

Network Working Group  
INTERNET-DRAFT  
Expires in: August 2003

Debra Stopp  
Ixia  
Brooks Hickman  
Spirent Communications  
February 2003

**Methodology for IP Multicast Benchmarking**  
<[draft-ietf-bmwg-mcastm-11.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 (2003). All Rights Reserved.

Abstract

The purpose of this document 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.....](#) [3](#)
- [2. KEY WORDS TO REFLECT REQUIREMENTS.....](#) [3](#)
- [3. TEST SET UP.....](#) [3](#)
  - [3.1. Test Considerations.....](#) [5](#)
    - [3.1.1. IGMP Support.....](#) [5](#)
    - [3.1.2. Group Addresses.....](#) [5](#)
    - [3.1.3. Frame Sizes.....](#) [6](#)
    - [3.1.4. TTL.....](#) [6](#)
    - [3.1.5. Trial Duration.....](#) [6](#)
- [4. FORWARDING AND THROUGHPUT.....](#) [6](#)
  - [4.1. Mixed Class Throughput.....](#) [7](#)
  - [4.2. Scaled Group Forwarding Matrix.....](#) [8](#)
  - [4.3. Aggregated Multicast Throughput.....](#) [9](#)
  - [4.4. Encapsulation/Decapsulation \(Tunneling\) Throughput.....](#) [10](#)
    - [4.4.1. Encapsulation Throughput.....](#) [10](#)
    - [4.4.2. Decapsulation Throughput.....](#) [12](#)
    - [4.4.3. Re-encapsulation Throughput.....](#) [13](#)
- [5. FORWARDING LATENCY.....](#) [15](#)
  - [5.1. Multicast Latency.....](#) [16](#)
  - [5.2. Min/Max Multicast Latency.....](#) [18](#)
- [6. OVERHEAD.....](#) [20](#)
  - [6.1. Group Join Delay.....](#) [20](#)
  - [6.2. Group Leave Delay.....](#) [21](#)
- [7. CAPACITY.....](#) [23](#)
  - [7.1. Multicast Group Capacity.....](#) [23](#)
- [8. INTERACTION.....](#) [24](#)
  - [8.1. Forwarding Burdened Multicast Latency.....](#) [24](#)
  - [8.2. Forwarding Burdened Group Join Delay.....](#) [25](#)
- [9. SECURITY CONSIDERATIONS.....](#) [26](#)
- [10. ACKNOWLEDGEMENTS.....](#) [27](#)
- [11. CONTRIBUTIONS.....](#) [27](#)
- [12. REFERENCES.....](#) [28](#)
- [13. AUTHOR'S ADDRESSES.....](#) [29](#)
- [14. FULL COPYRIGHT STATEMENT.....](#) [29](#)

## **1. Introduction**

This document defines tests for measuring and reporting the forwarding, latency and IGMP group membership characteristics of devices that support IP multicast routing protocols. The results of these tests will provide the user with meaningful data on multicast performance.

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 topologies.

## **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](#). [RFC 2119](#) defines the use of these key words to help make the intent of standards track documents as clear as possible. While this document uses these keywords, this document is not a standards track document.

## **3. Test set up**

The set of methodologies presented in this document are for single ingress, multiple egress scenarios as exemplified by Figures 1 and 2. Methodologies for multiple ingress and multiple egress scenarios are beyond the scope of this document.

Figure 1 shows a typical setup for an IP multicast test, with one source to multiple destinations.



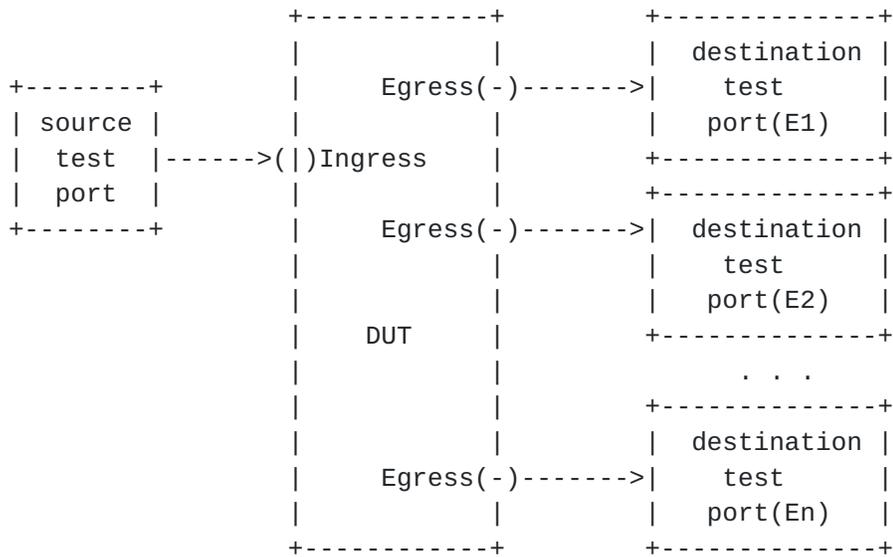


Figure 1

-----

If the multicast metrics are to be taken across multiple devices forming a System Under Test (SUT), then test frames are offered to a single ingress interface on a device of the SUT, subsequently forwarded across the SUT topology, and finally forwarded to the test apparatus' frame-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 test topology and all relevant configuration details MUST be disclosed with the corresponding test results.

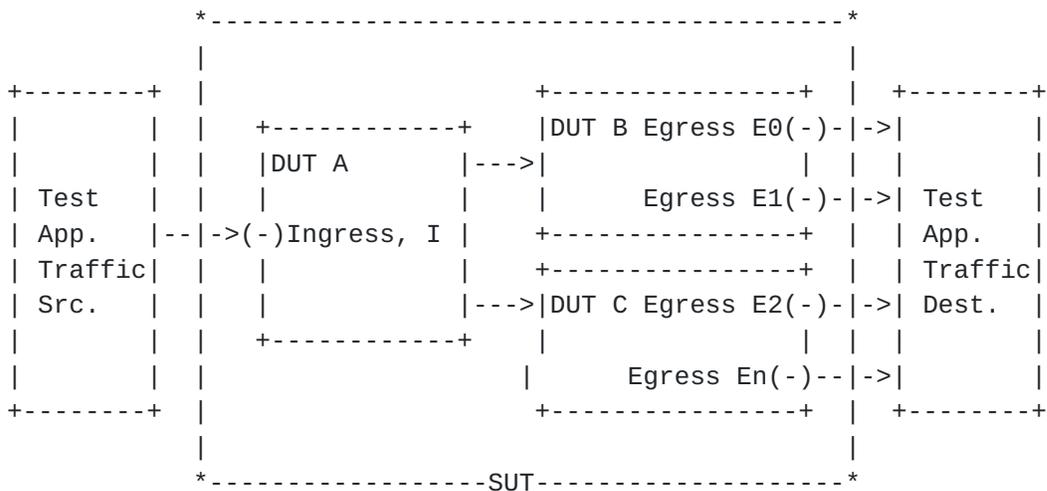


Figure 2

-----

Generally, the destination test ports first join the desired number of multicast groups by sending IGMP Group Report messages to the DUT/SUT. To verify that all destination test ports successfully joined the appropriate groups, the source port MUST transmit IP

multicast frames destined for these groups. The destination test 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, test equipment MUST validate the correct and proper forwarding actions of the devices they test in order to ensure the receipt of the frames that are involved in the test.

### **3.1. Test Considerations**

The methodology assumes a uniform medium topology. Issues regarding mixed transmission media, such as speed mismatch, headers differences, etc., are not specifically addressed. Flow control, QoS and other non-essential traffic or traffic-affecting mechanisms affecting the variable under test MUST be disabled. Modifications to the collection procedures might need to be made to accommodate the transmission media actually tested. These accommodations MUST be presented with the test results.

An actual flow of test traffic may be required to prime related mechanisms, (e.g., process RPF events, build device caches, etc.) to optimally forward subsequent traffic. Therefore, before an initial, measured forwarding test trial, the test apparatus MUST generate test traffic utilizing the same addressing characteristics to the DUT/SUT that will subsequently be used to measure the DUT/SUT response. The test monitor should ensure the correct forwarding of traffic by the DUT/SUT. The priming action need only be repeated to keep the associated information current.

#### **3.1.1. IGMP Support**

All of the ingress and egress interfaces MAY support any version of IGMP. The IGMP version on the ingress interface MUST be the same version of IGMP that is being tested on the egress interfaces.

Each of the ingress and egress interfaces SHOULD be able to respond to IGMP queries during the test.

Each of the ingress and egress interfaces SHOULD also send LEAVE (running IGMP version 2 or later) after each test.

#### **3.1.2. Group Addresses**

It is intended that the collection of benchmarks prescribed in this document be executed in an isolated lab environment. That is to say, the test traffic offered the tested devices MUST NOT traverse

a live internet, intranet, or other production network.

Assuming the above, there is no restriction to the use of multicast addresses to compose the test traffic other than those assignments imposed by IANA. The IANA assignments MUST be regarded for operational consistency. For multicast address assignments see:

<http://www.iana.org/assignments/multicast-addresses>

Address selection does not need to be restricted to Administratively Scoped IP Multicast addresses.

### **3.1.3. Frame Sizes**

Each test SHOULD be run with different multicast frame sizes. For Ethernet, the recommended sizes are 64, 128, 256, 512, 1024, 1280, and 1518 byte frames.

Other link layer technologies MAY be used. The minimum and maximum frame lengths of the link layer technology in use SHOULD be tested.

When testing with different frame sizes, the DUT/SUT configuration MUST remain the same.

### **3.1.4. TTL**

The data plane test traffic should have a TTL value large enough to traverse the DUT/SUT.

The TTL in IGMP control plane messages is in compliance with the version of IGMP in use.

### **3.1.5. Trial Duration**

The duration of the test portion of each trial SHOULD be at least 30 seconds. This parameter MUST be included as part of the results reporting for each methodology.

## **4. Forwarding and Throughput**

This section contains the description of the tests that are related to the characterization of the frame forwarding of a DUT/SUT in a multicast environment. Some metrics extend the concept of throughput presented in [RFC 1242](#). Forwarding Rate is cited in [RFC 2285](#).



#### **4.1. Mixed Class Throughput**

Objective:

To determine the throughput of a DUT/SUT when both unicast class frames and multicast class frames are offered simultaneously to a fixed number of interfaces 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.

The following events MUST occur before offering test traffic:

- o All DUT/SUT egress interfaces configured to receive multicast traffic MUST join all configured multicast groups;
- o The DUT/SUT MUST learn the appropriate unicast addresses; and
- o Group membership and unicast address learning MUST be verified through some externally observable method.

The intended load [[Ma98](#)] SHOULD be configured as alternating multicast frames and unicast frames to a single ingress interface in a 50-50 ratio. The unicast frames MUST be configured to transmit in a round-robin fashion to all of the egress interfaces. The multicast frames MUST be configured to transmit to all of the egress interfaces.

Mixed class throughput measurement is defined in [RFC2432](#) [[Du98](#)]. A search algorithm MUST be utilized to determine the throughput for both unicast class and multicast class traffic in a mixed class environment.

Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:

- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Total number of multicast groups
- o Traffic distribution for unicast and multicast traffic

classes

o The ratio of multicast and unicast traffic must be declared

Stopp & Hickman

[Page 7]

The following results MUST be reflected in the results specific to this methodology:

- o Mixed Class Throughput as defined in [RFC2432](#) [Du98], including: Throughput per unicast and multicast traffic classes.

The Mixed Class Throughput results for each test SHOULD be reported in the form of a table with a row for each of the tested frame sizes per the recommendations in [section 3.1.3](#). Each row SHOULD specify the intended load, number of multicast frames offered, number of unicast frames offered and measured throughput per class.

#### **[4.2](#). Scaled Group Forwarding Matrix**

Objective:

To determine Forwarding Rate as a function of tested multicast groups for a fixed number of tested DUT/SUT ports.

Procedure:

This is an iterative procedure. The destination test port(s) MUST join an initial number of multicast groups on the first iteration. All DUT/SUT destination test port(s) configured to receive multicast traffic MUST join all configured multicast groups. The recommended number of groups to join on the first iteration is 10 groups. Multicast traffic is subsequently transmitted to all groups joined on this iteration.

The number of multicast groups joined by each destination test port is then incremented, or scaled, by an additional number of multicast groups. The recommended granularity of additional groups to join per iteration is 10, although the tester MAY choose a finer granularity. Multicast traffic is subsequently transmitted to all groups joined during this iteration.

The total number of multicast groups joined MUST not exceed the capacity of the DUT/SUT. Both Group Join Delay and Group Capacity results MUST be known prior to running this test.

Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:



- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version

The following results MUST be reflected in the results specific to this methodology:

- o The total number of multicast groups joined for that iteration
- o Total number of frames transmitted
- o Total number of frames received
- o Offered load
- o Forwarding rate determined for that iteration

The Scaled Group Forwarding results for each test SHOULD be reported in the form of a table with a row representing each iteration of the test. Each row or iteration SHOULD specify the total number of groups joined for that iteration, total number of frames transmitted, total number of frames received and the aggregate forwarding rate determined for that iteration.

#### **4.3. Aggregated Multicast Throughput**

Objective:

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

Procedure:

Offer multicast traffic at an initial fixed offered load to a fixed set of interfaces with a fixed number of groups at a fixed frame length for a fixed duration of time. All destination test ports MUST join all specified multicast groups.

If any frame loss is detected, the offered load is decreased and the sender will transmit again. An iterative search algorithm MUST be utilized to determine the maximum offered frame rate with a zero frame loss.

Each iteration will involve varying the offered load of the multicast traffic, while keeping the set of interfaces, number of multicast groups, frame length and test duration fixed, until the maximum rate at which none of the offered frames are dropped is determined.

Parameters to be measured MUST include the offered load at which no frame loss occurred.

Stopp & Hickman

[Page 9]

Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:

- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Total number of multicast groups

The following results MUST be reflected in the results specific to this methodology:

- o Aggregated Multicast Throughput as defined in [RFC2432](#) [[Du98](#)]

The Aggregated Multicast Throughput results SHOULD be reported in the format of a table with a row for each of the tested frame sizes per the recommendations in [section 3.1.3](#). Each row or iteration SHOULD specify offered load, total number of offered frames and the measured Aggregated Multicast Throughput.

#### **[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.

##### **[4.4.1. Encapsulation Throughput](#)**

Objective:

To determine the maximum rate at which frames offered to one ingress interface of a DUT/SUT are encapsulated and correctly forwarded on one or more egress interfaces of the DUT/SUT without loss.





- o Test duration
- o IGMP version
- o Total number of multicast groups
- o MTU size of DUT/SUT interfaces

Stopp & Hickman

[Page 11]

The following results MUST be reflected in the results specific to this methodology:

- o Measured Encapsulated Throughput as defined in [RFC2432](#) [Du98]
- o Encapsulated frame size
- o Originating un-encapsulated frame size
- o Number of tunnels
- o Number of multicast groups per tunnel

The Encapsulated Throughput results SHOULD be reported in the form of a table and specific to this test there SHOULD be rows for each originating un-encapsulated frame size. Each row or iteration SHOULD specify the offered load, encapsulation method, encapsulated frame size, total number of offered frames, and the encapsulation throughput.

**4.4.2. Decapsulation Throughput**

Objective:

To determine the maximum rate at which frames offered to one ingress interface of a DUT/SUT are decapsulated and correctly forwarded by the DUT/SUT on one or more egress interfaces without loss.

Procedure:

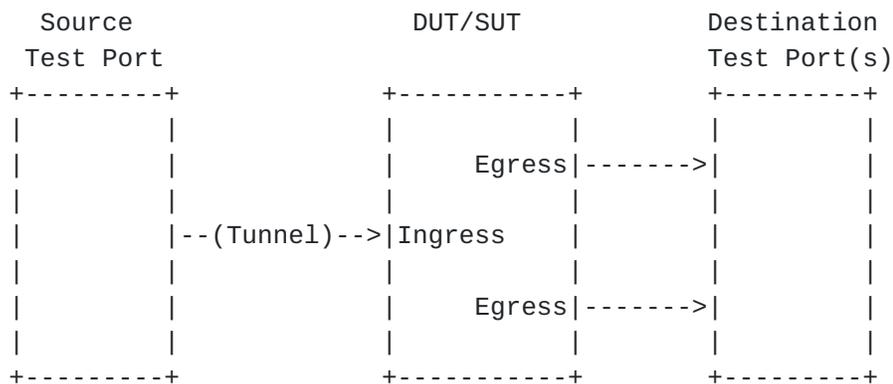


Figure 4  
-----

Figure 4 shows the setup for testing the decapsulation throughput of the DUT/SUT. One or more tunnels are created between the source test port and the DUT/SUT. Encapsulated multicast traffic will

then be offered by the source test port, decapsulated by the DUT/SUT and forwarded to the destination test port(s).

The DUT/SUT SHOULD be configured such that the traffic across each egress interface will consist of either:

- a) A single tunnel encapsulating one or more multicast address groups OR
- b) Multiple tunnels, each encapsulating one or more multicast address groups.

The number of multicast groups per tunnel MUST be the same when the DUT/SUT is configured in a multiple tunnel configuration. All destination test ports MUST join all multicast group addresses offered by the source test port. Each egress interface MUST be configured with the same MTU.

A search algorithm MUST be utilized to determine the decapsulation throughput as defined in [\[Du98\]](#).

Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:

- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Total number of multicast groups
- o MTU size of DUT/SUT interfaces

The following results MUST be reflected in the results specific to this methodology:

- o Measured Decapsulated Throughput as defined in [RFC2432](#) [\[Du98\]](#)
- o Originating encapsulation format
- o Decapsulated frame size
- o Originating encapsulated frame size
- o Number of tunnels
- o Number of multicast groups per tunnel

The Decapsulated Throughput results SHOULD be reported in the format of a table and specific to this test there SHOULD be rows for each originating encapsulated frame size. Each row or iteration SHOULD specify the offered load, decapsulated frame size, total number of offered frames and the decapsulation throughput.

#### **[4.4.3. Re-encapsulation Throughput](#)**

Objective:

To determine the maximum rate at which frames of one encapsulated format offered to one ingress interface of a DUT/SUT are converted

Stopp & Hickman

[Page 13]

to another encapsulated format and correctly forwarded by the DUT/SUT to one or more egress interfaces without loss.

Procedure:

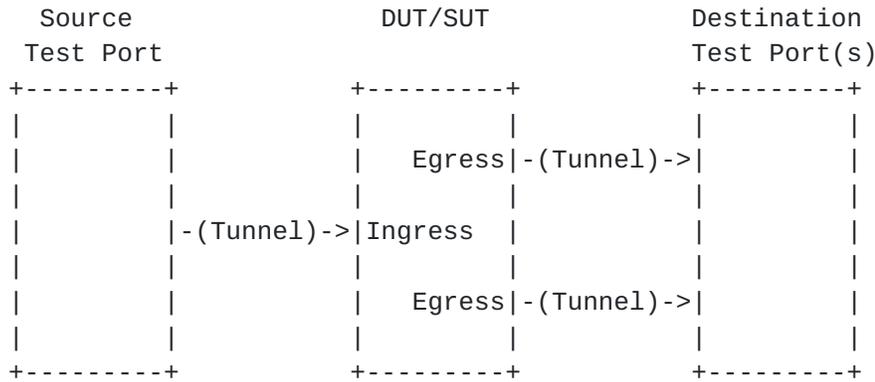


Figure 5

Figure 5 shows the setup for testing the Re-encapsulation throughput of the DUT/SUT. The source test port will offer encapsulated traffic of one type to the DUT/SUT, which has been configured to re-encapsulate the offered frames using a different encapsulation format. The DUT/SUT will then forward the re-encapsulated frames to the destination test port(s).

The DUT/SUT SHOULD be configured such that the traffic across each egress interface will consist of either:

- a) A single tunnel encapsulating one or more multicast address groups OR
- b) Multiple tunnels, each encapsulating one or more multicast address groups.

The number of multicast groups per tunnel MUST be the same when the DUT/SUT is configured in a multiple tunnel configuration.

In addition, the DUT/SUT SHOULD be configured such that the number of tunnels on the ingress and each egress interface are the same. All destination test ports MUST join all multicast group addresses offered by the source test port. Each egress interface MUST be configured with the same MTU.

A search algorithm MUST be utilized to determine the re-encapsulation throughput as defined in [Du98].

Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:

Stopp & Hickman

[Page 14]

- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Total number of multicast groups
- o MTU size of DUT/SUT interfaces

The following results MUST be reflected in the results specific to this methodology:

- o Measured Re-encapsulated Throughput as defined in [RFC2432](#) [[Du98](#)]
- o Originating encapsulation format
- o Decapsulated frame size
- o Originating encapsulated frame size
- o Number of tunnels
- o Number of multicast groups per tunnel

The Decapsulated Throughput results SHOULD be reported in the format of a table and specific to this test there SHOULD be rows for each originating encapsulated frame size. Each row or iteration SHOULD specify the offered load, decapsulated frame size, total number of offered frames and the decapsulation throughput

## 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](#).

To lessen the effect of frame buffering in the DUT/SUT, the latency tests MUST be run at the measured multicast throughput level of the DUT; multicast latency at other offered loads is optional.

Lastly, [RFC 1242](#) and [RFC 2544](#) draw a distinction between device types: "store and forward" and "bit-forwarding." Each type 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 type 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 frame, bit times to the timestamp trigger within the frame, etc.



## **5.1. Multicast Latency**

Objective:

To produce a set of multicast latency measurements from a single, multicast ingress interface of a DUT/SUT through multiple, egress multicast interfaces of that same DUT/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 for 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 frame-transmitting component of the test apparatus and n egress interfaces on the same DUT forwarding the multicast test traffic back to the frame-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 to be consistent with [RFC 2544](#). 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 MUST 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. It is RECOMMENDED to offer traffic at the measured aggregated multicast throughput rate

([Section 4.3](#)). At the mid-point of the trial's duration, the test apparatus MUST inject a uniquely identifiable ("tagged") frame into the test traffic frames being presented. This tagged frame 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" frame from the other frames comprising the test traffic set. A frame generation timestamp, Timestamp A, reflecting the completion of the transmission of the tagged frame by the test apparatus, MUST be determined.

The test apparatus then monitors frames from the DUT's tested egress interface(s) for the expected tagged frame(s) until the cessation of traffic generation at the end of the configured trial duration.

The test apparatus MUST record the time of the successful detection of a tagged frame 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. See [RFC 1242 \[Br91\]](#) for additional discussion regarding store and forward devices and bit forwarding devices.

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

- o Forwarded test frames directed to improper destinations.
- o 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.
- o Forwarded test frames improperly formed or frame header fields improperly manipulated.
- o Failure to forward required tagged frame(s) on all expected egress interfaces.
- o Reception of a tagged frame 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:

$$M = \{ (\text{Timestamp B}(E_0) - \text{Timestamp A}), \\ (\text{Timestamp B}(E_1) - \text{Timestamp A}), \dots \\ (\text{Timestamp B}(E_n) - \text{Timestamp A}) \}$$



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

A more continuous profile MAY be built from a series of individual measurements.

#### Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:

- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Offered load
- o Total number of multicast groups

The following results MUST be reflected in the results specific to this methodology:

- o The time units of the presented latency MUST be uniform and with sufficient precision for the medium or media being tested.
- o Specifically, when reporting the results of a valid test trial, the set of all latencies related to the tested ingress and each tested egress DUT/SUT interface of MUST be presented.

The latency results for each test SHOULD be reported in the form of a table, with a row for each of the tested frame sizes per the recommended frame sizes in [section 3.1.3](#), and SHOULD preserve the relationship of latency to ingress/egress interface(s) to assist in trending across multiple trials.

## **5.2. Min/Max Multicast Latency**

### Objective:

To determine 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 over a single test duration, as prescribed in [section 5.1](#). This will produce a set of

multicast latencies,  $M$ , where  $M$  is composed of individual forwarding latencies between DUT frame ingress and DUT frame egress port pairs. E.g.:

$$M = \{L(I,E1), L(I,E2), \dots, L(I,En)\}$$

where  $L$  is the latency between a tested ingress interface,  $I$ , of the DUT, and  $E_x$  a specific, tested multicast egress interface of the DUT.  $E_1$  through  $E_n$  are unique egress interfaces 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 function that yields the smallest latency value from set  $M$ .

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

$$\text{Result} = \text{MAX}(M) - \text{MIN}(M)$$

A more continuous profile MAY be built from a series of individual measurements.

#### Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:

- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Offered load
- o Total number of multicast groups

The following results MUST be reflected in the results specific to this methodology:

- o The result of the min/max value represented as a single numerical value in time units consistent with the corresponding latency measurements.
- o Specifically, when reporting the results of a valid test trial, the set of all latencies related to the tested ingress interface MUST be reported.

The time units of the presented latency MUST be uniform and with sufficient precision for the medium or media being tested. The

latency results for each test SHOULD be reported in the form of a table, with a row for each of the tested frame sizes per the recommendations in [section 3.1.3](#), and SHOULD preserve the

relationship of latency to ingress/egress interface(s) to assist in trending across multiple trials.

## **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:

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

Procedure:

Prior to sending any IGMP Group Membership Reports used to calculate the group join delay, it MUST be verified through externally observable means that the destination test ports are not currently a member of any of the specified multicast groups. If any of the egress interfaces forward multicast frames, the test is not valid.

Once verification is complete, multicast traffic for all relevant multicast group addresses MUST be offered to the ingress interface prior to receipt or processing of any IGMP Group Membership Report messages. It is RECOMMENDED to offer traffic at the measured aggregated multicast throughput rate ([Section 4.3](#)).

After the multicast traffic has been started, each destination test port (See Figure 1) SHOULD send one IGMP Group Membership Report with one or more (IGMP V3) multicast group(s) specified. All destination test ports MUST join all multicast groups offered on the ingress interface of the DUT/SUT. The test MUST be performed with one multicast group and SHOULD be performed with multiple groups.

The join delay is the difference in time from when the IGMP Group Membership message is sent (timestamp A) and the first frame of the multicast group is forwarded to a receiving egress interface (timestamp B).

Group Join delay time = timestamp B - timestamp A

Timestamp A MUST be the time the last bit of the IGMP group membership report is sent from the destination test port; timestamp

Stopp & Hickman

[Page 20]

B MUST be the time the first bit of the first valid multicast frame is forwarded on the egress interface of the DUT/SUT.

Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:

- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Total number of multicast groups

The following results MUST be reflected in the results specific to this methodology:

- o The group join delay time per multicast group address
- o The group join delay time per egress interface(s)

The Group Join Delay results for each test SHOULD be reported in the form of a table, with a row for each of the tested frame sizes per the recommendations in [section 3.1.3](#). Each row or iteration SHOULD specify the group join delay time for each multicast group per destination interface, number of frames transmitted and number of frames received for that iteration.

## **[6.2. Group Leave Delay](#)**

Objective:

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

Procedure:

Prior to sending any IGMP Group Leave Group messages used to calculate the group leave delay, it MUST be verified through externally observable means that the destination test ports are currently a member of all the specified multicast groups. If any of the destination test ports do not receive multicast frames, the test is not valid.

Once verification is complete, multicast traffic for all relevant multicast group addresses MUST be offered to the ingress interface

prior to receipt or processing of any IGMP Leave Group messages.  
It is RECOMMENDED to offer traffic at the measured aggregated  
multicast throughput rate ([Section 4.3](#)).

After the multicast traffic has been started, each destination test port (See Figure 1) MUST send one IGMP Leave Group message for each multicast group specified. All destination test ports MUST leave all relevant multicast groups offered on the ingress interface of the DUT/SUT. The test MUST be performed with one multicast group and SHOULD be performed with multiple groups.

The leave delay is the difference in time from when the IGMP Leave Group message is sent (timestamp A) and the last frame of the multicast group is forwarded to a receiving egress interface (timestamp B).

$$\text{Group Leave delay time} = \text{timestamp B} - \text{timestamp A}$$

Timestamp A MUST be the time the last bit of the IGMP Leave Group message is sent from the destination test port; timestamp B MUST be the time the last bit of the last valid multicast frame is forwarded on the egress interface of the DUT/SUT.

#### Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:

- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Total number of multicast groups

The following results MUST be reflected in the results specific to this methodology:

- o The group leave delay time per multicast group address
- o The group leave delay time per egress interface(s)

The Group Leave Delay results for each test SHOULD be reported in the form of a table, with a row for each of the tested frame sizes per the recommendations in [section 3.1.3](#). Each row or iteration SHOULD specify the group leave delay time for each multicast group per destination interface, number of frames transmitted and number of frames received for that iteration.



## **7. Capacity**

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

### **7.1. Multicast Group Capacity**

Objective:

To determine 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 egress interfaces of DUT/SUT will join an initial number of multicast groups.

After a delay as determined by [section 6.1](#), the ingress interface MUST transmit to each group at a specified offered load.

If at least one frame for each multicast group is forwarded properly by the DUT/SUT to each participating egress interface, the iteration is said to pass at the current capacity.

If the iteration passes, the test will add a user defined incremental value of groups to each egress interface. At the new group level and resultant capacity as stated above, run the iteration again.

Once the test fails, the last/previous iteration capacity that passed is the stated Maximum Group Capacity result.

Reporting Format:

The following configuration parameters MUST be reflected in the results specific to this methodology:

- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Offered load

The following results MUST be reflected in the results specific to this methodology:

- o The total number of multicast group addresses that were successfully forwarded through the DUT/SUT

The Multicast Group Capacity results for each test SHOULD be reported in the form of a table, with a row for each of the tested frame sizes per the recommendations in [section 3.1.3](#). Each row or iteration SHOULD specify the number of multicast groups joined per destination interface, number of frames transmitted and number of frames received for that iteration.

## **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 procedures to assist in the characterization of DUT/SUT behavior in consideration of potentially interacting factors.

### **[8.1. Forwarding Burdened Multicast Latency](#)**

Objective:

To produce a set of multicast latency measurements from a single multicast ingress interface of a DUT/SUT through multiple egress multicast interfaces of that same DUT/SUT as provided for by the metric "Multicast Latency" in [RFC 2432](#) while under the influence of a traffic forwarding requirement.

Procedure:

The Multicast Latency metrics can be influenced by forcing the DUT/SUT to perform extra processing of packets while multicast class traffic is being forwarded for latency measurements.

The Burdened Forwarding Latency test MUST follow the described setup in [Section 5.1](#).

Perform a baseline measurement of latency as described in [Section 5.1](#). After the baseline measurement is obtained, the test is repeated with the ingress interface offering an additional set of user specified multicast group addresses which have not been joined by the destination test port(s). The offered load MUST be the same as was used in the baseline measurement.

By sending such multicast class traffic, the DUT/SUT may perform a lookup on the frames that may affect the processing of traffic destined for the egress interface(s).



Reporting Format:

Similar to [Section 5.1](#), the following configuration parameters MUST be reflected in the results specific to this methodology:

- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Total number of multicast groups in the baseline setup
- o Total number of additional multicast groups used to burden the setup

The following results MUST be reflected in the results specific to this methodology:

- o The time units of the presented latency MUST be uniform and with sufficient precision for the medium or media being tested
- o Specifically, when reporting the results of a valid test trial, the set of all latencies related to the tested ingress and each tested egress DUT/SUT interface of MUST be presented
- o Reported results from baseline measurement; [section 5.1](#)

The latency results for each test SHOULD be reported in the form of a table, with a row for each of the tested frame sizes per the recommended frame sizes in [section 3.1.3](#), and SHOULD preserve the relationship of latency to ingress/egress interface(s) to assist in trending across multiple trials.

## **[8.2](#). Forwarding Burdened Group Join Delay**

Objective:

To determine the time duration it takes a DUT/SUT to start forwarding multicast frames from the time a successful IGMP Group Membership Report has been issued to the DUT/SUT while under the influence of a traffic forwarding requirement.

Procedure:

The Group Join Delay metrics can be influenced by forcing the DUT/SUT to perform extra process of packets while attempting to update and maintain the IP multicast address forwarding table.

The Forwarding Burdened Group Join Delay test MUST follow the

described setup in [Section 6.1](#).

Perform a baseline measurement of group join delay as described in [Section 6.1](#). After the baseline measurement is obtained, the test is repeated with the ingress interface offering an additional set of user specified multicast group address which have not been joined by the destination test port(s). The offered load MUST be the same as was used in the baseline measurement.

By sending such multicast class traffic, the DUT/SUT may perform a lookup on the frames that may affect the processing of the IGMP Group Report messages.

Reporting Format:

Similar to [Section 6.1](#), the following configuration parameters MUST be reflected in the results specific to this methodology:

- o Frame size(s)
- o Number of tested egress interfaces on the DUT/SUT
- o Test duration
- o IGMP version
- o Total number of multicast groups in the baseline setup
- o Total number of additional multicast groups used to burden the setup

The following results MUST be reflected in the results specific to this methodology:

- o The group join delay time per multicast group address
- o The group join delay time per egress interface(s)
- o Reported results from baseline measurement; [section 6.1](#)

The Group Join Delay results for each test SHOULD be reported in the form of a table, with a row for each of the tested frame sizes per the recommendations in [section 3.1.3](#). Each row or iteration SHOULD specify the group join delay time for each multicast group per destination interface, number of frames transmitted and number of frames received for that iteration.

## **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 Benchmarking Methodology Working Group of the IETF and particularly Kevin Dubray, Juniper Networks, are to be thanked for the many suggestions they collectively made to help complete this document.

## **11. Contributions**

The authors would like to acknowledge the following individuals for their help and participation of the compilation of this document: Hardev Soor, Ixia, and Ralph Daniels, Spirent Communications, both who made significant contributions to the earlier versions of this document. In addition, the authors would like to acknowledge the members of the task team who helped bring this document to fruition: Michele Bustos, Tony De La Rosa, David Newman and Jerry Perser.



## **12. References**

### Normative 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.
- [Ma98] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", [RFC 2285](#), February 1998.

### Informative References

- [Ca02] Cain, B., et al., "Internet Group Management Protocol, Version 3", [RFC 3376](#), October 2002.
- [De89] Deering, S., "Host Extensions for IP Multicasting", STD 5, [RFC 1112](#), August 1989.
- [Fe97] Fenner, W., "Internet Group Management Protocol, Version 2", [RFC 2236](#), November 1997.
- [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.
- [Mt98] Maufer, T. "Deploying IP Multicast in the Enterprise." Prentice-Hall, 1998.



### **13. Author's Addresses**

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

Phone: + 1 818 871 1800  
EMail: debby@ixiacom.com

Brooks Hickman  
Spirent Communications  
26750 Agoura Rd.  
Calabasas, CA 91302  
USA

Phone: + 1 818 676 2412  
EMail: brooks.hickman@spirentcom.com

### **14. Full Copyright Statement**

"Copyright (C) The Internet Society (2003). 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.ö

