Measuring IPv6 Extension Headers Survivability with James

Raphaël Léas[†], Justin Iurman[†], Éric Vyncke[*], Benoit Donnet[†]
[†] Université de Liège (Belgium), [*] Cisco

ABSTRACT

This extended abstract introduces James, a new tool for measuring how IPv6 Extension Headers (IPv6 EHs) are processed in the network. James sends specially crafted Paris traceroute packets between a set of controlled vantage points. Early measurement results show that IPv6 EHs may be dropped in the network, depending on their type and the size of the Extension Header.

CCS CONCEPTS

• **Networks** → **Network layer protocols**; **Network measurement**; *Routers*.

ACM Reference Format:
Raphaël Léas[†], Justin Iurman[†], Éric Vyncke[*], Benoit Donnet[†], [†] Université de Liège (Belgium), [*] Cisco. 2022. Measuring IPv6 Extension Headers Survivability with James. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22), October 25–27, 2022, Nice, France.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3517745.3563019

1 CONTEXT

During the last decade, IPv6 has been more and more adopted [7]. The initial goal of IPv6 was to deal with IPv4 address exhaustion. But it also comes with a mechanism called IPv6 Extension Headers (IPv6 EHs) [3] that leads to more flexibility and innovation. Examples of such innovations based on IPv6 EHs are Segment Routing [4] and In-Situ Operations, Administration, and Maintenance (IOAM) [1]. The purpose of IPv6 EHs is to extend IPv6 without any modification to the core protocol. IPv6 EHs form a chain, using the IPv6 *Next Header* field, and are placed between the IPv6 header and the upper-layer protocol header. While new IPv6 EHs might be defined in the future, the current list mainly includes the `Hop-by-Hop` Option, the `Destination` Option, the `Routing Header`, the `Fragment Header`, the `Encapsulating Security Payload`, and the `Authentication Header`.

Up to now, a few efforts have been made in assessing how operators process IPv6 EHs. RFC7045 [2] provides guidelines on how IPv6 EHs should be transmitted, also with a focus on middleboxes influence on the traffic. Gont and Liu [6] analyze the security implication of IPv6 EHs and the implications of discarding or filtering packets. Further, Hendriks et al. [8] state that dropping all traffic containing any IPv6 EHs is the de facto rule applied by any operator, for security reasons. To support their claim, they perform limited measurement campaigns on a national research network (CSNET) and a campus network (UTNET). RFC7872 [5] provides measurements on a subset of IPv6 EHs, while `APNIC Labs` shares IPv6 EHs measurements [9] on the `Fragment Header` and, more recently [10], also on the `Hop-by-Hop` and `Destination` Options.

This extended abstract introduces James [11], *Just Another Measurement of Extension header Survivability*, as another approach to perform IPv6 EHs measurements. James works in two steps: (*i*) the probing phase, where probes are sent and responses recorded in a *pcap* file; and (*ii*) the processing phase, where *pcap* data is processed and exported as human-readable results. On the contrary to aforementioned techniques, James relies on full mesh measurements between controlled vantage points. Currently, James measurement infrastructure is based on 21 vantage points spread across the world, i.e., 7 in Europe, 5 in North America, 4 in Asia, 3 in Africa, 1 in Oceania, and 1 in South America. All those vantage points are virtual machines rented from multiple cloud providers, so that the traffic generated by James runs through the legacy Internet and does not fall in a Datacenter-to-Datacenter type of traffic. This limited testbed is a good start but we already started probing random prefixes in the wild to discover as many ASes as possible. Indeed, the more probes we send from and to different vantage points, the more ASes we test and the more complete the view we have. At the end of the day, the only solution is large-scale measurements.

2 PRELIMINARY RESULTS

We run James over the measurement infrastructure six times between November 2021 and June 2022. In this section, we discuss the last results obtained by James. Note that previous campaigns showed the same trend in results [11].

James works by sending probes between each pair of vantage points on a Paris `traceroute` basis to limit load balancing issues. Each pair is firstly tested without IPv6 EHs for comparison purposes. For each experiment with IPv6 EHs, the Hop Limit (the equivalent of the TTL field in IPv4) starts at 1 and finishes at N, where N is the number of hops to reach the destination (obtained from the first trace without IPv6 EHs) plus a safety margin of several hops in case the path increases due to the presence of an Extension Header (e.g., slow path deviation). We run different experiments separately for each type of Extension Header, namely (*i*) `Hop-by-Hop` and `Destination` Options with different sizes, (*ii*) `Routing Header` from type 0 to 6 included, (*iii*) atomic and non-atomic `Fragment Header`, and (*iv*) some other protocols as next header. Each experiment is duplicated, once with `UDP` and then with `TCP`. The reason is obviously to observe if packets receive different treatment based on Layer-4 and also to avoid them being flagged as exotic. For all these experiments on all pairs, 44 ASes were traversed. The full list is available for interested readers [12].

This work is supported by the CyberExcellence project funded by the Walloon Region, under number 2110186.

IMC '22, October 25–27, 2022, Nice, France

ACM ISBN 978-1-4503-9259-4/22/10…$15.00
https://doi.org/10.1145/3517745.3563019

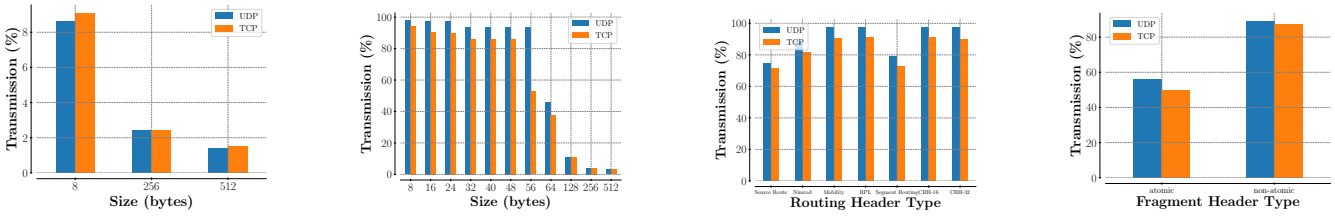

(a) `Hop-by-Hop` Option (b) `Destination` Option (c) `Routing Header` (d) `Fragment Header`

Figure 1: IPv6 EHs survivability results. "Transmission" means packets were able to reach their destination.

Fig. 1a shows the transmission percentage for the `Hop-by-Hop` Option, with the IPv6 EH size varying between 8 and 512 bytes. The purpose of a `Hop-by-Hop` Option is to be processed by all devices along the path. Roughly, only a low 10% of sent packets (with the minimal 8-byte size) traverse the whole path. Larger `Hop-by-Hop` Options (256 and 512 bytes) are more aggressively dropped by the network (respectively only 2.4% and 1.4% of the packets survive along the path). Our intuition is that `Hop-by-Hop`'s are heavily dropped by operators for security and performance reasons, whatever the size. It might change as soon as real use-cases or needs emerge (e.g., IPSec, see below). As a consequence, one cannot rely on `Hop-by-Hop` Options over the global Internet. Note that there is no difference between `UDP` and `TCP` here, probably due to the very low percentage of survivability.

Fig. 1b shows the transmission percentage for the `Destination` Option, with the IPv6 EH size varying between 8 and 512 bytes. The purpose of a `Destination` Option is to be processed only by the packet destination. With `UDP`, the `Destination` Option is largely reliable until 32 bytes. Between 32 and 64 bytes, results are still good but not enough anymore to be called reliable. Starting at 64 bytes, the transmission percentage is halved and it goes even worse when the size is doubled to 128 bytes, where only a low 10% of packets reach the destination. Fig. 1b also shows that the `Destination` Option treatment differs according to the transport protocol considered. Indeed, `TCP` performs worse than `UDP`. The reason behind such a difference is still under investigation but possible reasons are (*i*) some router's buffer limits for hardware lookup mechanisms since Layer-4 is pushed further, (*ii*) the difference between header sizes (8 bytes – `UDP` vs. 20 bytes – `TCP`), or (*iii*), middleboxes applying different treatment based on Layer-4 protocol. Overall, one could rely on `Destination` Options over the global Internet, but one must pay attention to the IPv6 EH size.

Fig. 1c shows the transmission percentage for the `Routing Header`, depending on the `Routing Header` Type. The purpose of a `Routing Header` is to route or steer a packet. Types 2 (Mobility), 3 (RPL), 5 (CHR-16), and 6 (CHR-32) are reliable (at least with `UDP`), as almost all packets (minus a 2%-margin error) reach the destination. Types 0 (Source Route), 1 (Nimrod), and 4 (Segment Routing) are all above 70%, but still cannot be called reliable. Actually, the fact that Types 0 and 1 suffer from drops is not important. Indeed, both are supposed to be deprecated. The situation is a little bit more complex for Type 4 (Segment Routing) as it is specified to be run only in limited domains. Unlike other Extension Headers where a perfect situation would be zero drop (although it is not the case), a good practice for Routing Headers would be to drop by default for security reasons. Regarding the difference in results between `UDP` and `TCP`, the same reasoning as previously can be applied.

Fig. 1d shows the transmission percentage for `Fragment Header`, both for atomic (M-flag = 0) and non-atomic (M-flag = 1) fragments. None of them is reliable, although the non-atomic `Fragment Header` seems to survive more easily. We are still investigating the reasons. Our intuition is that some stateful middleboxes drop atomic fragments because they are unexpected whenever no previous fragment with the same `Identification` number has been seen. Note that the size does not matter here, as different sizes were tested without any difference.

Finally, both the `Encapsulating Security Payload` and the `Authentication Header` were also tested and both survive, which is a good thing for IPSec.

3 REPOSITORY

The source code of James and its measurement results are freely available here: https://gitlab.uliege.be/Benoit.Donnet/james

REFERENCES

[1] F. Brockners, S. Bhandari, and T. Mizrahi. 2022. *Data Fields for In Situ Operations, Administration, and Maintenance (IOAM).* RFC 9197. Internet Engineering Task Force.

[2] B. Carpenter and S. Jiang. 2013. *Transmission and Processing of IPv6 Extension Headers.* RFC 7045. Internet Engineering Task Force.

[3] S. Deering and R. Hinden. 2017. *Internet Protocol, Version 6 (IPv6) Specification.* RFC 8200. Internet Engineering Task Force.

[4] C. Filsfils, S. Previdi, L. Grinsberg, B. Decraene, S. Likowski, and R. Shakir. 2018. *Segment Routing Architecture.* RFC 8402. Internet Engineering Task Force.

[5] F. Gont, J. Linkova, T. Chown, and W. Liu. 2016. *Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World.* RFC 7872. Internet Engineering Task Force.

[6] F. Gont and W. Liu. 2022. *Recommendations on the Filtering of IPv6 Packets Containing IPv6 Extension Headers at Transit Routers.* Internet Draft (Work in Progress) draft-ietf-opsec-ipv6-eh-filtering-10. Internet Engineering Task Force.

[7] Google. [n. d.]. IPv6 Statistics. ([n. d.]). https://www.google.com/intl/en/ipv6/statistics.html Last Access: June, 13th 2022.

[8] L. Hendrikx, P. Velan, R. Schmidts, P. T. De Boer, and A. Pras. 2017. Threats and Surprises Behind IPv6 Extension Headers. In *Proc. IFIP Network Traffic Measurement and Analysis (TMA).*

[9] G. Huston and J. Damas. 2022. IPv6 Fragmentation and EH behaviours. (March 2022). https://www.potaroo.net/presentations/2022-03-20-iepg-v6frag.pdf Last Access: August, 12th 2022.

[10] APNIC Labs. [n. d.]. IPv6 Fragmentation Drop Rate World Map. ([n. d.]). https://stats.labs.apnic.net/v6frag Last Access: August, 12th 2022.

[11] R. Léas. 2022. James Source Code and Dataset. (June 2022). See https://gitlab.uliege.be/Benoit.Donnet/james.

[12] E. Vyncke, R. Léas, and J. Iurman. 2022. *Just Another Measurement of Extension header Survivability (JAMES).* Internet Draft (Work in Progress) draft-vyncke-v6ops-james-02. Internet Engineering Task Force.