

Moonv6 Test Suite

*Protocol Independent Multicast – Sparse Mode
InterOperability Test Suite*

Technical Document

Revision 0.1



*IPv6 Consortium
InterOperability Laboratory
Research Computing Center
University of New Hampshire*

*121 Technology Drive, Suite 2
Durham, NH 03824-3525
Phone: (603) 862-2804
Fax: (603) 862-4181*

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	4
INTRODUCTION	5
TIMERS AND DEFAULT VALUES	7
TEST ORGANIZATION.....	8
REFERENCES	9
GROUP 1: Hello Messages and Designated Router Election	10
Test PIM.1.1: Sending Hello Messages.....	11
Test PIM.1.2: DR Election.....	12
Test PIM.1.3: Register Stop.....	14
GROUP 2: RP and SPT Tree	16
Test PIM.2.1: RP Tree	17
Test PIM.2.2: Shortest Path Tree	19
GROUP 3: Assert Messages	21
Test PIM.3.1: Assert Messages.....	22

MODIFICATION RECORD

Draft Version Complete

January 29, 2004

ACKNOWLEDGEMENTS

The University of New Hampshire would like to acknowledge the efforts of the following individuals in the development of this test suite.

Timothy Winters University of New Hampshire

INTRODUCTION

Overview

The University of New Hampshire's InterOperability Laboratory (IOL) is an institution designed to improve the interoperability of standards-based products by providing an environment where a product can be tested against other implementations of a standard. This suite of tests has been developed to help implementers evaluate the functioning of their Internet Protocol, version 6 products. The tests do not determine if a product conforms to the IPv6 specifications, nor are they purely interoperability tests. Rather, they provide one method to isolate problems within a device. Successful completion of all tests contained in this suite does not guarantee that the tested device will interoperate with other IPv6 devices. However, combined with satisfactory operation in the IOL's semi-production environment, these tests provide a reasonable level of confidence that the Node Under Test will function well in most multi-vendor IPv6 environments.

Acronyms

RUT: Router Under Test

TR: Testing Router

RP: Rendezvous Point

DR: Designated Router

IGMP: Internet Group Management Protocol

MRIB: Multicast Routing Information Base

RPF: Reverse Path Forwarding

TIB: Tree Information Base

MFIB: Multicast Forwarding Information Base

When several entities of the same type are present in a test configuration, a number is appended to the acronym to yield a label for each entity. For example, if there were three testing routers in the test configuration, they would be labeled TR1, TR2 and TR3.

Definitions

Rendezvous Point (RP):

An RP is a router that has been configured to be used as the root of the non-source-specific distribution tree for a multicast group. Join messages from receivers for a group are sent towards the RP, and data from senders is sent to the RP so that receivers can discover who the senders are, and start to receive traffic destined for the group.

Designated Router (DR):

A shared-media LAN like Ethernet may have multiple PIM-SM routers connected to it. If the LAN has directly connected hosts, then a single one of these routers, the DR, will act on behalf of those hosts with respect to the PIM-SM protocol. A single DR is elected per interface (LAN or otherwise) using a simple election process.

MRIB Multicast Routing Information Base:

This is the multicast topology table, which is typically derived from the unicast routing table, or routing protocols such as MBGP that carry multicast-specific topology information. In PIM-SM, the MRIB is used to decide where to send Join/Prune messages. A secondary function of the MRIB is to provide routing metrics for destination addresses, these metrics are used when sending and processing Assert messages.

RPF Neighbor:

RPF stands for "Reverse Path Forwarding". The RPF Neighbor of a router with respect to an address is the neighbor that the MRIB indicates should be used to forward packets to that address. In the case of a PIM-SM multicast group, the RPF neighbor is the router that a Join message for that group would be directed to, in the absence of modifying Assert state.

TIB Tree Information Base:

This is the collection of state at a PIM router that has been created by receiving PIM Join/Prune messages, PIM Assert messages, and IGMP or MLD information from local hosts. It essentially stores the state of all multicast distribution trees at that router.

MFIB Multicast Forwarding Information Base:

The TIB holds all the state that is necessary to forward multicast packets at a router. However, although this specification defines forwarding in terms of the TIB, to actually forward packets using the TIB is very inefficient. Instead a real router implementation will normally build an efficient MFIB from the TIB state to perform forwarding. How this is done is implementation-specific, and is not discussed in this document.

Upstream

Towards the root of the tree. The root of tree may either be the source or the RP depending on the context.

Downstream

Away from the root of the tree.

TIMERS AND DEFAULT VALUES

PIM-SM defines several timers and default values. For the purpose of testing, all configurable timers and values are set to their defaults, unless otherwise noted in the test description. These defaults are given here for reference, taken or calculated from draft-ietf-pim-sm-v2-new-05; Section 4.12:

Hello_Period:	30 seconds
Triggered_Hello_Delay:	5 seconds
Holdtime:	105 seconds
Assert_Time:	180 seconds
Assert_Override_Interval:	3 seconds
Register_Suppression_Time:	60 seconds
t_override:	2.5 seconds
t_periodic:	60 seconds

TEST ORGANIZATION

This document organizes tests by group based on related test methodology or goals. Each group begins with a brief set of comments pertaining to all tests within that group. This is followed by a series of description blocks; each block describes a single test. The format of the description block is as follows:

Test Label:	The test label and title comprise the first line of the test block. The test label is composed by concatenating the short test suite name, the group number, and the test number within the group, separated by periods. The Test Number is the group and test number, also separated by a period. So, test label PIM.1.2 refers to the second test of the first test group in the PIM-SM suite. The test number is 1.2.
Purpose:	The Purpose is a short statement describing what the test attempts to achieve. It is usually phrased as a simple assertion of the feature or capability to be tested.
References:	The References section lists cross-references to the specifications and documentation that might be helpful in understanding and evaluating the test and results.
Resource Requirements:	The Resource Requirements section specifies the software, hardware, and test equipment that will be needed to perform the test.
Discussion:	The Discussion is a general discussion of the test and relevant section of the specification, including any assumptions made in the design or implementation of the test as well as known limitations.
Test Setup:	The Test Setup section describes the configuration of all devices prior to the start of the test. Different parts of the procedure may involve configuration steps that deviate from what is given in the test setup. If a value is not provided for a protocol parameter, then the protocol's default is used for that parameter.
Procedure:	This section of the test description contains the step-by-step instructions for carrying out the test. These steps include such things as enabling interfaces, unplugging devices from the network, or sending packet from a test station. The test procedure also cues the tester to make observations, which are interpreted in accordance with the observable results given for that test part.
Observable Results:	This section lists observable results that can be examined by the tester to verify that the RUT is operating properly. When multiple observable results are possible, this section provides a short discussion on how to interpret them. The determination of a pass or fail for each test is usually based on how the RUT's behavior compares to the results described in this section.
Possible Problems:	This section contains a description of known issues with the test procedure, which may affect test results in certain situations.

REFERENCES

The following documents are referenced in this text:

- draft-ietf-pim-sm-v2-new-05 – Protocol Independent Multicast – Sparse Mode (PIM-SM): protocol specification
- S. Deering, W. Fenner, B. Haberman, Multicast Listener Discovery (MLD) for IPv6, RFC 2710, October 1999.

GROUP 1: Hello Messages and Designated Router Election

Scope:

These tests are designed to verify conformance with Hello Messaging and DR Election for PIM-SM.

Overview:

PIM routers transmit Hello Messages to notify other PIM-SM routers of their presence on a network. Designated Router Election is performed in order to elect a router to forward multicast data on a given subnet. Hello messages are also the way that option negotiation takes place in PIM, so that additional functionality can be enabled, or parameters tuned.

Test PIM.1.1: Sending Hello Messages

Purpose: To verify that a router properly transmits Hello messages.

References: [draft-ietf-pim-sm-v2-new-08] Sections 4.3.1 and 4.12

Resource Requirements:


- Monitor to capture packets

Last Modification: January 29, 2004

Discussion: PIM-SM-Hello messages are sent periodically on each PIM-enabled interface. Hello messages MUST be sent on all active interfaces, including physical point-to-point links, and are multicast to address ff02::d (the ALL-PIM-ROUTERS group). A per interface hello timer (HT(I)) is used to trigger sending Hello messages on each active interface. Hello messages must be sent every Hello_Period seconds. The hello timer should not be reset except when it expires.

Test Setup:

- Ensure the TR1 and TR2 has a static configuration RP.
- The multicast group of the RP is FF02::1:3.
- PIM-SM should not be running on any router.



Procedure:

Part A: Default Hello_Period

1. Enable PIM-SM on the RP, TR1, and TR2.
2. Observe the packets transmitted by the RUT on network 0.

Observable Results:

- In Part A, the RUT should transmit a Hello Message to the ALL-PIM-ROUTERS multicast address ff02::d on network 0 every 30 seconds (Hello_Period).

Possible Problems: None

Test PIM.1.2: DR Election

Purpose: To verify that a router properly performs DR Election

References: [draft-ietf-pim-sm-v2-new-08] Sections 4.3.1 and 4.3.2

Resource Requirements:

- Monitor to capture packets
- Packet Generator

Last Modification: January 30, 2004


Discussion: The `DR_Election_Priority` Option allows the network administrator to allow a particular router to win the DR election process by giving it a numerically larger DR Election Priority. The `DR_Election_Priority` Option SHOULD be included in every Hello message, even if no DR election priority is explicitly configured on that interface. This is necessary because priority-based DR election is only enabled when all neighbors on an interface advertise that they are capable of using the DR Election Priority Option. The default priority is 1.

The router with the highest DR Priority becomes the DR for the subnet. The router with the highest IP address becomes DR if there is a tie in DR Priority, the DR Priority is not configured, or the DR option is not present. Priority based DR election is only enabled when all neighbors on an interface advertise they are capable of using DR Election Priority Option.

The DR election priority is a 32-bit unsigned number and the numerically larger priority is always preferred. A router's idea of the current DR on an interface can change when a PIM-Hello message is received, when a neighbor times out, or when a router's own DR priority changes. If the router becomes the DR or ceases to be the DR, this will normally cause the DR Register state machine to change state. Subsequent actions are determined by that state-machine.

Test Setup:

- TR1 and TR2 has a static configuration for the address of the RP.
- The multicast group of the RP is FF02::1:3.
- G1 is a packet generator.
- OSPF is running on all interfaces with a metric of 1.
- If a network is not assigned to a particular area, the default area of 0.0.0.0 is used.



Procedure:

Part A: TR1 has higher DR Priority

1. TR1 is configured to have a DR priority of 3.

2. TR2 is configured to have a DR priority of 1.
3. On network 0, the RP, TR1, and TR2 transmit Hello messages.
4. On network 1, TR1 and TR2 transmit Hello messages.
5. On network 1, the G1 transmits data packets with a multicast destination of FF02::1:3.
6. Observe packets transmitted on all networks.

Part B: TR2 has lower DR Priority

7. TR1 is configured to have a DR priority of 1.
8. TR2 is configured to have a DR priority of 3.
9. On network 0, the RP, TR1, and TR2 transmit Hello messages.
10. On network 1, TR1 and TR2 transmit Hello messages.
11. On network 1, the G1 transmits data packets with a multicast destination of FF02::1:3.
12. Observe packets transmitted on all networks.

Part C: Change DR

13. TR1 is configured to have a DR priority of 2.
14. TR2 is configured to have a DR priority of 3.
15. On network 0, the RP, TR1, and TR2 transmit Hello messages.
16. On network 1, TR1 and TR2 transmit Hello messages.
17. On network 1, the G1 transmits data packets with a multicast destination of FF02::1:3.
18. Configure TR2 to have a DR priority of 1.
19. Observe packets transmitted on all networks.

Observable Results:

- In Part A, TR1 should be elected DR and encapsulate the multicast data packets on network 1 and forward them to the RP on network 0.
- In Part B, TR2 should be elected DR and encapsulate the multicast data packets on network 1 and forward them to the RP on network 0.
- In Part C, TR1 should be elected DR and encapsulate the multicast data packets on network 1 and forward them to the RP on network 0. After TR2 changes DR priority and is elected DR, TR1 should stop encapsulating data. TR2 should begin encapsulating data and forwarding data to the RP on network 0.

Possible Problems: A router may not implement the DR_Election_Priority Option.

Test PIM.1.3: Register Stop

Purpose: To verify that a router properly accepts a RegisterStop message.

References: [draft-ietf-pim-sm-v2-new-08] Section 3

Resource Requirements:


- Monitor to capture packets
- Packet Generator

Last Modification: January 30, 2004

Discussion: When the RP receives a register-encapsulated data packet from source S on group G, it will normally initiate an (S,G) source-specific Join towards S. This Join message will travel hop-by-hop towards S, instantiating (S,G) multicast tree state in the routers along the path. (S,G) multicast tree state is used only to forward packets for group G to those packets that come from source S. While the RP is in the process of joining the source-specific tree for S, the data packets will continue being encapsulated to the RP. When packets from S also start to arrive natively at the RP, the RP will be receiving two copies of each of these packets. At this point, the RP starts to discard the encapsulated copy of these packets, and it sends a RegisterStop message back to S's DR to prevent the DR unnecessarily encapsulating the packets.

Test Setup:

- TR1 and TR2 has a static configuration for the address of the RP.
- The multicast group of the RP is FF02::1:3.
- G1 is a packet generator.
- OSPF is running on all interfaces.
- If a network is not assigned to a particular area, the default area of 0.0.0.0 is used.



Procedure:

1. TR1 is configured to have a DR priority of 1.
2. TR2 is configured to have a DR priority of 3.
3. On network 0, TR1, TR2 and the RP transmit Hello messages.
4. On network 1, TR1 and TR2 transmit Hello messages.
5. On network 0, the G1 transmits data packets with a multicast group address of ff02::1:3. The TR1 should encapsulate these packets and forward them to the RP.
6. The RP should transmit a Join (S,G) with the source address of G1 and the group address of FF02::1:3.
7. On network 0, TR1 should start to transmit data packets from network 1.

8. After the RP receives both the encapsulated and native data packets, the RP transmits a RegisterStop to the TR1 on network 1.
9. Observe the packets transmitted by the RUT on all networks.

Observable Results:

- Upon receiving the RegisterStop, the RUT should stop forwarding the encapsulated data packets. The native data packets should continue to be forwarded onto network 1.

Possible Problems: The RP may never send a Join (S,G).

GROUP 2: RP and SPT Tree

Scope:

These tests are designed to verify conformance in forming RP and Shortest Path Trees using PIM-SM.

Overview:

In phase one, a multicast receiver expresses its interest in receiving traffic destined for a multicast group. Typically it does this using IGMP [6] or MLD [4], but other mechanisms might also serve this purpose. One of the receiver's local routers is elected as the Designated Router (DR) for that subnet. On receiving the receiver's expression of interest, the DR then sends a PIM Join message towards the RP for that multicast group. This Join message is known as a (*,G) Join because it joins group G for all sources to that group. The (*,G) Join travels hop-by-hop towards the RP for the group, and in each router it passes through, multicast tree state for group G is instantiated. Eventually the (*,G) Join either reaches the RP, or reaches a router that already has (*,G) Join state for that group. When many receivers join the group, their Join messages converge on the RP, and form a distribution tree for group G that is rooted at the RP. This is known as the RP Tree (RPT), and is also known as the shared tree because it is shared by all sources ending to that group. Join messages are resent periodically so long as the receiver remains in the group. When all receivers on a leaf-network leave the group, the DR will send a PIM (*,G) Prune message towards the RP for that multicast group. However if the Prune message is not sent for any reason, the state will eventually time out.

A multicast data sender just starts sending data destined for a multicast group. The sender's local router (DR) takes those data packets, unicast-encapsulates them, and sends them directly to the RP. The RP receives these encapsulated data packets, decapsulates them, and forwards them onto the shared tree. The packets then follow the (*,G) multicast tree state in the routers on the RP Tree, being replicated wherever the RP Tree branches, and eventually reaching all the receivers for that multicast group. The process of encapsulating data packets to the RP is called registering, and the encapsulation packets are known as PIM Register packets.

At the end of phase one, multicast traffic is flowing encapsulated to the RP, and then natively over the RP tree to the multicast receivers

Test PIM.2.1: RP Tree

Purpose: To verify that a routers properly form an RP tree

References: [draft-ietf-pim-sm-v2-new-08] Section 3

Resource Requirements:

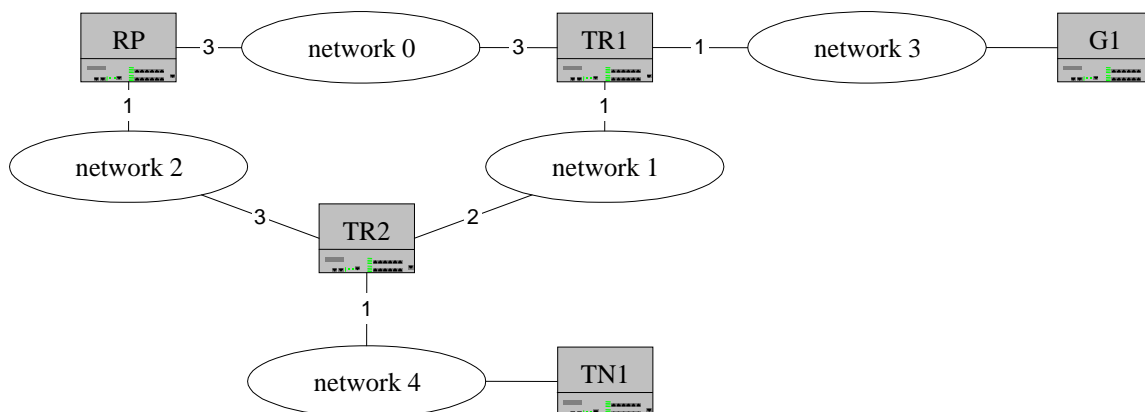
- Monitor to capture packets
- Packet Generator

Last Modification: January 30, 2004

Discussion: One of the receiver's local routers is elected as the Designated Router (DR) for that subnet. On receiving the receiver's expression of interest, the DR then sends a PIM Join message towards the RP for that multicast group. This Join message is known as a (*,G) Join because it joins group G for all sources to that group. The (*,G) Join travels hop-by-hop towards the RP for the group, and in each router it passes through, multicast tree state for group G is instantiated. Eventually the (*,G) Join either reaches the RP, or reaches a router that already has (*,G) Join state for that group. When many receivers join the group, their Join messages converge on the RP, and form a distribution tree for group G that is rooted at the RP. This is known as the RP Tree (RPT), and is also known as the shared tree because it is shared by all sources ending to that group.

Test Setup:

- TR1 and TR2 has a static configuration for the address of the RP.
- The multicast group of the RP is FF02::1:3.
- OSPF is running on all interfaces.
- If a network is not assigned to a particular area, the default area of 0.0.0.0 is used.



Procedure:

1. On network 0, TR1 and the RP transmit Hello messages.
2. On network 1, TR1 and TR2 transmit Hello messages.

3. On network 2, TR2 and the RP transmit Hello messages.
4. On network 3, TR1 transmits Hello messages.
5. On network 4, TR2 transmits Hello messages.
6. On network 3, G1 transmit multicast data for the group ff02::1:3.
7. TR1 encapsulates the multicast data and forward it to the RP on network 0.
8. On network 0, the RP transmits a Join (S,G) with an upstream neighbor of TR1, a source address of G1 and a multicast address of G1.
9. TR1 transmits multicast data packets from network 3 to network 0.
10. On network 0, the RP transmit a Register Stop with a source of G1 and a multicast group ff02::1:3.
11. On network 4, TN1 transmits a MLD report for the multicast group ff02::1:3.
12. Observe the packets transmitted by the RUT on all networks.

Observable Results:

- TR2 should transmit a Join (*,G) with an upstream neighbor address of RP and a multicast group of ff02::1:3. The RP should then forward the multicast data onto network 2. TR2 should forward that data from network 2 to network 4.

Possible Problems: None.

Test PIM.2.2: Shortest Path Tree

Purpose: To verify that a routers properly form an SPT tree

References: [draft-ietf-pim-sm-v2-new-08] Section 3

Resource Requirements:

- Monitor to capture packets
- Packet Generator

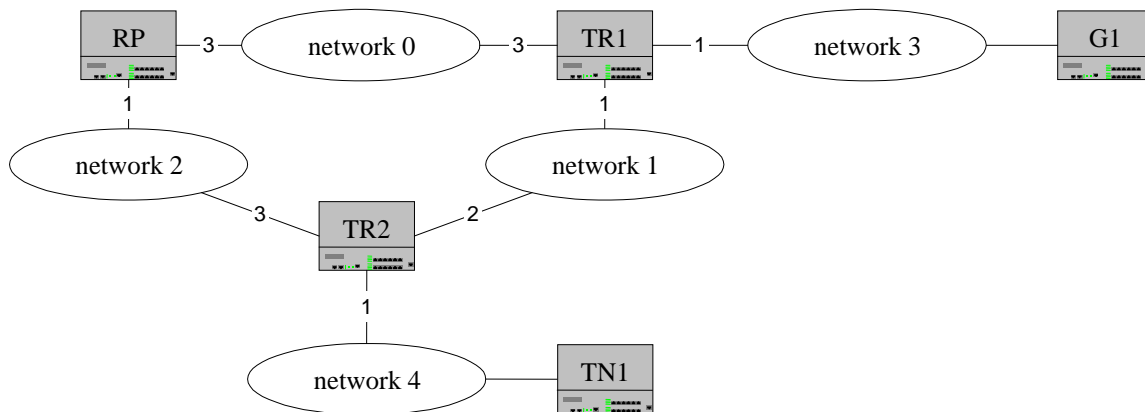
Last Modification: January 30, 2004

Discussion: To obtain lower latencies, a router on the receiver's LAN, typically the DR, may optionally initiate a transfer from the shared tree to a source-specific shortest-path tree (SPT). To do this, it issues an (S,G) Join towards S. This instantiates state in the routers along the path to S. Eventually this join either reaches S's subnet, or reaches a router that already has (S,G) state. When this happens, data packets from S start to flow following the (S,G) state until they reach the receiver. At this point the receiver (or a router upstream of the receiver) will be receiving two copies of the data - one from the SPT and one from the RPT. When the first traffic starts to arrive from the SPT, the DR or upstream router starts to drop the packets for G from S that arrive via the RP tree. In addition, it sends an (S,G) Prune message towards the RP. This is known as an (S,G,rpt) Prune. The Prune message travels hop-by-hop, instantiating state along the path towards the RP indicating that traffic from S for G should NOT be forwarded in this direction. The prune is propagated until it reaches the RP or a router that still needs the traffic from S for other receivers.

By now, the receiver will be receiving traffic from S along the shortest-path tree between the receiver and S. In addition, the RP is receiving the traffic from S, but this traffic is no longer reaching the receiver along the RP tree. As far as the receiver is concerned, this is the final distribution tree.

Test Setup:

- TR1 and TR2 has a static configuration for the address of the RP.
- The multicast group of the RP is FF02::1:3.
- OSPF is running on all interfaces.
- If a network is not assigned to a particular area, the default area of 0.0.0.0 is used.



Procedure:

1. On network 0, TR1 and the RP transmit Hello messages.
2. On network 1, TR1 and TR2 transmit Hello messages.
3. On network 2, TR2 and the RP transmit Hello messages.
4. On network 3, TR1 transmits Hello messages.
5. On network 4, TR2 transmits Hello messages.
6. On network 3, G1 transmit multicast data for the group ff02::1:3.
7. TR1 encapsulates the multicast data and forward it to the RP on network 0.
8. On network 0, the RP transmits a Join (S, G) with an upstream neighbor of TR1, a source address of G1 and a multicast address of G1.
9. TR1 transmits multicast data packets from network 3 to network 0.
10. On network 0, the RP transmit a Register Stop with a source of G1 and a multicast group ff02::1:3.
11. On network 4, TN1 transmits a MLD report for the multicast group ff02::1:3.
12. On network 2, TR1 transmits an Join (*, G) Join with an upstream neighbor of RP and a multicast group ff02::1:3. RP forwards multicast data on network 1.
13. On network 1, TR2 transmits an Join (S, G) with an upstream neighbor of TR1 and a source address of G1 with a multicast address of ff02::1:3.
14. Observe the packets transmitted by the RUT on all networks.

Observable Results:

- TR1 should forward multicast on network 1 after receiving the join message. TR2 should transmit a Prune (S,G, rpt) to the RP. The RP should stop transmitting data packets on network 2.

Possible Problems: TR2 may never send the Join (S,G) to create the SPT.

GROUP 3: Assert Messages

Scope:

These tests are designed to verify the Assert process in PIM-SM.

Overview:

If more than one upstream router has valid forwarding state for a packet there can be duplication on a network. When this occurs the Assert process begins and a single forwarder is elected for the duplicated packets.

Test PIM.3.1: Assert Messages

Purpose: To verify that a routers properly Assert.

References: [draft-ietf-pim-sm-v2-new-08] Section 4.6

Resource Requirements:

- Monitor to capture packets
- Packet Generator

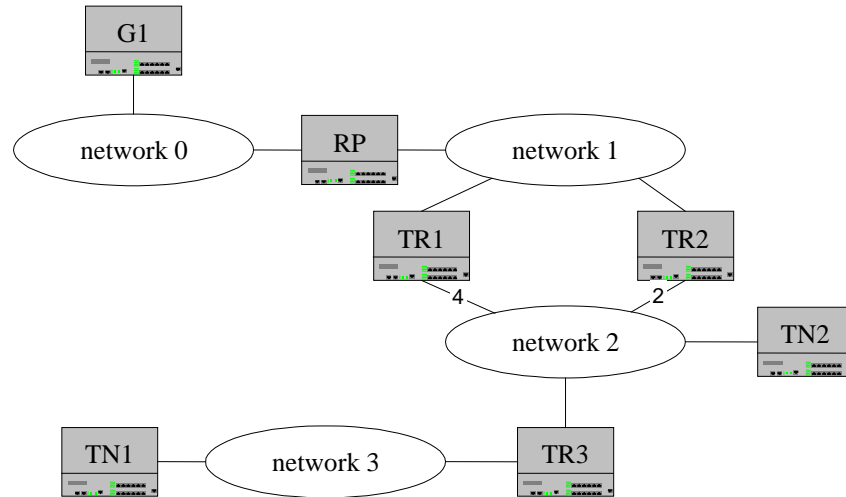
Last Modification: January 30, 2004

Discussion: Where multiple PIM routers peer over a shared LAN it is possible for more than one upstream router to have valid forwarding state for a packet, which can lead to packet duplication (see Section 3 "Multi-access LANs"). PIM does not attempt to prevent this from occurring. Instead it detects when this has happened and elects a single forwarder amongst the upstream outers to prevent further duplication. This election is performed using PIM Assert messages. Assert messages are also received by downstream routers on the LAN, and these cause subsequent Join/Prune messages to be sent to the upstream router that won the Assert.

In general, a PIM Assert message should only be accepted for processing if it comes from a known PIM neighbor. A PIM router hears about PIM neighbors through PIM Hello messages. If a router receives an Assert message from a particular IP source address and it has not seen a PIM Hello message from that source address, then the Assert message SHOULD be discarded without further processing.

Test Setup:

- TR1 and TR2 has a static configuration for the address of the RP.
- The multicast group of the RP is FF02::1:3.
- TR1 is configured with a higher DR priority then TR2 on network 2.
- OSPF is running on all interfaces.
- If a network is not assigned to a particular area, the default area of 0.0.0.0 is used.



Procedure:

1. On network 0, the RP transmits Hello messages.
2. On network 1, TR1, TR2 and the RP transmit Hello messages.
3. On network 2, TR1, TR2 and TR3 transmit Hello messages.
4. On network 3, TR3 transmits Hello messages.
5. On network 3, TN1 transmits a MLD report for the multicast group ff02::1:3.
6. On network 2, TR3 transmits a Join (*,G) with an upstream neighbor of TR2 and a group address of ff02::1:3.
7. On network 1, TR2 transmits a Join (*,G) with an upstream neighbor of RP and a group address of ff02::1:3.
8. G1 transmits multicast data with a address of ff02::1:3.
9. The RP should forward the multicast data from network 0 to network 1. TR2 should forward the data from network 1 to network 2. TR3 should forward the data onto network 3.
10. On network 2, TN2 transmits a MLD report for the multicast group ff02::1:3.
11. Observe the packets transmitted by the RUT on all networks.

Observable Results:

- TR1 should transmit a Join (*,G) with an upstream neighbor address of RP and a multicast group of ff02::1:3 on network 1. TR1 should then begin to forward data onto network 2. TR1 and TR2 should begin the Assert process. Because TR2 has a better metric it should be the Assert Winner and should continue to forward the data. TR1 should stop forwarding the data.

Possible Problems: None.