

Internet Measurements: IPv6 Extension Header Edition


Ana Custura*
Gorry Fairhurst*

IEPG, Yokohama March 2023



* University of Aberdeen
with funding from RIPE NCC

Active IPv6 Measurement is Difficult



- Many access networks do not support IPv6
 - e.g., mobile networks, broadband in Europe
- IPv6 server edge - hosting companies, do not always proxy IPv6 request to an IPv6 origin server
- IPv6 Top domains lists - not very diverse
- Hard to scan, but there are IPv6 hitlists

EH Measurement is *Harder*

- Some devices might not support EH to begin with
- Some hardware does not allow reading deep into a packet
 - Brokenness can be subtle, for network devices that inspect upper layer protocol information
- Network devices can be configured to filter EHs
 - e.g., at access edge or server edge
- Transit networks could be configured to filter EHs

Lots of Differing Results over the Years

- e.g., RFC 7872 [1] data showed traversal as:
 - Destination Options EH: 80-90%
 - Hop-by-Hop Options EH: 45-60%
- But, my own data [2] showed traversal as:
 - Destination Options EH: 70-75%
 - Hop-by-Hop Options EH: 15-20%
- APNIC [3] showed traversal as:
 - Destination Options EH: 30-80%
 - Hop-by-Hop Options EH: near 0%
- And then, JAMES [4] showed traversal as:
 - Destination Options EH: 94-97%
 - Hop-by-Hop Options EH: 8-9%

EH Measurements using a range of tools and vantage points

	Core	Access networks	Server Edge
Internet Core	JAMES - traceroute N. Elkins - custom FTP measurements	APNIC - Custom measurements	UoA - Pathspider RFC 7872 - traceroute N. Elkins - custom cloud measurements
Access Networks	UoA @RIPE Atlas - traceroute	N/A	Jen Linkova @RIPE Atlas - traceroute
Server Edge	N/A	N/A	N/A

Could the measurements, in fact, agree?

Examples of Different Measurement Results

Examples of Different Measurement Results

- Example 1: choice of cloud provider can influence results
- Example 2: measuring from the edge or the core does too

**Where we
measure
*from***

Examples of Different Measurement Results

- Example 1: choice of cloud provider can influence results
- Example 2: measuring from the edge or the core does too

**Where we
measure
from**

- Example 3: Top 1M lists need a per-AS breakdown
- Example 4: different target server types = different results
- Example 5: crowd-sourcing targets = different results
- Example 6: cloud provider targets = different results again

**Where we
measure to**

Examples of Different Measurement Results

- Example 1: choice of cloud provider can influence results
 - Example 2: measuring from the edge or the core does too
 - Example 3: Top 1M lists need a per-AS breakdown
 - Example 4: different target server types = different results
 - Example 5: crowd-sourcing targets = different results
 - Example 6: cloud provider targets = different results again
 - Example 7: different protocols can reveal path info
 - Example 8: the same path can reveal unexpected results
- Where we measure from**
- Where we measure to**
- How we measure**

Example 1: Vantage Points

	Hop-by-Hop Options EH UDP	Hop-by-Hop Options EH TCP
UK (JANET)	11.9%	11.5%
Canada (OVH)	19%	19.9%
Singapore (OVH)	17.4%	25.2%
Netherlands, Belarus, US, Singapore, UK, Canada (DigitalOcean)	0	0
US, Canada, Singapore, Japan, India (Linode)	0	0

Percentage End-to-End traversal for an 8 Byte **HBH Options EH**, measured in 2022

- Diverse vantage points do tell better stories!
- Digital Ocean, AWS, Linode - did not support HbH options
 - Still a valid measurement point!
 - But, cannot do wide scale measurements from here

Example 2: Vantage Points

	Hop-by-Hop Options EH UDP	Destinations Options EH UDP
Access Networks: RIPE Atlas	7-16%	77-96%
Internet core: various cloud providers	11-25%	92-97%

Percentage traversal for an 8 Byte **HBH Options EH**, from ~1000 RIPE Atlas vantage points vs 30 cloud provider vantage points, to cloud/R&E destinations, measured in 2022

- Edge v. core networks reveal a difference:
 - Networks can also differ: e.g., mobile, satellite, ...

Example 3: Destinations

	Per-Host	Per-AS
UK (JANET)	71%	92%
UK (JANET)	12%	38%
Canada (OVH)	72%	94%
Canada (OVH)	19%	59%

- Top 1X lists: Looking at just hosts can make things look better or worse that they actually are!

End-to-End percentage traversal for an 8 Byte [Destination Options EH](#), to the authoritative DNS servers for n=20082 destinations in 2867 different ASes.

- One third of all destinations are hosted by a small number of major hosting companies that do not support some EH types.
- Per-AS difference versus per-host difference for the same dataset

Examples 4-6: Destinations

Dataset	D08	HBH8
Web servers	11.88% (17.60%/20.80%)	40.70% (31.43%/40.00%)
Mail servers	17.07% (6.35%/26.98%)	48.86% (40.50%/65.42%)
Name servers	15.37% (14.29%/33.46%)	43.25% (42.49%/72.07%)

- RFC 7872: different destination infrastructure = *different results*
- Crowd sourcing destinations (APNIC): *a different story*
- FTP measurements (Nalini Elkins): *a different story*

Examples 4-6: Destinations




Figure 8 — DST option drop rate, October 2022.

- RFC 7872: different destination infrastructure = *different results*
- Crowd sourcing destinations (APNIC): *a different story*
- FTP measurements (Nalini Elkins): *a different story*

Examples 4-6: Destinations



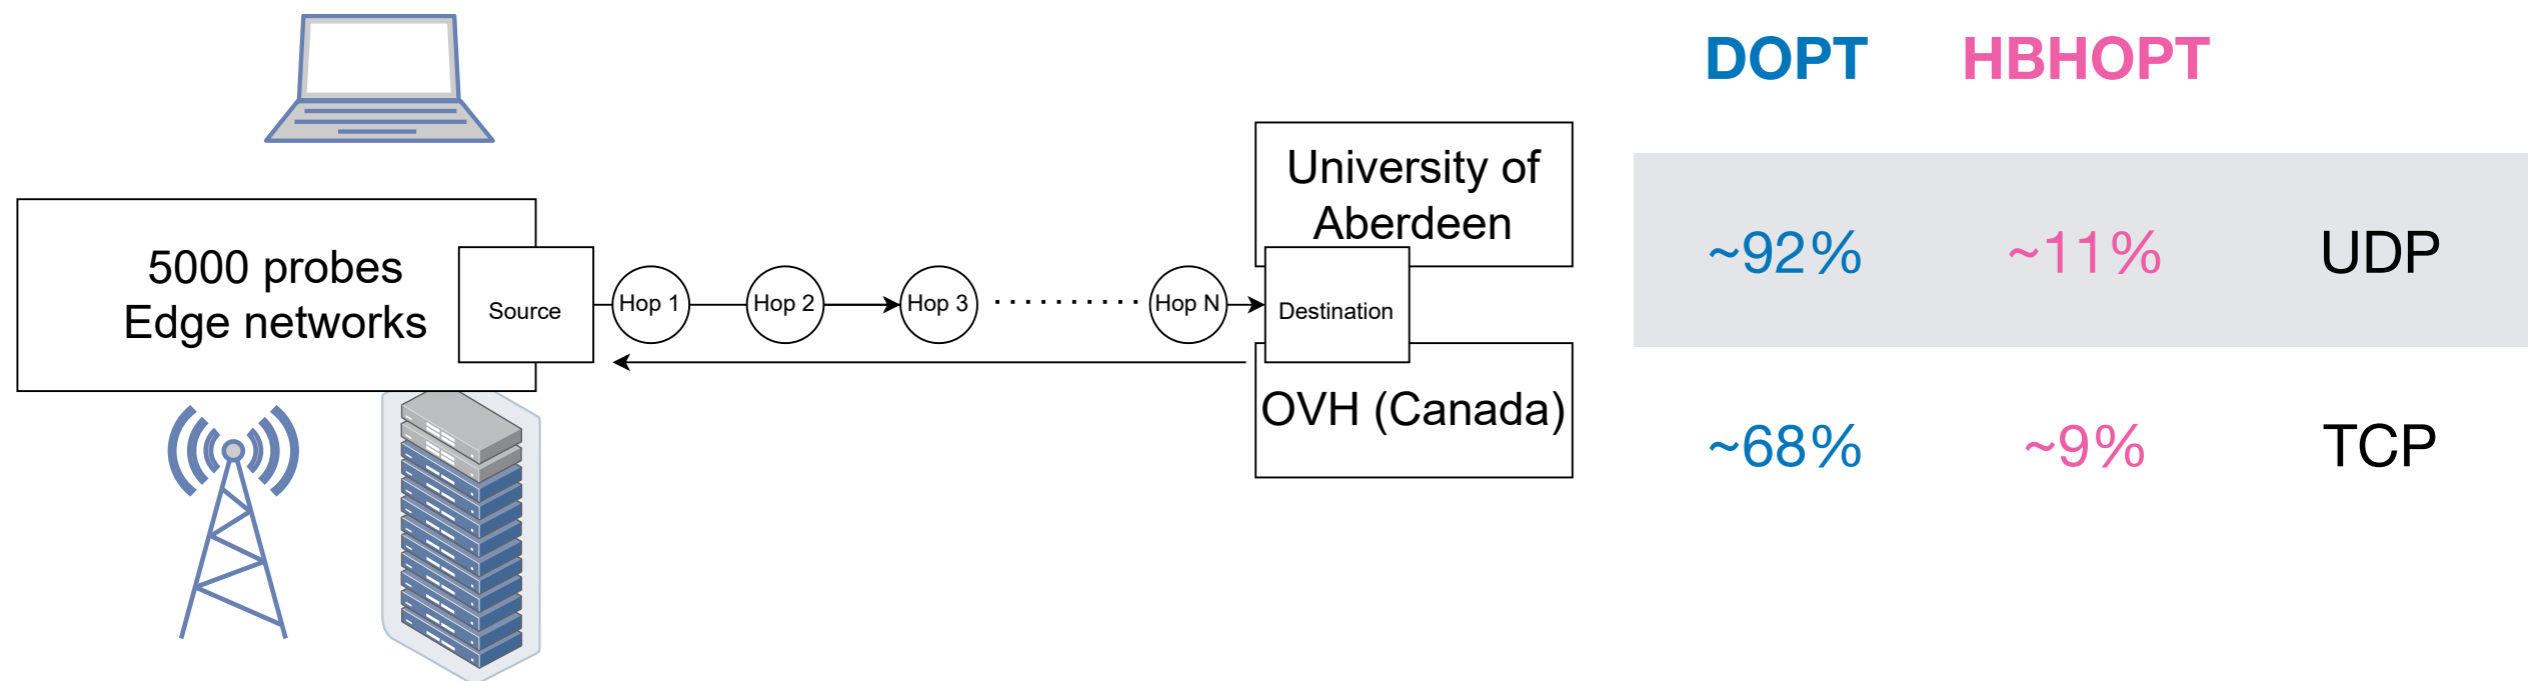
1. PDM-FTP Toronto to Warsaw - worked
2. PDM-FTP Toronto to Seattle - worked
3. PDM-FTP Toronto to Mumbai - worked
4. PDM-FTP Toronto to Melbourne - worked
5. PDM-FTP Toronto to Frankfurt - worked

Figure 6 — DSI option drop rate, October 2022.

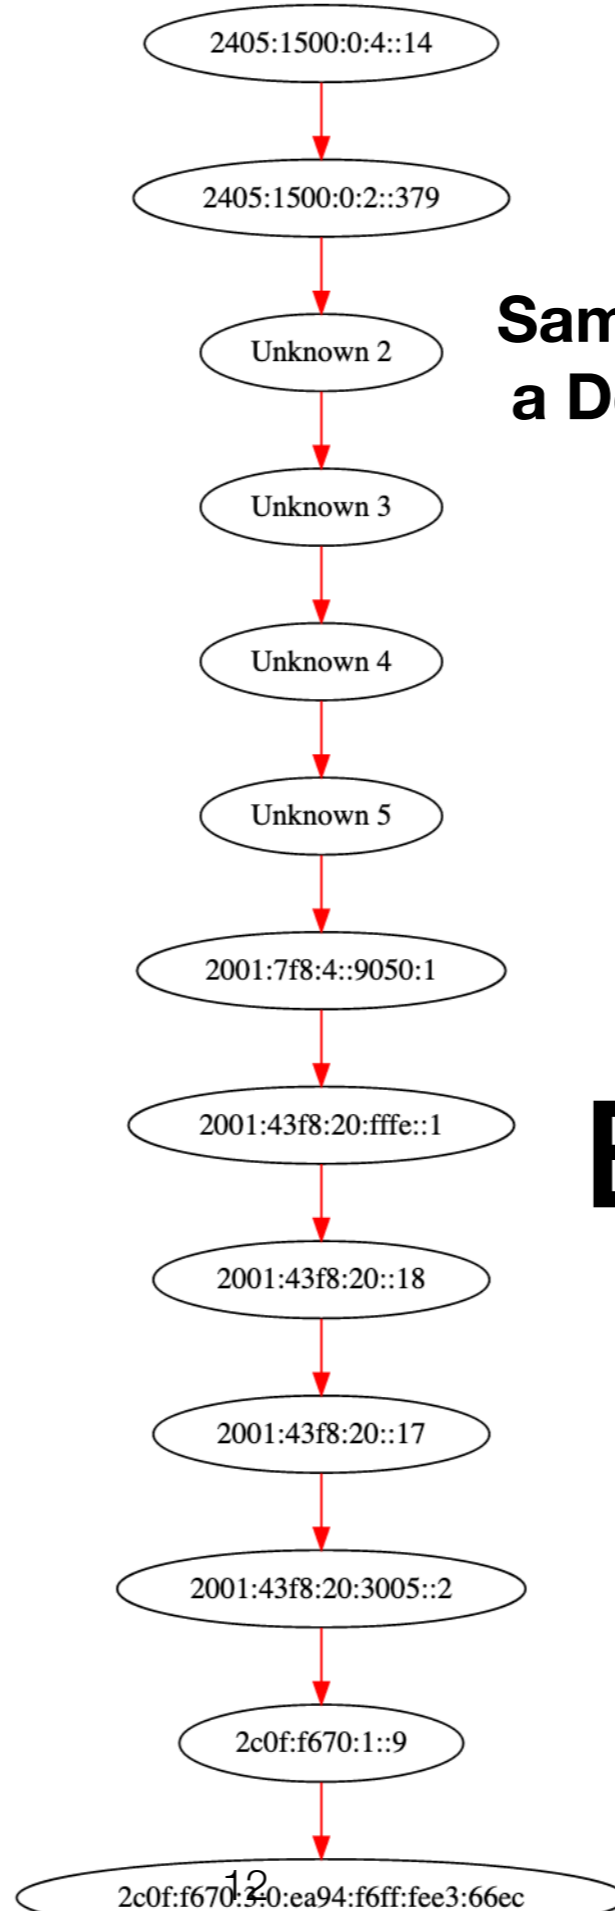
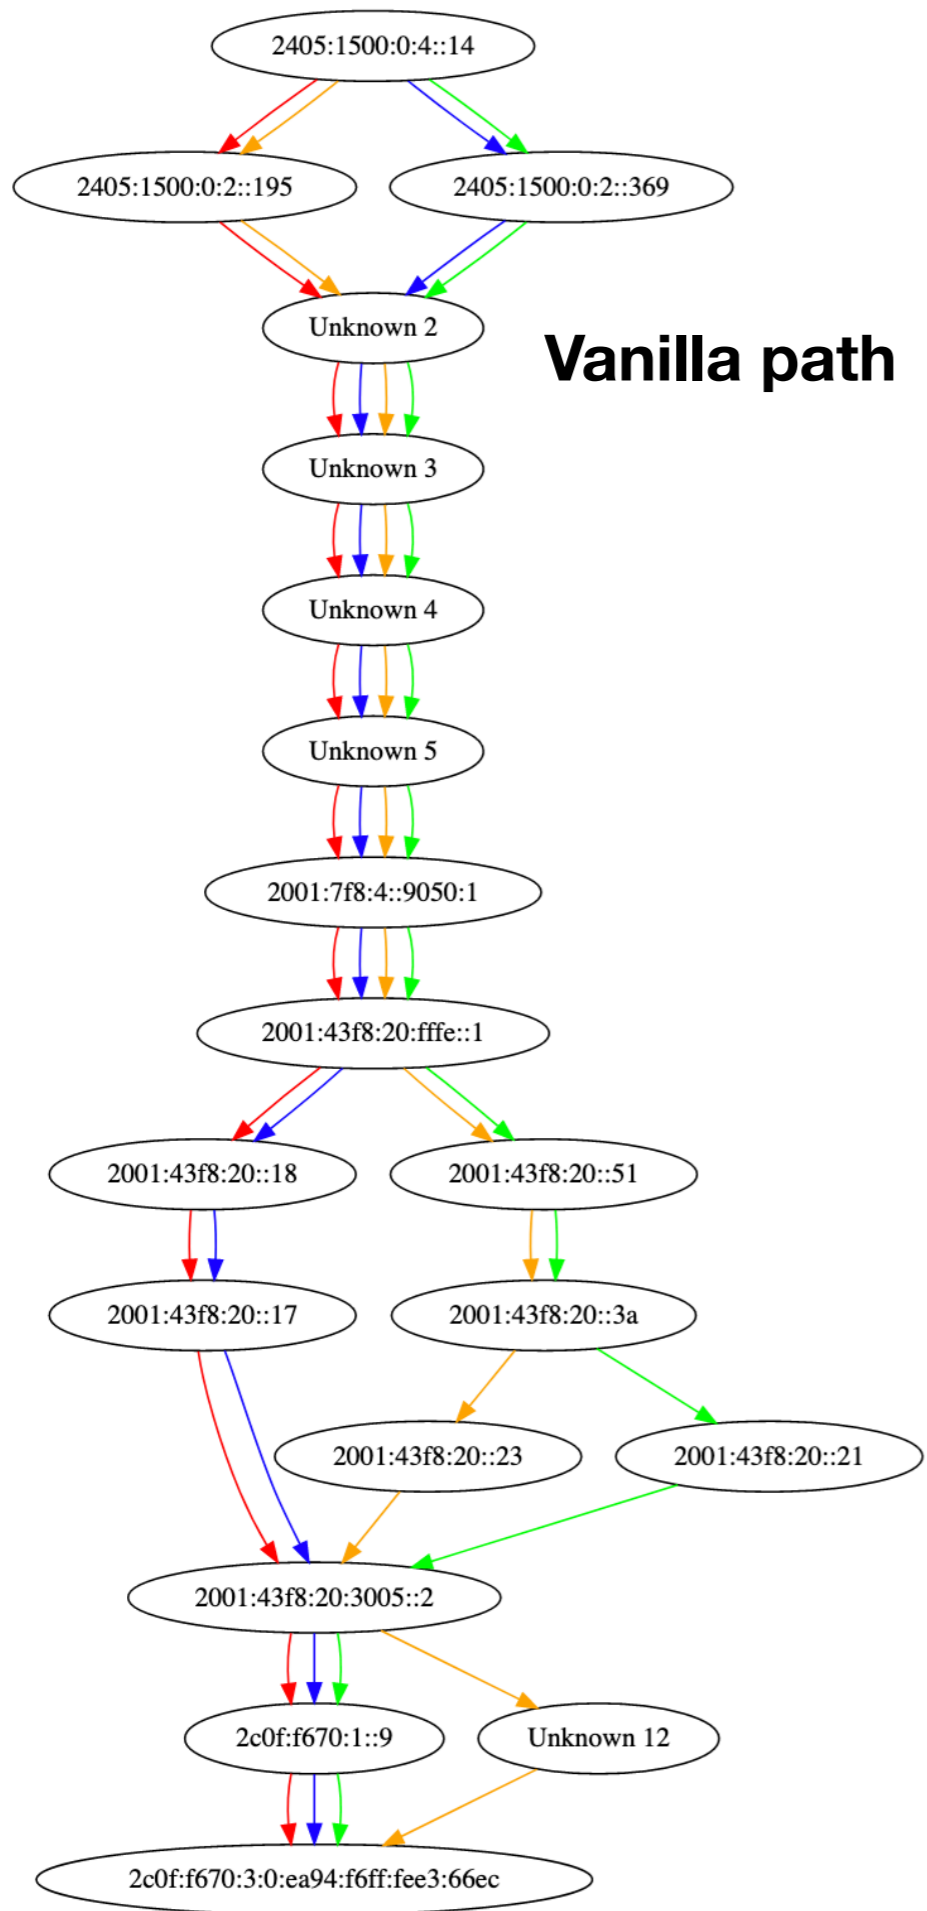
- RFC 7872: different destination infrastructure = *different results*
- Crowd sourcing destinations (APNIC): *a different story*
- FTP measurements (Nalini Elkins): *a different story*

Example 7: Protocol Differences

EH Traversal can depend on the transport protocol carried



- TCP/UDP difference for EH traversal in edge networks
 - Lots of edge devices mess with TCP; could there be a link between those devices and traversal?



Example 8: LB

Recap


	Core	Access networks	Server Edge
Core	JAMES - traceroute N. Elkins - custom FTP measurements	Apnic - Custom measurements	UoA - Pathspider RFC 7872 - traceroute N. Elkins - custom cloud measurements
Access networks	UoA @RIPE Atlas - traceroute	N/A	Jen Linkova @RIPE Atlas - traceroute
Server Edge	N/A	N/A	N/A


- Wide-scale measurement can be difficult; **our data works together!**

What next?

- Traversal depends on many factors
 - How, where, to, and when you measure influences the final result
- End-to-end + path measurements + diverse targets, destinations and protocols mitigates limitations of each way to measure !
- Still several areas for exploration:
 - Example 7: Do network devices read EH?
 - Example 8: To what extent are forwarding decisions influenced by the presence of EH?

Could you provide a home for our measurement probe?





Questions?

- [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/>