

Cisco.com

Interconnecting PIM & DVMRP Multicast Networks

Module 9

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

1

Module Objectives

Cisco.com

- **Learn that Cisco routers don't run DVMRP!**
- **Explain Cisco routers' PIM-DVMRP interoperability features & functions.**
- **Demonstrate an ability to configure PIM-DVMRP interoperability**
- **Demonstrate an ability to debug PIM-DVMRP interoperability**

Module Agenda



Cisco.com

- **PIM-DVMRP Interoperability**
- **Advanced PIM-DVMRP Features**
- **Debugging Tips**

PIM-DVMRP Interoperability

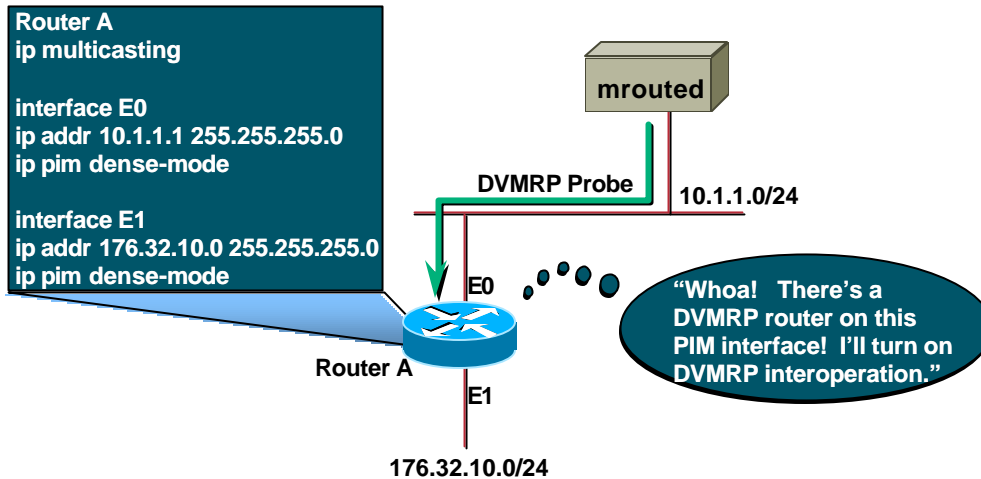
Cisco.com

- **DVMRP router discovery**
- **Basic PIM-DVMRP interaction**
- **PIM-DVMRP route exchange**
- **PIM-DVMRP RPF checking**
- **PIM-DVMRP congruency**
- **PIM-DM/DVMRP Boundaries**
- **PIM-SM/DVMRP Boundaries**

DVMRP Router Discovery

Cisco.com

Automatic Detection of DVMRP Routers on PIM Enabled Interfaces



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

5

- **DVMRP Router Discovery**

- A Cisco router detects the existence of a DVMRP router on an interface whenever it receives a DVMRP Probe packet.
- PIM-DVMRP Interoperability is enabled automatically whenever a DVMRP router is detected on an interface that has PIM enabled.

PIM-DVMRP Interoperability

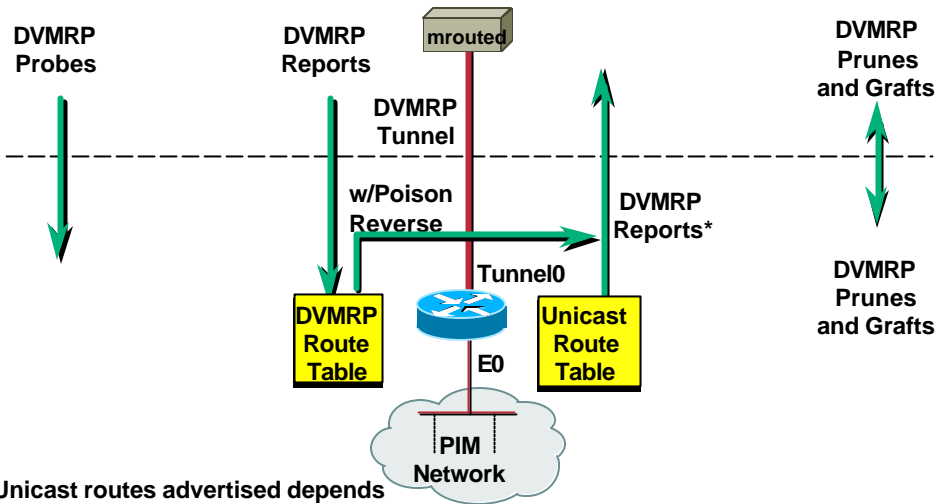
Cisco.com

- DVMRP router discovery
- **Basic PIM-DVMRP interaction**
- PIM-DVMRP route exchange
- PIM-DVMRP RPF checking
- PIM-DVMRP congruency
- PIM-DM/DVMRP Boundaries
- PIM-SM/DVMRP Boundaries

Basic PIM-DVMRP Interaction

Cisco.com

PIM-DVMRP Interaction over p2p interfaces



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

7

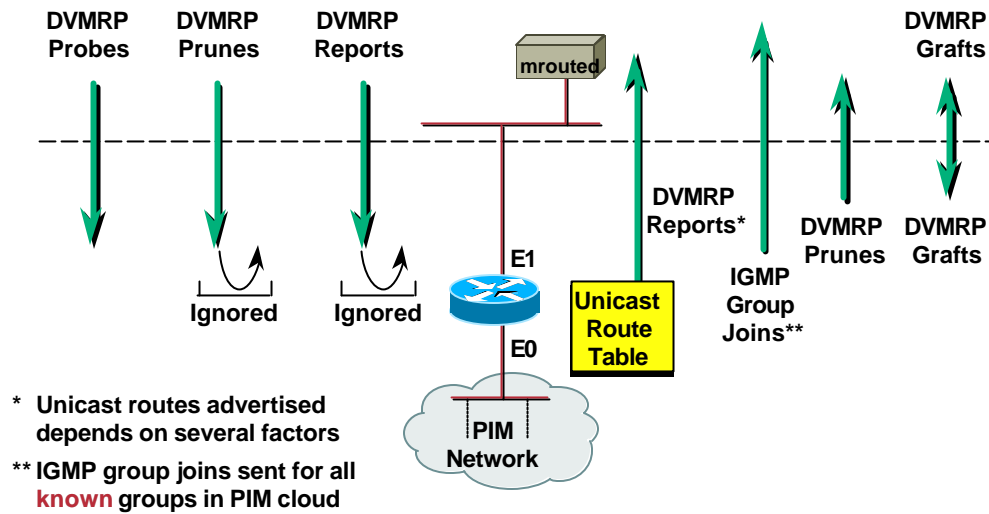
• Interaction over P2P interfaces

- Cisco routers receive the DVMRP Probes which notifies the Cisco router that a DVMRP neighbor is on the link. However, Cisco routers do not send DVMRP probes.
- Cisco routers receive DVMRP Route Reports over p2p links and stores them in a separate DVMRP routing table. The Cisco router sends back Poison-Reverse routes to the DVMRP neighbor so that the DVMRP neighbor will forward any traffic from sources in these networks across the link. (This effectively puts the Cisco router on the Truncated Broadcast Tree for sources in network that was Poison-Reversed.)
- Cisco routers redistribute certain Unicast routes to the DVMRP neighbor as DVMRP Route Reports.
- Cisco routers receive and process DVMRP Grafts and Prunes that arrive via p2p links.
- Cisco routers also send DVMRP Prunes and Grafts (as well as Graft-Acks) via p2p links.

Basic PIM-DVMRP Interaction

Cisco.com

Old PIM-DVMRP interaction over non-p2p interfaces (Ethernets, Etc.) prior to 11.2(13)



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

8

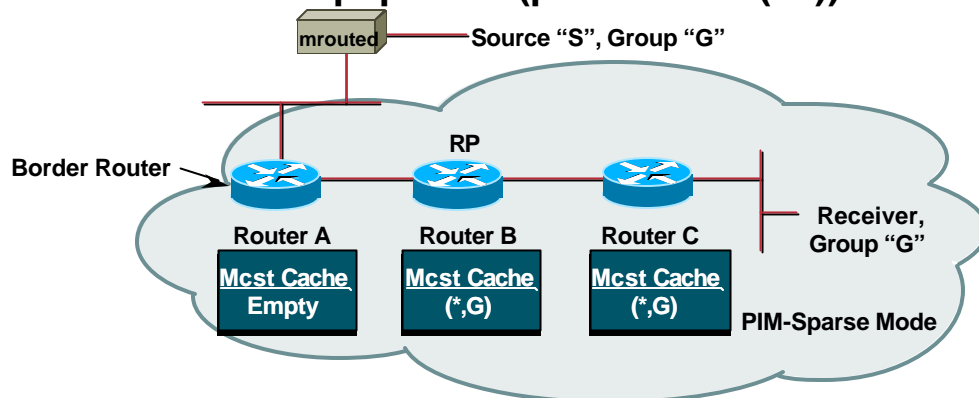
• Interaction over non-P2P interfaces (prior to 11.2(13))

- Cisco routers receive the DVMRP Probes which notifies the Cisco router that a DVMRP neighbor is on the link. However, Cisco routers do not send DVMRP probes.
- Cisco routers ignore DVMRP Route Reports received over p2p links prior to IOS version 11.2(13).
- Cisco routers redistribute certain Unicast routes to the DVMRP neighbor as DVMRP Route Reports.
- Cisco routers ignore any DVMRP Prunes received via a non-p2p link prior to IOS version 11.2(13).
- Prior to IOS version 11.2(13), Cisco routers would send IGMP Joins via non-p2p links for any active groups in the mroute table.
- Cisco routers ignore DVMRP Prunes received via non-p2p links. This is because Cisco routers do not keep track of how many DVMRP neighbors are on a multi-access link and which ones have sent prunes and which haven't. (Note: DVMRP Prunes are supported on multi-access links in later releases of IOS.)
- Cisco routers will send DVMRP Prunes and Grafts (as well as Graft-Acks) via non-p2p links prior to IOS version 11.2(13).
- Cisco routers receive and process DVMRP Grafts that arrive via non-p2p links prior to IOS version 11.2(13).

PIM-DVMRP Interaction

Cisco.com

Problem with old Default PM-DVMRP interaction over non-p2p links (prior to 11.2(13))



- Border router doesn't know about group "G" and doesn't send an IGMP group join for "G"
- Therefore, DVMRP router will not forward packets from source "S"

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

9

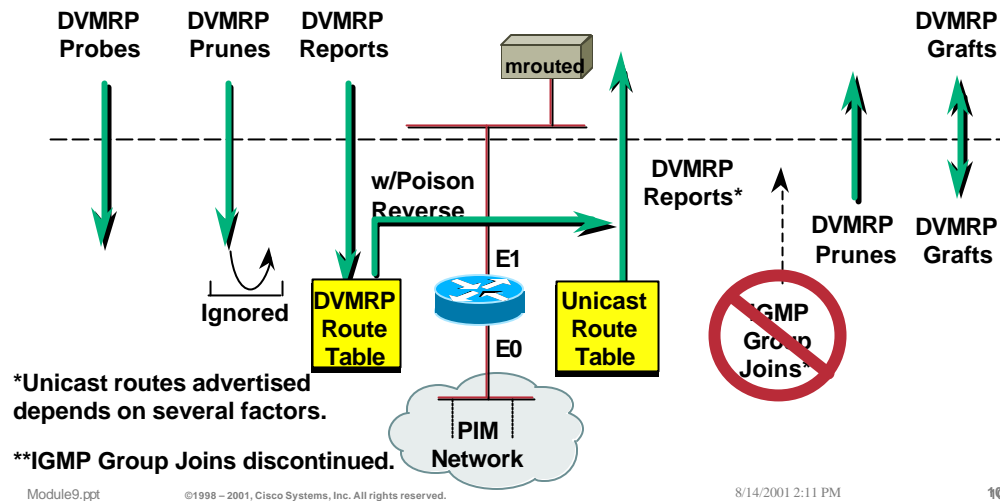
• Problem with the "IGMP Join" interaction

- Consider the example PIM-SM network in the drawing above. Here we have a receiver for group "G" which has resulted in the creation of a branch of the Shared Tree from the RP (Router B) down to the receiver.
- Router A is the border router connected to the DVMRP router. Unfortunately, Router A has no state for group "G" since it is not on the Shared Tree. As a result, Router A will not send IGMP Joins and source "S" traffic will not be forwarded across the boundary.

Basic PIM-DVMRP Interaction

Cisco.com

Using “ip dvmrp unicast-routing” to modify old default PIM-DVMRP interaction over non-p2p (prior to version 11.2(13))



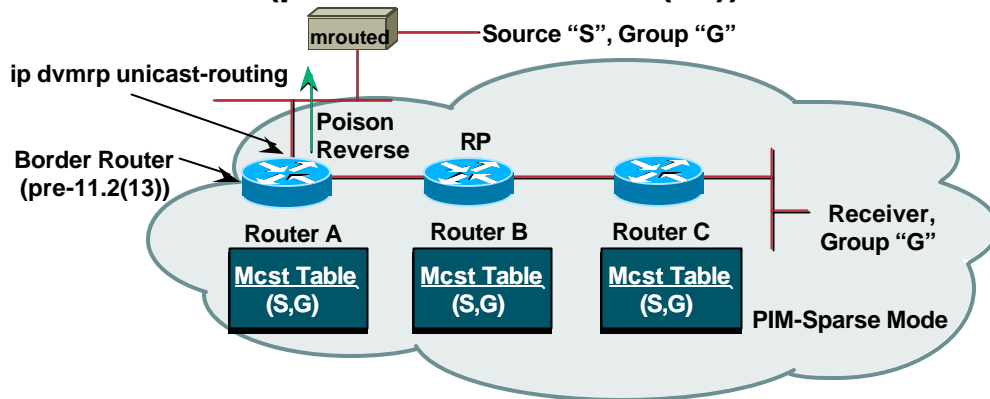
• Using “ip dvmrp unicast-routing” on non p2p links prior to 11.2(13)

- Cisco routers receive the DVMRP Probes which notifies the Cisco router that a DVMRP neighbor is on the link. However, Cisco routers do not send DVMRP probes.
- Cisco routers receive DVMRP Route Reports over p2p links and stores them in a separate DVMRP routing table when “ip dvmrp unicast-routing” is configured on the link. The Cisco router will then also begin sending back Poison-Reverse routes to the DVMRP neighbor so that the DVMRP neighbor will forward any traffic from sources in these networks across the link. (This effectively puts the Cisco router on the Truncated Broadcast Tree for sources in network that was Poison-Reversed.)
- Cisco routers redistribute certain Unicast routes to the DVMRP neighbor as DVMRP Route Reports.
- Prior to IOS version 11.2(13), Cisco routers would send IGMP Joins via non-p2p links for any active groups in the mroute table. However, this behavior is suppressed when “ip dvmrp unicast-routing” is configured on the link. (The router instead sends Poison-Reverse routes for source networks that it wishes to receive as mentioned in the previous paragraph.)
- Cisco routers ignore DVMRP Prunes received via non-p2p links. This is because Cisco routers do not keep track of how many DVMRP neighbors are on a multi-access link and which ones have sent prunes and which haven't. (Note: DVMRP Prunes are supported on multi-access links in later releases of IOS.)
- Cisco routers will send DVMRP Prunes and Grafts (as well as Graft-Acks) via non-p2p links prior to IOS version 11.2(13).
- Cisco routers receive and process DVMRP Grafts that arrive via non-p2p links prior to IOS version 11.2(13).

PIM-DVMRP Interaction

Cisco.com

Using “ip dvmrp unicast-routing” over non-p2p links (prior to version 11.2(13))



- Border router now sends poison reverse routes so DVMRP router knows to forward packets from “S”

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

11

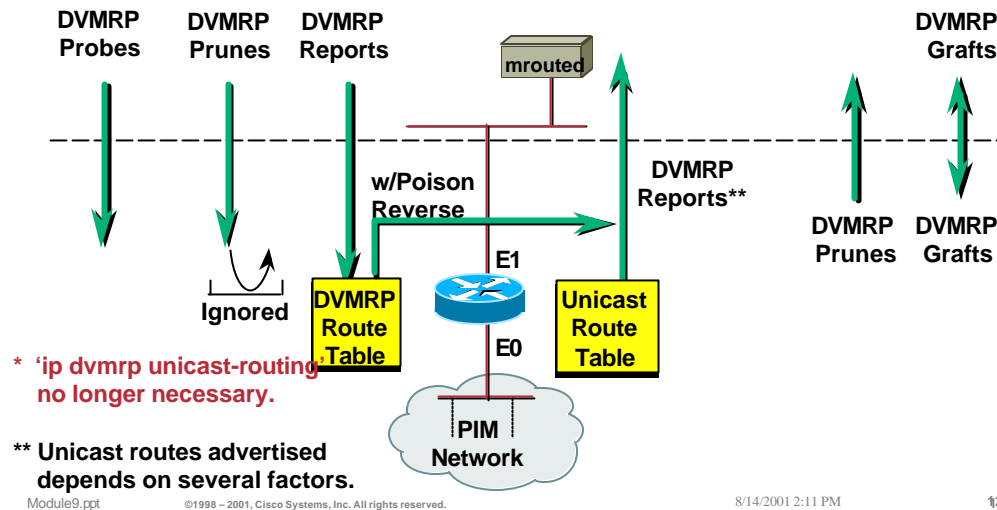
- **Workaround: Using “ip dvmrp unicast-routing” prior to 11.2(13)**

- Prior to IOS version 11.2(13), Cisco routers would not send Poison-Reverse routes over non-p2p links. Instead, it would send IGMP Joins if and only if the router had state for a group.
- The “ip dvmrp unicast-routing” command can be used to change this behavior and force the Cisco router to exchange DVMRP Route updates including Poison-Reverse updates.
- The Poison-Reverse updates inform the DVMRP router that the Cisco router should be placed on the Truncated Broadcast Tree for sources on the network specified in the Poison-Reversed route. This will cause the traffic from source “S” in this example, to be flooded across the boundary to Router A.
- When Router A receives this traffic from source “S”, it will then automatically proxy-register “S” to the RP (as if it were a directly connected source). In addition, Router A will join the Shared Tree for group “G”. (This insures that traffic for group “G” that is originated by a source in the PIM-SM cloud will be forwarded to the DVMRP router.

Basic PIM-DVMRP Interaction

Cisco.com

New PIM-DVMRP interaction over non-p2p (Ethernet, FDDI, etc.) beginning with 11.2(13)*



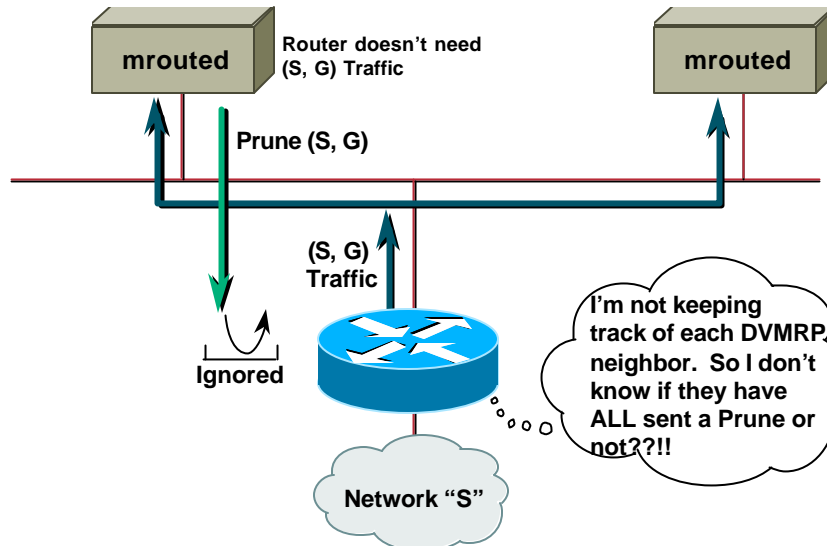
• Interaction over non p2p links after IOS version 11.2(13)

- Cisco routers receive the DVMRP Probes which notifies the Cisco router that a DVMRP neighbor is on the link. However, Cisco routers do not send DVMRP probes.
- Cisco routers automatically receive DVMRP Route Reports over p2p links and stores them in a separate DVMRP routing table. (It is no longer necessary to configure the "ip dvmrp unicast-routing" command on the link.) The Cisco router will also automatically send back Poison-Reverse routes to the DVMRP neighbor so that the DVMRP neighbor will forward any traffic from sources in these networks across the link. (This effectively puts the Cisco router on the Truncated Broadcast Tree for sources in network that was Poison-Reversed.)
- Cisco routers redistribute certain Unicast routes to the DVMRP neighbor as DVMRP Route Reports.
- Cisco routers will send DVMRP Prunes and Grafts (as well as Graft-Acks) via non-p2p links prior to IOS version 11.2(13).
- Cisco routers ignore DVMRP Prunes received via non-p2p links. This is because Cisco routers do not keep track of how many DVMRP neighbors are on a multi-access link and which ones have sent prunes and which haven't. (Note: DVMRP Prunes are supported on multi-access links in later releases of IOS.)
- Cisco routers receive and process DVMRP Grafts that arrive via non-p2p links prior to IOS version 11.2(13).

Basic PIM-DVMRP Interaction

Cisco.com

Prune problem on non-p2p links prior to 12.1



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

13

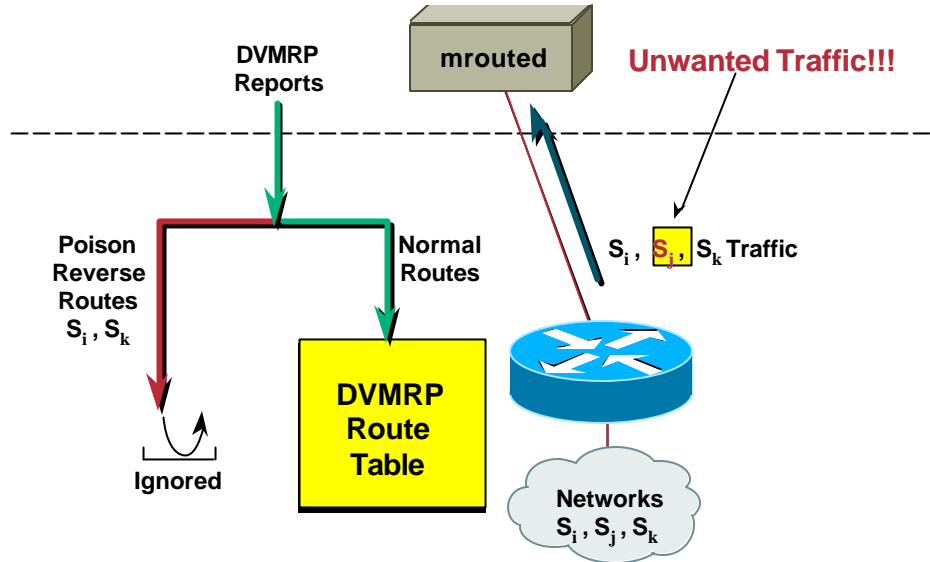
• Prune problem on non-p2p links prior to IOS version 12.1

- The DVMRP pruning mechanisms on non-p2p links work differently than the pruning mechanisms of PIM. In PIM, the upstream router waits for 3-4 seconds before pruning a non-p2p link (i.e. a multi-access link) over which it has received a PIM Prune message. This provides other PIM routers on the link to override the Prune with a Join message. In DVMRP, there is no Prune override operation. Instead, the upstream DVMRP router is expected to keep track of all downstream DVMRP neighbors on the link and only prune the traffic when ALL DVMRP routers on the link had sent a Prune message. (In other words, the upstream router had to get a unanimous vote to Prune from all DVMRP routers on the link before it would prune.)
- Prior to release 12.1, Cisco routers did not keep track of the individual DVMRP routers on a multi-access (non-p2p) link. Therefore, it was unable to ascertain that ALL DVMRP routers on the link had sent a DVMRP Prune for the traffic. As a result, the Cisco router would ignore all DVMRP Prune messages and continue to send traffic under the assumption that there was still some other DVMRP router on the link that needed the traffic.
- Beginning with IOS version 12.1, Cisco routers now keep track of all DVMRP routers on the link and will prune the link when it receives a DVMRP prune from all DVMRP routers on the wire.

Basic PIM-DVMRP Interaction

Cisco.com

Poison-Reverse ignored prior to 11.2(13)



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

14

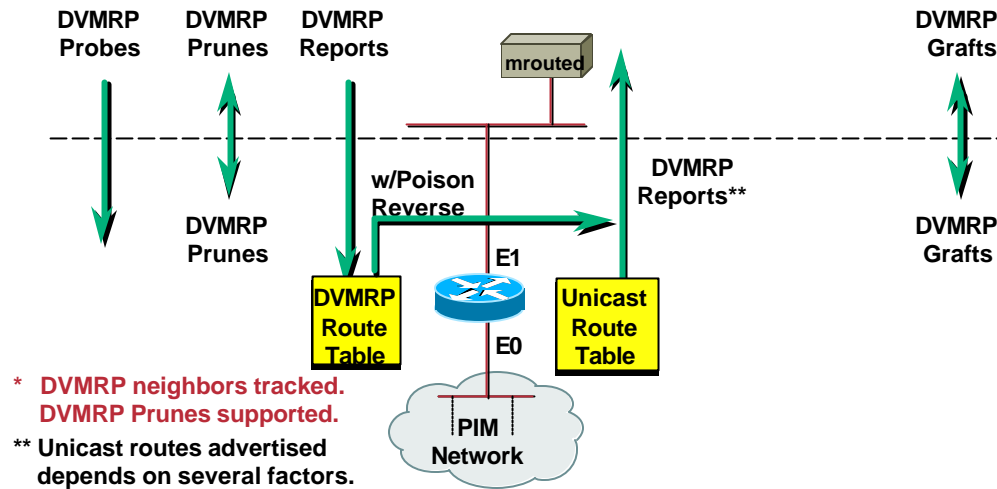
- **Poison-Reverse messages ignored prior to 11.2(13)**

- Prior to IOS version 11.2(13), Cisco routers ignored all Poison-Reverse routes from DVMRP routers and just *assumed* that the DVMRP routers wanted all traffic sourced from the PIM cloud. (This was equivalent to assuming every route advertised to a DVMRP neighbor resulted in a Poison-Reverse route being sent back from the DVMRP neighbor.)
- This behavior was sub-optimal as the DVMRP router would sometimes have a better route to a particular source network via some other DVMRP router in the DVMRP cloud. This would result in unwanted traffic to flow to the DVMRP router. This was fixed after release 11.2(13).

Basic PIM-DVMRP Interaction

Cisco.com

New PIM-DVMRP interaction over non-p2p (Ethernet, FDDI, etc.) beginning with 12.1*



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

15

• Interaction on non-p2p links after IOS version 12.1

Beginning with IOS version 12.1 and on, DVMRP prunes are no longer ignored on multi-access (non-p2p) links. This leads to the following PIM-DVMRP interaction over all links.

- Cisco routers receive the DVMRP Probes which notifies the Cisco router that a DVMRP neighbor is on the link. However, Cisco routers do not send DVMRP probes.
- Cisco routers automatically receive DVMRP Route Reports and stores them in a separate DVMRP routing table. The Cisco router will also automatically send back Poison-Reverse routes to the DVMRP neighbor so that the DVMRP neighbor will forward any traffic from sources in these networks across the link. (This effectively puts the Cisco router on the Truncated Broadcast Tree for sources in network that was Poison-Reversed.)
- Cisco routers redistribute certain Unicast routes to the DVMRP neighbor as DVMRP Route Reports.
- Cisco routers will send DVMRP Prunes and Grafts (as well as Graft-Acks).
- Cisco routers receive and process DVMRP Prunes and Grafts.

Recommendations

- Use IOS release 12.1 or later.

OR

- Use DVMRP tunnels or p2p interfaces whenever possible.

- **PIM-DVMRP Recommendation**

- Use IOS version 12.1 or later in order to more closely emulate the functions of a DVMRP router and to avoid several of the issues discussed on the previous pages.
- If this is not possible, the use of p2p links or DVMRP Tunnels is recommended in order to avoid issues such as the Pruning issue.
 - Note: DVMRP Tunnels can be used across a multi-access link to avoid many of the problems outline previously. This requires a separate DVMRP Tunnel to each DVMRP router on the multi-access network. In addition, the PIM router and the DVMRP router *SHOULD* disable multicast on this multi-access link and only permit multicast traffic to flow via the DVMRP Tunnels.

PIM-DVMRP Interoperability

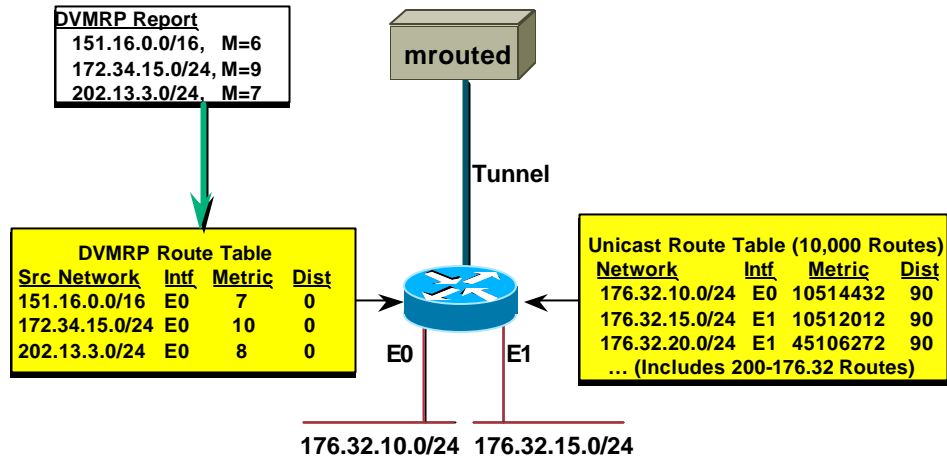
Cisco.com

- DVMRP router discovery
- Basic PIM-DVMRP interaction
- **PIM-DVMRP route exchange**
- PIM-DVMRP RPF checking
- PIM-DVMRP congruency
- PIM-DM/DVMRP Boundaries
- PIM-SM/DVMRP Boundaries

PIM-DVMRP Route Exchange

Cisco.com

Received DVMRP Routes Are Installed in a Separate DVMRP Route Table



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

18

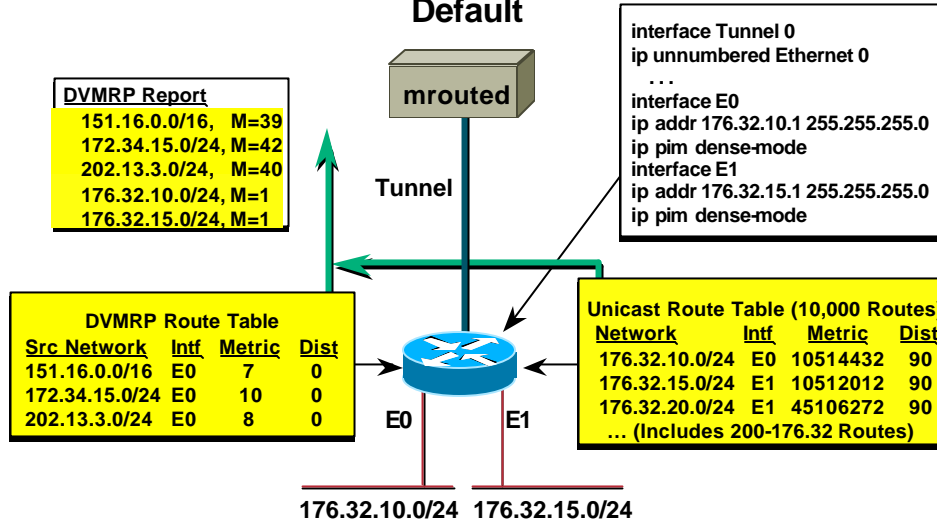
- **PIM-DVMRP Route Exchange**

- When a Cisco router receives a DVMRP Route update, the routes are stored in a separate DVMRP route table. The routes in the DVMRP route table have a default Admin. Distance of zero and are therefore preferred over routes in the Unicast Route table when performing the RPF calculation.

PIM-DVMRP Route Exchange

Cisco.com

Only "Connected" Unicast Routes Are Advertised by Default



Module9.ppt

©1998 - 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

19

• PIM-DVMRP Route Exchange

- Cisco routers will advertise the routes in the DVMRP route table to DVMRP neighbors. The Cisco router will use the normal DVMRP Poison-Reverse mechanism to indicate those networks that it expects to receive traffic from the DVMRP router.
- In addition to the DVMRP routes mentioned above, the Cisco router will by default, redistribute (advertise) all connected routes to the DVMRP neighbor as DVMRP routes with a metric of 1.
- In the example above, the Cisco router has received three DVMRP routes from the DVMRP neighbor:

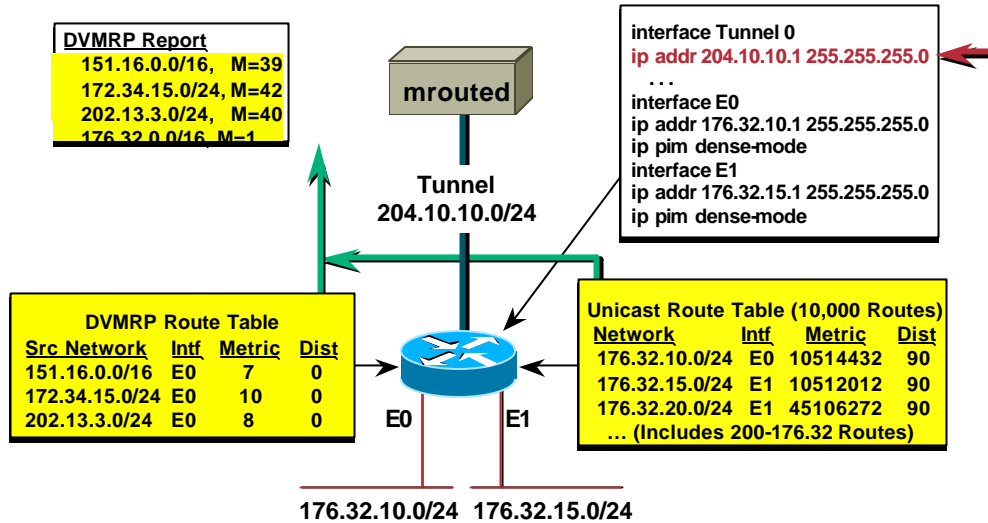
151.16.0.0/16	Metric=7
172.34.15.0/24	Metric=10
202.13.3.0/24	Metric=8
- The Cisco router Poison-Reverses these routes by adding "Infinity" (32) to these metrics and sending them back to the DVMRP router. This informs the DVMRP router that the Cisco router should be included on the Truncated Broadcast Tree for these source networks. This will cause traffic from any source in these networks to be flooded across the interface.
- The Cisco router is automatically redistributing (advertising) all connected routes from its unicast route table with a DVMRP metric of one. This results in the following DVMRP route advertisements being sent to the DVMRP neighbor:

176.32.10.0/24	Metric=1
176.32.15.0/24	Metric=1

PIM-DVMRP Route Exchange

Cisco.com

Classful Summarization of Unicast Routes



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

20

• Classful Summarization of redistributed Unicast Routes

Cisco routers automatically perform classful summarization of routes that are advertised to a DVMRP neighbor.

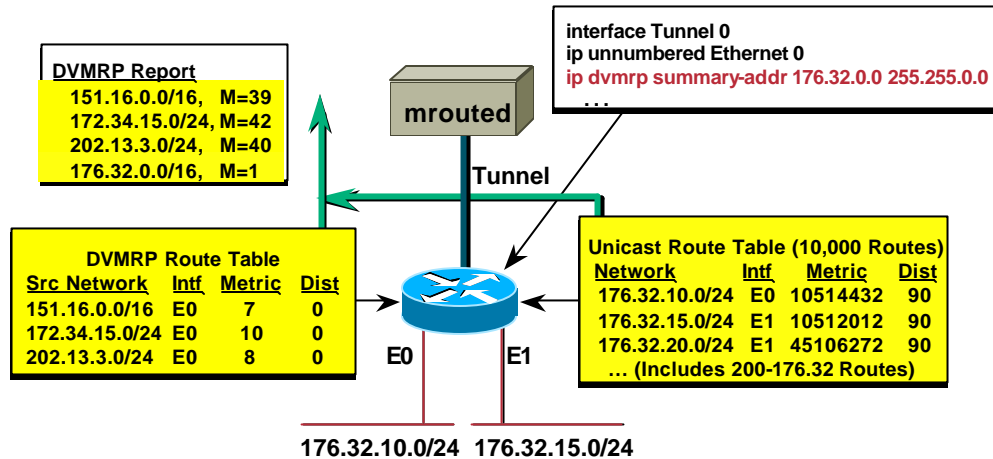
- In the example above, the interface between the Cisco router and the DVMRP router is not within the same classful subnet as the connected routes. This causes the connected networks to be summarized into a Class B network advertisement resulting in the following advertised route:

176.32.0.0/16 Metric=1

PIM-DVMRP Route Exchange

Cisco.com

Forcing DVMRP Route Summarization (11.2(5))



Module9.ppt

©1998 - 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

21

• Forcing Route Summarization

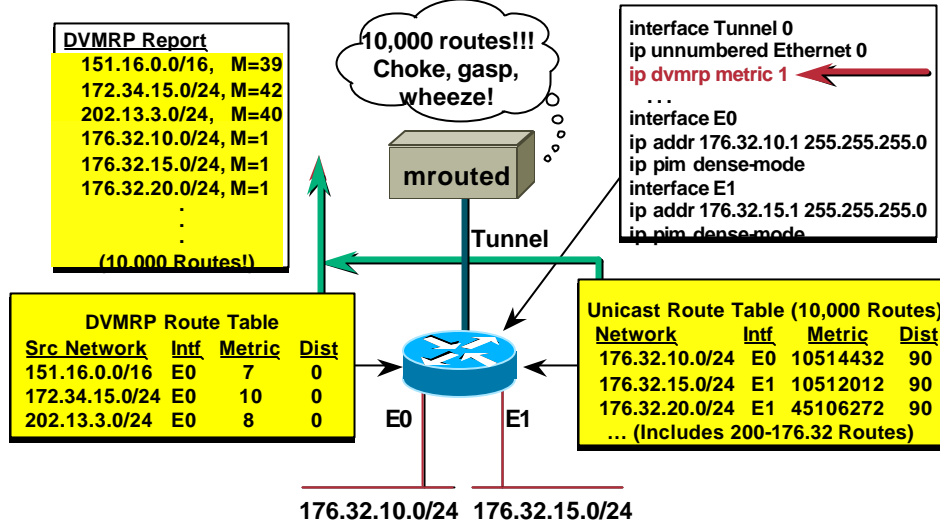
- Beginning with IOS version 11.2(5), classful summarization can be forced even though the Tunnel interface falls within the same classful network as the connected routes.
- In the above example, the 'ip dvmrp summary-address' command is being used to summarize the two connected networks into the following DVMRP route advertisement:

176.32.0.0/16 Metric=1

PIM-DVMRP Route Exchange

Cisco.com

The Deadly “ip dvmrp metric n” Command



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

22

• Redistributing other Unicast Routes

- If it is desired to redistribute more routes from the Unicast Routing table than the default “connected” routes, the interface may be configured with the ‘ip dvmrp metric’ command.
- In the example above, the Tunnel interface is configured with “ip dvmrp metric 1” command. Because no option ACL was specified in this command, a “permit any” is assumed which results in all 10,000 routes from the Unicast Routing table being redistributed to the DVMRP neighbor with a DVMRP metric of 1.
 - Note: The ‘ip dvmrp metric’ command should normally be used with an optional ACL in order to limit the routes being redistributed to the DVMRP neighbor so that route loops are not formed. The use of an ACL is especially important if the size of the Unicast Routing table is quite large as it may be possible to overwhelm the DVMRP router with too many routes.

PIM-DVMRP Interoperability

Cisco.com

- DVMRP router discovery
- Basic PIM-DVMRP interaction
- PIM-DVMRP route exchange
- **PIM-DVMRP RPF checking**
- PIM-DVMRP congruency
- PIM-DM/DVMRP Boundaries
- PIM-SM/DVMRP Boundaries

PIM-DVMRP RPF Checking

Cisco.com

- **Sources of RPF Information**
 - Static Mroutes
 - MBGP Multicast NLRI (M-RIB)
 - DVMRP Routes
 - Unicast Routes
- **RPF Information source selection**
 - Based on Admin. Distance
 - If equal, preferred in order listed above
 - i.e. Static Mroute is most preferred source
 - Default DVMRP Admin. Distance = 0

Module9.ppt

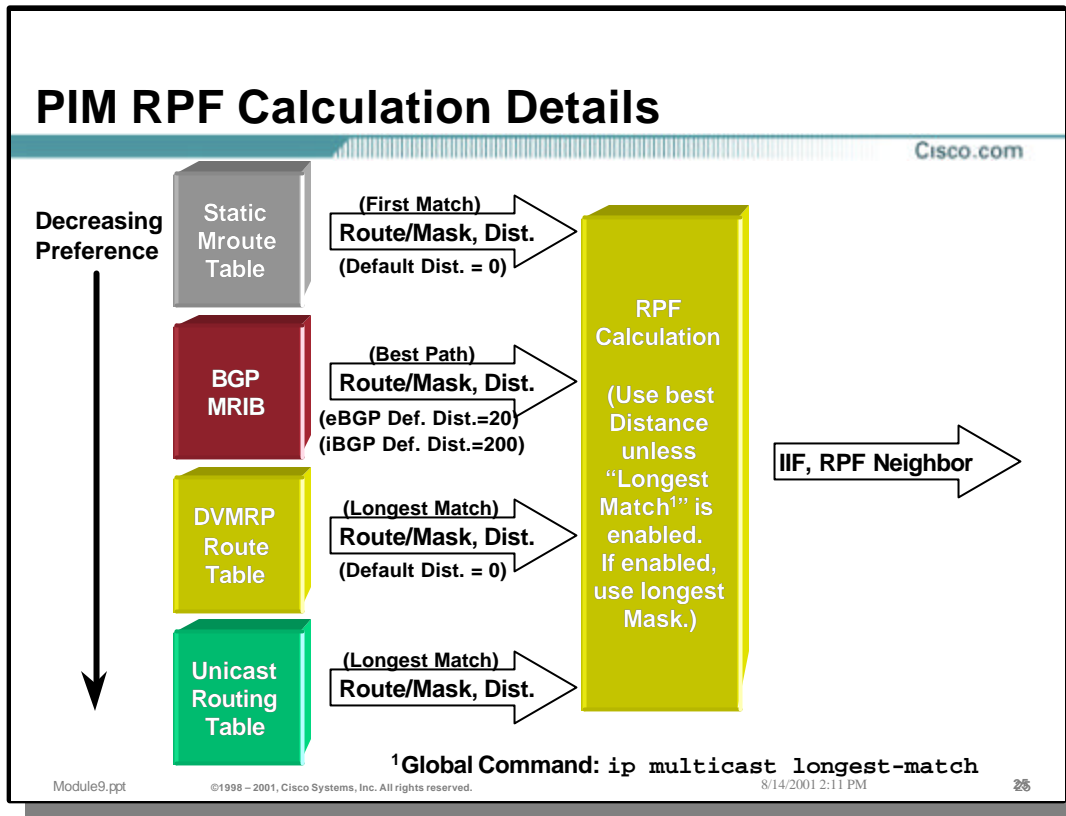
©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

24

• Sources of RPF Information

- Cisco routers calculate RPF Information (Incoming Interface and RPF Neighbor) using several sources. These sources are:
 - Static Mroutes
 - MBGP Multicast NLRI (M-RIB)
 - DVMRP Routes
 - Unicast Routes
- The selection as to which source to use is based on Administrative Distance. If the Admin. Distances are equal, the tie is broken using the order listed above. That is to say, Static Mroutes are preferred over MBGP Multicast NLRI which is preferred over DVMRP routes which are preferred over Unicast routes.
 - Note: The default Admin. Distance of routes in the DVMRP route table is zero.



• PIM RPF Calculations

Cisco IOS permits other sources of information to be used in the RPF calculation other than the unicast routing table. In general, these other sources are preferred based on their Admin. Distance. If Admin. Distance values are equal, the sources are preferred in the order listed below:

- Static Mroute Table

Static Mroutes may be defined that are local to the router on which they are defined. If a matching Static Mroute is defined, its default Admin. Distance is zero and is therefore preferred over other sources. (If another source also has a distance of zero, the Static Mroute takes precedence.)

- BGP Multicast RIB (M-RIB)

If MBGP is in use and a matching prefix exists in the MBGP M-RIB, it will be used as long as its Admin. Distance is the lowest of the other sources. (MBGP M-RIB prefixes are preferred over DVMRP or Unicast routes if the Admin. Distances are the same.)

- DVMRP Route Table

If DVMRP routes are being exchanged and there exists a matching route in the DVMRP route table, the default Admin. Distance of this route is zero. DVMRP routes are preferred over Unicast routes if their Admin. Distances are equal.

- Unicast Route Table

This is least preferred source of information. If no other source has a matching route with a lower Admin. Distance, then this information is used.

Note: The above behavior can be modified so that the longest match route is used from the available sources. This is configured with the 'ip multicast longest-match' hidden command.

PIM-DVMRP Interoperability

Cisco.com

- DVMRP router discovery
- Basic PIM-DVMRP interaction
- PIM-DVMRP route exchange
- PIM-DVMRP RPF checking
- **PIM-DVMRP congruency**
- PIM-DM/DVMRP Boundaries
- PIM-SM/DVMRP Boundaries

PIM-DVMRP Congruency

Cisco.com

- **RPF Information Source**
 - **Should be consistent in all routers**
 - **Otherwise, RPF checks may fail at some routers**
 - **This results in congruency problems**
 - **DVMRP often causes inconsistencies**
 - **Some routers use DVMRP routes for RPF check**
 - **Others use Unicast routes for RPF check**

Module9.ppt

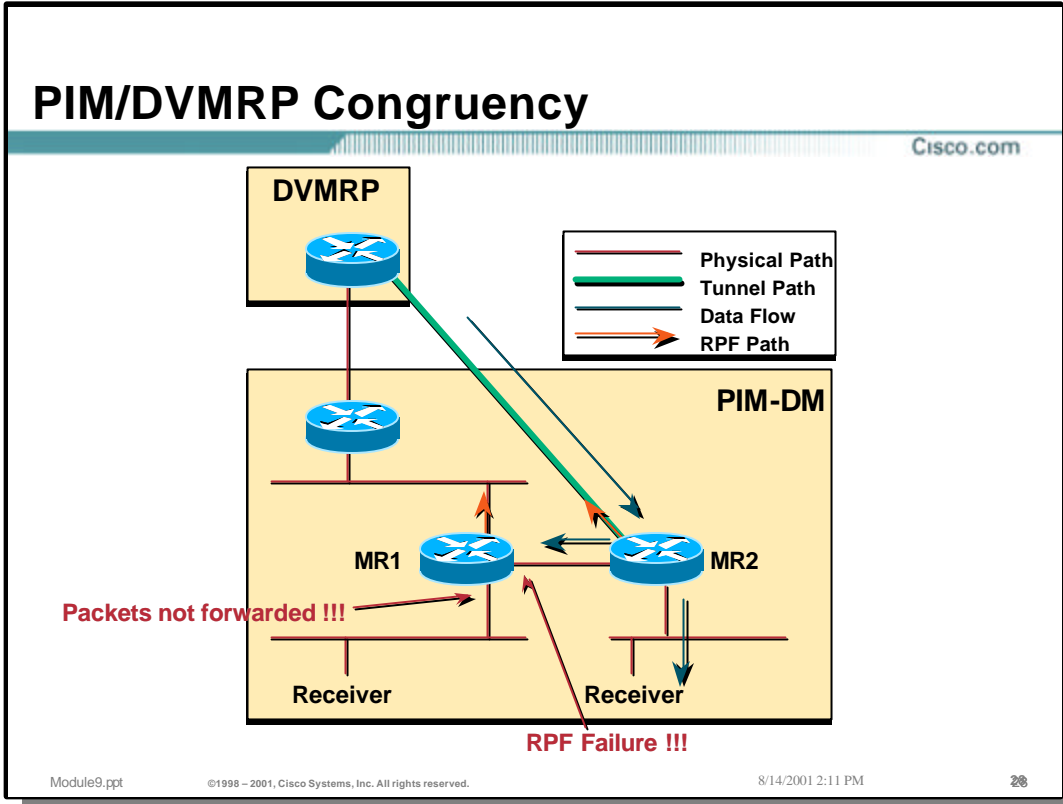
©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

27

- **PIM-DVMRP Congruency**

- The source of RPF information normally should be consistent in all routers in the network or the possibility of congruency problems can occur that result in multicast traffic failing the RPF check in some routers.
- The use of DVMRP routes can often cause these inconsistencies to occur if all routers in the network are not maintaining DVMRP route tables. This can result in some routers using the DVMRP route table for RPF calculations while the other routers are still using the Unicast route table.



- **PIM-DVMRP Congruency**

- The slide above shows an example of how the use of DVMRP routes can cause RPF failures as a result of Unicast/DVMRP incongruity.
 - Router “MR2” is connected to a router in the DVMRP cloud via a DVMRP tunnel over which it is receiving DVMRP routes for sources in the DVMRP network. Since the default distance of DVMRP routes is zero, the DVMRP routes will be used by MR2 to perform the RPF calculation for sources in the DVMRP cloud.
 - Multicast traffic from sources in the DVMRP cloud arrive via the DVMRP Tunnel and pass the RPF check. (Based on the DVMRP Routing table in MR2 which RPF’s up the Tunnel.)
 - When the multicast traffic flow reaches router MR1, it does not have any DVMRP routes and therefore uses the Unicast Route table to calculate the RPF interface which is not the interface where the multicast traffic is arriving. As a result, the multicast traffic fails the RPF check and is discarded.

PIM/DVMRP Congruency

Cisco.com

- **Don't run any DVMRP!**
 - Sometimes necessary to transition to PIM
- **Make the topologies congruent**
 - Not always an option (although the best)
- **Propagate DVMRP routes to other routers as necessary**
 - Transition method

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM


29

- **PIM/DVMRP Congruency**

- Avoid the use of DVMRP in your network if at all possible. However, this is sometimes unavoidable when transitioning a legacy DVMRP network over to PIM.
- The best solution is to maintain congruency between PIM routing and DVMRP routing. Since PIM normally uses the unicast routing table for the RPF calculation, this means that the DVMRP routing information would have to match the Unicast routing information in order to make the topologies congruent. This is not always an option.
- Another transition method is to propagate DVMRP routes to other routers in the PIM domain to insure that the topologies remain congruent. This should only be considered as a *transition* method since it may be necessary to propagate DVMRP routes to *EVERY* router in the network in order to accomplish this.

PIM/DVMRP Congruency

Cisco.com



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

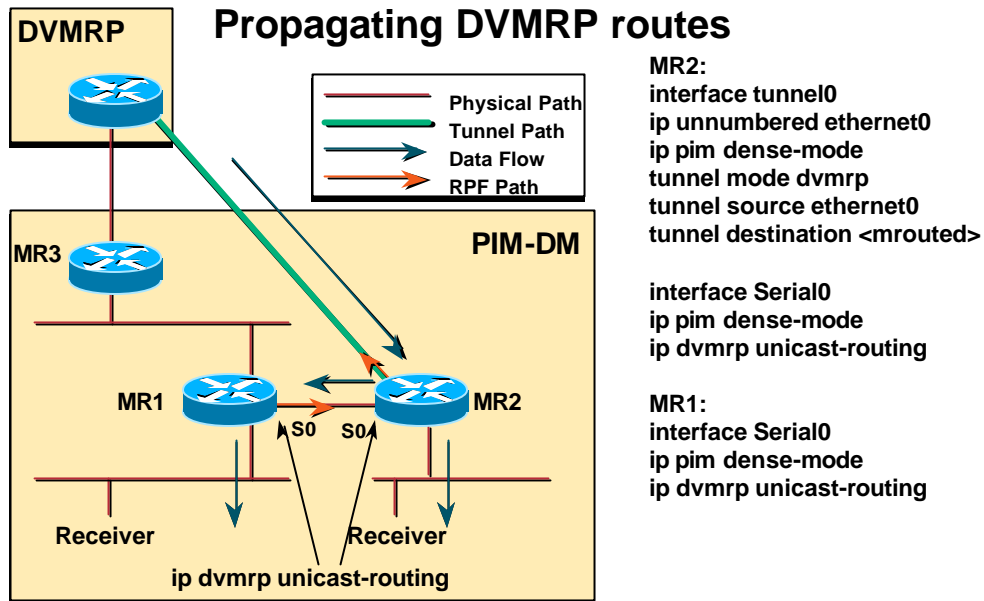
30

- **PIM/DVMRP Congruency**

- In the example above, the DVMRP and PIM (Unicast) topologies are physically congruent. This is because the DVMRP tunnel is terminated on router “MR3” which happens to also be the path out of the PIM domain. Hence, the RPF calculations done on MR1 and MR2 (based on the unicast route table) will RPF correctly.

PIM/DVMRP Congruency

Cisco.com



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

31

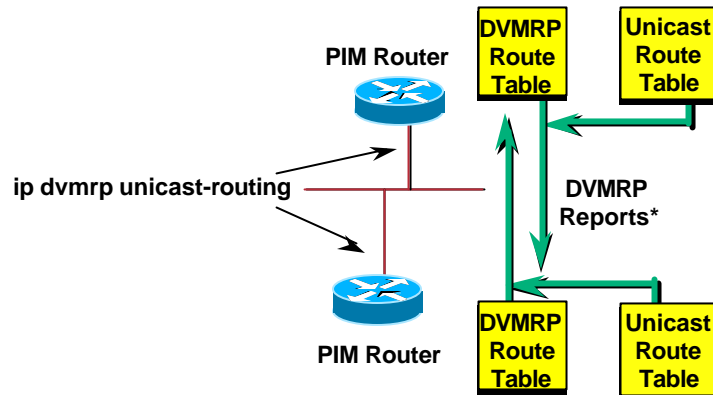
• PIM/DVMRP Congruency

- In the example above, the DVMRP and PIM (unicast) topologies are *not* physically congruent on routers MR1 and MR3. This is because the DVMRP tunnel is terminated further down inside of the PIM domain. (In this case, it is terminated on MR2). As a result, router MR1 would normally have an RPF path (for sources outside the PIM domain) via router MR3. This would cause RPF failures at router MR1 for arriving traffic received via router MR2 and the DVMRP tunnel.
- A transition strategy would be to have MR2 and MR1 exchange DVMRP routes using the 'ip dvmrp unicast-routing' command. This will permit MR2 to send DVMRP routes to MR1 which will in turn, result in MR1 using these routes to RPF for sources in the DVMRP network.
 - Note: It would also be necessary to exchange DVMRP routes from MR1 to MR3 if there are any other interfaces (not shown) on MR3 where there are hosts or PIM neighbor routers. Otherwise, MR3 would also RPF fail any traffic arriving from MR1.

Propagating DVMRP Routes

Cisco.com

Using “ip dvmrp unicast-routing” between two PIM neighbors causes DVMRP routes to be exchanged.



*Split-Horizon is used between two PIM neighbors instead of Poison Reverse.

Module9.ppt

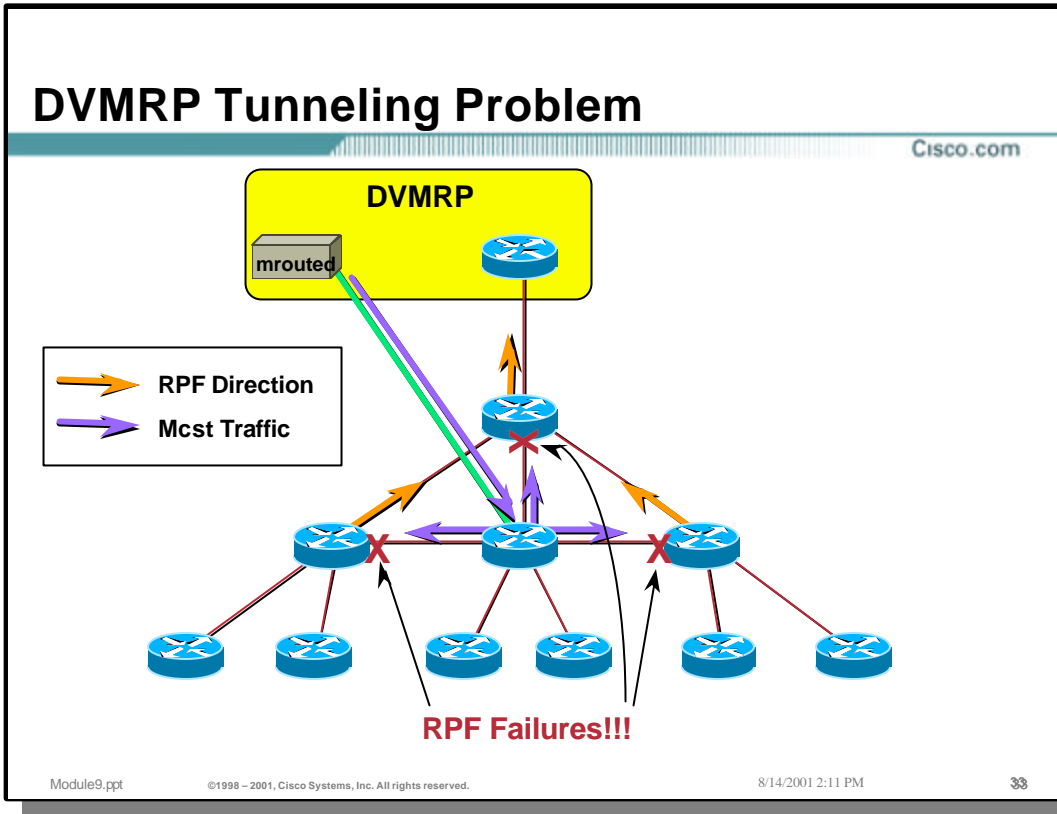
©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

32

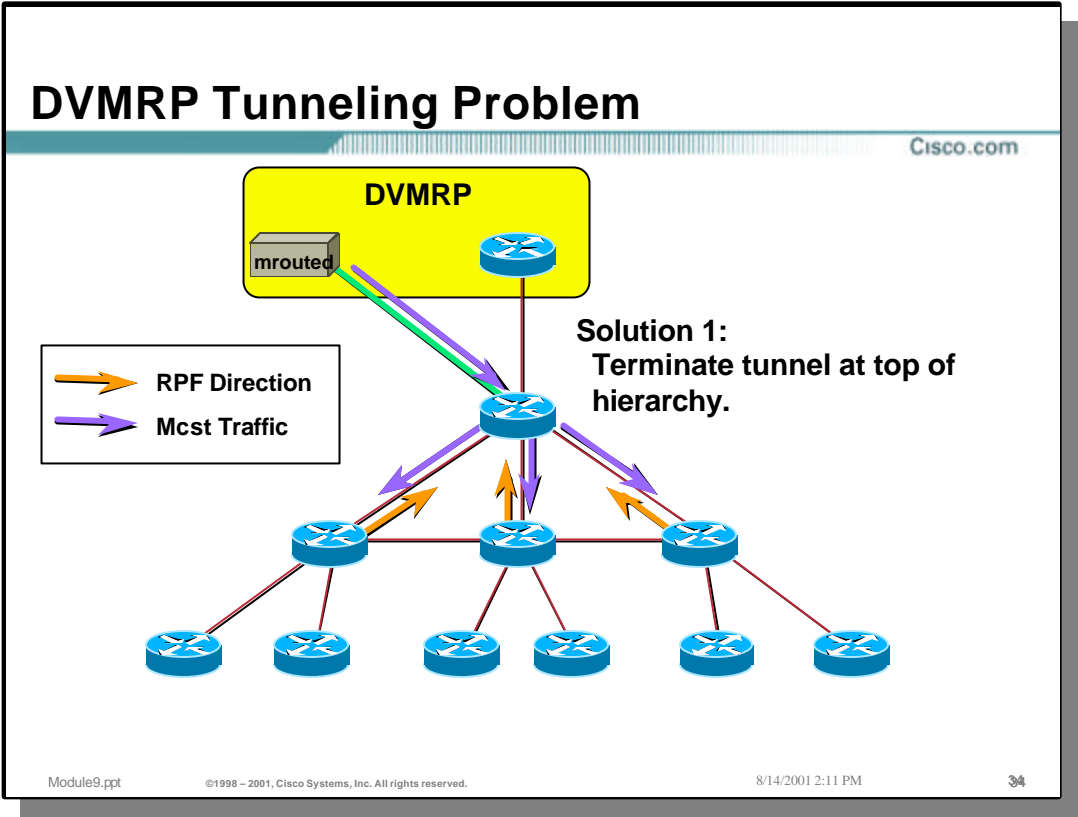
- **Propagating DVMRP Routes between Cisco routers**

- By configuring the 'ip dvmrp unicast-routing' interface command on the interfaces connecting two Cisco routers, DVMRP route updates will be exchanged in the same fashion as if there were a DVMRP router on the interface.
 - Exception: Poison-Reverse is not used between Cisco routers. Split-Horizon is used instead.



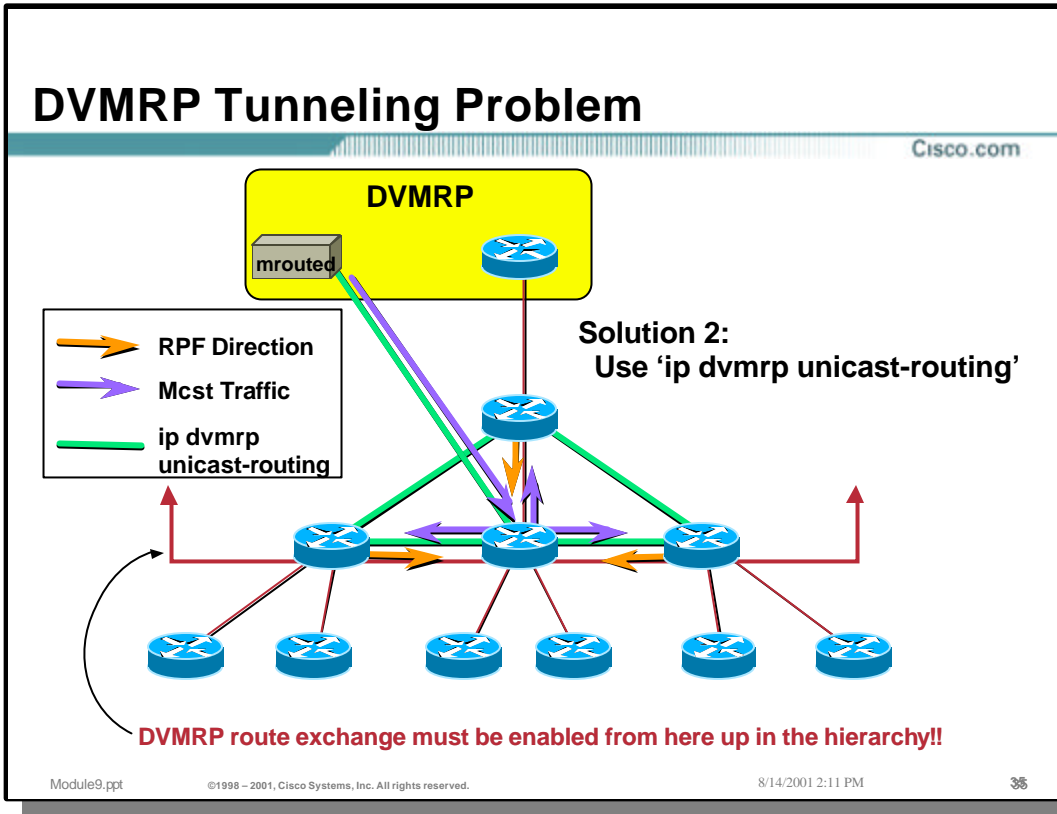
- **DVMRP Tunneling Problem**

- Care must be used when terminating a DVMRP tunnel into a PIM domain as the routers between the tunnel termination router and the normal network exit, will result in RPF failures.
 - An example of this is shown above. Notice that the routers at and above the tunnel termination point are using the unicast routing table to RPF to sources in the DVMRP network. This results in RPF failures as shown.



- **DVMRP Tunneling Problem**

- The best way to terminate a DVMRP tunnel in a PIM domain is to terminate the tunnel on the border router where the PIM domain and the DVMRP domain connect. This results in a congruent topology and all routers down the hierarchy will RPF correctly. (Using the normal Unicast routing table information.)



- **DVMRP Tunneling Problem**

- While it is possible to use 'ip dvmrp unicast-routing' to get the Cisco routers to exchange DVMRP routes (and hence RPF correctly for sources in the DVMRP network), the number of routers that must exchange DVMRP routes can quickly get out of hand.
- The basic rule of thumb is that 'ip dvmrp unicast-routing' must be run on all routers from the tunnel termination point up the hierarchy as shown in the drawing above.
 - Therefore, one should terminate the DVMRP tunnel as close to the PIM-DVMRP network border as possible to minimize the number of routers that must exchange DVMRP routes.

PIM-DVMRP Interoperability

Cisco.com

- DVMRP router discovery
- Basic PIM-DVMRP interaction
- PIM-DVMRP route exchange
- PIM-DVMRP RPF checking
- PIM-DVMRP congruency
- **PIM-DM/DVMRP Boundaries**
- PIM-SM/DVMRP Boundaries

PIM-DM/DVMRP Boundaries

Cisco.com

- **DVMRP Network uses “Push” model**
 - Traffic flooded everywhere
- **PIM-DM Network uses same model**
 - Traffic flooded everywhere
- **Common model makes interface easy.**
 - Use PIM-DM “flood and prune” mechanism.
 - Must also observe key DVMRP signals
 - Poison Reverse
 - DVMRP Prunes

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

37

- **PIM-DM/DVMRP Boundaries**

- DVMRP is a dense mode protocol which uses the “Push” model to flood traffic to all points in the network. (Routers that don’t have receivers for the traffic will then Prune the traffic flow.)
- PIM-DM uses the same “Push” model to “Flood-and-Prune” traffic to all parts of the network.
- Because the two protocols use the same basic “Push” model, traffic can be “flooded” across the PIM-DM/DVMRP boundary without using any complex mechanisms.
 - The PIM-DM boundary router will use the DVMRP Poison-Reverse mechanism as well as normal DVMRP Prunes and Grafts to maintain traffic flow between the two networks.

PIM-DM/DVMRP Boundaries

Cisco.com

- **Traffic sourced from DVMRP cloud**
 - **When packet arrives, create (S,G) state**
(If it didn't already exist.)
 - **Use normal PIM-DM state rules**
 - **Perform normal RPF check**
 - **If passed, flood out all (S,G) OIL entries**
 - **Observing any Pruned (S,G) OIL state**

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

38

• PIM-DM/DVMRP Boundaries

- Traffic sourced from sources in the DVMRP cloud:
 - When the first packet arrives from source “S” in the DVMRP cloud at the PIM-DM router, the router will create (S,G) state using the normal PIM-DM state creation rules.
 - The incoming interface for the (S,G) entry will be computed and assuming that a route exists for this source in the DVMRP routing table (it should since the PIM router is receiving DVMRP routes), the Incoming Interface will point to the DVMRP neighbor.
 - Packets arriving from the DVMRP neighbor will now RPF correctly and be flooded out all other interfaces on the PIM-DM router where there are other PIM-DM neighbors. This permits the “push” model to work across the boundary from the DVMRP network towards the PIM network.

PIM-DM/DVMRP Boundaries

Cisco.com

- **Traffic sourced from PIM-DM cloud**
 - **When packet arrives, create (S,G) state**
(if it doesn't already exist)
 - **Use normal PIM-DM state rules plus**
 - If DVMRP neighbor has sent Poison-Reverse for route
 - » Add DVMRP neighbor's interface to (*,G) OIL.
 - » Results in interface being added to (S,G) OIL.
 - Prune interface if all DVMRP neighbors send prune.
 - **Perform normal RPF check. If passed then**
 - **Flood out all (S,G) OIL entries.**
 - **Observing any Pruned (S,G) OIL state**

(Above behavior assumes latest IOS code.)

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

39

• PIM-DM/DVMRP Boundaries

- Traffic sourced from sources in the PIM-DM cloud:
 - When the first packet arrives from source “S” in the PIM-DM cloud at the PIM/DVMRP border router, the router will create (S,G) state using the normal PIM-DM state creation rules.
 - In addition, if the border router has received a DVMRP Poison-Reverse route for this particular source network, the border router will add the DVMRP neighbor's interface to the outgoing interface list of the (*, G) and (S,G) forwarding entry. If a Prune message is received from all DVMRP neighbors on this interface, the PIM router will “Prune” the interface in the (S, G) outgoing interface list.
 - The incoming interface for the (S,G) entry will be computed based on the unicast routing table resulting in the Incoming Interface pointing to the PIM-DM neighbor in the direction of the source in the PIM cloud.
 - Packets arriving from this PIM-DM neighbor will now RPF correctly and be flooded out all interfaces in the (S, G) outgoing interface list including the interface on which the DVMRP neighbor resides. This permits the “push” model to work across the boundary from the PIM network towards the DVMRP network.

PIM-DVMRP Interoperability

Cisco.com

- DVMRP router discovery
- Basic PIM-DVMRP interaction
- PIM-DVMRP route exchange
- PIM-DVMRP RPF checking
- PIM-DVMRP congruency
- PIM-DM/DVMRP Boundaries
- **PIM-SM/DVMRP Boundaries**

PIM-SM/DVMRP Boundaries

Cisco.com

- **DVMRP Network uses “Push” model**
 - Traffic flooded everywhere
- **PIM-SM Network uses “Pull” model**
 - Traffic only sent when requested
- **Differences in models can result in problems at the PIM-SM/DVMRP boundary.**
 - Requires DVMRP “Rcvr-is-Sender” hack

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

41

- **PIM-SM/DVMRP Boundaries**

- DVMRP is a dense mode protocol which uses the “Push” model to flood traffic to all points in the network. (Routers that don’t have receivers for the traffic will then Prune the traffic flow.)
- PIM-SM uses the “Pull” model to forward traffic to only those parts of the network where it has been explicitly requested.
- Because DVMRP is using the “Push” model, while the PIM-SM network is using the “Pull” model, special considerations must be used to get traffic to flow across the PIM-SM/DVMRP boundary.
 - The PIM-SM/DVMRP solution requires the use of the “Receiver-is-a-Sender” hack in order to get traffic to flow across the PIM-SM/DVMRP boundary.

PIM-SM/DVMRP Boundaries

Cisco.com

- **Traffic sourced from DVMRP cloud**
 - **Treat as directly connected source/receiver.**
 - **Create (S,G) state (if it doesn't exist)**
 - Using normal PIM-SM state rules
 - **Proxy-Register initial (S,G) packets with RP.**
 - **Join Shared-Tree for group "G".**
 - **Operate using normal PIM-SM forwarding rules.**

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

42

- **PIM-SM/DVMRP Boundaries**

- Traffic sourced by sources in the DVMRP cloud.
 - Since the DVMRP cloud uses the "Push" model, multicast traffic will be flooded to all points in the network including the PIM-SM/DVMRP border.
 - When the first packet arrives from source "S" in the DVMRP cloud at the PIM-SM router, the router will create (S,G) state using the normal PIM-SM state creation rules.
 - The incoming interface for the (S,G) entry will be computed and assuming that a route exists for this source in the DVMRP routing table (it should since the PIM router is receiving DVMRP routes), the Incoming Interface will point to the DVMRP neighbor.
 - The PIM-SM border router will then "Register" this traffic to the RP as if it was coming from a directly connected source. This will ultimately result in the traffic flowing to the RP and on down the Shared Tree to any interested receivers for group "G" traffic in the PIM-SM network.
 - The PIM-SM border router will also include the DVMRP neighbor's interface in the (*,G) outgoing interface list and trigger a (*,G) Join toward the RP to build a branch of the Shared Tree. This causes any traffic being sent by sources in the PIM-SM network to flow down the Shared Tree to the PIM border router and be forwarded to the DVMRP router.

PIM-SM/DVMRP Boundaries

Cisco.com

- **Traffic sourced from PIM-SM cloud.**
 - Depends on DVMRP “Rcvr-is-Sender” hack.
 - A host in the DVMRP cloud must send first.
 - Causes PIM-SM border router to join Shared Tree.
 - Traffic arriving via Shared Tree is flooded to DVMRP neighbor.

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

43

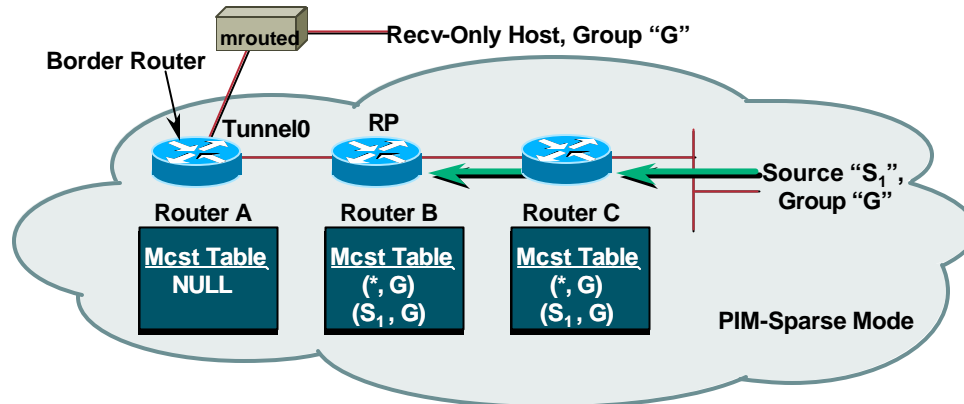
• PIM-SM/DVMRP Boundaries

- Traffic sourced by sources in the PIM-SM cloud.
 - Since the DVMRP cloud uses the “Push” model instead of the “Pull” model, the DVMRP border router has no way of knowing if there are any receivers in the DVMRP network and therefore cannot inform the PIM-SM router of this fact.
 - As a result of the differences in these two models, the only way to insure that traffic flows from PIM-SM network to the DVMRP network is by using the “Receiver-is-a-Sender” hack. This hack relies on the receiver to first send multicast traffic to the group. This traffic will be flooded to the PIM-SM/DVMRP border and will cause the border router to join the Shared Tree for the group. (See explanation on the previous page.)
 - Once a branch of the Shared Tree is setup, any traffic being sent by sources in the PIM-SM network can then flow down the Shared Tree to the PIM border router and be forwarded to the DVMRP router.

Sparse Mode Boundary Issue

Cisco.com

Problem: Receive-Only Hosts in DVMRP Cloud.



- Border router is unaware of Recv-Only Host in DVMRP cloud and therefore has no state for Group "G".
- (S₁, G) Traffic doesn't make it to Border router nor DVMRP Cloud.

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

44

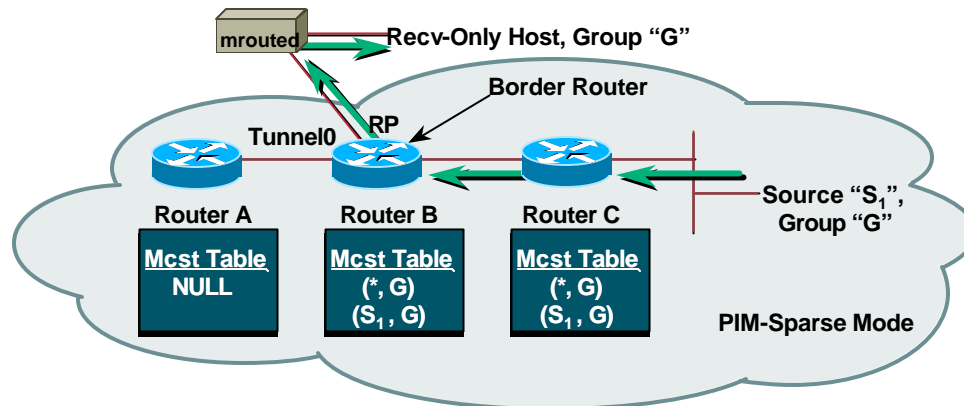
• Sparse Mode Boundary Issue

- The problem with "Receive-Only" hosts is shown in the drawing above.
 - Although the traffic from source "S" is flowing to the RP in the PIM-SM network, it is not reaching the border router, Router A. This is because Router A has no way of knowing about the "receive-only" host in the DVMRP network.

Sparse Mode Boundary Issue

Cisco.com

Solution 1: Terminate Tunnel on RP.



- Not always possible. (e.g. Multiple RPs, Multi-homing)

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

45


• Solution 1:

- One of the easiest solutions to this problem is to always terminate the DVMRP tunnel (or border) on the RP.
 - This will cause any traffic sent by sources in the PIM-SM cloud to flow to the RP (using normal PIM-SM forwarding rules) and then be “flooded” across the interface to the DVMRP neighbor.
 - Unfortunately, this solution is not always possible as there may be multiple candidate RP’s in the network or other network topology issues may preclude this approach.

Sparse Mode Boundary Issue

Cisco.com

Solution 2: Use Dense mode between Border & RP



- Need to minimize distance between Border & RP to keep Dense-mode cloud small.
- Requires carefully planning of topology.

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

46


• Solution 2:

- Another possible (but rather ugly) solution is to use PIM-DM to extend the dense mode flooding from the RP to the DVMRP neighbor.
 - This has some very obvious drawbacks including dense-mode flooding across potentially large portions of the network if the distance from the RP to the border are great.

Sparse Mode Boundary Issue

Cisco.com

Solution 3: Receiver is Sender Hack.



- Works fine for most Mbone applications since they both send and receive.
- Not applicable for applications that are Rcv-Only.

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

47

• Solution 3:

- The most often used solution is to employ the "Receiver-is-a-Sender" hack.
 - This solution normally works well for legacy multicast applications such as the Mbone multimedia conferencing applications (vic, vat, rat, wb, etc.) since these applications always send back-channel RTCP traffic to the multicast group.
 - Unfortunately, this solution can not be used reliably for applications that are truly "Receive-only".

Module Agenda

Cisco.com

- PIM-DVMRP Interoperability
- **Advanced PIM-DVMRP Features**
- Debugging Tips

Route Redistribution

Cisco.com

- **Command for route injection:**

```
ip dvmrp metric <metric> [list <acl>]
[<protocol> | dvmrp] [route-map <map>]
```

- **Metric 0 means don't inject**
- **Can select routes based on:**
 - Routing protocol
 - route-map specification
 - Enumeration using access-lists

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

49

- **Route Redistribution**

- By default, only "Connected" routes are advertised by a router in DVMRP route updates. This behavior can be overridden with the following interface command:

```
ip dvmrp metric <metric> [list <acl>]
                        [<protocol> | dvmrp]
                        [route-map <map>]
```

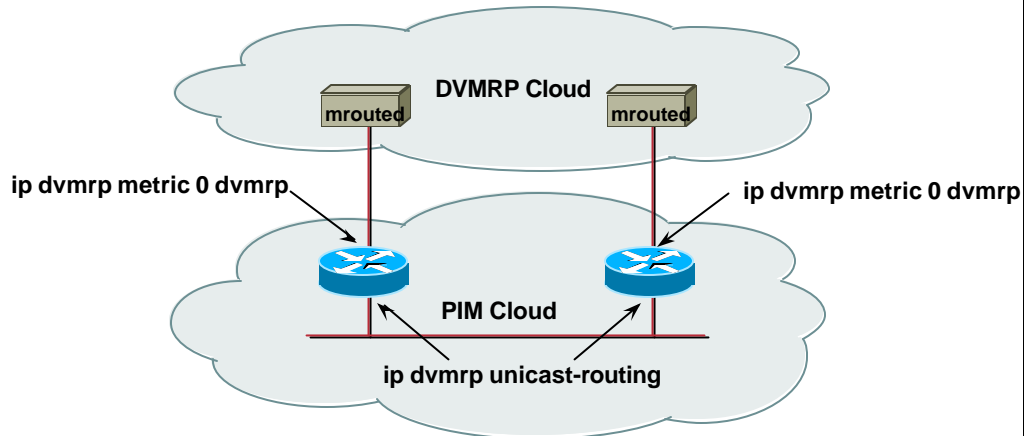
This command can be used multiple times on an interface.

- The <metric> field is in "hops" and is normally set to a value of 1. A metric of "0" has a special meaning which means any matching routes are **not** to be advertised.
 - Either ACL's or Route-Maps may be used to match the desired routes in the Unicast routing table that are to be advertised.
 - If the <protocol><process-id> is configured, only routes learned by the specified routing protocol will be advertised in DVMRP Report messages. This parameter can be used in conjunction with <access-list> so a selective list of destinations learned from a given routing protocol may be reported. If this command is not used, only directly connected networks are advertised when DVMRP neighbors are discovered.
 - If the "dvmrp" keyword is configured, only routes from the DVMRP routing table will be selected to be advertised with <metric>.
- Warning: Care should be used when configuring this command as it is easy to configure route loops whenever route redistribution is used.

Route Redistribution

Cisco.com

Preventing redistribution of DVMRP routes back into DVMRP cloud to avoid being transit network.



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

50

- **Avoid being a Transit DVMRP Network**

- It is possible to avoid becoming a DVMRP transit network by using the “**metric 0**” form of the ‘ip dvmrp metric’ interface command.
 - In the example above, the two Cisco routers in the bottom network are exchanging DVMRP routes by the use of the ‘ip dvmrp unicast-routing’ command. This would normally result in DVMRP routes learned via one router being advertised back into the DVMRP cloud which could possibly turn the bottom network into a transit DVMRP network.
 - By configuring the ‘ip dvmrp metric 0 dvmrp’ command on the border interfaces as shown above, the Cisco routers will not advertise routes from their DVMRP routing table to the DVMRP cloud. This prevents the network from becoming a possible transit network for the DVMRP network.

Route Redistribution

Cisco.com

- You can specify what you want to receive:

```
ip dvmrp accept-filter <acl> [<distance>]
```

- And from whom:

```
ip dvmrp accept-filter <acl> [neighbor-list <acl>]
 [<distance>]
```

Note: DVMRP Probes are ignored from DVMRP neighbors that are denied in the neighbor-list <acl>.

- This can be used to disable automatic PIM-DVMRP interoperability on an interface.

- And what metric to add in:

```
ip dvmrp metric-offset [in | out] <increment>
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

51

• Route Redistribution Tuning

```
ip dvmrp accept-filter <acl> [<distance>]
```

- This interface command controls which DVMRP routes will be accepted based on the specified ACL. An optional <distance> value may be used to set the Admin. Distance of the matching DVMRP routes to some other value than the default of zero.

```
ip dvmrp accept-filter <acl> [neighbor-list <acl>] [<distance>]
```

- This form of the 'accept-filter' interface command works identical to the previous version with the addition that it can be used to control which DVMRP routes will be accepted from which DVMRP neighbor. The "neighbor-list" ACL is used to accept or ignore updates from certain DVMRP neighbors.

Note: DVMRP Probes are also ignored from any DVMRP neighbors that are "denied" (either explicitly or implicitly) by The "neighbor-list" ACL. This feature can be used to disable the automatic PIM-DVMRP interoperability on an interface by configuring a "deny all" ACL for the "neighbor-list".

```
ip dvmrp metric-offset [in | out] <increment>
```

- This interface command can be used to offset the metrics of the incoming or outgoing DVMRP route updates by the amount specified by <increment>.

Route Redistribution

Cisco.com

- **Control rate at which reports are sent.**

```
ip dvmrp output-report-delay <delay> [<burst>]
```

- **Defaults: delay = 100ms; burst = 2**
- **Added IOS 11.2(9)**
- **Prior releases: delay = 0; burst = infinity**

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

52

• Route Redistribution Tuning

```
ip dvmrp output-report-delay <delay> [<burst>]
```

- This interface command was first introduced in IOS 11.2(9) and can be used to control the rate at which DVMRP route updates are transmitted. Prior to this release, Cisco routers would simply transmit all DVMRP routes in one continuous burst every DVMRP update interval (60 seconds). This often resulted in the overrun of the DVMRP neighbor's input buffers which would result in lost update messages. This in turn, would result in instabilities in the DVMRP network as routes would timeout and go into holddown.

In order to solve this problem, the 'ip dvmrp output-report-delay' command was introduced to provide pacing of the DVMRP route updates. If this command is not configured, the default values for the delay and burst size is 100ms and 2 packets.

Default Route Origination

Cisco.com

- You can originate the DVMRP default route with other routes

```
ip dvmrp default-information originate
```

- You can originate the DVMRP default route only

```
ip dvmrp default-information only
```

Good for stub networks and low-speed links—saves bandwidth and state

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

53

- **Default Route Origination**

```
ip dvmrp default-information originate
```

- This interface command controls the origination of a default DVMRP route. If this command is configured, the router will originate a DVMRP route for 0.0.0.0 out this interface in addition to the other DVMRP routes that would normally be sent.

```
ip dvmrp default-information only
```

- This interface command will cause the router to **only** originate the DVMRP default route out this interface. No other DVMRP routes will be sent.

Note: This is very useful when connecting to DVMRP stub networks as it saves on the amount of bandwidth and routing state consumed in the stub network router(s).

Summary Origination

Cisco.com

- **Classful summarization is on by default**

As long as the subnets are from a different network number than the network number of the interface the advertisement is being sent out on

- **You can turn auto summarization off**

```
no ip dvmrp auto-summary
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

54

- **Summary Origination**

- Classful summarization is on by default for DVMRP. This means that subnets from a network that are different than the network number of the interface will be automatically summarized into their classful summarization. This behavior may be disabled using the following command:

```
no ip dvmrp auto-summary
```

- This interface command turns off the classful summarization of routes across an interface.

Summary Origination

Cisco.com

- **Custom classless summarization**

```
ip dvmrp summary-address <address> <mask>  
metric <metric>
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

55

- **Summary Origination**

```
ip dvmrp summary-address <address> <mask> metric <metric>
```

- This interface command configures a summary address to be advertised out the interface. If there is at least one more specific route in the unicast routing table that matches the <address>/<mask>, the summary will be advertised. (Note: Routes in the DVMRP routing table are **not** candidates for summarization.)

When the “metric” keyword is supplied, the summary will be advertised with the metric specified by <value>. The default metric <value> is 1.

Multiple summary addresses can be configured on an interface. When multiple overlapping summary addresses are configured on an interface, the one with the longest mask takes preference. (This command was first introduced in IOS version 11.2.)

Legacy MBONE Features

Cisco.com

- **Limiting number of routes advertised**

```
ip dvmrp route-limit <route-count>
```

- **Generating “route-hog” Warnings**

```
ip dvmrp routehog-notification <route-count>
```

- **Ignore DVMRP neighbors that don't support Pruning.**

```
ip dvmrp reject-non-pruners
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

56

- **Legacy MBONE Features**

The following commands are considered “legacy” features that were primarily used when the DVMRP-based MBONE was used as the primary transit for multicast traffic across the Internet.

```
ip dvmrp route-limit <route-count>
```

- This global command limits the number of DVMRP routes advertised over an interface enabled to run DVMRP. That is a DVMRP tunnel, an interface where a DVMRP neighbor has been discovered, or an interface configured to run "ip dvmrp unicast-routing". The default value is 7000. This command will be automatically generated to the configuration file when at least one interface is enabled for multicast routing.

This command is necessary so misconfigured "ip dvmrp metric" commands don't cause massive route injection into the MBONE. The "no" version of the command configures no limit. [11.0]

```
ip dvmrp routehog-notification <route-count>
```

- This global command configures the number of routes allowed within an approximate one minute interval before a syslog message is issued warning that there maybe a route surge going on in the MBONE. This is typically used to detect quickly when someone has misconfigured their routers to inject a large number of routes into the MBONE. The default value is 10,000. You can find a running count in the "show ip igmp interface" display. When the count is exceeded, you'll see an "*** ALERT ***" string appended to the line. [10.2]

```
ip dvmrp reject-non-pruners
```

- This interface command will cause the router not to peer with a DVMRP neighbor if the neighbor doesn't support DVMRP Pruning/Grafting. This command was added so that the policy of no-support for non-Pruning DVMRP versions could be enforced in the Internet.

Module Agenda

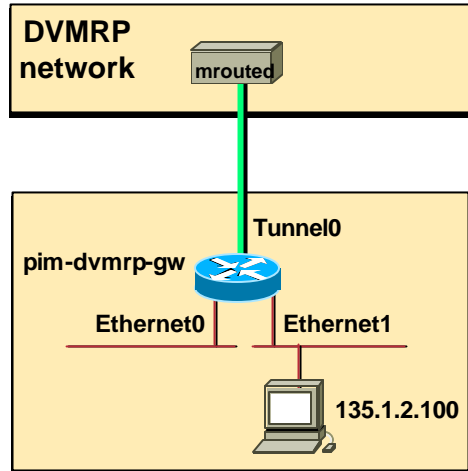
Cisco.com

- PIM-DVMRP Interoperability
- Advanced PIM-DVMRP Features
- **Debugging Tips**

Debugging Tips

Cisco.com

Example Network



pim-dvmrp-gw:

```
interface tunnel0
ip unnumbered ethernet0
ip pim dense-mode
tunnel mode dvmrp
tunnel source ethernet0
tunnel destination 135.1.22.98
```

```
interface ethernet0
ip addr 135.1.3.102 255.255.255.0
ip pim dense-mode
```

```
interface ethernet1
ip addr 135.1.2.102 255.255.255.0
ip pim dense-mode
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

58

• Example Network

- The network shown in the above drawing will be used through-out this section on debugging tips. Take particular note of **Tunnel0** source and destination addresses as well as the address of the host workstation. These addresses will be important in the upcoming pages.

Debugging Tips

Cisco.com

- **Verifying the DVMRP tunnel**
- **Verifying DVMRP route exchange**
- **Verifying Multicast reception**
- **Verifying Multicast transmission**

Verifying the DVMRP Tunnel

Cisco.com

Using the “show interface” Command

```
pim-dvmrp-gw> show int tunnel 0
Tunnel0 is up, line protocol is up
Hardware is Tunnel
Interface is unnumbered. Using address of Ethernet0 (135.1.3.102)
MTU 1500 bytes, BW 9 Kbit, DLY 500000 usec, rely 255/255, load 1/255
Encapsulation TUNNEL, loopback not set, keepalive set (10 sec)
Tunnel source 135.1.3.102 (Ethernet0), destination 135.1.22.98
Tunnel protocol/transport IP/IP (DVMRP), key disabled, sequencing
disabled
Checksumming of packets disabled, fast tunneling enabled
Last input 00:00:05, output 00:00:08, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0 (size/max/drops); Total output drops: 0
.
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

60

- **Verifying the DVMRP Tunnel**

- Use the **show interface Tunnel0** command to verify that Tunnel0 is **up** and the line protocol is also **up**. (This is normally the case as soon as the tunnel has been configured. However, it does not necessarily mean that the tunnel is operational.)

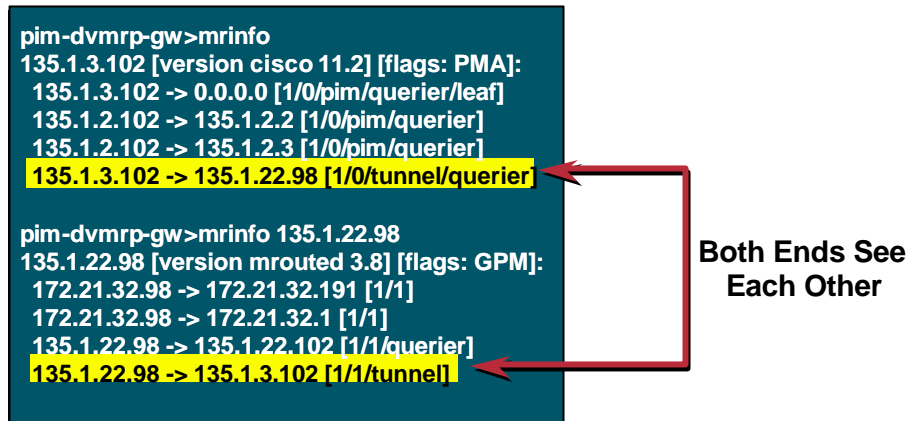
Verifying the DVMRP Tunnel

Cisco.com

Using the “mrinfo” Command

```
pim-dvmrp-gw>mrinfo
135.1.3.102 [version cisco 11.2] [flags: PMA]:
135.1.3.102 -> 0.0.0.0 [1/0/pim/querier/leaf]
135.1.2.102 -> 135.1.2.2 [1/0/pim/querier]
135.1.2.102 -> 135.1.2.3 [1/0/pim/querier]
135.1.3.102 -> 135.1.22.98 [1/0/tunnel/querier]

pim-dvmrp-gw>mrinfo 135.1.22.98
135.1.22.98 [version mrouterd 3.8] [flags: GPM]:
172.21.32.98 -> 172.21.32.191 [1/1]
172.21.32.98 -> 172.21.32.1 [1/1]
135.1.22.98 -> 135.1.22.102 [1/1/querier]
135.1.22.98 -> 135.1.3.102 [1/1/tunnel]
```



Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

61

• Verifying the DVMRP Tunnel

- The best way to verify that the tunnel is up in both directions is to use the “mrinfo” command on the Cisco PIM-DVMRP gateway router.
 - In the first example above, the “mrinfo” command is entered on the Cisco gateway router *without* any parameters. This causes the Cisco gateway router to display its own multicast interface status information.

We are looking for the tunnel interface status from the Cisco gateway router to the DVMRP router at the other end of the tunnel. This is shown on the line that reads:

```
135.1.3.120 -> 135.1.22.98 [1/0/tunnel/querier]
```

The key item to look for is the absence of the word “down” in the status. This indicates that the tunnel is up in the direction from the Cisco to the DVMRP router.

- In the second example above, the “mrinfo” command is entered on the PIM-DVMRP gateway router *with* the IP address of the DVMRP router. (I.e the tunnel destination address, 135.1.22.98.) This causes the DVMRP router at the other end of the tunnel to report back its multicast interface status information.

We are looking for the tunnel interface status from the DVMRP router at the other end of the tunnel back toward the Cisco gateway router. This is shown on the line that reads:

```
135.1.22.98 -> 135.1.3.102 [1/1/tunnel]
```

Again, the key item to look for is the absence of the word “down” in the status. This indicates that the tunnel is up in the direction from the DVMRP router to the Cisco gateway router.

- Since both ends of the tunnel are up (i.e. neither is displaying “down”), the tunnel is up in both directions.

Debugging Tips

Cisco.com

- Verifying the DVMRP tunnel
- **Verifying DVMRP route exchange**
- Verifying Multicast reception
- Verifying Multicast transmission

Verifying DVMRP Route Exchange

Cisco.com

Using the “show ip dvmrp route” Command

```
pim-dvmrp-gw# show ip dvmrp route
DVMRP Routing Table - 8 entries
130.1.0.0/16 [0/3] uptime 00:19:03, expires 00:02:13
  via 135.1.22.98, Tunnel0, [version mroute 3.8] [flags: GPM]
135.1.0.0/16 [0/3] uptime 00:19:03, expires 00:02:13
  via 135.1.22.98, Tunnel0, [version mroute 3.8] [flags: GPM]
135.1.22.0/24 [0/2] uptime 00:19:03, expires 00:02:13
  via 135.1.22.98, Tunnel0, [version mroute 3.8] [flags: GPM]
171.69.0.0/16 [0/3] uptime 00:19:03, expires 00:02:13
  via 135.1.22.98, Tunnel0, [version mroute 3.8] [flags: GPM]
172.21.27.0/24 [0/3] uptime 00:19:04, expires 00:02:12
  via 135.1.22.98, Tunnel0, [version mroute 3.8] [flags: GPM]
172.21.32.0/24 [0/2] uptime 00:19:04, expires 00:02:12
  via 135.1.22.98, Tunnel0, [version mroute 3.8] [flags: GPM]
172.21.33.0/24 [0/3] uptime 00:19:04, expires 00:02:12
  via 135.1.22.98, Tunnel0, [version mroute 3.8] [flags: GPM]
172.21.120.0/24 [0/3] uptime 00:19:04, expires 00:02:12
  via 135.1.22.98, Tunnel0, [version mroute 3.8] [flags: GPM]
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

63

- **Verifying DVMRP route exchange**

- The easiest way to verify that DVMRP routes are being exchanged is to use the “**show ip dvmrp route**” command on the Cisco PIM-DVMRP gateway router.
 - In the example above, the “**show ip dvmrp route**” command is entered on the Cisco gateway router. This causes the Cisco gateway router to display the contents of its DVMRP route table. In this case, we see that the Cisco gateway router has learned 8 DVMRP routes from the DVMRP neighbor at the other end of the tunnel. This is a pretty good indication that DVMRP routes are being exchanged via the tunnel.

Verifying DVMRP Route Exchange

Cisco.com

Using the “debug ip dvmrp” Command

```
pim-dvmrp-gw# debug ip dvmrp
DVMRP debugging is on
pim-dvmrp-gw#
Mar 20 11:39:36.335: DVMRP: Aging routes, 0 entries expired
Mar 20 11:39:41.271: DVMRP: Received Probe on Tunnel0 from 135.1.22.98
Mar 20 11:39:45.335: DVMRP: Building Report for Tunnel0 224.0.0.4
Mar 20 11:39:45.335: DVMRP: Send Report on Tunnel0 to 135.1.22.98
Mar 20 11:39:45.335: DVMRP: 2 unicast, 8 DVMRP routes advertised
Mar 20 11:39:47.335: DVMRP: Aging routes, 0 entries expired
Mar 20 11:39:51.371: DVMRP: Received Probe on Tunnel0 from 135.1.22.98
Mar 20 11:39:52.379: DVMRP: Received Report on Tunnel0 from 135.1.22.98
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

64

• Verifying DVMRP route exchange

- Another way to verify that DVMRP routes are being exchanged is to use the “**debug ip dvmrp**” command on the Cisco PIM-DVMRP gateway router.
 - In the example above, the “**debug ip dvmrp**” command is entered on the Cisco gateway router. This causes the Cisco gateway router to display the DVMRP routing protocol events.
 - In the first highlighted section, we see that the Cisco gateway router build and send a DVMRP Report message that contains 8 DVMRP routes and 2 unicast routes to the DVMRP router at the other end of the tunnel.

Note: The 8 DVMRP routes being sent are actually Poison Reverse routes for the 8 DVMRP routes received from the DVMRP neighbor. The 2 unicast routes are local unicast routes (in this case the directly connected networks on the Cisco gateway router) that are being advertised to the DVMRP neighbor in the DVMRP Report. However, this information is not obvious unless we turn on more debugging detail.

- In the second highlighted section, we see that the Cisco gateway router has received a DVMRP Report via Tunnel0 from the DVMRP neighbor, 135.1.22.98. Unfortunately, with this level of debug we do not know how many routes were contained in the Report.

Verifying DVMRP Route Exchange

Cisco.com

Checking DVMRP Routes being Advertised

```
pim-dvmrp-gw# debug ip dvmrp detail
DVMRP debugging is on
Mar 20 11:42:45.337: DVMRP: Building Report for Tunnel0 224.0.0.4
Mar 20 11:42:45.337: DVMRP: Report 130.1.0.0/16, metric 35, from DVMRP table
Mar 20 11:42:45.337: DVMRP: Report 135.1.0.0/16, metric 35, from DVMRP table
Mar 20 11:42:45.337: DVMRP: Report 135.1.22.0/24, metric 34, from DVMRP table
Mar 20 11:42:45.337: DVMRP: Report 171.69.0.0/16, metric 35, from DVMRP table
Mar 20 11:42:45.337: DVMRP: Report 172.21.27.0/24, metric 35, from DVMRP table
Mar 20 11:42:45.337: DVMRP: Report 172.21.32.0/24, metric 34, from DVMRP table
Mar 20 11:42:45.337: DVMRP: Report 172.21.33.0/24, metric 35, from DVMRP table
Mar 20 11:42:45.337: DVMRP: Report 172.21.120.0/24, metric 35, from DVMRP table
Mar 20 11:42:45.337: DVMRP: Report 135.1.2.0/24, metric 1
Mar 20 11:42:45.337: DVMRP: Report 135.1.3.0/24, metric 1
Mar 20 11:42:45.337: DVMRP: Send Report on Tunnel0 to 135.1.22.98
Mar 20 11:42:45.337: DVMRP: 2 unicast, 8 DVMRP routes advertised
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

65

• Verifying DVMRP route exchange

- In order to get more detailed information, it is necessary to enter the “**debug ip dvmrp detail**” command on the Cisco PIM-DVMRP gateway router. This will result in the Cisco router reporting detailed contents of each sent and received DVMRP route report. (Warning: Care should be taken when using this debug command as the router may be overloaded if the number of DVMRP routes in these reports is large.)
 - In the example above, the “**debug ip dvmrp detail**” command is entered on the Cisco gateway router. This causes the Cisco gateway router to display the DVMRP routing protocol events as well as the contents of the DVMRP Report messages.
 - In the example above, we see the Cisco gateway router build and send a DVMRP Report message that contains 8 DVMRP routes and 2 unicast routes to the DVMRP router at the other end of the tunnel.

Notice that the 8 DVMRP routes have a metric greater than 32 (infinity) which indicates that these routes are Poison-Reverse routes. The two unicast routes being advertised in this DVMRP report are highlighted in the above example. Notice that these are the two “connected” networks on the Cisco gateway router. (Refer to the Example network drawing in a previous slide.)

Verifying DVMRP Route Exchange

Cisco.com

Checking DVMRP Routes being Received

```
pim-dvmrp-gw# debug ip dvmrp detail
DVMRP debugging is on
... : DVMRP: Received Report on Tunnel0 from 135.1.22.98
... : DVMRP: Origin 130.1.0.0/16, metric 2, metric-offset 1, distance 0
... : DVMRP: Origin 135.1.0.0/16, metric 2, metric-offset 1, distance 0
... : DVMRP: Origin 171.69.0.0/16, metric 2, metric-offset 1, distance 0
... : DVMRP: Origin 135.1.2.0/24, metric 34, metric-offset 1, infinity
... : DVMRP: Origin 135.1.3.0/24, metric 34, metric-offset 1, infinity
... : DVMRP: Origin 135.1.22.0/24, metric 1, metric-offset 1, distance 0
... : DVMRP: Origin 172.21.27.0/24, metric 2, metric-offset 1, distance 0
... : DVMRP: Origin 172.21.32.0/24, metric 1, metric-offset 1, distance 0
... : DVMRP: Origin 172.21.33.0/24, metric 2, metric-offset 1, distance 0
... : DVMRP: Origin 172.21.120.0/24, metric 2, metric-offset 1, distance 0
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

66

- **Verifying DVMRP route exchange**

- In the example above, we again see an example of some “**debug ip dvmrp detail**” output from the Cisco PIM-DVMRP gateway router.
 - In this example, we see the Cisco gateway router has received a DVMRP Report message that contains a total of 10 routes. Notice that the 8 of the routes have a metric less than 32 (infinity) which indicates that these routes are normal DVMRP routes that are being advertised by the DVMRP neighbor. Notice also the two routes with a metric of 34. These are Poison-Reverse routes being sent from the DVMRP neighbor back toward the Cisco router. These are the two unicast routes that were originally advertised to the DVMRP neighbor by the Cisco gateway router. (Refer to the Example network drawing in a previous slide.)

Debugging Tips

Cisco.com

- Verifying the DVMRP tunnel
- Verifying DVMRP route exchange
- **Verifying Multicast reception**
- Verifying Multicast transmission

Verifying Multicast Reception

Cisco.com

- **Use SDR multicasts (224.2.127.254) as a test signal**
 - **Enable “ip sdr listen” on an interface**
 - Use “ip sap listen” on newer versions of IOS
 - **Should begin seeing entries in mroute table**
 - Use “show ip mroute summary” to list
 - **Should begin seeing SDR cache entries**
 - Use “show ip sdr” to list
 - Use “debug ip sd” to observe SDR packets

Note: Assumes SDR is being used in the old DVMRP cloud to announce sessions.

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM



• Verifying Multicast Reception

- The key to verifying multicast reception is to first know of an active source somewhere in the network that can be used as a multicast beacon. As far as applications go, SDR is almost always active somewhere in the Internet and often is active in other multicast networks as well. Therefore, if there are no other known active multicast applications, the SDR application is often a good choice.
 - In order to use the SDR application as a multicast test beacon, the **ip sdr listen** interface command can be configured on one of the router interfaces. (Note: On newer versions of IOS, the command has been renamed to **ip sap listen**.) This will cause the router to join the SDR group (224.2.127.254) and begin receiving any possible SDR sources in the network.
 - Once the router has joined the SDR multicast group, you should begin see (S,G) entries in the mroute table for sources sending to this group. This will confirm that you are receiving multicast traffic.
 - In addition to the mroute entries, the router will cache SDR Session Announcements for advertised multimedia sessions. The contents of this cache can be displayed by using the **show ip sdr** command. (Alternatively, you can enable **debug ip sd** and see these Sessions Announcements being processed by the router.) This will also confirm that you are receiving multicast traffic.

Verifying Multicast Reception

Cisco.com

Should begin seeing entries in mroute table

```
pim-dvmrp-gw# show ip mroute summary 224.2.127.254
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P -
Pruned
      R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.2.127.254), 00:08:07/00:02:56, RP 0.0.0.0, flags: SJC
(128.32.131.87/32, 224.2.127.254), 00:02:30/00:00:26, flags: CT
(129.89.142.50/32, 224.2.127.254), 00:02:35/00:00:19, flags: CT
(171.69.58.109/32, 224.2.127.254), 00:02:40/00:00:21, flags: CT
.
.
.
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM



• Verifying Multicast Reception

- The slide above is an example of the **show ip mroute summary** command being used to display a list of (S,G) entries for the SDR multicast group (224.2.127.254). This tells us that we have successfully received multicast traffic from these sources which in turn verifies that we have multicast reception.

Verifying Multicast Reception

Cisco.com

Should begin seeing SDR cache entries

```
pim-dvmrp-gw# show ip sdr
SDR Cache - 249 entries

Michigan State University Instructional Television
!CannesCast '97 - 50th anniversary
Aberdeen University, Scotland
ACM 97
Alan Kay: Georgia Tech Distinguished Lecture
AmiNet
Argonne Petroleum Seminar Series
as/cd discussions
ATM setup Mannheim-Bonn-Berkeley
Audio for Sunday
Basler Fasnacht 1997 !
.
.
.
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

70

- **Verifying Multicast Reception**

- The slide above is an example of the **show ip sdr** command being used to display a list of the names of the multi-media sessions learned via the SDR (multicast group 224.2.127.254) entries for the SDR multicast group (224.2.127.254). This tells us that we have successfully received Session Announcement multicast traffic from the sources advertising the sessions. This in turn, verifies that we have multicast reception.

Debugging Tips

Cisco.com

- Verifying the DVMRP tunnel
- Verifying DVMRP route exchange
- Verifying Multicast reception
- **Verifying Multicast transmission**

Verifying Multicast Transmission

Cisco.com

- **Use “sdr” & “rat” MBONE applications**
 - Run “sdr” on WS or PC
 - Join a “Well-Known” rat session
- **Check mroute table at “pim-dvmrp-gw”**
 - Should have (S, G) entry for WS or PC

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

72

• Verifying Multicast Transmission

- Verifying multicast transmission via the DVMRP tunnel is a little more complex. First it is necessary to activate a source somewhere inside your network for a multicast group that you know has active receivers in the DVMRP network. Again, SDR is almost always active somewhere in the Internet and often is active in other multicast networks as well. Therefore, if there are no other known active multicast applications, the SDR application is often a good choice. Another choice is to use the SDR application to identify an active multimedia session that may have active receivers. The next step is to activate a workstation or PC and run the SDR application or the RAT (Robust Audio Tool) to source multicast traffic from your network.
 - In order to use the SDR application as a multicast test source you need to activate the SDR application on a PC or workstation and *create* a new multimedia session. This will cause the PC/WS to begin sourcing SDR Session Announcements to the SDR multicast group (224.2.127.254). Assuming that there are active SDR receivers in the DVMRP network (or beyond), this SDR multicast traffic will begin flowing across the tunnel. This should be making sure that there is an (S,G) mroute entry for the PC/WS on the PIM-DVMRP gateway router. Furthermore, the Tunnel interface should appear in the outgoing interface list in an unpruned (Forwarding) state.
 - An alternative to using SDR is to launch the RAT (Robust Audio Tool) audio conferencing tool for an active multi-media audio session that is being advertised by SDR. Assuming that there are active receivers in the DVMRP network (or beyond) for this session this RAT multicast traffic will begin flowing across the tunnel. This should be making sure that there is an (S,G) mroute entry for the PC/WS on the PIM-DVMRP gateway router. Furthermore, the Tunnel interface should appear in the outgoing interface list in an unpruned (Forwarding) state.

Verifying Multicast Transmission

Cisco.com

Should Have (S, G) Entry for WS or PC

```
pim-dvmrp-gw> show ip mroute summary
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
       R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.2.0.1), 00:08:07/00:02:56, RP 0.0.0.0, flags: DJC
(13.2.116.11/32, 224.2.0.1), 08:11:47/00:02:55, flags: PCT
(128.16.64.19/32, 224.2.0.1), 16:05:41/00:02:52, flags: PCT
(129.99.50.40/32, 224.2.0.1), 1d08h/00:02:57, flags: PCT
(134.164.1.2/32, 224.2.0.1), 01:42:58/00:02:57, flags: PCT
(135.1.2.100/32, 224.2.0.1), 00:05:40/00:00:43, flags: CLT
(138.25.8.74/32, 224.2.0.1), 13:16:05/00:02:56, flags: PCT
(171.69.58.109/32, 224.2.127.254), 00:02:40/00:00:21, flags: PCT
```

Module9.ppt

©1998 – 2001, Cisco Systems, Inc. All rights reserved.

8/14/2001 2:11 PM

73

• Verifying Multicast Transmission

- The slide above is an example of the **show ip mroute summary** command being used to verify that an (S,G) entry for the sending PC/WS for the RAT multi-media session exists in the PIM-DVMRP gateway.
 - Notice the highlighted entry which corresponds to the IP address of the WS in the initial example network drawing. This tells us that we are successfully sending multicast traffic from this WS to the PIM-DVMRP gateway. In order to verify that the WS traffic is flowing across the Tunnel, the outgoing interface list of the highlighted entry would be checked. If the interface Tunnel0 is in the outgoing interface list in “Forwarding” state, it verifies that we are sending multicast traffic across the tunnel.

