user would 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 reach the DVMRP configuration sheet as seen in **Figure 10.7**.
2. Enter a value between 5 and 3600 in the **graft retransmit time** field.
3. Select the **apply** button to assign the changes.

Modifying Probe Interval

This 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 routing switches from which it has received probes. Possible values are 5 to 30 seconds. The default value is 10 seconds.

USING THE CLI

To modify the probe interval setting to 10, the user would 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 reach the DVMRP configuration sheet as seen in **Figure 10.7**.
2. Enter a value between 5 and 30 in the **probe interval** field.
3. Select the **apply** button to assign the changes.

Modifying Report Interval

This defines how often routing switches will propagate their complete routing tables to other neighbor DVMRP routing switches. Possible values are 10 to 2000 seconds. The default value is 60 seconds.

USING THE CLI

To support propagation of DVMRP routing information to the network every 90 seconds, the user would 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 reach the DVMRP configuration sheet as seen in **Figure 10.7**.
2. Enter a value between 10 and 2000 in the **report interval** field.
3. Select the **apply** button to assign the changes.

Modifying Trigger Interval

This 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 5 to 30 seconds. The default value is 5 seconds.

USING THE CLI

To support the sending of trigger updates every 20 seconds, the user would 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 reach the DVMRP configuration sheet as seen in **Figure 10.7**.
2. Enter a value between 5and 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, the user would 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 reach the DVMRP configuration sheet as seen in **Figure 10.7**.
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

All DVMRP interface parameters come with default settings. There is no need to modify these parameters unless network needs require it.

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

Modifying the TTL

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 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 to 64. The default value is 1.

USING THE CLI

To set a TTL of 64, the user would 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 DVMRP virtual interface link from the main menu. The panel seen in **Figure 10.8** will appear.
2. Select the type of interface that is being configured from the pull down menu—**IP tunnel** or **sub-net**.
3. Select the IP address of the interface being configured from the **local address** pull down menu.

NOTE: If configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the **remote address** field of the screen. 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 between 1 and 64 in the **time to live threshold (TTL)** field.
5. Select the **add** button to assign the change.

Modifying the Metric

Sets the default metric for a directly connected interface, when operating with DVMRP multicast. The router will use this in establishing the reverse paths to some networks. Possible values are 1 to 31 hops. The default is 1.

USING THE CLI

To set a metric of 15 for a DVMRP interface, the user would 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 DVMRP virtual interface link from the main menu. The panel seen in **Figure 10.8** will appear.
2. Select the type of interface that is being configured from the pull down menu—**IP tunnel** or **sub-net**.
3. Select the IP address of the interface being configured from the **local address** pull down menu.

NOTE: If configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the **remote address** field of the screen. 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 between 1 and 31 in the **metric** field.
5. Select the **add** button to assign the change.

Enabling Advertising

This parameter turns the advertisement of a local route on (enable) or off (disable) on the interface. By default the feature is enabled.

USING THE CLI

To enable advertising on an interface, the user would 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 seen in **Figure 10.8** will appear.
2. Select the type of interface that is being configured from the pull down menu—**IP tunnel** or **sub-net**.
3. Select the IP address of the interface being configured from the **local address** pull down menu.

NOTE: If configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the **remote address** field of the screen. 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.

Enabling Encapsulation

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

USING THE CLI

To enable and define encapsulation type for DVMRP, the user would 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 seen in **Figure 10.8** will appear.
2. Select the type of interface that is being configured from the pull down menu—**IP tunnel** or **sub-net**.
3. Select the IP address of the interface being configured from the **local address** pull down menu.

NOTE: If configuring an IP Tunnel, enter the IP address of the destination interface, the end point of the IP Tunnel, in the **remote address** field of the screen. 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.

Configuring an IP Tunnel

This feature is used to send traffic through routing switches that do not support PIM or DVMRP multicasting. IP multicast datagrams are encapsulated within an IP packet using IP in IP encapsulation and then sent to the remote address. When the destination router receives the encapsulated packet, it strips off the encapsulation and forwards them. An IP tunnel must be configured with a remote IP at each end of the tunnel.

NOTE: For IP tunneling to work, the remote routing switches must be reachable via an IP routing protocol.

NOTE: Multiple tunnels configured on a router cannot share the same remote address

---

EXAMPLE: To configure an IP tunnel as seen in **Figure 10.9**, the user would enter the IP tunnel destination address on an interface of the router. To configure an IP address on Router A, the user would enter the following:

USING THE CLI

```
HP9300(config)# int e 3/1

HP9300(config-if-3/1)# 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, the user would enter the following:

```
HP9300(config-if-3/1)# ip tunnel 192.58.4.1
```

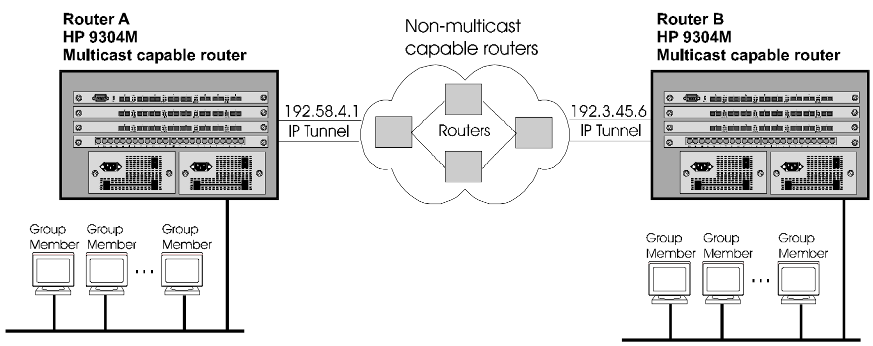

Figure 10.9 IP in IP tunneling on multicast packets in a unicast network

USING THE WEB MANAGEMENT INTERFACE

1. Select the <u>DVMRP virtual interface</u> link from the main menu. The panel seen in **Figure 10.8** will appear.

   NOTE: IP tunnels can be configured for PIM and DVMRP. If the user were configuring an IP tunnel for PIM, he or she would 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** pull down 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, the user can set a TTL value, and enable advertisement of routes and encapsulation on the IP tunnel. For PIM, the user can set a TTL value.

---

6. Select the **add** button.
7. Repeat steps 1 through 6 for the interface on the remote end of the IP tunnel.

Configuring IGMP

This feature allows HP routing switches to limit the multicast of IGMP packets to only those ports on the router that are identified as IP Multicast members.

The router will actively send out host queries to identify IP Multicast groups on the network, insert this information in the IGMP packet and forward it accordingly.

This feature is automatically activated when the DVMRP routing protocol is enabled on the router.

IGMP has two parameters that can be modified:

- IGMP query interval period
- IGMP membership time

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 to 3,600 seconds and the default value is 60 seconds.

USING THE CLI

To modify the default value for the IGMP query interval, the user 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, the user would do the following:

1. Select IGMP from the DVMRP configuration sheet. The panel shown in **Figure 10.10** will appear.
2. Enter a value between 1 and 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 1 to 7,200 seconds and the default value is 140 seconds.

USING THE CLI

To define an IGMP membership time of 240 seconds, the user would 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, the user would do the following:

1. Select IGMP from the DVMRP configuration sheet. The panel shown in **Figure 10.10** will appear.
2. Enter a value between 1 and 7200.
3. Select the **apply** button to assign the changes.

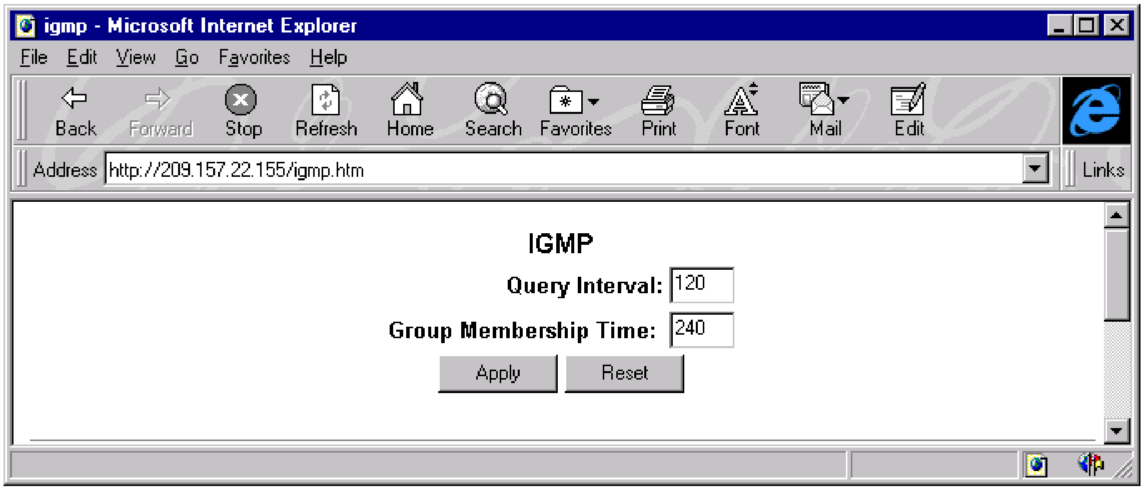

Figure 10.10 IGMP configuration panel