

Network Working Group
Request for Comments: 3222
Category: Informational

G. Trotter
Agilent Technologies
December 2001

Terminology for Forwarding Information Base (FIB) based Router
Performance

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 (2001). All Rights Reserved.

Abstract

This document describes the terms 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 a router. The forwarding performance of an IP router may be dependent upon or may be linked to the composition and size of the forwarding information base installed within a router.

Table of Contents

1. Introduction.....	2
2. Overview.....	3
3. Existing Definitions.....	3
4. Definition Format.....	3
5. Definitions - parameters.....	4
5.1 Network Prefix.....	4
5.2 Network Prefix Length.....	4
5.3 Forwarding Information Base (FIB).....	5
5.4 Forwarding Information Base Entry.....	6
5.5 Forwarding Information Base Size.....	6
5.6 Longest Length Prefix Match Algorithm.....	7
5.7 Forwarding Information Base Prefix Distribution.....	7
5.8 Per-Interface or Per-Card Forwarding Information Base.....	8
5.9 Per-Interface Forwarding Information Base Cache.....	9
5.10 Route Aggregation.....	10
6. Definitions - metrics.....	10
6.1 Maximum Forwarding Information Base Size.....	11
6.2 Forwarding Information Base Learning Time.....	11
6.3 Forwarding Information Base-dependent Throughput.....	12
6.4 Forwarding Information Base-dependent Latency.....	12
6.5 Forwarding Information Base-dependent Frame Loss Rate.....	13
7. Security Considerations.....	13
8. References.....	13
9. Author's Address.....	14
10. Full Copyright Statement.....	15

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 (5.7). The forwarding performance of a router may be dependent upon these two primary factors, particularly if FIB prefix distributions tend towards longer network prefixes (5.1). The FIB-dependent throughput, latency and frame loss rate (6.3, 6.4, 6.5), measured with fully meshed traffic flows [2], will reflect the change in performance of the router. Tests may need to be performed up to the maximum FIB size (6.1).

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 (6.2)) 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 (5.10) does not occur. Also, the effect of a per-interface FIB cache (5.9) 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 network interconnect 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 - parameters

This section defines parameters that would dictate the execution of methodology to determine the FIB based forwarding performance of a router.

5.1 Network Prefix

Definition:

"A network prefix is . . . a contiguous set of bits at the more significant end of the address that defines a set of systems; host numbers select among those systems."

(This definition is taken directly from section 2.2.5.2, "Classless Inter Domain Routing (CIDR)", in [4].)

Discussion:

In the CIDR context, the network prefix is the network component of an IP address. A common alternative to using a bitwise mask to communicate this component is the use of "slash (/) notation." Slash notation binds the notion of network prefix length (see 5.2) in bits to an IP address. E.g., 141.184.128.0/17 indicates the network component of this IPv4 address is 17 bits wide.

Measurement units:

<n/a>

Issues:**See Also:**

Network Prefix Length (5.2)

5.2 Network Prefix Length

Definition:

The number of bits used to define the network prefix. Network prefixes, using CIDR terminology, are typically referred to as 15.35.128.0 /17, indicating that the network prefix is 17 bits long.

Discussion:

When referring to groups of addresses, the network prefix length is often used as a means of describing groups of addresses as an equivalence class. For example, 100 /16 addresses refers to 100 addresses whose network prefix length is 16 bits.

Measurement units:

bits

Issues:**See Also:**

network prefix (5.1)

forwarding information base prefix distribution (5.7)

5.3 Forwarding Information Base (FIB)

Definition:

As according to the definition in Appendix B of [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."

Discussion:

The forwarding information base describes a database indexing network prefixes versus router port identifiers.

A forwarding information base consists of [FIB size (5.5)] FIB entries (5.4).

The forwarding information base is distinct from the "routing table" (or, the Routing Information Base), which holds all routing information received from routing peers.

The forwarding information base contains unique paths only (i.e. does not contain secondary paths).

Measurement units:

<none>

Issues:

See Also:

- forwarding information base entry (5.4)
- forwarding information base size (5.5)
- forwarding information base prefix distribution (5.7)
- maximum forwarding information base size (6.1)

5.4 Forwarding Information Base Entry

Definition:

A single entry within a forwarding information base. This entry consists of the minimum amount of information necessary to make a forwarding decision on a particular packet. The typical components within a forwarding information base entry are a network prefix, a router port identifier and next hop information. This is an entry that the router can and does use to forward packets.

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 (5.7)
- maximum forwarding information base size (6.1)

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:

Number of routes

Issues:

See Also:

- forwarding information base (5.3)
- forwarding information base entry (5.4)
- forwarding information base prefix distribution (5.7)
- maximum forwarding information base size (6.1)

5.6 Longest Length Prefix Match Algorithm

Definition:

An algorithm that a router uses to quickly match destination addresses within received IP packets to exit interfaces on the router.

Discussion:

Measurement Units:

<none>

Issues:

See Also:

5.7 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 may 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, 100], [/20, 360], [/24, 540]}
```

This indicates that 100 of the entries within the forwarding information base have a 16 bit network prefix length, 360 have a 20 bit network prefix length, and 540 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 (6.1)

5.8 Per-Interface or Per-Card Forwarding Information Base

Definition:

A complete copy of the forwarding information base, installed on a router's card or individual physical interface 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 or interface card on the router. By doing this, destination address / network prefix lookups can be performed on the interface or card, unloading a router's CPU.

Measurement units:

<n/a>

Issues:**See Also:**

- forwarding information base (5.3)
- per-interface forwarding information base cache (5.9)

5.9 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 (5.8)

5.10 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.

Route aggregation on a router can be turned off, but routes should still be advertised into the router in such a manner as to avoid route aggregation.

Measurement units:

<none>

Issues:

See Also:

6. Definitions - metrics

This section defines the metrics, or results, that would characterized the FIB based forwarding performance of a router.

6.1 Maximum Forwarding Information Base Size

Definition:

The maximum number of forwarding information base entries that can be supported within the forwarding information base. The Maximum Forwarding Information Base Size is the size over which all entries can and are used to forward traffic.

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:

Number of routes

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 (5.7)

6.2 Forwarding Information Base Learning Time

Definition:

The time a router takes to process received routing messages, and to construct (and, possibly to distribute amongst the interface cards in the router) the forwarding information base. This is measured from the time at which a router is presented with the first routing message, through to when it can forward packets using any entry in 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:
Prefixes per second.

Issues:

See Also:
forwarding information base (5.3)

6.3 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 (6.4)
forwarding information base-dependent frame loss rate (6.5)

6.4 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 (6.3)

forwarding information base-dependent frame loss rate (6.5)

6.5 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 (6.3)

forwarding information base-dependent latency (6.4)

7. 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.

8. References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [2] Bradner, S. and J. McQuaid, "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.

9. Author's Address

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

10. Full Copyright Statement

Copyright (C) The Internet Society (2001). 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

