

v6ops
Internet-Draft
Intended status: Informational
Expires: 3 November 2022

G. Lencse
Szechenyi Istvan University
2 May 2022

Performance Analysis of IPv6 Transition Technologies for IPv4aaS
draft-lencse-v6ops-transition-benchmarking-01

Abstract

Several IPv6 transition technologies have been developed to provide customers with IPv4-as-a-Service (IPv4aaS) for ISPs with an IPv6-only access and/or core network. All these technologies have their advantages and disadvantages, and depending on existing topology, skills, strategy and other preferences, one of these technologies may be the most appropriate solution for a network operator.

This document examines and compares the performance of some free software implementations of the five most prominent IPv4aaS technologies (464XLAT [[RFC6877](#)], Dual Stack Lite [[RFC6333](#)], Lightweight 4over6 [[RFC7596](#)], MAP-E [[RFC7597](#)] and MAP-T [[RFC7599](#)]) and DNS64 [[RFC6147](#)].

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 3 November 2022.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	2
2.	Performance Analysis and Comparison of DNS64 Implementations	2
2.1.	Purpose and Significance of DNS64	3
2.2.	Measurement Method and Parameters	3
2.3.	Measurement Results	4
3.	Acknowledgements	5
4.	IANA Considerations	5
5.	Security Considerations	5
6.	References	5
6.1.	Normative References	5
6.2.	Informative References	6
Appendix A.	Change Log	7
A.1.	00	7
A.2.	01	7
	Author's Address	7

[1.](#) Introduction

IETF has standardized several IPv6 transition technologies [[LEN2019](#)] and occupied a neutral position trusting the selection of the most appropriate ones to the market.

[[I-D.ietf-v6ops-transition-comparison](#)] provides a comprehensive comparative analysis of the five most prominent IPv4aaS technologies to assist operators with this problem. This document adds one more detail: performance analysis and comparison of the most important free software implementations of the examined IPv4aaS technologies and DNS64.

[2.](#) Performance Analysis and Comparison of DNS64 Implementations

[2.1.](#) Purpose and Significance of DNS64

DNS64 [[RFC6147](#)] can be used together with stateful NAT64 [[RFC6146](#)] to facilitate the communication of an IPv6-only client with an IPv4-only server. In typical deployments, 464XLAT is used together with DNS64 [[RFC6147](#)], see [Section 3.1.2 of \[RFC8683\]](#). In this case, the stateful NAT64 gateway is granted as the PLAT of 464XLAT. Theoretically, DNS64 could be used together with any of the other IPv4aaS technologies, but then it would require the operation of a stateful NAT64 gateway. So DNS64 performance is important for those, who choose to use 464XLAT.

As elaborated in Sections [4.5.2](#) and [4.5.3](#) of [[I-D.ietf-v6ops-transition-comparison](#)], the usage of DNS64 reduces the amount of the traffic that undergoes double translation at least by an order of magnitude, because the traffic of an IPv6-only client and an IPv4-only server undergoes only a single translation.

[2.2.](#) Measurement Method and Parameters

The benchmarking method for DNS64 servers is defined in [Section 9 of \[RFC8219\]](#). Further details of the method and its design considerations are elaborated in [[LEN2017](#)].

In short, the performance of the DNS64 servers is characterized by the number of queries resolved per second, provided that the resolution happens within (one second) timeout time and all the sent queries are resolved. (This is the usual "zero loss" criterion traditionally used in [[RFC2544](#)] and its derivatives.)

The worst case test, when there is no cache hit and no "AAAA" record exists, thus the DNS64 server needs to send two queries (one query for the "AAAA" record and another one for the "A" record) for each received query is compulsory, and additional tests with various rates of cache hits and existing "AAAA" records are optional.

All the details of the actual measurements are described in [\[LEN2018\]](#), where all the information given in the rest of [Section 2](#) is taken from.

As for implementations, we have benchmarked all usable free software DNS64 implementations we were aware of, namely: BIND (9.10.3-P4-Debian), PowerDNS Recursor (4.0.4) and Unbound (1.6.0). Their performance was examined as the function of the number of active CPU cores using 1, 2, 4, 8 and 16 CPU cores. Besides the compulsory and optional tests defined by [\[RFC8219\]](#), we have also examined the computing power relative DNS64 performance, that is the number of processed queries per second divided by the CPU utilization of the computer.

The measurements were performed using the resources of NICT StarBED, Japan. The used Dell PowerEdge C6220 servers had two Intel Xeon E5-2650 2GHz CPUs, having 8 cores each, and 16x8GB 1333MHz DDR3 SDRAM. They were interconnected by 1Gbps speed VLANs. Debian GNU/Linux 9.2 operating system with kernel version 4.9.0-4-amd64 was used on all computers.

All measurements were executed 20 times, and then median, minimum and maximum of the 20 results were calculated.

[2.3.](#) Measurement Results

Only the most important results (median values of the worst case performance) are summarized here. All the details can be found in [\[LEN2018\]](#).

The performance of BIND was 2,425qps (queries per second) at a single core, and it scaled up to 4,788qps and 6,659qps at 2 and 4 cores, respectively. However its performance did not increase at 8 cores

and it was practically the same (6,661qps) at 16 cores. Strangely, the minimum and maximum performance was exactly the same as the median at 4, 8 and 16 cores. We have reported this strange issue to the developers of BIND, but we did not receive any reply or bugfix.

The performance of PowerDNS was 3,077qps at a single core, and it scaled up quite well up to 26,620qps at 16 cores. The results were consistent: the difference between the maximum and minimum values was acceptable (under 13% of the median) at any number of CPU cores.

The performance of Unbound was 8,708qps at a single core with only a low difference between the maximum and minimum values (about 4% of the median). Its median performance showed an acceptable scale-up to 31,502qps at 8 cores, but then it dropped to 17,131qps at 16 cores. And the results were rather inconsistent from 2 to 16 cores. (E.g. the difference of the maximum and minimum values was more than 58% of the median at 16 cores.)

All-in-all, PowerDNS has shown the best scale-up and the most consistent and predictable results, therefore we recommend its usage in the case of a modern server with 16 or more CPU cores.

In the case of a single core server (e.g. a virtual machine) Unbound can give the highest performance.

BIND has shown the lowest single core performance and poor scale-up. It can only be used, if very moderate performance is needed.

[3.](#) Acknowledgements

TBD.

[4.](#) IANA Considerations

This document does not make any request to IANA.

[5.](#) Security Considerations

TBD.

6. References

6.1. Normative References

- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", [RFC 2544](#), DOI 10.17487/RFC2544, March 1999, <<https://www.rfc-editor.org/info/rfc2544>>.
- [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, <<https://www.rfc-editor.org/info/rfc6146>>.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", [RFC 6147](#), DOI 10.17487/RFC6147, April 2011, <<https://www.rfc-editor.org/info/rfc6147>>.
- [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, <<https://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, <<https://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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc7599>>.
- [RFC8219] Georgescu, M., Pislaru, L., and G. Lencse, "Benchmarking Methodology for IPv6 Transition Technologies", [RFC 8219](#), DOI 10.17487/RFC8219, August 2017, <<https://www.rfc-editor.org/info/rfc8219>>.
- [RFC8683] Palet Martinez, J., "Additional Deployment Guidelines for NAT64/464XLAT in Operator and Enterprise Networks", [RFC 8683](#), DOI 10.17487/RFC8683, November 2019, <<https://www.rfc-editor.org/info/rfc8683>>.

6.2. Informative References

- [I-D.ietf-v6ops-transition-comparison] Lencse, G., Martinez, J. P., Howard, L., Patterson, R., and I. Farrer, "Pros and Cons of IPv6 Transition Technologies for IPv4aaS", Work in Progress, Internet-Draft, [draft-ietf-v6ops-transition-comparison-03](#), 5 April 2022, <<https://www.ietf.org/archive/id/draft-ietf-v6ops-transition-comparison-03.txt>>.
- [LEN2017] Lencse, G. and Y. Kadobayashi, "Benchmarking Methodology for DNS64 Servers", Computer Communications, vol. 109, no. 1, pp. 162–175, DOI: 10.1016/j.comcom.2017.06.004, 1 September 2017, <<http://www.hit.bme.hu/~lencse/publications/ECC-2018-DNS64-BM-for-review.pdf>>.

- [LEN2018] Lencse, G. and Y. Kadobayashi, "Benchmarking DNS64 Implementations: Theory and Practice", Computer Communications, vol. 127, no. 1, pp. 61–74, 10.1016/j.comcom.2018.05.005, 1 September 2018, <<http://www.hit.bme.hu/~lencse/publications/ECC-2018-DNS64-BM-for-review.pdf>>.
- [LEN2019] Lencse, G. and Y. Kadobayashi, "Comprehensive Survey of

IPv6 Transition Technologies: A Subjective Classification for Security Analysis", IEICE Transactions on Communications, vol. E102-B, no.10, pp. 2021-2035., DOI: 10.1587/transcom.2018EBR0002, 1 October 2019, <http://www.hit.bme.hu/~lencse/publications/e102-b_10_2021.pdf>.

Appendix A. Change Log

A.1. 00

Initial version.

A.2. 01

DNS64 benchmarking was added.

Author's Address

Gabor Lencse
Szechenyi Istvan University
Gyor
Egyetem ter 1.
H-9026
Hungary
Email: lencse@sze.hu