

Exploring IPv6 Extension Header Deployment - Updates 2020

Gorry Fairhurst
Ana Custura
University of Aberdeen



Our ERG Probes

We did experiments too!

- Physical probes
 - APU boards
 - Ansible configured
 - California, Slough and Aberdeen (hosted by Tim Chown, Bob Hinden and Warren Kumari)
- Other **virtual probes** in Digital Ocean Data Centres
 - *Could YOU also host one?*



Experiments

Hop-by-Hop Options Extension Header (Type 0)

Destination Options Extension Header (Type 60)

	End2End traversal tests	Traceroute-style tests
Protocol	DNS/UDP/IPv6	ICMP/IPv6
Test	Send a DNS query and tests for valid responses	Records ICMP replies from routers along the path
Targets?	IPv6 authoritative NS-es for all domains in the latest Alexa Top 1M list (=~20,000 targets)	

Experiments

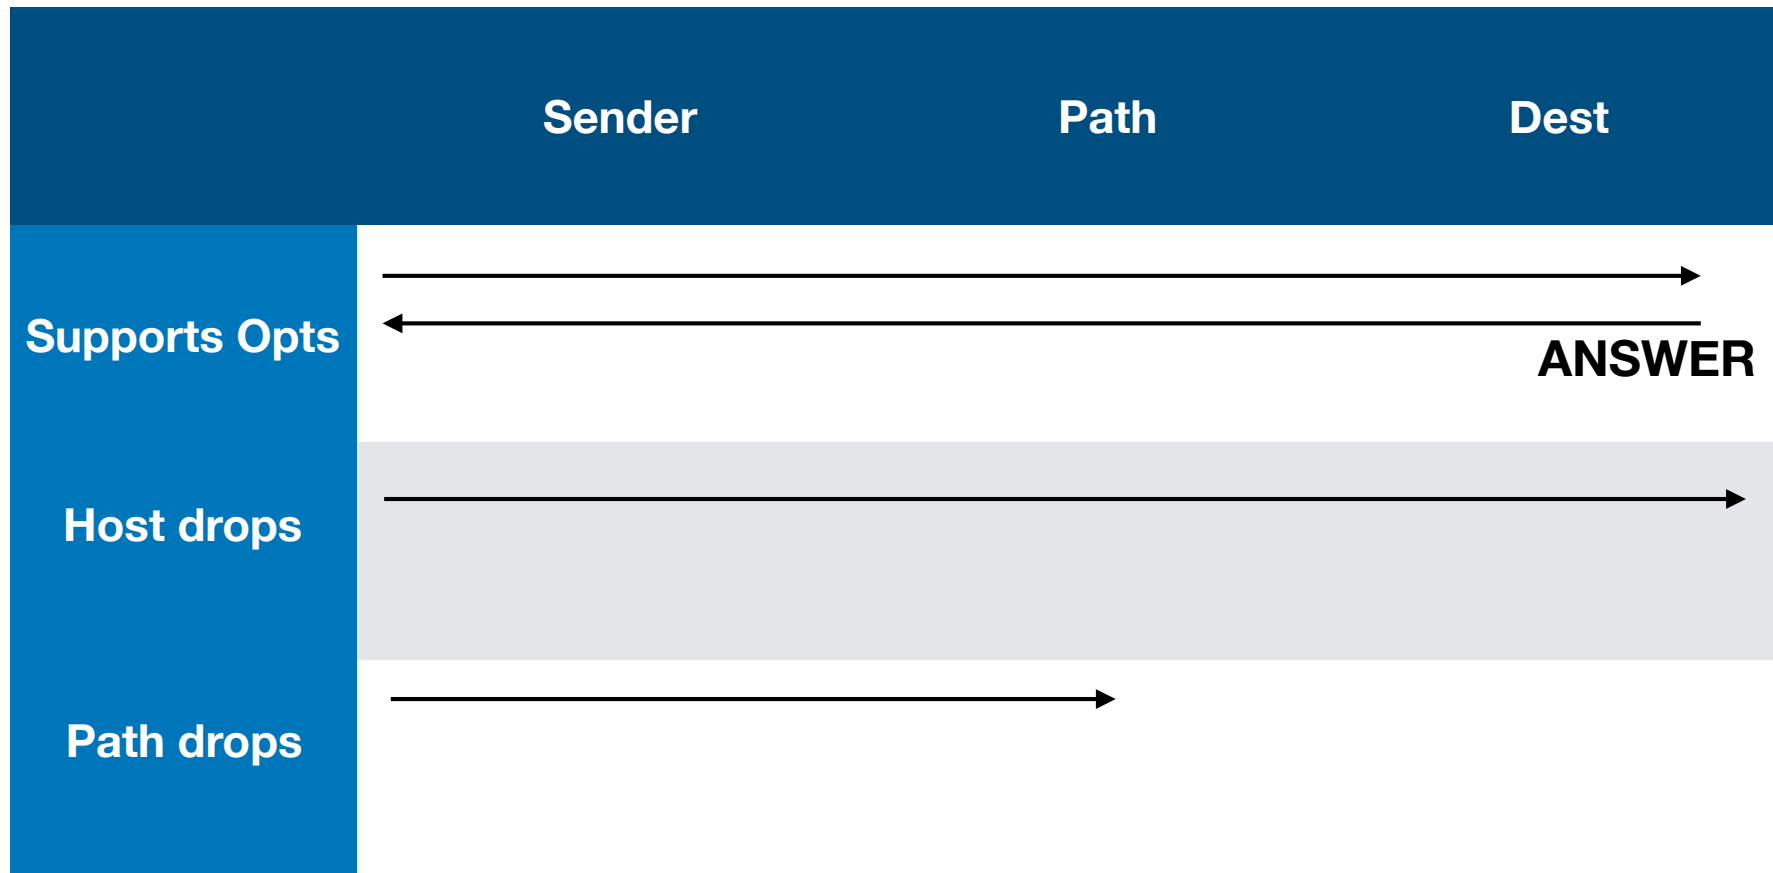
Hop-by-Hop Options Extension Header (Type 0)

Destination Options Extension Header (Type 60)

	End2End traversal tests	Traceroute-style tests
Protocol	DNS/UDP/IPv6	ICMP/IPv6
Test	Send a DNS query and tests for valid responses	Records ICMP replies from routers along the path
Targets?	IPv6 authoritative NS-es for all domains in the latest Alexa Top 1M list (=~20,000 targets)	
Additional tests	<ul style="list-style-type: none">• Valid PadN• Invalid PadN• Empty/all zeroes• Minimum PMTU (draft-ietf-6man-mtu-option)	

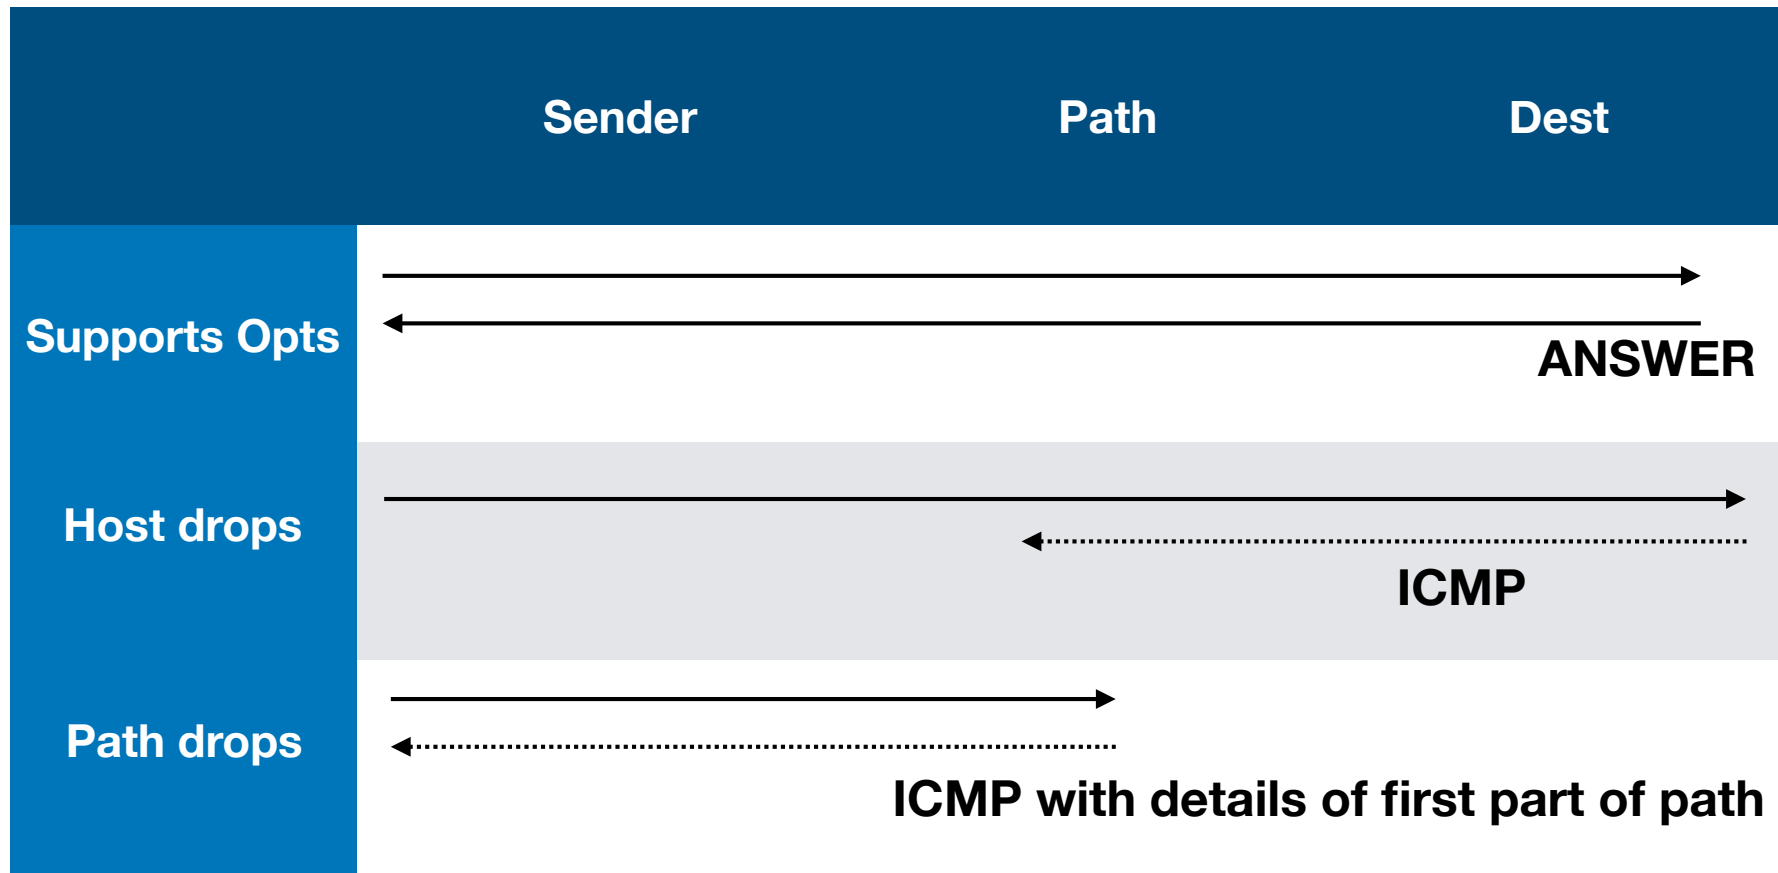
Tests - E2E Traversal

Traversal considered successful if we receive a valid DNS reply to our query



Tests - E2E Traversal

Traversal considered successful if we receive a valid DNS reply to our query



Results of E2E Server Tests

- Around 50% of DNS servers support IPv6 Destination options, reliably from all sites
- HbH traversal varies greatly depending on the vantage point! Bob's house does not pass them (sorry Comcast!) We did not look at other home access networks.
- A padN option that advertises an invalid length only traverses **0.2% of paths** - we conclude people do check!
- A new/unknown option (Min PMTU) has the same traversal rate as a well-known one (PadN)



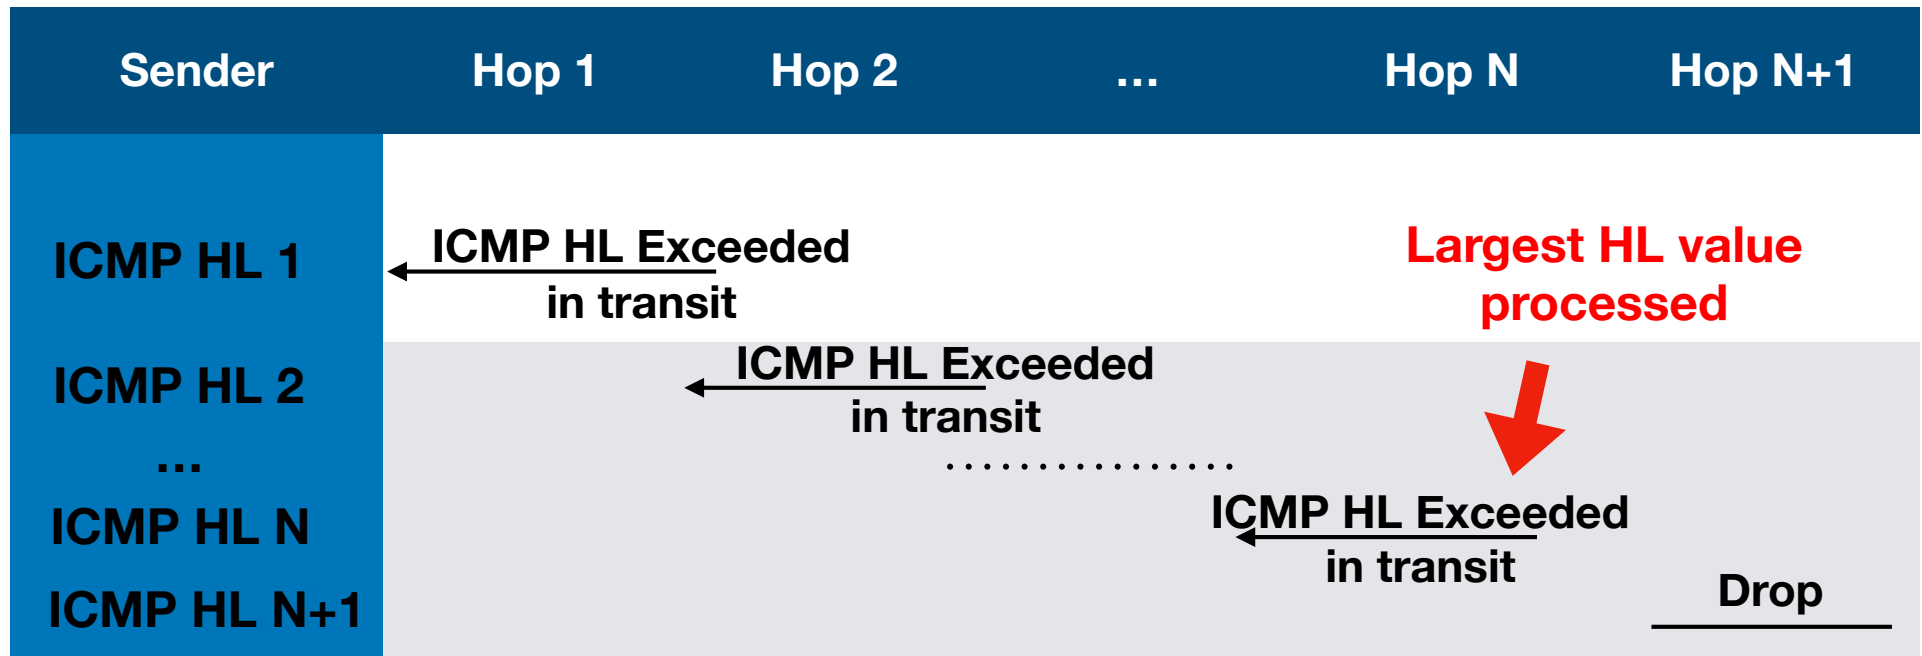
Results - E2E Traversal

Feb2020 -> Jul2020

	Aberdeen	Slough	California (Bob)	California (Warren)
Supports HBH Opts	14.9%->13.6%	22.4%->21.8%	0%->0%	36%->42.4%
Supports Destination Opts	52.7%->47.9%	53.7%->49.5%	51.9%->47.1%	52%->46.5%
None Supported	35.1%->41%	26%->34.3%	38.5%->42%	28.2%->26.2%
Both supported	14.6%->13.4%	12.6%->14.1%	0%->0%	26%->24.1%
Total servers	18296->19519	18252->19657	18261 ->19647	18250->19652

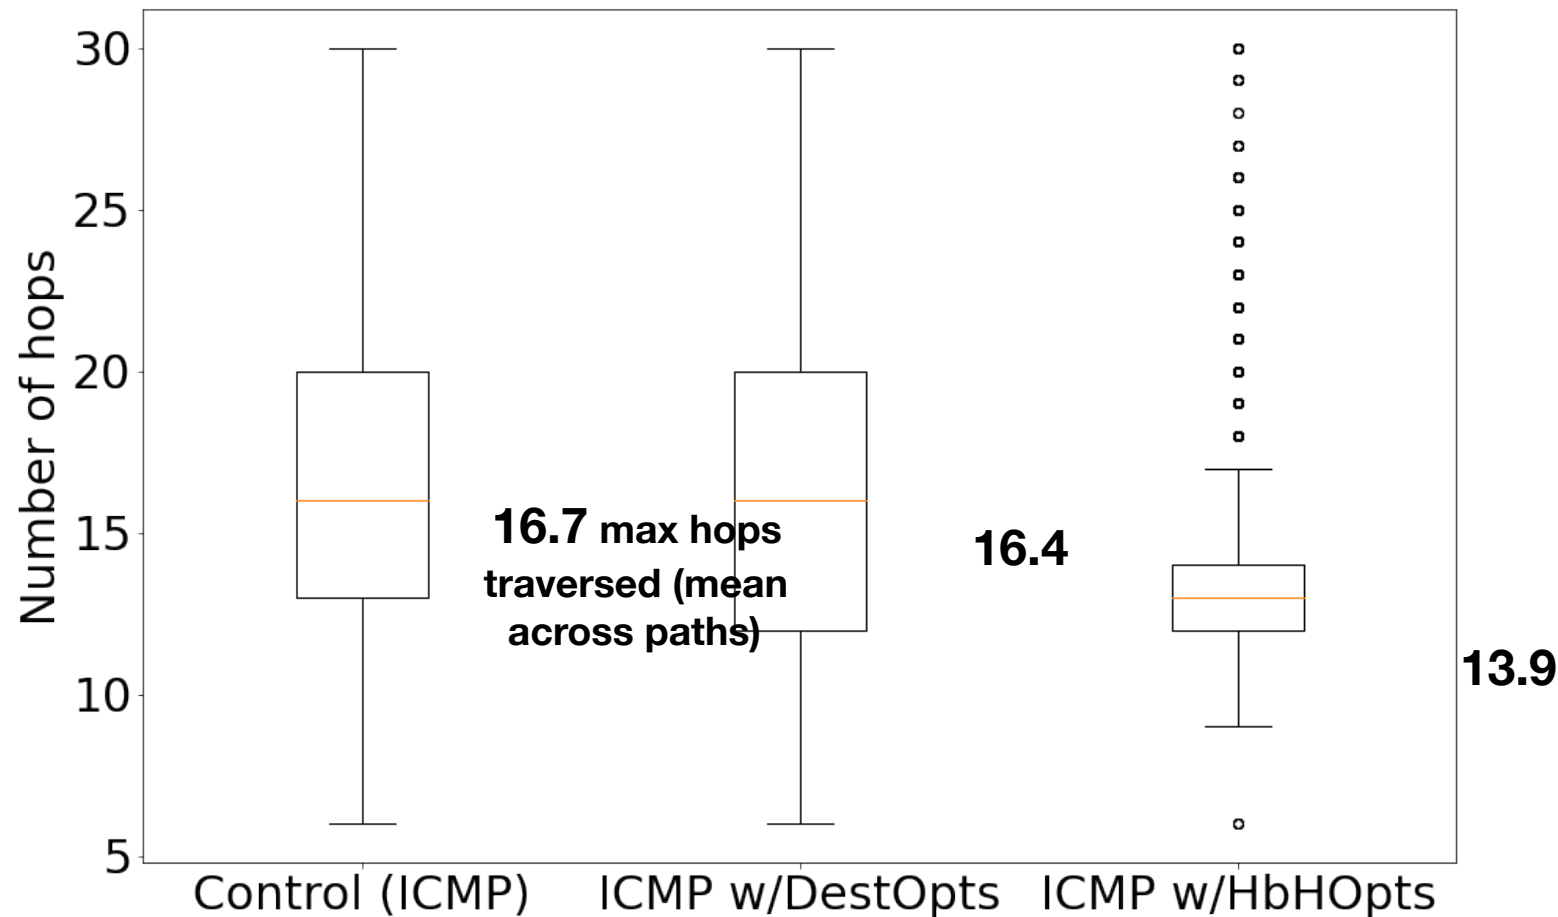
Results from testing with a valid PadN option

Results of Path Tests using ICMP



- Out of 1.8 M packets examined, none had their PadN option altered
- One hop removed the fields for both Extension Headers entirely
- Largest HL value correctly processed = Maximum hops traversed until discarded or reached destination
- We averaged maximum hops traversed from a vantage point across all paths

Number of hops traversed (max hops, from Aberdeen)



Destination Option packets are probably discarded at server edge (16.7 vs 16.4)

HbH Option packets are dropped 2-3 hops earlier on average (13.9)

Results validated from (almost) all vantage points. Sorry Comcast!

What did we Learn?

- Dest and HbH options travel “far” along a path,
 - We did not test access networks...more measurements needed!!
- In many cases options are forwarded to the destination: the Min PMTU option can be used by the DPLPMTUD algorithm
- In many cases the packet carrying the option is dropped...best to send any option on a probe packet that doesn't impact normal traffic!
 - When dropped, 2% of paths get an ICMPv6 response that quotes EH
 - Mostly, ICMP and v6 Options are silently dropped
 - The default behaviour for many firewalls (pf, ipfw & friends)!
- Nothing to suggest using this option in a DC would break anything

Next Steps for this Draft

We plan to revise the draft with what we learned.

- (1) Advise setting option on a probe packet.
- (2) Other aspects are the same as any other HbH Option.
- (3) No current API in protocol stacks to set/receive this information.
- (4) Could this ID target PS?

Virtual Probe Details

- Help us do experiments from the edge!
 - e.g. **RIPE Atlas?**
- Can you host a virtual probe?
 - Image file for a VM with 2 cores, 4096 MB of RAM and 50 GB disk, format (e.g., qcow2) can be arranged