



IXP2800 Intel Network Processor IP Forwarding Benchmark Full Disclosure Report for OC192-POS

Oct 30, 2003

Editor(s):
Contact:

David Meng, Ravi Gunturi & Manohar Castelino
Eric Dudley
Intel Corporation
CHP3-108
2200 Mission College Blvd
Santa Clara, CA 95052 USA
tel: +1 408 545 6210
e-mail: eric.dudley@intel.com

ABSTRACT:

This document describes the IP Forwarding Application-Level Benchmark Implementation Agreement (Revision 1.0) results obtained using 2 Intel IXP2800 Network Processors in the data path for IPv4 and IPv6.

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.

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

1	Revision History	4
2	Executive Summary	5
3	Reference Design Details.....	7
	3.1 Component List.....	7
	3.2 Mechanical dimensions.....	7
	3.3 Media, Fabric and Control interfaces.....	7
	3.4 Traffic Matrix.....	7
	3.5 Total Power Consumption	8
	3.6 IP Control Memory	8
	3.7 Test Equipment	8
	3.8 Forwarding Table and Traffic Details.....	8
	3.9 NPU Software	8
4	Data Plane Benchmark Tests	10
	4.1 Forwarding Rate.....	10
	4.2 Throughput.....	14
	4.3 Latency.....	18
	4.4 Loss Rate.....	26
	4.5 Overload Forwarding Rate.....	28
	4.6 System Power Consumption	29
5	Control Plane Benchmark Tests.....	30
	5.1 Forwarding Table Update Rate.....	30
	5.2 Forwarding Rate With Concurrent Forwarding Table Update Rate	31

1 Revision History

Revision	Date	Reason for Changes
1.0	10/30/03	Created Rev 1.0 of the OC-192 POS IP Benchmarking Report on the IXP2800


2 Executive Summary

Intel achieved a forwarding rate of 20 Gbps when forwarding IPv4 and IPv6 traffic using the IP Forwarding Application-Level Benchmark Implementation Agreement (Revision 1.0) employing 2 IXP2800 NPUs.


Summary of results is as follows

- Forwarding Rate of 20 Gbps for all packet sizes
- Throughput of 20 Gbps for all packet sizes
- Zero loss for all packet sizes
- Maximum Route update rate of 58,883 updated/sec for IPv4 and 24,000 updates/sec for IPv6

The IPv4 route table used had 135,432 route table entries with the prefix length distribution shown below



The IPv6 route table used had 1,200 route table entries with the prefix length distribution shown below



3 Reference Design Details

The IP forwarding benchmark was run on a dual IXP2800 reference design with 2 OC-192 POS media interfaces. The details of the reference design are given below.

3.1 Component List

- IXDP2800 Platform with
 - 2 IXP2800 A2 1.4 GHz Network Processors with the XScale running at 700 Mhz
 - 3 Channels of 256 MB 127.3 MHz RDRAM per Network Processor
 - 4 Channels of 8 MB 200 MHz QDR II SRAM per Network Processor
- 2 IXF18101 media cards each with an OC-192 POS Port

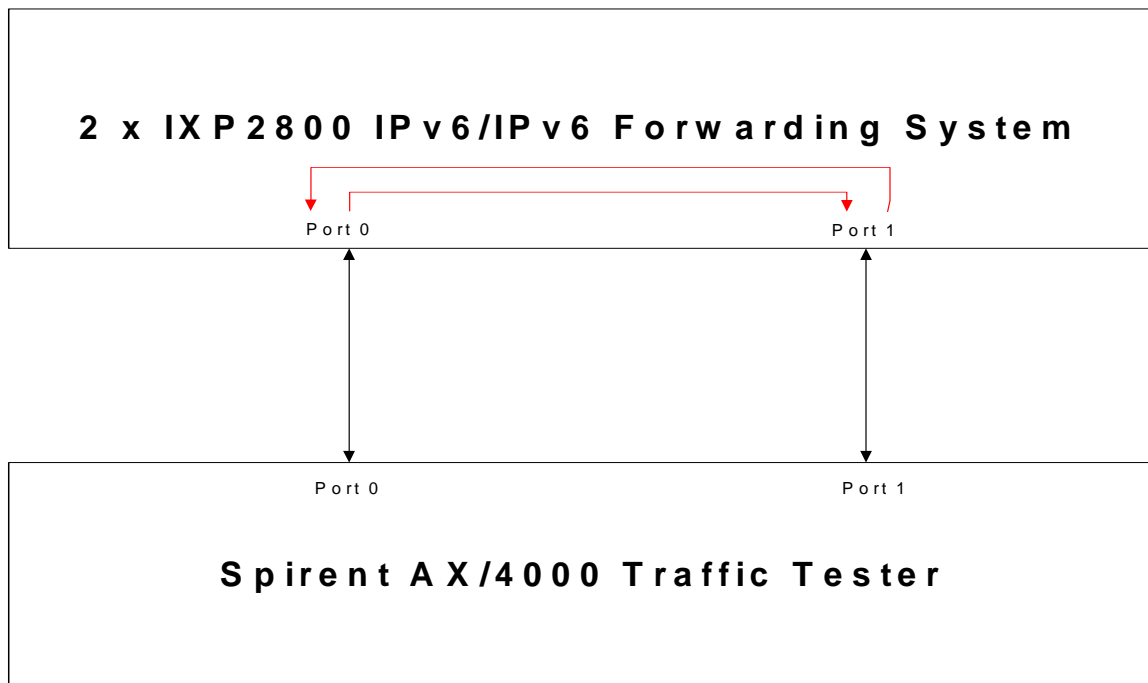
3.2 Mechanical dimensions

- Height 3.5 in.; Width 17.5 in.; Depth 22.5 in.

3.3 Media, Fabric and Control interfaces

- Media interfaces: 2 x OC-192 POS ports
- Control interfaces: On-chip XScale control processor communicates with the data plane via shared memory.
- Fabric interface: None

3.4 Traffic Matrix



3.5 Total Power Consumption

- Each IXP2800 A2 @ 1400 MHz consumes 25.5 Watts
 - This does not include the power consumed by the I/O interfaces on the IXP2800

3.6 IP Control Memory

- IPv4 and IPv6
 - 7.3 MB of SRAM for IPv4 and IPv6 routing lookup table (FIB), next hop forwarding tables and packet meta data per IXP2800
- 123.9 MB DRAM per IXP2800 for store and forward packet data buffers.

3.7 Test Equipment

- Data plane tester: Spirent AX/4000 with 2 OC-192 POS Ports.
- Control plane tester: Control plane application running on Intel XScale Core.

3.8 Forwarding Table and Traffic Details

- Total Number of forwarding table entries present on the DUT
 - IPv4
 - Across the whole system : 135,432
 - Per DUT Test Port : 67,176
 - IPv6
 - Across the whole system : 1,200
 - Per DUT Test Port : 600
- Number of route table entries exercised by traffic for all packet sizes are as follows
 - IPv4
 - Across the whole system : 2,000
 - Per DUT Test Port : 1,000
 - IPv6
 - Across the whole system : 400
 - Per DUT Test Port : 200
- Traffic Flow description
 - The Traffic Generator to DUT flows are setup as follows
 - Each DUT Port(i) is connected to Traffic Generator Port(i)
 - Each DUT and Traffic Generator port is setup in full duplex mode, and hence and receive and transmit packets at the same time
 - The Traffic and route tables are setup as follows
 - All Traffic send in on DUT Port 0 (by Traffic Generator Port 0) is routed to DUT Port 1 (to Traffic Generator Port 1) and vice versa

3.9 NPU Software

- IPv4, IPv6, Ethernet Rx/Tx production microblocks from IXA SDK 3.1

- Implemented RFCs: RFC1812, RFC2544, RFC2460 (MUST checks) and RFC2373 (SHOULD checks)

Note:

1. Latency is measured first bit in to last bit out

4 Data Plane Benchmark Tests

4.1 Forwarding Rate

4.1.1 IPv4 Forwarding Rate

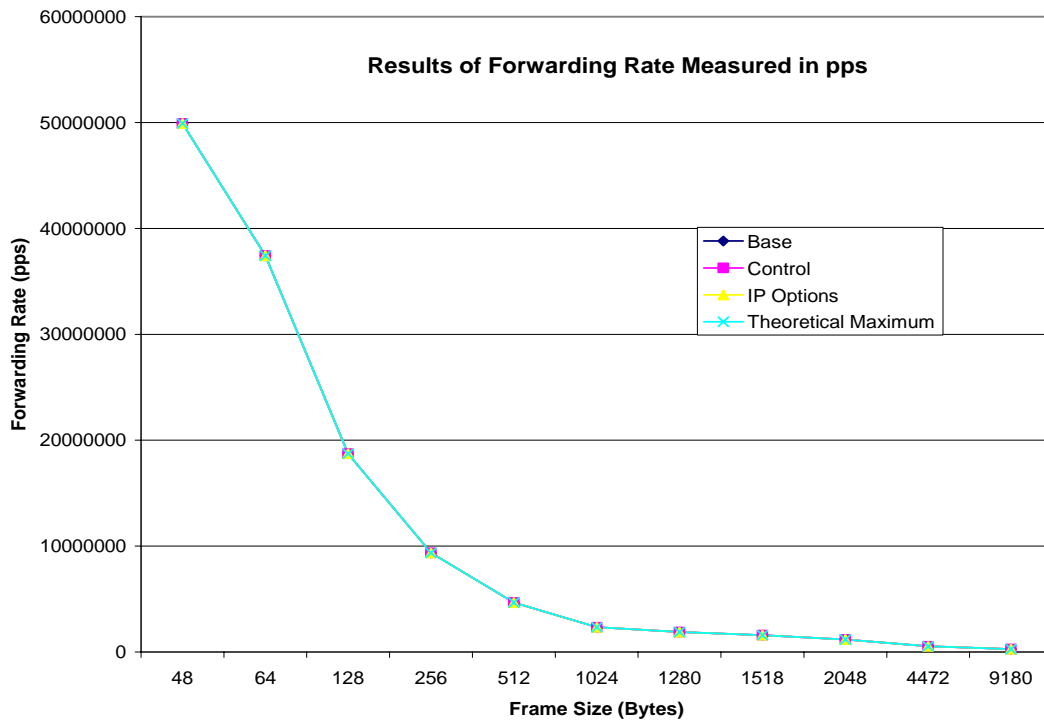


Figure 1: Results of IPv4 Forwarding Rate Benchmark Measured in pps

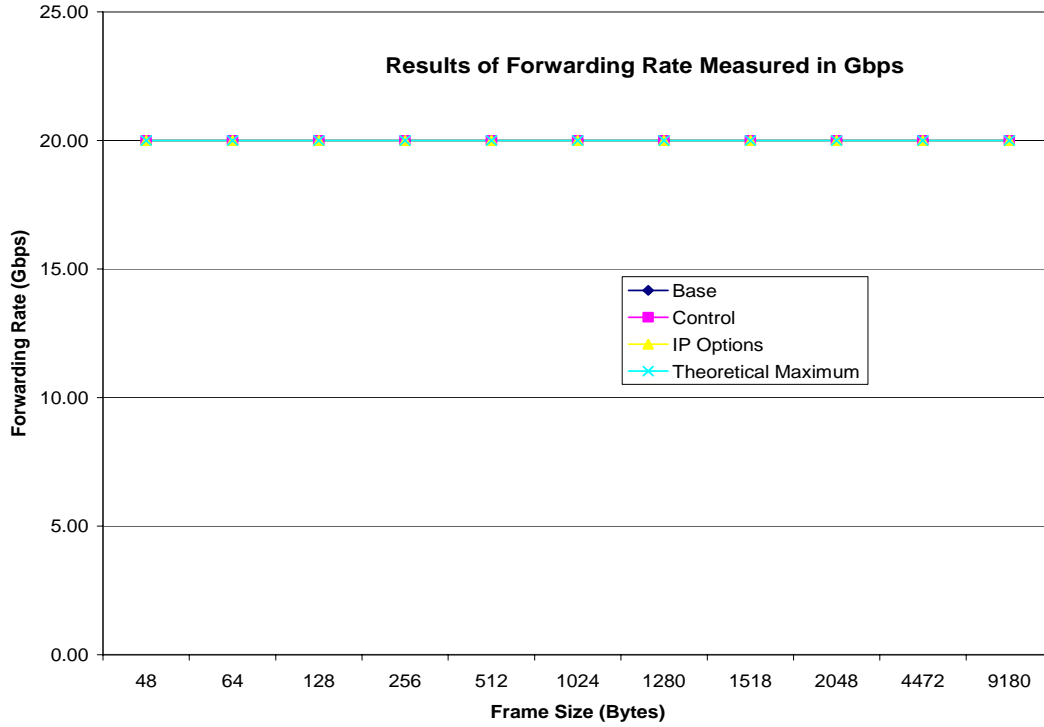


Figure 2: Results of IPv4 Forwarding Rate Benchmark Measured in Gbps

IPv4 Forwarding Rate with Internet Mix Traffic = 20 Gbps

4.1.2 IPv6 Forwarding Rate

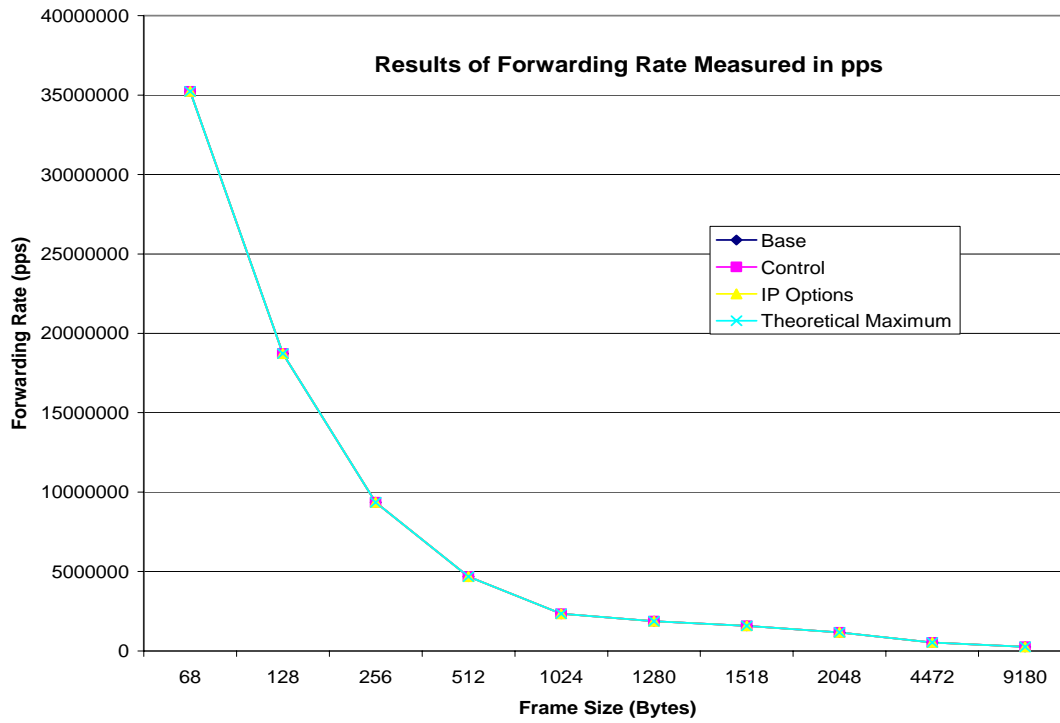


Figure 3: Results of IPv6 Forwarding Rate Benchmark Measured in pps

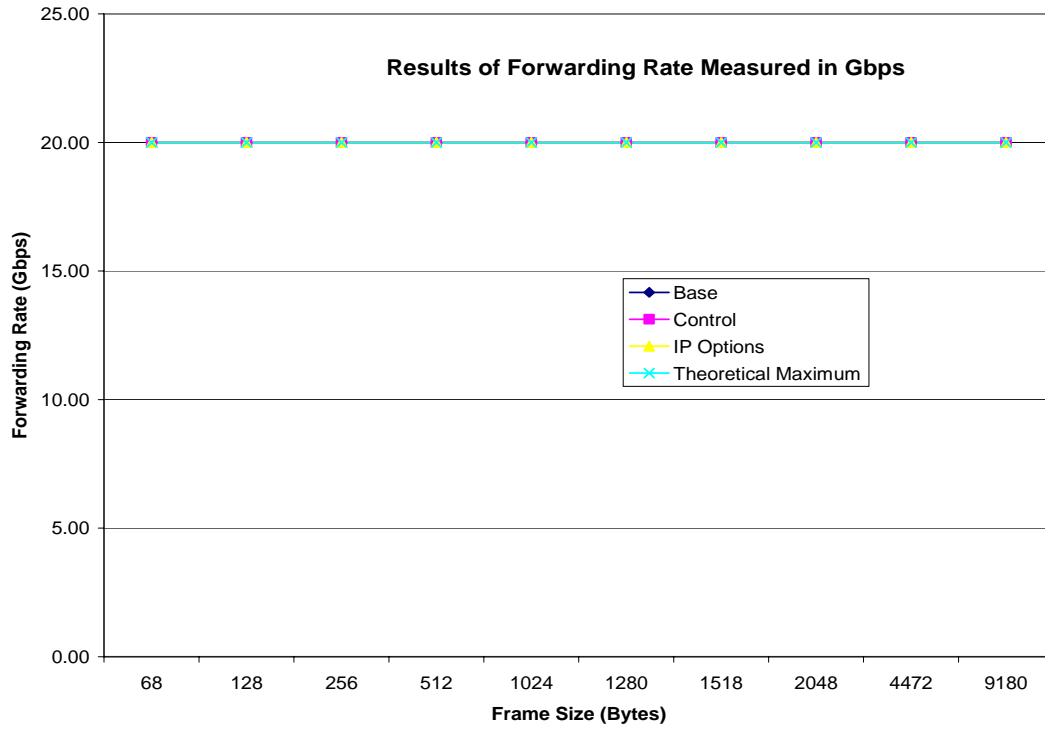


Figure 4: Results of IPv6 Forwarding Rate Benchmark Measured in Gbps

IPv6 Forwarding Rate with Internet Mix Traffic = 20 Gbps

4.2 Throughput

4.2.1 IPv4 Throughput

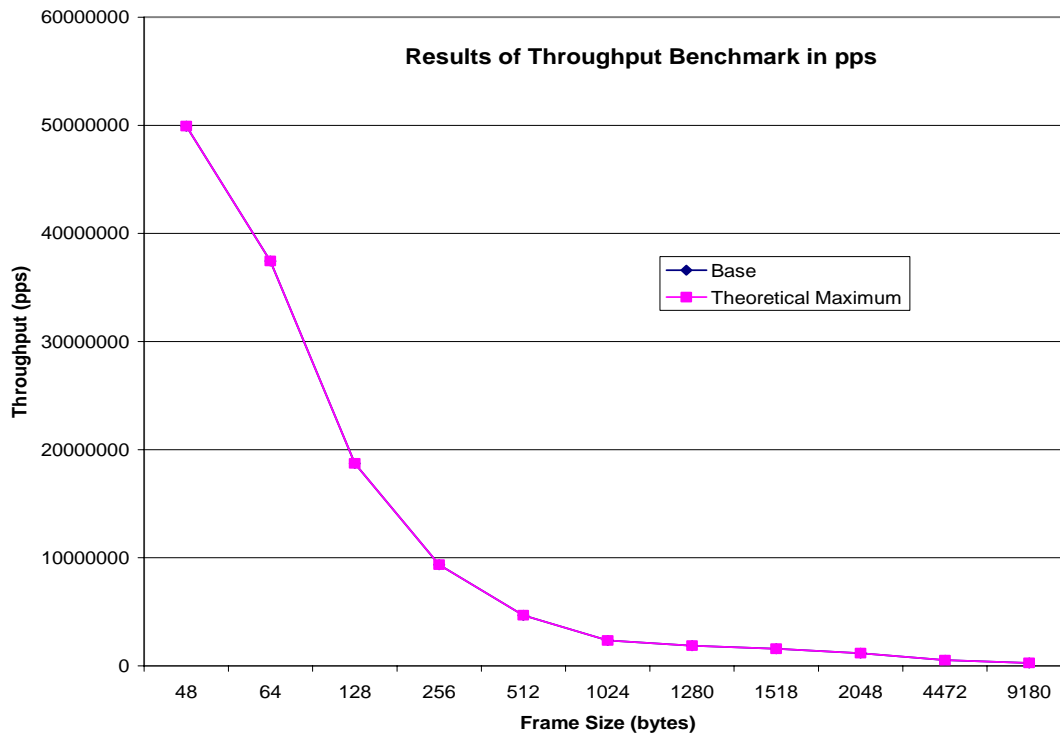


Figure 5: Results of IPv4 Throughput Benchmark Measured in pps




Figure 6: Results of IPv4 Throughput Benchmark Measured in Gbps

IPv4 Throughput with Internet Mix Traffic = 20 Gbps

4.2.2 IPv6 Throughput

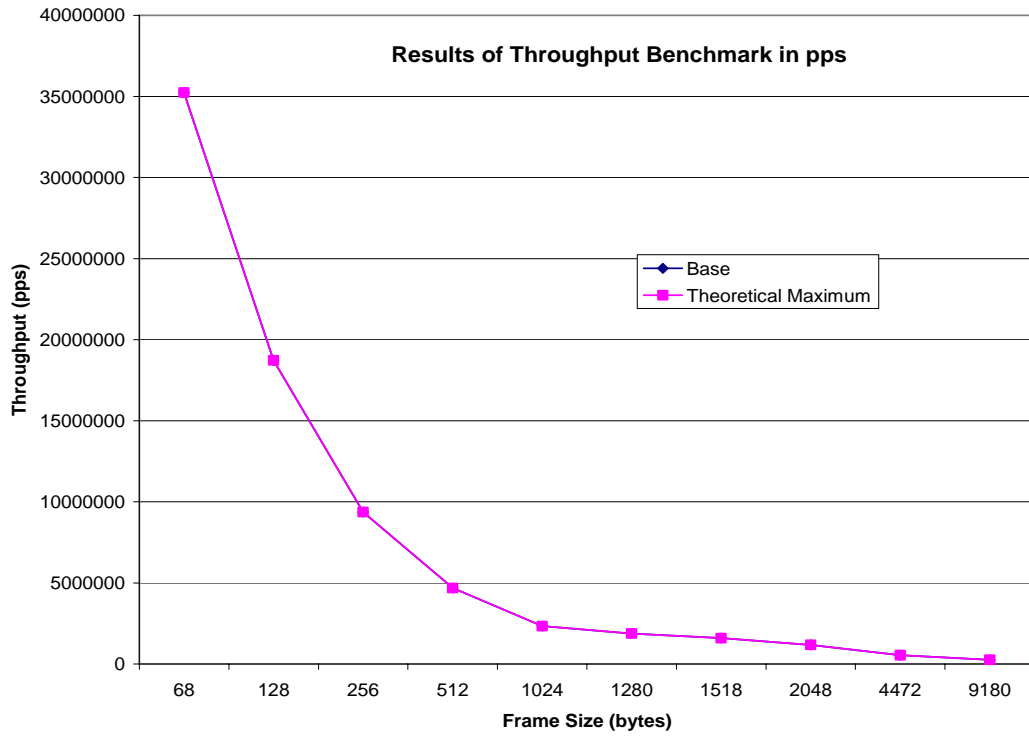


Figure 7: Results of IPv6 Throughput Benchmark Measured in pps

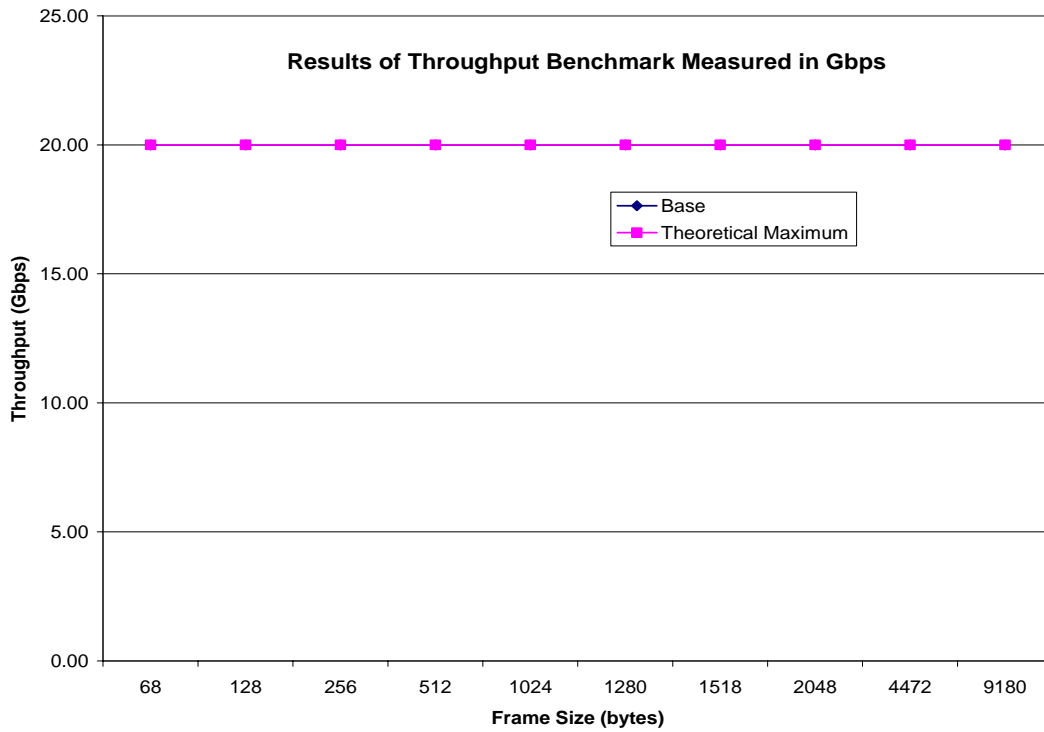


Figure 8: Results of IPv6 Throughput Benchmark Measured in Gbps

IPv6 Throughput with Internet Mix Traffic = 20 Gbps

4.3 Latency

4.3.1 IPv4 Latency




Figure 9: Results of IPv4 Latency at 100% Throughput

IPv4 Latency at 100% Throughput with Internet Mix Traffic Average = 10.80 μ s

IPv4 Latency at 100% Throughput with Internet Mix Traffic Minimum = 8.85 μ s

IPv4 Latency at 100% Throughput with Internet Mix Traffic Maximum = 13.85 μ s

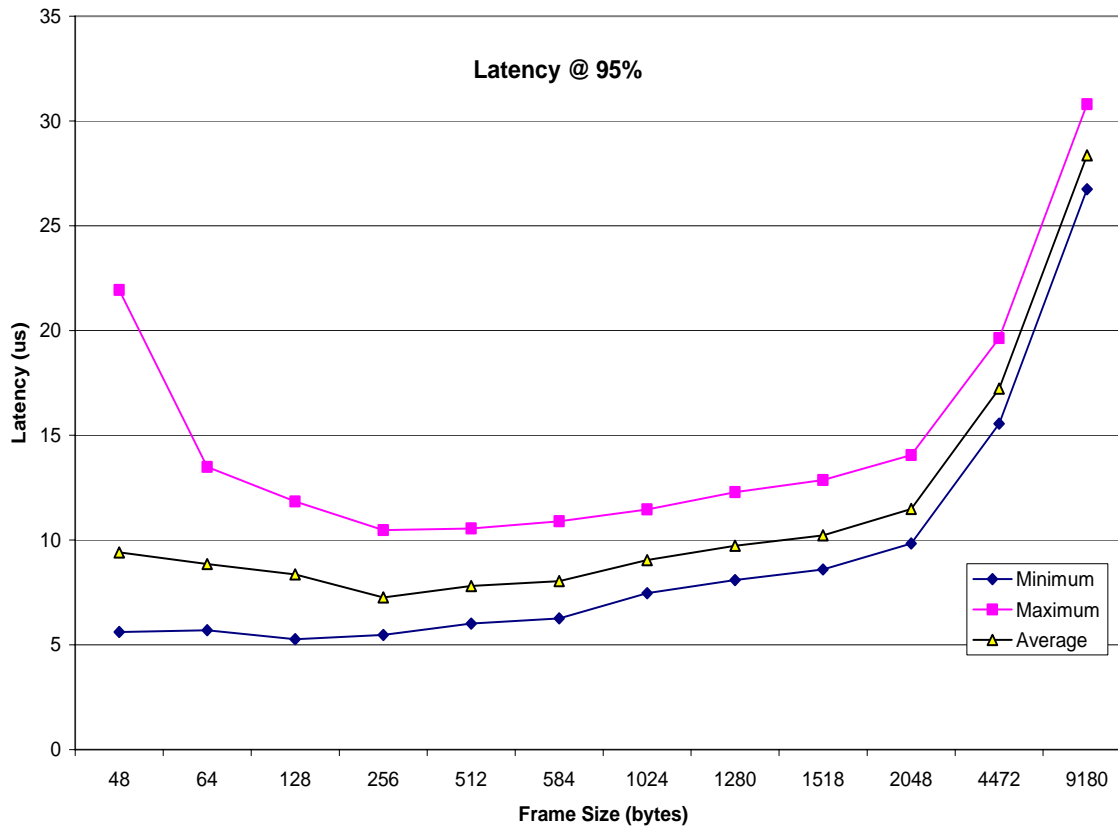


Figure 10: Results of IPv4 Latency at 95% Throughput

IPv4 Latency at 95% Throughput with Internet Mix Traffic Average = 10.80 μs
 IPv4 Latency at 95% Throughput with Internet Mix Traffic Minimum = 6.09 μs
 IPv4 Latency at 95% Throughput with Internet Mix Traffic Maximum = 13.74 μs




Figure 11: Results of IPv4 Latency at 90% Throughput

IPv4 Latency at 90% Throughput with Internet Mix Traffic Average = 10.79 μs
 IPv4 Latency at 90% Throughput with Internet Mix Traffic Minimum = 6.98 μs
 IPv4 Latency at 90% Throughput with Internet Mix Traffic Maximum = 13.92 μs




Figure 12: Results of IPv4 Latency at 50% Throughput

IPv4 Latency at 25% Throughput with Internet Mix Traffic Average = 10.80 μ s
 IPv4 Latency at 25% Throughput with Internet Mix Traffic Minimum = 7.8 μ s
 IPv4 Latency at 25% Throughput with Internet Mix Traffic Maximum = 14.18 μ s

4.3.2 IPv6 Latency

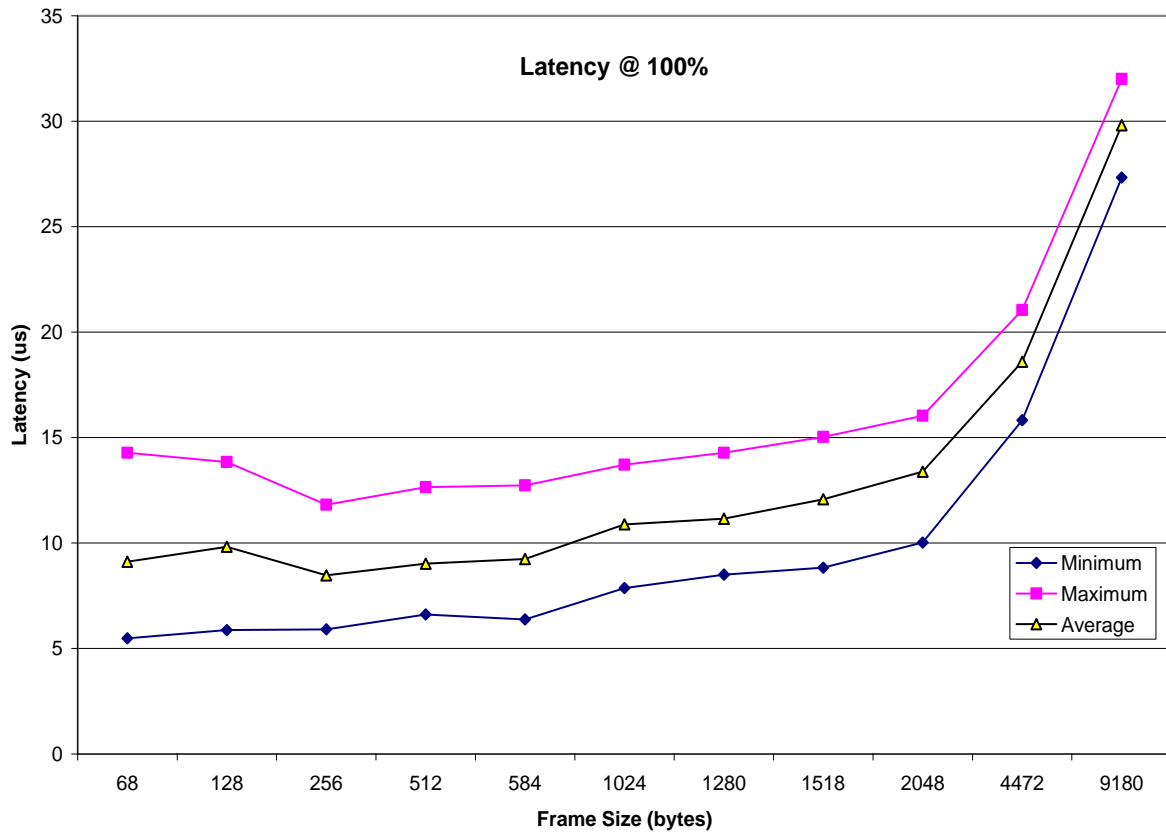


Figure 13: Results of IPv6 Latency at 100% Throughput

IPv6 Latency at 100% Throughput with Internet Mix Traffic Average = 12.15 μ s
IPv6 Latency at 100% Throughput with Internet Mix Traffic Minimum = 9.22 μ s
IPv6 Latency at 100% Throughput with Internet Mix Traffic Maximum = 16.18 μ s




Figure 14: Results of IPv6 Latency at 95% Throughput

IPv6 Latency at 95% Throughput with Internet Mix Traffic Average = 12.15 μ s
 IPv6 Latency at 95% Throughput with Internet Mix Traffic Minimum = 5.19 μ s
 IPv6 Latency at 95% Throughput with Internet Mix Traffic Maximum = 15.88 μ s

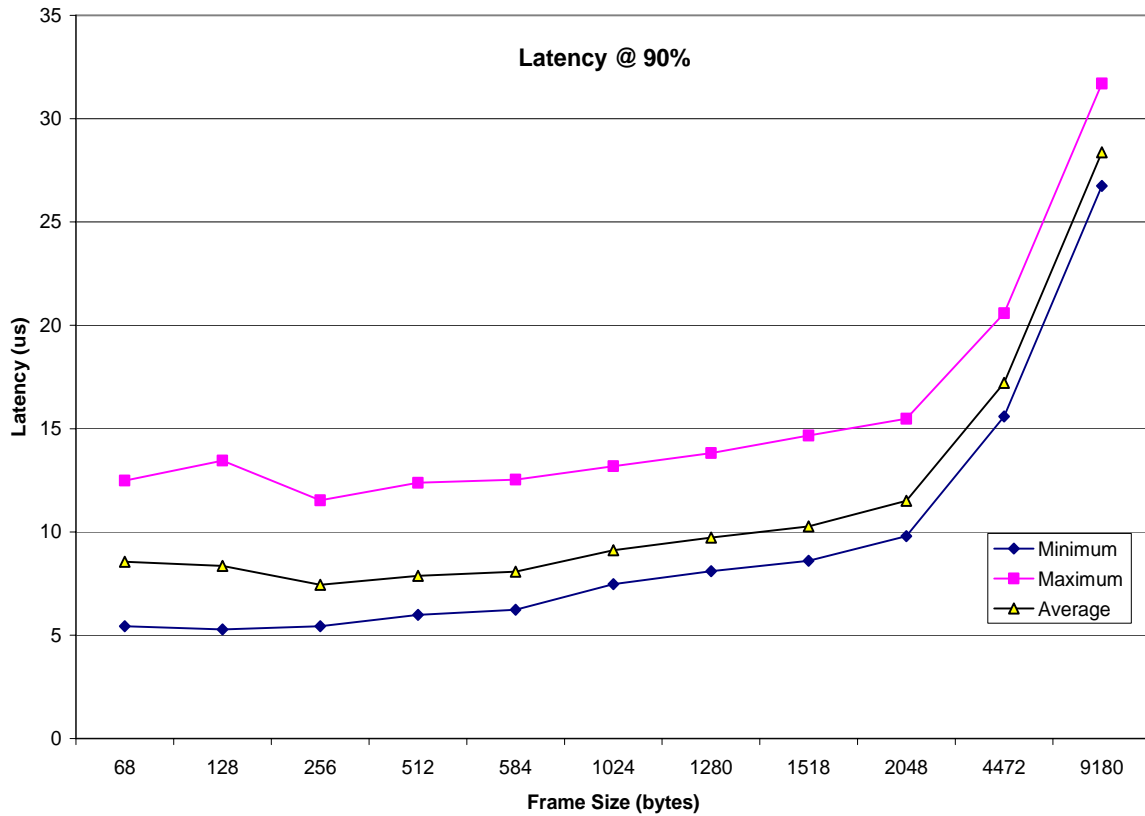


Figure 15: Results of IPv6 Latency at 90% Throughput

IPv6 Latency at 90% Throughput with Internet Mix Traffic Average = 12.15 μ s
 IPv6 Latency at 90% Throughput with Internet Mix Traffic Minimum = 7.89 μ s
 IPv6 Latency at 90% Throughput with Internet Mix Traffic Maximum = 16.12 μ s

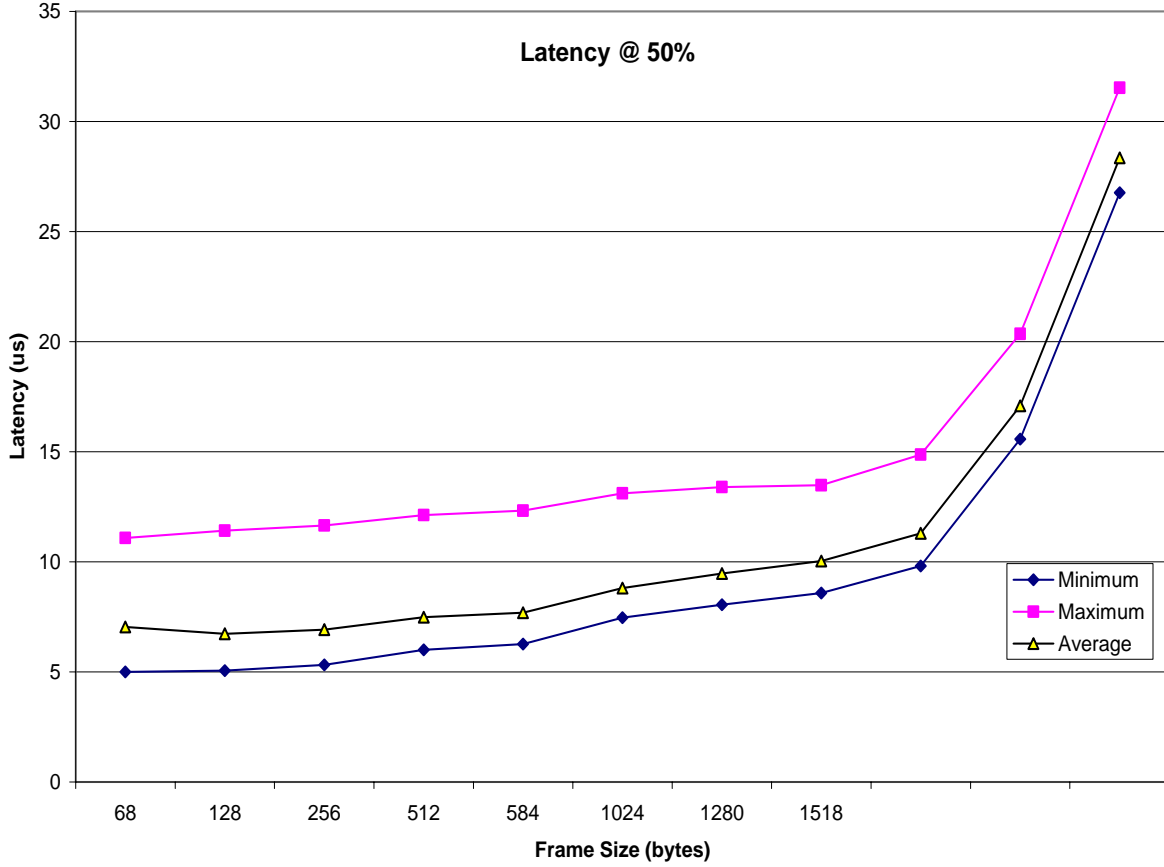


Figure 16: Results of IPv6 Latency at 50% Throughput

IPv6 Latency at 25% Throughput with Internet Mix Traffic Average = 12.14 μ s
 IPv6 Latency at 25% Throughput with Internet Mix Traffic Minimum = 6.81 μ s
 IPv6 Latency at 25% Throughput with Internet Mix Traffic Maximum = 16.39 μ s

4.4 Loss Rate

4.4.1 IPv4 Loss Rate

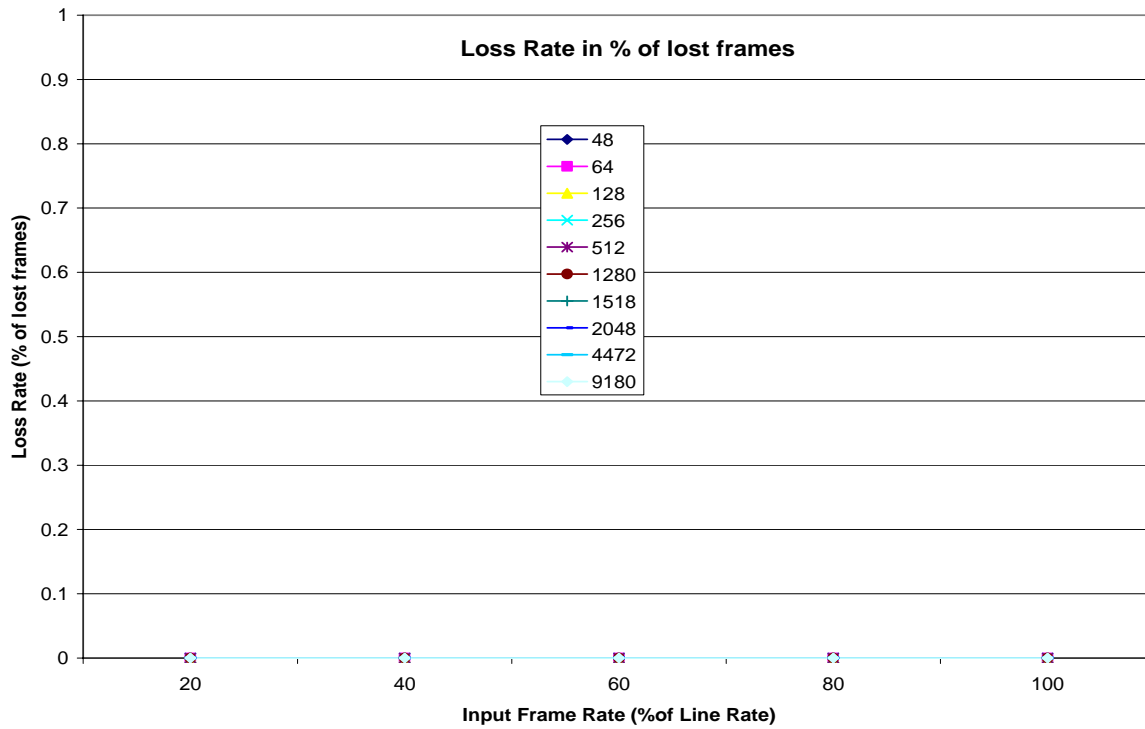


Figure 17: Results of IPv4 Loss Rate Benchmark

Note: There is no frame loss for any packet size at any rate.

4.4.2 IPv6 Loss Rate

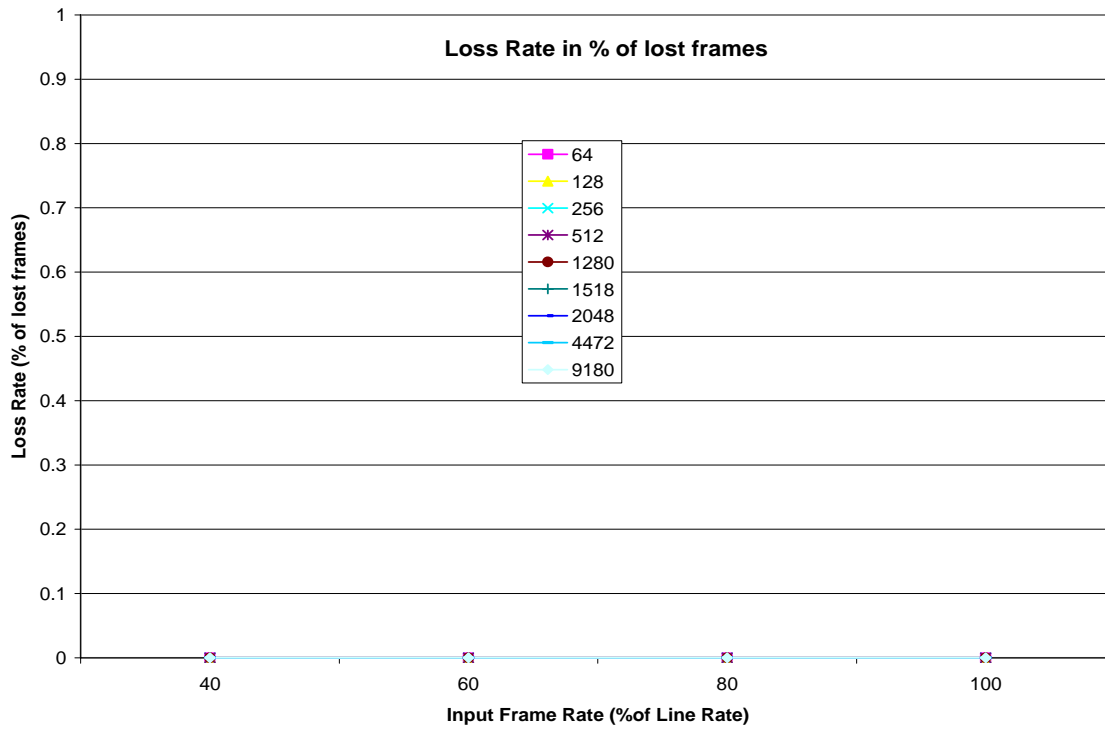


Figure 18: Results of IPv6 Loss Rate Benchmark

Note: There is no frame loss for any packet size at any rate.

4.5 Overload Forwarding Rate

The overload configuration is not valid in a two port forwarding setup. Hence the test was not carried out.

4.6 System Power Consumption

- At 100% of throughput rate each IXP2800 consumes upto a maximum of 30 Watts
- The typical power consumption of each IXP2800 is 25.5

5 Control Plane Benchmark Tests

5.1 Forwarding Table Update Rate

Maximum IPv4 Route update rate = 58,883 route updates/sec, with 1 route update per call.
Maximum IPv6 Route update rate = 24,000 route updates/sec, with 1 route update per call.

5.2 Forwarding Rate With Concurrent Forwarding Table Update Rate

5.2.1 IPv4 Forwarding Rate With Concurrent Forwarding Table Update Rate




Figure 19: Results of IPv4 Forwarding Rate With 100% of Concurrent Forwarding Table Update Rate Measured in pps

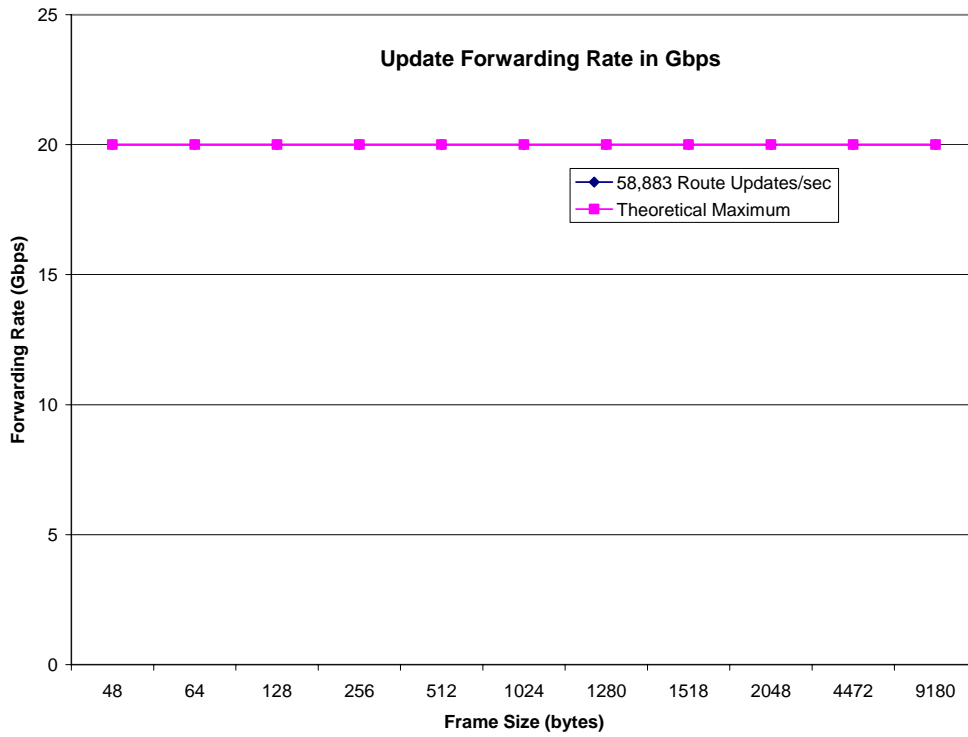


Figure 20: Results of IPv4 Forwarding Rate With 100% of Concurrent Forwarding Table Update Rate Measured in Gbps

IPv4 Forwarding Rate With 100% of Concurrent Forwarding Table Update Rate and Internet Mix Traffic = 20 Gbps

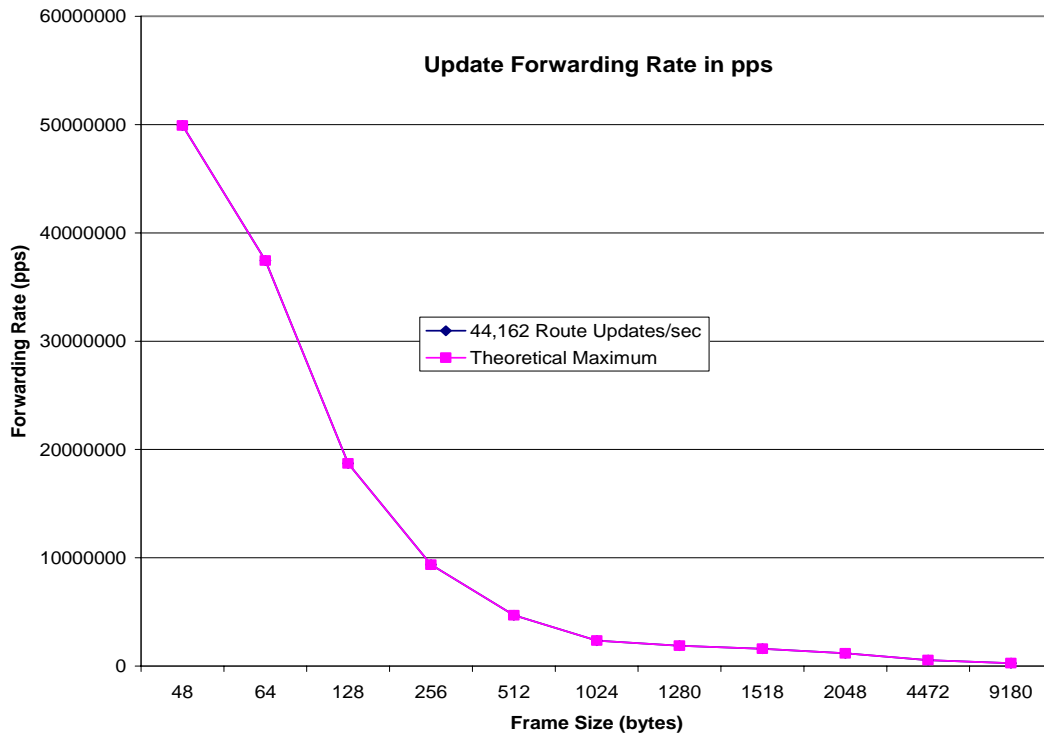


Figure 21: Results of IPv4 Forwarding Rate With 75% of Concurrent Forwarding Table Update Rate Measured in pps

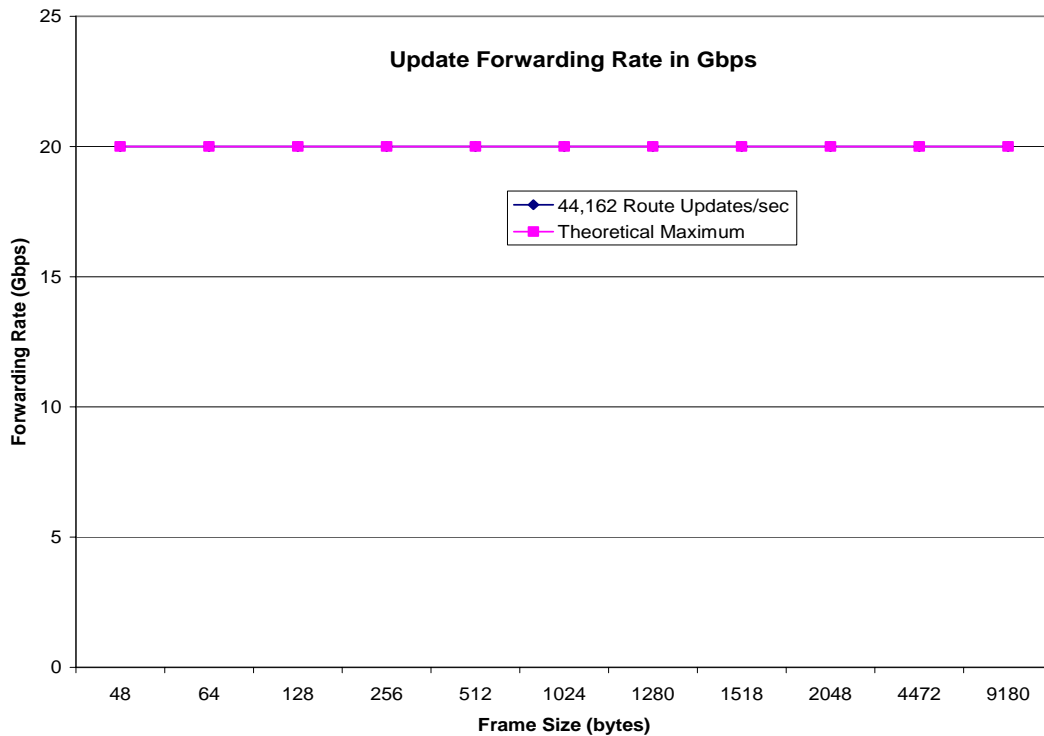


Figure 22: Results of IPv4 Forwarding Rate With 75% of Concurrent Forwarding Table Update Rate Measured in Gbps

IPv4 Forwarding Rate With 75% of Concurrent Forwarding Table Update Rate and Internet Mix Traffic = 20 Gbps




Figure 23: Results of IPv4 Forwarding Rate With 50% of Concurrent Forwarding Table Update Rate Measured in pps

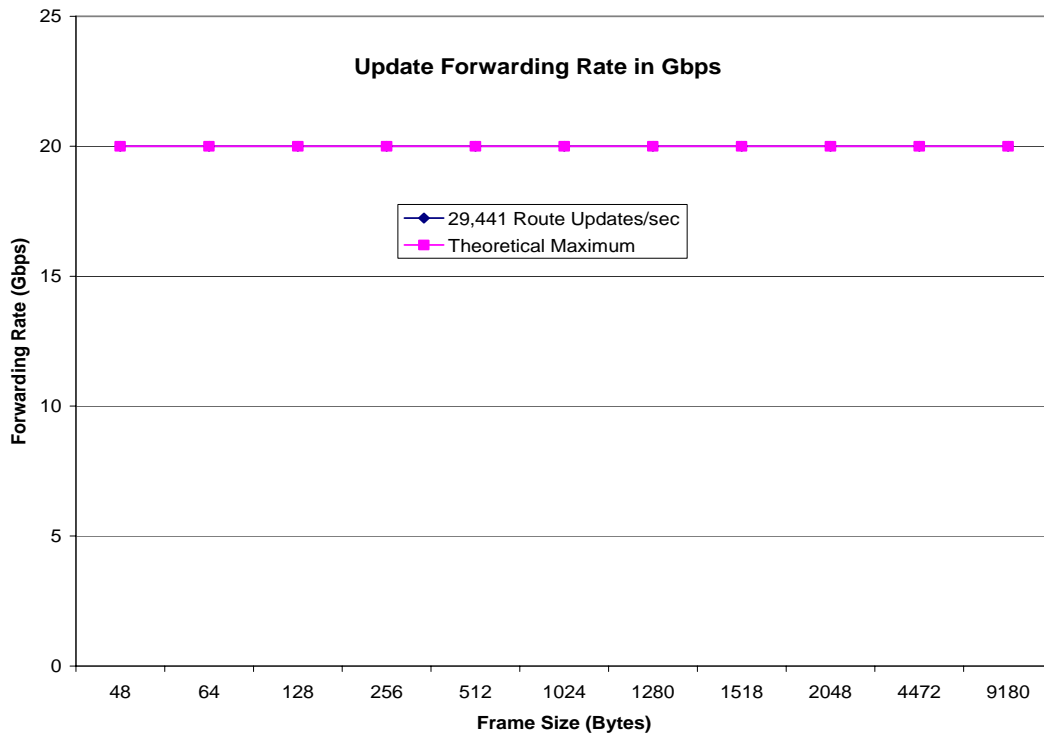


Figure 24: Results of IPv4 Forwarding Rate With 50% of Concurrent Forwarding Table Update Rate Measured in Gbps

IPv4 Forwarding Rate With 50% of Concurrent Forwarding Table Update Rate and Internet Mix Traffic = 20 Gbps

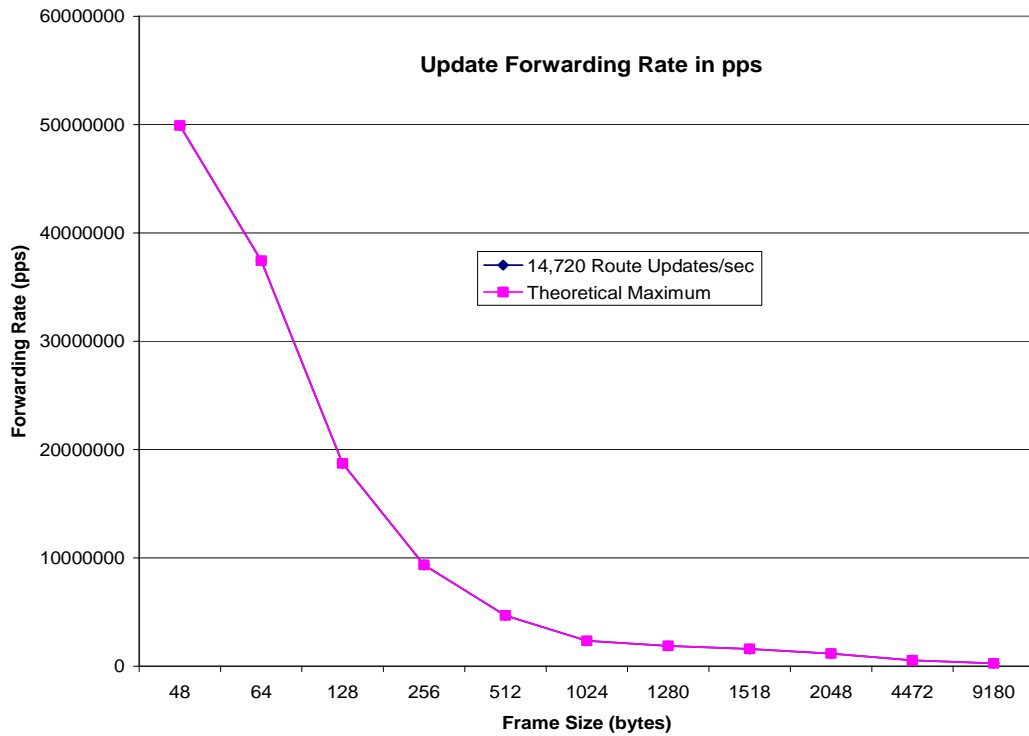


Figure 25: Results of Forwarding Rate With 25% of Concurrent Forwarding Table Update Rate Measured in pps

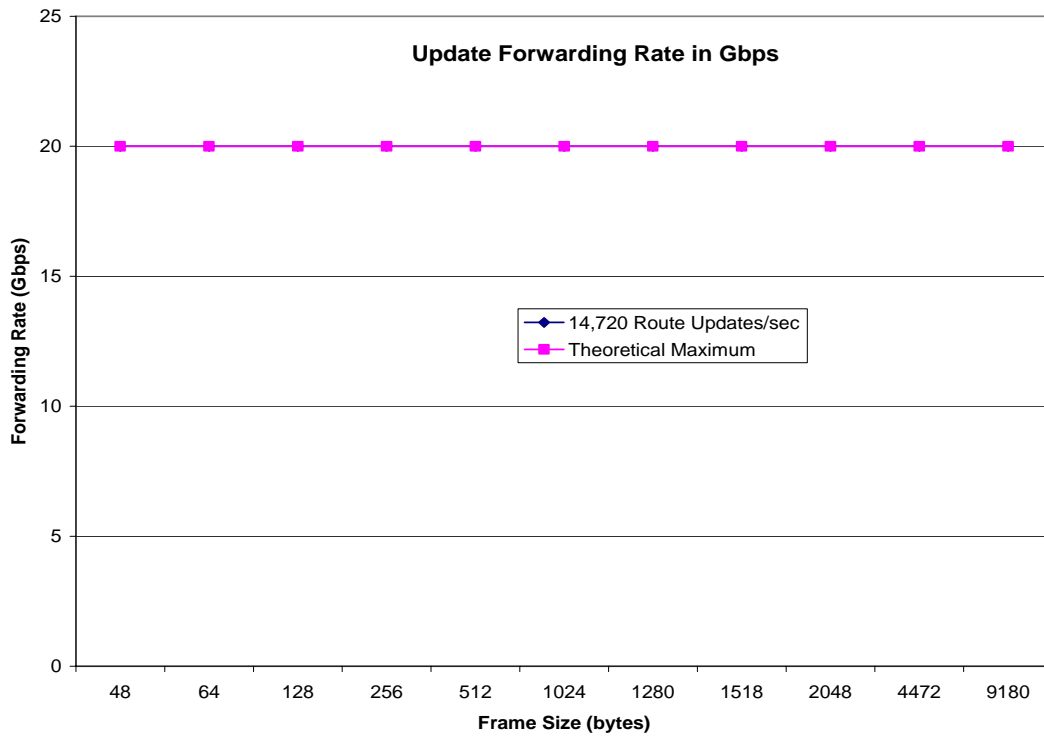


Figure 26: Results of Forwarding Rate With 25% of Concurrent Forwarding Table Update Rate Measured in Gbps

IPv4 Forwarding Rate With 25% of Concurrent Forwarding Table Update Rate and Internet Mix Traffic = 20 Gbps

5.2.2 IPv6 Forwarding Rate With Concurrent Forwarding Table Update Rate

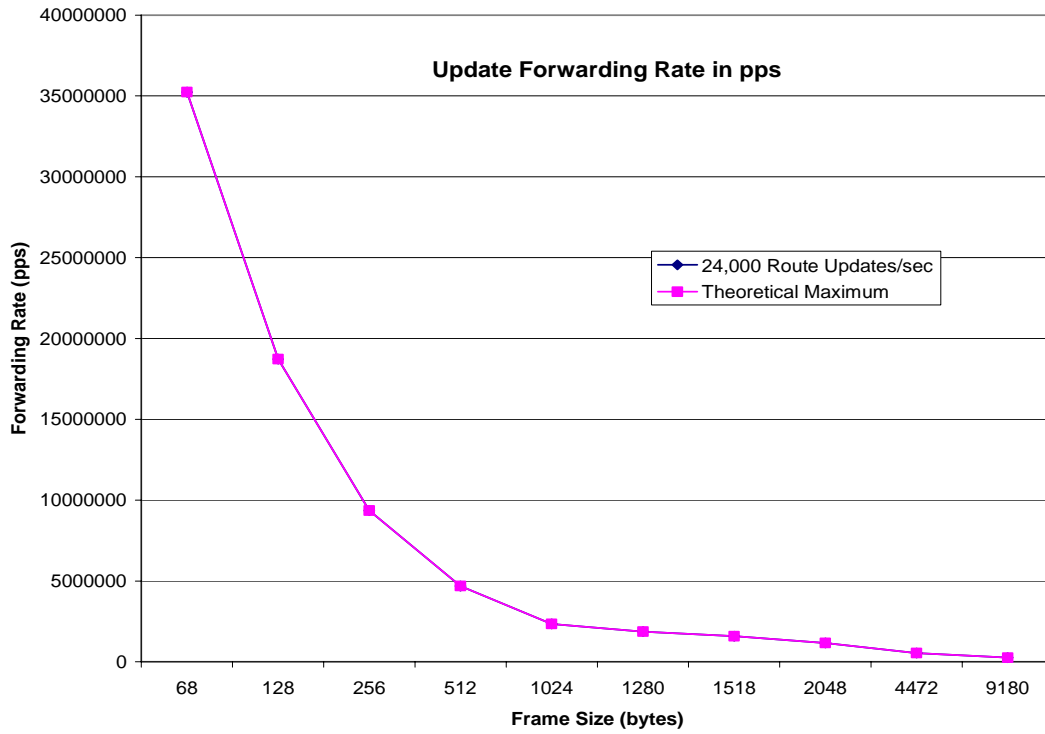


Figure 27: Results of IPv6 Forwarding Rate With 100% of Concurrent Forwarding Table Update Rate Measured in pps

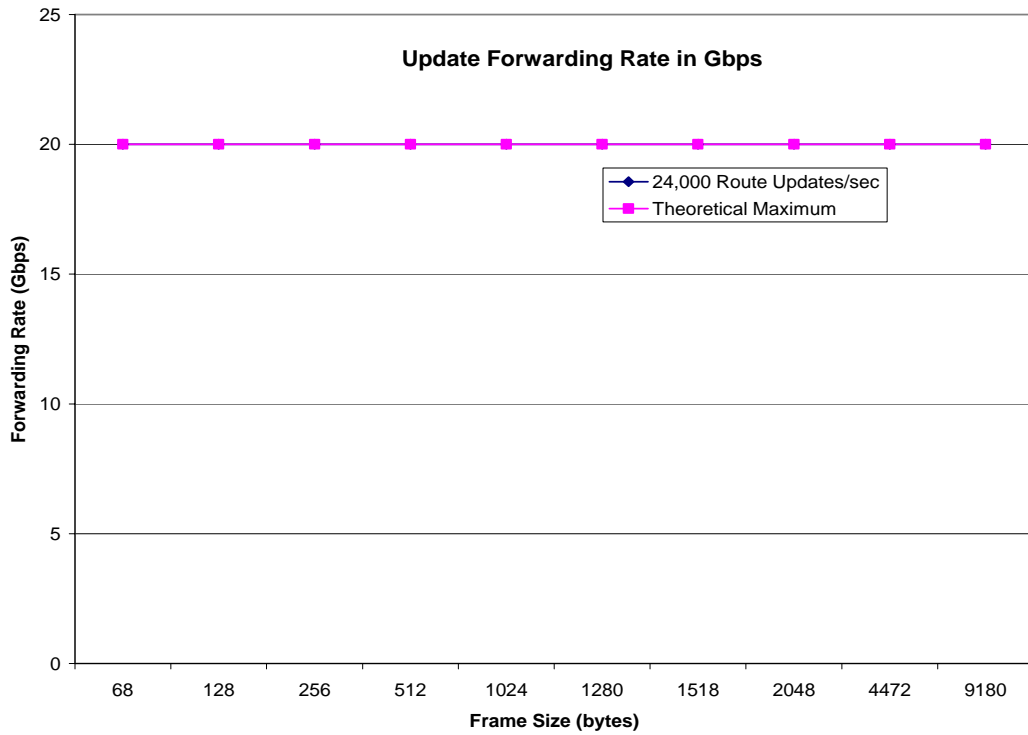


Figure 28: Results of IPv6 Forwarding Rate With 100% of Concurrent Forwarding Table Update Rate Measured in Gbps

IPv6 Forwarding Rate With 100% of Concurrent Forwarding Table Update Rate and Internet Mix Traffic = 20 Gbps




Figure 29: Results of IPv6 Forwarding Rate With 75% of Concurrent Forwarding Table Update Rate Measured in pps




Figure 30: Results of IPv6 Forwarding Rate With 75% of Concurrent Forwarding Table Update Rate Measured in Gbps

IPv6 Forwarding Rate With 75% of Concurrent Forwarding Table Update Rate and Internet Mix Traffic = 20 Gbps




Figure 31: Results of IPv6 Forwarding Rate With 50% of Concurrent Forwarding Table Update Rate Measured in pps

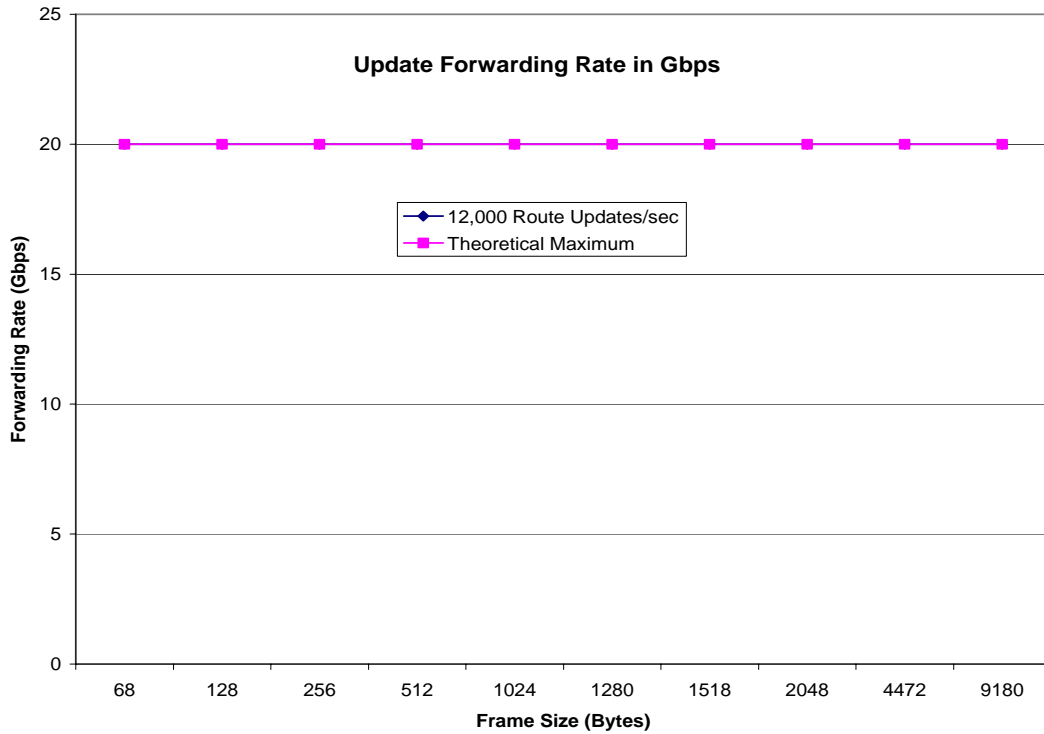


Figure 32: Results of IPv6 Forwarding Rate With 50% of Concurrent Forwarding Table Update Rate Measured in Gbps

IPv6 Forwarding Rate With 50% of Concurrent Forwarding Table Update Rate and Internet Mix Traffic = 20 Gbps

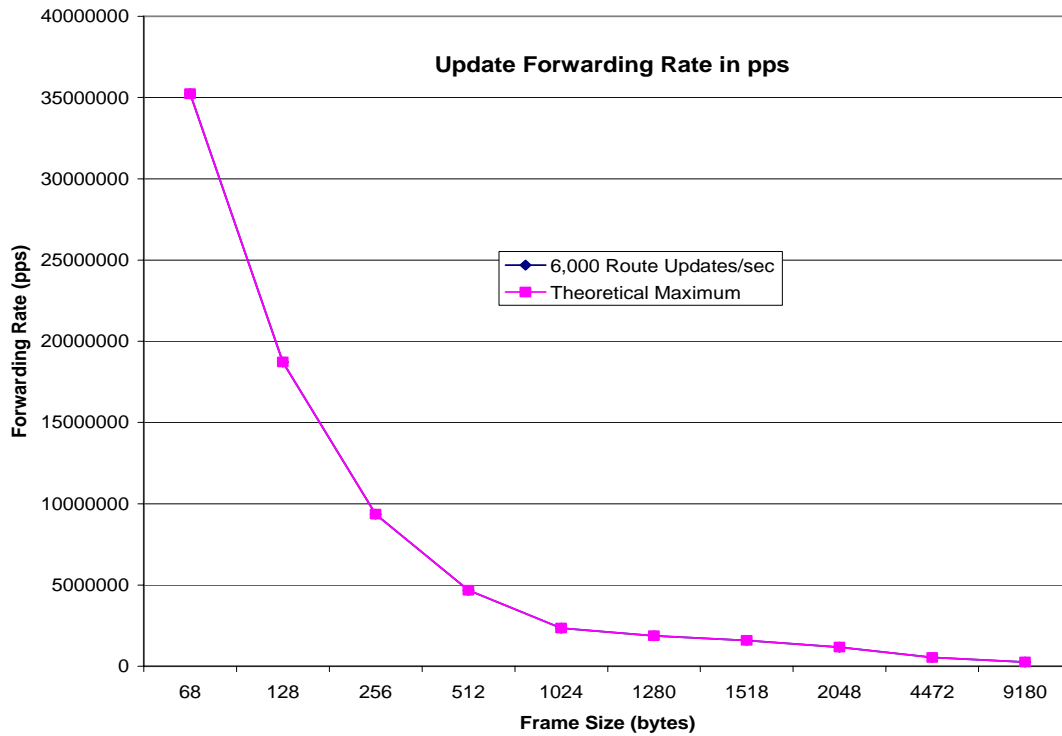


Figure 33: Results of IPv6 Forwarding Rate With 25% of Concurrent Forwarding Table Update Rate Measured in pps

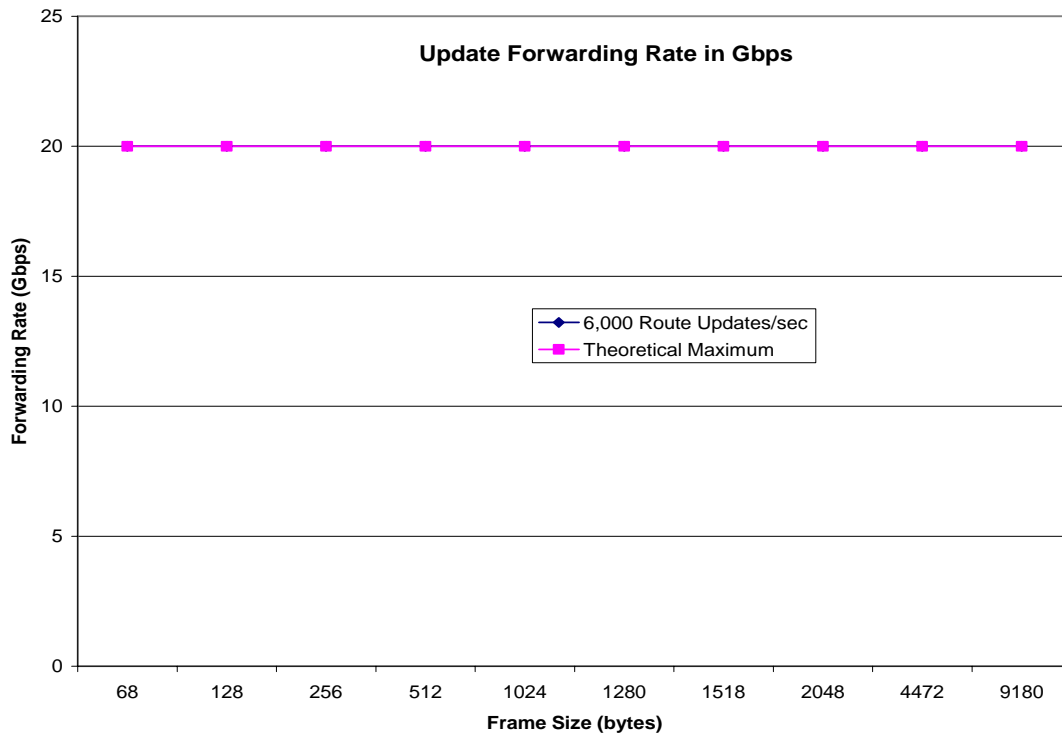


Figure 34: Results of IPv6 Forwarding Rate With 25% of Concurrent Forwarding Table Update Rate Measured in Gbps

IPv6 Forwarding Rate With 25% of Concurrent Forwarding Table Update Rate and Internet Mix Traffic = 20 Gbps