White Paper

Performance-Comparison Testing of IPv4 and IPv6 Throughput and Latency on Key Cisco Router Platforms

A Summary of Findings

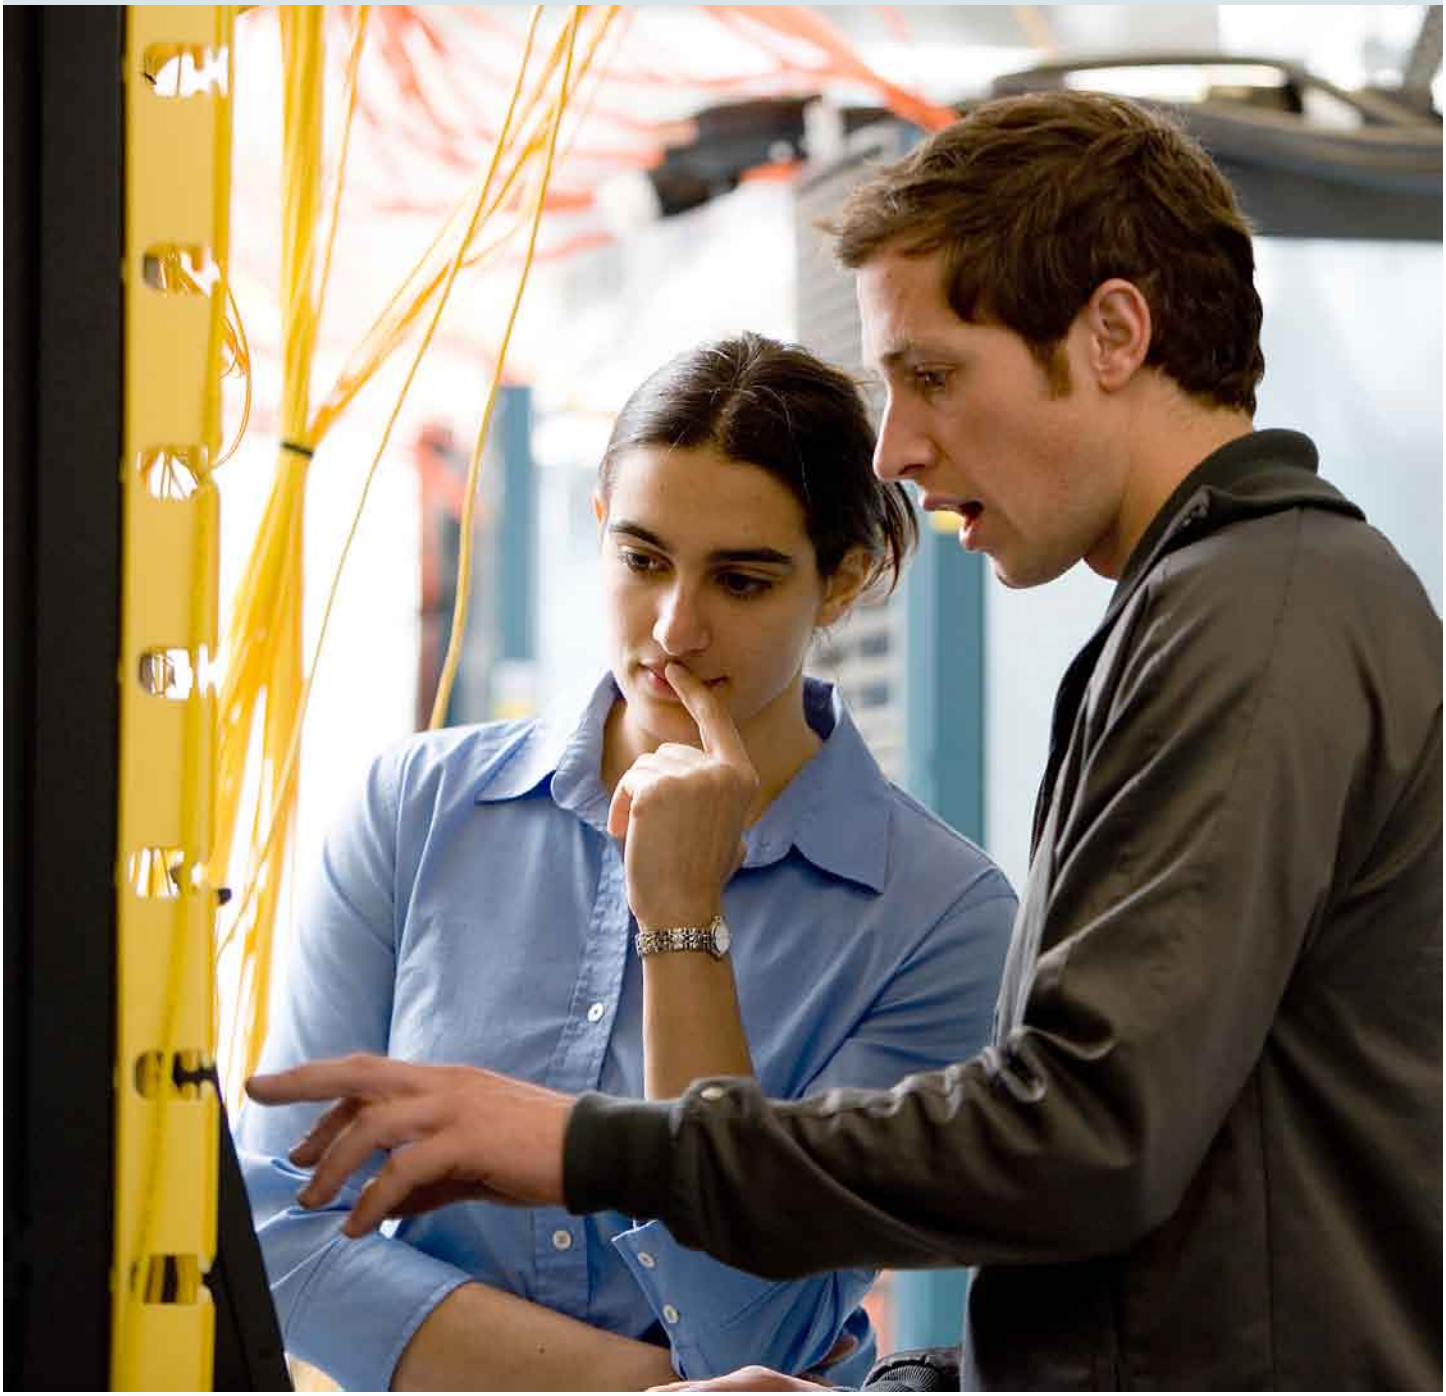

Executive Overview

In response to requests from federal civilian agencies like the Social Security Administration and the Department of Education, and from the Joint Chiefs of Staff — who had initiated a Statement of Work (SoW) with the Joint Interoperability Test Command (JITC) — the Government Systems Engineering (GSE) team recently executed tests to provide data on Cisco router platforms in a dual-stack IPv4 and IPv6 environment.

The organizations' major concern was that they operated in five-year purchase cycles, meaning the equipment they purchase today remains in their networks for five years before it can be refreshed/replaced. Therefore, decision-makers wanted to know what, if any, impact or performance degradation turning on IPv6 in their networks would have on new equipment. Specifically, they feared that, while IPv6 packet size had increased from 32 to 128 bits, networking gear bus-width and CPU-lookup sizes remained predominately at 32 or 64 bits.

The purpose of this testing effort was to provide information to Cisco's federal sales team and Cisco customers regarding IPv4 and IPv6 performance in a realistic network environment. With this information, sales personnel and customers alike can make reasonable, informed decisions on upcoming purchase cycles.

Overview of Results

As a prelude to the detailed results provided in this document, our testing showed that overall, across all platforms, IPv4 and IPv6 interface level throughput and latency results were remarkably similar. It was only at the smaller packet sizes — generally 256 bytes or less — that IPv6 showed a lower throughout compared to IPv4. At the larger frame sizes, IPv4 and IPv6 throughput is typically identical.

The data also verifies the difference in IPv4 and IPv6 throughput using small packet sizes was generally only seen on the smaller software switching platforms tested (e.g., Cisco1841 ISR). Larger hardware switching platforms, like the Cisco 7606, showed no throughout variance even at the smaller packet size.

For More Information

Testing results have been condensed in this document for easy readability. To review test results in their entirety, please email your request to IPv6_info@cisco.com.

Throughput and Latency Measurements

As the new IP protocol designed to replace IPv4, IPv6 quadruples the number of address bits from 32 bits (in IPv4) to 128 bits or approximately 3.4 x 1038 addressable nodes, which provides more than enough globally unique IP addresses for every network device on the planet.

The purposes of these tests were to explore and document the throughput and latency measurements of various IPv4 and IPv6 traffic mixes. This document summarizes the findings of extensive processing-performance comparison testing for the following Cisco router platforms:

- Cisco 1841 ISR
- Cisco 2811 ISR
- Cisco 3825 ISR
- Cisco 7206 Router
- Cisco 7301 Router
- Cisco 7606 Router

Test Parameters

In conducting this test, a range of IP traffic was injected into the Device Under Test (DUT) via an independent test tool using a pair of Ethernet interfaces on the DUT. To provide a range of data that accurately characterized performance divergences, the DUT was analyzed under the following conditions:

1. 100% IPv4
2. 100% IPv6
3. A variety of dual-stack mode configurations representing various IPv4 and IPv6 traffic mixes

Important Considerations when Reviewing this Document

It is important to distinguish between performance-comparison testing and full-capacity platform-performance testing typically done by a business unit. The purpose of this testing was to provide an IPv4 to IPv6 comparison for performance on a router "fresh out of the box." When reviewing the results, please keep the following in mind:

- These tests were performed with a default router configuration.
- Many additional enhancements could be made to improve performance of the device in specific situations (for voice, video, etc.).
- These numbers show the relative performance gap between 100% IPv4 performance, 100% IPv6 performance and several dual stack configurations.
- The term "throughput" as used in this document refers to interface level throughput (e.g. interface to interface) as measured on a device using a default configuration under various test loads. It does not refer to the maximum rate at which the entire system can forward..

Table of Contents

Section 1 – Methodology

Automated Multiple Platform Testing

Testing Approach

The purpose of this testing was to assess basic traffic throughput, frame loss and latency variances. Reference data was first taken running IPv4-only traffic, then tests were run with the DUT having various percentages of IPv4 and IPv6 traffic running simultaneously with bidirectional flows.

The test setup used the Spirent Test Center (STC) to inject traffic into the DUT and to collect and analyze the results. The main measurements taken from the testing were:

- **Throughput:** Maximum number of frames per second with no frame loss
- **Frame Loss:** Number of lost frames at specific line rates
- **Latency:** Delay of traffic through a router
- **CPU Utilization:** Percentage of CPU time being used by the DUT

Description of Devices Under Test

DUT	Software	Model	Processor Board ID	CPU	Midplane	FastEthernet/ Gigabit Interface(s)	NVRAM
1841	– Cisco IOS – 1841 (C1841-ADVENTERPRISEK9-M) – Version 12.4(11)T – RELEASE (fc2)	Revision 5.0 with 355328K/37888K bytes of memory	FTX0934W0MF	NA	NA	2/NA	191K bytes
2811	– Cisco IOS – 2800 (C2800NM-ADVENTERPRISEK9-M) – Version 12.4(11)T – RELEASE (fc2)	Revision 53.51with 774144K/12288K bytes of memory	FTX1020A0HY	NA	NA	2/NA	239K bytes
3825	– Cisco IOS – 3800 (C3825-ADVIPSERVICESK9-M) – Version 12.4(11)T – RELEASE (fc2)	Revision 1.1 with 1009664K/38912K bytes of memory	FTX1024A406	NA	NA	2/2	479K bytes
7206	– Cisco IOS – 7200 (C7200P-ADVENTERPRISEK9-M) – Version 12.4(11)T – RELEASE (fc2)	7206VXR (NPE-G2) processor (revision A) with 917504K/65536K bytes of memory	34979327	– MPC7448 at 1666Mhz – Implementation 0 – Rev 2.1	– 6 slot VXR – Version 2.11	1/3	2045K bytes
7301	– Cisco IOS – 7301 (C7301-ADVENTERPRISEK9-M) – Version 12.4(11)T – RELEASE (fc2)	NPE processor (revision D) with 491520K/32768K bytes of memory	0	– SB-1 at 700MHz – Implementation 1025 – Rev 0.2, 512KB L2 Cache	– 1 slot – Version 2.0	NA/3	509K bytes
7606	– Cisco IOS – c7600s72033_rp (c7600s72033_rp-ADVENTERPRISEK9-M) – Version 12.2(33)SRB – RELEASE (fc6)	CISCO7606 (R7000) processor (revision 1.0) with 983008K/65536K bytes of memory	FOX104612JY	– SR71000 at 600Mhz – Implementation 0x504 – Rev 1.2 – 512KB L2 Cache	NA	– 2 SIP-600 controllers (2 TenGigabitEthernet). – 1 Virtual Ethernet interface – 2 Gigabit Ethernet interfaces – 2 10-Gigabit Ethernet interfaces	– 1917K bytes – 8192K bytes of packet buffer memory

Test Topology

The test automation used a very basic test topology, shown below. The interface speed (100M, 1G or 10G) was based on the DUT and each test lasted 120 seconds.

Throughput Test

Finds the highest packet rate that can be switched through a given interface type and for a given packet size without packet loss. If a single frame is dropped, the test fails. It is then repeated at a lower rate.

- Measures throughput by sending a series of frames with particular source and destination MAC addresses to the DUT.
- Frames are sent by one of the Spirent TestCenter test ports and are intended to be received by a second test port. The number of frames is determined by the length of time the test is to run.
- Received frames are validated by counting only the frames generated by the transmitting test port, not those sent by the DUT. Keep-alive and routing update frames are not counted in the test as received frames.

Latency Test

Determines the delay (latency) imposed by the DUT on forwarded traffic.

- As the test runs, the transmitting test port sends a burst of frames, at a user-specified frame size, to the DUT, at a user-specified throughput rate.
- In the middle of the burst stream, the port inserts one frame with an identifying trigger. The time when the trigger frame is fully transmitted represents the Transmit Timestamp.
- The time the receiving test port recognizes the trigger frame is the Receive Timestamp.
- (Receive Timestamp) – (Transmit Timestamp) = Latency

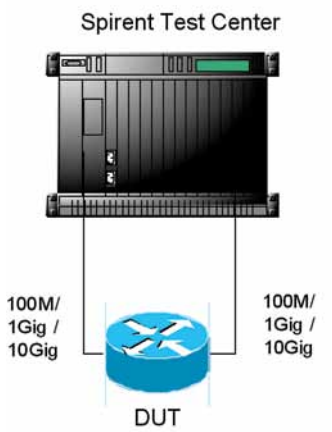

Testing Guidelines

We constructed our test using the following reference material:

1. RFC 2544 - Benchmarking Methodology for Network Interconnect Devices
2. IPv6 Benchmarking Methodology (IETF draft)
3. DoD IPv6 JCS Decomposed Criteria

Test Variables

Throughput, frame loss, latency and DUT health was measured at varying test conditions, including every combination of the following:

- IPv4/IPv6 Ratio (100% / 0% , 90% / 10% , 50% / 50% , 10% / 90% , 0% / 100%)
- Frame Sizes (86, 128, 256, 512, 768, 1024, 1280, 1518)
- Line Rate (10% 20% 30% 40% 50% 60% 70% 80% 90% 100%)

Note

1. All traffic size numbers represent frames in bytes, not packets.
2. 86 bytes was chosen as the minimum frame size tested because it is the smallest UDP frame size that will take advantage of the Spirent Test Center capabilities:
   - 66 bytes (smallest IPv6 frame with no Layer 4 header and no STC signature field)
   - 78 bytes (smallest IPv6 frame with no Layer 4 header & including the STC signature field)
   - 86 bytes (smallest IPv6 frame with UDP header and including the STC signature field)
3. Only Ethernet interfaces were tested

Section 2 – Results

Automated Multiple Platform Testing Results

At a Glance: Summary of Test Results

- The 7606 platform lost zero frames at 100 percent of 10Gbps line rate (bidirectional) regardless of frame size tested or amount of IPv6 traffic in the test stream.
- On the other test platforms, smaller frame sizes showed a degradation of throughput directly proportional to the amount of IPv6 traffic. The severity of the degradation differed between each platform, with the 7200 being the least severe and the 2811 being the most severe. This is expected because of the relative sizing of the centralized CPU on the respective routers.
- Increasing the amount of IPv6 in the test stream had little impact on the average latency.
- The CPU usage was essentially the same for a dual stack environment as compared to an IPv4 only environment, indicating that dual stack does not impact CPU usage.

Chart Legend

Configuration	Translation
100-0	100% IPv4 traffic 0% IPv6 traffic
90-10	90% IPv4 traffic 10% IPv6 traffic
50-50	50% IPv4 traffic 50% IPv6 traffic
10-90	10% IPv4 traffic 90% IPv6 traffic
0-100	0% IPv4 traffic 100% IPv6 traffic

Cisco 1841 ISR

Throughput

1841 Throughput Rate (100M bidirectional traffic)

The following table shows the throughput rate (bidirectional traffic via the 2 onboard 100M Ethernet interfaces) for the 1841.

1841 Throughput Rate (100M bidirectional traffic)

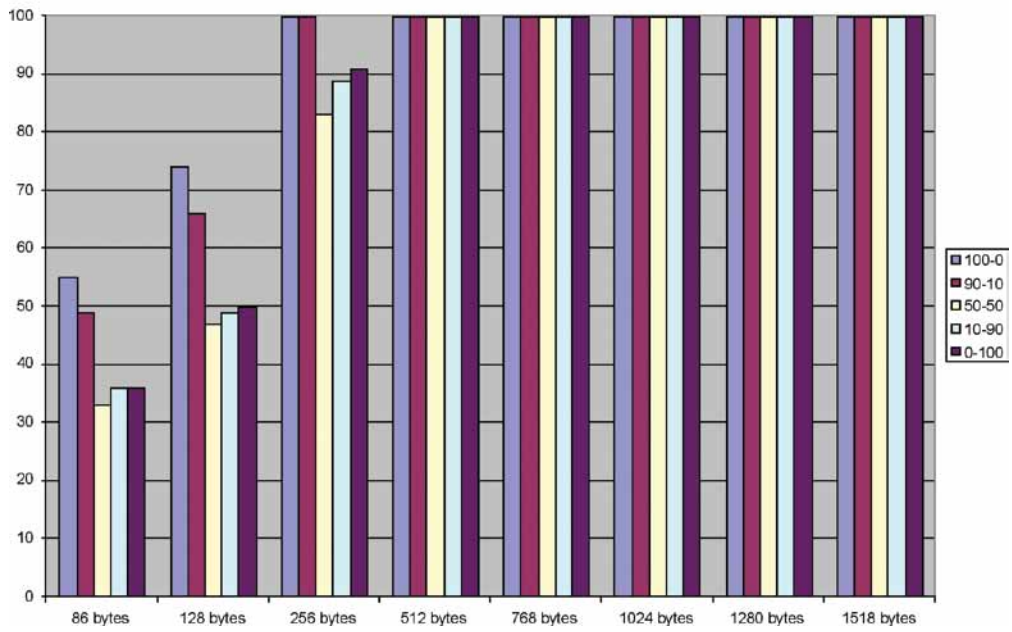

1841 Relative Throughput with Respect to IPv4

The following table shows dual stack and 100% IPv6 throughput numbers compared as a percentage of 100% IPv4 numbers. (For example, for 86 bytes, the throughput percentage with no loss is 55%. The 90% IPv4/10% IPv6 throughput percentage for 86 bytes is 49% - 49/55 x 100 = 89.091%.) This is the relative performance of the 90/10 traffic stream to the 100/0 traffic stream.

The smaller frame sizes of 86, 128 and 256 bytes showed a degradation of throughput as compared to IPv4 (with the worst performance at the 50% / 50% traffic ratio). At the larger frame sizes, throughput equaled that of 100% IPv4 traffic.

1841 Relative Max Throughput with Respect to IPv4

Frame Size (bytes)	Rel Max Throughput (%) 90% IPv4 / 10% IPv6	Rel Max Throughput (%) 50% IPv4 / 50% IPv6	Rel Max Throughput (%) 10% IPv4 / 90% IPv6	Rel Max Throughput (%) 0% IPv4 / 100% IPv6
86	89.091	60	65.455	65.455
128	89.189	63.514	66.216	67.568
256	100	83	89	91
512	100	100	100	100
768	100	100	100	100
1024	100	100	100	100
1280	100	100	100	100
1518	100	100	100	100

Latency

When comparing each frame size individually, there was a slight increase to frame latency with the introduction of IPv6 traffic. The upward trend in latency as the frame size increased mimics the trend that exists for 100% IPv4 traffic. (Note: smaller frame sizes are excluded from the 90% line rate chart due to frame loss.)

1841 Average Latency at the 30% Line Rate

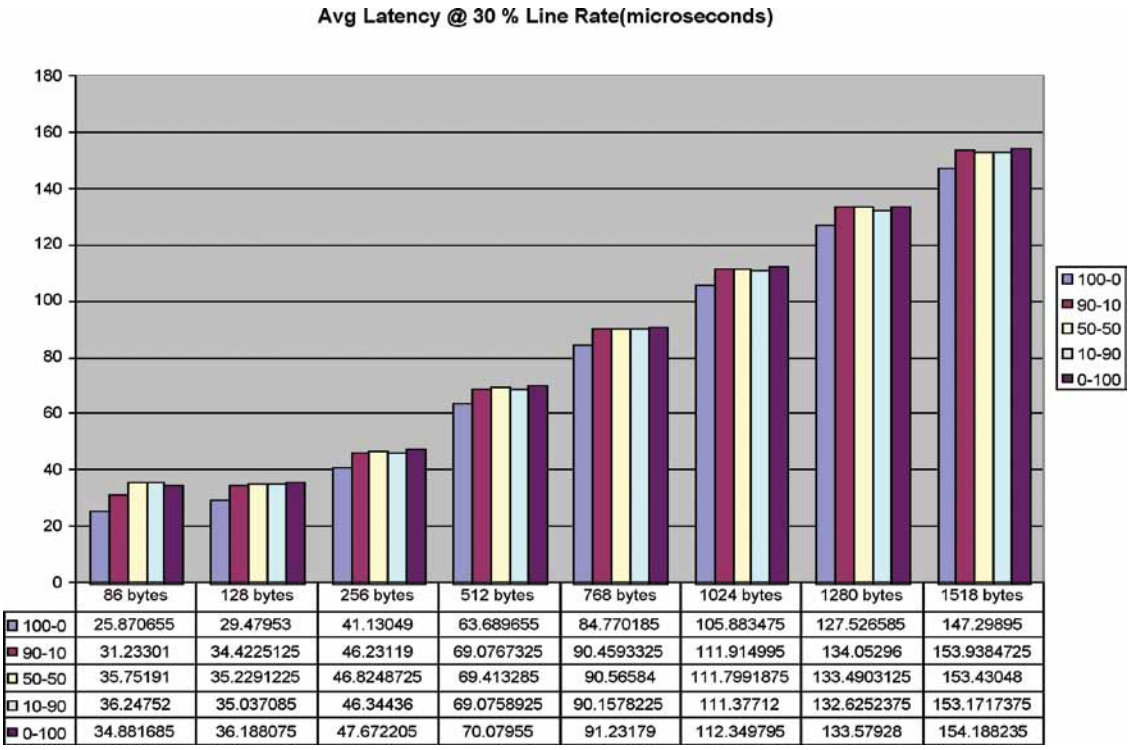

	86 bytes	128 bytes	256 bytes	512 bytes	768 bytes	1024 bytes	1280 bytes	1518 bytes
100-0	25.870655	29.47953	41.13049	63.689655	84.770185	105.883475	127.526585	147.29895
90-10	31.23301	34.4225125	46.23119	69.0767325	90.4593325	111.914995	134.05296	153.9384725
50-50	35.75191	35.2291225	46.8248725	69.413285	90.56584	111.7991875	133.4903125	153.43048
10-90	36.24752	35.037085	46.34436	69.0758925	90.1578225	111.37712	132.6252375	153.1717375
0-100	34.881685	36.188075	47.672205	70.07955	91.23179	112.349795	133.57928	154.188235

1841 Average Latency at the 90% Line Rate

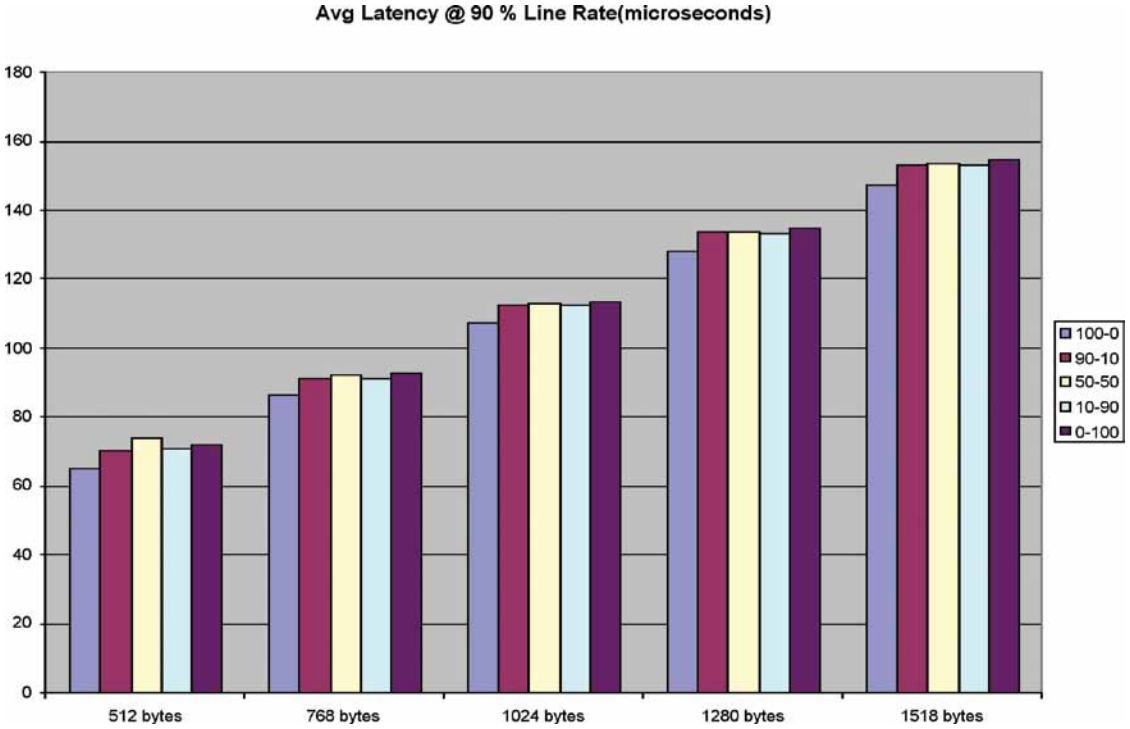

CPU Utilization

The chart below is representative of frame sizes with a throughput value of 100% for all IPv4/IPv6 traffic ratios. The CPU steadily increases as the line rate increases, and the largest utilization per line rate are those traffic mixes with the most IPv6 traffic (100% IPv6 and 10% IPv4 / 90% IPv6).

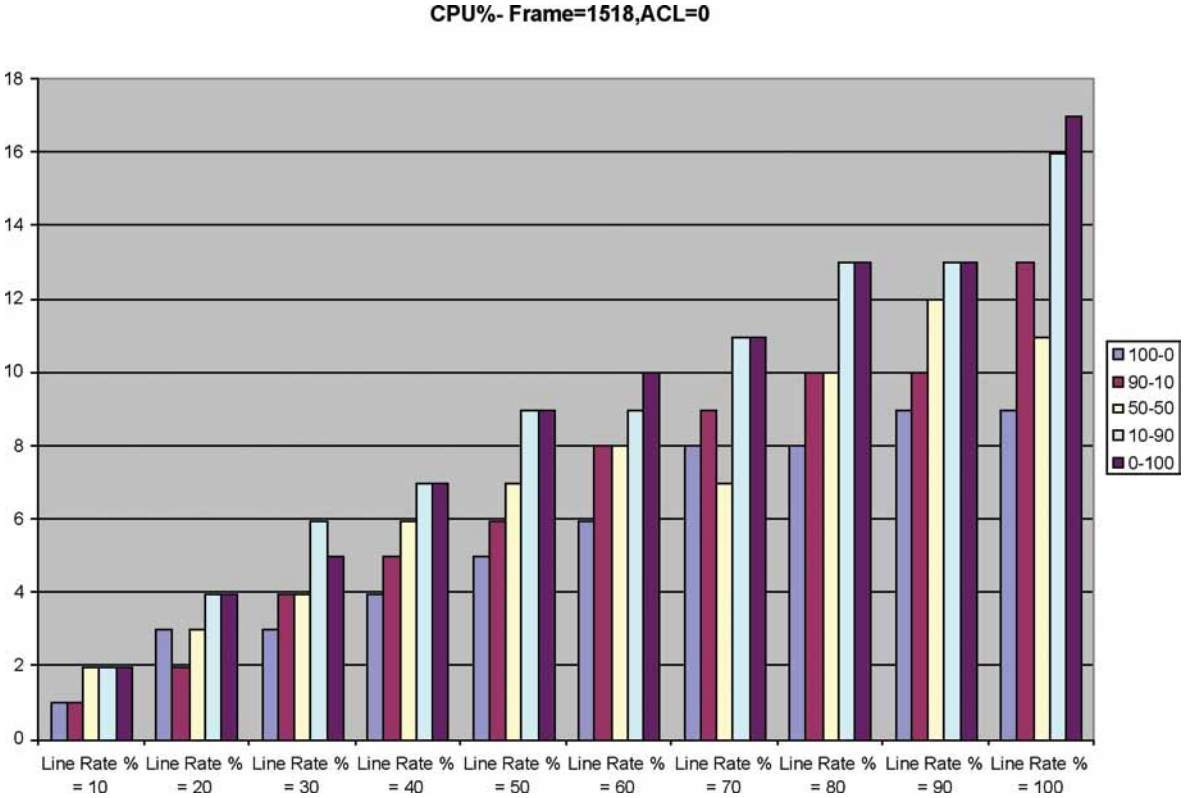

Cisco 2811 ISR

Throughput

2811 Throughput (100M bidirectional traffic)

The chart below shows the throughput rate (bidirectional traffic via the 2 onboard 100M Ethernet interfaces) for the 2811.

2811 Throughput Rate (100M bidirectional traffic)

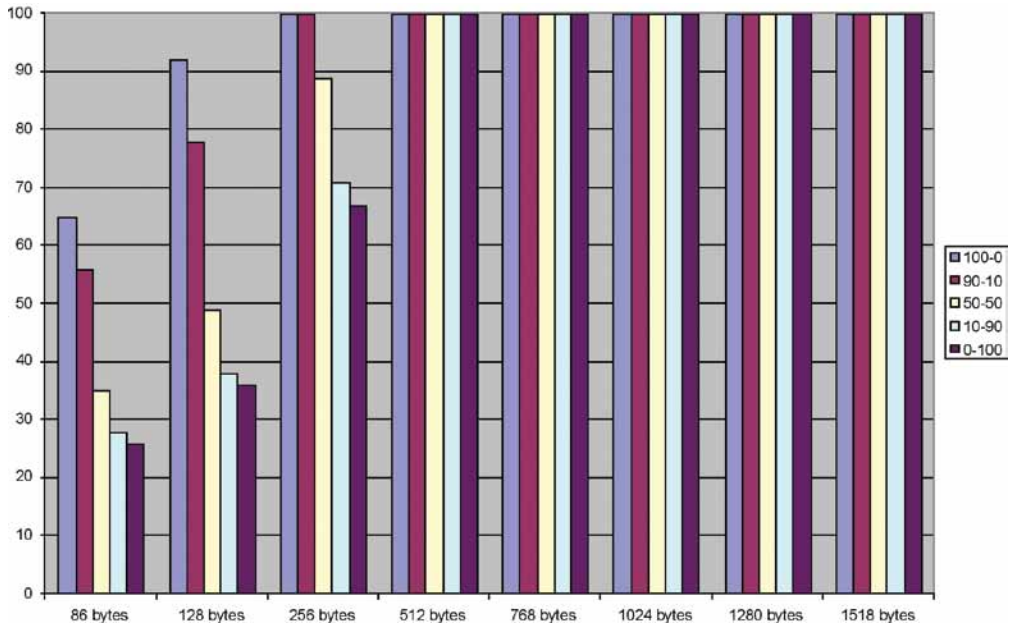

2811 Relative Throughput with Respect to IPv4

The following table shows dual stack and 100% IPv6 relative throughput numbers compared as a percentage of 100% IPv4 numbers. (For example, for 86 bytes the throughput % with no loss is 65%. The 90% IPv4/10% IPv6 throughput % for 86 bytes is 56% - 56/65 x 100 = 86.15384615%.) This is the relative performance of the 90/10 traffic stream to the 100/0 traffic stream.

The smaller frame sizes of 86, 128 and 256 bytes showed a degradation of throughput as compared to IPv4. The degradation appeared directly proportional to the IPv6 traffic in the traffic stream. At the larger frame sizes, throughput equaled that of 100% IPv4 traffic.

2811 Relative Throughput with Respect to IPv4

Frame Size (bytes)	Rel Max Throughput (%) 90% IPv4 / 10% IPv6	Rel Max Throughput (%) 50% IPv4 / 50% IPv6	Rel Max Throughput (%) 10% IPv4 / 90% IPv6	Rel Max Throughput (%) 0% IPv4 / 100% IPv6
86	86.15384615	53.84615385	43.07692308	40
128	84.7826087	53.26086957	41.30434783	39.13043478
256	100	89	71	67
512	100	100	100	100
768	100	100	100	100
1024	100	100	100	100
1280	100	100	100	100
1518	100	100	100	100

Latency

When comparing each frame size individually, there was a slight increase in frame latency with the introduction of IPv6 traffic. The upward trend in latency as the frame size increased mimics the existing trend for 100% IPv4 traffic. There was a substantial increase in 100% IPv6 latency as compared to IPv4 latency for the 512-byte frame size. (Note: smaller frame sizes are excluded from the 90% line rate chart due to frame loss.)

2811 Average Latency at the 20% Line Rate

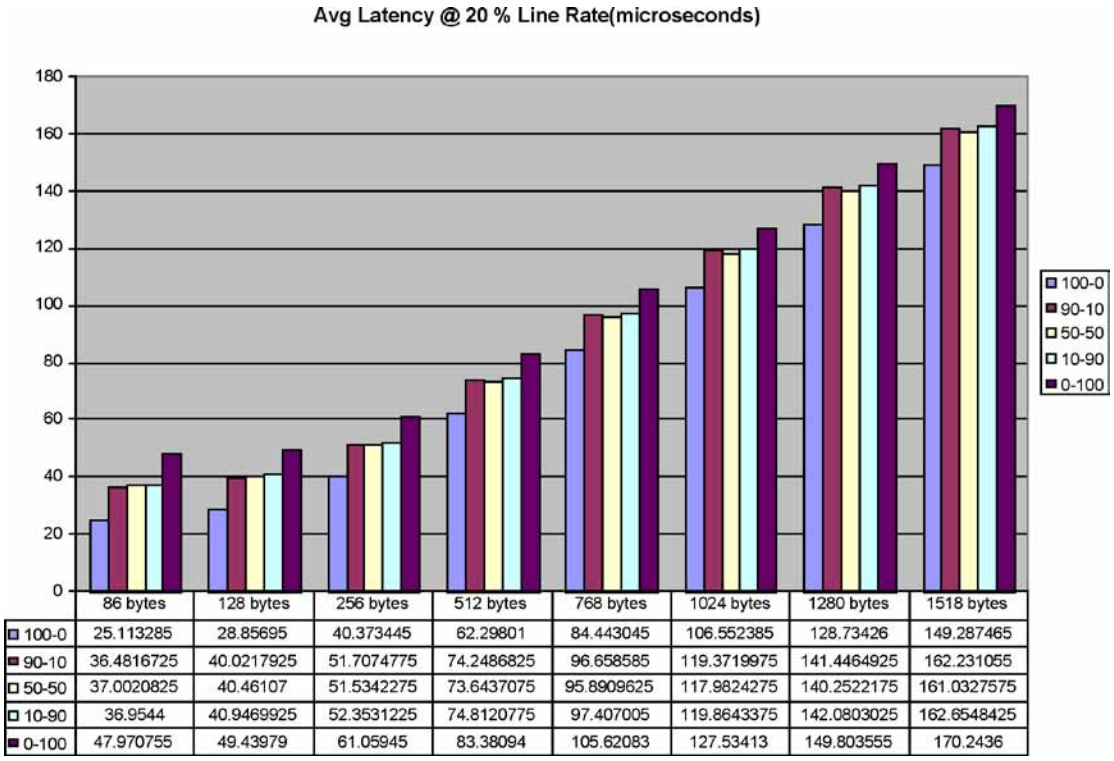

	86 bytes	128 bytes	256 bytes	512 bytes	768 bytes	1024 bytes	1280 bytes	1518 bytes
100-0	25.113285	28.85695	40.373445	62.29801	84.443045	106.552385	128.73426	149.287465
90-10	36.4816725	40.0217925	51.7074775	74.2486825	96.658585	119.3719975	141.4464925	162.231055
50-50	37.0020825	40.46107	51.5342275	73.6437075	95.8909625	117.9824275	140.2522175	161.0327575
10-90	36.9544	40.9469925	52.3531225	74.8120775	97.407005	119.8643375	142.0803025	162.6548425
0-100	47.970755	49.43979	61.05945	83.38094	105.62083	127.53413	149.803555	170.2436

2811 Average Latency at the 90% Line Rate

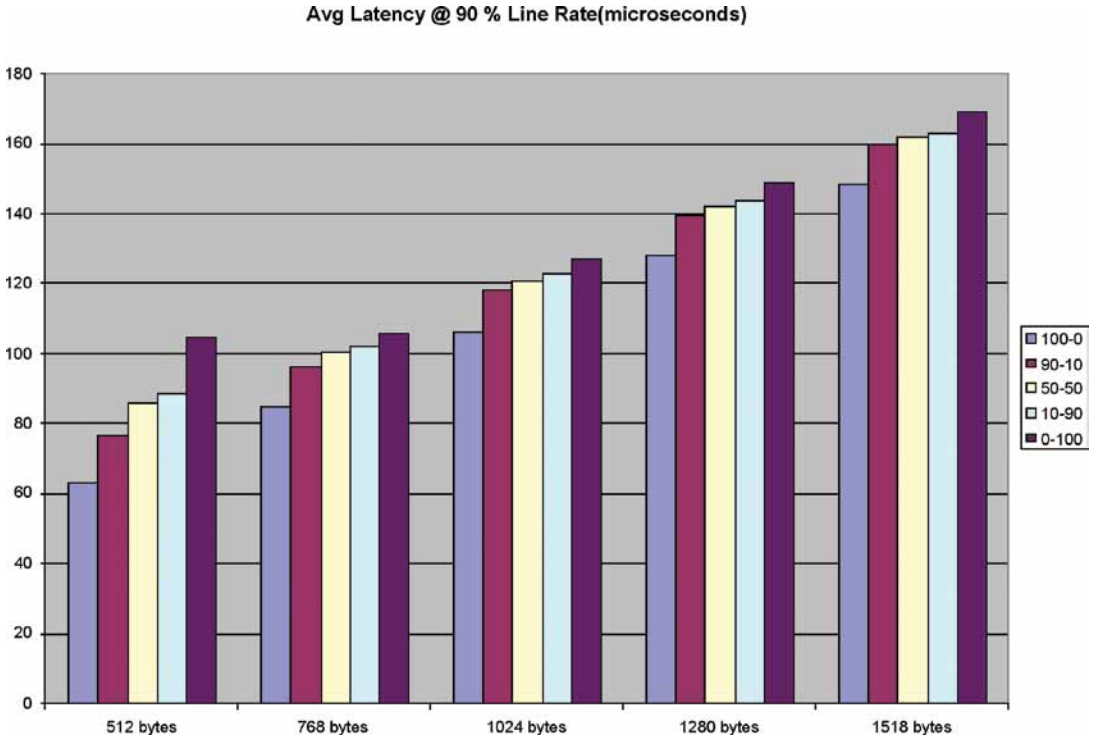

CPU Utilization

The chart below is representative of frame sizes with a throughput value of 100% for all IPv4/IPv6 traffic. The CPU steadily increases as the line rate increases, and the largest utilization per line rate are those traffic mixes with the most IPv6 traffic (100% IPv6 and 10% IPv4 / 90% IPv6).

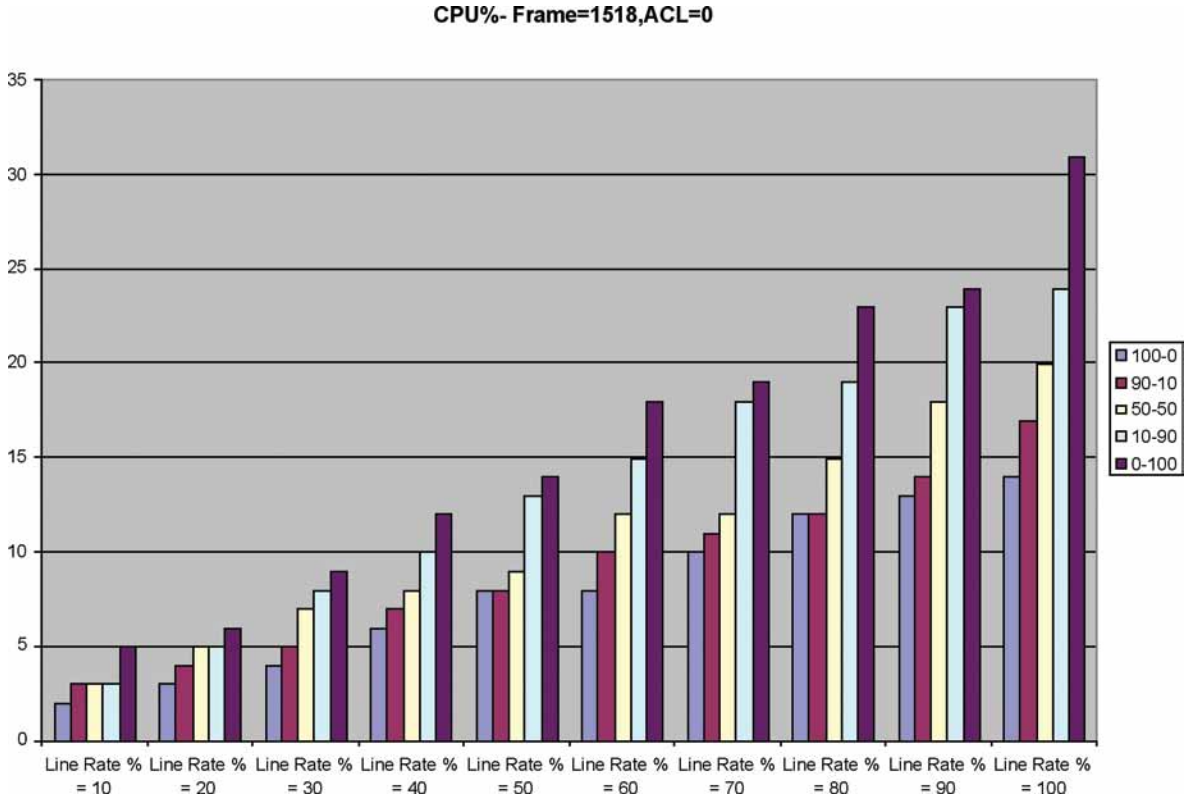

Cisco 3825 ISR

Throughput

3825 Throughput (100M bidirectional traffic)

The chart below shows the throughput rate (bidirectional traffic via the 2 onboard 1000M Ethernet interfaces) for the 3825.

3825 Throughput Rate (100M bidirectional traffic)

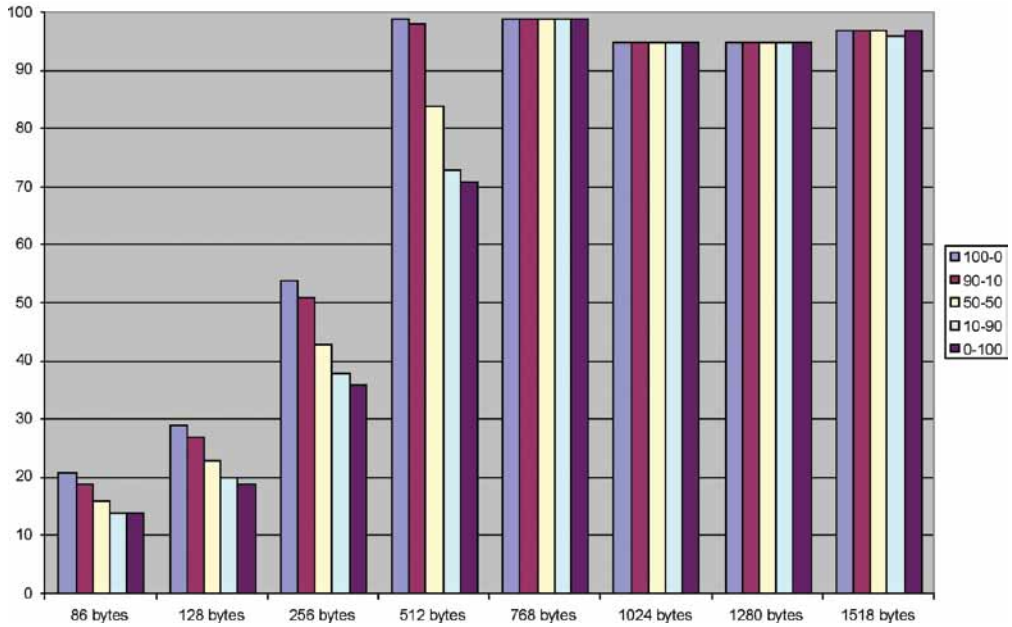

3825 Relative Throughput with Respect to IPv4

The following table shows dual stack and 100% IPv6 relative throughput numbers compared as a percentage of 100% IPv4 numbers. (For example, for 86 bytes the throughput % with no loss is 21%. The 90% IPv4/10% IPv6 throughput % for 86 bytes is 19% - 19 / 21 x 100 = 90.47619048%.) This is the relative performance of the 90/10 traffic stream to the 100/0 traffic stream.

The frame sizes of 86, 128, 256 and 512 bytes showed a degradation of throughput as compared to IPv4. The degradation appeared directly proportional to the IPv6 traffic in the traffic stream. At the larger frame sizes, throughput equaled that of 100% IPv4 traffic, with one exception at the 1518 frame size (177 lost IPv6 frames caused the 1518 frame size data point to drop 1% of line rate, resulting in a less-than-100% value in the table).

3825 Relative Throughput with Respectto IPv4

Frame Size (bytes)	Rel Max Throughput (%) 90% IPv4 / 10% IPv6	Rel Max Throughput (%) 50% IPv4 / 50% IPv6	Rel Max Throughput (%) 10% IPv4 / 90% IPv6	Rel Max Throughput (%) 0% IPv4 / 100% IPv6
86	90.47619048	76.19047619	66.66666667	66.66666667
128	93.10344828	79.31034483	68.96551724	65.51724138
256	94.44444444	79.62962963	70.37037037	66.66666667
512	98.98989899	84.84848485	73.73737374	71.71717172
768	100	100	100	100
1024	100	100	100	100
1280	100	100	100	100
1518	100	100	98.96907216	100

It's important to note these results are based on the 3825's onboard Gigabit Ethernet interfaces. ISR platforms detailed previously in this document (1841 and 2811 ISRs) used 100M interfaces.

Latency

When comparing each frame size individually, there was a slight increase to frame latency with the introduction of IPv6 traffic. The upward trend in latency as the frame size increases mimics the trend that exists for 100% IPv4 traffic. With 10% line rate, the 1024 frame size latency deviated from the upward trend typically seen. At 70%, the 1024 frame size resumed the upward trend. (Note: smaller frame sizes are excluded from the 70% line rate chart due to frame loss.)

3825 Average Latency at the 10% Line Rate

	86 bytes	128 bytes	256 bytes	512 bytes	768 bytes	1024 bytes	1280 bytes	1518 bytes
100-0	32.022335	41.29509	92.60988	176.39263	207.605275	155.93334	214.381115	248.06983
90-10	32.85535	42.1566925	92.84914	176.7137075	208.1128275	156.2955425	213.37796	247.8907825
50-50	28.1821775	43.7124575	93.27667	175.93006	207.737265	156.4117625	214.84754	246.964255
10-90	27.39926	44.873305	83.73825	176.09053	208.2121525	156.786885	219.31296	247.7202875
0-100	27.61137	44.992705	83.98058	175.99254	209.19877	156.699615	214.370085	247.43846

3825 Average Latency at the 70% Line Rate

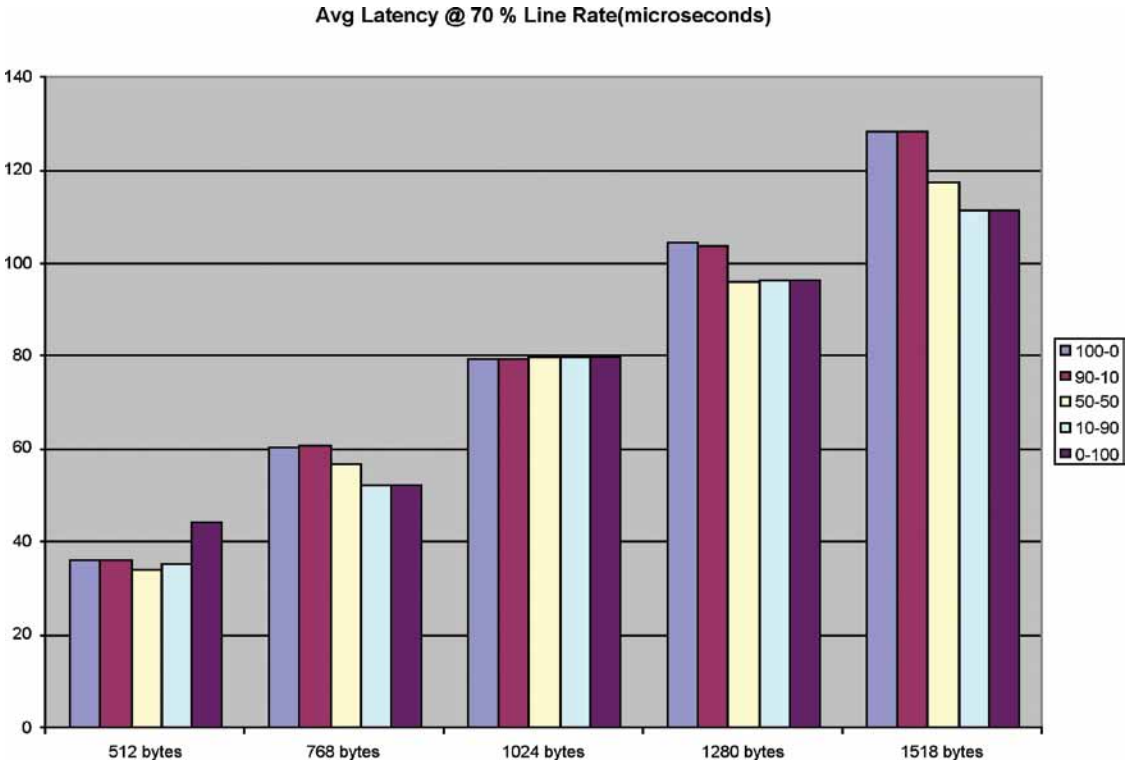

CPU Utilization

At line rates up to 90%, the CPU utilization for the 1518 frame size steadily increased as the line rate increased, with no unexpected spikes in utilization from the various IPv6-mixed traffic loads. The 100% line rate of Fast Ethernet jumped significantly, and at the 100% line rate there was frame loss.

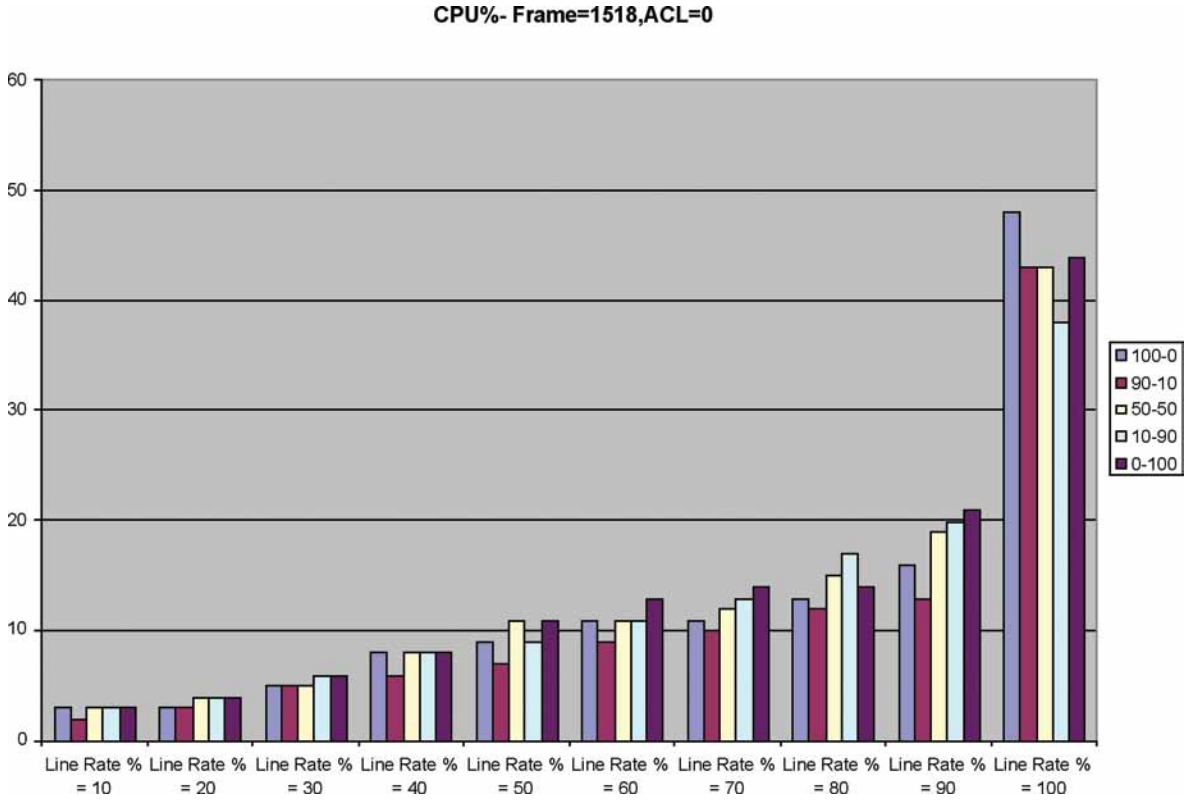

Cisco 7200 Router

Throughput

7200 Throughput (100M bidirectional traffic)

The chart below shows the throughput rate (bidirectional traffic via 2 1000M Ethernet interfaces on the NPE-G2) for the 7200.

7200 Throughput Rate (100M bidirectional traffic)

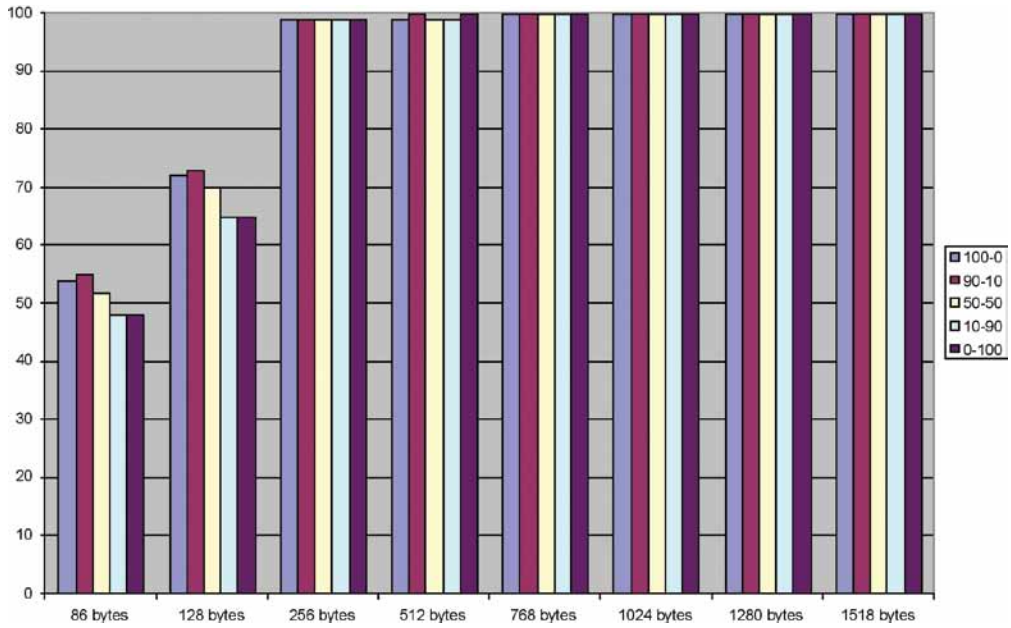

7200 Relative Max Throughput with Respect to IPv4

The table below shows dual stack and 100% IPv6 relative throughput numbers as compared to a percentage of 100% IPv4 numbers. (For example, for 86 bytes the throughput % with no loss is 54%. The 90% IPv4/10% IPv6 throughput % for 86 bytes is 55% - 55/54 x 100 = 101.8518519%.) This is the relative performance of the 90/10 traffic stream to the 100/0 traffic stream.

While these numbers are consistent with what was observed with other ISRs (equal max throughput at higher packet sizes - descending throughput at lower packet sizes) some measurements indicate dual-stack traffic performed better than 100% IPv4 traffic. Due to the strict design of the throughput testing, even a single lost frame meant the line rate failed and was to be decremented to try again.

For the 86, 128 and 512 byte tests, there were observances of as little as one (a single) frame lost for the 100% IPv4 test. This resulted in a slightly lower throughput number and, consequently, improved the perception of superior dual-stack performance.

7200 Relative Throughput with respect to IPv4

Frame Size (bytes)	Rel Max Throughput (%) 90% IPv4 / 10% IPv6	Rel Max Throughput (%) 50% IPv4 / 50% IPv6	Rel Max Throughput (%) 10% IPv4 / 90% IPv6	Rel Max Throughput (%) 0% IPv4 / 100% IPv6
86	101.8518519	96.2962963	88.88888889	88.88888889
128	101.3888889	97.22222222	90.27777778	90.27777778
256	100	100	100	100
512	101.010101	100	100	101.010101
768	100	100	100	100
1024	100	100	100	100
1280	100	100	100	100
1518	100	100	100	100

Latency

The following charts show the average latency at the 40% and 90% line rate for each frame size, respectively. When comparing each frame size individually, there was negligible variation in latency with the introduction of IPv6 traffic. (Note: smaller frame sizes are excluded from the 90% line rate chart due to frame loss.)

7200 Average Latency at the 40% Line Rate

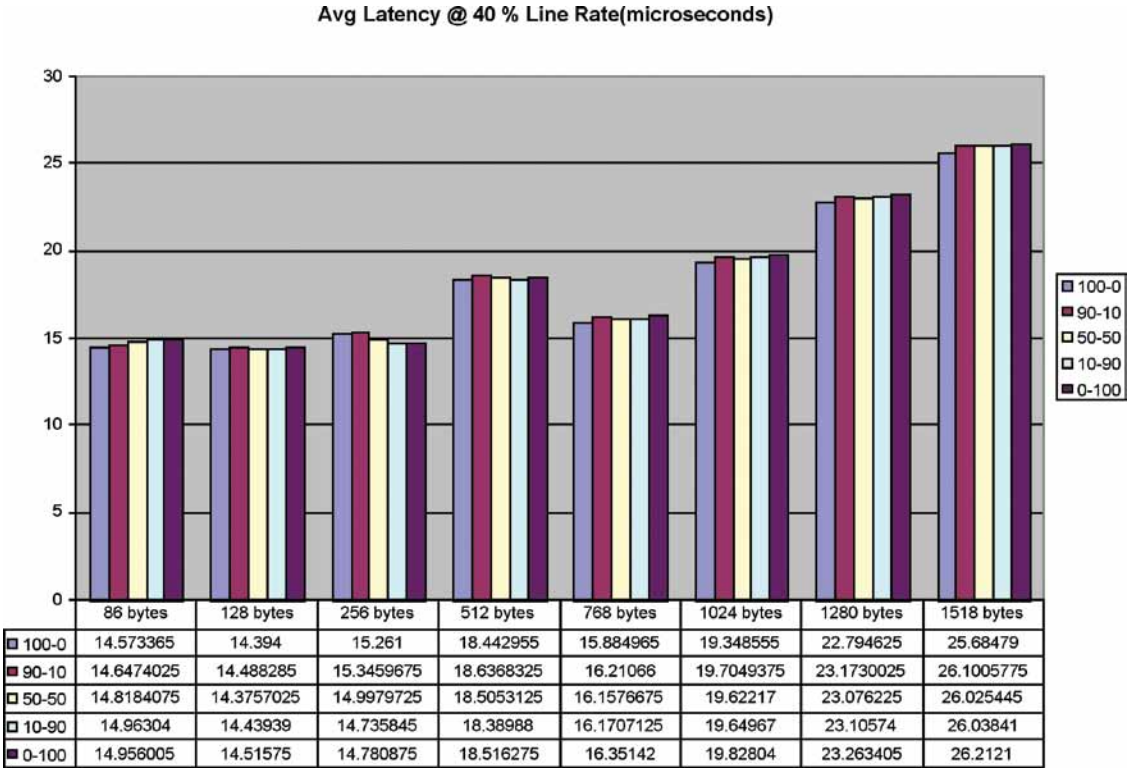

	86 bytes	128 bytes	256 bytes	512 bytes	768 bytes	1024 bytes	1280 bytes	1518 bytes
100-0	14.573365	14.394	15.261	18.442955	15.884965	19.348555	22.794625	25.68479
90-10	14.6474025	14.488285	15.3459675	18.6368325	16.21066	19.7049375	23.1730025	26.1005775
50-50	14.8184075	14.3757025	14.9979725	18.5053125	16.1576675	19.62217	23.076225	26.025445
10-90	14.96304	14.43939	14.735845	18.38988	16.1707125	19.64967	23.10574	26.03841
0-100	14.956005	14.51575	14.780875	18.516275	16.35142	19.82804	23.263405	26.2121

7200 Average Latency at the 90% Line Rate

CPU Utilization

The chart below is representative of frame sizes with a throughput value of 100% for all IPv4/IPv6 traffic ratios. The CPU steadily increased as the line rate increased.

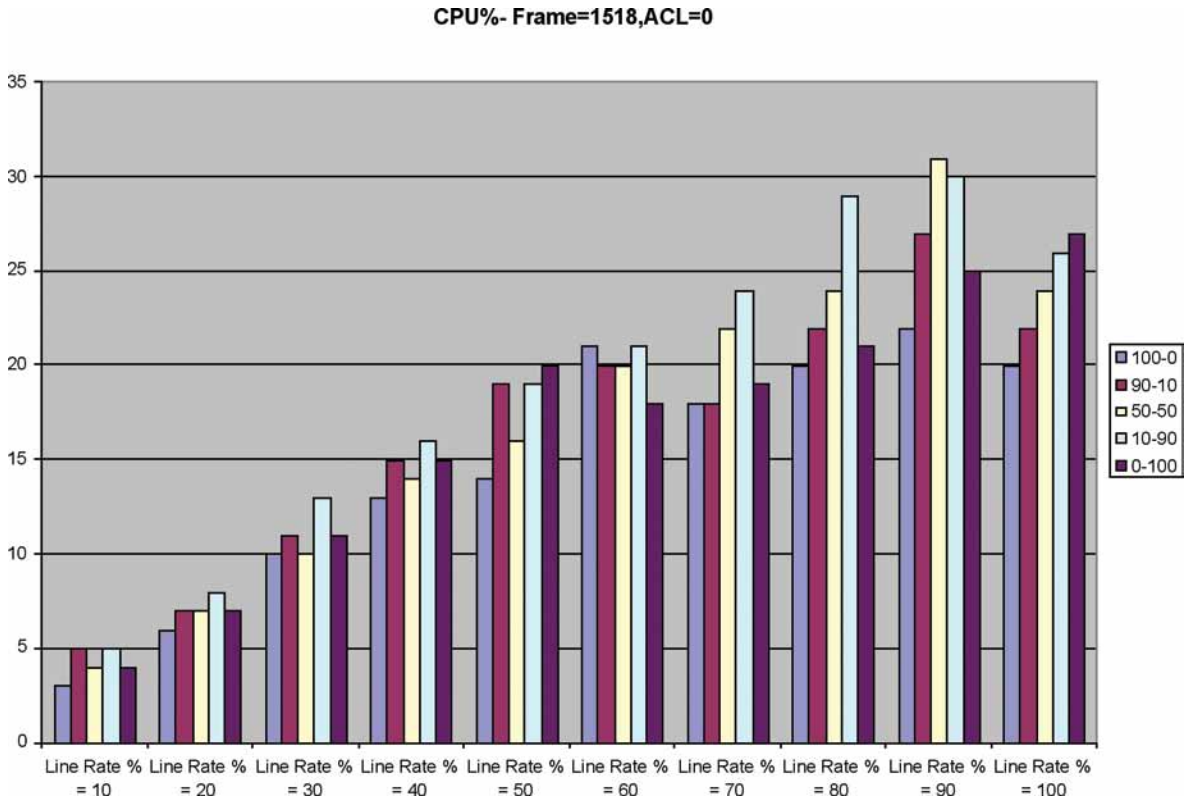

Cisco 7301 Router

Throughput

7301 Throughput (1000M bidirectional traffic)

The chart below shows the throughput rate (bidirectional traffic via 2 1000M Ethernet interface) for the 7301.

7301 Throughput Rate (100M bidirectional traffic)

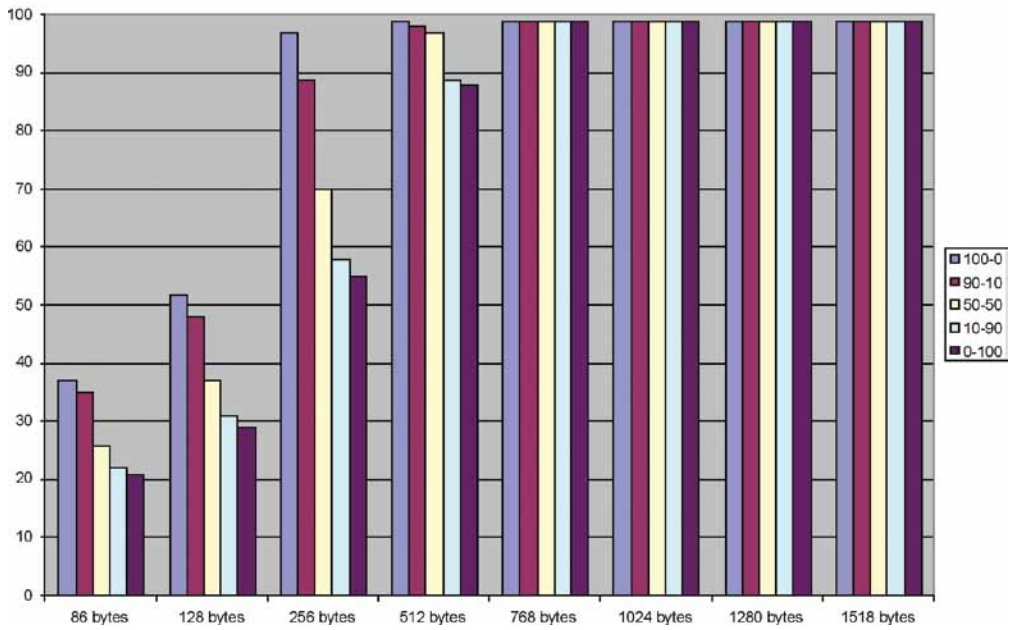

7301 Relative Throughput with Respect to IPv4

The table below shows dual stack and 100% IPv6 throughput numbers compared as a percentage of 100% IPv4 numbers. (For example, for 86 bytes the throughput % with no loss is 37%. The 90% IPv4/10% IPv6 throughput % for 86 bytes is 35% - 35 / 37 x 100 = 94.59459459%.) This is the relative performance of the 90/10 traffic stream to the 100/0 traffic stream.

The frame sizes of 86, 128, 256 and 512 bytes showed a degradation of throughput as compared to IPv4. The degradation appeared directly proportional to the IPv6 traffic in the traffic stream, and appeared equally severe at the 86, 128 and 256 frame sizes. However, it got notably better at the 512 frame size. At the larger frame sizes, throughput equaled that of 100% IPv4 traffic.

7301 Relative Throughput with Respect to IPv4

Frame Size (bytes)	Rel Max Throughput (%) 90% IPv4 / 10% IPv6	Rel Max Throughput (%) 50% IPv4 / 50% IPv6	Rel Max Throughput (%) 10% IPv4 / 90% IPv6	Rel Max Throughput (%) 0% IPv4 / 100% IPv6
86	94.59459459	70.27027027	59.45945946	56.75675676
128	92.30769231	71.15384615	59.61538462	55.76923077
256	91.75257732	72.16494845	59.79381443	56.70103093
512	98.98989899	97.97979798	89.8989899	88.88888889
768	100	100	100	100
1024	100	100	100	100
1280	100	100	100	100
1518	100	100	100	100

Latency

The charts below show that, when comparing each frame size individually, there appeared to be no latency degradation with IPv6 traffic and, in some cases, even showed a slight latency improvement. This trend was seen at both the 20% and 80% line rates. (Note: smaller frame sizes are excluded from the 80% line rate chart due to frame loss.)

7301 Average Latency at the 20% Line Rate

	86 bytes	128 bytes	256 bytes	512 bytes	768 bytes	1024 bytes	1280 bytes	1518 bytes
100-0	17.45318	23.24003	52.630255	99.82861	147.45619	190.259585	236.22067	280.19041
90-10	16.445775	23.6273175	52.94851	99.8027475	144.9369425	191.92854	234.4327725	276.8296175
50-50	15.4273475	21.93628	45.181885	100.447	145.448175	190.84626	235.8773425	277.9478875
10-90	12.8600175	20.235355	42.2536975	100.879895	145.7494025	191.261055	236.40359	275.74361
0-100	14.4714	20.64959	42.519705	100.713235	145.81308	191.46904	236.384575	278.347945

7301 Average Latency at the 80% Line Rate

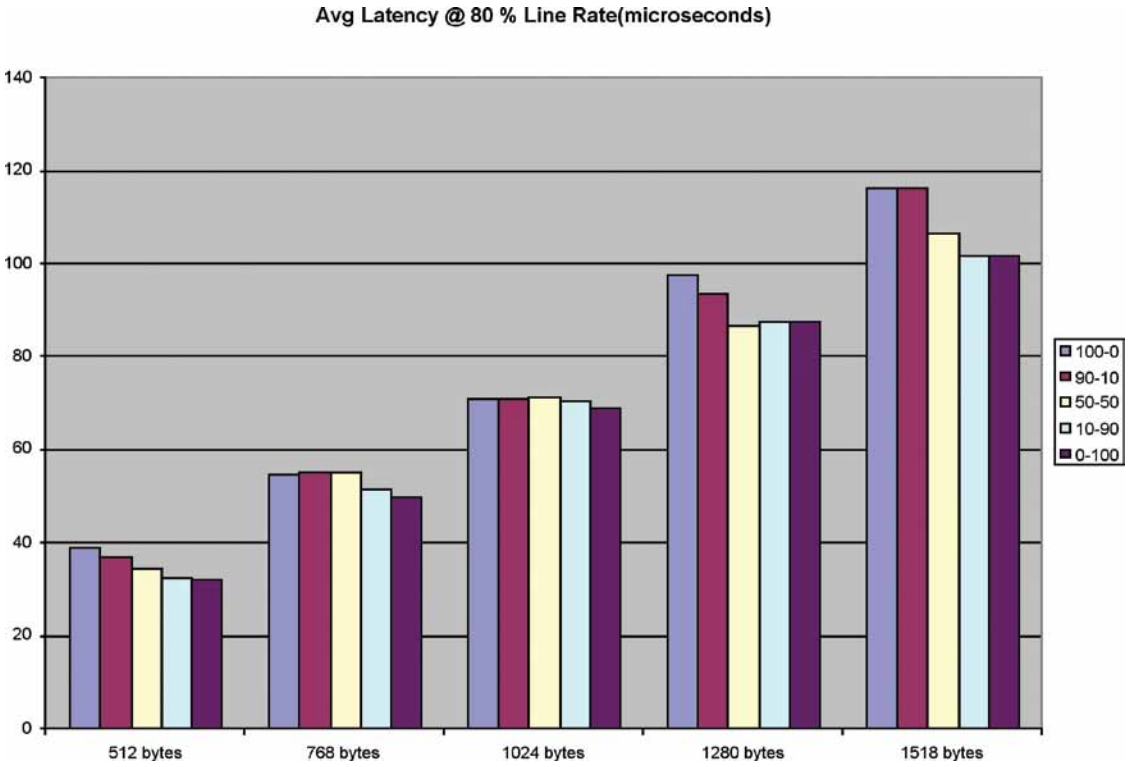

CPU Utilization

The chart below is representative of frame sizes with zero or minimal frame loss at 100% of line rate. The CPU steadily increases as the line rate increases.

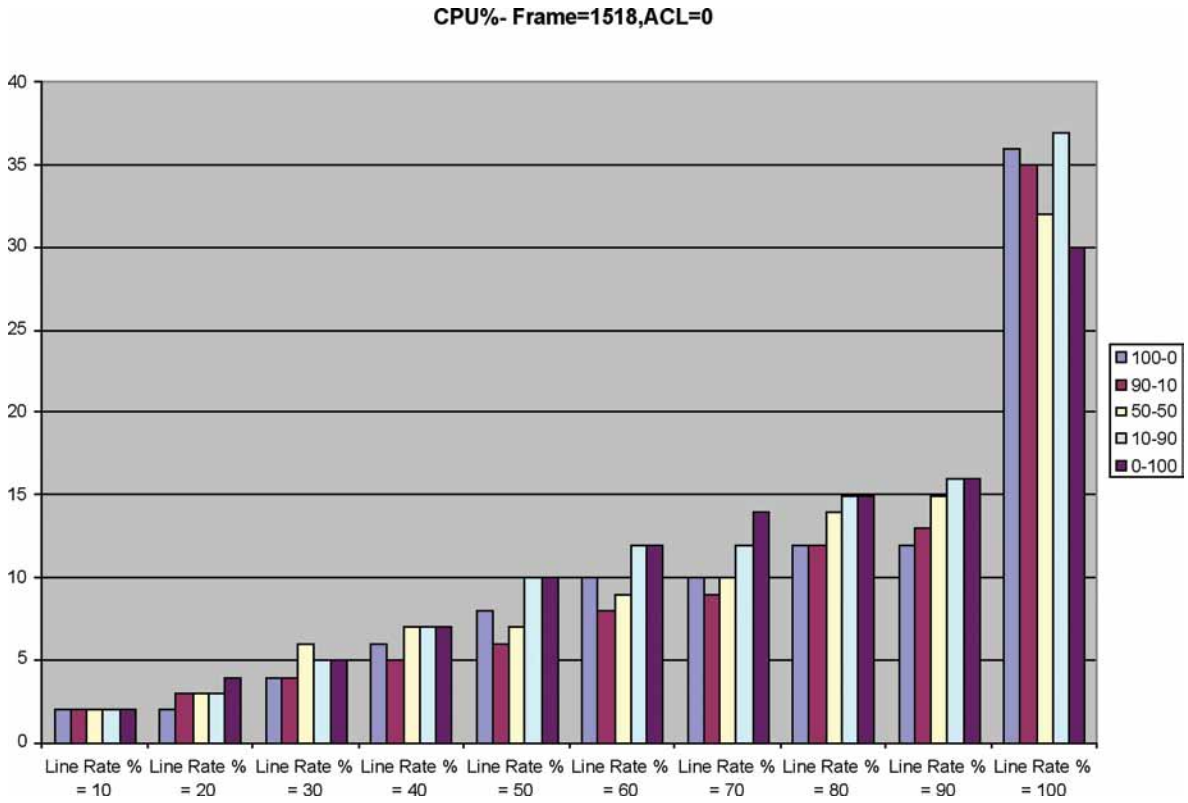

Cisco 7606 Router

Throughput

7606 Throughput (10G bidirectional traffic)

The figure below shows the throughput rate (bidirectional traffic via 2 10G Ethernet SPAs with SIP-600s) for the 7606.

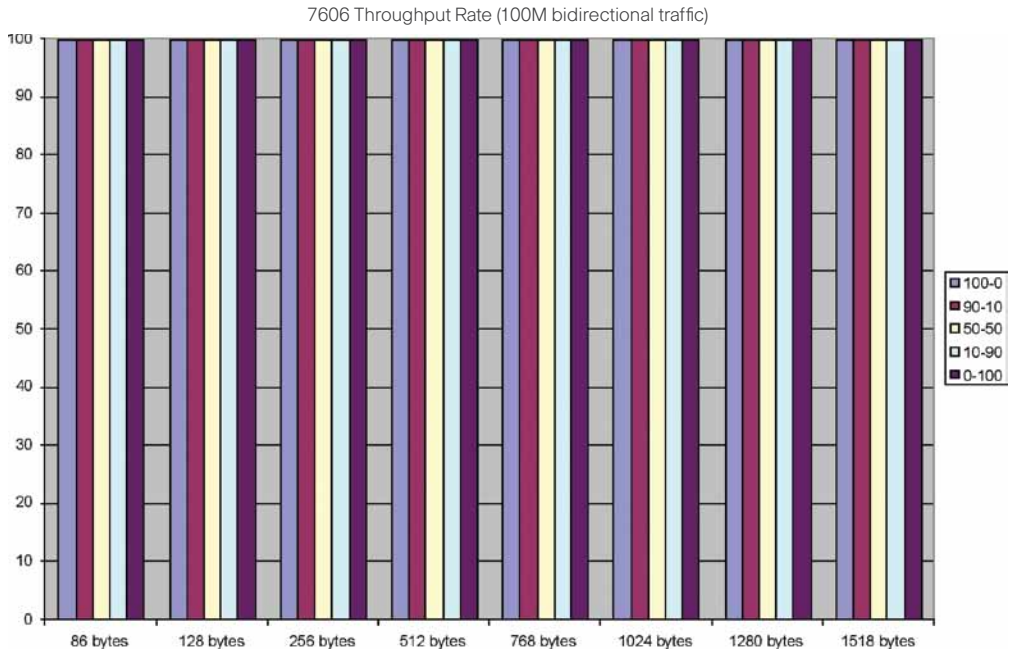

7606 Relative Throughput with Respect to IPv4

The table below shows dual stack and 100% IPv6 relative throughput numbers compared as a percentage of 100% IPv4 numbers. (For example, for 86 bytes the throughput % with no loss is 100%. The 90% IPv4/10% IPv6 throughput % for 86 bytes is 100% - 100/100 x 100 = 100%.) This is the relative performance of the 90/10 traffic stream to the 100/0 traffic stream.

All of these values are 100%, as throughput in all cases was 100% of line rate.

Frame Size (bytes)	Rel Max Throughput (%) 90% IPv4 / 10% IPv6	Rel Max Throughput (%) 50% IPv4 / 50% IPv6	Rel Max Throughput (%) 10% IPv4 / 90% IPv6	Rel Max Throughput (%) 0% IPv4 / 100% IPv6
86	100	100	100	100
128	100	100	100	100
256	100	100	100	100
512	100	100	100	100
768	100	100	100	100
1024	100	100	100	100
1280	100	100	100	100
1518	100	100	100	100

Latency

7606 Average Latency at the 90% Line Rate

When comparing each frame size individually, there was negligible difference in frame latency with the introduction of IPv6 traffic. The very slight upward trend (from 24 microseconds to 28 microseconds) in latency as the frame size increased mimics the existing trend for 100 percent IPv4 traffic.

	86 bytes	128 bytes	256 bytes	512 bytes	768 bytes	1024 bytes	1280 bytes	1518 bytes
100-0	24.0088	24.7092	24.90528	25.353395	25.76872	27.009965	27.06372	28.154975
90-10	24.020535	24.74055	24.82388	25.6615925	25.76212	27.05688	27.08326	28.1215975
50-50	24.037315	24.84784	24.8498175	25.7937175	25.79877	27.0288875	27.1281675	28.23158
10-90	24.20529	24.9048725	24.8333675	25.588365	25.8557875	27.2440825	27.2040925	28.3024525
0-100	24.16114	24.88197	24.86985	25.85178	25.840315	27.290425	27.23211	28.296

CPU Utilization

The 7606 platform switches packets in hardware, so CPU impact is negligible.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**