



Comparison Between IPv4 to IPv6 Transition Techniques by Edwin Cordeiro is licensed under a Creative Commons Attribution 4.0 International License.

EXPERIENCE REPORT

Comparison Between IPv4 to IPv6 Transition Techniques

Edwin S Cordeiro^{1*†}, Rodrigo M Carnier² and Wagner L Zucchi²

*Correspondence:

edwin@scordeiro.net

¹Institute for Technological Research, São Paulo, Av. Prof. Almeida Prado 532, Cid. Universitária - Butantã, 05508-901, São Paulo, Brazil

Full list of author information is available at the end of the article

[†]Principal corresponding author

Abstract

The IPv4 addresses exhaustion demands a protocol transition from IPv4 to IPv6. The original transition technique, the dual stack, is not widely deployed yet and it demanded the creation of new transition techniques to extend the transition period.

This work makes an experimental comparison of techniques that use dual stack with a limited IPv4 address. This limited address might be a RFC 1918 address with a NAT at the Internet Service Provider (ISP) gateway, also known as Carrier Grade NAT (CGN), or an Address Plus Port (A+P) shared IPv4 address. The chosen techniques also consider an IPv6 only ISP network. The transport of the IPv4 packets through the IPv6 only networks may use IPv4 packets encapsulated on IPv6 packets or a double translation, by making one IPv4 to IPv6 translation to enter the IPv6 only network and one IPv6 to IPv4 translation to return to the IPv4 network. The chosen techniques were DS-Lite, 464XLAT, MAP-E and MAP-T.

The first part of the test is to check some of the most common usages of the Internet by a home user and the impacts of the transition techniques on the user experience. The second part is a measured comparison considering bandwidth, jitter and latency introduced by the techniques and processor usage on the network equipment.

Keywords: IPv4 to IPv6 Transition; A+P; CGN; Encapsulation; Double Translation

1 Introduction

Vinton Cerf was asked in an interview to Slashdot what he would change about TCP/IP if he could travel back in time. His reply was: "I wish I had realized we'd need more than 32 bits of address space! At the time, I thought this was still an experiment and that, if successful, we would develop a production version. I guess IPv6 is the production version!" [1].

The IPv6 main change is the 128 bits addresses, but it also added new functionalities like the neighbor discovery protocol and router advertisement. It also introduced a simpler fixed size header and extension headers for additional information that aren't needed for packet routing.

The transition from the IPv4 to the IPv6 was planned to avoid problems on the currently working IPv4 networks. This would be accomplished by creating dual stack networks and, when all the devices had IPv6 addresses, the IPv4 would become obsolete. Unfortunately this plan wasn't fully executed yet and the IPv4 addresses are ending before the IPv6 arrives to all devices on the Internet [2]. As the IPv4 and the IPv6 are not compatible, other transition techniques were created on IETF Softwires group [3].

This work provides an experimental comparison of existing techniques that transport IPv4 packets through IPv6 only networks.

2 Related Work

Some of the first transition techniques considered that the Internet still were mainly IPv4 and because of that the created techniques tried to enable IPv6 communication through IPv4 networks, like 6to4 [4] and 6rd [5]. Currently the techniques are being developed to do the opposite, to enable IPv4 communication through IPv6 only networks.

The first option to achieve this is by tunneling IPv4 packets through the IPv6 network, using encapsulated IPv4 packets as data on IPv6 packets. The second option is a double translation, in which IPv4 packets are translated to IPv6 packets as specified in RFC6145 [6] and on the edge of the ISP network a second translation from IPv6 back to IPv4 is made.

Another problem faced for those techniques is the lack of IPv4 addresses, so they have considered ways to reduce the need of them, by sharing IPv4 using NAT at the ISP network (CGN) [7] or by using the source port as an extension of the address, A+P (Address Plus Port Sharing) [8] restricting the ports that each customer can use.

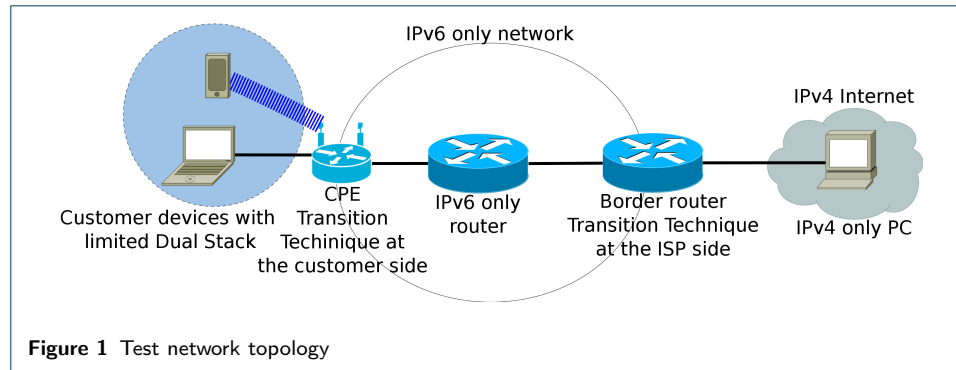
Considering those options, some techniques were created and others are still being proposed. The four techniques compared in this article are examples of possible combination of the described techniques for an IPv6 only ISP network with shared IPv4 addresses: DS-Lite [9], MAP-E [10] e MAP-T [11], 464XLAT [12]. The basic comparison of the techniques is available on table 1.

Table 1 Comparison of IPv6 only ISP transition techniques

	Tunneling	Double Translation	CGN	A+P
DS-Lite	X		X	
MAP-E	X			X
464XLAT		X	X	
MAP-T		X		X

3 Network Topology

The network topology is the same for all the tested techniques and it is explained on figure 1 and includes a customer device to access Internet, a customer premises equipment (CPE) implementing the transition technique on the customer side, an IPv6 only router emulating an IPv6 only ISP network and a provider edge router implementing the ISP part of the transition technique. The CPE and the IPv6 only router are a TP-Link 1043ND Wireless Router with OpenWRT [13] firmware and the provider edge router is a Linux desktop with Ubuntu or Fedora (depending on the technique).



The software used to implement the DS-Lite network was developed by the Internet Systems Consortium [14], a public benefit corporation dedicated to supporting the infrastructure of the Internet by developing software, protocols, and operations. The DS-Lite, the oldest and most mature of the analyzed techniques, is implemented on a lot of network equipment from the most important manufacturers.

MAP-E and MAP-T implementation [15] were developed by China Education and Research Network (CERNET) as proof of concept for the protocol they were proposing at the IETF. Originally MAP-E and MAP-T were only different operating modes for a single protocol called MAP, because of that MAP-E and MAP-T share the same source code. Later IETF decided MAP should be divided in two, that MAP-E should proceed in the standard track and MAP-T should proceed in the experimental track. Both MAP-E or MAP-T are still IETF's drafts and still aren't implemented in large scale.

Despite being approved in the IETF as an informational document and currently in use at the T-Mobile USA's network [16] the 464XLAT is currently only supported at Android devices [17]. As 464XLAT doesn't have a Linux compatible open source implementation for the CPE, like DS-Lite, MAP-E and MAP-T, it was used Tayga [18], a NAT64 [19] software, to implement the CPE's IPv4 to IPv6 translation. Tayga was also used in its original function as NAT64 Border Router for IPv6 to IPv4 translations needed to implement the 464XLAT. The last Tayga update was on June 2011 and as the software isn't updated in a long time it wasn't optimized for small routers nor to the CPE's IPv4 to IPv6 translation.

4 Comparison Results

The comparison was made through a qualitative analysis, testing for noticeable problems or performance issues on common applications and uses of the Internet, and through a quantitative analysis that considered measurement results on networks that implemented the selected techniques.

4.1 Qualitative Analysis

Ten of the most used applications and services [20, 21] were chosen to be tested on networks running DS-Lite, MAP-E, 464XLAT and MAP-T: Internet search, Email access (webmail and smtp), Social networks, Video streaming, Peer-to-peer download (bit torrent), News websites, Internet home banking, Video call, Online gaming and L2TP / IPsec VPN.

Those applications were tested to check if they worked correctly and if any performance or functionality issue could be noticed by the end user. The tests were made on desktops

with Windows 7, Ubuntu Linux 12.04 and MAC OS 10.8.3. Mobile devices were also tested with Android 4.1.2 and Apple IOS 6.1.3.

The peer-to-peer download using the bit torrent protocol failed on all techniques, because none of them allowed incoming connections for the devices. Despite being able to download the target file on all techniques it failed to receive incoming connections after the download of the file is finished, so the client fails to work as a seed which is a fundamental part of the protocol.

The L2TP / IPsec VPN also had very bad results as it failed to work on all operating systems when running through MAP-E, 464XLAT and MAP-T networks. When running on the DS-Lite network it worked correctly on Windows 7 and Android, totally failed on MAC OS and Apple IOS and had a strange result on Linux as the VPN session was successfully established but no traffic was being sent or received through the VPN.

The online game used on the testing was League of Legends [22] so it was only tested on Windows 7 and MAC OS. The game worked fine on the DS-Lite and 464XLAT networks for both operating systems. On the networks with MAP-E and MAP-T it worked fine at MAC OS version, but the Windows version constantly lost connection with the game servers. No significant delay was noticed during the gameplay.

The Skype software [23] was used for the video call and it clearly suffered a quality loss when running on all the tested networks and all operating systems. The video call suffered a quality loss that was noticeable on the transmitted and received videos that became pixelated and blurry. The video call was particularly worse on Android at all the tested networks, but all the operating systems suffered some quality loss on all tested networks. The following figures 2 and 3 are samples acquired to demonstrate the differences. On the Windows sample (figure 2) the difference is very easy to notice as the quality is really bad on the 464XLAT network, showing a very blurred image. On the Linux sample (figure 3) the difference is still perceptible, but it isn't so bad on the DS-Lite test network as it were on the 464XLAT test network.


The rest of the tested services (Internet search, Email access (webmail and smtp), Social networks, Video streaming, News websites and Internet home banking) had no noticeable problem.

4.2 Quantitative Analysis

The software iperf [24] was used to generate network traffic and to measure network quality parameters like jitter, packet loss and throughput. ICMP Echo Request and Echo Reply with different packet sizes were used to measure latency. Others parameters were measured, using sar [25], to help to determinate the technique impact on network performance through the measurement of the processor usage on the network equipment running the transition technique.


On those performance tests the size of the packets were also a variable using packets of 100, 500, 1000 and 1400 bytes and the throughput was measured using TCP and UDP. On the UDP test the iperf was configured to send 100 Mbps and if the routers were unable to transmit all the packets it would cause some packet loss and result on a smaller effective speed.

Analyzing the CPU usage at the CPE's running the transition technique it was possible to evaluate the impact the of the technique in the CPE, as presented on the figure 4. The results demonstrated that the CPE's CPU were completed used on all 464XLAT and DS-Lite tests. In the MAP-E and MAP-T networks the CPU was at maximum capacity for



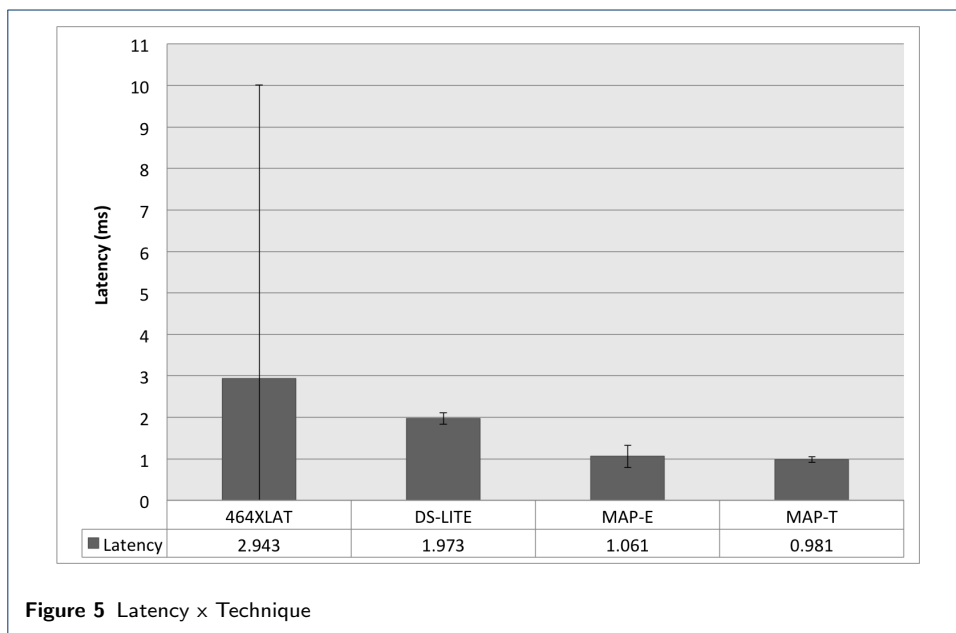
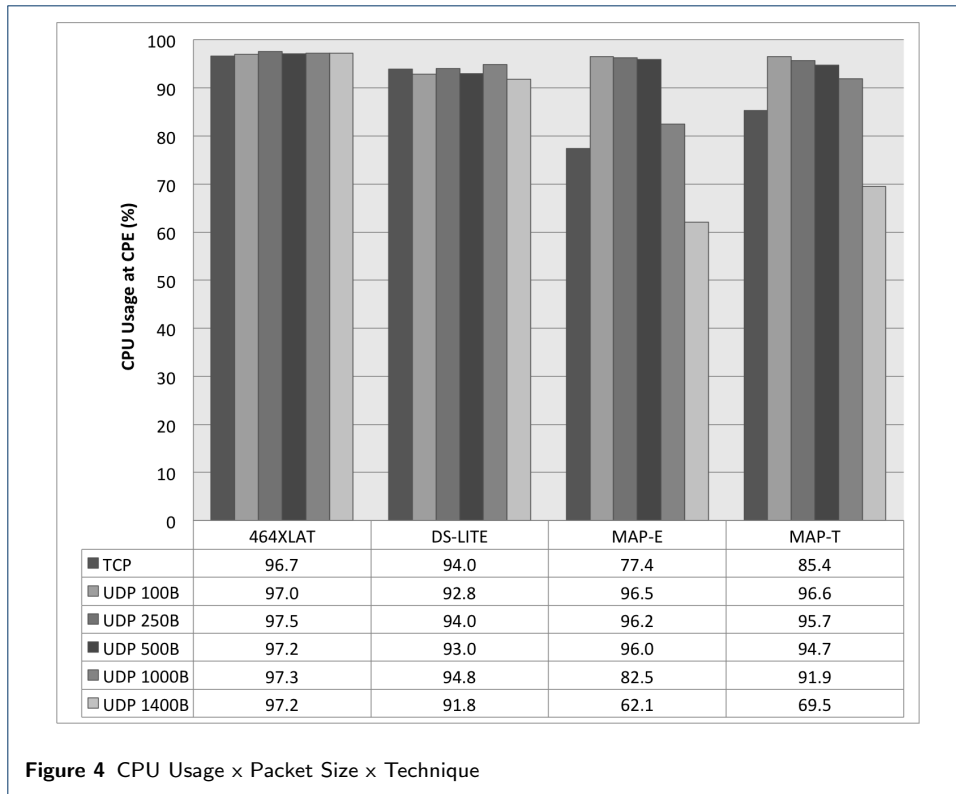
the 100, 250 and 500 bytes UDP packets, but for the TCP and the 1000 and 1400 bytes UDP tests a decrease on CPU usage was noticeable. This result was considered unexpected because neither 464XLAT or DS-Lite make stateful translations on the CPE, they make respectively a stateless IPv4 to IPv6 translation and encapsulation and desencapsulation of IPv4 packets to and from IPv6 packets.

The other analyzed parameters results seemed to be related with the CPU usage in the CPE, and with lower CPU usage better results appeared. The latency results, presented on figure 5, shows that 464XLAT had a latency three times bigger than MAP-E and MAP-T



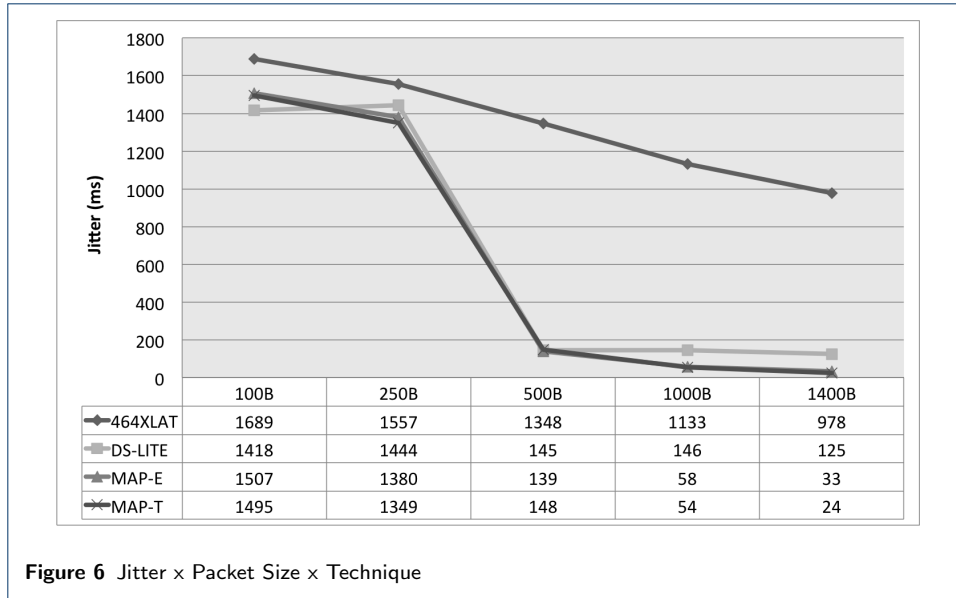
that were less CPU intensive. DS-Lite had an intermediate result being 1 ms faster than 464XLAT and 1 ms slower than MAP-E and MAP-T, this result being compatible with the CPU usage that was smaller than 464XLAT, but bigger than MAP-E and MAP-T.

Analyzing the jitter tests MAP-E, MAP-T and DS-Lite had similar results for the cases with high CPU usage. They had a very bad result for 100 and 250 bytes packets, when the CPU usage was very high. On the 500 bytes packets a considerable better result was noticeable, but while DS-Lite jitter stabilized at bigger than 100 ms values, MAP-E and MAP-T continued to have better results with bigger packets. 464XLAT was the worst of all

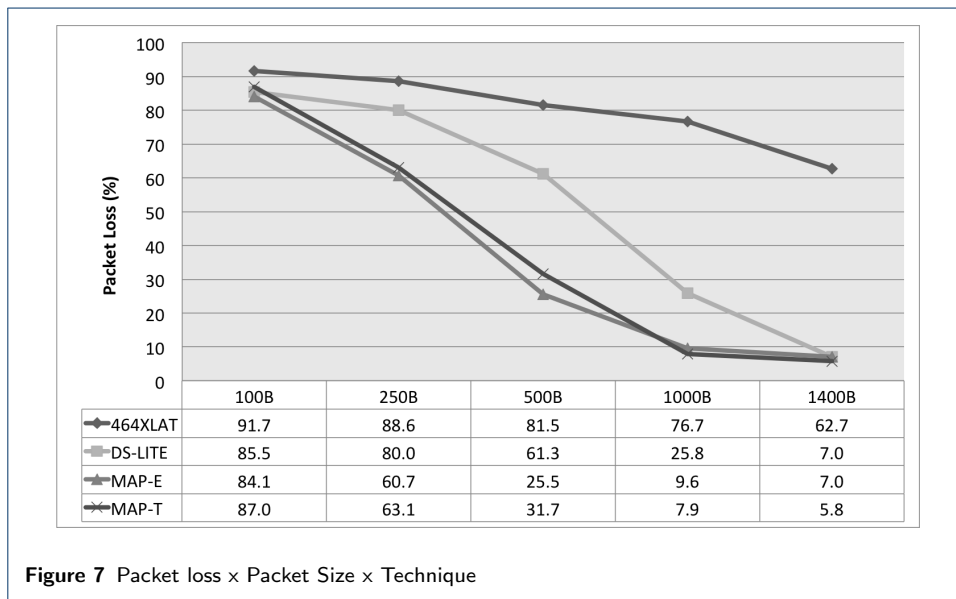


in this test and besides the improvement for bigger packets the almost 1000 ms of its best result, still is a very bad jitter for a network. The results are presented on figure 6.

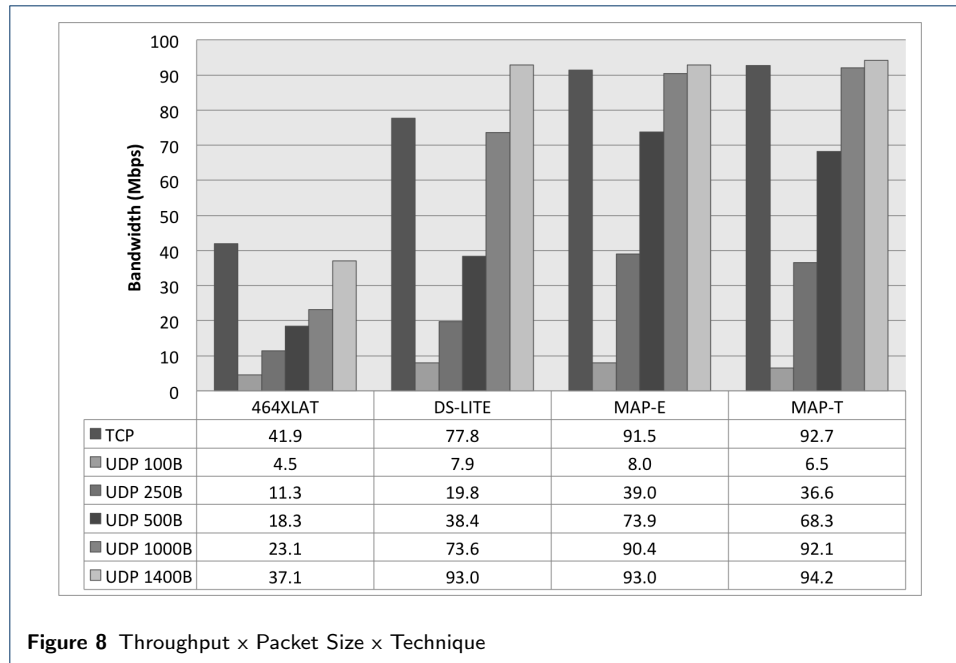
Looking at the packet loss no technique was able to provide an acceptable result for small packets as showed on figure 7. For the 100 bytes packets the best result was at MAP-E test with 84% of packets lost and the worst was 464XLAT with more than 91% of packets lost. MAP-E, MAP-T and DS-Lite had significant improvement for the bigger packets, but



despite similar results for the 1400 bytes packets, DS-Lite had worse results for the intermediate sizes. 464XLAT had an improvement for big packets, but even with big packets it had more than 60% of packet lost.



Because of the packet loss the real UDP throughput was always smaller than the 100 Mbps generated, as showed on figure 8. For UDP packets of 1000 or 1400 bytes MAP-T and MAP-E results were similar to the TCP performance, but as the packet loss was very big for smaller packets the throughput was also smaller. DS-Lite had an increasing throughput results as the packet sizes were getting bigger and had similar performance to MAP-E and MAP-T for the bigger UDP packets, despite being slower on the TCP test. 464XLAT had the worse performance in all tests with a maximum throughput of 40 Mbps on the TCP test.



The encapsulation techniques, DS-Lite and MAP-E, added a 40 bytes overhead for each packet. This overhead was the IPv6 header added to the IPv4 packet to allow the packet to go through the IPv6 only network. It means a 40% overhead for the smaller UDP packet, but only a 2.8% overhead for the bigger packet. For the double translation techniques, 464XLAT and MAP-T, the added size to each packet was 20 bytes, half of the value of the encapsulation. It happens because the original 20 bytes IPv4 header is not transmitted on the IPv6 only network, but it is replaced by the 20 bytes bigger IPv6 header. Considering those results the double translation techniques would be expected to have better performance, but as the IPv6 only network was a Gigabit Ethernet, ten times bigger than the 100 Mbps used on the UDP tests, the lower overhead wasn't a advantage for the double translation techniques.

5 Conclusion

As mentioned on the Network Topology section, the 464XLAT doesn't have a CPE dedicated open source implementation, as DS-Lite, MAP-E and MAP-T and the software used to implement it was last updated on June 2011. Because of that the translation used almost all CPU for every packet size and still were unable to achieve 40 Mbps with the biggest UDP packets. As the 464XLAT is currently being used at T-Mobile's production network the bad result obtained is probably related to the software implementation not to the technique.

On the DS-Lite, MAP-E and MAP-T networks the low speed obtained for the smaller packets seems directly related to a packets per second limitation on the CPE, as the packet loss for the smaller packets was bigger than 80%, while the packet loss was smaller than 8% for the 1400 bytes packets.

The DS-Lite CPU usage at the CPE is almost constant for all the packet sizes and this high CPU usage, as it also happened in the 464XLAT, is the reason for the two times bigger latency than on MAP-E and MAP-T networks.

By comparing encapsulation versus double translation of MAP-E against MAP-T, as both were implemented for the same group and share most of the source code, it seems that encapsulation is less processor intensive and this trend clearly visible when using bigger packets. The smaller processor usage also reflect on MAP-E having a higher throughput on the small packets.

All techniques proved to be a reasonable alternative for the IPv4 to IPv6 transition on an IPv6 only network, but the available open source must be optimized or a CPE with higher computational power must be used to enable large scale deployments with any of the techniques analyzed in this comparison.

Competing interests

The authors declare that they have no competing interests.

Author's contributions

EC designed and built the test networks and defined the tests, made some data acquisition, analyzed and interpreted the data and drafted the manuscript. RC made some data acquisition and preliminary data analysis and revision of the manuscript. WZ contributed to the conception of the test networks, definition of the tests, reviewed and have given final approval to publish the manuscript.

Acknowledgements

We would like to acknowledge the support from NIC.br (Brazilian Network Information Center) that borrowed the necessary equipment and provided the network access to the development of the tests.

Author details

¹Institute for Technological Research, São Paulo, Av. Prof. Almeida Prado 532, Cid. Universitária - Butantã, 05508-901, São Paulo, Brazil. ²University of São Paulo, USP, Av. Prof. Luciano Gualberto, travessa 3 n° 380, Cid. Universitária - Butantã, 05508-010 São Paulo, Brazil.

References

1. Cerf, V.: Vint Cerf Answers Your Questions About IPv6 and More on Slashdot.org. <http://interviews.slashdot.org/story/11/10/25/1532213/vint-cerf-answers-your-questions-about-ipv6-and-more> (2011)
2. Huston, G.: IPv4 Address Report. <http://www.potaroo.net/tools/ipv4/> (2014)
3. IETF: Softwire email group at IETF. <http://www.ietf.org/mail-archive/web/softwires/current/maillist.html> (2012)
4. IETF - RFC 3056: Connection of IPv6 Domains via IPv4 Clouds. <http://tools.ietf.org/html/rfc3056> (2001)
5. IETF - RFC 5969: IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) – Protocol Specification. <http://tools.ietf.org/html/rfc5969> (2010)
6. IETF - RFC 6145: IP/ICMP Translation Algorithm. <http://tools.ietf.org/html/rfc6145> (2011)
7. IETF - draft-shirasaki-nat444-06: NAT444. <http://tools.ietf.org/html/draft-shirasaki-nat444-06> (2012)
8. Maennel, O., Bush, R., Cittadini, L., Bellovin, S.M.: A Better Approach than Carrier-Grade-NAT. Technical report, Department of Computer Science, Columbia University (2008)
9. IETF - RFC 6333: Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion. <http://tools.ietf.org/html/rfc6333> (2011)
10. IETF - draft-ietf-softwire-map-10: Mapping of Address and Port with Encapsulation (MAP). <http://tools.ietf.org/html/draft-ietf-softwire-map-10> (2014)
11. IETF - draft-ietf-softwire-map-t-05: Mapping of Address and Port using Translation (MAP-T). <http://tools.ietf.org/html/draft-ietf-softwire-map-t-05> (2014)
12. IETF - RFC 6877: 464XLAT: Combination of Stateful and Stateless Translation. <http://tools.ietf.org/html/rfc6877> (2013)
13. Baker, M. and Rozema, G. and Kaloz, I. and et al: OpenWRT - Wireless Freedom. <https://openwrt.org/> (2006)
14. Internet Systems Consortium, Inc.: DS-Lite by ISC. <http://www.isc.org/downloads/ds-lite/> (2011)
15. China Education and Research Network - CERNET: MAP-E and MAP-T by CERNET. <https://github.com/cernet/MAP> (2013)
16. T-Mobile USA, Inc.: T-Mobile IPv6 is Here and Now. <https://sites.google.com/site/tmoipv6/lg-mytouch> (2013)
17. Byrne, C.: 464XLAT: Breaking Free of IPv4. In: Asia Pacific Regional Internet Conference on Operational Technologies (APRICOT) (2014)
18. Lutchansky, N.: Tayga. <http://www.litech.org/tayga/> (2012)
19. IETF - RFC 6146: Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers. <http://tools.ietf.org/html/rfc6146> (2011)
20. CETIC.br / NIC.br: Survey on the use of Information and Communication Technologies - ICT Households and Enterprises 2011. <http://op.ceptro.br/cgi-bin/cetic/tic-domicilios-e-empresas-2011.pdf> (2012)
21. Pew Internet & American Life Project: Trend Data (Adults) - What Internet Users do Online. [http://www.pewinternet.org/Trend-Data-\(Adults\)/Online-Activites-Total.aspx](http://www.pewinternet.org/Trend-Data-(Adults)/Online-Activites-Total.aspx) (2013)
22. Riot Games: League of Legends. <http://leagueoflegends.com> (2012)

23. Microsoft: Skype. <http://www.skype.com/> (2012)
24. Dugan, J.: Software for measuring maximum TCP and UDP bandwidth performance - iperf. <http://sourceforge.net/projects/iperf/?source=navbar> (2011)
25. Godard, S.: sar. <http://linux.die.net/man/1/sar> (2012)



Comparison Between IPv4 to IPv6 Transition Techniques by Edwin Cordeiro is licensed under a Creative Commons Attribution 4.0 International License.