Is it possible to extend IPv6?

Ana Custura
Gorry Fairhurst

* University of Aberdeen
with funding from RIPE NCC

IPv6 promises

- Larger Address Space

Fix to lack of IPv4 address space

- More Efficient Forwarding/Routing

- Improved IP Packet Fragmentation*

**After some refinements*

- Multicast

?

- End-to-end Security (aka IPSEC)

?

Other ways have emerged, such as QUIC

- Extensibility

Fix to lack of extension in IPv4

?

This project!

Our Project

- Extend Pathspider to support different types and sizes of IPv6 Extension Headers
- Acquire and distribute 5 hardware measurement probes to be used as Core vantage points
- **Measure** EHs using Pathspider and RIPE Atlas
- Disseminate results @RIPE 86 and @IETF 116
- Publish a peer-reviewed measurement paper to understand IPv6 Extension Header deployment

Extensibility - EHs

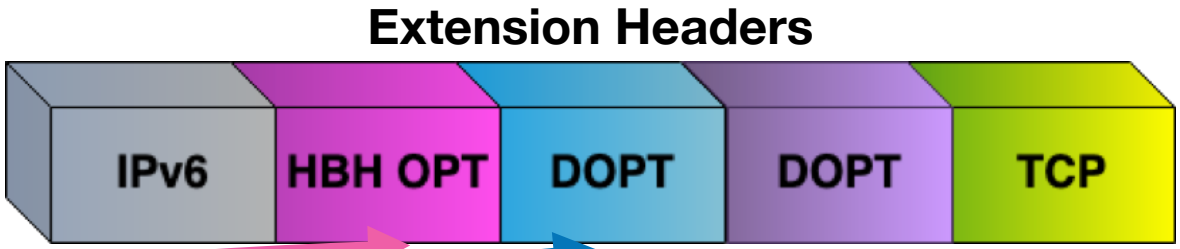

Protocol Number	Description	References
0	IPv6 Hop-by-Hop Option	C8200]
43	Routing Header for IPv6	200] [RFC5095]
44	Fragment Header for IPv6	[RFC8200]
50	Encapsulating Security Payload	[RFC4303]
51	Authentication Header	[RFC4302]
60	Destination Options for IPv6	[RFC8200]
135	Mobility Header	[RFC6275]
139	Host Identity Protocol	[RFC7401]
140	Shim6 Protocol	[RFC5533]
253,254	Use for experimentation and testing	[RFC3692] [RFC4727]

Some EHs carry ‘Options’

https://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml

EH concerns in RFC 9098 (2021)

- Slow-path processing of EHs
- Buggy implementations* -> DoS
- Complexity not bounded: can reduce router forwarding rate
- Large EH can exceed router parsing buffer

Some EHs had a rocky start

Measurements in RFC 7872 show many networks drop packets with EH

* To this date, vulnerabilities still found: https://www.interruptlabs.co.uk/articles/linux-ipv6-route-of-death

Renewed Interest in EHs

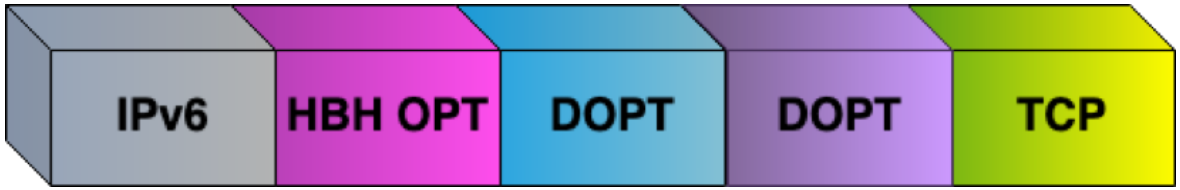

- IPv6 Segment Routing type (SRv6) [RFC8986]
- Service Management and Performance Measurement using PDM [RFC8250]
- In-situ Operations, Administration, and Maintenance [RFC9378]
- AltMark Measurement DO and HbH Options [RFC9343]
- minPMTU HBH Option [RFC9268]

ASICs and Options being standardized are providing renewed interest!

Existing Measurements

- Focus on Destination Options (DOPT) and Hop-by-Hop Options (HBHOPT) EHs
- Let's measure survival of packets with EH

	Destination Option EH	Hop-by-Hop Option EH
RFC 7872 (2016) [1] - server edge	80-90%	45-60%
My own (2018) data [2] - server edge	70-75%	15-20%
APNIC (2022) [3] - client edge	30-80%	0%
JAMES (2022) [4] - core	94-97%	8-9%

Experiment 1: Survival

- ~5500 IPv6-enabled probes in RIPE, globally distributed
- Testing survival by sending packets to 7 targets (UK, US, Canada, Australia, Zambia, Kazakhstan, France)
  - {TCP, UDP} to port 443
    - {**DOPT**, **HBHOPT**} + control IPv6 packets
  - Survives if packet reaches destination AS

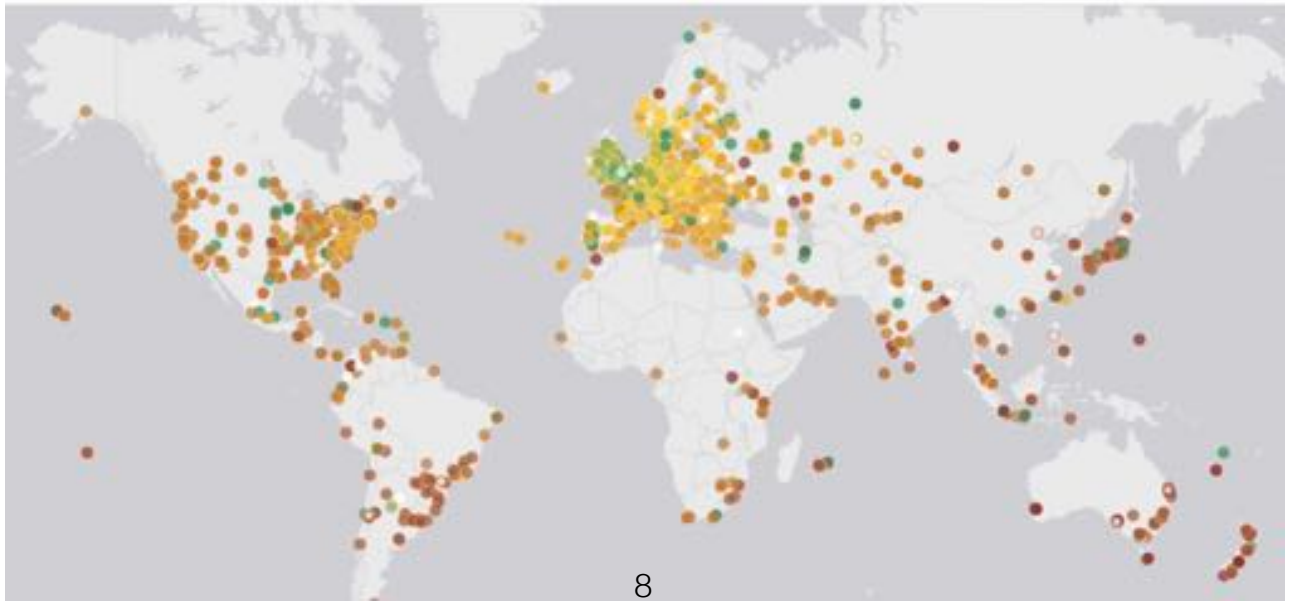

Survival at a Glance
DOPTs

- 8B PadN option
- High survival for **DOPTs**
- Difference between TCP and UDP

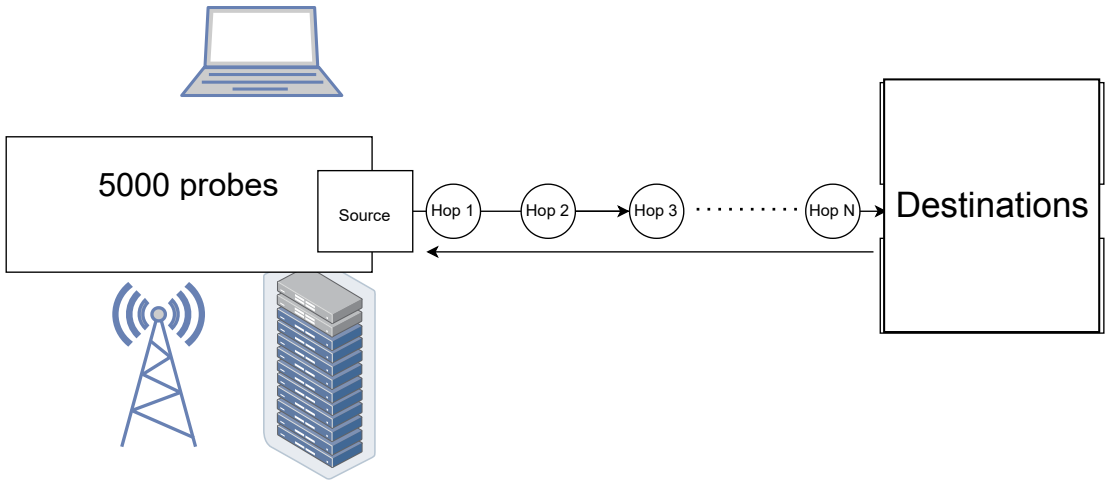

DOPT	
~92%	UDP
~68%	TCP

Survival at a Glance HBHOPTs

- 8B PadN option
- **HBHOPTs** survive some paths
- Difference between TCP and UDP

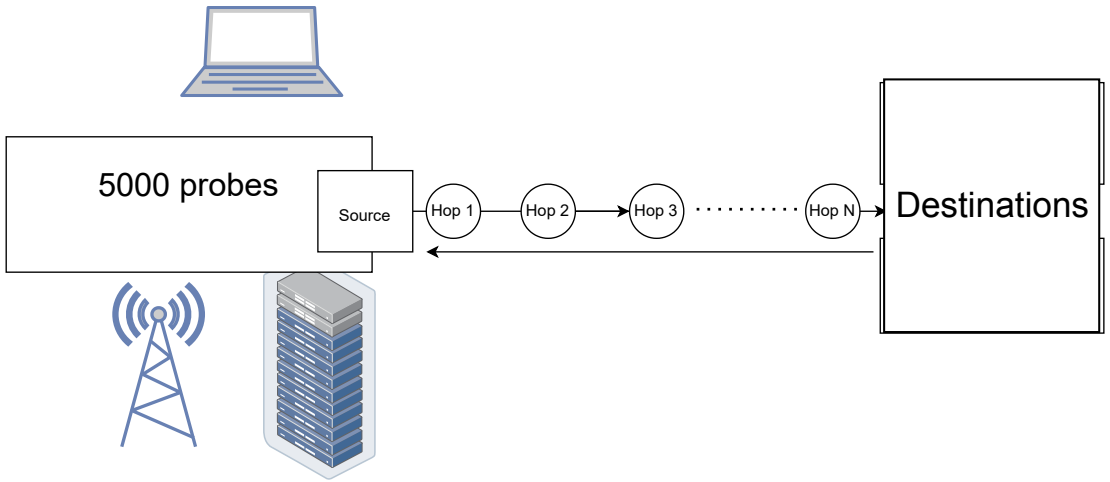

HBHOPT	
~11%	UDP
~9%	TCP

Per-AS Survival (UK path)

DOPT

The **local AS** is responsible for most of the drops:

- 5% for UDP
- 25% for TCP

	1st AS	AS1>AS2	∞
DOPT UDP 8B	95.3%	93%	91.5%
DOPT TCP 8B	74.7%	70%	68.5%

HBHOPT

The **local AS** is responsible for most of the drops:

- 68% for UDP
- 74% for TCP

	1st AS	AS1>AS2	2nd AS	AS2>AS3	∞
HBHOPT UDP 8B	31.4%	20.1%	15%	12.2%	11.4%
HBHOPT TCP 8B	26.9%	16.3%	13.9%	9.7%	8.6%

Drops are considered to be within the AS if the next hop on a control measurement is also in that AS.
If the next hop would otherwise be in a different AS, then the drop is attributed to the AS boundary.

Why?

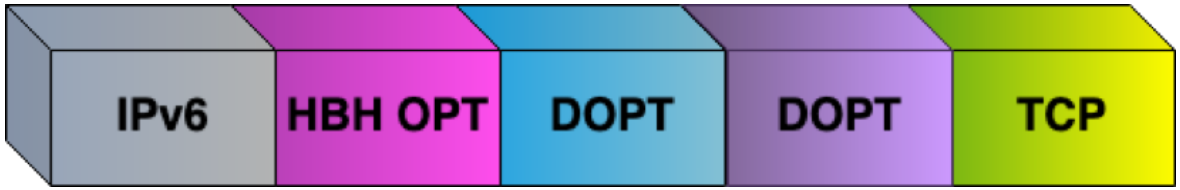

- Network/Firewall policy (e.g. Fastly)
- Different router designs
- Different devices (CPE, load balancers, firewalls, IDS) wanting access to upper layer protocols
- End-systems (NICs that do processing in hosts)
  - Is EH size a factor? Is full chain size a factor?

Experiment 2: Size

- {TCP, UDP} to port 443
  - {**DOPT**, **HBHOPT**} + control measurement
  - {8,16,32,40,48,56,64} B in size to one target
- Survival is successful if packet reaches destination AS

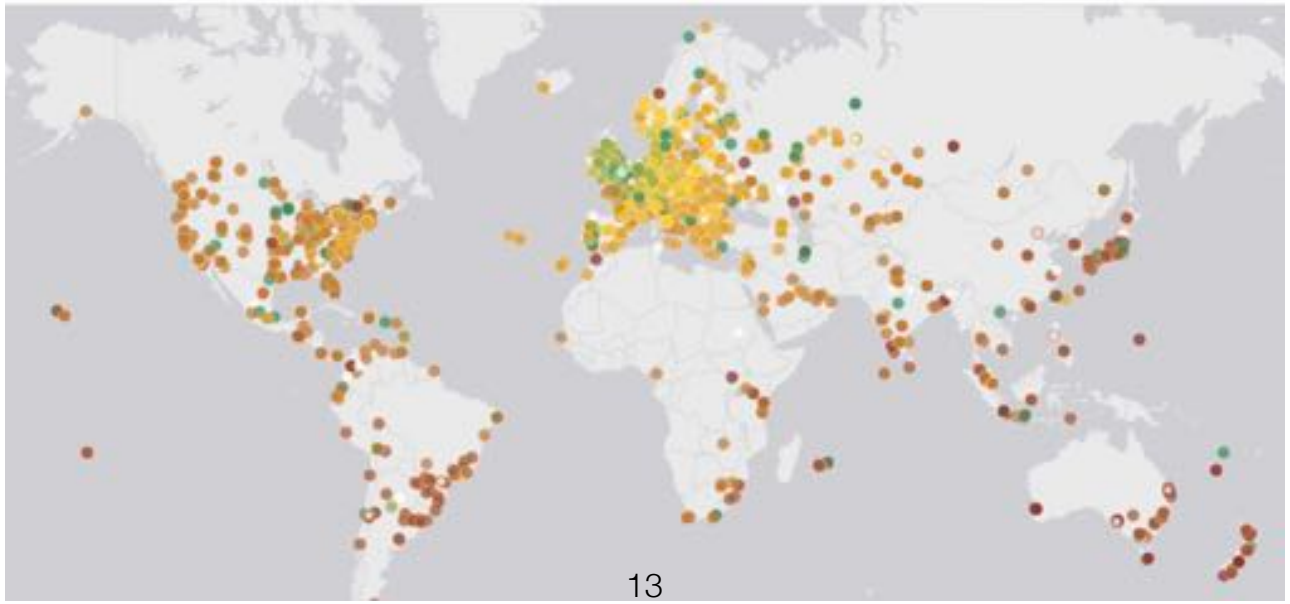

Traversal vs Size

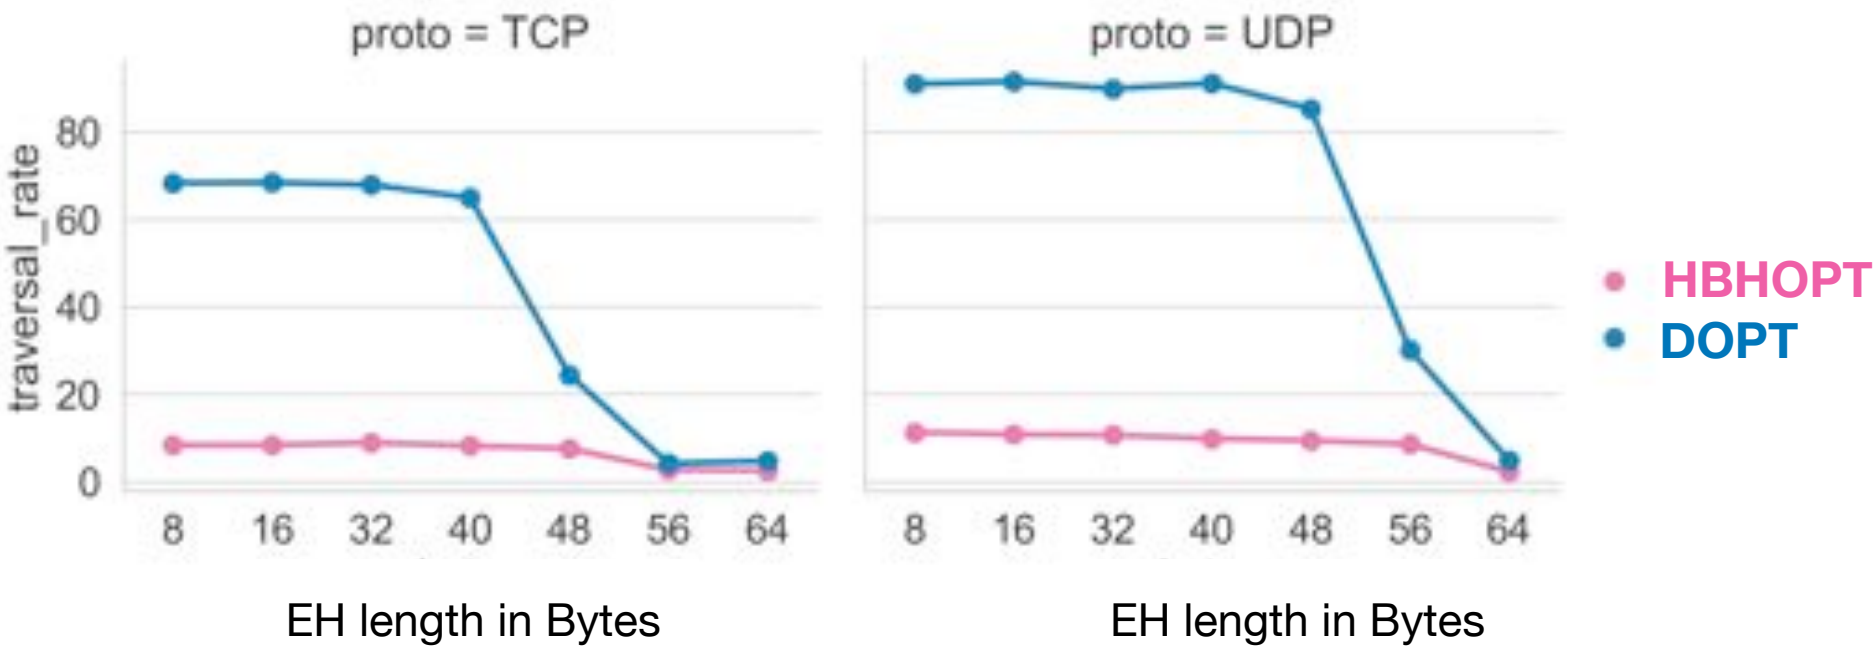

- TCP sees the biggest drop in traversal at 48B: 48 + 20 = 68B (108B total)
- UDP sees the biggest drop at 56B: 56 + 8 = 64B (104B total)
- Is this due to EH size or IPv6 total chain size?
- 40B is the max for IPv4 options

Where EHs can be used, 40B often works

Experiment 3: ECMP

- ECMP uses header information for load-balancing
- UDP to port 443 from ~850 probes
  - {**DOPT**, **HBHOPT**} + control measurement
  - **We measure 16 Paris ID variations to the same target (Flow Label + source port combinations)**

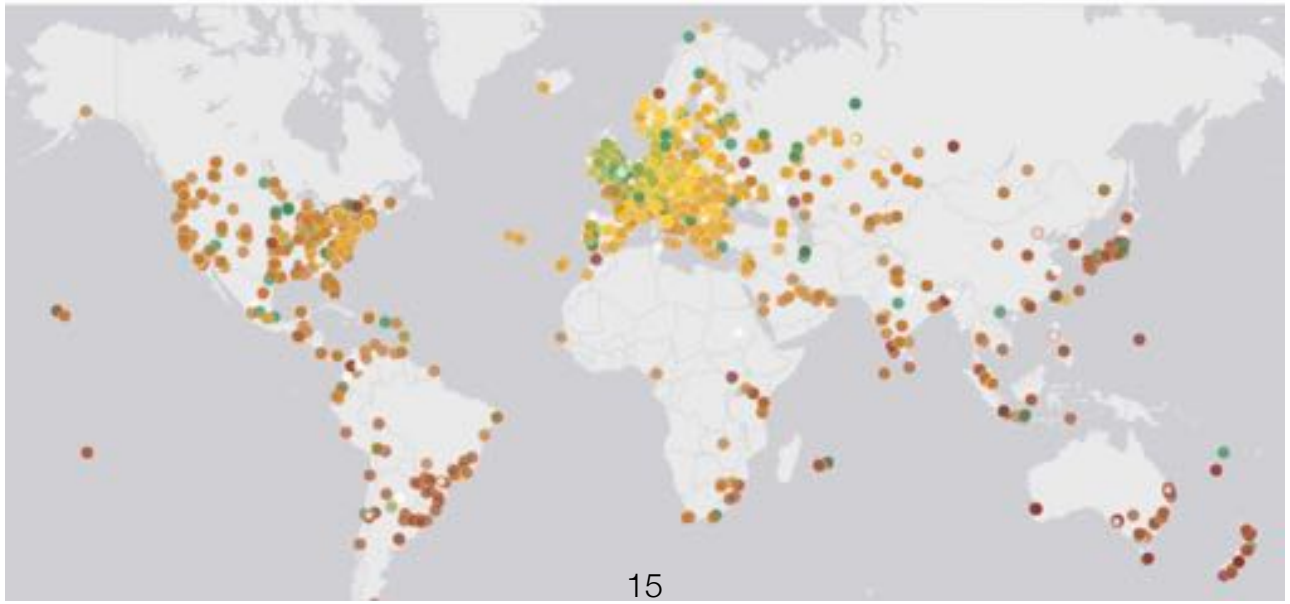

Statistics: ECMP

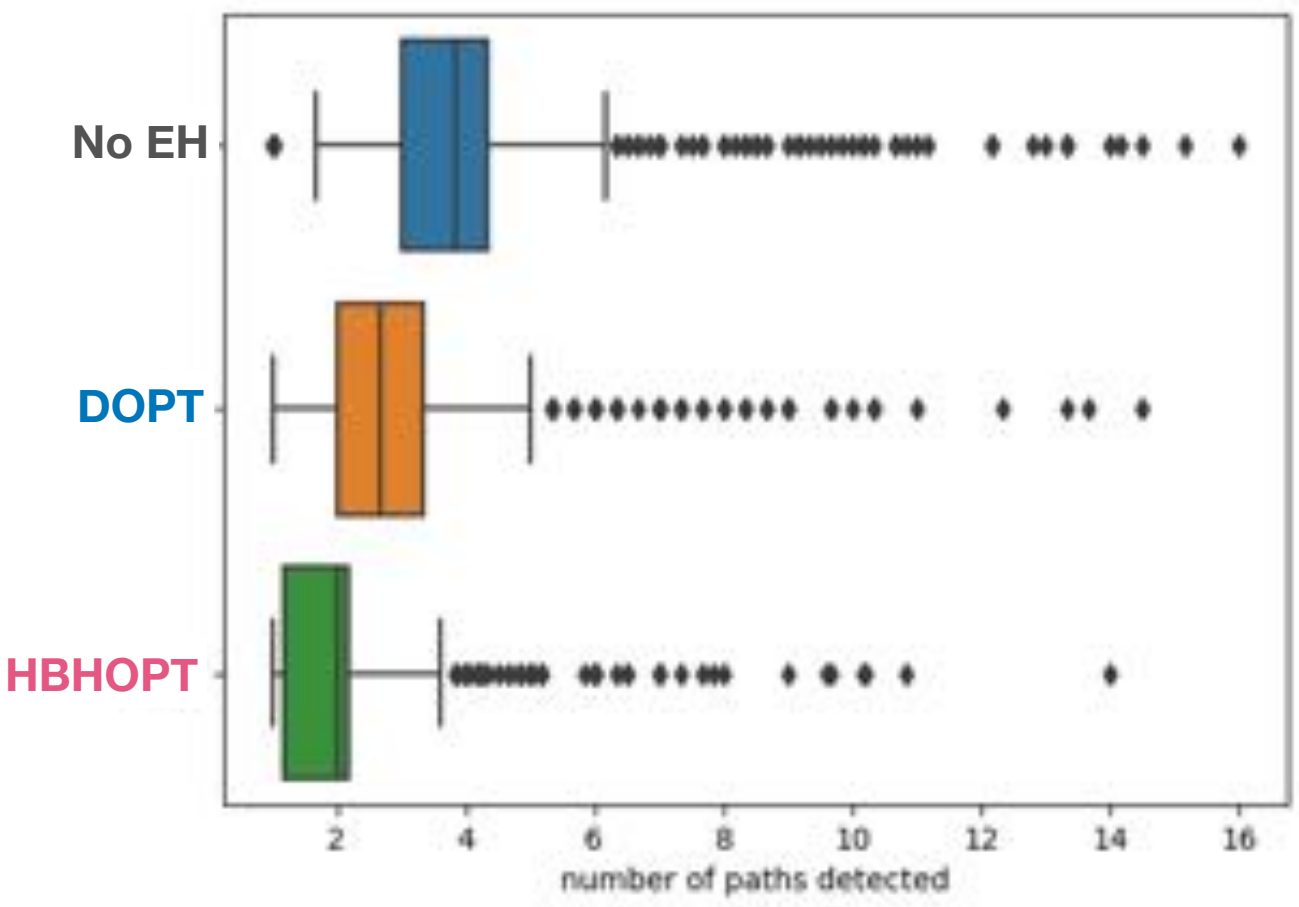

- Not all devices are equipped to handle flows that mix packets with and without EHs
  - Motivates the use of Flow Label for ECMP

Flow Label

- Can we investigate FL impact in Load Balancing
  - Do routers even look at this?
  - Does the FL help when using an EH?
- Thought a new test was needed in Ripe Atlas to control the FL in Paris measurements
  - Turns out we don’t need to: **half of all Atlas probes don’t set it!**

Project status

- Extend Pathspider to support different types and sizes of IPv6 extension headers - **Done**
- Measure EHs using Pathspider and RIPE Atlas - **Done**
- Acquire and distribute 5 hardware measurement probes to be used as Core vantage points - **Done**, software probes** due to hw supply issues :(
- Disseminate results @RIPE 86 and @IETF 116 - **Done**
- Publish a peer-reviewed measurement paper to understand IPv6 Extension Headers - **In review**

Outcomes

- Presentations at the IETF, RIPE, also in the UK @Networkshop: lots of useful discussion
- Data helped publish a paper (in review)
- Data helped progress an IETF draft
  - Found a new way to use Ripe Atlas data

..Flow Label results coming soon!

So, is it possible to extend IPv6?

- Options:

  …within a domain? It is low-risk, can be and IS done now

  …opportunistically in the Internet? DOPTs almost there

- Firewalls sometimes needed, but barriers bad for innovation

- More capable ASICs - > Forwarding + processing without impacting performance

- **Measurements help understand deployment challenges!**

References

- [1] https://www.rfc-editor.org/rfc/rfc7872
- [2] https://datatracker.ietf.org/meeting/108/materials/slides-108-6man-sessb-exploring-ipv6-extension-header-deployment-updates-2020-01
- [3] https://blog.apnic.net/2022/10/13/ipv6-extension-headers-revisited/
- [4] https://datatracker.ietf.org/doc/draft-vyncke-v6ops-james/