

## Terminology for IP Multicast Benchmarking

### Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

### Abstract

The purpose of this document is to define terminology specific to the benchmarking of multicast IP forwarding devices. It builds upon the tenets set forth in RFC 1242, RFC 2285, 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.

### 1. Introduction

Network forwarding devices are being required to take a single frame and support delivery to a number of destinations having membership to a particular group. As such, multicast support may place a different burden on the resources of these network forwarding devices than with unicast or broadcast traffic types.

Such burdens may not be readily apparent at first glance - the IP multicast packet's Class D address may be the only noticeable difference from an IP unicast packet. However, there are many factors that may impact the treatment of IP multicast packets.

Consider how a device's architecture may impact the handling of a multicast frame. For example, is the multicast packet subject to the same processing as its unicast analog? Or is the multicast packet treated as an exception and processed on a different data path?

Consider, too, how a shared memory architecture may demonstrate a different performance profile than an architecture which explicitly passes each individual packet between the processing entities.

In addition to forwarding device architecture, there are other factors that may impact a device's or system's multicast related performance. Protocol requirements may demand that routers and switches consider destination and source addressing in its multicast forwarding decisions. Capturing multicast source/destination addressing information may impact forwarding table size and lengthen lookups. Topological factors such as the degree of packet replication, the number of multicast groups being supported by the system, or the placement of multicast packets in unicast wrappers to span non-multicast network paths may all potentially affect a system's multicast related performance. For an overall understanding of IP multicasting, the reader is directed to [Se98], [Hu95], and [Mt98].

By clearly identifying IP multicast benchmarks and related terminology in this document, it is hoped that detailed methodologies can be generated in subsequent documents. Taken in tandem, these two efforts endeavor to assist the clinical, empirical, and consistent characterization of certain aspects of multicast technologies and their individual implementations. Understanding the operational profile of multicast forwarding devices may assist the network designer to better deploy multicast in his or her networking environment.

Moreover, this document focuses on one source to many destinations profiling. Elements of this document may require extension when considering multiple source to multiple destination IP multicast communication.

## 2. Definition Format

This section cites the template suggested by RFC 1242 in the specification of a term to be defined.

Term to be defined.

Definition:

The specific definition for the term.

Discussion:

A brief discussion of the term, its application, or other information that would build understanding.

**Measurement units:**

Units used to record measurements of this term, if applicable.

**[Issues:]**

List of issues or conditions that affect this term. This field can present items that may impact the term's related methodology or otherwise restrict its measurement procedures. This field is optional in this document.

**[See Also:]**

List of other terms that are relevant to the discussion of this term. This field is optional in this document.

## 2.1 Existing Terminology

This document draws on existing terminology defined in other BMWG work. Examples include, but are not limited to:

Throughput	[RFC 1242, section 3.17]
Latency	[RFC 1242, section 3.8]
Constant Load	[RFC 1242, section 3.4]
Frame Loss Rate	[RFC 1242, section 3.6]
Overhead behavior	[RFC 1242, section 3.11]
Forwarding Rates	[RFC 2285, section 3.6]
Loads	[RFC 2285, section 3.5]
Device Under Test (DUT)	[RFC 2285, section 3.1.1]
System Under Test (SUT)	[RFC 2285, section 3.1.2]

Note: "DUT/SUT" refers to a metric that may be applicable to a DUT or SUT.

## 3. Table of Defined Terms

### 3.1 General Nomenclature

- 3.1.1 Traffic Class. (TC)
- 3.1.2 Group Class. (GC)
- 3.1.3 Service Class. (SC)

### 3.2 Forwarding and Throughput

- 3.2.1 Mixed Class Throughput (MCT).
- 3.2.2 Scaled Group Forwarding Matrix (SGFM).
- 3.2.3 Aggregated Multicast Throughput (AMT)
- 3.2.4 Encapsulation Throughput (ET)
- 3.2.5 Decapsulation Throughput (DT)
- 3.2.6 Re-encapsulation Throughput (RET)

### 3.3 Forwarding Latency

#### 3.3.1 Multicast Latency (ML)

#### 3.3.2 Min/Max Multicast Latency (Min/Max ML)

### 3.4 Overhead

#### 3.4.1 Group Join Delay. (GJD)

#### 3.4.2 Group Leave Delay. (GLD)

### 3.5 Capacity

#### 3.5.1 Multicast Group Capacity. (MGC)

### 3.6 Interaction

#### 3.6.1 Burdened Response

#### 3.6.2 Forwarding Burdened Multicast Latency (FBML)

#### 3.6.3 Forwarding Burdened Join Delay (FBJD)

## 3.1 General Nomenclature

This section will present general terminology to be used in this and other documents.

### 3.1.1 Traffic Class. (TC)

#### Definition:

An equivalence class of packets comprising one or more data streams.

#### Discussion:

In the scope of this document, Traffic Class will be considered a logical identifier used to discriminate between a set or sets of packets offered the DUT.

For example, one Traffic Class may identify a set of unicast packets offered to the DUT. Another Traffic Class may differentiate the multicast packets destined to multicast group X. Yet another Class may distinguish the set of multicast packets destined to multicast group Y.

Unless otherwise qualified, the usage of the word "Class" in this document will refer simply to a Traffic Class.

#### Measurement units:

Not applicable.

### 3.1.2 Group Class. (GC)

Definition:

A specific type of Traffic Class where the packets comprising the Class are destined to a particular multicast group.

Discussion:

Measurement units:

Not applicable.

### 3.1.3 Service Class. (SC)

Definition:

A specific type of Traffic Class where the packets comprising the Class require particular treatment or treatments by the network forwarding devices along the path to the packets' destination(s).

Discussion:

Measurement units:

Not applicable.

## 3.2 Forwarding and Throughput.

This section presents terminology related to the characterization of the packet forwarding ability 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.

### 3.2.1 Mixed Class Throughput (MCT).

Definition:

The maximum 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.

Discussion:

Often times, throughput is collected on a homogenous traffic class - the offered load to the DUT is either singularly unicast or singularly multicast. In most networking environments, the traffic mix is seldom so uniformly distributed.

Based on the RFC 1242 definition for throughput, the Mixed Class Throughput benchmark attempts to characterize the DUT's ability to process both unicast and multicast frames in the same aggregated traffic stream.

**Measurement units:**

Frames per second

**Issues:**

Related methodology may have to address the ratio of unicast packets to multicast packets.

Since frame size can sometimes be a factor in frame forwarding benchmarks, the corresponding methodology for this metric will need to consider frame size distribution(s).

**3.2.2 Scaled Group Forwarding Matrix (SGFM).****Definition:**

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

**Discussion:**

A desirable attribute of many Internet mechanisms is the ability to "scale." This benchmark seeks to demonstrate the ability of a SUT to forward as the number of multicast groups is scaled upwards.

**Measurement units:**

Packets per second, with corresponding tested multicast group and port configurations.

**Issues:**

The corresponding methodology may have to reflect the impact that the pairing (source, group) has on many multicast routing protocols.

Since frame size can sometimes be a factor in frame forwarding benchmarks, the corresponding methodology for this metric will need to consider frame size distribution(s).

**3.2.3 Aggregated Multicast Throughput (AMT)****Definition:**

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

**Discussion:**

Another "scaling" type of exercise, designed to identify the DUT/SUT's ability to handle traffic as a function of the multicast destination ports it is required to support.

**Measurement units:**

The ordered pair (N,t) where,

N = the number of destination ports of the multicast group.

t = the throughput, in frames per second, relative to the source stream.

**Issues:**

Since frame size can sometimes be a factor in frame forwarding benchmarks, the corresponding methodology for this metric will need to consider frame size distribution(s).

**3.2.4 Encapsulation Throughput (ET)****Definition:**

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

**Discussion:**

A popular technique in presenting a frame to a device that may not support a protocol feature is to encapsulate, or tunnel, the packet containing the unsupported feature in a format that is supported by that device.

More specifically, encapsulation refers to the act of taking a frame or part of a frame and embedding it as a payload of another frame. This benchmark attempts to characterize the overhead behavior associated with that translational process.

**Measurement units:**

Frames per second.

**Issues:**

Consideration may need to be given with respect to the impact of different frame formats on usable bandwidth.

Since frame size can sometimes be a factor in frame forwarding benchmarks, the corresponding methodology for this metric will need to consider frame size distribution(s).

**3.2.5 Decapsulation Throughput (DT)****Definition:**

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

**Discussion:**

A popular technique in presenting a frame to a device that may not support a protocol feature is to encapsulate, or tunnel, the packet containing the unsupported feature in a format that is supported by that device. At some point, the frame may be required to be returned its original format from its encapsulation wrapper for use by the frame's next destination.

More specifically, decapsulation refers to the act of taking a frame or part of a frame embedded as a payload of another frame and returning it to the payload's appropriate format. This benchmark attempts to characterize the overhead behavior associated with that translational process.

**Measurement units:**

Frames per second.

**Issues:**

Consideration may need to be given with respect to the impact of different frame formats on usable bandwidth.

Since frame size can sometimes be a factor in frame forwarding benchmarks, the corresponding methodology for this metric will need to consider frame size distribution(s).

**3.2.6 Re-encapsulation Throughput (RET)****Definition:**

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

**Discussion:**

A popular technique in presenting a frame to a device that may not support a protocol feature is to encapsulate, or tunnel, the packet containing the unsupported feature in a format that is supported by that device. At some point, the frame may be required to be converted from one encapsulation format to another encapsulation format.

More specifically, re-encapsulation refers to the act of taking an encapsulated payload of one format and replacing it with another encapsulated format - all the while preserving the original payload's contents. This benchmark attempts to characterize the overhead behavior associated with that translational process.

**Measurement units:**

Frames per second.



**Issues:**

Consideration may need to be given with respect to the impact of different frame formats on usable bandwidth.

Since frame size can sometimes be a factor in frame forwarding benchmarks, the corresponding methodology for this metric will need to consider frame size distribution(s).

**3.3 Forwarding Latency.**

This section presents terminology relating to the characterization of the forwarding latency of a DUT/SUT in a multicast environment. It extends the concept of latency presented in RFC 1242.

**3.3.1 Multicast Latency. (ML)****Definition:**

The set of individual latencies from a single input port on the DUT or SUT to all tested ports belonging to the destination multicast group.

**Discussion:**

This benchmark is based on the RFC 1242 definition of latency. While it is useful to collect latency between a pair of source and destination multicast ports, it may be insightful to collect the same type of measurements across a range of ports supporting that Group Class.

A variety of statistical exercises can be applied to the set of latencies measurements.

**Measurement units:**

Time units with enough precision to reflect a latency measurement.

**3.3.2 Min/Max Multicast Latency. (Min/Max ML)****Definition:**

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

**Discussion:**

This statistic may yield some insight into how a particular implementation handles its multicast traffic. This may be useful to users of multicast synchronization types of applications.

**Measurement units:**

Time units with enough precision to reflect latency measurement.

### 3.4 Overhead

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

#### 3.4.1 Group Join Delay. (GJD)

Definition:

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

Discussion:

Many factors can contribute to different results, such as the number or type of multicast-related protocols configured on the device under test. Other factors are physical topology and "tree" configuration.

Because of the number of variables that could impact this metric, the metric may be a better characterization tool for a device rather than a basis for comparisons with other devices.

Issues:

A consideration for the related methodology: possible need to differentiate a specifically-forwarded multicast frame from those sprayed by protocols implementing a flooding tactic to solicit prune feedback.

While this metric attempts to identify a simple delay, the underlying and contributing delay components (e.g., propagation delay, frame processing delay, etc.) make this a less than simple measurement. The corresponding methodology will need to consider this and similar factors to ensure a consistent and precise metric result.

Measurement units:

Microseconds.

#### 3.4.2 Group Leave Delay. (GLD)

Definition:

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

**Discussion:**

While it is important to understand how quickly a device can process multicast frames; it may be beneficial to understand how quickly that same device can stop the process as well.

Because of the number of variables that could impact this metric, the metric may be a better characterization tool for a device rather than a basis for comparisons with other devices.

**Measurement units:**

Microseconds.

**Issues:**

The Methodology may need to consider protocol-specific timeout values.

While this metric attempts to identify a simple delay, the underlying and contributing delay components (e.g., propagation delay, frame processing delay, etc.) make this a less than simple measurement. Moreover, the cessation of traffic is a rather unobservable event (i.e., at what point is the multicast forwarded considered stopped on the DUT interface processing the Leave?). The corresponding methodology will need to consider this and similar factors to ensure a consistent and precise metric result.

### 3.5 Capacity

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

#### 3.5.1 Multicast Group Capacity. (MGC)

**Definition:**

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

**Discussion:****Measurement units:**

Multicast groups.

**Issues:**

The related methodology may have to consider the impact of multicast sources per group on the ability of a SUT/DUT to "scale up" the number of supportable multicast groups.

### 3.6 Interaction

Network forwarding devices are generally required to provide more functionality than than 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.

#### 3.6.1 Burdened Response.

**Definition:**

A measured response collected from a DUT/SUT in light of interacting, or potentially interacting, distinct stimuli.

**Discussion:**

Many metrics provide a one dimensional view into an operating characteristic of a tested system. For example, the forwarding rate metric may yield information about the packet processing ability of a device. Collecting that same metric in view of another control variable can oftentimes be very insightful. Taking that same forwarding rate measurement, for instance, while the device's address table is injected with an additional 50,000 entries may yield a different perspective.

**Measurement units:**

A burdened response is a type of metric. Metrics of this this type must follow guidelines when reporting results.

The metric's principal result **MUST** be reported in conjunction with the contributing factors.

For example, in reporting a Forwarding Burdened Latency, the latency measurement should be reported with respect to corresponding Offered Load and Forwarding Rates.

**Issues:** A Burdened response may be very illuminating when trying to characterize a single device or system. Extreme care must be exercised when attempting to use that characterization as a basis of comparison with other devices or systems. Test agents must ensure that the measured response is a function of the controlled stimuli, and not secondary factors. An example of of such an interfering factor would be configuration mismatch of a timer impacting a response process.

### 3.6.2 Forwarding Burdened Multicast Latency. (FBML)

Definition:

A multicast latency taken from a DUT/SUT in the presence of a traffic forwarding requirement.

Discussion:

This burdened response metric builds on the Multicast Latency definition offered in section 3.3.1. It mandates that the DUT be subjected to an additional measure of traffic not required by the non-burdened metric.

This metric attempts to provide a means by which to evaluate how traffic load may or may not impact a device's or system's packet processing delay.

Measurement units:

Time units with enough precision to reflect the latencies measurements.

Latency measurements MUST be reported with the corresponding sustained Forwarding Rate and associated Offered Load.

### 3.6.3 Forwarding Burdened Group Join Delay. (FBGJD)

Definition:

A multicast Group Join Delay taken from a DUT in the presence of a traffic forwarding requirement.

Discussion:

This burdened response metric builds on the Group Join Delay definition offered in section 3.4.1. It mandates that the DUT be subjected to an additional measure of traffic not required by the non-burdened metric.

Many factors can contribute to different results, such as the number or type of multicast-related protocols configured on the device under test. Other factors could be physical topology or the logical multicast "tree" configuration.

Because of the number of variables that could impact this metric, the metric may be a better characterization tool for a device rather than a basis for comparisons with other devices.

Measurement units:

Time units with enough precision to reflect the delay measurements.

Delay measurements MUST be reported with the corresponding sustained Forwarding Rate and associated Offered Load.

Issues:

While this metric attempts to identify a simple delay, the underlying and contributing delay components (e.g., propagation delay, frame processing delay, etc.) make this a less than simple measurement. The corresponding methodology will need to consider this and similar factors to ensure a consistent and precise metric result.

#### 4. Security Considerations

This document addresses metrics and terminology relating to the performance benchmarking of IP Multicast forwarding devices. The information contained in this document does not impact the security of the Internet.

Methodologies regarding the collection of the metrics described within this document may need to cite security considerations. This document does not address methodological issues.

#### 5. Acknowledgments

The IETF BMWG participants have made several comments and suggestions regarding this work. Particular thanks goes to Harald Alvestrand, Scott Bradner, Brad Cain, Eric Crawley, Bob Mandeville, David Newman, Shuching Sheih, Dave Thaler, Chuck Winter, Zhaohui Zhang, and John Galgay for their insightful review and assistance.

## 6. 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 1944, May 1996.
- [Hu95] Huitema, C. "Routing in the Internet." Prentice-Hall, 1995.
- [Se98] Semeria, C. and Maufer, T. "Introduction to IP Multicast Routing." <http://www.3com.com/nsc/501303.html> 3Com Corp., 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.

## 7. Author's Address

Kevin Dubray  
IronBridge Networks  
55 Hayden Avenue  
Lexington, MA 02421  
USA

Phone: 781 372 8118  
EMail: [kdubray@ironbridgenetworks.com](mailto:kdubray@ironbridgenetworks.com)

## 8. Full Copyright Statement

Copyright (C) The Internet Society (1998). 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 languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.