Application Note

PIM-SM Multicast Performance Testing

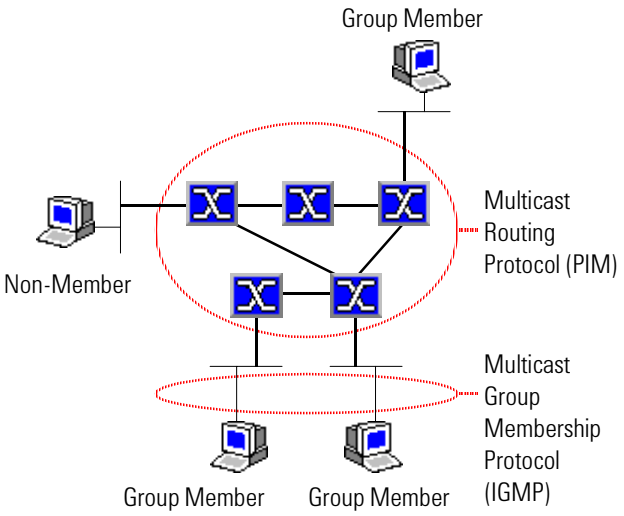

Introduction

Multicasting allows a host to send data packets across the Internet to a set of hosts that can be on different, geographically dispersed subnets. The source host sends data to a pseudo destination called a *multicast group*, and does so efficiently, using less bandwidth than unicast or broadcast traffic. Unlike unicast transmission, which would copy a packet to send it to multiple destinations, multicast sources send a packet only once.

Multicast-aware routers on the Internet use multicast *routing* protocols like PIM to deliver packets across the Internet to subnets that have hosts in the multicast group. These routers build and maintain distribution trees to forward multicast traffic.

Multicast routers connected to subnets use multicast *group membership* protocols like IGMP to discover which local hosts are members of which multicast groups, and to deliver multicasted packets to member hosts.

Current applications of multicasting include email distribution lists, routing information flooding, and web-based training seminars and voice/video conferences.

Agilent Technologies

About PIM

PIM is protocol-independent in that it can use existing unicast routing tables populated using BGP-4, OSPF, IS-IS, etc. to forward multicast traffic.

PIM-DM (Dense Mode) is ideal for groups whose members are densely distributed through a network (e.g., a corporate email group whose hosts are on the same LAN). This mode employs a push model. When a source host sends data to a multicast group, its Designated Router (DR) uses the Shortest Path First algorithm to build its own source distribution tree to each member host in the multicast group. This tree is designated (S,G), where S is the IP address of the source, and G is the IP address of the multicast group.

PIM-SM (Sparse Mode) is more efficient for groups whose members are sparsely distributed through a network (e.g., a WAN). This mode employs a pull model. When a source host sends data to a multicast group, its DR simply sends the data to a central router called a Rendezvous Point (RP). The RP maintains the only, shared distribution tree and forwards the data to each member host in the group. The onus is on each router to find the optimal path to the RP. This tree is designated (*,G), where * indicates that it is used by multiple sources.

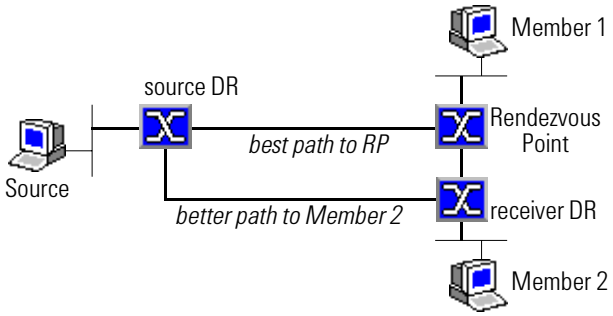

If the path through the RP is not the best path from a particular source to a host, a router can switch over to a source distribution tree using the better path.

Test Challenges

Routers supporting PIM-SM Version 2 must correctly implement these features:

As a receiver DR, it must send Join messages to the RP when a host in its subnet joins a multicast group, and Prune messages when the host leaves.

On receiving a Join/Prune message of type:

(*,G): A router must update the multicast group's shared distribution tree rooted at the RP. It must forward only the multicast traffic received on the interface that has the shortest path from the RP, using the Reverse Path Forwarding (RPF) check to avoid forwar ding loops.

(S,G): A router must update the source distribution tree rooted at the source of the multicast traffic and switch over from the RP to the source. In this case, the router uses the RPF check against the source.

As a source DR, the router must send Register messages when it receives multicast traffic from a host and does not have multicast forwarding information on the group.

As a Rendezvous Point, the router must decapsulate Register messages and forward multicast traffic.

A router must be capable of sending Assert messages to prevent the forwarding of duplicate multicast messages.

A router must function independent of the unicast routing protocol used.

Other tests:

Group Join/Prune latency: The time it takes a router to update its distribution tree after receiving a Join/Prune message.

Latency in switching from a shared to a source distribution tree.

Whether the performance of unicast traffic suffers while multicast traffic is being propagated.

Scaling to find the maximum number of multicast groups a SUT supports before packet loss or excess latency occurs.

Test Descriptions

(*,G) distribution

This section describes how to send a Join message to a System Under Test (SUT) to see how fast it can update an upstream shared distribution tree and start forwarding multicast traffic back to a multicast group member:

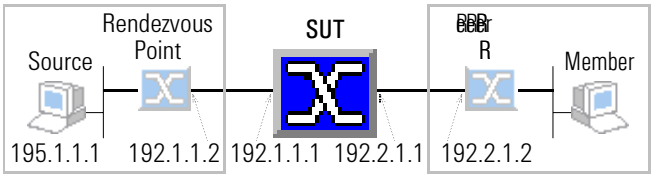

Test Port 1A
Simulates a host sending traffic to a multicast group through its RP.

Test Port 1B
Simulates a receiver DR that sends a Join message to the SUT so that it can receive multicast data.

Test steps

1. Enable PIM-SM emulation on test ports 1A and 1B.
2. Simulate a multicast group and enable test port 1B to become a member.
3. Set up test port 1A to send traffic to the multicast group.
4. Send a Join message from port 1B and check the latency before receiving multicast packets (then check the Prune latency).

(S,G) distribution

This section describes how to send a S,G Join message to a System Under Test (SUT) to see how fast it can create a source-specific distribution tree and start forwarding multicast traffic from the source to a multicast group member:

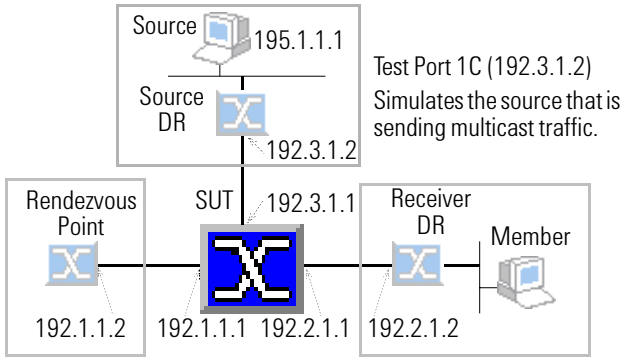

Test Port 1C (192.3.1.2)
Simulates the source that is sending multicast traffic.

Test Port 1A
Simulates a RP router.

Test Port 1B
Simulates a receiver DR that sends a Join message to the SUT so that it can receive multicast data.

Test steps

1. Enable PIM-SM emulation on test ports 1A and 1B.
2. Simulate a multicast group and enable test port 1B to become a member.
3. Set up test port 1A to be the rendezvous point (RP) router.
4. Set up test port 1C to send multicast traffic from a specific source (195.1.1.1).
5. Send multicast traffic from test port 1C to the multicast group being simulated on port 1B.
6. Send a (*,G) Join message from port 1B and verify that no traffic is being received.
7. Send a (S,G) Join message from port 1B and verify that traffic is being received.

NOTE: This note does not provide detailed instructions for the (S,G) distribution test steps. The following pages show instructions for the (*,G) distribution test scenario.

Preamble steps

This note does not illustrate these test preamble steps:

Select test ports 1A and 1B.

Configure the IP addresses of the test ports and connected SUT interfaces.

Bring up the physical and link layers.

SUT setup

Configure the SUT as follows:

Enable PIM-SM Version 2 on the SUT interfaces.

Set the RP to test port 1A's interface address (192.1.1.2).

For the (S,G) distribution scenario, add a route for the multicast source address (195.1.1.1), with test port 1C (192.3.1.2) as the forwarding router.

References

RFC 2362: PIM-SM

draft-ietf-pim-sm-v2-new-nn.txt: PIM-SM Version 2

draft-ietf-mboned-anycast-rp-nn.txt: Anycast RP using PIM and MSDP

Step 1: Enable PIM-SM emulation on test ports 1A and 1B

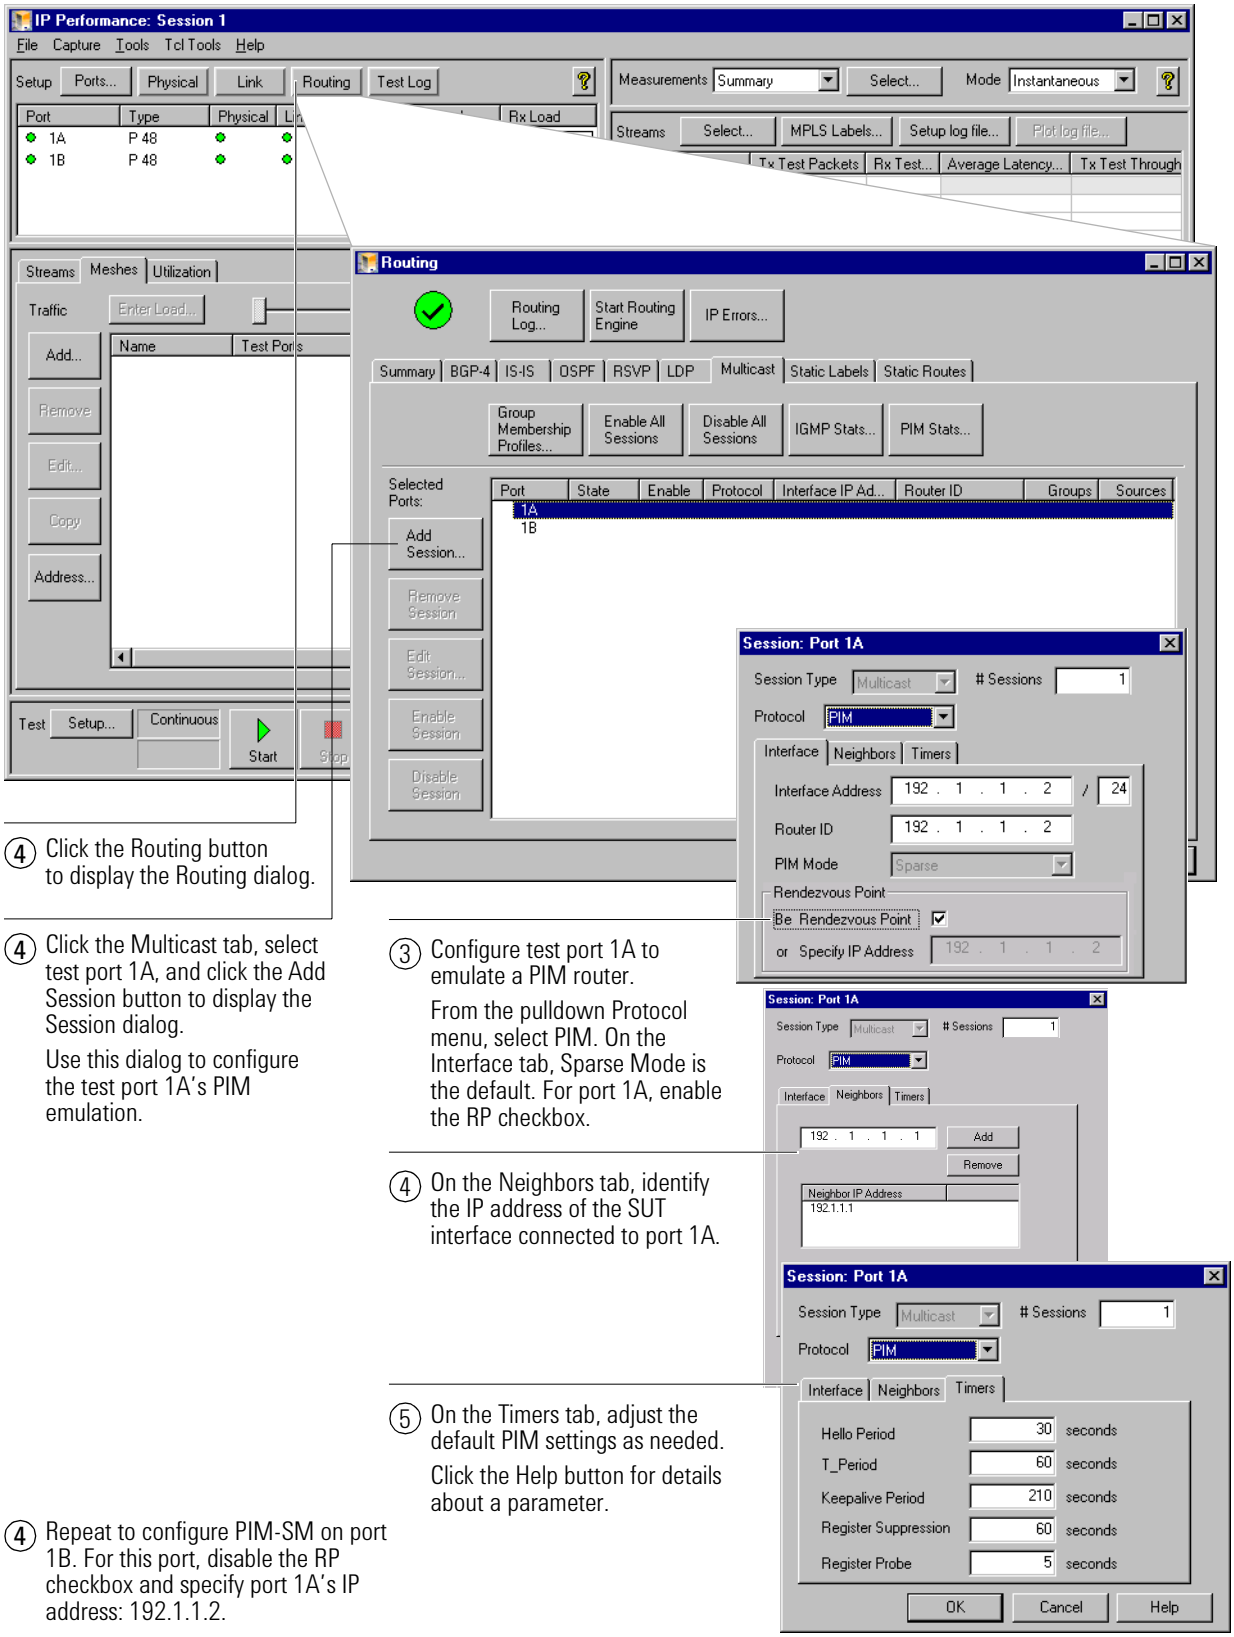

4 Click the Routing button to display the Routing dialog.

4 Click the Multicast tab, select test port 1A, and click the Add Session button to display the Session dialog.

Use this dialog to configure the test port 1A's PIM emulation.

3 Configure test port 1A to emulate a PIM router.

From the pulldown Protocol menu, select PIM. On the Interface tab, Sparse Mode is the default. For port 1A, enable the RP checkbox.

4 On the Neighbors tab, identify the IP address of the SUT interface connected to port 1A.

5 On the Timers tab, adjust the default PIM settings as needed.

Click the Help button for details about a parameter.

4 Repeat to configure PIM-SM on port 1B. For this port, disable the RP checkbox and specify port 1A's IP address: 192.1.1.2.

Step 2: Simulate a multicast group and enable test port 1B to become a member later

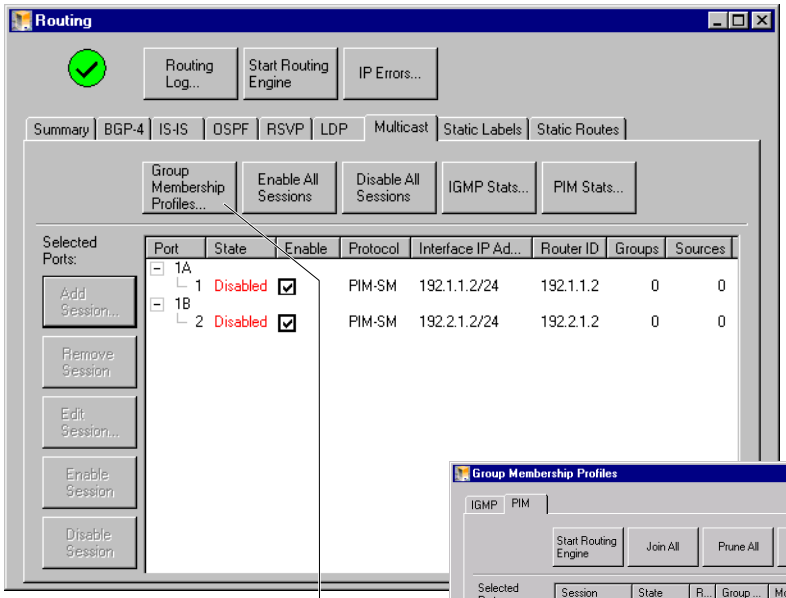

1 Back on the Routing dialog, click the Group Membership Profiles button.

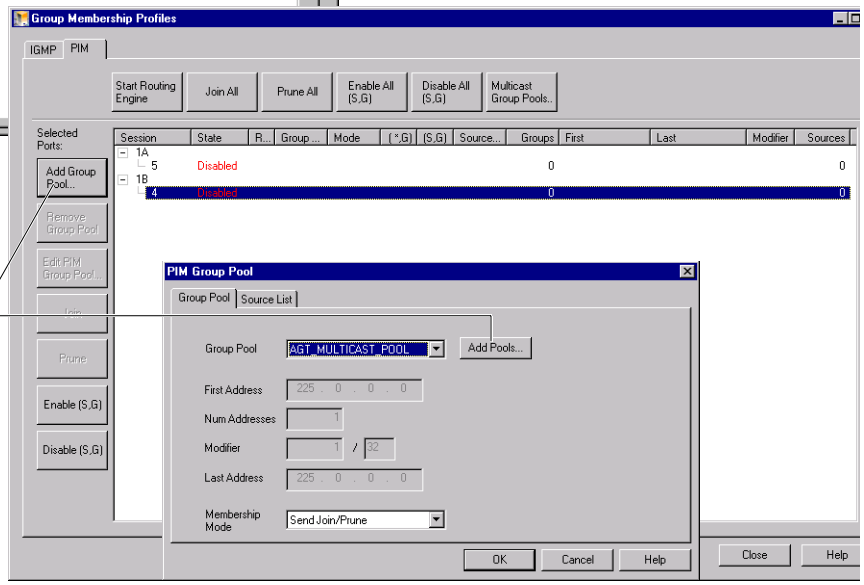

2 On the Group Membership Profiles dialog, select port 1B and click Add Group Pool.

On the PIM Group Pool dialog, click Add Pools.

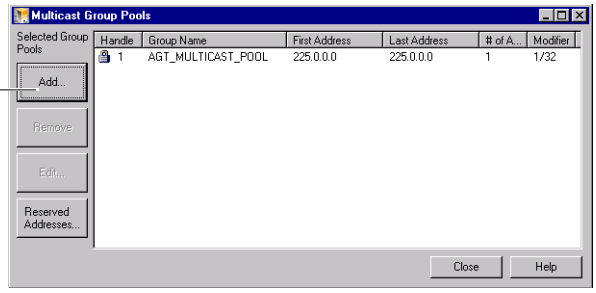

3 On the Multicast Group Pools dialog, click the Add button to define a new multicast group.

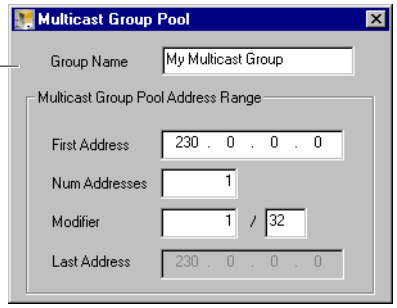

4 On the Multicast Group Pool dialog, define the multicast group address.

You can define a "pool" of several addresses to scale the test and see how the SUT handles up to 200,000 different multicast groups.

5 Back on the PIM Group Pool dialog, select the newly defined multicast group from the pulldown menu.

Back on the Group Membership Profiles dialog, under port 1B, this multicast group is shown with a checkbox so that you can dynamically join and leave the group.

Step 3: Set up test port 1A to send traffic to the multicast group

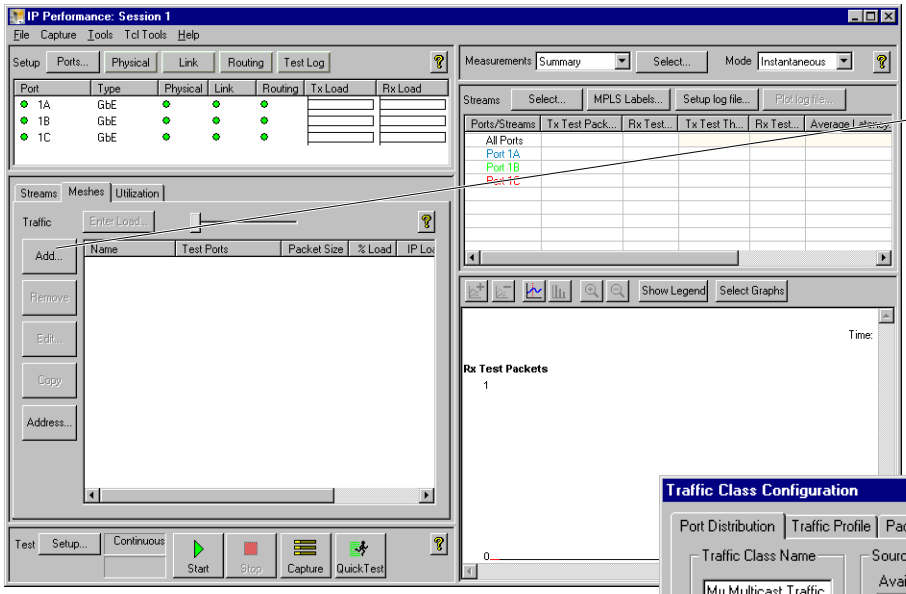

1. In the Traffic area of the IP Performance window, select the Meshes tab, then click the Add button.

2. On the Traffic Class Configuration dialog, select the port from which to send multicast traffic.

   Click on a port in the Available Ports list and click the Add button to send traffic from the selected test port.

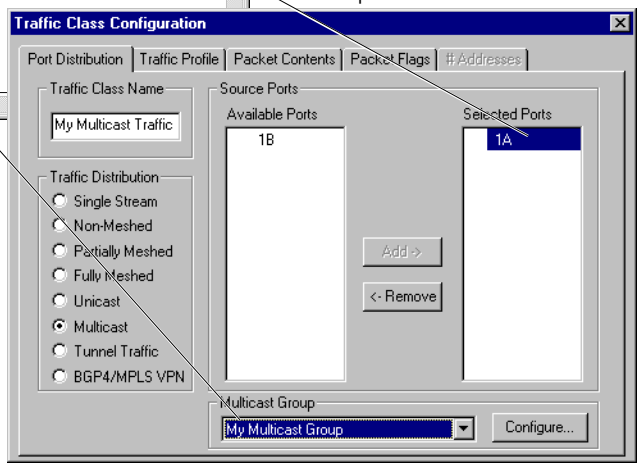

3. On the Traffic Class Configuration dialog, select the multicast group to receive traffic.

   The default AGT_MULTICAST_GROUP cannot be used to receive traffic. Select a multicast group you configured on the previous page from the pull-down menu.

4. Back on the IP Performance window's Traffic area, click the Address button.

5. Specify the source address to use in multicast packets.

   Select the multicast traffic stream, select the destination "route" at the bottom, then click the Edit Source button.

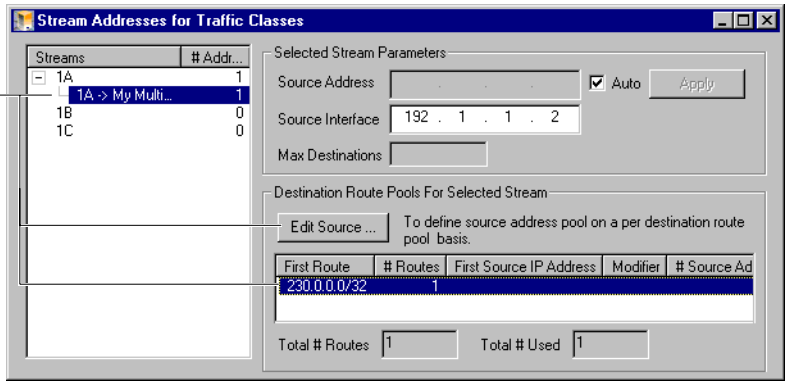

6. Add a new source pool.

   Click Add Pools. On the Multicast Source Address Pools dialog, click Add. Then, on the Multicast Source Address Pool dialog, enter the source address(es) to use.

7. Back on the IP Performance window's Test area, click the Start button to start generating the multicast traffic and measuring statistics.

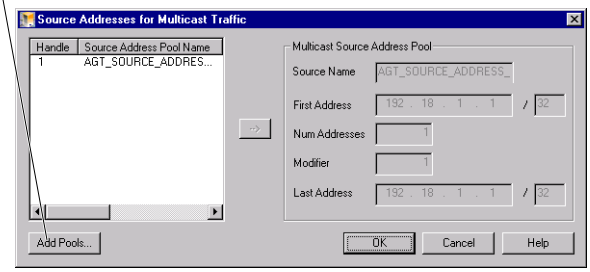

Agilent's RouterTester System

Agilent's RouterTester System offers a powerful and versatile test platform to address the evolving test needs of metro/edge platforms, core routers and optical switches. RouterTester provides Network Equipment Manufacturers and Service Providers with the industry's leading tools for wire speed, multiport traffic generation and performance analysis of today's networking devices.

Warranty and Support

Hardware Warranty

All RouterTester and QA Robot hardware is warranted against defects in materials and workmanship for a period of 3 years from the date of shipment.

Software Warranty

All RouterTester and QA Robot software is warranted for a period of 90 days. The applications are warranted to execute and install properly from the media provided. This warranty only covers physical defects in the media, whereby the media is replaced at no charge during the warranty period.

Software Updates

With the purchase of any new system controller Agilent will provide 1 year of complimentary software updates. At the end of the first year you can enroll into the Software Enhancement Service (SES) for continuing software product enhancements.

Support

Technical support is available throughout the support life of the product. Support is available to verify that the equipment works properly, to help with product operation, and to provide basic measurement assistance for the use of the specified capabilities, at no extra cost, upon request.

Ordering Information

To order and configure the test system consult your local Agilent field engineer.

United States:

Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

Canada:

Agilent Technologies Canada Inc.
5150 Spectrum Way
Mississauga, Ontario
L4W 5G1
1-877-894-4414

Europe:

Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-2323

United Kingdom
07004 666666

Japan:

Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:

Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 269-7500
Fax: (305) 267-4286

Asia Pacific:

Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 3197-7777
Fax: (852) 2506-9233

Australia/New Zealand:

Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881

www.agilent.com/comms/RouterTester

Specifications subject to change
April 15, 2002
5988-6560EN

Agilent Technologies