



## Is it possible to extend IPv6?

Ana Custura<sup>a,\*</sup>, Raffaello Secchi<sup>a</sup>, Elizabeth Boswell<sup>b</sup>, Gorry Fairhurst<sup>a</sup>

<sup>a</sup> University of Aberdeen, United Kingdom

<sup>b</sup> University of Glasgow, United Kingdom

### ARTICLE INFO

#### Keywords:

IPv6 protocol  
Extension headers  
Protocol evolution  
Destination options  
Hop-by-Hop options

### ABSTRACT

The IPv6 Hop-by-Hop Options and Destination Options Extension Headers have historically faced challenges in deployment due to a lack of router support coupled with concerns around potential for denial-of-service attacks. However, there has been a renewed interest within the standards community both in simplifying their processing, and in using these extension headers for new applications. Through a wide-scale measurement campaign, we show that many autonomous systems in both access networks and the core of the Internet do permit the traversal of packets that include options, and that the path traversal currently depends on the type of network, size of the option and the transport protocol used, but does not usually depend on the type of included option. This is an encouraging result when considering the extensibility of IPv6. We show that packets that include an extension header can also impact the function of load balancing network devices, and present evidence of equipment mis-configuration, noting that a different path to the same destination can result in a different traversal result. Finally, we outline the current deployment challenges and provide recommendations for how extension headers can utilise options to extend IPv6.

### 1. Introduction

IPv6 Extension Headers (EHs) [1] are optional headers that an IPv6 source node can add after the base IPv6 header. They can extend IPv6 by introducing new functionality and add features as a packet traverses a network path. IPv6 EHs are already widely used to implement specific functions (e.g., end-to-end use of IPsec, or within a network to perform source routing). In this paper, we focus on network support for two EHs: the Destination Options (DST) header and the Hop-by-Hop Options (HBH) header [2], as these are the primary means to introduce new end-to-end IPv6 functions.

Recent presentations to the networking community have commented on the limited path traversal of packets that include an EH and noted that network devices, such as firewalls, routers, load balancers and intrusion detection systems [3,4] do not properly handle packets that include an EH. Plausible reasons for the limited traversal are documented in [5], where the authors note that early IPv6 routers processed EHs in software. This processing typically utilises the slow-path, rather than an optimised fast-path (e.g., using hardware forwarding), resulting in a decreased router forwarding rate. In some designs, this processing consumes control plane resources, opening up critical router functions to a denial-of-service (DoS) attack, reducing its ability to perform routing or management [6]. This could have motivated network operators to implement policies that drop packets that include EHs [2]. To date this has discouraged the use of EHs.

Additionally, some network administrators use firewalls to implement Access Control Lists (ACLs) at the outer edge of access and enterprise networks, that discard packets that include an EH. This can mitigate security concerns, such as bypassing security mechanisms or DoS attacks, but also results in packet drops.

The desire to add functionality motivates a fresh look at the usability of EHs as a mechanism to extend IPv6: modern high-speed routers are being introduced with flexible forwarding hardware capable of parsing and processing simple headers within the fast-path [7–9]; and specific use-cases have emerged where there is an operational demand for features that can be effectively implemented through EHs, such as performance metrics [10]. Our paper provides insight into whether these changes in operations and equipment have impacted the forwarding of packets that include EHs, and seeks to understand the opportunity to use these EHs to introduce new functions.

The remainder of this paper is organised as follows: Section 2 presents the required background for IPv6 EHs and describes the historical challenges related to their deployment. The literature describing measurement of paths is surveyed in Section 2.1. Sections 3–5 presents our methodology using a broad dataset to explore key aspects (e.g., the size of the EH, the choice of transport, and choice and composition of the EH Options), revealing a more diverse and nuanced picture of Internet paths than was previously reported. The results are organised by the type of network path and the analysis provides insight into why

\* Corresponding author.

E-mail address: [ana@erg.abdn.ac.uk](mailto:ana@erg.abdn.ac.uk) (A. Custura).

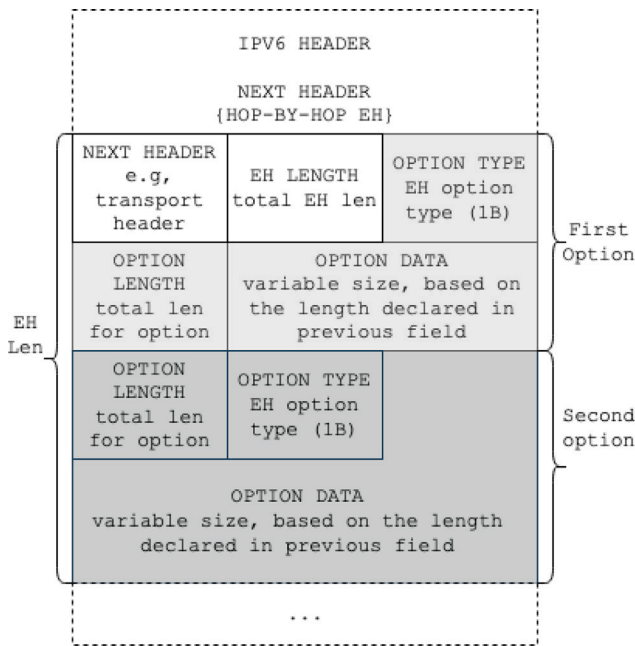


Fig. 1. IPv6 packet with a base header that includes an EH containing two Options.

previous independent measurements reported a variety of results [3, 11–13]. The implications of our results are discussed in Section 6, along with recommendations for how EHs might be utilised in the future. The conclusion summarises our findings, and seeks to answer the question about whether EHs can be used to extend IPv6.

## 2. Extending IPv6

IPv6 introduced a flexible header structure consisting of a fixed-length base header followed by one or more optional EHs. When standardised [14], it was assumed that all routers would process an EH. Initially, relatively few EHs were standardised and deployed. When IPv6 became a full standard in 2017 [1], the processing rules for EHs were modified to align with the prevailing operational practices at that time.

The Next Header field of the IPv6 base header indicates whether a packet includes an EH. Besides HBH and DST, IPv6 specifies EHs for Routing, Fragmentation, Authentication and Security Encapsulation (Encapsulation Security Payload, ESP). The Fragment, Authentication and ESP headers operate end-to-end and follow the Routing header, if present. Each consecutive EH contains a Next Header field to specify the type of the following EH, forming a chain terminated by the IPv6 payload. Each EH also contains a Length field. IPv6 [1] does not specify a maximum size of the EH chain, but does require it to be less than the first fragment in the case of fragmentation.

The DST and HBH EHs are the primary means by which IPv6 functions can be extended, by introducing new Options. Options are encoded using a Type–Length–Value (TLV) encoding [1] with 1 byte for each of the Type and Length fields, and a variable-sized Value field that carries the option data. The total option length must be a multiple of 8 bytes to guarantee alignment in router memory.

When an HBH EH is included, it must be placed immediately after the base header [1]. Fig. 1 shows the structure of an IPv6 packet with a base header followed by an HBH EH containing two Options. A Routing EH can follow the HBH EH to perform Source Routing (SRv6). SRv6 is used to ensure a path includes specified intermediate routers, usually within a single domain [15,16].

The DST EH is typically placed immediately prior to the payload, although can also be included before the Routing EH [1]. An HBH or

Table 1  
Currently standardised DST and HBH options.

Hex	Action bits	Type	Description
0x00	000	HBH, Dest	Pad1 (padding)
0x01	000	HBH, Dest	PadN (padding)
0xC2	110	HBH	Jumbo payload
0x23	001	HBH	Low-Power and Lossy Networks Routing
0x04	000	Dest	IPv6 Encapsulation
0x05	000	HBH	Router Alert Option
0xC9	110	Dest	Mobility Support in IPv6
0x8C	100	Dest	Identification of Broadband Subscribers
0x6D	011	HBH	Multicast Protocol for Low-Power and Lossy Networks
0x0F	000	Dest	Delay Measurement
0x30	001	HBH	MinPathMTU
0x11	000	HBH, Dest	On-path Operational Info
0x31	001	HBH, Dest	On-path Telemetry
0x12	000	HBH, Dest	On-path Telemetry

DST EH placed before a Routing EH has to be processed or skipped when a router is an SRv6 intermediate node, to permit updating the next destination in the SRv6 header.

A router can skip any Option that it does not recognise or it is not configured to support within an EH [1]. The two most significant bits (MSB) of the Option Type field specify the action to be performed when encountering an unrecognised header. When set to “00”, a router should ignore the Option and continue processing the header. If any action bits are non-zero, the packet should be discarded. This specifies that when the action bits are “01”, an ICMP message is returned to the sender. Similarly, if the bits are “11”, an ICMP message is returned, but only if the destination address of the original message was not multicast. Table 1 presents the currently standardised Options. Note that most Options set the two action bits to “00”.

Setting the third MSB allows the data field of a HBH Option to be modified by routers on the path. This can be used, for example, to provide and collect data for traffic measurement [17,18], or collect operational and telemetry data using the recently-proposed Option 0x31 [10].

Originally, all routers on a path were required to examine and process the HBH EH [14]. This requirement was relaxed by [1] to only require processing when configured — for example, option 0x0F [19] can be used to measure performance and provide diagnostic metrics such as round-trip delay.

### 2.1. Previous path traversal studies for IPv6 with EHs

Since its standardisation, the IPv6 protocol has seen widespread adoption [20,21]. However, the Internet community has long been aware of the limited traversal faced by packets containing EHs. In 2015, an Informational IETF RFC presented traceroute measurements to destinations within the Alexa top 1M domains [11] and revealed a significantly higher drop rate over the Internet for packets that include EHs compared to packets without them. Other studies [3,12,13] have also supported this claim. However, the level and nature of the loss varied significantly from report to report. This motivated further analysis and different investigation methods to understand the causes of this loss [13,22].

Another IETF draft [13] presented results using traceroute over a mesh network with 21 vantage points located in a set of globally distributed Autonomous Systems (ASes). This study tested all standard EHs in a setting where both endpoints are under the control of the researcher concluding that only 8%–9% of paths in Internet can be traversed by an 8 Byte (8B) HBH EH, and 97% of paths by an 8B DST EH. These percentages decrease as the size of the EH increases [23]. It

should be noted, however, that 6 of the 21 vantage points were hosted by Digital Ocean™, an Internet service provider that drops packets that include an HBH EH.

An innovative measurement methodology was proposed by engineers in APNIC [12] to analyse end-to-end traversal for Fragmentation, HBH and DST EHs. This technique uses IPv6 packets including an EH to open TCP connections from a crowd-sourced pool of clients, and evaluates the number of successful connection establishments. The results from 4M measurements/day from clients across the Internet found that 50% of attempts to open a connection using a packet that included a DST EH were successful, but close to 0% when an HBH EH was included. This test required both the Internet path to forward the EH and the endpoint to reply to a packet that included the tested EH.

A large-scale passive measurement campaign used the Czech Republic national research and education network to analyse IPv6 traffic over a period of one month in 2016 [24]. It found that 0.1% of IPv6 flows included an EH, out of which 40.9% packets included an HBH EH with an ICMPv6 payload, primarily multicast (although not specified by the original authors, we identify this as Multicast for Low-Power and Lossy Networks [25]). The study also noted that dropping of ICMPv6 traffic that included EHs could result in loss of essential network control information.

Our large-scale measurement study complements and extends these previous analyses. It not only looks at the end-to-end support in servers, it also provides comparative path analysis and observation of longitudinal changes in the traversal for HBH and DST.

## 2.2. Challenges and operational considerations

The Internet hosts a wide range of router designs, spanning from Customer Premises Equipment (CPE) access routers to high-speed transit routers with the capacity to handle thousands of GB/s. Many high-speed routers use an architecture where packets are processed on the “fast-path” utilising hardware support (e.g., an Application Specific Integrated Circuit, ASIC). Packets that cannot be processed on this path use the “slow-path” in software, possibly utilising the control plane processor [5,26]. Using the slow-path exposes the routers to DoS attacks [6], where traffic processing is forced on the control plane reducing resources to manage the router [27]. This could be mitigated by reducing the rate of packets entering the control plane. Awareness of this problem [24] motivated network operators to configure routers to discard packets that include an EH, in particular HBH EHs. The authors also noted that some routers discard packets that include an EH due to flawed implementations of the IPv6 stack [24]. Our paper aims to analyse whether this practice remains prevalent in the current Internet.

Certain network nodes also have a need to inspect the transport protocol information, for instance when an ACL inspects packet ports. Use of ACLs is common at a network domain edge, including the edge of enterprise and home access networks, to implement functions such as firewalls, multi-field Quality of Service classifiers, deep packet inspection and DoS attack mitigation [28]. Some access-network routers also modify upper layer protocol headers to avoid issues related to encapsulation, e.g., by performing TCP Maximum Segment Size (MSS) Clamping [29]. When an EH is present, the router must parse the entire IPv6 header chain and locate the payload to read or modify the TCP header.

Routers that operate in transit networks typically do not require access to upper-layer information. A notable exception are the devices performing Equal Cost Multipath Routing (ECMP) or application-layer load balancing, which can use transport-layer information to drive utilisation of multiple alternative paths. RFC 9288 [30] recommends that transit routers forward packets only on the fast-path, or employ a mechanism to limit the rate of packets appearing on the slow-path. Whenever no rate mitigations are available over the slow-path, discarding packets is recommended.

**Table 2**  
Experiments and datasets.

Purpose	Tool	Name	Date	Trans.
Test traversal of 8B Opts in access networks	Traceroute	R1- Access	Oct '22–Jan '23	UDP TCP
Test traversal and EH size in access networks	Traceroute	R2- Size	Oct '22	UDP TCP
Test whether a consistent path is used	Paris Traceroute	R3- Paris	Jan '23 Aug '23	UDP
Test traversal of Opts to the server edge	PATHSpider	P1- Server	Jul '20–Jan '23	UDP TCP
Test variations in Opt type or content	PATHSpider	P2- Opts	Jul '22–Dec '22	UDP

## 3. Description of datasets

This paper employs a combination of tools and experiments, described in Table 2, to explore how packets that include an HBH or DST EH traverse Internet paths. Datasets R1-R3 are traceroute-based access network measurements from a distributed measurement platform, while Datasets P1 and P2 describe end-to-end measurements to edge servers using PATHSpider [31]. The methodology for each test is discussed in the next two subsections.

### 3.1. Access network paths using RIPE Atlas

Datasets R1-R3 in Table 2 were collected using the RIPE Atlas measurement platform [32]. In January 2023, this provided 5464 IPv6 vantage points (probes) across 644 unique AS Numbers (ASNs), spanning a range of commercial ISPs and R&E access networks. Traceroute packets were sent that included a PadN Option inserted in a DST and/or HBH EH. This Option was defined in the original IPv6 standard to provide 8B alignment within an EH and is expected to be recognised by all IPv6 implementations.

Data was collected using both UDP and TCP transports from all 5464 vantage points, targeting seven globally distributed servers (dataset R1 - Access). These targets are a mix of servers under the control of the authors and other Atlas vantage points, chosen for geographical diversity. In a separate experiment (dataset R2 - Size), we varied the transport and EH size from 8B to 64B to four destinations, for a total of 129,585 measurements, with a mean of 4628 measurements ( $\sigma=351$ ) for each combination of transport, size and EH. The variation in the number of measurements results from availability and connectivity of probes with time. The time required for a complete set of tests (with two different transports, two types of EHs and several sizes) limited the total number of targets.

Finally, we used Paris Traceroute [33] to detect whether the presence of an EH influences the path taken, in dataset R3 - Paris. For this experiment we only selected the vantage points where traversal was successful over UDP for both types of tested EH to a specific target, for a total of 766 vantage points. We measured the paths from these vantage points to our target server using IPv6 packets with no EH, and packets carrying 8B DST and HBH EHs. Each measurement was repeated 16 times (each identified by a Paris ID), as in [33]. Each repetition varied the source port and, in the case of approximately half of the vantage points, the Flow Label (FL). Each set of 16 measurements was repeated five times. We also repeated this test in August 2023 using 32 Paris variations from a subset of approximately 380 vantage points for which the IPv6 FL setting behaviour was known.

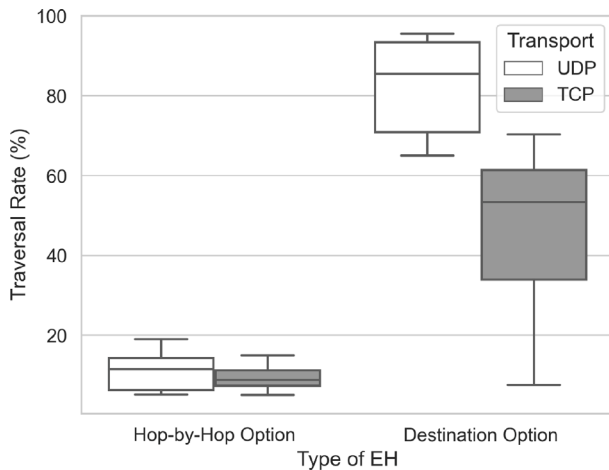


Fig. 2. Traversal for packets that include HBH and DST, from Atlas vantage points to target servers located in seven different countries (dataset R1 - Access).

### 3.2. Measuring server edge paths using PATHSpider

Datasets P1 and P2 were collected using PATHSpider [31], a path transparency testing tool. This works by performing consecutive tests to the same target server, first with a packet that has no EH, and then with packets that include an EH. Compared to Atlas, which performs an in-depth hop-by-hop measurement, PATHSpider does not test individual routers on the path. This instead allows for the rapid measurement of a very large number of target servers.

PATHSpider was used to survey Domain Name System (DNS) servers between 2019 and 2023. This used a vantage point within the University of Aberdeen. The set of DNS servers were initially chosen as they could be tested with both TCP and UDP. The experiment targeted IPv6 authoritative Name Servers (NS) for the then-current Alexa Top 1M domains list. This longitudinal measurement used a consistent set of domains to avoid list changes. Each domain was resolved, removing duplicates and unreachable addresses, resulting in 19,000–22,000 unique IPv6 addresses per test. These tests are included in dataset P1 - Server.

In 2023, we also measured DNS (using UDP and TCP) and Web servers (using TCP) from five global vantage points. The server list was extracted from the Cisco Umbrella Top 1M Domains. ICMP message reception was recorded to analyse router behaviour with EH packets. Together with the longitudinal measurements, these tests form dataset P1 - Server.

Dataset P2 - Opts explores the effects of varying the Option Type and Option Length fields, to observe the impact of different types of options, as well as incorrectly declared lengths. Additionally, we recorded any received ICMP messages for each source–destination pair to determine the frequency of ICMP Type 3 (Destination Unreachable) or ICMP Type 4 (Parameter Problem) messages sent by routers when dropping packets. One tested Option Type has the action bits set to 11, as indicated in Table 1.

## 4. Measurement results using RIPE Atlas probes

This section presents the results of experiments R1 through R3, conducted on the Atlas platform. The primary performance metric analysed is path “traversal”. This is the proportion of paths where probe packets that include an EH successfully reached the destination AS, represented as a percentage of the total paths tested.

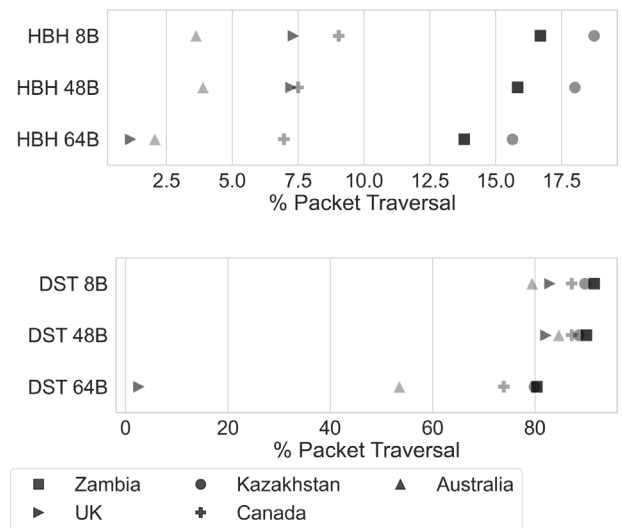


Fig. 3. Traversal for packets that include HBH and DST EHS of three different sizes from the Atlas vantage points to four target servers, with UDP transport.

### 4.1. Traversal to destination AS

Fig. 2 shows the distribution of the traversal with four header compositions (using DST or HBH extension headers, and using TCP or UDP), while consistently employing the 8B PadN Option in EHs. The target destinations were located in seven countries: the United States (US), the United Kingdom (UK), Australia, Poland, Zambia, Kazakhstan, and Singapore, using an average of 4750 vantage points per destination.

The figure shows a 83% and 57% median for path traversal respectively using a packet that includes the DST EH and carries a UDP and TCP payload. The traversal is lower for an HBH EH, with a median of 12% (UDP) and 9% (TCP). The traversal for TCP has much greater variability than with UDP, ranging from 8% for the Zambian destination to 67% for the destination in the UK. The lower traversal for HBH EHs, and for packets carrying TCP was linked to the behaviour and configuration of routers within access networks, more specifically ISP ingress routers (discussed in sub-sections 4.3 and 4.4).

The impact of EH size on the traversal is explored in dataset R2 - Size. Packets that include an EH between 16B and 64B were sent from the ATLAS vantage points to 5 target destinations. This experiment was repeated by including the HBH and DST EHs, and using both UDP and TCP.

The results depicted in Figs. 3 and 4 show the relationship between header size and traversal, demonstrating a decrease in the traversal as the header size increases between 16 and 64B. Fig. 4 only considers results where a packet that includes an 8B EH successfully traversed the path.

Although the decrease is visible for all destinations, the amount of decrease is destination-dependent. These findings resemble results in [23], which examined traversal for DST EHs of sizes 32B and 64B, and identified a lower traversal for packets that include an EH larger than 64B.

For one destination, presented in Fig. 4, packets containing a DST EH over UDP have the most substantial decrease in traversal between 48B and 56B. A comparable pattern is observed in the TCP experiments, where the most significant drop occurs between 40B and 48B.

The difference in traversal between UDP and TCP can be attributed to the overall size of the transport header. The combined header size of TCP (20B) and IPv6 (40B) plus a 48B EH is 108B, while the combined header size of UDP (8B) and IPv6 plus a 56B EH is 104B. This is consistent with a router parsing buffer of 128B, of which several bytes are reserved for parsing packet Layer 2 header information [34,35].

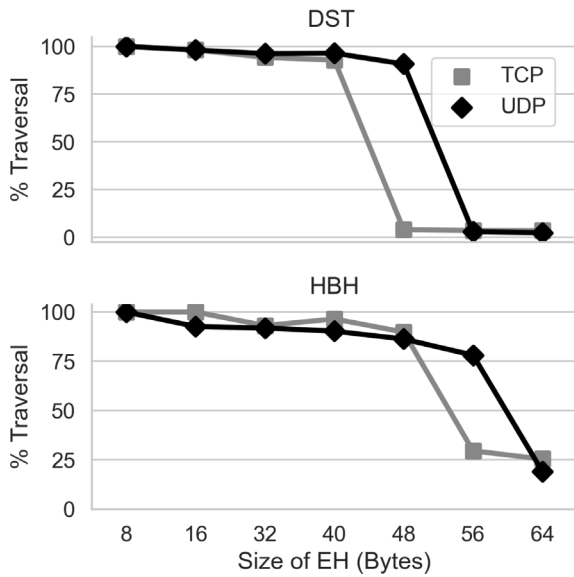


Fig. 4. Traversal for packets that include HBH and DST EHs from the Atlas vantage points to a target server in the JANET network (AS876), showing size of EH and split by transport. These results are compared to other destinations in Fig. 3.

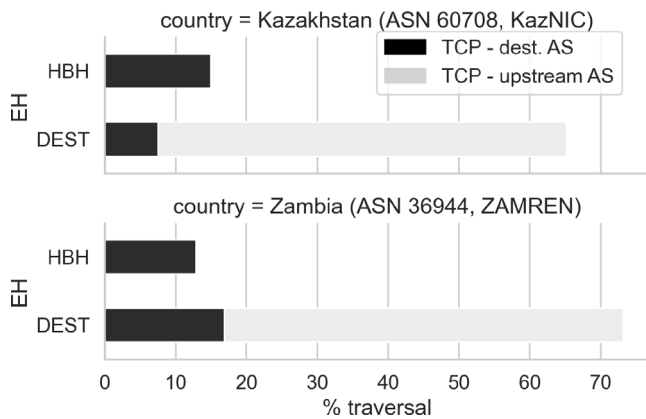


Fig. 5. Packet traversal for TCP payloads from Atlas vantage points to both destination and upstream AS for targets in Kazakhstan (n = 5075) and Zambia (n = 4462).

The size of the parsing buffer in currently deployed routers imposes a constraint on the current usability of a large IPv6 EH. This size could increase with time.

#### 4.2. Locating the point of drop along the path

While there has been prior data on path traversal, less attention has been given to identifying the specific router responsible for packet drops along the path.

Table 3 presents the traversal along the paths between the Atlas vantage points and the UK destination. Within this table, the columns labelled as 1st, 2nd, and ∞ represent the traversal seen at the first, second, and last AS. Additionally, the columns labelled 1st → 2nd and 2nd → 3rd indicate the traversal between ASes, where packet drops could not be attributed to either AS.

#### 4.3. Drops within the first AS

In many cases, a packet that includes an HBH EH sent from a vantage point is dropped within the AS where the vantage point is located, i.e. the initial AS (68% UDP, 74% TCP). Similarly, a notable

fraction of packets that include a DST EH (5% UDP, 25% TCP) are also dropped within the initial AS. This drop rate is irrespective of the destination. Most of these packets are dropped by the first router on the path: for HBH EHs this drop rate is 54% for UDP and 56% for TCP (the type of transport has minimal influence), and 2.5% for UDP, 10% for TCP for DST EHs. Unlike for the overall path, this drop rate is not dependent on the EH length (Fig. 3), and consistent traversal is seen across the entire range of tested EH sizes. We attribute this to the architecture of access routers, which in many cases is not constrained by the use of the parsing buffer in higher-speed router architectures.

UDP packets that include a DST EH experience less than 1% drop rate as they travel across further ASes. This suggests that, once the first AS is traversed, these packets travel to the destination with minimal disruption.

To further understand the impact of the initial AS, we examined the relationship between EH traversal and MSS Clamping. MSS Clamping inserts or modifies a TCP MSS Option into the TCP handshake segments to “clamp” the MSS for a connection to a suitable value to compensate for network encapsulation overhead [29]. Given that Atlas probes do not send a TCP MSS Option by default, the presence of a TCP MSS Option at the destination indicates that an intermediate router inserted it. This implies that a router had to parse the EH chain to analyse the complete TCP/IP header to identify the insertion point. If parsing fails, it is likely to result in a packet drop. In our traces, we identified 853 paths to a UK destination where the TCP MSS option was inserted. Within this subset of paths, the traversal for a packet that includes an 8B HBH EH is 2.6%, while for a DST EH, it is 48.1%. The chi-square test ( $p$ -value  $< 10^{-43}$ ) provides strong evidence of a correlation between EH drops and MSS Clamping, indicating that drops occur more frequently on paths where MSS Clamping is used. This problem is expected to reduce when router protocol stacks are updated to parse EHs.

#### 4.4. Effects of operational configuration

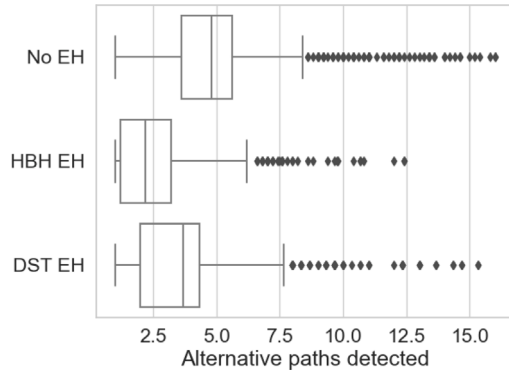
In some cases, the traversal for packets that include an HBH EH is significantly affected by strict traffic aggregation policies enforced by network operators. Two notable examples in our traces are the Kazakhstan and Zambian networks. Both destinations are reachable only through a single Border Gateway Protocol (BGP) peer. When using TCP, the majority of packets were dropped at the second to last AS on their path (corresponding to the destination’s only upstream AS). While the traversal to the destination’s upstream AS shows a behaviour similar to other destinations (see Fig. 5), there was a significantly lower traversal to the target AS.

The tested Kazakhstan network employs a brokering service that tunnels IPv6 traffic to an endpoint situated in Düsseldorf (Germany). Upon closer examination, nearly all the TCP paths to this network that allow a DST EH originate from ASes located in Australia or New Zealand. Conversely, packets originating from other geographical areas are filtered at the tunnel endpoint. This is a specific result arising from a mis-configuration or policy within this operator’s transit network. Similarly, the only BGP peer connecting the target AS in Zambia is UbuntuNet Alliance for Research and Education Networking. Notably, there is no shared origin for the paths on which packets successfully traverse to this AS, indicating that the drops associated with this destination are likely due to an operator policy.

These two examples show how configuration and policy decisions can result in non-delivery of packets that include an EH. We expect that an increase in EH traffic could drive resolution of such issues in the longer term.

**Table 3**  
Packet traversal for the ASes along each path.

AS	1st	1st→2nd	2nd	2nd→3rd	∞
DST UDP	95.3%	93%	–	–	91.5%
DST TCP	74.7%	70%			68.5%
HBH UDP	31.4%	20.1%	15%	12.2%	11.4%
HBH TCP	26.9%	16.3%	13.9%	9.7%	8.6%



**Fig. 6.** Number of paths detected by Paris Traceroute in 766 source–destination pairs, averaged over five measurement runs, each using the same 16 Paris variations (dataset R3).


#### 4.5. EHs and router forwarding

Routers using ECMP (Equal Cost Multi Path) can distribute traffic on different paths based on entropy that includes the Next Header field in IPv6 [28]. Many routers allow the ECMP entropy to be calculated in different ways, e.g. using a simple hash calculation that relies on extracting the transport port from a fixed offset after the EH chain. Some routers also utilise the IPv6 FL. If entropy is not extracted in a way that accounts for the EH chain, this could have two implications: it could mean that packets that include an EH to measure a path, or to detect a path property, will not observe the same path compared to packets that do not. This could also potentially result in packet reordering within a flow. We therefore investigate using dataset R3 whether the inclusion of an EH and/or setting a non-zero FL results in any discernible change in forwarding behaviour compared to packets with no EH.

Paris Traceroute observes the presence of a multipath forwarding by performing several measurements between the same source–destination pair and varying a predetermined set of fields called a *Paris variation* [33]. Packets belonging to the same Paris variation are identified by either a range of sequence numbers in TCP or the checksum in UDP.

Paris Traceroute was used from the vantage points under our control in Atlas to the second to last hop on the path to the Zambian destination. This destination was chosen based on measurements in dataset R1, as it saw the highest traversal rate for both HBH and DST EH packets. An Atlas vantage point was included in this test only if traversal from it was consistently successful in previous tests. In total, packets from 766 source vantage points were measured. Each run between a source and a destination was repeated using 16 Paris variations. Each variation was repeated five times to eliminate the possibility that the forwarding path decision was influenced by an internal source of randomness within the router. The test was subsequently repeated with 32 Paris variations from approximately 380 Atlas vantage points, for which the FL setting behaviour was previously measured.

Fig. 6 compares the distribution of the number of alternative paths discovered by Paris Traceroute including all source–destination pairs and Paris variations. A baseline measurement using packets with no EH is compared with packets that include a DST or HBH EH. A median of 4.7 alternative paths in the baseline experiment is consistent with the



**Fig. 7.** Role of the FL in the number of paths detected by Paris Traceroute in 380 source–destination pairs, averaged over five measurement runs, each using 32 Paris variations (dataset R3). Setting the FL for a measurement results in between 0.5 and 1.2 additional paths detected on average by Paris Traceroute.

results in [33]. However, when DST EHs and HBH EHs are included, the median respectively reduces to 3.6 and 2.1 alternative paths. This variation in the number of identified paths suggests the inclusion of an EH influences the forwarding behaviour. When analysing individual source–destination pairs, the measurement that included a DST EH detected the same number of alternative paths as the baseline (within 1 path) in 60% of cases and fewer paths (by less than 1) in 38% of the cases. The measurement with HBH EHs observed 13.4% of source–destination pairs with the same number of alternative paths and 69.7% discovering less alternative paths than the baseline. Results indicate that inclusion of an EH causes some routers to make different forwarding decisions. In two cases, the inclusion of an EH caused the packet to be dropped on some of the alternative paths; these cases are not included in the analysis.

Fig. 7 shows that setting the FL increases the number of discovered alternative paths. This was measured using 32 Paris variations, and dividing the results depending on whether the Atlas vantage point sets a non-zero FL.

If ECMP causes a packet that includes an EH to traverse a different path, it becomes crucial to determine whether these path variations occur within a single AS (where consistent configuration policies may be in place), or across multiple ASes. This distinction is important as it could result in packets that include Options being processed by different sets of routers than packets without an EH. The implications of this depend on the specific type of extension being introduced. This effect does not reduce the probability of transmission across the path.

## 5. Measuring EH traversal using PATHspider

This section presents results from the analysis of the PATHspider datasets (P1 and P2). These experiments measure the end-to-end traversal from a small pool of vantage points (5 worldwide locations) to a large number of web and DNS servers. Unlike for the Atlas dataset, these measurements require the server to process the packet that includes the EH and reply.

The IPv6 addresses of the target web servers in these tests were collected by resolving the domain names in the Cisco Umbrella top 1M list. The IPv6 addresses of the DNS servers were obtained considering the list of authoritative name servers for the same list of domains. Each experiment sent an IPv6 probe packet (either a DNS query over UDP or TCP to a DNS server, or a TCP SYN to a web server) with a PadN option included in either an HBH or DST EH. If server replies to packets that include an EH were observed, the test then considers the path to successfully forward that EH.

### 5.1. End-to-end support

Table 4 shows the traversal to DNS servers (UDP) and web servers (TCP) in February 2023. Support for DST EH (69%–77%) is higher than

**Table 4**  
Server support for DST and HBH EHs (Feb 2023).

Vantage point location	DST support		HBH support	
	TCP	UDP	TCP	UDP
UK	69.1	69.3	12.5	15.8
Canada	76.3	76	23.3	24.2
Australia	72.5	72.2	17.7	17.5
Singapore	72.8	72.7	17.4	17.4
Poland	76.5	76.8	24.4	24.7
Avg	73.4	73.4	19	19.9

**Table 5**  
Support for DST and HBH EH from DNS providers (Dec 2022).

	% of dataset	DST support	HBH support
Cloudflare	18	Yes	No
Amazon	11	No	No
Hetzner	3	Yes	No
Gandi	4	No	No
Ionos	3	Yes	No
Total	39		

**Table 6**  
Support for DST and HBH from web providers (Dec 2022).

	% of dataset	DST support	HBH support
Amazon	52	Yes	No
Cloudflare	23	Yes	No
Akamai	2.7	Yes	No
Google	2.3	No	No
Total	80		

HBH EH (12%–25%). These results are only slightly affected by the choice of transport protocol, attributed to the absence of access routers.

The table also reveals significant variations of the HBH EH traversal than DST EHs when servers were probed from different vantage points. For instance, the support for HBH EHs varies from 12% from a UK location to almost 25% from a Polish site. This indicates that the transit network may have a greater impact on dropping of HBH EHs than for DST EHs.

The majority of web servers and over one-third of the DNS servers were managed by a few major hosting companies, such as Cloudflare™ and Amazon™. Tables 5 and 6 provide a ranking of the hosting companies based on the share of hosted IP addresses for web and DNS servers, respectively. The tables also report the policy adopted by these companies regarding the propagation of packets that include a DST or HBH EH. This is indicative of how large companies tend to enforce stringent filtering policies to packets that include EHs, possibly to reduce potential risks of incorrect handling of EHs. We found that the policies implemented by the larger hosting providers have indeed the greatest impact on the global traversal of packets that include an EH.


To illustrate this impact, consider that in early December 2022, a change of policy in Cloudflare™ enabled servers to respond to DNS queries carrying EHs. As a result, there was a dataset-wide increase in traversal from 57% to 70%, as currently reported in the table. Extrapolating from this, if all the major providers were to enable support, we estimate that the traversal of this test would exceed 90% for DST and 60% for HBH.

## 5.2. Analysis of AS support for EHs

If we consider the traversal of EHs to target ASes rather than to individual servers, the outlook is different. Because an AS may contain multiple target servers, the number of measurements obtained for each AS will be 5 times number of destinations in that AS, as described in Fig. 8. In its second column, Table 7 reports the percentage of target ASes that could be reached over at least one path from any vantage

**Table 7**  
Reachable AS by DST or HBH EHs (Dec 2022).

Supported EH	Paths per AS $\geq$ 1	Paths per AS $\geq$ 10
Total ASes	2787	1606
DST on at least 1 path	2575 (92.4%)	1496 (93.2%)
DST on at least 50% paths	2476 (88.8%)	1437 (89.4%)
HBH on at least 1 path	1500 (53.8%)	897 (55.9%)
HBH on at least 50% paths	1037 (37.2%)	580 (36.1%)



**Fig. 8.** ASes containing  $n$  destinations will be measured over  $5 \cdot n$  paths, once from each of the 5 vantage points used.

**Table 8**  
Support for EH Options in DNS queries.

Test	DST support	HBH support	Opt. MSBs
Pad N Option (1)	69.3	15.1	000
PMTU Discovery (48)	69.5	15.8	000
Experimental Option (30)	69.4	15.1	000
Experimental Option (254)	0.4	0	110
Incorrect Option Length	0.5	0.05	000

point with the tested EH. Because each AS was measured over multiple paths, rows 3 and 5 also report the percentage where the test was successful for more than half of the tested paths. Only 1606 ASes out of 2787 were measured over more than 10 paths, and show results in-line with those obtained over the pool of all ASes. This is presented in the third column of Table 7. The estimates of the actual support of these EHs in ASes are conservative (under-estimated). A destination AS not reachable from any vantage point may well support them, but could be masked by an upstream AS which drops packets that include them.

The results in Table 7 highlight that 90% of the tested ASes forward packets that include an 8B DST EH and about half forward packets that include an 8B HBH EH. There is little variation when considering the ASes (1606) tested over 10 or more paths.

Analysing whether traversal to an AS was successful over multiple paths suggests that many packets could be dropped before reaching the destination AS. Compared to the percentage of ASes that show support over at least one path, 3.4% fewer ASes allow DST traversal on more than half the paths, whereas for HBH, the difference is 16.6%. Again, this demonstrates the need for transit networks to forward HBH EHs.

## 5.3. Support for IPv6 options

Measurements were performed to determine whether the results are influenced by option data carried in the EH. We first evaluated the effect of the two higher ordered bits of the option type that indicate how router should behave when the option is unknown [1]. Table 8 reports the traversal for various option types in tests towards DNS servers. In addition to the already considered PadN Option, we tested the recently standardised MinPathMTU HBH [17], and two experimental options: 30 and 254 [36]. Option 254 is specified for testing only, packets which include it are ought to be dropped by all routers, since both action bits are set and the option is unrecognised.

The measurements show that if the action bits are zero (i.e., option type  $\leq$  63), the type of option has no effect on EH traversal. When the

**Table 9**  
Percentage of probes triggering ICMP messages.

		UK	Can	Aus	Sgp	Pol
ICMP rcvd from local AS	HBH	0	0	0	0	0
	DST	100	51.6	51.9	51.9	51.5
ICMP rcvd from other AS	HBH	72.8	52.5	68.2	69.2	73
	DST	0	0	0	0	0
ICMP rcvd & packet fwd	HBH	0	0	0	0	0
	DST	0.52	0.48	0.46	0.24	0.46
ICMP not received	HBH	27.2	47.5	31.8	30.8	27
	DST	0	48.4	48.1	48.1	48.5

action bits are set, the packet was expected to be dropped [1]. Instead, we observe responses for 0.4% of paths. This means that the action bits have been ignored by all routers on a small number of paths. Finally, we tested an incorrectly set Option Length field. Any node parsing this EH field should validate the Option Length and discard the packet [1]. However, also in this case, we found a small number of paths (0.5%) where all routers on the path ignore the field. These routers could have been configured to ignore this Option [1].

#### 5.4. ICMP parameter problem messages

A router unable to process an option with a non-zero value for the action bits ought to return an ICMP message. Option type  $\geq 192$  should cause the packet containing the Option to be discarded and return an ICMP “Parameter Problem” message [1] to the source. This expected behaviour was observed from all the vantage points.

The local router returned an ICMP message in response to a packet that includes Option 254 in a DST EH (Table 9). However, depending on the vantage point, only between 50 and 100% of probes generated an ICMP message. Where messages were returned on fewer than 100% of paths, this is attributed to ICMP rate-limiting. ICMP rate-limiting was also observed from other routers in response to a packet that includes HBH Option 254. The widespread presence of rate-limiting makes the use of ICMP notifications an unreliable indicator of packet drops due to an unknown Option.

We encountered a few paths where packets were consistently forwarded regardless of the action bits, or instances (for 0.2–0.5% of packets that include DST) where an ICMP message was generated and the packet was still forwarded to the destination. On these paths, ICMP messages were exclusively received from routers in destination ASes and are a result of incorrect processing of the EH [1] by previous routers along the path.

#### 5.5. ICMP destination unreachable

When a packet is discarded due to an EH, an ICMP “Destination Unreachable” message could be generated back to the sender, either from the destination or another router on the path. We found that for all destinations, an ICMP “Destination Unreachable” message is received only in up to 2% of the paths even if the test succeeds.

For tests that included a HBH EH over paths towards DNS servers, the ICMP messages are returned for 0.2% of the paths. For a DST EH, these messages are also infrequent, ranging from 0.3 to 8.8% of paths depending on the vantage point. This indicates that ICMP messages cannot be reliably used to determine whether a packet was dropped in transit due to the presence of Options.

#### 5.6. Longitudinal analysis of support for EH

Table 10 shows a longitudinal analysis across a set of domains collected over three years, between Jan 2020 and Dec 2022. The table shows successful traversal for a packet that included an 8B Pad N Option for both DST and HBH EHs, from a single vantage point to

**Table 10**  
Support for an PadN Option for DST and HBH EHs towards DNS servers.

	Jan 2020	Jul 2020	July 2022	Dec 2022
DST support	59.9%	54.3%	57.4%	71.7%
HBH support	25.7%	23.8%	16.4%	11.9%
Unique IP addresses	18 296	19 690	19 553	20 050

the authoritative NSes in dataset P1. Each tested domain was resolved at the time of the measurement, resulting in a different pool of IP addresses in each session. This shows a trend for decreasing support for the HBH EH. However, the DST EH support remained constant until December 2022 when Cloudflare™ enabled support on their network boosting overall support, as previously mentioned in Section 5.1.

RFC 7872 [11] describes the traversal to the authoritative name servers for the Alexa Top 1M domains in 2014. This observes that packets that include an 8B PadN DST Option traverse paths to 78.6% of server destinations and packets that include an 8B PadN HBH Option traverse paths to 45.9% of destinations. These results were measured from a single vantage point and are not grouped per AS, and therefore can only be compared with results in Table 10. The comparison indicates a 5%–9% decrease in support for the DST EH and a 25%–30% decrease in support for the HBH EH, although we note this could reflect the choice of vantage point or changes within the top 1M domain list itself between 2014 and 2023.

## 6. Discussion

IPv6 hardware and software continue to mature as adoption increases [20]. Some designs based on hardware and re-configurable logic have enabled the introduction of new features [8], and packet parsing capability in routers is improving [9,37].

Many deployment scenarios for HBH Options are currently within a single domain, while some DST Options [10] are being proposed for Internet-wide deployment. This is driving interest within the standards community [5,35,38] to develop new Options.

The next sections discuss the usability of EHs on Internet paths and the barriers to introducing new Options.

### 6.1. Usability of EH across internet paths

It is timely to ask what is the prospect for using EHs to extend IPv6 across adjacent domains, or across end-to-end Internet paths. This leads to a series of questions, which we seek to answer:

#### 6.1.1. What is the expected traversal for a packet that includes an EH sent on an internet path?

First, we consider whether we found evidence that packets which include EHs are expected to traverse an Internet path. We find that packets that include the DST EH do traverse up to 96% of Internet paths to the destination AS (Fig. 2), and that over 92% of server edge ASes (Table 7) also support the DST EH.

We find that packets that include the HBH EH are supported on paths from some vantage points, although many are currently dropped by transit networks (Table 7) and by access networks (see Fig. 2). Misconfiguration or other network policies can also result in anomalies within transit networks, shown in Section 4.4. Server-side, a longitudinal test to DNS servers reveals the support for the 8B PadN HBH Option has decreased over time when considering individual destinations (Table 10), because servers have become centralised under a few ASes that do not yet support the HBH EH. However, more than half of the tested ASes allow packets that include an HBH EH. In some cases, a low traversal was attributed to policy-based dropping. As noted, configured ACLs may be necessary in some networks today to protect routers (e.g., where EH processing cannot be disabled and leads to DoS vulnerabilities or undesirable side-effects [24]). In cases where this is



not needed, we suggest that such a policy is not desirable, because it results in ossification that will obstruct new uses of EHs.

We also find that traversal reduces significantly for packets that include EHs (both DST and HBH) when the path contains an access network router that inserts a TCP transport option. This is a barrier to deployment that we expect requires updates to these routers. This effect is not an obstacle when using other transport protocols.

#### 6.1.2. What size of options can be safely used in the current internet?

To understand if traversal can be improved by limiting the size of the total EH chain, we explored using different size of EHs. We found that packets that include either a HBH or a DST EH that is less than 40B have a higher probability of traversing an access network path with a UDP transport, shown in Fig. 4. A smaller EH chain currently provides more consistent Internet forwarding, and we suggest that until routers are updated, new options ought to be sized so that packets include an EH chain of less than 40B.

#### 6.1.3. What is the impact on forwarding behaviour when a packet includes an option?

Results presented in Fig. 6 show that including an EH will in some cases change the packet's forwarding path. We attribute this to the position of the EH between the IPv6 header and the transport header (which contains the transport port). We infer some ECMP routers do not process, or skip, the header chain to find the actual port information, but might instead wrongly use a byte offset to the expected position of the source port, reducing the entropy available for selecting the path.

Traffic flows which mix packets that do and do not include EHs must anticipate that these packets might not take the same Internet path. This motivates re-considering the use of the FL for load-balancing [39]. We note that modern operating systems set the FL on packets belonging to a traffic flow [40], and some routers do already use this to perform load balancing [28], as shown in Fig. 7. However, the FL has also been used for other purposes [41] including mobility and traffic engineering [42]. We argue, following the recommendation in [39], that using the FL for entropy in load-balancing devices would mitigate the need to parse the entire IPv6 header chain and would also prevent packet reordering. This will enable new use-cases for EH.

#### 6.1.4. Can new options be defined and used across the internet?

Our data shows that packets that include a DST can already traverse many paths across the Internet, and also at the server and access network edge. This traversal does not depend on the type of Option (see Table 8), which is important because it suggests that IPv6 can be extended by defining new Options that can be used on any path that allows EH processing.

The functionality to process new HBH Options needs to be implemented in routers, therefore it is unlikely that all routers on an Internet path will support any specific new HBH Option. We recommend that any functions that use an Option need to be designed to be robust to routers skipping HBH processing (e.g., as described in the MinPMTU Option [17,18]).

### 6.2. Designing and deploying new options

We assert that it is possible to incrementally extend IPv6 by only utilising an EH when a path is found to forward it, using a method similar to [17]: An application can be designed to first send a test packet that includes an EH with an Option, or combination of Options, and not send additional packets that include this Option until this test packet is acknowledged. The process of sending packets both with and without a header to discover whether a path can support that specific header is sometimes called “racing” (e.g., transport protocol racing is explained in [43]; this resembles “A/B protocol feature testing”, as used in PATHSpider [31]). Our proposed racing is performed after

connection setup, in contrast to methods that race the IP protocol version or choice of transport protocol.

Our results show that for up to three-quarters of current access networks, the first AS on the path will drop packets that include an HBH Option (Table 3). In this case, racing would not find support for EH on the path. However, on the remaining 1335 access networks, racing would discover support on between 31% and 66% of paths. Use of racing would allow new extensions to opportunistically take advantage of any support as networks evolve.

This method could also be used to extend the use of the EHs that are currently restricted to controlled domains (e.g., within an AS), such as [10,19], enabling their use across consecutive multiple domains. Since the set of routers forming a path can change with time, the design of the discovery process ought to repeat the racing from time-to-time.

#### 6.2.1. How useful is ICMP message processing for EH?

Since [14], we observe there are important changes in the way that EHs are being used. Modern routers only process an EH when support is explicitly configured. This has the effect of diminishing the usefulness of any ICMP messages generated when an EH cannot be processed. We also find some routers which (correctly) send an ICMP message in response to a packet that included a DST Option with its action bits set, but nevertheless incorrectly forward the packet. We also find that packets which include an HBH Option with both action bits set are commonly forwarded without sending an ICMP message. When considering whether or not a new Option needs to set these action bits, protocol designers need to take into account that ICMP does not provide a reliable mechanism for indicating whether a function is supported by the path. It is therefore expected that future Options will set a 00 value for the action bits.

## 7. Conclusion

This paper presents an extensive measurement campaign and provides novel results indicating the traversal across Internet paths of packets that include either a HBH or a DST EH. It is the first detailed study of treatment by the routers along an IPv6 path, considering both access and server edge networks. These results provide understanding into whether EHs can be used to extend IPv6.

The results indicate that successful reception across an IPv6 path can currently depend on the type of included EH, its size, and on the transport protocol used. This traversal is strongly influenced by the type of network at the source and destination. We also observed that the inclusion of an EH in a packet can impact the set of forwarding paths when the network layer performs load-balancing, although we assert that this can be mitigated by senders setting the FL, and ECMP routers using it.

The results show there are opportunities to use the IPv6 HBH and DST EHs beyond a single controlled domain, with the expectation that applications incrementally utilise new features using these EHs.

Similarly, the results show that existing methods currently utilised in controlled domains could in future be used to extend IPv6 across multiple domains.

The paper finally provides recommendations for the design of new IPv6 extensions using Options. New designs need to be aware that not all routers process EHs and that some paths will drop packets that include EHs. To overcome these challenges, we motivate the use of racing, which will facilitate incremental deployment that can enable new IPv6 functionality.

### Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Ana Custura's experiment work was funded by RIPE NCC Community Fund Project ID 619935

## Data availability

Data will be made available on request.

## Acknowledgements

The authors appreciate the valuable comments provided by Justin Iurman and Benoit Donnet and Bob Hinden. Elizabeth Boswell received funding from the University of Aberdeen to help analyse paths using Paris Traceroute. This work was supported by the University of Aberdeen's School of Engineering Department, and experiments using Atlas probes were funded by the RIPE NCC Community Fund, Project ID 619935.

## References

- [1] S. Deering, R. Hinden, Internet Protocol, Version 6 (IPv6) Specification, RFC 8200, IETF, 2017.
- [2] F. Gont, et al., Operational Implications of IPv6 Packets with Extension Headers, RFC 9098, IETF, 2021.
- [3] N. Elkins, et al., Performance and diagnostic metrics (PDM) destination option testing across the internet, 2022, IEPG, IETF 114.
- [4] F. Gont, IPv6 Security Mythbusting, UK IPv6 Council Enterprise Workshop, 2023.
- [5] S. Peng, et al., Operational Issues with Processing of the Hop-by-Hop Options Header, Internet-draft-ietf-v6ops-hbh-06, IETF, 2023, in preparation.
- [6] M. Naagas, et al., DEH-DoSv6: A defendable security model against IPv6 extension headers denial of service attack, Bull. Electr. Eng. Inf. 10 (1) (2021) 274–282.
- [7] R. Bifulco, G. Rétvári, A survey on the programmable data plane: Abstractions, architectures, and open problems, in: IEEE HPSR'19, 2018, pp. 1–7.
- [8] Cisco Systems, Cisco Silicon One Product Family, Technical Report, Cisco Systems Inc., 2022.
- [9] F. Hauser, et al., A survey on data plane programming with P4: Fundamentals, advances, and applied research, J. Netw. Comput. Appl. 212 (2023) 103561.
- [10] S. Bhandari, F. Brockners, IPv6 Options for In Situ Operations, Administration, and Maintenance (IOAM), RFC 9486, IETF, 2023.
- [11] F. Gont, et al., Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World, RFC 7872, IETF, 2016.
- [12] G. Huston, IPv6 extension headers revisited, 2022, APNIC Blog.
- [13] É. Vyncke, R. Léas, J. Iurman, Just Another Measurement of Extension header Survivability (JAMES), Internet-draft-vyncke-v6ops-james-03, IETF, 2023, in preparation.
- [14] S. Deering, R. Hinden, Internet Protocol, Version 6 (IPv6) Specification, RFC 2460, IETF, 1998.
- [15] P.L. Ventre, et al., Segment routing: A comprehensive survey of research activities, standardization efforts, and implementation results, IEEE Commun. Surv. Tutor. 23 (1) (2021) 182–221.
- [16] A. Abdelsalam, et al., SRPerf: A performance evaluation framework for IPv6 segment routing, IEEE Trans. Netw. Serv. Manag. 18 (2) (2021) 2320–2333.
- [17] R. Hinden, G. Fairhurst, IPv6 Minimum Path MTU Hop-by-Hop Option, RFC 9268, IETF, 2022.
- [18] G. Fioccola, et al., IPv6 Application of the Alternate-Marking Method, RFC 9343, IETF, 2022.
- [19] N. Elkins, et al., IPv6 Performance and Diagnostic Metrics (PDM) Destination Option, RFC 8250, IETF, 2017.
- [20] M. Nikkhah, R. Guérin, Migrating the internet to IPv6: An exploration of the when and why, IEEE Trans. Network. 24 (4) (2016) 2291–2304.
- [21] Google, IPv6 Adoption Statistics, Technical Report, Google LLC, 2023, <https://www.google.com/intl/en/ipv6/statistics.html>.
- [22] N. Elkins, et al., Deep Dive into IPv6 Extension Header Testing, Internet-draft-elkins-v6ops-eh-deepdive-fw-01, IETF, 2022, in preparation.
- [23] R. Léas, et al., Measuring IPv6 extension headers survivability with James, in: ACM Internet Measurement Conference, 2022, pp. 746–747.
- [24] L. Hendriks, et al., Threats and surprises behind IPv6 extension headers, in: Traffic Measurement and Analysis Conference, 2017, pp. 1–9.
- [25] J. Hui, R. Kelsey, Multicast Protocol for Low-Power and Lossy Networks (MPL), RFC 7731, IETF, 2016.
- [26] H. Khosravi, T. Anderson, Requirements for Separation of IP Control and Forwarding, RFC 3654, IETF, 2003.
- [27] J. Aweya, IP router architectures: An overview, Int. J. Commun. Syst. 14 (5) (2001) 447–475.
- [28] R. Almeida, et al., A characterization of load balancing on the IPv6 Internet, in: Passive and Active Measurement Conference, PAM, Springer, 2017, pp. 242–254.
- [29] A. Custura, et al., Exploring usable path MTU in the internet, in: Traffic Measurement Analysis Conference, 2018, pp. 1–8.
- [30] F. Gont, W.S. Liu, Recommendations on the Filtering of IPv6 Packets Containing IPv6 Extension Headers at Transit Routers, RFC 9288, IETF, 2022.
- [31] I. Learmonth, et al., Pathspider: A tool for active measurement of path transparency, in: Applied Networking Research Workshop, 2016, pp. 62–64.
- [32] V. Bajpai, et al., Lessons learned from using the RIPE Atlas platform for measurement research, ACM SIGCOMM Comput. Commun. Rev. 45 (3) (2015) 35–42.
- [33] B. Augustin, et al., Avoiding traceroute anomalies with Paris traceroute, in: ACM Internet Measurement Conference, ACM, NY, USA, 2006, pp. 153–158.
- [34] S. Yeluri, Sizing router buffers — small is the new big... 2023, APNIC Blog.
- [35] T. Herbert, Limits on Sending and Processing IPv6 Extension Headers, Internet-draft-ietf-6man-eh-limits-09, IETF, 2023, in preparation.
- [36] B. Fenner, Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers, RFC 4727, IETF, 2006.
- [37] P. Bosshart, et al., Forwarding metamorphosis: Fast programmable match-action processing in hardware for SDN, ACM SIGCOMM Comput. Commun. Rev. 43 (4) (2013) 99–110.
- [38] B. Hinden, G. Fairhurst, IPv6 Hop-by-Hop Options Processing Procedures, Internet-draft-ietf-6man-hbh-processing-12, IETF, 2023, in preparation.
- [39] S. Amante, B. Carpenter, S. Jiang, J. Rajahalme, IPv6 Flow Label Specification, RFC 6437, IETF, 2011.
- [40] J. Berger, et al., Flaw label: Exploiting IPv6 flow label, in: IEEE Symposium on Security and Privacy, 2020, pp. 1259–1276.
- [41] L. Becerra Sanchez, J. Padilla Aguilar, Review of approaches for the use of the label flow of IPv6 header, IEEE Latin America Trans. 12 (8) (2014) 1602–1607.
- [42] L. Becerra Sánchez, et al., An approach to support traffic engineering in IPv6 networks based on IPv6 facilities, Telecommun. Syst. 72 (1) (2019) 11–27.
- [43] T. Pauly, et al., An Architecture for Transport Services, Internet-draft-ietf-taps-arch-19, IETF, 2023, in preparation.