

Performance Comparison of IPv4 and IPv6 in Peer-Peer and Client Server Local Area Networks

Samad S. Kolahi, Burjiz K. Soorty, Navneet Chand and Zhang Qu

UNITEC New Zealand

skolahi@unitec.ac.nz

burjizsoorty@hotmail.com

Abstract-In this paper, the performance of IPv4 and IPv6 are compared for both Client-Server and Peer-Peer networks. For both networks, IPv4 produced higher bandwidth for TCP protocol. For UDP, IPv4 and IPv6 showed insignificant bandwidth differences except for packet size of 384 bytes where IPv4 had better performance in client-server environment.

I. INTRODUCTION

The inevitable growth in networks has significantly increased the number of IP addresses required that has in-turn led to transition into the IPv6 era of addressing. As exhaustion of IPv4 addresses materialize, this inadequacy is successfully overcome by IPv6 which supports a total of 2^{128} addresses as opposed to 2^{32} for IPv4. As this transition grows actively closer each year, newer operating systems aim to put a higher emphasis into enhancing IPv6 performance. Several works have been carried out on evaluating IPv4 and IPv6 that have shown their performance to largely vary depending on the operating system used on the network [1].

In 1998, Draves et al. [2] conducted a throughput evaluation of IPv6 using Windows NT over a network where two clients were directly connected through a crossover cable. Their study ran TCP throughput tests for both IPv4 and IPv6 over 10 Mb/s and 100 Mb/s Ethernet. The analysis however did not consider any other additional protocols such as UDP or any parameters such as different packet sizes that could potentially affect network performance. Their findings showed a 1.9% downgrade on IPv6 performance over the IPv4.

In 2000, Ariga et al. [3] conducted a performance evaluation of data transmission through IPv6 and IPv4 using IPSec on a network implementing the Unix-like open source operating system called FreeBSD. The authors took into account the TCP and UDP protocols however did not consider the different packet sizes. Their studies concluded that the TCP/UDP throughput for IPv6 was almost equivalent to that of IPv4.

In 2003, Zeadally and Raicu [4] conducted a performance evaluation of IPv4 and IPv6 on Windows 2000 and Solaris 8. Two identical workstations were connected using a point-to-point link in-order to eliminate most variables such as router processing, from the experiments. This study included an additional metric of RTT (Round Trip Time) along with the TCP/UDP throughput calculated. Their analysis concluded that the performance of IPv4 and IPv6 varied significantly, especially a notably large difference in throughput for packet

sizes smaller than 256 bytes in Solaris. The difference, an apparent three times higher throughput for IPv4 mainly due to the rather large increase in socket-creation and connection time for IPv6 because of its high overhead [2].

In 2004, Zeadally et al. [1] conducted another empirical performance evaluation of IPv4 and IPv6 protocol stack. They used Linux as the operating system and compared the results to their 2003 study on Windows 2000 and Solaris 8. Throughput and RTT were measured for TCP and UDP protocols. Their evaluation concluded that both IPv4 and IPv6 on Linux outperformed Windows 2000 & Solaris 8 for both the metrics used, the latter having a lower margin of difference compared to Windows 2000.

In 2006, Mohammed et al. [5] carried out an evaluation on the performance of IPv6 on two different operating systems, namely, Windows 2003 and Linux based Red Hat 9. The study implemented three different test beds each representing two different environments. The first, a "Normal View Environment" where the payload size was less than the MTU (Maximum Transmission Unit) and the second, a "Global View Environment" where the payload size was greater than the MTU. The MTU can best be defined as the highest amount of data that can be transferred in one physical frame on the network. If a packet that has a smaller MTU than the packet's frame length is sent, fragmentation would occur. Their studies concluded that the performance of IPv6 was far better in Red Hat 9 than in Windows Server 2003 as the latter fragmented packets larger than 1440 bytes as opposed to the former where fragmentation occurred at 16384 bytes.

The authors in [6] carried out a performance comparison of IPv4 and IPv6 using Windows XP operating system over a network where two clients were directly connected through a standard Category 5e crossover cable. The study performed an empirical evaluation on IPv4 and IPv6 for TCP and UDP protocols by measuring throughput and RTT for various packet sizes. Their research concluded that IPv6 performed better for TCP and IPv4 performed better for UDP with Windows XP. It also proved that the performance of IPv6 was significantly better on Windows XP than on Windows Server 2000 [4, 6].

At the time of writing, there has been no work published on comparison of the performance of IPv4 and IPv6 on two commonly implemented network environments, one running a client-server environment of Windows Vista (plus Service Pack One) with Windows Server 2008 and the other running a peer

to peer network of Windows Vista (plus Service Pack One). The former being commonly implemented in large corporate networks and the latter a popular setup in small branch offices. The IPv4 and IPv6 evaluation incorporates different parameters that affect network performance. Those parameters study the throughput and round trip time of TCP and UDP on a range of packet sizes varying from 128 to 1408 bytes.

The organization of this paper is as follows. In the next section the network setup is discussed. Section three covers information regarding the data generating and traffic measurement tool. Section four covers the results and the last sections include the conclusion, future works and acknowledgments followed by the references.

II. NETWORK SETUP

The hardware benchmark comprised of an Intel® Core™ 2 Duo 6300 1.87 GHz processor with 2.00 GB RAM and an Intel Pro/100 S Desktop Adapter NIC and a Western Digital Caviar SE 160 GB hard-drive on the two workstations.

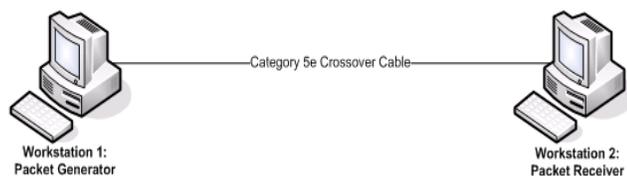


Figure 1: Network Test-Bed

The proposed network test-bed was setup through a direct connection via standard fast Ethernet Category 5e cabling between two workstations, as shown in figure 1. This was done in order to calculate the raw throughput and RTT without the use of a hub, switch or a router that could create latency and degrade the actual speed of the network. It was also done so as to maintain consistency with similar research shown in the past.

A client and a server were connected through a Category 5e crossover cable in TIA/EIA 568-B wiring thereby maintaining global industrial networking standards for use over 100Base-TX. Workstation one operated as the server which generated packets and Workstation two operated as the client assigned to receiving the data.

III. DATA GENERATION AND TRAFFIC MEASUREMENT TOOL

IP Traffic [7] was selected as the traffic generating and measurement tool for its compatibility with Windows, and for its powerful analysis of a wide range of quality of service parameters to acquire accurate results. IP Traffic has extensively been used for researches including performance evaluation of network security [8] and impact of encryption effects on network performance [9].

IV. RESULTS

The TCP and UDP throughput and RTT were measured for IPv4 and IPv6 for various packet sizes. The range of packet sizes varied from 128 to 1408 bytes over the two network environments.

This evaluation methodology comprised of performing 40 test runs for every protocol individually (TCP and UDP) and for each specific packet size (128 to 1408 bytes) in-order to get rid of any inconsistencies shown in the results. One run included sending 1 million packets of one particular packet size and protocol. The results were then averaged and a standard deviation of results was recorded.

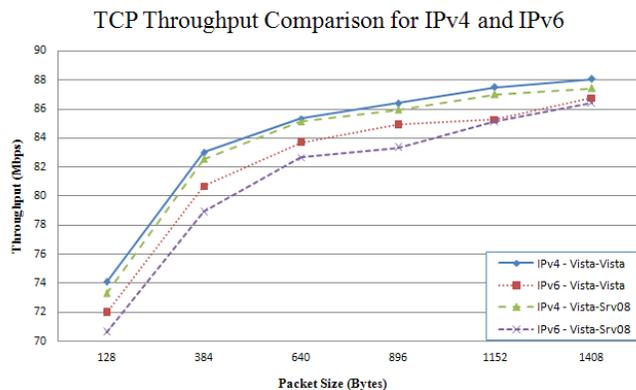


Figure 2: TCP Throughput Comparison for IPv4 and IPv6 on Peer-to-Peer vs. Client-Server

Figure 2 shows the TCP throughput of IPv4 and IPv6 on a peer-to-peer network and client server. On client-server, throughput with IPv4 is the highest throughput all packet sizes however IPv6 throughput is largely the lowest on most packet sizes. On the peer-to-peer network, throughput with IPv4 again is higher on all packet sizes than the throughput observed with IPv6.

Comparing the client-server to the peer to peer network, as depicted, the performance for IPv4 and IPv6 both seem slightly higher on the peer to peer network. This is mainly due to the absence of server applications such as DNS, Active Directory Domain Services and DHCP that utilize a small margin of bandwidth thereby compromising overall throughput rate. Throughput on a peer to peer network however would experience a drop in bandwidth with the increase in nodes, something which wouldn't occur on a client-server environment [4].

Figure 3 shows the UDP throughput of IPv4 and IPv6 on a peer-to-peer network and client-server network. On both the network environments the difference in IPv4 and IPv6 throughput remain largely insignificant with the highest difference at packet size 384 bytes where throughput for IPv4 on client-server is the highest and throughput for IPv6 the lowest. Results reveal IPv4 to perform better than IPv6 for client server however that difference is by far insignificant as depicted in figure 3.

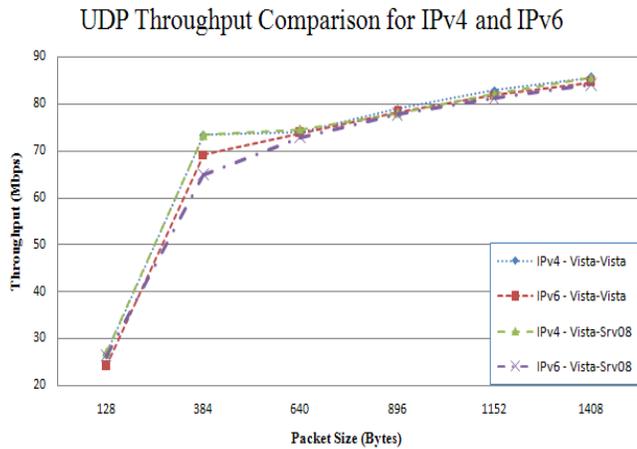


Figure 3: UDP Throughput Comparison for IPv4 and IPv6 on Peer-to-Peer vs. Client-Server.

The Peer to Peer results showed TCP throughput to be higher on IPv4 than on IPv6 with a maximum difference of 2.34 Mbps at a 2.74% increase in IPv4 for packet size 1152 bytes (87.48 Mbps for IPv4 vs. 85.27 Mbps for IPv6). UDP throughput was slightly higher on IPv4 than on IPv6 with a maximum difference of 4.16 Mbps at a 6.01% increase in IPv4 for packet size 384 bytes (73.29 Mbps for IPv4 vs. 69.13 Mbps for IPv6).

The Client Server results showed TCP throughput to be higher on IPv4 with a maximum difference of 3.65 Mbps at a 4.62% increase in IPv4 for packet size 384 bytes (82.56 Mbps for IPv4 vs. 78.91 Mbps for IPv6). UDP results showed an insignificant difference on IPv4 and IPv6 however UDP throughput was marginally higher on IPv4 with a maximum difference of 8.39 Mbps at 12.92% increase in IPv4 for packet size 384 bytes (73.31 Mbps for IPv4 vs. 64.92 Mbps for IPv6).

Comparing the two networks, IPv4 and IPv6, both resulted in better TCP throughput on the peer to peer network than on the client-server. For the lowest packet size of 128 bytes the peer to peer network showed a 0.79 Mbps at 1.07% increase in IPv4 throughput from the client-server network (74.1 Mbps for peer-peer vs. 73.31 Mbps for client-server) and a 1.34 Mbps at 1.89% increase in IPv6 throughput from client-server network (71.96 Mbps for peer-peer vs. 70.62 Mbps for client server). For the highest packet size of 1408 bytes the peer to peer network showed a 0.61 Mbps at 0.69% increase in IPv4 throughput (88.05 Mbps for peer-peer vs. 87.44 Mbps for client server) and a 0.38 Mbps at 1.89% increase in IPv6 throughput from the client-server network (86.77 Mbps for peer-peer vs. 86.39 Mbps for client server). In terms of UDP throughput for IPv6, the Windows peer to peer network performed marginally better than the client-server network with a 4.21 Mbps at a 6.48% maximum increase noticed at the packet size of 384 bytes (69.13 Mbps for peer-peer vs. 64.92 Mbps for client-server). UDP throughput for IPv4 resulted with the peer to peer network performing approximately equal to the throughput produced for the client-server network with a 1.19 Mbps at a 1.62% maximum increase noticed at the packet size of 384 bytes

(73.29 Mbps for peer-peer vs. 73.31 Mbps for client-server). Overall UDP throughput however was by far insignificant on both networks.

The gain in TCP and UDP throughput as packet size increase is likely to the amortization of overheads associated with larger user packet sizes (larger user payloads) [3]. Both network environments recorded IPv4 throughput to be higher than IPv6 for TCP. This is mainly due to IPv6 having a relatively larger overhead and increased address space to IPv4 [3]. Similarly because UDP does not use any form of error correction and due to it being a connectionless protocol it thereby explains the insignificant difference in throughput noticed between IPv4 and IPv6 on UDP.

The standard deviation for the above throughput results are recorded in Table 1.

TABLE I
STANDARD DEVIATION FOR THROUGHPUT

Packet size Bytes	Windows Peer-Peer				Windows Client-Server			
	TCP		UDP		TCP		UDP	
	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6
128	1.97	3.85	0.43	0.46	2.55	2.48	0.43	0.52
384	1.35	2.13	1.06	1.5	0.91	1.09	1.3	0.78
640	1.13	1.7	0.54	0.71	1.11	0.8	0.73	0.46
896	0.99	1.99	0.97	0.34	0.69	0.56	0.73	0.6
1152	0.81	0.64	0.67	0.55	0.56	0.45	0.29	0.74
1408	0.79	0.4	0.38	0.38	0.42	0.43	0.65	0.53

Figure 4 shows the TCP Round Trip Time for IPv4 and IPv6 on the peer-to-peer network and client-server networks. The TCP results show a gain in delay for IPv4 and IPv6 with the increase in each packet size.

As seen in the figure, IPv6 has a slightly higher delay than IPv4 on the client-server network whereas IPv4 has a higher delay than IPv6 on the peer to peer network. The highest point of difference between IPv4 and IPv6 for client-server network stands at the packet size of 1408 bytes where IPv4 has a lower delay rate by 1.03 ms at 3.55% compared to IPv6 (28.97 ms for IPv4 vs. 30 ms for IPv6). The highest point of difference between IPv4 and IPv6 for peer-to-peer network stands at the packet size of 640 bytes where IPv6 has a lower delay rate by 2.27 ms at 6.03% compared to IPv4 (37.6 ms for IPv4 vs. 35.3 ms for IPv6).

RTT Comparison of TCP for IPv4 and IPv6

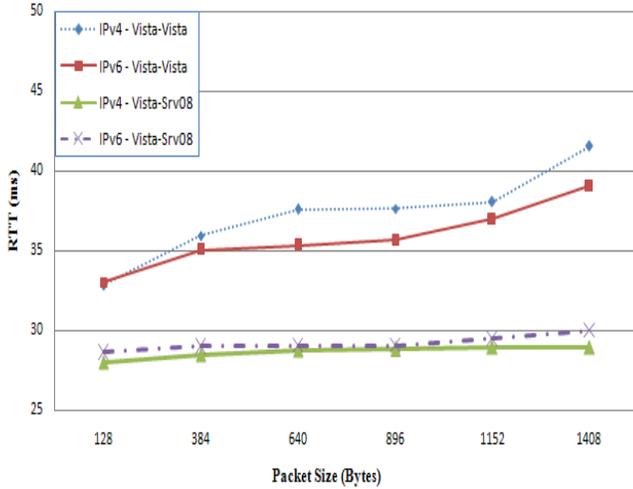


Figure 4: RTT Comparison of TCP for IPv4 and IPv6 on Peer-to-Peer vs. Client-Server

Comparing the two client-server environments, as depicted, the RTT for TCP is lowest client-server for IPv4. IPv4 resulted in a significantly lower delay rate on the client-server network than on the peer to peer network with the maximum difference of 12.6 ms at 30.31% noticed at the packet size of 1408 bytes (41.57 ms for peer-peer vs. 28.97 ms for client-server). IPv6 also resulted in a significantly lower delay rate on the client-server network than on the peer to peer network with the maximum difference of 9.03 ms at 23.13% noticed at the packet size of 1408 bytes (39.03 ms for peer-peer vs. 30 ms for client-server).

Figure 5 shows the UDP Round Trip Time for IPv4 and IPv6 on client-server network and peer-to-peer network. The UDP results once again portray a gain in delay for IPv4 and IPv6 as each packet size increases methodically. The UDP Round Trip Time for IPv4 and IPv6 appear significantly close. IPv6 however, interestingly enough has a marginally lower delay rate this time as by average RTT is slightly lower on IPv6 than on IPv4 for packet sizes 384-640.

The highest point of difference stands at packet size 384 bytes where IPv6 has a lower delay rate by 12.57 ms at 92.56% compared to IPv4 on the client-server network (13.58 ms for IPv4 vs. 1.18 ms for IPv6) and a lower delay rate by 11.85 ms at 83.86% compared to IPv4 on the peer to peer network (14.43 ms for IPv4 vs. 2.28 ms for IPv6).

Comparing the network environments, as depicted, the RTT for UDP is lowest on Windows client-server with IPv6 however this difference is largely insignificant. Similar to the RTT for TCP IPv4 where client server performed better, RTT results for UDP show IPv4 having a slightly lower delay rate on client-server than on the peer-to-peer network.

RTT Comparison of UDP for IPv4 and IPv6

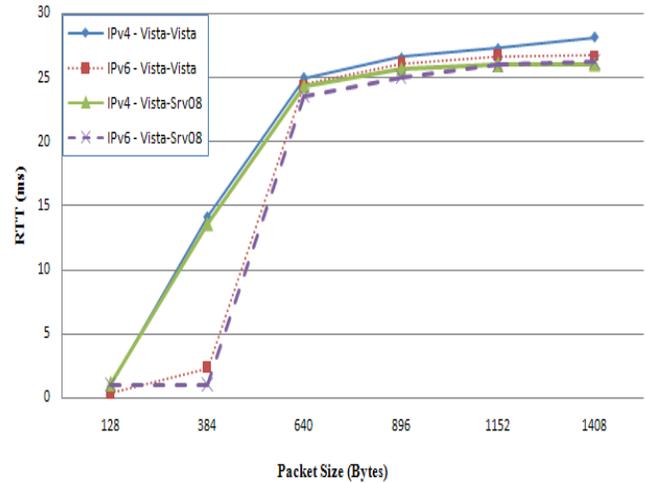


Figure 5: RTT Comparison of UDP for IPv4 and IPv6 on Peer-to-Peer vs. Client-Server

Similarly the RTT for TCP IPv6 client-server performed better, RTT for UDP IPv6 also showed a slightly lower delay rate on client-server than on the peer-to-peer network and this margin of difference decreased as packet sizes increased.

The standard deviation for the above RTT results is recorded in Table 2:

TABLE II
STANDARD DEVIATION FOR RTT

The gain in delay with the increase in each packet size is

Packet size (Bytes)	Windows Peer-Peer				Windows Client-Server			
	TCP		UDP		TCP		UDP	
	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6
128	1.14	2.78	0.08	0.15	0.69	0.73	0.09	0.04
384	2.99	1.95	1.33	1.26	0.42	0.04	1.91	0
640	2.67	3.33	0.4	0.73	0.25	0.09	0.38	0.51
896	3.07	3.57	0.18	0.5	0.15	0.43	0.17	0.05
1152	3.91	0.37	0.24	0.05	0.12	0	0.02	0.05
1408	3.96	0.15	0.33	0.11	0.5	0.07	0.01	0.09

likely due to the amortization of overheads associated with larger packet sizes (larger user payloads) thus higher transmission time. RTT was slightly lower with TCP in IPv6. This is possibly since IPv6 has a simplified although relatively larger header compared to IPv4 and is likely also due to the much increased throughput produced on IPv4. Similarly because UDP does not use any form of error correction and due to it being a connectionless protocol it thereby explains the insignificant difference noticed between IPv4 and IPv6 on UDP.

V. CONCLUSION

For both the Peer-Peer and client-server networks studied, the results indicated that TCP throughput is higher in IPv4 than on IPv6. In Peer-Peer the maximum difference was 2.3Mbps for packet size of 1152bytes. For Client Server the maximum difference was 3.6 Mbps at packet size 384 bytes. For both Peer-Peer and Client Server, IPv4 and IPv6 had same performance for UDP protocol with some differences for Client-Server at packet size of 384 Bytes where IPv4 performed better than IPv6.

VI. FUTURE WORKS

Future works include performing and comparing an empirical evaluation of IPv4 and IPv6 over Gigabit Ethernet LAN's. Forthcoming projects also include further performance evaluation of IPv4 and IPv6 over Category 5e vs. Category 6 Cabling Systems.

ACKNOWLEDGMENT

The authors would like to thank UNITEC Institute of Technology for funding the research team and providing the inventory needed. They are also thankful to Chris Manford for his valuable input and suggestions that helped improve this project.

REFERENCES

- [1] S. Zeadally, R. Wasseem, and I. Raicu, "Comparison of end-system IPv6 protocol stacks," IEE Proceedings Communications, vol. 151, no. 3, 2004, pp. 238-242.
- [2] R.P. Draves, A. Mankin, and B.D. Zill, "Implementing IPv6 for Windows NT," Proc. Proceedings, of the 2nd USENIX Windows NT Symposium, 1998.
- [3] S. Ariga, K. Nagahashi, M. Minami, H. Esaki, and J. Murai, "Performance evaluation of data transmission using IPSec over IPv6 networks," Proc. Proceedings of the INET, 2000.
- [4] S. Zeadally and I. Raicu, "Evaluating IPv6 on Windows and Solaris," Internet Computing, IEEE, vol. 7, no. 3, 2003, pp. 51-57.
- [5] S.S. Mohamed, A.Y.M. Abusin, and D. Chieng, "Evaluation of IPv6 and comparative study with different operating systems," Proc. Third International Conference on Information Technology and Applications, 2005, pp. 665-670 vol.662.
- [6] S. Narayan, S. S. Kolahi, Y. Sonanto, D. T. Nguyen, and P. Mani, "Performance comparison of IPv4 and IPv6 on various windows operating systems," Proc. 11th International Conference on Computer and Information Technology, 2008, pp. 663-668.
- [7] ZTI Telecom, "IP Traffic - test & measure," <http://www.zti-telecom.com>.
- [8] N. Baghaei and R. Hunt, "IEEE 802.11 wireless LAN security performance using multiple clients," Proc. Proceedings, The 12th IEEE International Conference on Networks, 2004, pp. 299-303 vol.291.
- [9] B. Ezedin, B. Mohammed, A. Amal, S. Hanadi Al, K. Huda, and M. Meera Al, "Impact of Security on the Performance of Wireless-Local Area Networks," Proc. Innovations in Information Technology, 2006, pp. 1-5.