Network Working Group
INTERNET-DRAFT
Expires in: August 2001

Debra Stopp
Hardev Soor
IXIA

Methodology for IP Multicast Benchmarking

<draft-ietf-bmwg-mcastm-07.txt>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet- Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet- Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

The purpose of this draft is to describe methodology specific to the benchmarking of multicast IP forwarding devices. It builds upon the tenets set forth in RFC 2544, RFC 2432 and other IETF Benchmarking Methodology Working Group (BMWG) efforts. This document seeks to extend these efforts to the multicast paradigm.

The BMWG produces two major classes of documents: Benchmarking Terminology documents and Benchmarking Methodology documents. The Terminology documents present the benchmarks and other related terms. The Methodology documents define the procedures required to collect the benchmarks cited in the

corresponding Terminology documents.

Table of Contents

1. Introduction

This document defines a specific set of tests that vendors can use to measure and report the performance characteristics and forwarding capabilities of network devices that support IP multicast protocols. The results of these tests will provide the user comparable data from different vendors with which to evaluate these devices.

A previous document, " Terminology for IP Multicast Benchmarking" (RFC 2432), defined many of the terms that are used in this document. The terminology document should be consulted before attempting to make use of this document.

This methodology will focus on one source to many destinations, although many of the tests described may be extended to use multiple source to multiple destination IP multicast communication.

2. Key Words to Reflect Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

3. Test set up

Figure 1 shows a typical setup for an IP multicast test, with one source to multiple destinations, although this MAY be extended to multiple source to multiple destinations.

```
                                              +----------------+
                       +------------+         |    Egress      |
     +--------+        |         (-)-------->| destination(E1)|
     |        |        |            |         |                |
     | source |------->(|)Ingress    |         +----------------+
     |        |        |            |         +----------------+
     +--------+        |   D U T   (-)-------->|    Egress      |
                       |            |         | destination(E2)|
                       |            |         |                |
                       |            |         +----------------+
                       |            |              . . .
                       |            |         +----------------+
                       |            |         |    Egress      |
                       |         (-)-------->| destination(En)|
                       |            |         |                |
                       +------------+         +----------------+
```

Figure 1

If the multicast metrics are to be taken across multiple devices forming a System Under Test (SUT), then test packets are offered to a single ingress interface on a device of the SUT, subsequently routed across the SUT topology, and finally forwarded to the test apparatus' packet-receiving components by the test egress interface(s) of devices in the SUT. Figure 2 offers an example SUT test topology. If a SUT is tested, the details of the test topology MUST be disclosed with the corresponding test results.

```
+--------+                       +----------------+    +--------+
|        |     +------------+    |DUT B Egress E0(-)-->|        |
|        |     |DUT A       |--->|                |    |        |
| Test   |     |            |    |      Egress E1(-)-->| Test   |
| App.   |--->(-)Ingress, I |    +----------------+    | App.   |
| Traffic|     |            |    +----------------+    | Traffic|
| Src.   |     |            |--->|DUT C Egress E2(-)-->| Dest.  |
|        |     +------------+    |                |    |        |
|        |                       |      Egress En(-)-->|        |
+--------+                       +----------------+    +--------+
```

Figure 2

Generally , the destination ports first join the desired number of multicast groups by sending IGMP Join Group messages to the DUT/SUT. To verify that all destination ports successfully joined the appropriate groups, the source port MUST transmit IP multicast frames destined for these groups. The destination ports MAY send IGMP Leave Group messages after the transmission of IP Multicast frames to clear the IGMP table of the DUT/SUT.

In addition, all transmitted frames MUST contain a recognizable pattern that can be filtered on in order to ensure the receipt of only the frames that are involved in the test.

3.1. Test Considerations

The procedures outlined below are written without regard for specific physical layer or link layer protocols. The methodology further assumes a uniform medium topology. Issues regarding mixed transmission media, such as speed mismatch, headers differences, etc., are not specifically addressed. Moreover, no provisions are made for traffic-affecting factors, such as congestion control or service differentiation mechanisms. Modifications to the specified collection procedures might need to be made to accommodate the transmission media actually tested. These accommodations MUST be presented with the test results.

3.1.1. IGMP Support

Each of the destination ports should support and be able to test all IGMP versions 1, 2 and 3. The minimum requirement, however, is IGMP version 2.

Each destination port should be able to respond to IGMP queries during the test.

Each destination port should also send LEAVE (running IGMP version 2) after each test.

3.1.2. Group Addresses

The Class D Group address SHOULD be changed between tests. Many DUTs have memory or cache that is not cleared properly and can bias the results.

The following group addresses are recommended by use in a test:

```
224.0.1.27-224.0.1.255
224.0.5.128-224.0.5.255
224.0.6.128-224.0.6.255
```

If the number of group addresses accommodated by these ranges do not satisfy the requirements of the test, then these ranges may be overlapped. The total number of configured group addresses must be less than or equal to the IGMP table size of the DUT/SUT.

3.1.3. Frame Sizes

Each test SHOULD be run with different Multicast Frame Sizes. The recommended frame sizes are 64, 128, 256, 512, 1024, 1280, and 1518 byte frames.

3.1.4. TTL

The source frames should have a TTL value large enough to accommodate the DUT/SUT.

3.2. Layer 2 Support

Each of the destination ports should support GARP/GMRP protocols to join groups on Layer 2 DUTs/SUTs.

4. Forwarding and Throughput

This section contains the description of the tests that are related to the characterization of the packet forwarding of a DUT/SUT in a multicast environment. Some metrics extend the concept of throughput presented in RFC 1242. The notion of Forwarding Rate is cited in RFC 2285.

4.1. Mixed Class Throughput

Objective

To determine the maximum throughput rate at which none of the offered frames, comprised from a unicast Class and a multicast Class, to be forwarded are dropped by the device across a fixed number of ports as defined in RFC 2432.

Procedure

Multicast and unicast traffic are mixed together in the same aggregated traffic stream in order to simulate the non-homogenous networking environment. While the multicast traffic is transmitted from one source to multiple destinations, the unicast traffic MAY be evenly distributed across the DUT/SUT architecture. In addition, the DUT/SUT MUST learn the appropriate unicast IP addresses, either by sending ARP frames from each unicast address, sending a RIP packet or by assigning static entries into the DUT/SUT address table.

The mixture of multicast and unicast traffic MUST be set up in one of two ways:

a) As a percentage of the total traffic flow employing maximum bandwidth utilization. Thus, each type of traffic is transmitted at the maximum available bandwidth. This also implies that the intended load, regardless of the type of traffic, remains constant.

b) As a percentage of the total traffic flow employing a proportionate bandwidth utilization. Thus, each type of traffic is transmitted at a fraction of the available bandwidth proportional to the specified ratio. This also implies that the intended load for each traffic type varies in proportion to its specified ratio.

The transmission of the frames MUST be set up so that they form a deterministic distribution while still maintaining the specified forwarding rates. See Appendix A for a discussion on non-homogenous vs. homogenous packet distribution.

Similar to the Frame loss rate test in RFC 2544, the first trial SHOULD be run for the frame rate that corresponds to 100% of the maximum rate for the frame size on the input media. Repeat the procedure for the rate that corresponds to 90% of the maximum rate used and then for 80% of this rate. This sequence SHOULD be continued (at reducing 10% intervals) until there are two successive trials in which no frames are lost. The maximum granularity of the trials MUST be 10% of the maximum rate, a finer granularity is encouraged.

Result

Parameters to be measured SHOULD include the frame loss and percent loss for each class of traffic per destination port. The ratio of unicast traffic to multicast traffic MUST be reported.

The nature of the traffic stream contributing to the result MUST be reported. All required reporting parameters of mixed class throughput MUST be reflected in the results report, such as the transmitted packet size(s) and offered load of the packet stream.

4.2. Scaled Group Forwarding Matrix

Objective

A table that demonstrates Forwarding Rate as a function of tested multicast groups for a fixed number of tested DUT/SUT ports.

Procedure

Multicast traffic is sent at a fixed percent of maximum offered load with a fixed number of receive ports of the tester at a fixed frame length.

The receive ports SHOULD continue joining incrementally by 10 multicast groups until a user defined maximum is reached.

The receive ports will continue joining in the incremental fashion until a user defined maximum is reached.

Results

Parameters to be measured SHOULD include the frame loss and percent loss per destination port for each multicast group address.

The nature of the traffic stream contributing to the result MUST be reported. All required reporting parameters MUST be reflected in the results report, such as the transmitted packet size(s) and offered load of the packet stream.

4.3. Aggregated Multicast Throughput

Objective

The maximum rate at which none of the offered frames to be forwarded through N destination interfaces of the same multicast group are dropped.

Procedure

Multicast traffic is sent at a fixed percent of maximum offered load with a fixed number of groups at a fixed frame length for a fixed duration of time.

The initial number of receive ports of the tester will join the group(s) and the sender will transmit to the same groups after a certain delay (a few seconds).

Then the an incremental number of receive ports will join the same groups and then the Multicast traffic is sent as stated.

The receive ports will continue to be added and multicast traffic sent until a user defined maximum number of ports is reached.

Results

Parameters to be measured SHOULD include the frame loss and percent loss per destination port for each multicast group address.

The nature of the traffic stream contributing to the result MUST be reported. All required reporting parameters of aggregated throughput MUST be reflected in the results report, such as the transmitted packet size(s) and offered load of the packet stream.

4.4. Encapsulation/Decapsulation (Tunneling) Throughput

This sub-section provides the description of tests that help in obtaining throughput measurements when a DUT/SUT or a set of DUTs are acting as tunnel endpoints. The following Figure 3 presents the a tunneled network.

```
Client A      DUT/SUT A      Network      DUT/SUT B      Client B

             ----------                   ----------
             |        |      ------       |        |
-----(a)  (b)|        |(c)  (      )   (d)|        |(e) (f)-----
||||| -----> |        |---->(      )----->|        |-----> |||||
-----        |        |      ------       |        |       -----
             |        |                   |        |
             ----------                   ----------

                           Figure 3
                           --------
```

A tunnel is created between DUT/SUT A (the encapsulator) and DUT/SUT B (the decapsulator). Client A is acting as a source and

Client B is the destination. Client B joins a multicast group (for example, 224.0.1.1) by sending an IGMP Join message to DUT/SUT B to join that group. Client A now wants to transmit some traffic to Client B. It will send the multicast traffic to DUT/SUT A which encapsulates the multicast frames, sends it to DUT/SUT B which will decapsulate the same frames and forward them to Client B.

4.4.1. Encapsulation Throughput

Objective

The maximum rate at which frames offered a DUT/SUT are encapsulated and correctly forwarded by the DUT/SUT without loss.

Procedure

Traffic is sent through a DUT/SUT that has been configured to encapsulate the frames. Traffic is received on a test port prior to decapsulation and throughput is calculated based on RFC2544.

Results

Parameters to be measured SHOULD include the measured throughput per tunnel.

The nature of the traffic stream contributing to the result MUST be reported. All required reporting parameters of encapsulation throughput MUST be reflected in the results report, such as the transmitted packet size(s) and offered load of the packet stream.

4.4.2. Decapsulation Throughput

Objective

The maximum rate at which frames offered a DUT/SUT are decapsulated and correctly forwarded by the DUT/SUT without loss.

Procedure

Encapsulated traffic is sent through a DUT/SUT that has been configured to decapsulate the frames. Traffic is received on a test port after decapsulation and throughput is calculated based on RFC2544.

Results

Parameters to be measured SHOULD include the measured throughput per tunnel.

The nature of the traffic stream contributing to the result MUST be reported. All required reporting parameters of decapsulation throughput MUST be reflected in the results report, such as the transmitted packet size(s) and offered load of the packet stream.

4.4.3. Re-encapsulation Throughput

Objective

The maximum rate at which frames of one encapsulated format offered a DUT/SUT are converted to another encapsulated format and correctly forwarded by the DUT/SUT without loss.

Procedure

Traffic is sent through a DUT/SUT that has been configured to encapsulate frames into one format, then re-encapsulate the frames into another format. Traffic is received on a test port after all decapsulation is complete and throughput is calculated based on RFC2544.

Results

Parameters to be measured SHOULD include the measured throughput per tunnel.

The nature of the traffic stream contributing to the result MUST be reported. All required reporting parameters of re-encapsulation throughput MUST be reflected in the results report, such as the transmitted packet size(s) and offered load of the packet stream.

5. Forwarding Latency

This section presents methodologies relating to the characterization of the forwarding latency of a DUT/SUT in a multicast environment. It extends the concept of latency characterization presented in RFC 2544.

In order to lessen the effect of packet buffering in the DUT/SUT, the latency tests MUST be run such that the offered load is less than the multicast throughput of the DUT/SUT as determined in the previous section. The tests should also take into account the DUT's/SUT's need to cache the traffic in its IP cache, fastpath cache or shortcut tables since the initial part of the traffic will be utilized to build these tables.

Lastly, RFC 1242 and RFC 2544 draws distinction between two classes of devices: "store and forward" and "bit-forwarding." Each class impacts how latency is collected and subsequently presented. See the related RFCs for more information. In practice, much of the test equipment will collect the latency measurement for one class or the other, and, if needed, mathematically derive the reported value by the addition or subtraction of values accounting for medium propagation delay of the packet, bit times to the timestamp trigger within the packet, etc. Test equipment vendors SHOULD provide documentation regarding the composition and calculation latency values being reported. The user of this data SHOULD understand the nature of the latency values being reported, especially when comparing results collected from multiple test vendors. (E.g., If test vendor A presents a "store and forward" latency result and test vendor B presents a "bit-forwarding" latency result, the user may erroneously conclude the DUT has two differing sets of latency values.)

5.1. Multicast Latency

Objective

To produce a set of multicast latency measurements from a single, multicast ingress port of a DUT or SUT through multiple, egress multicast ports of that same DUT or SUT as provided for by the metric "Multicast Latency" in RFC 2432.

The procedures highlighted below attempt to draw from the collection methodology for latency in RFC 2544 to the degree possible. The methodology addresses two topological scenarios: one for a single device (DUT) characterization; a second scenario is presented or multiple device (SUT) characterization.

Procedure

If the test trial is to characterize latency across a single Device Under Test (DUT), an example test topology might take the form of Figure 1 in section 3. That is, a single DUT with one ingress interface receiving the multicast test traffic from packet-transmitting component of the test apparatus and n egress interfaces on the same DUT forwarding the multicast test traffic back to the packet-receiving component of the test apparatus. Note that n reflects the number of TESTED egress interfaces on the DUT actually expected to forward the test traffic (as opposed to configured but untested, non-forwarding interfaces, for example).

If the multicast latencies are to be taken across multiple devices forming a System Under Test (SUT), an example test topology might take the form of Figure 2 in section 3.

The trial duration SHOULD be 120 seconds. Departures to the suggested traffic class guidelines MUST be disclosed with the respective trial results. The nature of the latency measurement, "store and forward" or "bit forwarding," MUST be associated with the related test trial(s) and disclosed in the results report.

End-to-end reachability of the test traffic path SHOULD be verified prior to the engagement of a test trial. This implies that subsequent measurements are intended to characterize the latency across the tested device's or devices' normal traffic forwarding path (e.g., faster hardware-based engines) of the device(s) as opposed a non-standard traffic processing path (e.g. slower, software-based exception handlers). If the test trial is to be executed with the intent of characterizing a non-optimal, forwarding condition, then a description of the exception processing conditions being characterized MUST be included with the trial's results.

A test traffic stream is presented to the DUT. At the mid-point of the trial's duration, the test apparatus MUST inject a uniquely identifiable ("tagged") packet into the test traffic packets being presented. This tagged packet will be the basis for the latency measurements. By "uniquely identifiable," it is meant that the test apparatus MUST be able to discern the "tagged" packet from the other packets comprising the test traffic set. A packet generation timestamp, Timestamp A, reflecting the completion of the transmission of the tagged packet by the test apparatus, MUST be determined.

The test apparatus then monitors packets from the DUT's tested egress port(s) for the expected tagged packet(s) until the cessation of traffic generation at the end of the configured trial duration.A value of the Offered Load presented the DUT/SUT MUST be noted.

The test apparatus MUST record the time of the successful detection of a tagged packet from a tested egress interface with a timestamp, Timestamp B. A set of Timestamp B values MUST be collected for all tested egress interfaces of the DUT/SUT.

A trial MUST be considered INVALID should any of the following conditions occur in the collection of the trial data:

- Forwarded test packets directed to improper destinations.
- Unexpected differences between Intended Load and Offered Load or unexpected differences between Offered Load and the resulting Forwarding Rate(s) on the DUT/SUT egress ports.
- Forwarded test packets improperly formed or packet header fields improperly manipulated.
- Failure to forward required tagged packet(s) on all expected egress interfaces.
- Reception of a tagged packet by the test apparatus outside the configured test duration interval or 5 seconds, whichever is greater.

Data from invalid trials SHOULD be considered inconclusive. Data from invalid trials MUST not form the basis of comparison.

The set of latency measurements, M, composed from each latency measurement taken from every ingress/tested egress interface pairing MUST be determined from a valid test trial:

```
      M = { (Timestamp B(E0) - Timestamp A),
            (Timestamp B(E1) - Timestamp A), ...
            (Timestamp B(En) - Timestamp A) }
```

where (E0 ... En) represents the range of all tested egress interfaces and Timestamp B represents a tagged packet detection event for a given DUT/SUT tested egress interface.

Results

Two types of information MUST be reported: 1) the set of latency measurements and 2) the significant environmental, methodological, or device particulars giving insight into the test or its results.

Specifically, when reporting the results of a VALID test trial, the set of ALL latencies related to the tested ingress interface and each tested egress DUT/SUT interface of MUST be presented. The time units of the presented latency MUST be uniform and with sufficient precision for the medium or media being tested. Results MAY be offered in tabular format and SHOULD preserve the relationship of latency to ingress/egress interface to assist in trending across multiple trials.

The Offered Load of the test traffic presented the DUT/SUT, size of the "tagged" packet, trial duration, and nature (i.e., store-and-forward or bit-forwarding) of the trial's measurement MUST be associated with any reported test trial's result.

5.2. Min/Max Multicast Latency

Objective

The difference between the maximum latency measurement and the minimum latency measurement from a collected set of latencies produced by the Multicast Latency benchmark.

Procedure

Collect a set of multicast latency measurements, as prescribed in section 5.1. This will produce a set of multicase latencies, M, where M is composed of individual forwarding altencies between DUT packet ingress and DUT packet egress port pairs. E.g.:

```
    M = {L(I,E1),L(I,E2), à, L(I,En)}
```

where L is the latency between a tested ingress port, I, of the DUT, and Ex a specific, tested multicast egress port of the DUT. E1 through En are unique egress ports on the DUT.

From the collected multicast latency measurements in set M, identify MAX(M), where MAX is a function that yields the largest latency value from set M.

Identify MIN(M), when MIN is a funtion that yields the smallest latency value from set M.

The Max/Min value is determined from the following formula:

    Result = MAX(M) û MIN(M)

Results

The result MUST be represented as a single numerical value in time units consistent with the corresponding latency measurements. In addition the number of tested egress ports on the DUT MUST be reported.

The nature of the traffic stream contributing to the result MUST be reported. All required reporting parameters of multicast latency MUST be reflected in the min/max results report, such as the transmitted packet size(s) and offered load of the packet stream in which the tagged packet was presented to the DUT.

6. Overhead

This section presents methodology relating to the characterization of the overhead delays associated with explicit operations found in multicast environments.

6.1. Group Join Delay

Objective

The time duration it takes a DUT/SUT to start forwarding multicast packets from the time a successful IGMP group membership report has been issued to the DUT/SUT.

Procedure

Traffic is sent on the source port at the same time as the IGMP JOIN Group message is transmitted from the destination ports. The join delay is the difference in time from when the IGMP Join is sent (timestamp A) and the first frame is forwarded to a receiving member port (timestamp B).

Group Join delay = timestamp B - timestamp A

One of the keys is to transmit at the fastest rate the DUT/SUT can handle multicast frames. This is to get the best resolution and the least margin of error in the Join Delay.

However, you do not want to transmit the frames so fast that frames are dropped by the DUT/SUT. Traffic should be sent at the throughput rate determined by the forwarding tests of section 4.

Results

The parameter to be measured is the join delay time for each multicast group address per destination port. In addition, the number of frames transmitted and received and percent loss may be reported.

6.2. Group Leave Delay

Objective

The time duration it takes a DUT/SUT to cease forwarding multicast packets after a corresponding IGMP "Leave Group" message has been successfully offered to the DUT/SUT.

Procedure

Traffic is sent on the source port at the same time as the IGMP Leave Group messages are transmitted from the destination ports. The leave delay is the difference in time from when the IGMP leave is sent (timestamp A) and the last frame is forwarded to a receiving member port (timestamp B).

Group Leave delay = timestamp B - timestamp A

One of the keys is to transmit at the fastest rate the DUT/SUT can handle multicast frames. This is to get the best resolution and least margin of error in the Leave Delay. However, you do not want to transmit the frames too fast that frames are dropped by the DUT/SUT. Traffic should be sent at the throughput rate determined by the forwarding tests of section 4.

Results

The parameter to be measured is the leave delay time for each multicast group address per destination port. In addition, the number of frames transmitted and received and percent loss may be reported.

7. Capacity

This section offers terms relating to the identification of multicast group limits of a DUT/SUT.

7.1. Multicast Group Capacity

Objective

The maximum number of multicast groups a DUT/SUT can support while maintaining the ability to forward multicast frames to all multicast groups registered to that DUT/SUT.

Procedure

One or more destination ports of DUT/SUT will join an initial number of groups.

Then after a delay (enough time for all ports to join) the source port will transmit to each group at a transmission rate that the DUT/SUT can handle without dropping IP Multicast frames.

If all frames sent are forwarded by the DUT/SUT and received the test iteration is said to pass at the current capacity.

If the iteration passes at the capacity the test will add an user defined incremental value of groups to each receive port.

The iteration is to run again at the new group level and capacity tested as stated above.

Once the test fails at a capacity the capacity is stated to be the last Iteration that pass at a giving capacity.

Results

The parameter to be measured is the total number of group addresses that were successfully forwarded with no loss.

In addition, the nature of the traffic stream contributing to the result MUST be reported. All required reporting parameters MUST be reflected in the results report, such as the transmitted packet size(s) and offered load of the packet stream.

8. Interaction

Network forwarding devices are generally required to provide more functionality than just the forwarding of traffic. Moreover, network forwarding devices may be asked to provide those functions in a variety of environments. This section offers terms to assist in the characterization of DUT/SUT behavior in consideration of potentially interacting factors.

8.1. Forwarding Burdened Multicast Latency

The Multicast Latency metrics can be influenced by forcing the DUT/SUT to perform extra processing of packets while multicast traffic is being forwarded for latency measurements. In this test, a set of ports on the tester will be designated to be source and destination similar to the generic IP Multicast test setup. In addition to this setup, another set of ports will be selected to transmit some multicast traffic that is destined to multicast group addresses that have not been joined by these additional set of ports.

For example, if ports 1,2, 3, and 4 form the burdened response setup (setup A) which is used to obtain the latency metrics and ports 5, 6, 7, and 8 form the non-burdened response setup (setup B) which will afflict the burdened response setup, then setup B traffic will join multicast group addresses not joined by the ports in this setup. By sending such multicast traffic, the DUT/SUT will perform a lookup on the packets that will affect the processing of setup A traffic.

8.2. Forwarding Burdened Group Join Delay

The port configuration in this test is similar to the one described in section 8.1, but in this test, the multicast traffic is not sent by the ports in setup B. In this test, the setup A traffic must be influenced in such a way that will affect the DUT's/SUT's ability to process Group Join messages. Therefore, in this test, the ports in setup B will send a set of IGMP Group Join messages while the ports in setup A are also joining its own set of group addresses. Since the two sets of group addresses are independent of each other, the group join delay for setup A may be different than in the case when there were no other group addresses being joined.

9. Security Considerations

As this document is solely for the purpose of providing metric methodology and describes neither a protocol nor a protocol's implementation, there are no security considerations associated with this document.

10. Acknowledgements

The authors would like to acknowledge the following individuals for

their help and participation of the compilation and editing of this document û Ralph Daniels, Netcom Systems, who made significant

contributions to earlier versions of this draft and Kevin Dubray, Juniper Networks.

11. References

[Br91] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, July 1991.

[Br96] Bradner, S., and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.

[Br97] Bradner, S. "Use of Keywords in RFCs to Reflect Requirement Levels, RFC 2119, March 1997

[Du98] Dubray, K., "Terminology for IP Multicast Benchmarking", RFC 2432, October 1998.

[Hu95] Huitema, C. "Routing in the Internet." Prentice-Hall, 1995.

[Ka98] Kosiur, D., "IP Multicasting: the Complete Guide to Interactive Corporate Networks", John Wiley & Sons, Inc, 1998.

[Ma98] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.

[Mt98] Maufer, T. "Deploying IP Multicast in the Enterprise." Prentice-Hall, 1998.

[Se98] Semeria, C. and Maufer, T. "Introduction to IP Multicast Routing." http://www.3com.com/nsc/501303.html 3Com Corp., 1998.

12. Author's Addresses

Debra Stopp
IXIA
26601 W. Agoura Rd.
Calabasas, CA 91302
USA

Phone: 818 871 1800
EMail: debby@ixiacom.com

Hardev Soor
IXIA
26601 W. Agoura Rd.
Calabasas, CA 91302
USA

Phone: 818 871 1800
EMail: hardev@ixiacom.com

13. Full Copyright Statement

"Copyright (C) The Internet Society (date). All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into.

Appendix A: Determining an even distribution

It is important to understand and fully define the distribution of frames among all multicast and unicast destinations. If the distribution is not well defined or understood, the throughput and forwarding metrics are not meaningful.

In a homogeneous environment, a large single burst of multicast frames may be followed by a large burst of unicast frames. This is a very different distribution than that of a non-homogeneous environment, where the multicast and unicast frames are intermingled throughout the entire transmission.

The recommended distribution is that of the non-homogeneous environment because it more closely represents a real-world scenario. The distribution is modeled by calculating the number of multicast frames per destination port as a burst, then calculating the number of unicast frames to transmit as a percentage of the total frames transmitted. The overall effect of the distribution is small bursts of multicast frames intermingled with small bursts of unicast frames.