Performance of Various Throughput, Latency & CPU Utilization on IPv4 & IPv6

Mandeep Singh [1], Vikrant Kumar[2], Navdeep Singh [3]
[1] Associate Professor, [2]Assistant Professor, [3]Assistant Professor

Abstract- In this paper we have discussed the IP protocol designed to replace IPv4, IPv6 quadruples address number of bits from 32 bits (in IPv4) to 128 bits addressable nodes, which will provides more unique IP addresses .These tests were explore the throughput and latency measurement of various IPv4 and IPv6 traffic mixes.

I. INTRODUCTION

Internet Protocol version 6 (IPv6),the protocol communications that provides an location system for computers networks Internet.IPv6 is replacing IPv4, which carries the vast Internet traffic. Each device on the Internet have assigned an IP address in order to communicate with other devices. With the ever-increasing number of new devices being connected to the Internet, the need arose for more addresses than IPv4. IPv6 uses a 128-bit address, allowing 2^{128}, or approximately 3.4×10^{38} addresses, or more than 7.9×10^{28} times as many as IPv4, which uses 32-bit addresses. . The two protocols are not designed to be complicating the transition to IPv6.On the Internet, data is transmitted in the form of network packets. IPv6 specifies a new packet format, designed to minimize packet header processing by routers.[3][11] Because the headers of IPv4 packets and IPv6 packets are significantly different, the two protocols are not interoperable. However, in most respects, IPv6 is a conservative extension of IPv4. IPv6 over IPv4 is its larger address space. While these numbers are large, it wasn't the intent of the designers of the IPv6 address space to assure geographical saturation with usable addresses. Rather, the longer addresses simplify allocation of addresses, enable efficient route aggregation, and allow implementation of special addressing features. In IPv4, complex Classless Inter-Domain Routing (CIDR) methods were developed to make the best use of the small address space. The standard size of a subnet in IPv6 is 2^{64} addresses, the square of the size of the entire IPv4 address space. Thus, actual address space utilization rates will be small in IPv6, but network management and routing efficiency is improved by the large subnet space and hierarchical route aggregation. In IPv4 it is very difficult for an organization to get even one globally routable multicast group assignment, and the implementation of inter-domain solutions is very arcane.[19] Unicast address assignments by a local Internet registry for IPv6 have at least a 64-bit routing prefix, yielding the smallest subnet size available in IPv6 (also 64 bits). With such an assignment it is possible to embed the unicast address prefix into the IPv6 multicast address format, while still providing a 32-bit block, the least significant bits of the address, or approximately 4.2 billion multicast group identifiers. Thus each user of an IPv6 subnet automatically has available a set of globally routable source-specific multicast groups for multicast applications sources of auto configuration information, such as router and prefix advertisements. Stateless configuration of routers can be achieved with a special router renumbering protocol.[21]

II. METHODOLOGY

Automated Multiple Platform Testing

A. Testing Approach

The purpose of testing assess basic traffic throughput, frame loss and latency variances. Data was taken running IPv4-only traffic, and then tests were run with the DUT having various percentages of IPv4 and IPv6 traffic running simultaneously
The test setup used (STC) to inject traffic 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 DUTt

B. Throughput Test

The highest packet rate that can be switched through a given interface and for a given packet size without packet loss. If a single frame is dropped, the test fails Measures throughput by sending a series of frames with particular source and destination MAC addresses to the DUT. Frames are sent by oe of the Spirent Test Center test ports and are intended to be received by a second test port.

C. Latency Test

Determines the delay (latency) imposed on forwarded traffic.As the test runs, the transmitting test port sends a burst of frames, at a user-specified frame size
(Receive Timestamp) – (Transmit Timestamp) = Latency

D. Test Variables

Throughput, frame loss, latency was measured at varying test conditions, including every combination of the following:
IPv4/IPv6 Ratio
Frame Sizes
Line Rate
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)
• 78 bytes (smallest IPv6 frame with no Layer 4 header)
3. Only Ethernet interfaces were tested

III. RESULTS

7200 Router

IV. 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.

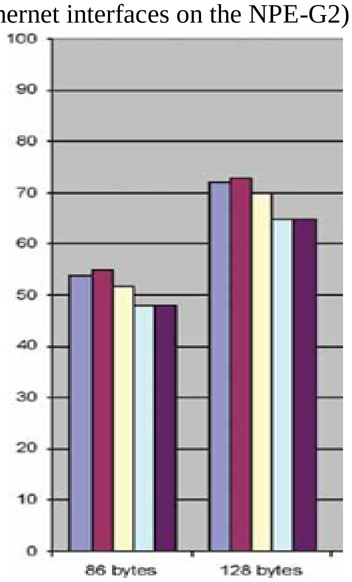

V. LATENCY

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.)

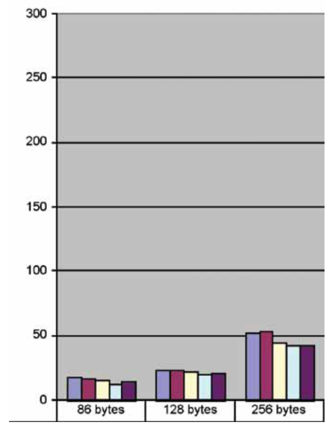

VI. CPU UTILIZATION

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.

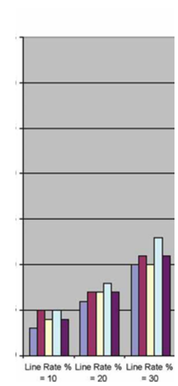

7301 ROUTER

VII. THROUGHPUT

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

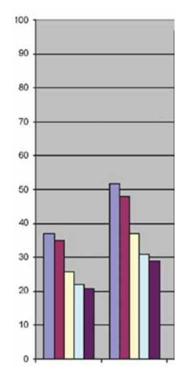

VIII. LATENCY

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
7301 Average Latency at the 80% Line Rate

IX. 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.

X. 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.

REFERENCES

[1] David Frost (20 April 2011). "Ipv6 traffic volumes going backwards". iTWire. Retrieved 19 February 2012.

[2] Roberts, Phil (24 September 2013). "IPv6 Deployment Hits 2%, Keeps Growing". Internet Society. Retrieved 27 September 2013.

[3] RFC 2460, Internet Protocol, Version 6 (IPv6) Specification, S. Deering, R. Hinden (December 1998)

[4] Google IPv6 Conference 2008: What will the IPv6 Internet look like?. Event occurs at 13:35.

[5] RFC 1752 The Recommendation for the IP Next Generation Protocol, S. Bradner, A. Mankin, January 1995.

[6] Rashid, Fahmida. "IPv4 Address Exhaustion Not Instant Cause for Concern with IPv6 in Wings". eWeek. Retrieved 23 June 2012.

[7] Ward, Mark. "Europe hits old internet address limits". BBC. Retrieved 15 September 2012.

[8] RFC 1550, IP: Next Generation (IPng) White Paper Solicitation, S. Bradner, A. Mankin (December 1993)

[9] "History of the IPng Effort". Sun.[dead link]

[10] The Internet Engineering Task Force.

[11] [b] RFC 1726, Technical Criteria for Choosing IP The Next Generation (IPng), Partridge C., Kastenholz F. (December 1994)

[12] "U.S. Census Bureau". Census.gov. Retrieved 20 January 2012.

[13] "Moving to IPv6: Now for the hard part (FAQ)". Deep Tech. CNET News. Retrieved 3 February 2011.

[14] RFC 2071, Network Renumbering Overview: Why would I want it and what is it anyway?, P. Ferguson, H. Berkowitz (January 1997)

[15] RFC 2072, Router Renumbering Guide, H. Berkowitz (January 1997)

[16] [b] RFC 4862, IPv6 Stateless Address Autoconfiguration, S. Thomson, T. Narten, T. Jinmei (September 2007)

[17] RFC 1112, Host extensions for IP multicasting, S. Deering (August 1989)

[18] RFC 3956, Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address, P. Savola, B. Haberman (November 2004)

[19] RFC 2908, The Internet Multicast Address Allocation Architecture, D. Thaler, M. Handley, D. Estrin (September 2000)

[20] RFC 3306, Unicast-Prefix-based IPv6 Multicast Addresses, B. Haberman, D. Thaler (August 2002)