

Benchmarking Terminology for LAN Switching Devices

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.

Table of Contents

1. Introduction	2
2. Existing definitions	2
3. Term definitions	3
3.1 Devices.	3
3.1.1 Device under test (DUT)	3
3.1.2 System under test (SUT)	3
3.2 Traffic orientation.	4
3.2.1 Unidirectional traffic.	4
3.2.2 Bidirectional traffic	5
3.3 Traffic distribution	6
3.3.1 Non-meshed traffic.	6
3.3.2 Partially meshed traffic.	7
3.3.3 Fully meshed traffic.	8
3.4 Bursts	9
3.4.1 Burst	9
3.4.2 Burst size	10
3.4.3 Inter-burst gap (IBG).	10
3.5 Loads	11
3.5.1 Intended load (Iload)	11
3.5.2 Offered load (Oload)	12
3.5.3 Maximum offered load (MOL)	13
3.5.4 Overloading	14
3.6 Forwarding rates	15
3.6.1 Forwarding rate (FR)	15
3.6.2 Forwarding rate at maximum offered load (FRMOL).	16
3.6.3 Maximum forwarding rate (MFR).	16
3.7 Congestion control	17
3.7.1 Backpressure	17
3.7.2 Forward pressure	18

3.7.3 Head of line blocking	19
3.8 Address handling	20
3.8.1 Address caching capacity	20
3.8.2 Address learning rate	20
3.8.3 Flood count	21
3.9 Errored frame filtering	21
3.9.1 Errored frames	22
3.10 Broadcasts	22
3.10.1 Broadcast forwarding rate at maximum load	22
3.10.2 Broadcast latency	23
4. Security Considerations	24
5. References.	24
6. Acknowledgements.	24
7. Author's Address.	24
8. Full Copyright Statement.	25

1. Introduction

This document is intended to provide terminology for the benchmarking of local area network (LAN) switching devices. It extends the terminology already defined for benchmarking network interconnect devices in RFCs 1242 and 1944 to switching devices.

Although it might be found useful to apply some of the terms defined here to a broader range of network interconnect devices, this RFC primarily deals with devices which switch frames at the Medium Access Control (MAC) layer. It defines terms in relation to the traffic put to use when benchmarking switching devices, forwarding performance, congestion control, latency, address handling and filtering.

2. Existing definitions

RFC 1242 "Benchmarking Terminology for Network Interconnect Devices" should be consulted before attempting to make use of this document. RFC 1944 "Benchmarking Methodology for Network Interconnect Devices" contains discussions of a number of terms relevant to the benchmarking of switching devices and should also be consulted.

For the sake of clarity and continuity this RFC adopts the template for definitions set out in Section 2 of RFC 1242. Definitions are indexed and grouped together in sections for ease of reference.

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. Term definitions

3.1 Devices

This group of definitions applies to all types of networking devices.

3.1.1 Device under test (DUT)

Definition:

The network forwarding device to which stimulus is offered and response measured.

Discussion:

A single stand-alone or modular unit which receives frames on one or more of its interfaces and then forwards them to one or more interfaces according to the addressing information contained in the frame.

Measurement units:

n/a

Issues:

See Also:

system under test (SUT) (3.1.2)

3.1.2 System Under Test (SUT)

Definition:

The collective set of network devices to which stimulus is offered as a single entity and response measured.

Discussion:

A system under test may be comprised of a variety of networking devices. Some devices may be active in the forwarding decision-making process, such as routers or switches; other devices may be passive such as a CSU/DSU. Regardless of constituent components, the system is treated as a singular entity to which stimulus is offered and response measured.

Measurement units:

n/a

Issues:

See Also:

device under test (DUT) (3.1.1)

3.2 Traffic orientation

This group of definitions applies to the traffic presented to the interfaces of a DUT/SUT and indicates whether the interfaces are receiving only, transmitting only, or both receiving and transmitting.

3.2.1 Unidirectional traffic

Definition:

When all frames presented to the input interfaces of a DUT/SUT are addressed to output interfaces which do not themselves receive any frames.

Discussion:

This definition conforms to the discussion in section 16 of RFC 1944 which describes how unidirectional traffic can be offered to a DUT/SUT to measure throughput. Unidirectional traffic is also appropriate for:

-the measurement of the minimum inter-frame gap -the creation of many-to-one or one-to-many interface overload -the detection of head of line blocking -the measurement of forwarding rates and throughput when congestion control mechanisms are active.

When a tester offers unidirectional traffic to a DUT/SUT reception and transmission are handled by different interfaces or sets of interfaces of the DUT/SUT. All frames received from the tester by the DUT/SUT are transmitted back to the tester from interfaces which do not themselves receive any frames.

It is useful to distinguish traffic orientation and traffic distribution when considering traffic patterns used in device testing. Unidirectional traffic, for example, is traffic orientated in a single direction between mutually exclusive sets of source and destination interfaces of a DUT/SUT. Such traffic,

however, can be distributed between interfaces in different ways. When traffic is sent to two or more interfaces from an external source and then forwarded by the DUT/SUT to a single output interface the traffic orientation is unidirectional and the traffic distribution between interfaces is many-to-one. Traffic can also be sent to a single input interface and forwarded by the DUT/SUT to two or more output interfaces to achieve a one-to-many distribution of traffic.

Such traffic distributions can also be combined to test for head of line blocking or to measure forwarding rates and throughput when congestion control mechanisms are active.

When a DUT/SUT is equipped with interfaces running at different media rates the number of input interfaces required to load or overload an output interface or interfaces will vary.

It should be noted that measurement of the minimum inter-frame gap serves to detect violations of the IEEE 802.3 standard.

Issues:

half duplex / full duplex

Measurement units:

n/a

See Also:

bidirectional traffic (3.2.2)
non-meshed traffic (3.3.1)
partially meshed traffic (3.3.2)
fully meshed traffic (3.3.3)
congestion control (3.7)
head of line blocking (3.7.3)

3.2.2 Bidirectional traffic

Definition:

Frames presented to a DUT/SUT such that every receiving interface also transmits.

Discussion:

This definition conforms to the discussion in section 14 of RFC 1944.

When a tester offers bidirectional traffic to a DUT/SUT all the interfaces which receive frames from the tester also transmit frames back to the tester.

Bidirectional traffic MUST be offered when measuring the throughput or forwarding rate of full duplex interfaces of a switching device.

Issues:

truncated binary exponential back-off algorithm

Measurement units:

n/a

See Also:

unidirectional traffic (3.2.1)
non-meshed traffic (3.3.1)
partially meshed traffic (3.3.2)
fully meshed traffic (3.3.3)

3.3 Traffic distribution

This group of definitions applies to the distribution of frames forwarded by a DUT/SUT.

3.3.1 Non-meshed traffic

Definition:

Frames offered to a single input interface and addressed to a single output interface of a DUT/SUT where input and output interfaces are grouped in mutually exclusive pairs.

Discussion:

In the simplest instance of non-meshed traffic all frames are offered to a single input interface and addressed to a single output interface. The one-to-one mapping of input to output interfaces required by non-meshed traffic can be extended to multiple mutually exclusive pairs of input and output interfaces.

Measurement units:

n/a

Issues:

half duplex / full duplex

See Also:

unidirectional traffic (3.2.1)
bidirectional traffic (3.2.2)
partially meshed traffic (3.3.2.)
fully meshed traffic (3.3.3)
burst (3.4.1)

3.3.2 Partially meshed traffic

Definition:

Frames offered to one or more input interfaces of a DUT/SUT and addressed to one or more output interfaces where input and output interfaces are mutually exclusive and mapped one-to-many, many-to-one or many-to-many.

Discussion:

This definition follows from the discussion in section 16 of RFC 1944 on multi-port testing. Partially meshed traffic allows for one-to-many, many-to-one or many-to-many mappings of input to output interfaces and readily extends to configurations with multiple switching devices linked together over backbone connections.

It should be noted that partially meshed traffic can load backbone connections linking together two switching devices or systems more than fully meshed traffic. When offered partially meshed traffic devices or systems can be set up to forward all of the frames they receive to the opposite side of the backbone connection whereas fully meshed traffic requires at least some of the offered frames to be forwarded locally, that is to the interfaces of the DUT/SUT receiving them. Such frames will not traverse the backbone connection.

Measurement units:

n/a

Issues:

half duplex / full duplex

See Also:

- unidirectional traffic (3.2.1)
- bidirectional traffic (3.2.2)
- non-meshed traffic (3.3.1)
- fully meshed traffic (3.3.3)
- burst (3.4.1)

3.3.3 Fully meshed traffic

Definition:

Frames offered to a designated number of interfaces of a DUT/SUT such that each one of the interfaces under test receives frames addressed to all of the other interfaces under test.

Discussion:

As with bidirectional partially meshed traffic, fully meshed traffic requires each one the interfaces of a DUT/SUT to both receive and transmit frames. But since the interfaces are not divided into groups as with partially meshed traffic every interface forwards frames to and receives frames from every other interface. The total number of individual input/output interface pairs when traffic is fully meshed over n switched interfaces equals $n \times (n - 1)$. This compares with $n \times (n / 2)$ such interface pairs when traffic is partially meshed.

Fully meshed traffic on half duplex interfaces is inherently bursty since interfaces must interrupt transmission whenever they receive frames. This kind of bursty meshed traffic is characteristic of real network traffic and can be advantageously used to diagnose a DUT/SUT by exercising many of its component parts simultaneously. Additional inspection may be warranted to correlate the frame forwarding capacity of a DUT/SUT when offered meshed traffic and the behavior of individual elements such as input or output buffers, buffer allocation mechanisms, aggregate switching capacity, processing speed or medium access control.

The analysis of forwarding rate measurements presents a challenge when offering bidirectional or fully meshed traffic since the rate at which the tester can be observed to transmit frames to the DUT/SUT may be smaller than the rate at which it intends to transmit due to collisions on half duplex media or the action of congestion control mechanisms. This makes it important to take account of both the intended and offered loads defined in sections 3.5.1. and 3.5.2 below when reporting the results of such forwarding rate measurements.

When offering bursty meshed traffic to a DUT/SUT a number of variables have to be considered. These include frame size, the number of frames within bursts, the interval between bursts as well as the distribution of load between incoming and outgoing traffic. Terms related to bursts are defined in section 3.4 below.

Measurement units:

n/a

Issues:

half duplex / full duplex

See Also:

unidirectional traffic (3.2.1)
bidirectional traffic (3.2.2)
non-meshed traffic (3.3.1)
partially meshed traffic (3.3.2)
burst (3.4.1)
intended load (3.5.1)
offered load (3.5.2)

3.4 Bursts

This group of definitions applies to the intervals between frames or groups of frames offered to the DUT/SUT.

3.4.1 Burst

Definition:

A sequence of frames transmitted with the minimum legal inter-frame gap.

Discussion:

This definition follows from discussions in section 3.16 of RFC 1242 and section 21 of RFC 1944 which describes cases where it is useful to consider isolated frames as single frame bursts.

Measurement units:

n/a

Issues:

See Also:

burst size (3.4.2)
inter-burst gap (IBG) (3.4.3)

3.4.2 Burst size

Definition:

The number of frames in a burst.

Discussion:

Burst size can range from one to infinity. In unidirectional traffic as well as in bidirectional or meshed traffic on full duplex interfaces there is no theoretical limit to burst length. When traffic is bidirectional or meshed bursts on half duplex media are finite since interfaces interrupt transmission intermittently to receive frames.

On real networks burst size will normally increase with window size. This makes it desirable to test devices with small as well as large burst sizes.

Measurement units:

number of N-octet frames

Issues:

See Also:

burst (3.4.1)
inter-burst gap (IBG) (3.4.3)

3.4.3 Inter-burst gap (IBG)

Definition:

The interval between two bursts.

Discussion:

This definition conforms to the discussion in section 20 of RFC 1944 on bursty traffic.

Bidirectional and meshed traffic are inherently bursty since interfaces share their time between receiving and transmitting frames. External sources offering bursty traffic for a given frame size and burst size must adjust the inter-burst gap to achieve a specified average rate of frame transmission.

Measurement units:

- nanoseconds
- microseconds
- milliseconds
- seconds

Issues:

See Also:

- burst (3.4.1)
- burst size (3.4.2)

3.5 Loads

This group of definitions applies to the rates at which traffic is offered to any DUT/SUT.

3.5.1 Intended load (Iload)

Definition:

The number of frames per second that an external source attempts to transmit to a DUT/SUT for forwarding to a specified output interface or interfaces.

Discussion:

Collisions on CSMA/CD links or the action of congestion control mechanisms can effect the rate at which an external source of traffic transmits frames to a DUT/SUT. This makes it useful to distinguish the load that an external source attempts to apply to a DUT/SUT and the load it is observed or measured to apply.

In the case of Ethernet an external source of traffic MUST implement the truncated binary exponential back-off algorithm to ensure that it is accessing the medium legally

Measurement units:

bits per second
N-octets per second
(N-octets per second / media_maximum-octets per second) x 100

Issues:

See Also:

burst (3.4.1)
inter-burst gap (3.4.3)
offered load (3.5.2)

3.5.2 Offered load (Oload)

Definition:

The number of frames per second that an external source can be observed or measured to transmit to a DUT/SUT for forwarding to a specified output interface or interfaces.

Discussion:

The load which an external device can be observed to apply to a DUT/SUT may be less than the intended load due to collisions on half duplex media or the action of congestion control mechanisms. This makes it important to distinguish intended and offered load when analyzing the results of forwarding rate measurements using bidirectional or fully meshed traffic.

Frames which are not successfully transmitted by an external source of traffic to a DUT/SUT MUST NOT be counted as transmitted frames when measuring forwarding rates.

The frame count on an interface of a DUT/SUT may exceed the rate at which an external device offers frames due to the presence of spanning tree BPDUs (Bridge Protocol Data Units) on 802.1D-compliant switches or SNMP frames. Such frames should be treated as modifiers as described in section 11 of RFC 1944.

Offered load MUST be indicated when reporting the results of forwarding rate measurements.

Measurement units:

bits per second
N-octets per second
(N-octets per second / media_maximum-octets per second) x 100

Issues:

token ring

See Also:

bidirectional traffic (3.2.2)
fully meshed traffic (3.3.3)
intended load (3.5.1)
forwarding rate (3.6.1)

3.5.3 Maximum offered load (MOL)

Definition:

The highest number of frames per second that an external source can transmit to a DUT/SUT for forwarding to a specified output interface or interfaces.

Discussion:

The maximum load that an external device can apply to a DUT/SUT may not equal the maximum load allowed by the medium. This will be the case when an external source lacks the resources to transmit frames at the minimum legal inter-frame gap or when it has sufficient resources to transmit frames below the minimum legal inter-frame gap. Moreover, maximum load may vary with respect to parameters other than a medium's maximum theoretical utilization. For example, on those media employing tokens, maximum load may vary as a function of Token Rotation Time, Token Holding Time, or the ability to chain multiple frames to a single token. The maximum load that an external device applies to a DUT/SUT MUST be specified when measuring forwarding rates.

Measurement units:

bits per second
N-octets per second
(N-octets per second / media_maximum-octets per second) x 100

Issues:

See Also:

offered load (3.5.2)

3.5.4 Overloading

Definition:

Attempting to load a DUT/SUT in excess of the maximum rate of transmission allowed by the medium.

Discussion:

Overloading can serve to exercise buffers and buffer allocation algorithms as well as congestion control mechanisms. The number of input interfaces required to overload one or more output interfaces of a DUT/SUT will vary according to the media rates of the interfaces involved.

An external source can also overload an interface by transmitting frames below the minimum inter-frame gap. A DUT/SUT MUST forward such frames at intervals equal to or above the minimum gap specified in standards.

A DUT/SUT using congestion control techniques such as backpressure or forward pressure may exhibit no frame loss when a tester attempts to overload one or more of its interfaces. This should not be exploited to suggest that the DUT/SUT supports rates of transmission in excess of the maximum rate allowed by the medium since both techniques reduce the rate at which the tester offers frames to prevent overloading. Backpressure achieves this purpose by jamming the transmission interfaces of the tester and forward pressure by hindering the tester from gaining fair access to the medium. Analysis of both cases should take the distinction between intended load (3.5.1) and offered load (3.5.2) into account.

Measurement units:

bits per second

N-octets per second

$(N\text{-octets per second} / \text{media_maximum-octets per second}) \times 100$

Issues:

See Also:

- unidirectional traffic (3.2.1)
- intended load (3.5.1)
- offered load (3.5.2)
- forwarding rate (3.6.1)
- backpressure (3.7.1)
- forward pressure (3.7.2)

3.6 Forwarding rates

This group of definitions applies to the rates at which traffic is forwarded by any DUT/SUT in response to a stimulus.

3.6.1 Forwarding rate (FR)

Definition:

The number of frames per second that a device can be observed to successfully transmit to the correct destination interface in response to a specified offered load.

Discussion:

Unlike throughput defined in section 3.17 of RFC 1242, forwarding rate makes no explicit reference to frame loss. Forwarding rate refers to the number of frames per second observed on the output side of the interface under test and MUST be reported in relation to the offered load. Forwarding rate can be measured with different traffic orientations and distributions.

It should be noted that the forwarding rate of a DUT/SUT may be sensitive to the action of congestion control mechanisms.

Measurement units:

N-octet frames per second

Issues:

See Also:

- offered load (3.5.2)
- forwarding rate at maximum offered load (3.6.2)
- maximum forwarding rate (3.6.3)

3.6.2 Forwarding rate at maximum offered load (FRMOL)

Definition:

The number of frames per second that a device can be observed to successfully transmit to the correct destination interface in response to the maximum offered load.

Discussion:

Forwarding rate at maximum offered load may be less than the maximum rate at which a device can be observed to successfully forward traffic. This will be the case when the ability of a device to forward frames degenerates when offered traffic at maximum load.

Maximum offered load **MUST** be cited when reporting forwarding rate at maximum offered load.

Measurement units:

N-octet frames per second

Issues:

See Also:

maximum offered load (3.5.3)
forwarding rate (3.6.1)
maximum forwarding rate (3.6.3)

3.6.3 Maximum forwarding rate (MFR)

Definition:

The highest forwarding rate of a DUT/SUT taken from an iterative set of forwarding rate measurements.

Discussion:

The forwarding rate of a device may degenerate before maximum load is reached. The load applied to a device must be cited when reporting maximum forwarding rate.

The following example illustrates how the terms relative to loading and forwarding rates are meant to be used. In particular it shows how the distinction between forwarding rate at maximum offered load (FRMOL) and maximum forwarding rate (MFR) can be used to characterize a DUT/SUT.

	(A) Test Device Offered Load -----	(B) DUT/SUT Forwarding Rate -----
(1)	14,880 fps - MOL	7,400 fps - FRMOL
(2)	13,880 fps	8,472 fps
(3)	12,880 fps	12,880 fps - MFR

Measurement units:

N-octet frames per second

Issues:

See Also:

offered load (3.5.2)
forwarding rates (3.6.1)
forwarding rate at maximum load (3.6.2)

3.7 Congestion control

This group of definitions applies to the behavior of a DUT/SUT when congestion or contention is present.

3.7.1 Backpressure

Definition:

Any technique used by a DUT/SUT to attempt to avoid frame loss by impeding external sources of traffic from transmitting frames to congested interfaces.

Discussion:

Some switches send jam signals, for example preamble bits, back to traffic sources when their transmit and/or receive buffers start to overfill. Switches implementing full duplex Ethernet links may use IEEE 802.3x Flow Control for the same purpose. Such devices may incur no frame loss when external sources attempt to offer traffic to congested or overloaded interfaces.

It should be noted that jamming and other flow control methods may slow all traffic transmitted to congested input interfaces including traffic intended for uncongested output interfaces.

A DUT/SUT applying backpressure may exhibit no frame loss when a tester attempts to overload one or more of its interfaces. This should not be interpreted to suggest that the interfaces of the DUT/SUT support forwarding rates above the maximum rate allowed by the medium. In these cases overloading is only apparent since through the application of backpressure the DUT/SUT avoids overloading by reducing the rate at which the tester can offer frames.

Measurement units:

frame loss on congested interface or interfaces N-octet frames per second between the interface applying backpressure and an uncongested destination interface

Issues:

jamming not explicitly described in standards

See Also:

intended load (3.5.1)
offered load (3.5.2)
overloading (3.5.4)
forwarding rate (3.6.1)
forward pressure (3.7.2)

3.7.2 Forward pressure

Definition:

Methods which depart from or otherwise violate a defined standardized protocol in an attempt to increase the forwarding performance of a DUT/SUT.

Discussion:

A DUT/SUT may be found to inhibit or abort back-off algorithms in order to force access to the medium when contention occurs. It should be noted that the back-off algorithm should be fair whether the DUT/SUT is in a congested or an uncongested state. Transmission below the minimum inter-frame gap or the disregard of flow control primitives fall into this category.

A DUT/SUT applying forward pressure may eliminate all or most frame loss when a tester attempts to overload one or more of its interfaces. This should not be interpreted to suggest that the interfaces of the DUT/SUT can sustain forwarding rates above the maximum rate allowed by the medium. Overloading in such cases is only apparent since the application of forward pressure by the DUT/SUT enables interfaces to relieve saturated output queues by forcing access to the medium and concomitantly inhibiting the tester from transmitting frames.

Measurement units:

intervals between frames in microseconds
intervals in microseconds between transmission retries during
16 successive collisions.

Issues:

truncated binary exponential back-off algorithm

See Also:

intended load (3.5.1)
offered load (3.5.2)
overloading (3.5.4)
forwarding rate (3.6.1)
backpressure (3.7.1)

3.7.3 Head of line blocking

Definition:

Frame loss or added delay observed on an uncongested output interface whenever frames are received from an input interface which is also attempting to forward frames to a congested output interface.

Discussion:

It is important to verify that a switch does not slow transmission or drop frames on interfaces which are not congested whenever overloading on one of its other interfaces occurs.

Measurement units:

forwarding rate and frame loss recorded on an uncongested interface when receiving frames from an interface which is also forwarding frames to a congested interface.

Issues:

input buffers

See Also:

unidirectional traffic (3.2.1)

3.8 Address handling

This group of definitions applies to the address resolution process enabling a DUT/SUT to forward frames to their correct destinations.

3.8.1 Address caching capacity

Definition:

The number of MAC addresses per n interfaces, per module or per device that a DUT/SUT can cache and successfully forward frames to without flooding or dropping frames.

Discussion:

Users building networks will want to know how many nodes they can connect to a switch. This makes it necessary to verify the number of MAC addresses that can be assigned per n interfaces, per module and per chassis before a DUT/SUT begins flooding frames.

Measurement units:

number of MAC addresses per n interfaces, modules, or chassis

Issues:

See Also:

address learning rate (3.8.2)

3.8.2 Address learning rate

Definition:

The maximum rate at which a switch can learn new MAC addresses without flooding or dropping frames.

Discussion:

Users may want to know how long it takes a switch to build its address tables. This information is useful to have when considering how long it takes a network to come up when many users log on in the morning or after a network crash.

Measurement units:

frames with different source addresses per second

Issues:**See Also:**

address caching capacity (3.8.1)

3.8.3 Flood count**Definition:**

Frames forwarded to interfaces which do not correspond to the destination MAC address information when traffic is offered to a DUT/SUT for forwarding.

Discussion:

When recording throughput statistics it is important to check that frames have been forwarded to their proper destinations. Flooded frames MUST NOT be counted as received frames. Both known and unknown unicast frames can be flooded.

Measurement units:

N-octet valid frames

Issues:

spanning tree BPDUs.

See Also:

address caching capacity (3.8.1)

3.9 Errored frame filtering

This group of definitions applies to frames with errors which a DUT/SUT may filter.

3.9.1 Errored frames

Definition:

Frames which are over-sized, under-sized, misaligned or with an errored Frame Check Sequence.

Discussion:

Switches, unlike IEEE 802.1d compliant bridges, do not necessarily filter all types of illegal frames. Some switches, for example, which do not store frames before forwarding them to their destination interfaces may not filter over-sized frames (jabbers) or verify the validity of the Frame Check Sequence field. Other illegal frames are under-sized frames (runts) and misaligned frames.

Measurement units:

n/a

Issues:

See Also:

3.10 Broadcasts

This group of definitions applies to MAC layer and network layer broadcast frames.

3.10.1 Broadcast forwarding rate

Definition:

The number of broadcast frames per second that a DUT/SUT can be observed to deliver to all interfaces located within a broadcast domain in response to a specified offered load of frames directed to the broadcast MAC address.

Discussion:

There is no standard forwarding mechanism used by switches to forward broadcast frames. It is useful to determine the broadcast forwarding rate for frames switched between interfaces on the same card, interfaces on different cards in the same chassis and

interfaces on different chassis linked together over backbone connections. The terms maximum broadcast forwarding rate and broadcast forwarding rate at maximum load follow directly from the terms already defined for forwarding rate measurements in section 3.6 above.

Measurement units:

N-octet frames per second

Issues:

See Also:

forwarding rate at maximum load (3.6.2)
maximum forwarding rate (3.6.3)
broadcast latency (3.10.2)

3.10.2 Broadcast latency

Definition:

The time required by a DUT/SUT to forward a broadcast frame to each interface located within a broadcast domain.

Discussion:

Since there is no standard way for switches to process broadcast frames, broadcast latency may not be the same on all receiving interfaces of a switching device. The latency measurements SHOULD be bit oriented as described in section 3.8 of RFC 1242. It is useful to determine broadcast latency for frames forwarded between interfaces on the same card, on different cards in the same chassis and on different chassis linked over backbone connections.

Measurement units:

nanoseconds
microseconds
milliseconds
seconds

Issues:

See Also:

broadcast forwarding rate (3.10.1)

4. Security Considerations

Documents of this type do not directly effect the security of the Internet or of corporate networks as long as benchmarking is not performed on devices or systems connected to operating networks.

The document points out that switching devices may violate the IEEE 802.3 standard by transmitting frames below the minimum interframe gap or unfairly accessing the medium by inhibiting the backoff algorithm. Although such violations do not directly engender breaches in security, they may perturb the normal functioning of other interworking devices by obstructing their access to the medium. Their use on the Internet or on corporate networks should be discouraged.

5. References

- [1] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, July 1991.
- [2] Bradner, S., and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 1944, May 1996.

6. Acknowledgments

The Benchmarking Methodology Working Group of the IETF and particularly Kevin Dubray (Bay Networks) are to be thanked for the many suggestions they collectively made to help complete this document. Ajay Shah (WG), Jean-Christophe Bestaux (ENL), Henry Hamon (Netcom Systems), Stan Kopek (Digital) and Doug Ruby (Prominet) all provided valuable input at various stages of this project.

Special thanks go to Scott Bradner for his seminal work in the field of benchmarking and his many encouraging remarks.

7. Author's Address

Robert Mandeville
European Network Laboratories (ENL)
2, rue Helene Boucher
78286 Guyancourt Cedex
France

Phone: + 33 1 39 44 12 05
Mobile Phone + 33 6 07 47 67 10
Fax: + 33 1 39 44 12 06
EMail: bob.mandeville@eunet.fr

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.

