Benchmarking Methodology Working Group
Internet-Draft
Expires: January 2001

G. Trotter
Agilent Technologies

July 2000

Terminology for Forwarding Information Base (FIB) based Router Performance Benchmarking

<draft-ietf-bmwg-fib-term-00.txt>

Status of this Memo

1. Introduction

This document defines terms that are to be used in a methodology that determines the IP packet forwarding performance of IP routers as a function of the forwarding information base installed within the router.

The objective of this methodology is to evaluate the performance levels of IP routers as forwarding information bases continue to grow in size and complexity of structure.

This methodology utilizes the packet forwarding performance measurements described in [2]; reference will also be made to the associated terminology document [3] for these terms.

2. Overview

In order to measure the forwarding information base-based router performance, different forwarding information bases (5.3) are installed in the router. The two key elements describing the FIB are the FIB size (5.5) and FIB prefix distribution (5.6). As FIB sizes increase, and as FIB prefix distributions tend towards longer network prefixes (5.1), then it will take more time to match the

destination address within an IP packet and its corresponding entry

within the FIB. The FIB-dependent throughput, latency and frame loss rate (5.11, 5.12, 5.13), measured with fully meshed traffic flows (5.15), will reflect the change in performance of the router. Tests may need to be performed up to the maximum FIB size (5.7).

When configuring the router for these measurements, the routes need to be manually entered into the router, or advertised via a routing protocol. It may take some period of time (the FIB learning time (5.8)) before the router learns all the routes.

When routes are advertised into the router, the routes should be advertised in such a way so that route aggregation (5.14) does not occur. Also, the effect of a per-interface FIB cache (5.10) needs to be taken into account.

3. Existing Definitions

[3] should be consulted before attempting to make use of this document. [2] contains discussions of a number of terms relevant to the benchmarking of switching devices and should also be consulted.

4. Definition Format

The definition format is the equivalent to that defined in [3], and is repeated here for convenience:

X.x Term to be defined. (e.g., Latency)

Definition:
The specific definition for the term.

Discussion:
A brief discussion about the term, it's application and any restrictions on measurement procedures.

Measurement units:
The units used to report measurements of this term, if applicable.

Issues:
List of issues or conditions that effect this term.

See Also:
List of other terms that are relevant to the discussion of this term.

5. Definitions

5.1 Network Prefix

Definition:

See section 2.2.5, "Addressing Architecture", in [4].

A network prefix refers to the high-order 'n' bits of an IP address, identifying a particular network within which an IP host is located.

Discussion:
Network prefixes are represented as a 32 bit IP address with a mask indicating the number of bits representing the network prefix. I.e. 141.184.128 /17 indicates that the network prefix is 17 bits long.

Measurement units:
<n/a>

Issues:

See Also:
Network Prefix Length (5.2)

5.2 Network Prefix Length

Definition:
Refers to the length in bits of a network prefix. Network prefixes, using CIDR terminology, are typically referred to as 15.35.128 /17, indicating that the network prefix is 17 bits long.

Discussion:

Measurement units:
length in bits

Issues:

See Also:
network prefix (5.1)
forwarding information base prefix distribution (5.6)

5.3 Forwarding Information Base (FIB)

Definition:
As according to the definition in [4]:

"The table containing the information necessary to forward IP Datagrams, in this document, is called the Forwarding Information Base. At minimum, this contains the interface identifier and next hop information for each reachable destination network prefix."

A forwarding information base consists of [FIB size (5.7)]

FIB entries (5.4).

Discussion:

The forwarding information base is an abstract concept used to describe a database indexing network prefixes versus router port identifiers. This database may, in reality, be constructed in any fashion that optimizes the execution of the longest prefix length match algorithm, which maps the destination addresses within received IP packets against the entries in the forwarding information base.

Measurement units:
<none>

Issues:

See Also:
forwarding information base entry (5.4)
forwarding information base size (5.5)
forwarding information base prefix distribution (5.6)
maximum forwarding information base size (5.7)

5.4 Forwarding Information Base Entry

Definition:
A single entry within a forwarding information base, consisting of a network prefix and a router port identifier.

Discussion:
See (5.3).

Measurement units:
<n/a>

Issues:

See Also:
forwarding information base (5.3)
forwarding information base size (5.5)
forwarding information base prefix distribution (5.6)
maximum forwarding information base size (5.7)

5.5 Forwarding Information Base Size

Definition:
Refers to the number of forwarding information base entries within a forwarding information base.

Discussion:
The number of entries within a forwarding information base is one of the key elements that may influence the forwarding performance of a router. Generally, the more entries within

the forwarding information base, the longer it could take to find the longest matching network prefix within the forwarding information base.

Measurement units:
positive integer

Issues:

See Also:
forwarding information base (5.3)
forwarding information base entry (5.4)
forwarding information base prefix distribution (5.6)
maximum forwarding information base size (5.7)

5.6 Forwarding Information Base Prefix Distribution

Definition:
The distribution of network prefix lengths within the forwarding information base.

Discussion:
Network prefixes within the forwarding information base could be all of a single network prefix length, but, more realistically, the network prefix lengths will be distributed across some range.

Individual performance measurements will be made against FIBs populated with the same network prefix length, as well as against FIBs with some distribution of network prefix lengths.

The distribution of network prefix lengths will have an impact on the forwarding performance of a router. The longer the network prefix length, the longer it will take for a router to perform the longest length prefix match algorithm, and potentially the lower the performance of the router.

Measurement units:
The forwarding information base prefix distribution is expressed by a list of network prefix lengths and the percentage of entries within the forwarding information base with a particular network prefix length. For example, a forwarding information base prefix distribution is represented as:

{[/16, 10%], [/20, 36%], [/24, 54%]}

This indicates that 10% of the entries within the forwarding information base have a 16 bit network prefix length, 36% have a 20 bit network prefix length, and 54% have a 24 bit network prefix length.

Issues:

See Also:
forwarding information base (5.3)

forwarding information base entry (5.4)
forwarding information base size (5.5)
maximum forwarding information base size (5.7)

5.7 Maximum Forwarding Information Base Size

Definition:
The maximum number of forwarding information base entries that can be supported within the forwarding information base.

Discussion:
It is useful to know the maximum forwarding information base size for a router as it will be an indicator of the ability of the router to function within the given application space, and whether the router will be able to handle projected network growth.

As a benchmarking value, it is necessary to discover this value so that performance measurements can be made up to the maximum possible forwarding information base size.

Measurement units:
Positive integer.

Issues:
Could this value vary with the forwarding information base prefix distribution?

See Also:
forwarding information base (5.3)
forwarding information base entry (5.4)
forwarding information base size (5.5)
forwarding information base prefix distribution (5.6)

5.8 Forwarding Information Base Learning Time

Definition:
The forwarding information base learning time is the time taken for a router to receive and process received routing messages, and to construct (and, possibly to distribute amongst the interface cards in the router) the forwarding information base.

Discussion:
It takes time for a router to construct its forwarding information base. A router needs to process received routing packets, build the routing information database, select the best paths, build the forwarding information base and then possibly distribute the forwarding information base or a

subset thereof to the interface cards. This entire process can take several minutes with very large forwarding information bases.

When performing benchmarking tests that take the forwarding information base into account, time must be allocated for the router to process the routing information and to install the complete forwarding information base within itself, before performance measurements are made.

Measurement units:

Issues:

See Also:
forwarding information base (5.3)

5.9 Per-Interface Forwarding Information Base

Definition:
A complete copy of the forwarding information base, installed on a router's interface card to speed the destination address to network prefix lookup process.

Discussion:
Router manufacturers have developed many optimizations for routers, of which one optimization is to copy the forwarding information base to every interface card on the router. By doing this, destination address / network prefix lookups can be performed on the interface, unloading a router's CPU.

Measurement units:
<n/a>

Issues:

See Also:
forwarding information base (5.3)
per-interface forwarding information base cache (5.10)

5.10 Per-Interface Forwarding Information Base Cache

Definition:
A subset of a forwarding information base, installed on a router's interface card to speed the destination address / network prefix lookup process.

Discussion:
Prior to installing a complete copy of the forwarding information base on each interface of a router, a popular technique for speeding destination address lookups is to install a cache of frequently used routes on a router's interface.

The most frequently used routes are placed in the forwarding information base cache. IP packets whose destination address does not match a network prefix within the per-interface

forwarding information base cache are forwarded to a router's central processor for lookup in the complete forwarding information base.

The implication for benchmarking the performance of a router as a function of the forwarding information base is significant. IP packets whose destination address matches an entry within the per-interface forwarding information base cache could be forwarded more quickly than packets whose destination address does not match an entry within the per-interface forwarding information base cache.

To create useful benchmarks, the role of a per-interface forwarding cache needs to be considered. The nature of benchmarking tests to measure the impact of the forwarding performance of a router requires that the destination addresses within IP packets transmitted into the router be distributed amongst the total set of network prefixes advertised into the router. This negates the role of a per-interface forwarding information base cache, but serves to stress the forwarding information base-based packet forwarding performance of the router.

Measurement units:
<n/a>

Issues:

See Also:
forwarding information base (5.3)
per-interface forwarding information base (5.9)

5.11 Forwarding Information Base-dependent Throughput

Definition:
Throughput, as defined in [3], used in a context where the forwarding information base influences the throughput.

Discussion:
This definition for FIB-dependent throughput is added to distinguish the context of this measurement from that defined in [3].

Measurement units:
See [3].

Issues:

See Also:

forwarding information base-dependent latency (5.12)
forwarding information base-dependent frame loss rate (5.13)

5.12 Forwarding Information Base-dependent Latency

Definition:
Latency, as defined in [3], used in a context where the forwarding information base influences the throughput.

Discussion:
This definition for FIB-dependent latency is added to distinguish the context of this measurement from that defined in [3].

Measurement units:
See [3].

Issues:

See Also:
forwarding information base-dependent throughput (5.11)
forwarding information base-dependent frame loss rate (5.13)

5.13 Forwarding Information Base-dependent Frame Loss Rate

Definition:
Frame Loss Rate, as defined in [3], used in a context where the forwarding information base influences the throughput.

Discussion:
This definition for FIB-dependent frame loss rate is added to distinguish the context of this measurement from that defined in [3].

Measurement units:
See [3].

Issues:

See Also:
forwarding information base-dependent throughput (5.11)
forwarding information base-dependent latency (5.12)

5.14 Route Aggregation

Definition:
The ability of a router to collapse many forwarding information base entries into a single entry.

Discussion:
A router may aggregate routes in a forwarding information base into a single entry to conserve space.

When advertising routes into a router to perform benchmarking tests as a function of the forwarding information base

installed within the router, it is necessary to ensure that a router does not aggregate routes.

Thus, when routes are advertised to the router or installed statically, care must be taken to ensure that the router does not aggregate routes.

For example, if advertising a set of /24 network prefixes into a particular port on the router, 256 consecutive /24 routes, sharing a common leading 16 bits, should not be advertised on a single port. If this is done, then the router will install a single entry within the forwarding information base indicating that all networks matching a particular /16 network prefix are accessible through one particular entry.

Measurement units:
<none>

Issues:

See Also:

5.15 Fully Meshed Traffic Flows

Definition:
A traffic mesh that fully exercises the forwarding table continuously in the router under test.

Discussion:
In order to stress the forwarding information base lookup mechanisms within the router, packets need to be delivered on every interface such that every entry within the forwarding information base is accessed.

This implies that the set of destination addresses used by IP packets delivered to each interface of the router matches the set of all network prefixes advertised into the router. That is, IP packets are sent into every interface such that IP packets are directed through to every network advertised on every other interface on the router.

Fully meshed traffic flows ensure that the forwarding information base or per-interface forwarding information bases are fully and continuously exercised.

Measurement units:
<n/a>

Issues:

See Also:
fully meshed traffic (section 3.3.3 in [5])

6. Security Considerations

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

7. References

1 Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
2 Bradner, S., McQuaid, J., "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999
3 Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, July 1991
4 Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995
5 Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998

8. Acknowledgments

9. Author's Addresses

Guy Trotter
Agilent Technologies (Canada) Inc.
#2500 4710 Kingsway
Burnaby, British Columbia
Canada
V5H 4M2
Phone: +1 604 454 3516
Email: Guy_Trotter@agilent.com

Full Copyright Statement

followed, or as required to translate it into languages other than English.