Network Working Group K. Dubray
INTERNET-DRAFT Bay Networks
Expiration Date: January 1998 July 1997

Terminology for IP Multicast Benchmarking
<draft-ietf-bmwg-mcast-02.txt>

Status of this Memo

This document is an Internet-Draft. 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.''

To view the entire list of current Internet-Drafts, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), ftp.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

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

Abstract

The purpose of this draft is to add terminology specific to the benchmarking of multicast IP forwarding devices. It builds upon the tenets set forth in RFC 1242, RFC 1944, and other IETF Benchmarking Methodology Working Group (BMWG) effort and extends them to the multicast paradigm.

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.

By clearly identifying 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.

[While primarily directed towards intermediate IP multicast forwarding devices on LANs, elements of this text may or may not be applicable to other media as well.]

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 and any restrictions on measurement procedures.

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

[Issues:]
List of issues or conditions that effect this term. This field is optional in this draft.

[See Also:]
List of other terms that are relevant to the discussion of this term. This field is optional in this draft.

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	([4], section 3.6)
Loads	([4], section 3.5)
Devices	([4], section 3.1)

3. Table of Defined Terms

3.2.2 Scaled Group Forwarding Matrix (SGFM).

3.1 General Nomenclature

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

3.1.1 Traffic Class.

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.

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.

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 relating 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.

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.

Discussion:
Often times, throughput is collected on a homogenous traffic type - though the packets' destinations may vary, the packets follow the same packet forwarding path through the DUT.

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.

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 (or even the definition itself) may have to reflect the impact that the pairing (source, group) has on many multicast routing protocols.

Refers to the concept of Forwarding Rate originally defined in this document. The definition of Forwarding Rate has been moved to [4].

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.

3.2.4 Translational Throughput (TT)

Definition:
The maximum rate at which none of the frames offered in an transitional format to the SUT are dropped in the process of converting those frames to their appropriate, final format and subsequent correct delivery.

Discussion:
A popular technique in presenting frames to devices 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. This benchmark attempts to characterize the overhead behavior associated with that transitional process.

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

Measurement units:
Frames per second.

3.3 Fairness.

Definition:
The ability of a SUT to fulfill the requirements of a Traffic Class without compromising the requirements, if any, of other Classes.

Discussion:

Measurement units:
Not applicable.

3.4 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.

3.4.1 Multicast Latency.

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 measurement.

3.4.2 Min/Max Multicast Latency.

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 measurement.

3.5 Overhead

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

3.5.1 Group Join Delay.

Definition:
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.

Discussion:
Many different factors can contribute to different results, such as the number or type of multicast-related protocols configured on the system under test.

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.

Measurement units:
Microseconds.

3.5.2 Group Leave Delay.

Definition:
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.

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

Measurement units:
Microseconds.

Issues: Methodology may need to consider protocol-specific timeout values.

3.6 Capacity

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

3.6.1 Multicast Group Capacity.

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.

4. Security Considerations

Security issues are not addressed in this memo.

5. References

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

[2] Bradner, S., McQuaid, J. Benchmarking Methodology for Network Interconnect Devices. RFC 1944. May, 1996.

[3] Craig, R. Terminology for Cell/Call Benchmarking. <draft-ietf-bmwg-call-01.txt> March, 1997. Work in progress.

[4] Mandeville, R. Benchmarking Terminology for LAN Switching Devices. <draft-ietf-bmwg-lanswitch-06.txt> July, 1997. Work in progress.

5. Author's Address

Kevin Dubray
Bay Networks, Inc.
2 Federal Street
Billerica, MA 01984
(508) 916-3862
kdubray@baynetworks.com

or direct discussion to the Benchmarking Methodology Working Group:
bmwg@harvard.edu