



IP Forwarding Application Level Benchmark

May 12, 2003

Revision 1.6

Editor(s):

Robert Peschi, Alcatel, Robert.Peschi@alcatel.be

Prashant Chandra, Intel Corporation, prashant.chandra@intel.com

Manohar R. Castelino, Intel Corporation, manohar.r.castelino@intel.com

Copyright © 2002 The Network Processing Forum (NPF). 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 other than the following, (1) the above copyright notice and this paragraph must be included on all such copies and derivative works, and (2) this document itself may not be modified in any way, such as by removing the copyright notice or references to the NPF, except as needed for the purpose of developing NPF Implementation Agreements.

By downloading, copying, or using this document in any manner, the user consents to the terms and conditions of this notice. Unless the terms and conditions of this notice are breached by the user, the limited permissions granted above are perpetual and will not be revoked by the NPF or its successors or assigns.

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN "AS IS" BASIS WITHOUT ANY WARRANTY OF ANY KIND. THE INFORMATION, CONCLUSIONS AND OPINIONS CONTAINED IN THE DOCUMENT ARE THOSE OF THE AUTHORS, AND NOT THOSE OF NPF. THE NPF DOES NOT WARRANT THE INFORMATION IN THIS DOCUMENT IS ACCURATE OR CORRECT. THE NPF DISCLAIMS ALL WARRANTIES, WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING BUT NOT LIMITED THE IMPLIED LIMITED WARRANTIES OF MERCHANTABILITY, TITLE OR FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OF THIRD PARTY RIGHTS.

The words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in the remainder of this document are to be interpreted as described in the NPF Software API Conventions Implementation Agreement revision 1.0.

For additional information contact:

**The Network Processing Forum, 39355 California Street,
Suite 307, Fremont, CA 94538
+1 510 608-5990 phone ♦ info@npforum.org**

Table of Contents

Table of Contents	3
1 Revision History	4
2 Scope And Purpose	5
3 Normative References.....	6
4 Terminology.....	7
4.1 Benchmarking Terminology	7
5 Test Configuration	9
5.1 Reference Design	9
5.2 Test Setup.....	12
6 Benchmark Tests.....	13
6.1 Data Plane Tests.....	13
6.1.1 Aggregate Forwarding Rate.....	15
6.1.2 Throughput.....	16
6.1.3 Latency.....	17
6.1.4 Loss Ratio	18
6.1.5 Overload Forwarding Rate.....	18
6.1.6 System Power Consumption	19
6.2 Control Plane Tests	20
6.2.1 Forwarding table Update Rate	20
6.3 Concurrent Data Plane and Control Plane Tests.....	21
6.3.1 Route Update and Forwarding Rate.....	21

1 Revision History

Revision	Date	Reason for Changes
1.1	12/4/01	Updated the IPv4 Forwarding Application Level Benchmark to include IPv6 and IPv4-IPv6 tunneling.
1.2	11/06/02	Updated via consensus at NP/CP Task Group face to face meeting
1.3	11/15/02	Updated via consensus at NP/CP Task Group conference call and face to face meeting
1.4	04/01/03	Updated with comments received via straw ballot
1.5	04/15/03	Deleted tunneling
1.6	05/12/03	Updated with comments from straw ballot.

2 Scope And Purpose

This document defines the methodology used to obtain network processor IP application level benchmarks. Specifically, this document describes the tests that may be used to obtain the IPv4 and IPv6 performance metrics, the test configurations that may be used and, the formats for reporting the results of the tests. This document is closely based on the IETF RFCs 1242 [1], 2544 [2], 2285[10], 2289[11] and NPF IPv4 Forwarding Application Level Benchmark Implementation Agreement [3].

Measurements of the tests described in this document can be used by customers to evaluate the performance of network processors versus their requirements and to compare the performance of different Network Processors.

3 Normative References

The following documents contain provisions, which through reference in this text constitute provisions of this specification. At the time of publication, the editions indicated were valid. All referenced documents are subject to revision, and parties to agreements based on this specification are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below.

1. *Benchmarking Terminology for Network Interconnect Devices*, IETF RFC 1242.
2. *Benchmarking Methodology for Network Interconnect Devices*, IETF RFC 2544.
3. *IPv4 Forwarding Application-Level Benchmark Implementation Agreement Revision 1.0*
4. *NPF IPv4 Unicast Forwarding API*, Network Processing Forum.
5. *Internet Routing Table Statistics*, http://www.merit.edu/ipma/routing_table/
6. *Internet IP Packet Size Distribution*, Network Processing Forum contribution npf2001.078.00
7. *Annex to IPv4 Forwarding Application Level Benchmark Specification*, Version 0.1, Network Processing Forum
8. *Transition Mechanisms for IPv6 Hosts and Routers*; IETF RFC 2893.
9. *Connection of IPv6 Domains via IPv4 Clouds*; IETF RFC 3056
10. *Benchmarking Terminology for LAN Switching Devices*; IETF RFC 2285
11. *Benchmarking Methodology for LAN Switching Devices*; IETF RFC 2289
12. *Internet Protocol, Version 6 (IPv6) Specification*; IETF RFC 2460
13. Route Table Study Group, Proposal for generating and using IP forwarding tables for NPF Benchmark Implementation Agreements, NPF2003.004
14. *Annex to IP Forwarding Application Level Benchmark Specification*, Version 0.1, Network Processing Forum
15. BGP Peering Routers. <http://bgp.potaroo.net>

4 Terminology

This section defines the terms used in the document. The goal is to define a specific terminology that every network processor vendor can use to measure and report the results of the benchmark tests described later on in this document. The benchmarking terminology is based on the IETF benchmarking terminology for network interconnection devices described in RFC1242 [1].

4.1 Benchmarking Terminology

Aggregate Forwarding Rate:

Definition: The maximum rate at which the received frames are forwarded by the IP forwarding functions.

Measurement units: Output frames per second at a frame size of N bytes.
Output bits per second.

Throughput:

Definition: The maximum rate at which none of the valid received frames are dropped by the IP forwarding functions.

Measurement units: Input frames per second at a frame size of N bytes.
Input bits per second.

Latency:

Definition for store and forward devices: The time interval starting when the last bit of the input frame reaches the input port of the DUT and ending when the first bit of the output frame is seen on the output port of the DUT.

Definition for bit-forwarding devices: The time interval starting when the end of the first bit of the input frame reaches the input port of the DUT and ending when the start of the first bit of the output frame is seen on the output port of the DUT.

Measurement units: Seconds.

Frame Loss Ratio:

Definition: Percentage of received frames that should have been forwarded by the IP forwarding function but were dropped instead.

Measurement units: Percentage of N byte input frames that are dropped.

Overload Forwarding Rate:

Definition: The maximum rate at which received frames are forwarded over a port, measured under the condition caused by a given sustained (constant) aggregate rate of frames destined to that port exceeding the theoretical line rate for that port.

Measurement units: Output frames per second at a frame size of N bytes.
Output bits per second.

Forwarding Table Update Latency:

Definition: The time interval starting when the request for the forwarding table update by [1..n] entries is issued and ending when a notification of that request being completed is received.

Measurement units: Seconds.

Forwarding Table Update Rate

Definition: The maximum rate at which forwarding table updates by [1..n] entries can be issued with the average forwarding table update latency below a threshold.

Measurement units: Route entries updated per second.

5 Test Configuration

The IP Forwarding Benchmark should be run on hardware. Simulation based numbers cannot be reported.

5.1 Reference Design

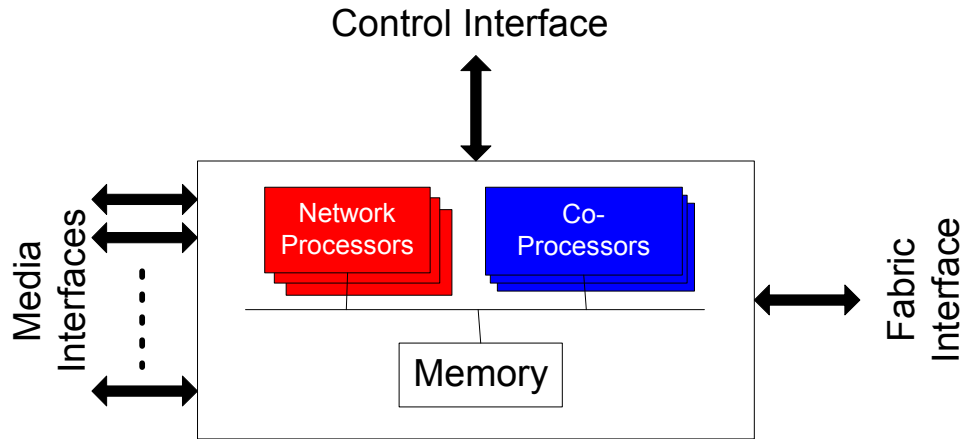


Figure 1: Reference Design for IP Function Level Benchmarks

The reference design [1] used for the IPv4 function level benchmarks is shown in Figure 1. The reference design is assumed to consist of one or more media interfaces and one control interface. This specification does not define the speed and type of the media and control interfaces. The choice of the media and control interfaces is left to network processor vendors or customers comparing different network processors. When reporting the results of the tests specified in this benchmark, it is required that the reference design be specified in terms of the following parameters in addition to the parameters listed in the Framework document [1]. This reference design is considered to be the device under test (DUT).

REQUIRED information

Block Diagram	Configuration block diagram of the reference design. This diagram should include details on <ul style="list-style-type: none"> • DUT to Traffic Tester connections • Number of PCBs used and mode of connectivity between PCBs
Component List	List of hardware components used on the reference design. Must include the NPs, CPs, memory chips, PHYs, framers fabric interface chips, and switch fabric chips.
Mechanical Size	Size of the base PCB and feature card/daughter card PCB of the reference design
Media Interface(s)	List of all media interface types on DUT For example:

	<p>10/100 Ethernet Gigabit Ethernet OC-3 POS OC-12 POS OC-48 POS OC-192 POS etc.</p>
Fabric Interface	<p>List of all fabric interface types on DUT For example: CSIX L1, Gigabit Ethernet, ATM, etc.</p>
Number of ports	<p>Total number of media interfaces supported in this reference design. Can be one of the above media interfaces, or a combination of a few of the above.</p>
Level of Channelization	<p>For example, a OC-48 Network Processor may be capable of processing only OC-48c; or capable of 4xOC-12, 16xOC-3, 48xSTS-1 channelization</p>
Network Processor details	<p>Type (part number) and number of NPs used, showing arrangement between NPs: paralleled, pipelined or compound.</p>
Coprocessor details	<p>Type (part number) and number of CPs used, showing arrangement between coprocessors: paralleled, pipelined or compound. For e.g., Type (part number) and number of TCAMs or other external search engines used .</p>
External Memory details	<p>External Memory interfaces – SRAM (what type) and/or DRAM (what type) Function of each External memory interface Amount of memory on the reference design for specific functions e.g., Total memory used for each of the IPv4 and IPv6. Total data memory used for packet buffers.</p>
Control Interface details	<p>PCI/Power PC etc Bandwidth of the interface</p>
Total power consumption	<p>Total power consumption of the reference design at idle condition. Should be accurate within $\pm 10\%$. Idle condition is defined as running the IP forwarding application with no traffic.</p>
Latency Definition used	<p>Latency definition used: bit-forwarding or store and forward. (per RFC 1242 Section 3.8).</p>
Forwarding Table and Traffic Details	<ul style="list-style-type: none"> • Number of forwarding table entries present on the DUT. <ul style="list-style-type: none"> ○ Across the system ○ Per DUT test port • Number of route table entries exercised by the traffic <ul style="list-style-type: none"> ○ Across the system ○ Per DUT test port

	<ul style="list-style-type: none"> • Description of the traffic flows from <ul style="list-style-type: none"> ○ The tester to the DUT ○ Between the ports of the DUT
--	--

DESIRABLE information

Schematic	Reference Design schematics in pdf format
Data Sheets	Data sheets for all the NPs and CPs used in the reference design.
Power consumption per NP/CP	Power consumption breakdown for each NP and CP used in the reference design under the following idle condition. Idle condition is defined as running the IP forwarding application with no traffic

5.2 Test Setup

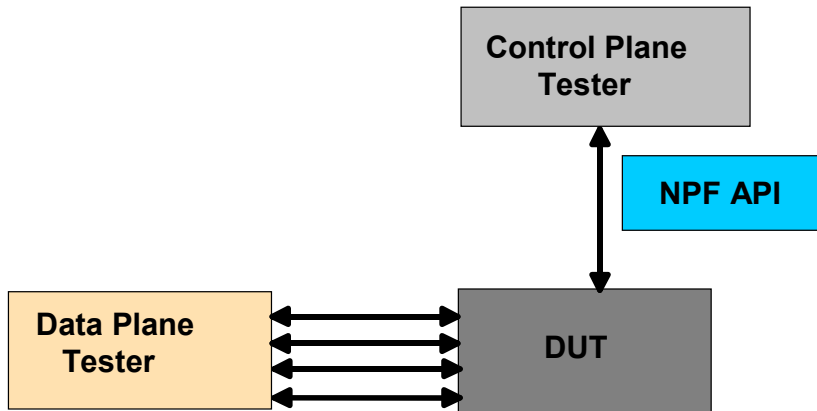


Figure 2: Test Setup for IP Function Level Benchmarks

The test setup used to obtain the results of the tests described in this document is shown in Figure 2. The media interfaces of the Device Under Test (DUT – same as the reference design shown in Figure 1) are connected to the media interfaces on a data plane tester. The data plane tester is a device capable of sending and receiving network traffic with configurable IP layer and link layer headers. The data plane tester should also be capable of calculating the number of packets transmitted and received, and should have the ability to measure the end-to-end latency through the DUT.

The control interface on the DUT is connected to a control plane tester. The control plane tester is a device capable of generating route update requests and other control messages required for IP forwarding. Route update requests may be generated either as API calls or by injecting control messages over the control interface of the DUT. The control messages or APIs may be based on the NPF IPv4 API [4], NPF IPv6 API or other equivalent APIs. At this time there are no commercially available control plane testers in the market. An external control plane processor that runs the NPF IPv4 API [4], NPF IPv6 API or equivalent software may be used as a control plane tester. Alternatively, data plane testers can also be used as control plane testers to inject route update requests and other control messages over the control interface of the DUT.

6 Benchmark Tests

This section describes the benchmark tests used to measure the metrics specified in Section 4.1. It also describes the format used for reporting the results of the tests. Each test has a set of “MUST” report results and may have a set of “SHOULD” report results. MUST report results means that reporting these results is an absolute requirement of the specification. SHOULD report results means that reporting these results is recommended but optional.

6.1 Data Plane Tests

The data plane tests will be carried out using different traffic combinations.

Traffic Combination	Native IPv4	Native IPv6
IPv4 only	100%	0
IPv4 Concurrent with IPv6	50%	50%
IPv6 Only	0	100%

Table 1:IP Traffic Configurations

All benchmark tests SHOULD be performed on all traffic combinations supported by the DUT. All supported IP forwarding tables (IPv4 and IPv6), ARP tables and Next Neighbor tables MUST be initialized prior to all tests even if that table is not being exercised by traffic. The IPv6 packets will not contain hop-by-hop options except when explicitly called for.

The packet sizes to be used for each packet type in the benchmark are shown below.

Traffic Type	Packet Sizes to be exercised																				
Native IPv4	<ul style="list-style-type: none"> • Ethernet: 64, 128, 256, 512, 594, 1024, 1280, 1518. • AAL5 ATM: 48, 64, 128, 256, 512, 584, 1024, 1518, 2048, 4472, 9180 • POS: 48, 64, 128, 256, 512, 584, 1024, 1518, 2048, 4472, 9180 • Internet Mix: This data is based on real world analysis of Internet IP packet size distributions [6]. <table border="1" data-bbox="678 680 1300 835" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>IP</th> <th>Ethernet</th> <th>ATM</th> <th>POS</th> <th>Probability</th> </tr> </thead> <tbody> <tr> <td>40</td> <td>64</td> <td>48</td> <td>48</td> <td>0.56</td> </tr> <tr> <td>576</td> <td>594</td> <td>584</td> <td>584</td> <td>0.20</td> </tr> <tr> <td>1500</td> <td>1518</td> <td>1508</td> <td>1508</td> <td>0.24</td> </tr> </tbody> </table>	IP	Ethernet	ATM	POS	Probability	40	64	48	48	0.56	576	594	584	584	0.20	1500	1518	1508	1508	0.24
IP	Ethernet	ATM	POS	Probability																	
40	64	48	48	0.56																	
576	594	584	584	0.20																	
1500	1518	1508	1508	0.24																	
Native IPv6	<ul style="list-style-type: none"> • Ethernet: 78, 128, 256, 512, 594, 1024, 1280, 1518. • AAL5 ATM: 68, 128, 256, 512, 584, 1024, 1518, 2048, 4472, 9180 • POS: 68, 128, 256, 512, 584, 1024, 1518, 2048, 4472, 9180 • Internet Mix: This data extrapolated is based on real world analysis of Internet IP packet size distributions [6]. <table border="1" data-bbox="678 1136 1300 1287" style="margin-left: auto; margin-right: auto;"> <thead> <tr> <th>IP</th> <th>Ethernet</th> <th>ATM</th> <th>POS</th> <th>Probability</th> </tr> </thead> <tbody> <tr> <td>60</td> <td>78</td> <td>68</td> <td>68</td> <td>0.56</td> </tr> <tr> <td>596</td> <td>614</td> <td>584</td> <td>584</td> <td>0.20</td> </tr> <tr> <td>1500</td> <td>1518</td> <td>1508</td> <td>1508</td> <td>0.24</td> </tr> </tbody> </table>	IP	Ethernet	ATM	POS	Probability	60	78	68	68	0.56	596	614	584	584	0.20	1500	1518	1508	1508	0.24
IP	Ethernet	ATM	POS	Probability																	
60	78	68	68	0.56																	
596	614	584	584	0.20																	
1500	1518	1508	1508	0.24																	

Table 2: Packet Size Mapping

6.1.1 Aggregate Forwarding Rate

Objective: To determine the DUT's maximum aggregate forwarding rate for the parameters described in Table 2.

Procedure: Send a specific number of frames at line rate on all of the media interfaces of the DUT and then count the number of frames output by the DUT. The maximum forwarding rate is the rate at which frames are received at the data plane tester measured over the time interval required to receive the specific number of frames. Each run of this test should last 120 seconds or more to allow a large number of frames to be received at the tester.

Test setup: All supported IP forwarding tables (IPv4 and IPv6), ARP tables and Next Neighbor tables **MUST** be initialized prior to the test even if that forwarding table is not being exercised by traffic. The test **MUST** be repeated for each of traffic patterns supported by the DUT and the forwarding tables specified below. IP forwarding tables should be initialized with the entries in the forwarding table under test. The next hops for the forwarding table entries should be uniformly distributed over the number of media interfaces.

Forwarding tables: The route tables used is an independent variable for this test.

- IPv4
 - Unmodified Telestra AS1221 snapshot as described in RTSG recommendation [15]
 - A smaller 10,000 entry subset of Telestra AS1221 extracted using the scheme described in section 3.4.1 of the RTSG recommendation [15]
 - And one or more of the larger 250,000; 500,000 and a 1 million entry synthetic route table generated from AS1221 using the scheme described in section 3.4.2 of the RTSG recommendation [15]
 - Optionally a vendor may choose to run the tests with a table size **in addition** to one of the sizes specified above.
- IPv6
 - Unmodified AS4554 snapshot as described in RTSG recommendation [15]
 - A larger 1,200 entry synthetic route table generated from AS4554 using the scheme described in RTSG recommendation, Section 4.4.1

Frame sizes: The frame sizes to be used depend on the traffic configuration used in the test. The traffic configuration to frame size mapping is described in Table 2. If the DUT supports media interfaces of more than one link-layer type with different MTUs, the frame sizes for the larger MTU must be used for the test. If the DUT does not support IP

fragmentation in the presence of MTU mismatch, the forwarding rate for that frame size should be reported as zero.

Traffic: All traffic patterns **MUST** be used for this test unless specified as optional. Each traffic pattern should be reported as a separate run of the test. The data plane tester should generate traffic evenly distributed on all the media ports with IP destination addresses drawn randomly from the forwarding table under test to assure uniformly distributed traffic over the number of output media interfaces. The algorithm used to generate traffic is described in detail in RTSG recommendation [15], Section 5.

- **Base case:** All frames should consist of datagrams with no IP options.
- **Control:** 5% of the frames should be destined to the control processor.
- **IP options:** Base case with 0.1% of the frames having options in the IP header. The IPv4 option to be used is the RECORD ROUTE option. The data plane tester should include one IP address in the option field in the traffic sent to the DUT. On receiving the traffic from the DUT, the data plane tester should verify that packets with the RECORD ROUTE option have a total of two IP addresses in the option field. In case of IPv6 0.1% of the frames should have the Hop-by-Hop options header. The Hop-by-Hop options should be used to record route as in the case of IPv6.

Optional

- **Bursty Case:** Base case but with bursts of length 10 for each destination IP.

Reporting format: The results of this test **MUST** be reported in the form of a graph. The x coordinate **MUST** be the frame size and the y coordinate **MUST** be the measured forwarding rate in frames per second or bits per second. There **MUST** be at least two lines on the graph. The first line **MUST** show the theoretical forwarding rate calculated at various frame sizes. The second line **MUST** show the measured forwarding rate obtained from the test. The results for the Internet Mix traffic case **MUST** be shown as an explicitly stated number on the graph. The results **MUST** report the aggregate forwarding rate measured over all the media interfaces and **SHOULD** report the forwarding rates over each individual media interface for each of the traffic configurations used. There should be separate graphs for each of the traffic configurations used. (Note: The theoretical upper limit for this test is the sum of all output interface line rates.). The size of the route table used and the seed used to generate and distribute traffic **MUST** be documented.

6.1.2 Throughput

Objective: To determine the DUT throughput for all inputs together (aggregate throughput).

Procedure: Send a specific number of frames at a specific rate on each of the media interfaces of the DUT and then count the number of frames output by the DUT. If the count of received frames is equal to the count of output frames, then the throughput is the

rate at which frames are received at the data plane tester measured over the time interval required to transmit the specific number of frames. If fewer frames are received than were transmitted, the rate of the frames generated by the data plane tester is reduced and the test is rerun. Each run of the test should last 120 seconds or more to allow a large number of frames to be received at the tester.

Reporting format: Same as in Section 6.1.1.

Test setup: Same as in Section 6.1.1.

Frame sizes: Same as in Section 6.1.1.

Traffic: The following traffic patterns must be used for this test. Each traffic pattern should be reported as a separate run of the test. For the aggregate throughput test, the data plane tester should generate traffic evenly distributed on all the media ports with IP destination addresses drawn such that traffic is uniformly distributed over the number of output media interfaces.

- Base case: All frames should consist of datagrams with no IP options.

6.1.3 Latency

Objective: To determine the latency of the DUT.

Procedure: Send a stream of frames at a particular frame size through the DUT at a specific rate. The stream should be at least 120 seconds in duration. The time at which the frame is fully transmitted is recorded inside the frame (timestamp A). The receiver logic in the data plane tester records the time at which the tagged frame was received (timestamp B). The latency is timestamp B minus timestamp A for each and every frame transmitted.

Reporting format: The latency results **MUST** be reported in the format of a graph. The x coordinate **MUST** be the frame size and the y coordinate **MUST** be the measured latency in seconds. Each graph **MUST** have three lines: one for the average latency, one for the minimum latency and one for the maximum latency at each frame size and protocol type. The latencies for the Internet Mix traffic case **MUST** be reported numerically on the graph. A separate graph **MUST** be drawn for each fraction of the throughput rate given below. The results **MUST** report both the latencies measured in an aggregate fashion over all the media interfaces and **SHOULD** report latencies measured over each individual media interface. There should be separate graphs for each of the traffic patterns used.

Test setup: Same as in Section 6.1.1.

Frame sizes: Same as in Section 6.1.1.

Traffic: Same as in Section 6.1.1. The latency test must be repeated for frame rates of 50, 90, 95 and 100% of the throughput rate determined in Section 5.1.2.

6.1.4 Loss Ratio

Objective: To determine the frame loss rate of the DUT over the entire range of input data rates and frame sizes.

Procedure: Send a specific number of frames at the specific rate through the DUT and count the number of frames received by the data plane tester from the DUT. Each run of this test should last 120 seconds or more to allow a large number of frames to be received at the tester. The frame loss rate is calculated using the following equation:

$$((\text{input_count} - \text{output_count}) * 100) / \text{input_count}$$

The first trial should be run by transmitting frames at line rate on all the media interfaces. Subsequent trials should be run by reducing the frame rate by 10% increments of the line rate until there are two successive trials in which no frames are lost.

Reporting format: The results of the frame loss rate **MUST** be reported in the form of a graph. The x axis must be the input frame rate as a percent of the line rate at the specific frame size. The y axis must be the percent loss at the particular input rate. The left end of the x axis and the bottom of the y axis must be 0 percent; the right end of the x axis and the top of the y axis must be 100 percent. Multiple lines on the graph **MUST** be used to report the frame loss rate for different frame sizes as well as the Internet Mix traffic case. The loss rate for the Loss results **MUST** be reported for both the aggregate loss rate measured over all media interfaces and **SHOULD** be reported for loss rates measured over each individual media interface for all traffic configurations. There should be separate graphs for each of the traffic configuration used.

Test setup: Same as in Section 6.1.1.

Frame sizes: Same as in Section 6.1.1.

Traffic: Same as in Section 6.1.1. Base case only

6.1.5 Overload Forwarding Rate

Objective: To determine the forwarding rate of the DUT under overload conditions. This may help determine how costly the discard of overflow packets is at the egress side.

Procedure: The fastest media port on the system, or any one of the media ports in case of a system where all media ports are of the same speed is designated as the output interface. All other media ports on the system are designated as the input ports. Send a specific number of frames at line rate on all of the input media ports to the single output port of the DUT and then count the number of frames received. Each run of this test should last 120 seconds or more to allow a large number of frames to be received at the

tester. The overload-forwarding rate is the rate at which frames are received at the data plane tester measured over the time interval required to transmit the specific number of frames.

Reporting format: Same as in Section 6.1.1.

Test setup: The IP forwarding tables and ARP/Next Neighbor tables should be initialized prior to the test. The IP forwarding tables should be initialized with route entries taken from the forwarding table under test. The next hops for all the forwarding table entries should be set to the same output media port on the DUT. The media port selected should be the port with the highest theoretical line rate. The ARP table should be initialized with the destination MAC address of the data plane tester port connected to the DUT output port selected.

Frame sizes: Same as in Section 6.1.1.

Traffic: Same as in Section 6.1.1. Base case only.

6.1.6 System Power Consumption

Procedure: Same as in Section 6.1.3. The average system power consumption **MUST** be measured for 50% and 100% of the throughput rate determined in Section 6.1.2. It is desirable to measure power consumption at 50% and 100% of the throughput rate for the major components in the system, including NPs, CPs and Network Search Engines (NSEs).

Reporting Format: Average system power consumption **MUST** be reported at 50% and 100% of the throughput rate determined in Section 6.1.2 for minimum packet size. It is desirable to measure and report power consumption at the minimum packet size for the major components in the system, including NPs, CPs and NSEs, at 50% and 100% of the throughput rate determined in Section 6.1.2.

Test Setup: Same as in Section 6.1.3.

Frame sizes: Only the minimum IP packet size for the link-layer **MUST** be sent into the DUT at the prescribed throughput rates.

Traffic: Same as in Section 6.1.3.

6.2 Control Plane Tests

6.2.1 Forwarding table Update Rate

Objective: To determine the maximum rate at which forwarding table entries can be added, deleted and updated subject to a maximum update latency of one millisecond.

Procedure: Using a control plane or data plane tester, issue a specific number of forwarding table updates (add, delete or update) at a specific rate and measure the average latency for the update requests. Each forwarding table update can update one or more route entries. The number of route entries updated in each call is an independent variable for this test. The forwarding table update API call may be blocking or non-blocking. For blocking update requests, the latency is defined as the time interval between the time the update API call is made and the time the call returns. For non-blocking requests, the latency is defined as the time interval between the time the update API call is made and the time the callback function is called after processing the update. If the average latency is less than the 1ms specified threshold, then the update rate is increased and the test is rerun. If the average latency is more than the specified threshold, then the update rate is decreased and the test rerun. The rate at which the average latency is equal to the threshold is the forwarding table update rate. Each run of this test should last 120 seconds or more to allow a large number of frames to be received at the tester.

Reporting format: The results of this test **MUST** be specified in the form of a graph. The x axis must be the number of route entries in each update call. The y axis must be the forwarding table update rate in entries/call. There **SHOULD** be three lines on the graph showing the update rates for the route add, route delete and route update operations. There should be a separate graph for each type of IP forwarding table.

Test setup: The route addresses used in the forwarding table update calls should cycle through the entire forwarding table(s) under test. Route update ports should be selected such that the traffic distribution on the output ports remains balanced. Forwarding table(s) under test are all those that are actively exercised.

6.3 Concurrent Data Plane and Control Plane Tests

6.3.1 Route Update and Forwarding Rate

Objective: To determine the effect of concurrent forwarding table updates on forwarding rate.

Procedure: Using a control plane or data plane tester, issue forwarding table updates at 25%, 50%, 75% and 100% of the rate determined in Section 6.2.1. Concurrently, send a specific number of frames at line rate on each of the media interfaces of the DUT. The forwarding rate is the rate at which frames are received at the data plane tester measured over the time interval required to transmit the specific number of frames. Each run of this test should last 120 seconds or more to allow a large number of frames to be received at the tester.

Reporting format: Same as in Section 6.1.1

Test setup: The IP forwarding tables and ARP/Next Neighbor tables should be initialized prior to the test. The IP forwarding tables should be initialized with route entries taken from the forwarding table under test. The next hops for the forwarding table entries should be uniformly distributed over the number of media interfaces. The forwarding table updates should change the next hops for the updated route table entries.