IPv6 EH Traversal
Edge measurements

Ana Custura
ana@erg.abdn.ac.uk

06-11-2022
IEPG IETF 115

Scope of Experiments

- Tests concerned with edge networks
- Consumer Edge: RIPE Atlas probes
- Server Edge: Web + Authoritative NSes for the Alexa Top 1M domains

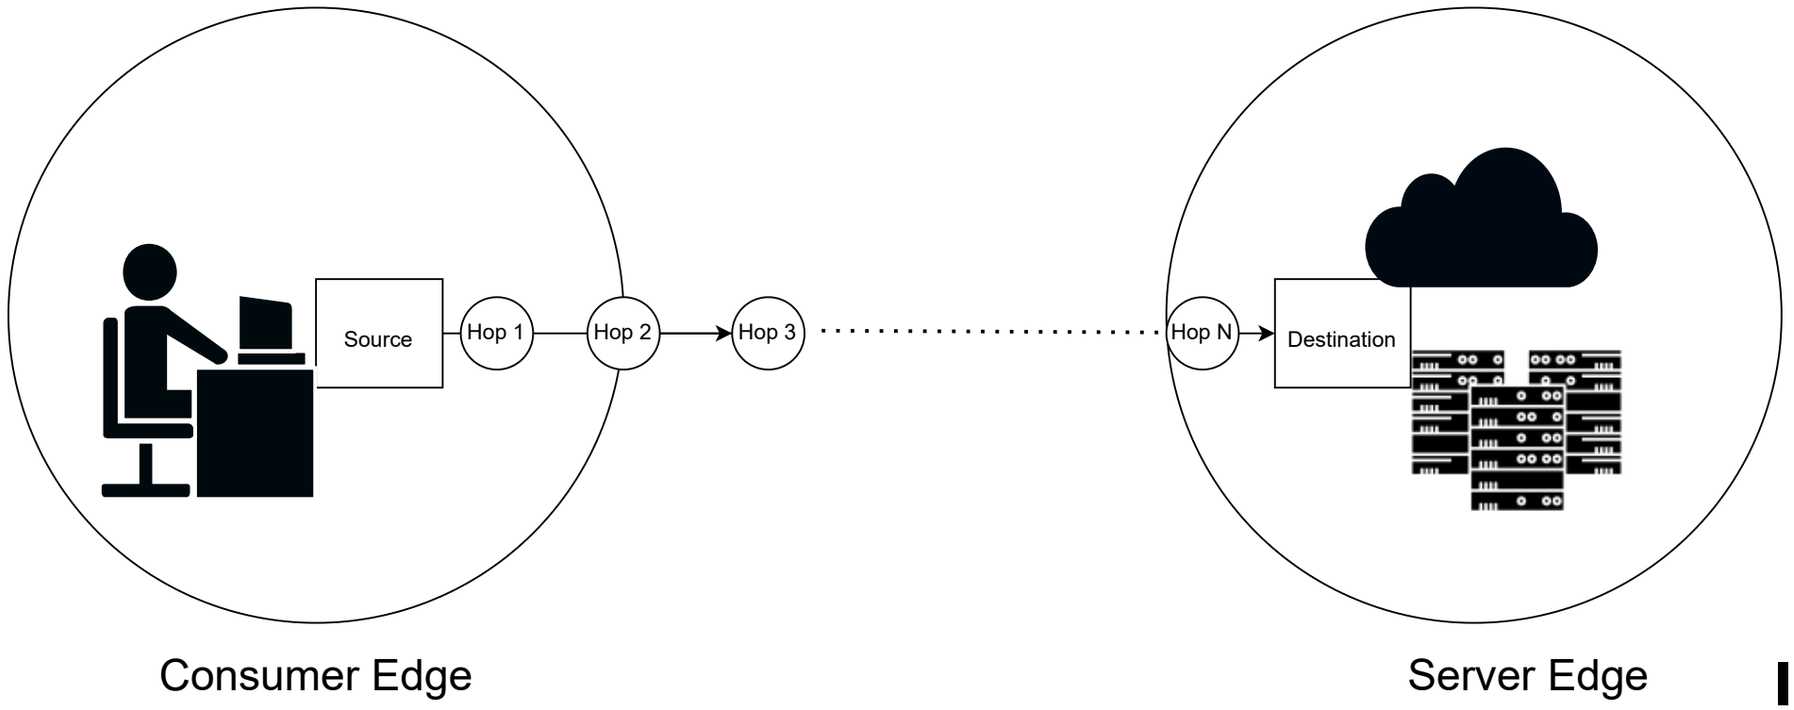

RIPE ATLAS experiments

- ~5500 IPv6-enabled probes in RIPE, globally distributed
- Tested traversal by sending packets to 2 target servers (in UK and Canada)
  - {TCP, UDP} to port 443
    - {**DOPT**, **HBHOPT**}
  - {8,16,32,40,48,56,64} B in size
  - Thanks Brian!

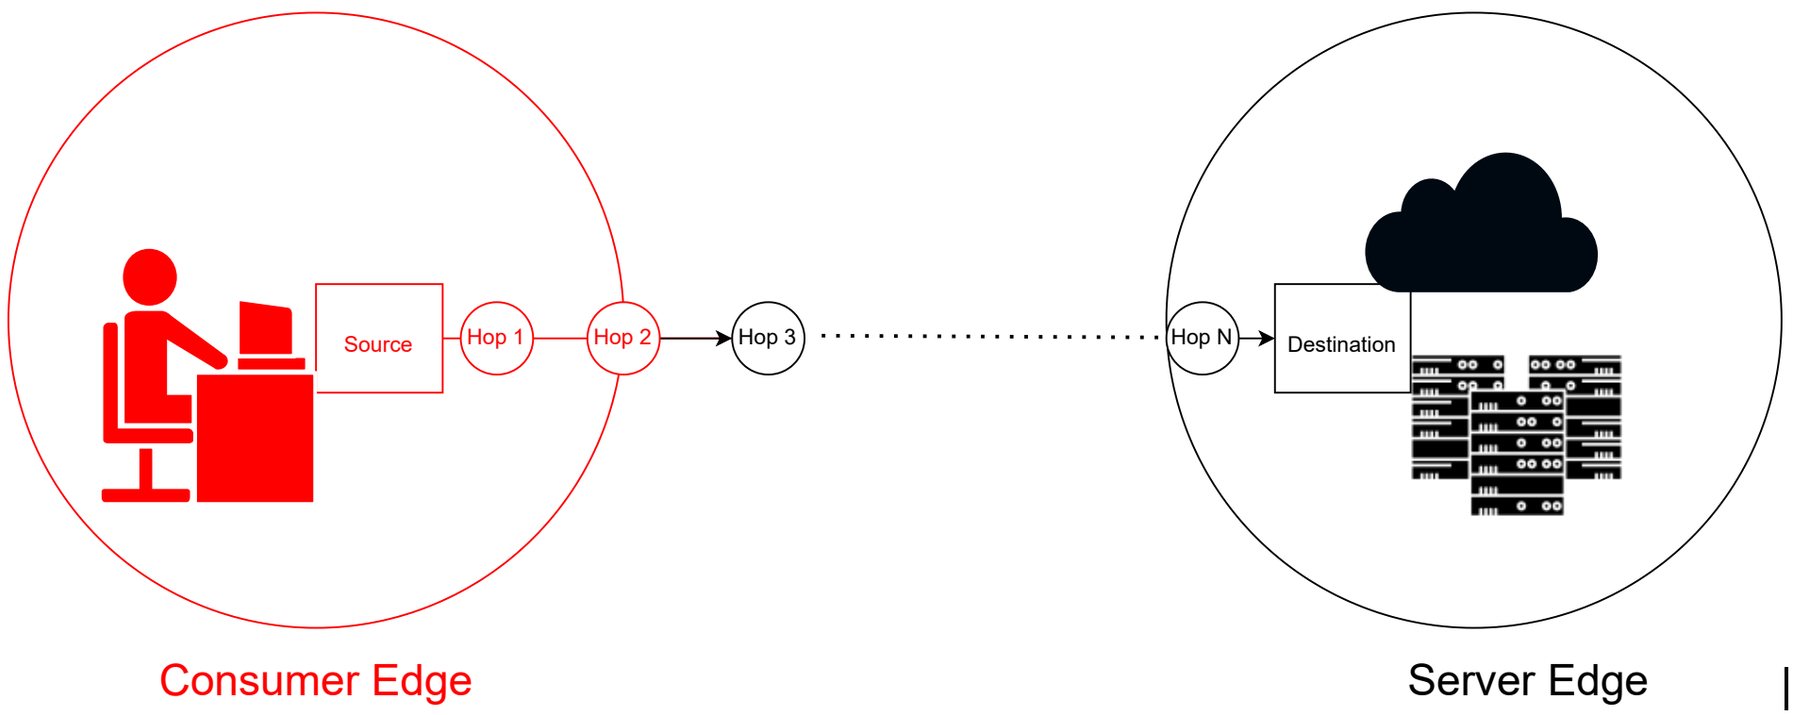

RIPE ATLAS experiments

- ~5500 IPv6-enabled probes in RIPE, globally distributed
- Tested traversal by sending packets to 2 target servers (in UK and Canada)
  - {TCP, UDP} to port 443
    - {**DOPT**, **HBHOPT**}
  - {8,16,32,40,48,56,64} B in size
- Thanks Brian!

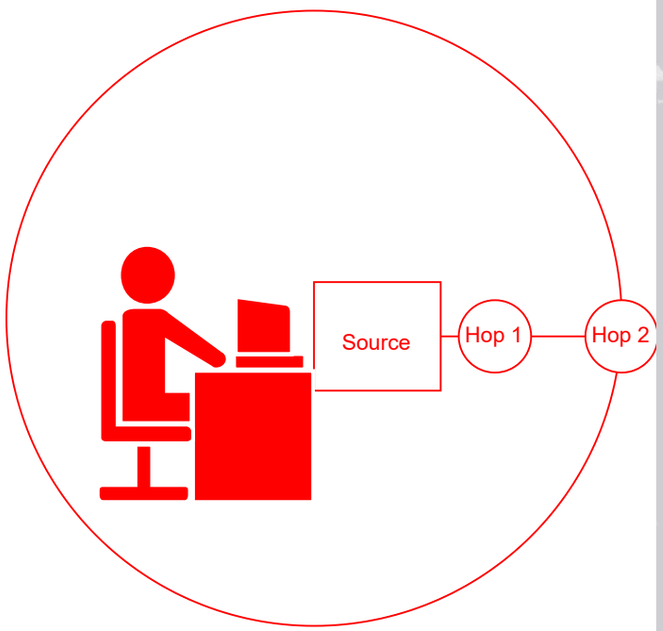

Consumer Edge

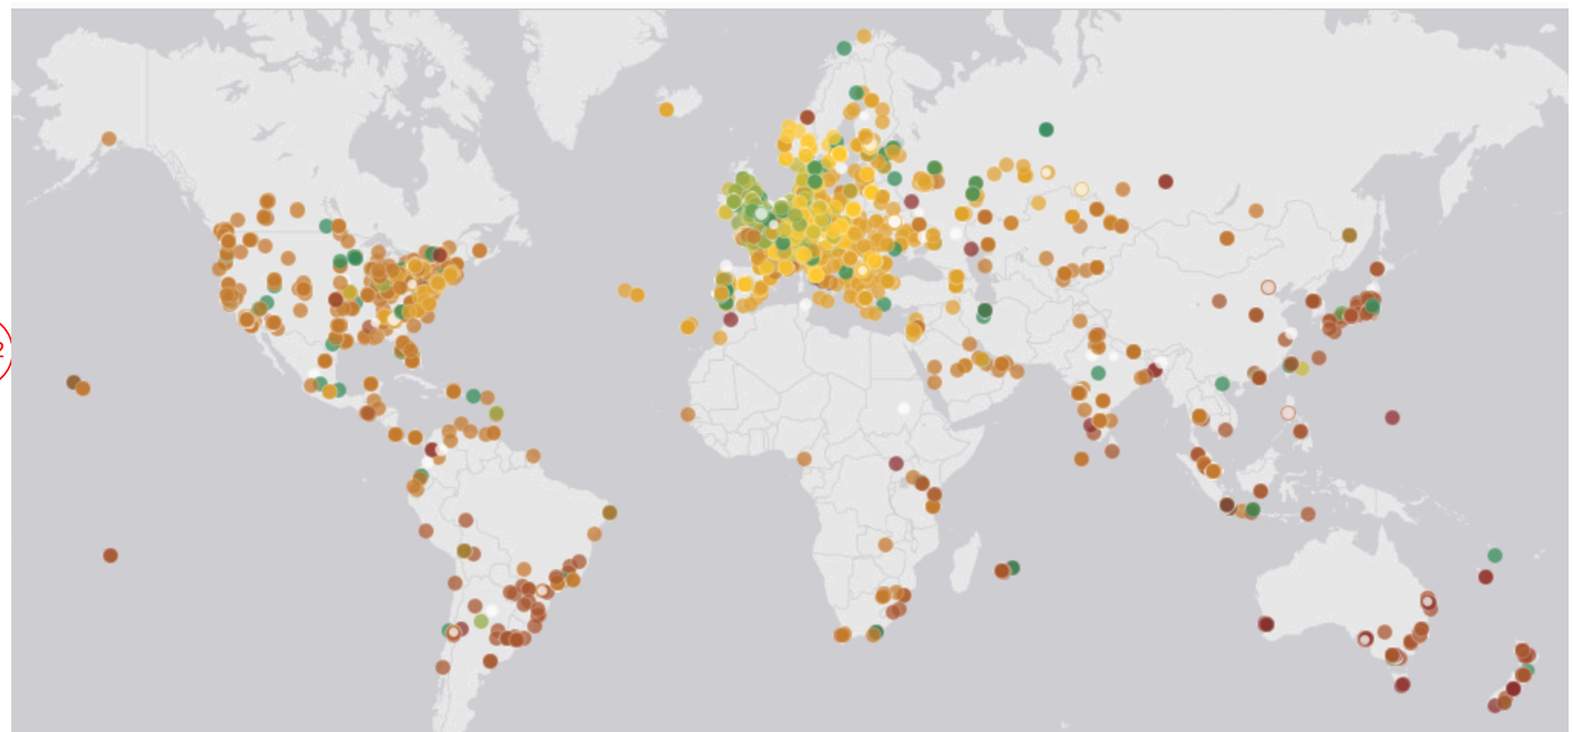

At a glance - traversal

- 8B PadN option
- High traversal for **Destination Options (DOPTs)**
- Some paths support **Hop-by-Hop Options (HBHOPTs)**
- Difference between UDP and TCP regardless of EH

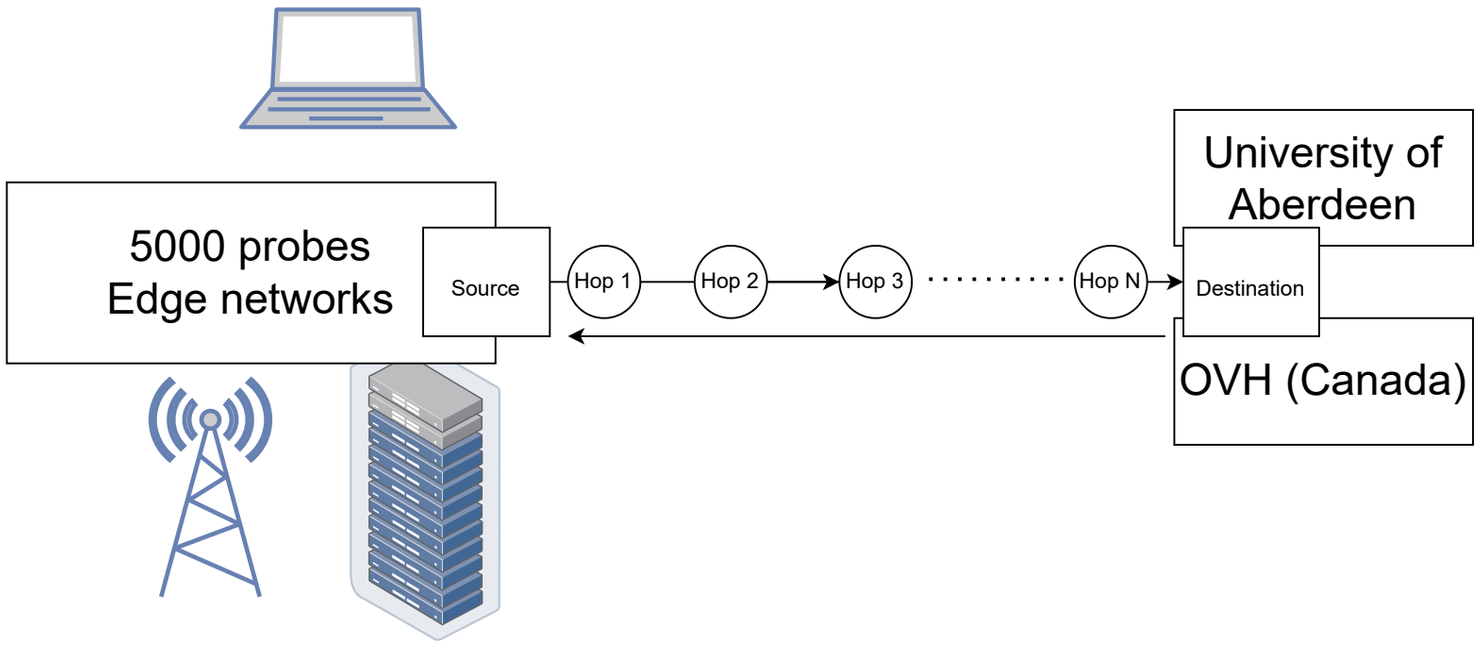

	DOPT	HBHOPT
UDP	~92%	~11%
TCP	~68%	~9%

Drops by 1st hop on the path

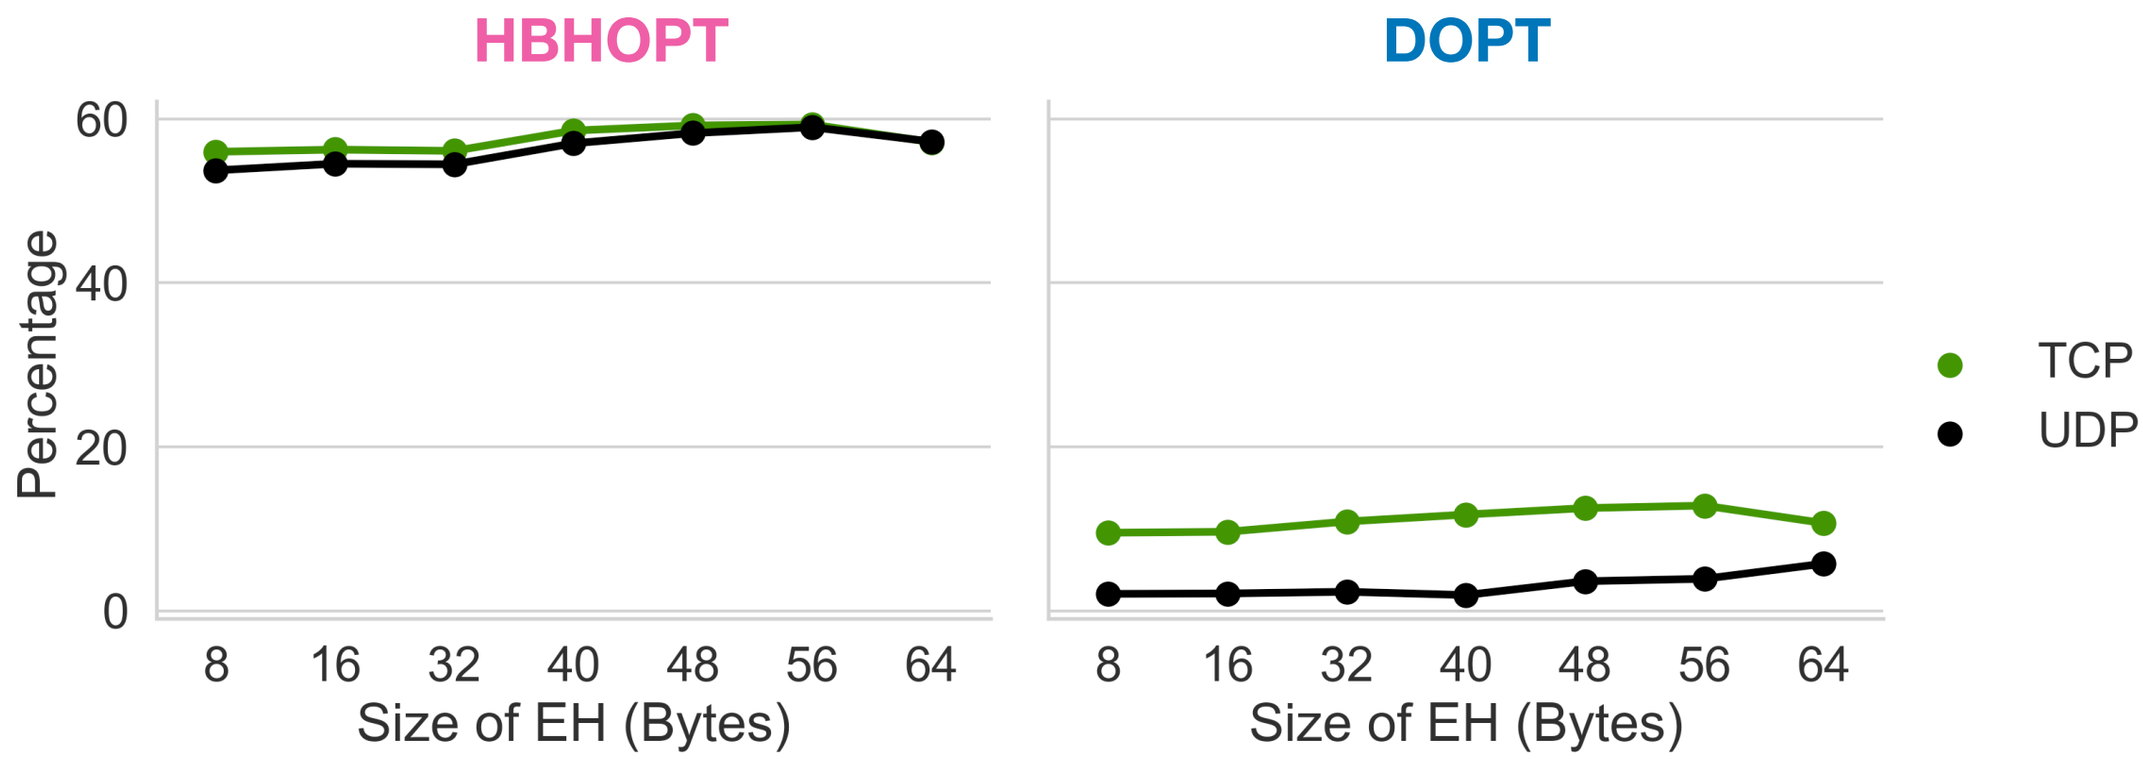

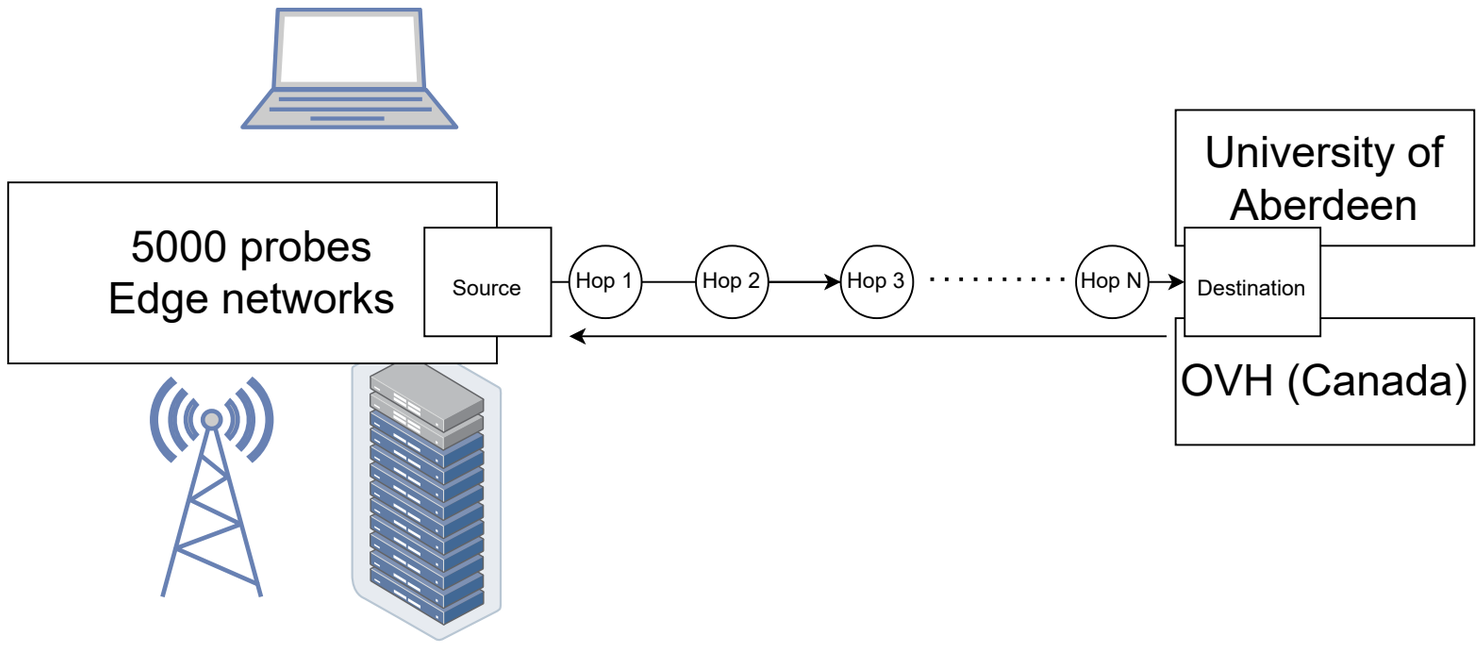

- On over 55% of the paths, HBHOPTs packets get dropped at the local router, no protocol difference;
- For DOPTs, w/TCP 10% of the paths discard packets at the first hop

Drops by 1st hop on the path

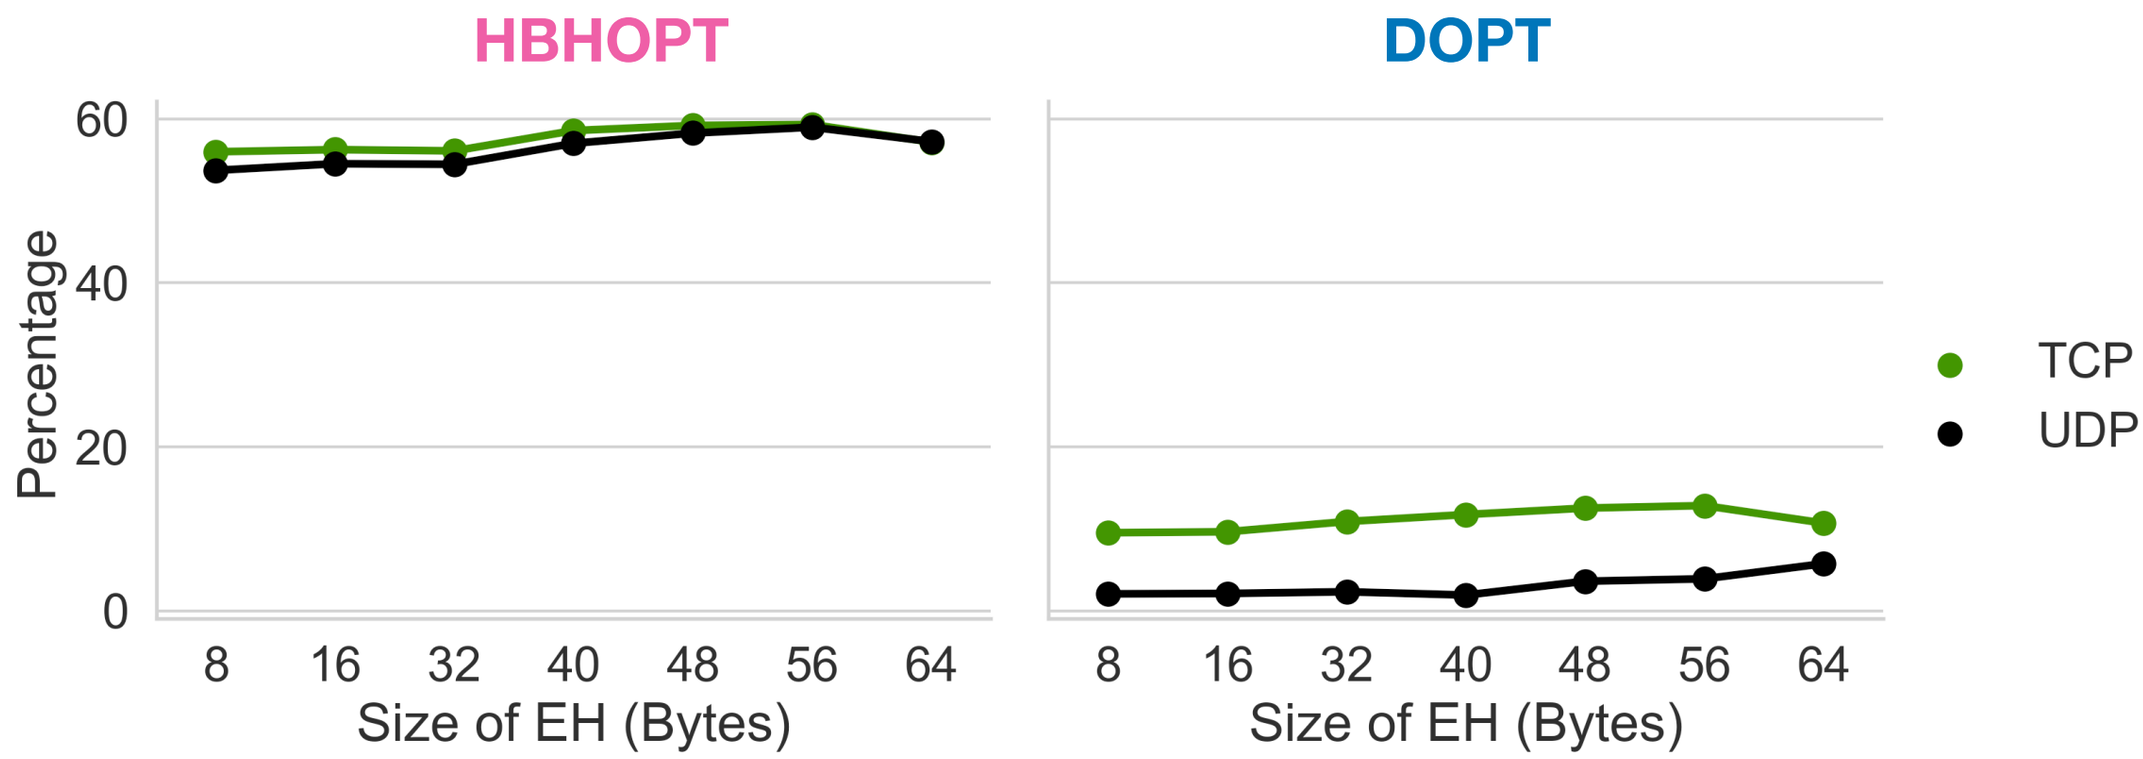

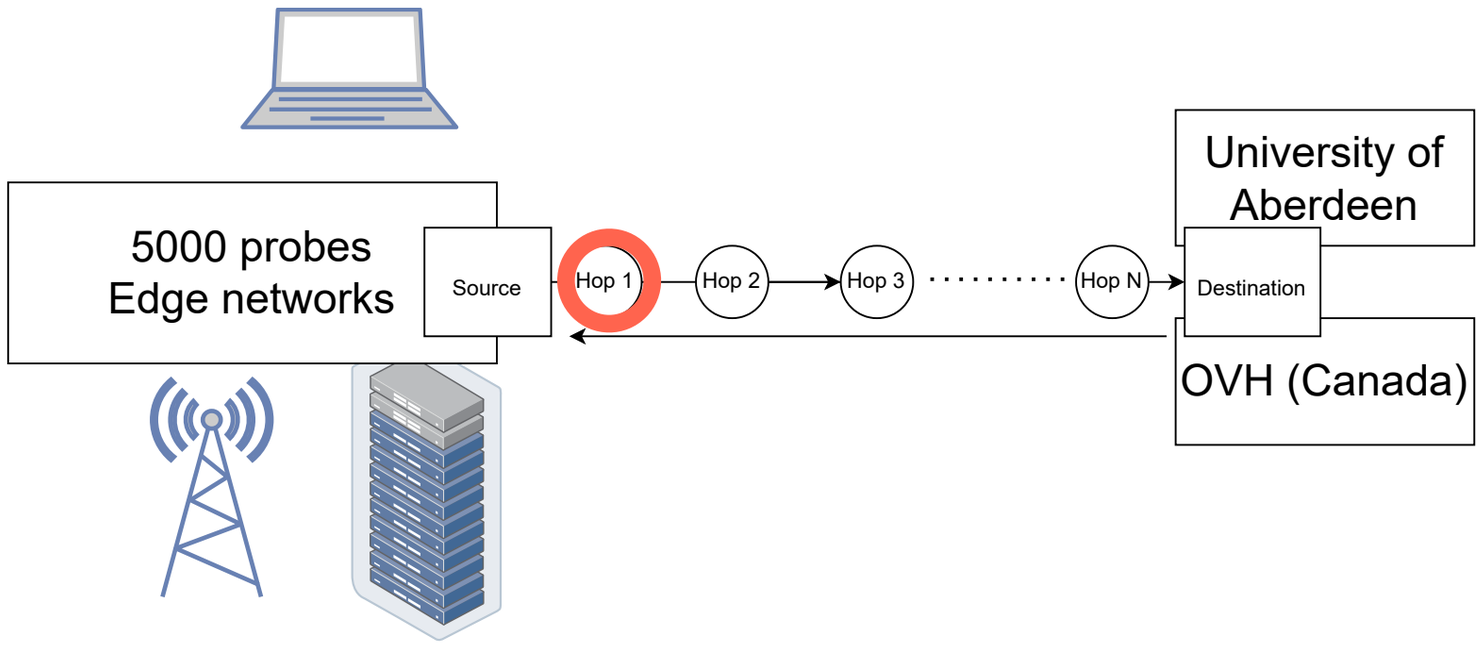

- On over 55% of the paths, HBHOPTs packets get dropped at the local router, no protocol difference;
- For DOPTs, w/TCP 10% of the paths discard packets at the first hop

IEPG IETF 115

Per-AS traversal (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.

What if packets would traverse the first AS?

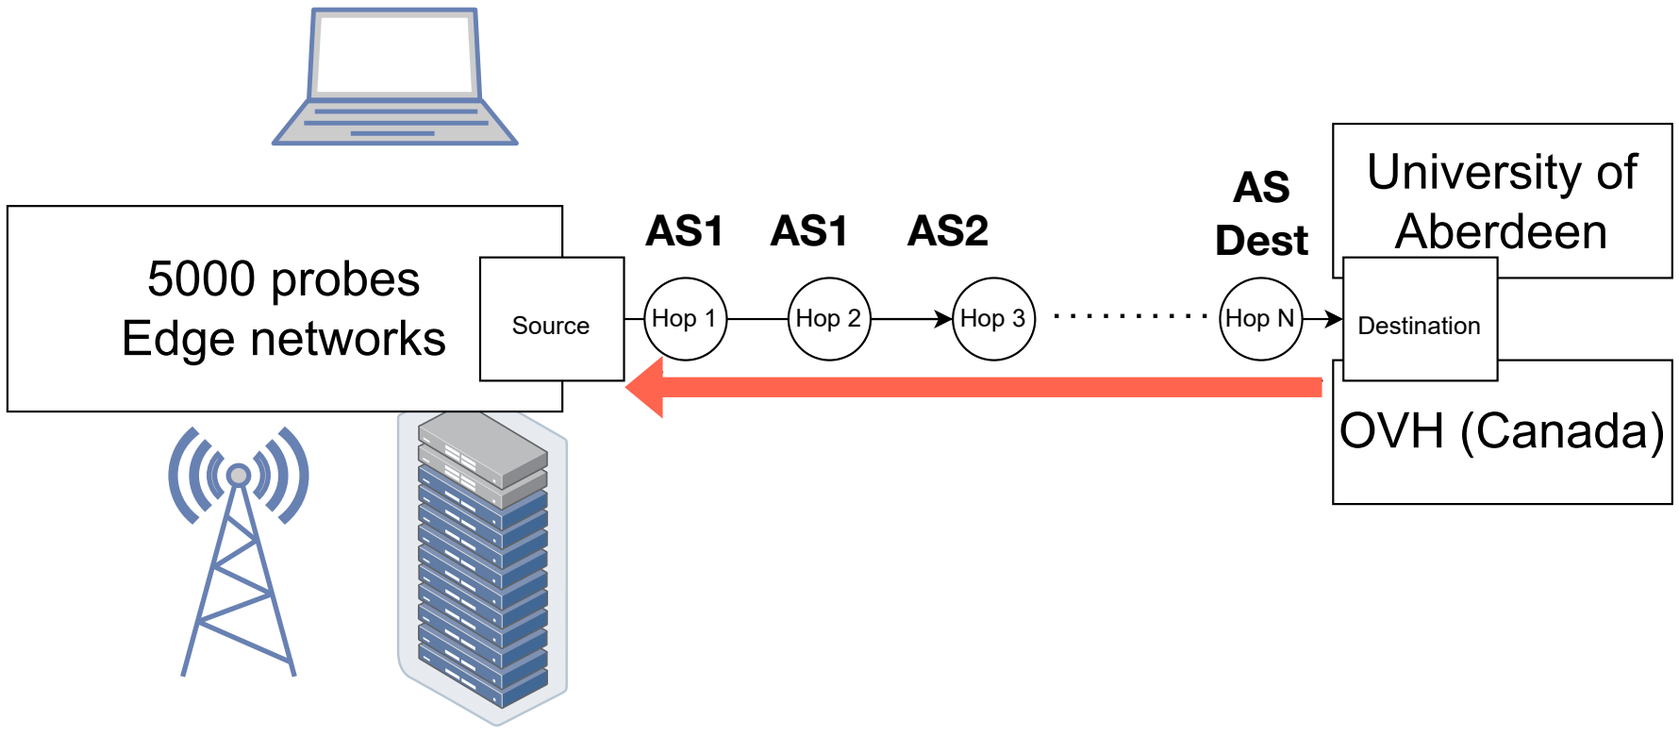

- Most probes have public IPv6 addresses
- Reverse traceroute on paths where drops happen in first AS
- Same protocol/port
- Does the packet reach original AS?

What if packets would traverse the first AS?

DOPTs

Reverse traceroute on paths where drops happen in first AS (n=271 paths for UDP)

	% make it back to original AS	%predicted traversal
DOPT UDP (UK)	97.6%	**~96%**
DOPT TCP (UK)	TBD	TBD
DOPT UDP (Canada)	95.3%	**~96%**
DOPT TCP (Canada)	TBD	TBD

HBOPTs

Reverse traceroute on paths where drops happen in first AS (n=3150 paths for UDP)

	% make it back to original AS	%predicted traversal	Notes
HBHOPT UDP (UK)	10%	**~17%**	60% packets get dropped at LINX peering
HBHOPT UDP (Canada)	17%	**~25%**	

Traversal vs size

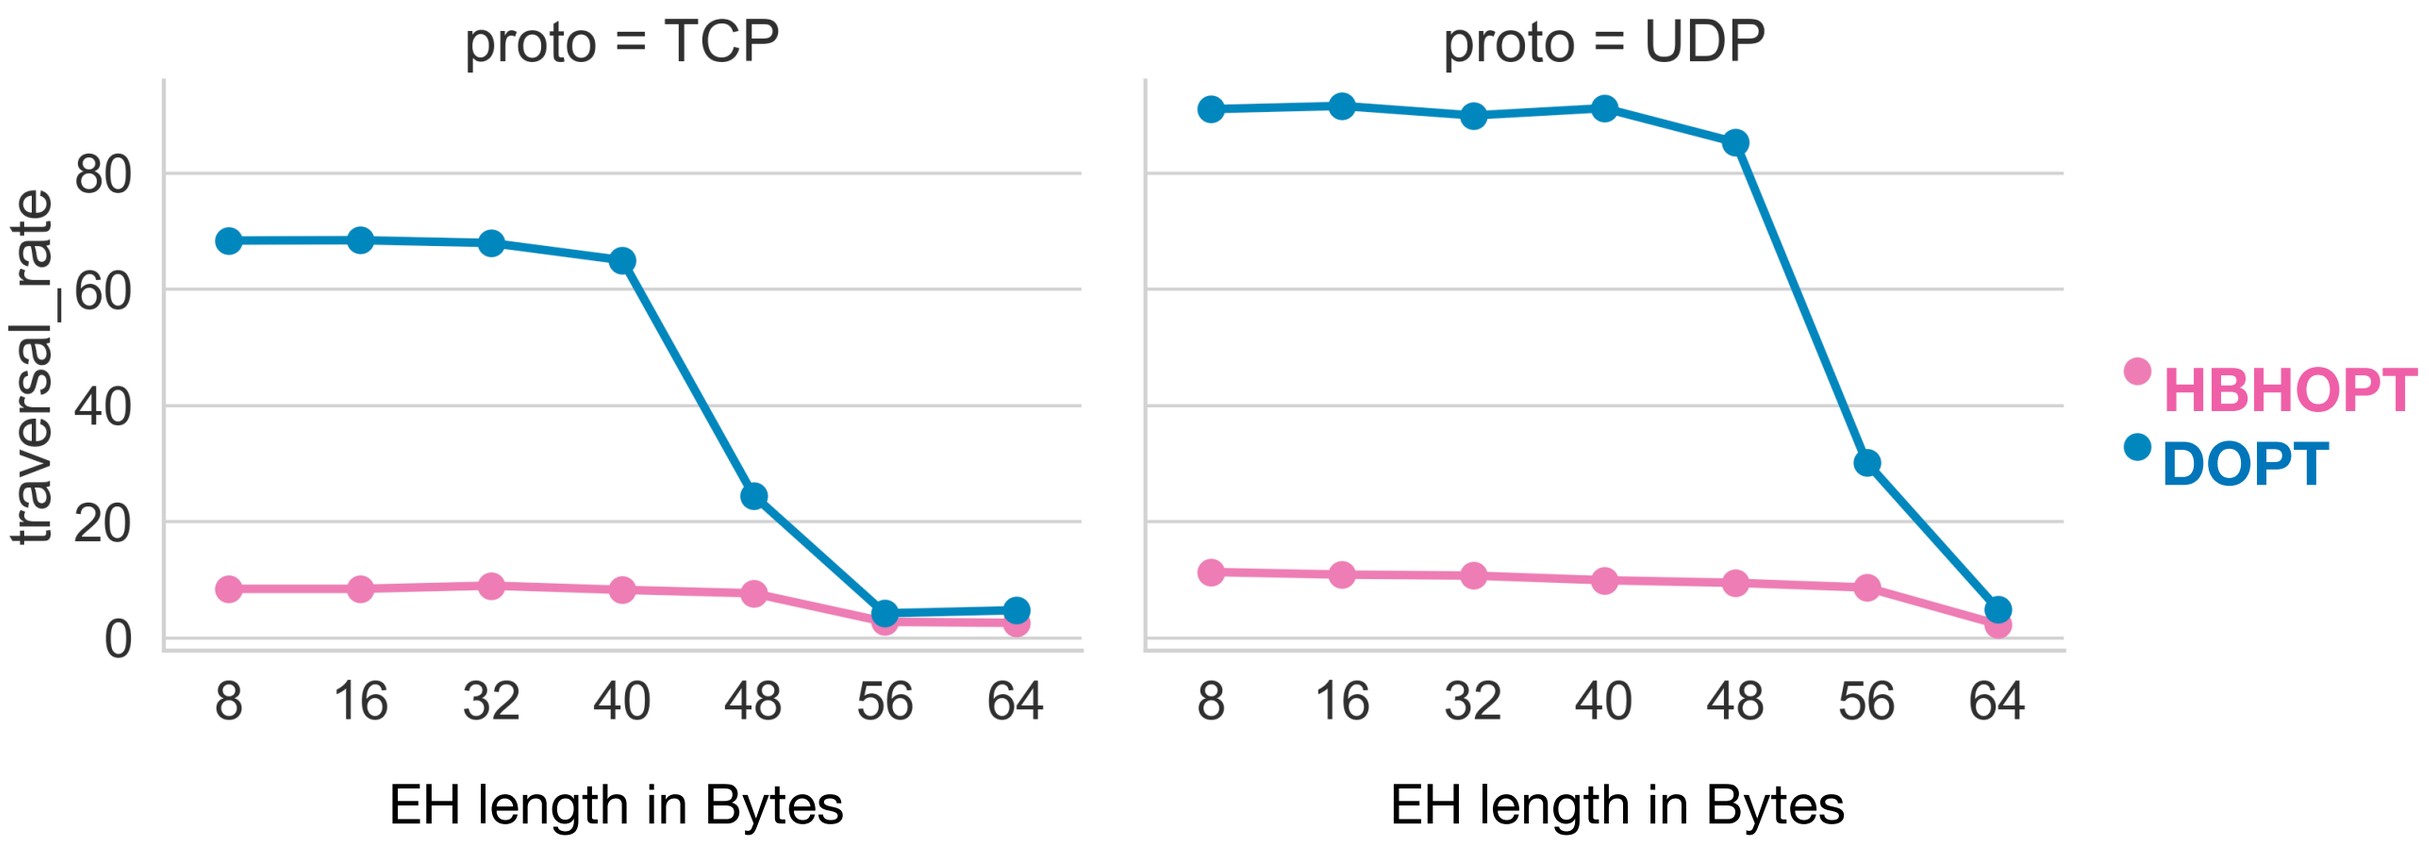

- TCP sees the biggest drop in traversal at 48B: 48 + 20 = 68B (108B total)
- UDP sees the biggest drop at 56B: 56 + 8 = 64B (104B total)
- 40B - the magic number?
- 40B is the max for IPv4 options

PathSpider Experiments

- 12 globally distributed vantage points
- DNS/{UDP,TCP}/IPv6
    - {**DOPT**, **HBHOPT**}
    - IPv6 authoritative NS-es for domains in Alexa Top 1M list (=~20,000 targets)
    - Test for valid DNS response!!

Tests:

- Valid PadN
- Invalid length
- Unknown option

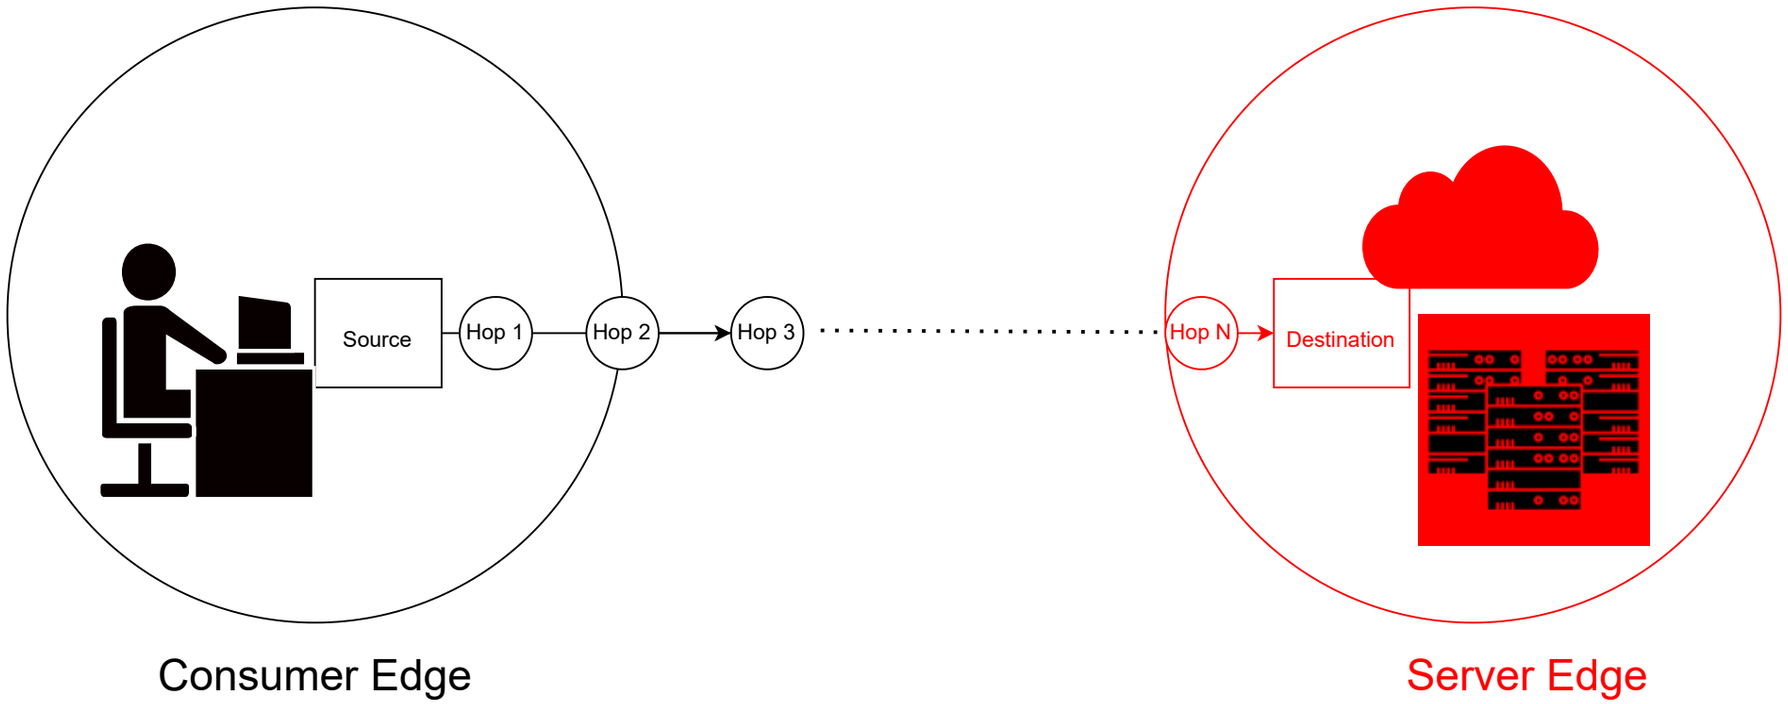

Not a “traversal” test! It is a “functional” test!

Protocol and size: DNS server edge

	DOPT - UDP	DOPT - TCP	HBHOPT - UDP	HBHOPT - TCP
8B	53%	51.16%	16.2%	15.7%
16B	9%	8.7%	2%	2%

N= 19966 DNS Servers autoritative for Alexa Top 1M domain
Test is “successful” if the server replies to a DNS Query

- Very small difference between TCP and UDP
- DOPT results validated from 12 locations
- HBHOPT results validated from 3 locations

Server Edge ASes not passing DOPTs

- AS 63911 NetActuate (880)
- AS 8075 Microsoft Corporation (926)
- AS 397238 Ultradns (4572)
- AS 209453, AS 29169 Gandi (6136)
- AS 16509 Amazon-02 (n=16668)
- AS13335 Cloudflare (n=28098 paths validated over 12 locations)

If these were transparent, E2E test success would be 87%

Test from 12 locations, for 20,000 servers (number of paths tested in brackets).

Which fields are inspected?

Invalid total EH length leads to 99.9% drops

Invalid option lengths lead to 99.9% drops

Invalid EH or Opt Length makes a difference to traversal

Invalid or unknown option type does not make a difference

Payload of option does not make a difference

What did we learn this time?

- DOPTs currently travel very far along a path, seen in both consumer and server edge networks
  - However, some edge paths still drop packets with DOPTs
- HBHOPTs are currently dropped on many types of paths - edge CPE, CDNs (Akamai, Cloudflare and friends), mobile networks and some transit networks
  - A diverse set of paths still support HBHOPTs
  - TCP sees higher drops at the consumer edge

QUESTION TIME

IEPG IETF 115

Experiments Overview

	End2End traversal tests	Traceroute-style tests
Protocol	DNS/{UDP,TCP}/IPv6	{UDP, TCP}/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 domains in Alexa Top 1M list (=~20,000 targets)	5000 RIPE Atlas probes to and from 2 vantage points
Additional tests	• Valid PadN • Invalid length • Unknown option	8, 16, 32, 48, 56, 64B

Both tests done for Hop By Hop and Destination Options

Anatomy of a HBHOPT/DOPT Ext. Header

8B HBHOPT or DOPT EH with a PadN Option

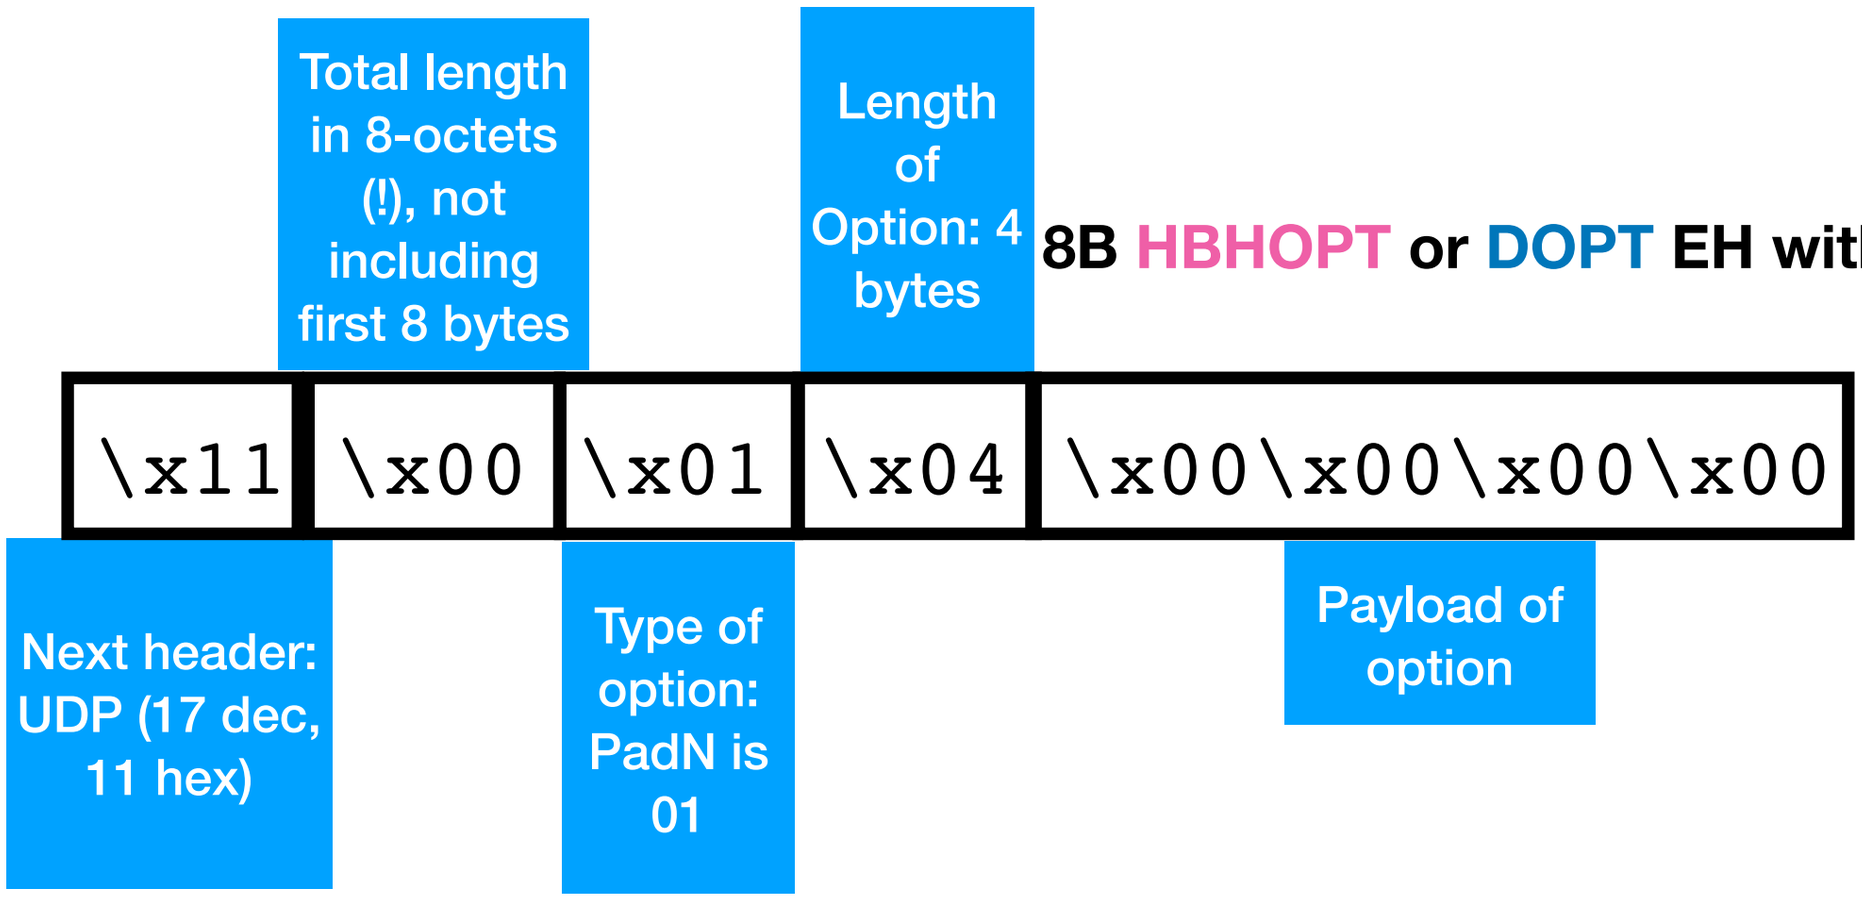

16B HBHOPT with a PadN Option and a PMTU Option

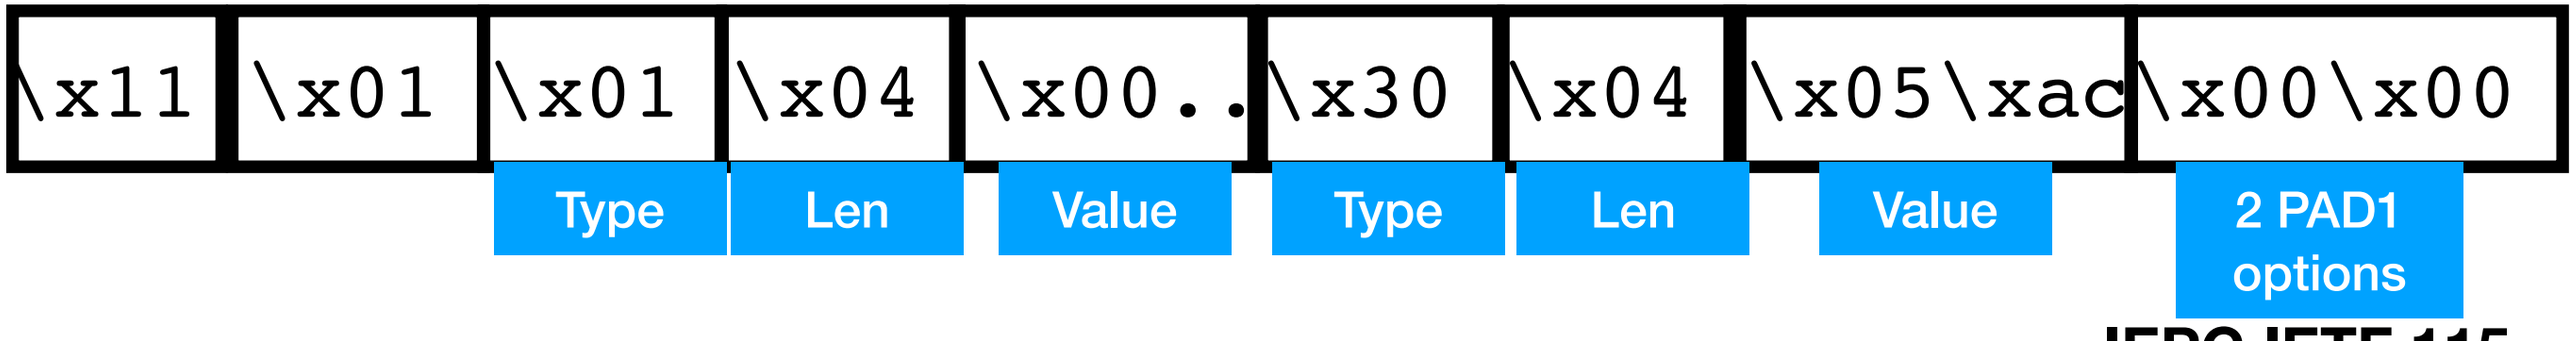

Comparison to RFC7872

- Drops at the server edge AS have increased compared to RFC7872 (result of a few major players), but drops in ASes other than the destination edge have decreased

- Transit networks see better traversal