

IPv6 Operations Working Group
Internet Draft
Intended status: Informational
Expires: April 2018

G. Lencse
BUTE
J. Palet Martinez
The IPv6 Company
L. Howard
Retevia
R. Patterson
Sky UK
October 6, 2018

**Pros and Cons of IPv6 Transition Technologies for IPv4aaS
draft-lmhp-v6ops-transition-comparison-00.txt**

Abstract

Several IPv6 transition technologies can be used to provide IPv4-as-a-service (IPv4aaS) to the customers, while the ISPs have only IPv6 in their access and or core network. All these technologies have their advantages and disadvantages. Depending on several various conditions and preferences, different technologies may prove to be the most appropriate solution. This document examines the five most prominent IPv4aaS technologies considering several different aspects in order to provide network operators with an easy to use guideline for selecting the technology that suit their needs the best.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on April 6, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	High-level Architectures and their Consequences	3
2.1.	Service Provider Network Traversal	3
2.2.	IPv4 Address Sharing	4
3.	More Detailed Analysis	4
3.1.	Details of Architectural Differences	4
3.1.1.	464XLAT	4
3.1.2.	DS-Lite	4
3.1.3.	Lw4o6	4
3.1.4.	MAP-E	5
3.1.5.	MAP-T	5
3.2.	Tradeoff between Port Number Efficiency and Stateless Operation	5
3.3.	Support for Server Operation	6
3.4.	Support and Implementations	6
3.4.1.	OS Support.....	6
3.4.2.	Support in Cellular and Broadband Networks.....	6
3.4.3.	Implementation Code Sizes	6
4.	Performance Comparison	7
5.	Security Considerations	7
6.	IANA Considerations	8
7.	Conclusions	8
8.	References	8
8.1.	Normative References	8
8.2.	Informative References	9
9.	Acknowledgments	10

1. Introduction

IETF has standardized a high number of IPv6 transition technologies [[Len2017](#)] and occupied a neutral position trusting the selection to the market. In the upcoming years, several network operators would like to get rid of the burden of maintaining IPv4 in their access and/or core networks. This document deals with five different solutions, each of which can be used to provide an IPv4 service using an IPv6-only access/core network. The following IPv6 transition technologies are covered: 464XLAT [[RFC6877](#)], DS-Lite (Dual Stack Lite) [[RFC6333](#)], lw4o6 (Lightweight 4over6) [[RFC7596](#)], MAP-E [[RFC7597](#)] and MAP-T [[RFC7599](#)].

2. High-level Architectures and their Consequences

2.1. Service Provider Network Traversal

As for the high-level solution of the IPv6 service provider network traversal, MAP-T use double translation (first at the CE from IPv4 to IPv6, NAT46, and then from IPv6 to IPv4, NAT64, at the service provider network), 464XLAT may use single (NAT64, IPv6 to IPv4) or double translation (as MAP-T), depending on different factors, such as the use of DNS by the applications and the availability of a DNS64. DS-Lite, lw4o6 and MAP-E encapsulate the IPv4 packets into IPv6 packets. Each solution has its own advantages and drawbacks. Double translation results in only 20 bytes overhead (the difference in the minimum header size between IPv4 and IPv6), whereas the overhead of the encapsulation is 40 bytes (because both, the IPv4 and IPv6 headers are needed, compared with only the IPv4 one). The difference may be significant in the case of small packet sizes or when the larger one results in fragmentation.

On the one hand, the first step of the double translation case is a stateless NAT from IPv4 to IPv6 implemented as SIIT (Stateless IP/ICMP Translation Algorithm) [[RFC7915](#)], which does not translate IPv4 options and/or multicast IP/ICMP packets, whereas with encapsulation-based technologies these remain intact.

On the other hand, single and double translation results in "normal" IPv6 traffic which can be inspected, e.g., by hashing algorithms, and firewalls, whereas encapsulation results in IPv4-embedded IPv6 packets and their interpretation requires special software/hardware for DPI (deep-packet-inspection).

The worst case is DS-Lite, which is also doing double stateful translation (NAT44 at the CE, NAT44 at the AFTR).

Another consequence is that the solutions using double translation can carry only TCP, UDP or ICMP over IP, when they are used with IPv4 address sharing (please refer to [section 3.3](#) for more details), whereas the solutions using encapsulation can carry any other protocols over IP, too.

[2.2.](#) IPv4 Address Sharing

All five technologies support IPv4 address sharing, which has severe consequences described in [[RFC6269](#)].

The efficiency of the address sharing of the five technologies is significantly different, it is discussed in [section 3.2](#).

We note that lw4o6, MAP-E and MAP-T may not necessarily be configured to do IPv4 address sharing, see the details in [Section 3.3](#), however in that case there is no advantage in terms of public IPv4 addresses saving.

[3.](#) More Detailed Analysis

[3.1.](#) Details of Architectural Differences

[3.1.1.](#) 464XLAT

CLAT performs a stateless translation from IPv4 to IPv6 [[RFC7915](#)]. It uses IPv4-embedded IPv6 addresses [[RFC6052](#)] for both source address and destination address. PLAT performs stateful NAT64 [[RFC6146](#)].

[3.1.2.](#) DS-Lite

The B4 (Basic Bridging BroadBand) element encapsulates the IPv4 packets into IPv6 packets. The AFTR (Address Family Transition Router) de-encapsulates the IPv4 packets from the IPv6 packets and performs stateful NAPT (Network Address and Port Translation).

[3.1.3.](#) Lw4o6

Lightweight 4over6 is a variant of DS-Lite. The difference is, that the stateful NAPT is moved from the AFTR to the B4 element. It uses a provisioning mechanism to determine the size of the limited port set per user.

3.1.4. MAP-E

The CE (Customer-Edge) router first encapsulates IPv4-in-IPv6. Packets must traverse a MAP BR to be [de-]encapsulated.

3.1.5. MAP-T

The CE (Customer Edge) router first performs a NAT44 to transform the source addresses and source port numbers of the IPv4 packets into a predefined range, the size of which is a design parameter. The CE router then performs stateless translation from IPv4 to IPv6 [[RFC7915](#)], which translates the IPv4 address and the port numbers into the IPv6 address space. The transformations may be fine-tuned by the mapping rules. The MAP BR (Border Relay) performs stateless translation from IPv4 to IPv6 [[RFC7915](#)].

3.2. Tradeoff between Port Number Efficiency and Stateless Operation

464XLAT and DS-Lite use stateful NAPT at the PLAT and AFTR devices, respectively. This may cause scalability issues. Lw4o6, MAP-E and MAP-T avoid using NAPT in the service provider network. Its cost is that they have to limit the port numbers available for a user, which is also the case for DS-Lite. Determining the optimal size of the port set is not an easy task. On the one hand, the lack of port situation may cause serious problems in the operation of certain programs (e.g. the consequences of the session number limitation due to port number shortage were impressively demonstrated using Google Map in [[Miy2010](#)]). The port number consumption of different applications is highly varying and e.g. in the case of web browsing it depends on several factors [[Rep2014](#)]. And there may be several users behind a CPE, especially in the wired case (e.g. Internet is used by different members of a family simultaneously), thus sometimes even a few thousands ports may not be enough. However, on the other hand, assigning too many ports per user will result in waste of public IPv4 addresses, which is a scarce and expensive resource. Therefore, 464XLAT is much more efficient in terms of port number and thus public IP address saving. The price is the stateful operation in the service provider network, which is allegedly does not scale up well. XXX MEASUREMENTS ARE PLANNED TO DECIDE IF IT IS TRUE. XXX

We also note that NAT64 does not pre-allocate ports for customers but allocates them "on the fly" (which means that even the same customer is using ports from different addresses, and most of the time, different customers get ports from any given addresses), and in fact this creates many application/service providers (Sony

PlayStation Network, OpenDNS, etc.) to permanently black-list the IPv4 ranges once they are detected to be used for address sharing.

3.3. Support for Server Operation

Lw4o6, MAP-E and MAP-T may be configured without IPv4 address sharing, allowing exclusive use of all ports, and non-port-based layer 4 protocols. Thus, they may also be used to support server operation, when public IPv4 addresses are assigned to the subscribers, however, then there is no advantage in terms of IPv4 public addresses saving.

It is also possible to configure specific ports mapping in 464XLAT/NAT64 using EAMT [[RFC7757](#)], which means that only those ports are "lost" from the pool of addresses, so there is a higher maximization of the total usage of IPv4/port resources.

3.4. Support and Implementations

3.4.1. OS Support

As for 464XLAT, client support exists in Windows 10, Linux (including Android), Windows Mobile, and Chrome OS, but it is missing from iOS and MacOS. For the other four solutions, we are not aware of any OS support.

3.4.2. Support in Cellular and Broadband Networks

Several cellular networks use 464XLAT, whereas we are not aware of any deployment of the four other technologies in cellular networks, as they are not supported.

In broadband networks, there are some deployments of 464XLAT, MAP-E and MAP-T. Both, lw4o6 and DS-Lite have more deployments, having been up now DS-Lite the most used one, but lw4o6 taking over in the last years.

3.4.3. Implementation Code Sizes

As for complexity hint, the code size reported from OpenWRT implementation is 17kB, 35kB, 15kB, 35kB, and 48kB for 464XLAT, lw4o6, DS-Lite, MAP-E, MAP-T, and lw4o6, respectively (<https://openwrt.org/packages/start>).

We note that the support for all five technologies requires much less code size than the total sum of the above quantities, because

they contain a lot of common functions (data plane is shared among several of them).

4. Performance Comparison

We plan to compare the performances of the most prominent free software implementations of the five IPv6 transition technologies using the methodology described in "Benchmarking Methodology for IPv6 Transition Technologies" [[RFC8219](#)].

On the one hand, the Dual DUT Setup of [RFC8219](#) makes it possible to use the existing "Benchmarking Methodology for Network Interconnect Devices" [[RFC2544](#)] compliant measurement devices, however, this setup has the drawback that the performances of the CE and of the ISP side device (e.g. the CLAT and the PLAT of 464XLAT) are measured together. In order to measure the performance of only one of them, we need to ensure that the desired one is the bottleneck.

On the other hand, the Single DUT Setup of [[RFC8219](#)] makes it possible to benchmark the selected device separately, however, no [[RFC8219](#)] compliant testers available yet. A DPDK-based software Tester for stateless NAT64 is currently under development and it is expected to be available this autumn as a free software. XXX FROM WHERE XXX

Any volunteers for developing [[RFC8219](#)] compliant measurement software?

5. Security Considerations

According to the simplest model, the number of bugs is proportional to the number of code lines. Please refer to [section 3.4.3](#) for code sizes of CE implementations.

For all five technologies, the CE device should contain a DNS proxy. However, the user may change DNS settings. If it happens and lw4o6, MAP-E and MAP-T are used with significantly restricted port set, which is required for an efficient public IPv4 address sharing, the entropy of the source ports is significantly lowered (e.g. from 16 bits to 10 bits, when 1024 port numbers are assigned to each subscriber) and thus these technologies are theoretically less resilient against cache poisoning, see [[RFC5452](#)]. However, an efficient cache poisoning attack requires that the subscriber operates an own caching DNS server and the attack is performed in the service provider network. Thus, we consider the chance of the successful exploitation of this vulnerability as low.

An in-depth security analysis of all five IPv6 transition technologies and their most prominent free software implementations according to the methodology defined in [Len2018] is planned.

6. IANA Considerations

TBD.

7. Conclusions

TBD.

8. References

8.1. Normative References

- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", [RFC 2544](#), DOI 10.17487/RFC2544, March 1999, <<http://www.rfc-editor.org/info/rfc2544>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", [RFC 6052](#), DOI 10.17487/RFC6052, October 2010, <<http://www.rfc-editor.org/info/rfc6052>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [RFC 6146](#), DOI 10.17487/RFC6146, April 2011, <<http://www.rfc-editor.org/info/rfc6146>>.
- [RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", [RFC 6269](#), DOI 10.17487/RFC6269, June 2011, <<http://www.rfc-editor.org/info/rfc6269>>.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", [RFC 6333](#), DOI 10.17487/RFC6333, August 2011, <<http://www.rfc-editor.org/info/rfc6333>>.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", [RFC 6877](#), DOI 10.17487/RFC6877, April 2013, <<http://www.rfc-editor.org/info/rfc6877>>.

- [RFC7596] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture", [RFC 7596](#), DOI 10.17487/RFC7596, July 2015, <<http://www.rfc-editor.org/info/rfc7596>>.
- [RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", [RFC 7597](#), DOI 10.17487/RFC7597, July 2015, <<http://www.rfc-editor.org/info/rfc7597>>.
- [RFC7599] Li, X., Bao, C., Dec, W., Ed., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", [RFC 7599](#), DOI 10.17487/RFC7599, July 2015, <<http://www.rfc-editor.org/info/rfc7599>>.
- [RFC7915] Bao, C., Li, X., Baker, F., Anderson, T., and F. Gont, "IP/ICMP translation algorithm", [RFC 7915](#), DOI: 10.17487/RFC7915, June 2016, <<http://www.rfc-editor.org/info/rfc7915>>.
- [RFC7757] Anderson, T., and A. Leiva Popper "Explicit Address Mappings for Stateless IP/ICMP Translation", [RFC 7757](#), DOI 10.17487/RFC7757, February 2016, <<http://www.rfc-editor.org/info/rfc7757>>.
- [RFC8219] Georgescu, M., Pislaru, L., and G. Lencse, "Benchmarking Methodology for IPv6 Transition Technologies", IETF [RFC 8219](#), DOI: 10.17487/RFC8219, Aug. 2017, <<http://www.rfc-editor.org/info/rfc8219>>.

8.2. Informative References

- [Len2017] Lencse, G., and Y. Kadobayashi, "Survey of IPv6 Transition Technologies for Security Analysis", IEICE Communications Society Internet Architecture Workshop, Tokyo, Japan, Aug. 28, 2017, IEICE Tech. Rep., vol. 117, no. 187, pp. 19-24. <http://www.hit.bme.hu/~lencse/publications/IEICE-IA-2017-survey.pdf>

- [Len2018] Lencse, G., and Y. Kadobayashi, "Methodology for the identification of potential security issues of different IPv6 transition technologies: Threat analysis of DNS64 and stateful NAT64", Computers & Security (Elsevier), vol. 77, no. 1, pp. 397-411, August 1, 2018, DOI: 10.1016/j.cose.2018.04.012, <http://www.hit.bme.hu/~lencse/publications/ECS-2018-Methodology-revised.pdf>
- [Miy2010] Miyakawa, S., "IPv4 to IPv6 transformation schemes", IEICE Trans. Commun., vol.E93-B, no.5, pp.1078-1084, May 2010. DOI:10.1587/transcom.E93.B.1078
- [Rep2014] Repas, S., Hajas, T., and G. Lencse, "Port number consumption of the NAT64 IPv6 transition technology", Proc. 37th Internat. Conf. on Telecommunications and Signal Processing (TSP 2014), Berlin, Germany, pp.66-72, Jul. 1-3, 2014. DOI: 10.1109/TSP.2015.7296411

9. Acknowledgments

The authors would like to acknowledge the inputs of Ole Troan, Mark Andrews, Edwin Cordeiro, Fred Baker, Alexandre Petrescu, and TBD.

This document was prepared using 2-Word-v2.0.template.dot.

Copyright (c) 2018 IETF Trust and the persons identified as authors of the code. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted pursuant to, and subject to the license terms contained in, the Simplified BSD License set forth in [Section 4.c](#) of the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>).

Authors' Addresses

Gabor Lencse
Budapest University of Technology and Economics
Magyar Tudosok korutja 2.
H-1117 Budapest, Hungary

Email: lencse@hit.bme.hu

Jordi Palet Martinez
The IPv6 Company
Molino de la Navata, 75
La Navata - Galapagar, Madrid - 28420
Spain

Email: jordi.palet@theipv6company.com

Lee Howard
Retevia
9940 Main St., Suite 200
Fairfax
Virginia
22031, USA
Email: lee@asgard.org

Richard Patterson
Sky UK
1 Brick Lane
London, E1 6PU
United Kingdom

Email: richard.patterson@sky.uk