

Methodology for Forwarding Information Base (FIB) based Router  
Performance  
<[draft-ietf-bmwg-fib-meth-01.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [1].

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

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. This document describes the methodology to be used to determine the IP packet forwarding performance of an IP router as a function of the forwarding information base installed within a router.

Table of Contents

<a href="#">1. Introduction.....</a>	<a href="#">2</a>
<a href="#">2. Terms Of Reference.....</a>	<a href="#">2</a>
<a href="#">3. Overview.....</a>	<a href="#">2</a>
<a href="#">4. Router Setup.....</a>	<a href="#">3</a>
<a href="#">5. Test Setup.....</a>	<a href="#">3</a>
<a href="#">6. Methodology Of Route Advertisements.....</a>	<a href="#">3</a>
<a href="#">7. Table Size.....</a>	<a href="#">4</a>
<a href="#">8. Table Composition.....</a>	<a href="#">5</a>
<a href="#">9. Offered Packet Load.....</a>	<a href="#">5</a>

<a href="#">10. Packet Format/Size.....</a>	<a href="#">6</a>
<a href="#">11. Determining The Maximum Fib Size.....</a>	<a href="#">6</a>
<a href="#">12. Methodology.....</a>	<a href="#">7</a>
<a href="#">12.1 Baseline Measurements.....</a>	<a href="#">7</a>

<a href="#">12.2 Maximum Fib Size Test.....</a>	<a href="#">7</a>
<a href="#">12.3 Fib Size Impact On Packet Forwarding Performance.....</a>	<a href="#">8</a>
<a href="#">13. Security Considerations.....</a>	<a href="#">9</a>
<a href="#">14. References.....</a>	<a href="#">9</a>
<a href="#">15. Acknowledgments.....</a>	<a href="#">9</a>
<a href="#">16. Author's Addresses.....</a>	<a href="#">9</a>

## [1. Introduction](#)

This document covers the measurement of the IP packet forwarding performance of IP routers on the basis of the forwarding information base (FIB) installed within a router. [2] describes the terminology associated with this document.

The purpose of this document is to describe a methodology for measuring the impact of a FIB of a particular size and composition on the forwarding performance of a router. The objective is to determine whether a router can maintain performance levels as FIBs continue to grow in size.

This document utilizes the methodology described in [3] for measuring the FIB-dependent throughput, FIB-dependent latency and FIB-dependent frame loss rate of IP packets as they traverse the router under test. The forwarding performance of a router should be observed under different FIB sizes and compositions.

In this document, the FIB is modified only when traffic is not traversing the router. The effect of dynamic changes on the forwarding performance of a router is outside the scope of this document.

## [2. Terms Of Reference](#)

This document utilizes the methodologies for packet throughput, latency and loss measurements described in [3].

Definitions unique to this test methodology are covered in [2].

This document covers only IPv4 packets traversing an IP router.

Applicability of this document to other packet types (such as MPLS-prefixed packets or non-IP packets) is not appropriate.

### [3.](#) Overview

This methodology takes into account two key features of a FIB - the FIB size and the FIB prefix distribution.

This methodology does not specify a particular means of populating a FIB - a routing protocol should be used, although a FIB could be manually entered into a router. One implication of this is that there needs to be protocol-independent means for determining when a router's FIB has been fully populated (described in 12.2).

Trotter

[Page 2]

---

FIB-based Router Performance

February, 2003

Once a FIB of a particular size and prefix distribution is entered into the router, [3] style measurements are performed in order to gauge the effect of the FIB on packet FIB-dependent throughput, FIB-dependent latency and FIB-dependent frame loss rate.

### [4.](#) Router Setup

This is an out-of-service methodology. The router should be configured realistically. That is, it should be configured with physical interface cards, switching fabrics, and route processors providing the basis for realistic product comparisons.

### [5.](#) Test Setup

The router is connected to test equipment capable of advertising network prefixes into the router under test and generating IP packets to these advertised networks. In the absence of test gear capable of advertising network prefixes into the router under test, the network prefixes could be manually entered into the router.

This methodology does not specify or require a particular physical interface type: all measurements are performed exclusively on IP packets.

### [6.](#) Methodology Of Route Advertisements

It is recommended that routes be advertised into the router under test by using a routing protocol to speed the advertisement of routes. Routes may be entered into the FIB manually, but this is

both tedious, and makes the correlation between entered routes and offered traffic awkward (see [section 9](#)).

Routes may be advertised into the router on a single interface, or over multiple interfaces. This methodology does not distinguish between either method.

### [6.1](#) Order of Route Advertisements

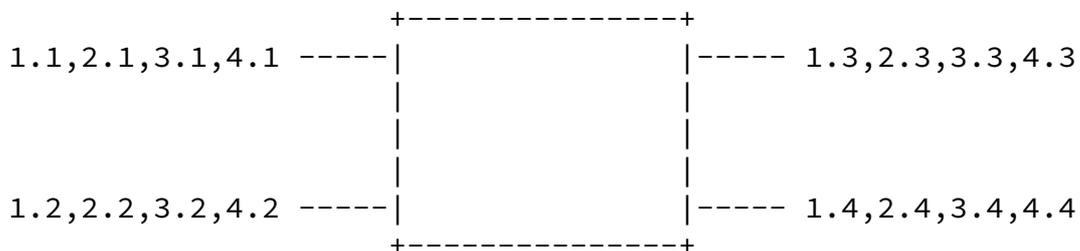
Routes can be advertised into the router in any order - a router may organize its FIB to maximize its performance - the impact of order of advertisement cannot be presupposed to have an impact on the structure of the FIB.

### [6.2](#) Prevention of Route Aggregation

Routes should be advertised into the router in such a manner as to prevent route aggregation. Routes may be aggregated as follows. If 100 routes are advertised on a single interface within a common prefix of 192.47.0.0, and no other advertisements on any other interface have a prefix of 192.47.0.0, then the router may enter a single /16 entry into the FIB, directing all traffic with a 192.47.0.0 prefix to a particular interface. Care must be taken to

prevent this from occurring, as it will result in smaller than intended FIBs being entered within the router under test.

As an example of a means of router advertisement designed to prevent route aggregation, consider the advertisement of 16 /16 prefixes into a 4-port router in the diagram below:



In this example, the router will not be able to aggregate the entries from 1.0, since an entry exists on every port - the same applies for the entries from 2.0, 3.0 and 4.0. This method of route advertisement into the router must be used.

### [6.3](#) Learning time

A router takes time to "learn" routes. The overhead associated with learning a route may include processing the routing protocol message containing a new FIB entry, placing the entry into the FIB, and then, in some modern routers, distributing the updated FIB to each interface in the router.

This methodology does not measure or dictate the time that a router takes to build a complete FIB. However, it is vital to the execution of this methodology to ensure that a router has completed learning routes and building the forwarding information base before tests proceed.

There are two methods that can be taken to ensure that a router has learned all routes. For the first method, the router can be queried via an attached console to reveal the current FIB size. If access to the console is not available, then test packets can be sent through the router at such rate that packet loss does not ordinarily occur. Test packets should test each entry in the FIB. No packet loss can indicate that the router has learned all routes, and is forwarding packets to all advertised destinations. However, this method is not foolproof - a router may instead have filled its FIB, and lost packets may indicate that the FIB is full, and not that the router is still learning routes.

### [7.](#) Table Size

This methodology does not specify FIB sizes to be installed within the router under test. However, it is expected that this methodology will be used through to the maximum possible FIB size purported to be supported by the router under test.

This methodology recommends starting with a small FIB size - a single network per physical interface - in order to perform a basic performance benchmark on the router under test. Subsequent tests will increase the FIB by size 'N'. The FIB will be increased through to the maximum size possible (see [section 11](#) for a discussion of determining the maximum possible FIB size).

### [8.](#) Table Composition

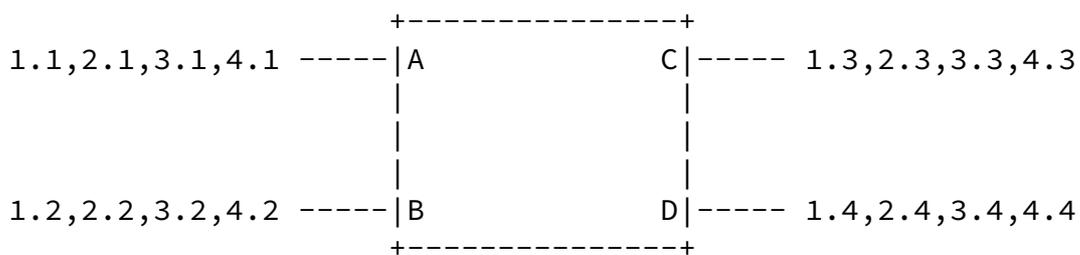
The composition of a FIB consists of a number FIB entries, with each FIB entry containing a network prefix. The network prefix length of each FIB entry may vary. These network prefixes typically range, in a typical Internet router's FIB, from 8 to 30 bits in length (/8 to /30) (this document does not attempt to pre-suppose the composition of a typical FIB within a router in the Internet, however). Destination address matches within FIBs consisting entirely of 8-bit network prefixes may be very rapid (few bits in each destination address/network prefix need to be matched), whereas it can be expected that destination address matches within a FIB populated entirely with 30-bit network prefixes (highly unlikely) will be slower.

## 9. Offered Packets

To gauge the impact of a FIB on the performance of a router, IP packets must be delivered to the router such that the FIB is fully exercised. That is, packets must be delivered into the router on every interface directed to stations whose network prefix is advertised on every other interface.

This method of traffic generation fairly exercises the FIB on the router. Modern routers may maintain a cache of the most frequently used routes on a per-interface (or, per-line card) basis, or they may maintain a complete copy of the FIB on a per-interface (or per-line card) basis. Generating packets from every interface to every advertised network ensures that each copy of the FIB is stressed, as well as ensuring that a router's FIB cannot be optimized (i.e. by installing summary routes).

Consider the diagram below:



In this diagram, packets sent into interface A must be directed to a destination address present in each network advertised on ports B, C and D. That is, the "transmit packet list" for packets sent into port A must be:

1.2.0.1  
1.3.0.1  
1.4.0.1  
2.2.0.1  
2.3.0.1  
2.4.0.1  
3.2.0.1  
3.3.0.1  
3.4.0.1  
4.2.0.1  
4.3.0.1  
4.4.0.1

and then repeated.

Note that packets are transmitted to destinations on each port in a round-robin manner. That is, A transmits a packet to B, then a packet to C and then a packet to D, and then repeats, using the next destination advertised on port A, and so on. This method ensures that all aspects of the router (switching fabric, interface cards, route processor, etc.) are stressed equally.

Also note that packets are sent from all interfaces simultaneously, again ensuring that all router components are stressed equally. See [section 16](#) in [3].

#### 10. Packet Format/Size

See [3] for a discussion of packet formats and sizes.

#### 11. Determining The Maximum Fib Size

In order to determine the maximum FIB size, a method similar to that used to determine whether all routes have been learned (see [section 6.3](#)) will be used.

Without observing the exchange of the routing protocol, it is necessary to ascertain the size of the FIB from observation of packets sent from the router.

The most methodical way of discovering the maximum FIB size is to advertise network prefixes to the FIB, test the addition of the entry by sending packets to the new network, ensure receipt of this packet and repeat until packets are no longer forwarded to the newly added network. Care must be taken to ensure that packets continue to be forwarded to all advertised networks, even while new networks are added - a router may "age-out" old entries in order to ensure that its FIB is filled with current entries.

The absence of packets received at a particular network is indicative of entries not being present within the FIB, and thus indicates that the FIB has filled.

If the approximate FIB size is known in advance, a larger number of routes can be advertised into the router, then the maximum table size can be found incrementally from this point.

When using IP packets to discover the maximum sized FIB supported by the DUT, IP packets should be sent in at a rate that is significantly below that of the maximum measured throughput of the router discovered during baseline measurements (see [section 12.1](#)).

## [12.](#) Methodology

This methodology makes use of the packet forwarding measurements described in [RFC 2544](#).

### [12.1](#) Baseline Measurements

**Objective:** The objective of the baseline measurement is to determine the nominal packet forwarding performance of the router, with a minimal sized FIB.

**Procedure:** A route to a single network should be advertised on each port of the router under test - a router with 64 ports would thus have an installed FIB of 64 entries.

As described in [section 9](#), packets should be offered to the router in order to create a full mesh between all advertised networks.

Using measurements described in [3] for the throughput, latency, and frame loss rates (as per sections [26.1](#), [26.2](#), [26.3](#) and [26.4](#) in [3]), the FIB-dependent throughput, FIB-dependent latency and FIB-dependent frame loss rate should be measured and the results duly noted.

**Reporting Format:** Results should be reported as per [3].

### [12.2](#) Maximum FIB Size Test

**Objective:** The objective of this test is to determine the maximum size of the FIB that can be accommodated within the DUT. Discovery

of the maximum limit of the size of the FIB is necessary for the establishment of limits in subsequent tests.

Procedure: Initially, advertise a single network on every port as per [section 12.1](#). Establish a traffic mesh between all advertised networks as per [section 9](#). Transmit IP packets through the DUT to all advertised networks at a rate well below the maximum throughput rate of the router, discovered in [section 12.1](#). Ensure that packet loss is zero (indicating that packets are being forwarded to all advertised networks).

Repeatedly advertise additional networks into the DUT and add these networks into the traffic mesh established through the DUT to the

advertised networks. Stop when packet loss is detected, or when an ICMP destination unreachable message with a code value of 0 (network unreachable) or 6 (destination network unknown) is received. This will indicate that the FIB has filled (or, that the DUT is unable to add additional entries into the FIB). The test should continue in order to establish increasing packet loss as a consequence of increasing networks advertised into the DUT.

Reporting Format: The results should be reported as a single number (the maximum FIB size). The results should be backed up with a graph of the packet loss (on the y-axis) versus the networks advertised into the DUT (on the x-axis).

### [12.3](#) FIB Size Impact On Packet Forwarding Performance

Objective: The objective of this test is to evaluate the impact of FIB size and composition upon the packet forwarding performance of the DUT.

Procedure: This iterative test changes both the size and composition of the FIB, and makes performance measurements for each iteration. That is, the FIB is set to a certain size and composition, then performance measurements are made for packet throughput, latency, and packet loss rate. The FIB is incremented, then the tests are repeated, through to the maximum sized FIB supported by the DUT.

The following code segment illustrates this methodology:

```
set forward_table_compositions [List 16 22 24 realistic]
```

```

for (forwarding_table_size = num_physical_interfaces;
    Forwarding_table_size <= maximum_forwarding_table_size;
    Forwarding_table_size += increment)
{
    foreach value forward_table_compositions_
    {
        set_forwarding_table_composition $value
        measure_FIB-dependent_throughput;
        measure_FIB-dependent_latency;
        measure_FIB-dependent_packet_loss_rate;
        measure_FIB-dependent_back-to-back_frames;
    }
}

```

Reporting Format: Results will be reported as a series of [3]-style graphs, with a separate set of graphs for each individual FIB size / FIB composition test. Results could be overlaid onto a single set of graphs.

It may be useful to create a set of graphs to aid in direct interpretation of this test. One recommended graph would be, for a particular frame size, a graph of the appropriate metric on the y-

axis (FIB-dependent throughput, FIB-dependent latency, FIB-dependent packet loss) versus the FIB size, with overlaid lines for each FIB composition.

### 13. Security Considerations

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

### 14. Informative References

- 1 Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- 2 Trotter, G., "Terminology for Forwarding Information Base (FIB) based Router Performance", [RFC 3222](#), December, 2001.
- 3 Bradner, S., McQuaid, J., "Benchmarking Methodology for Network Interconnect Devices", [RFC 2544](#), March 1999

[15. Acknowledgements](#)

[16. 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

"Copyright (C) The Internet Society (2002). 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.

