Designing Experiments to Avoid Internet Measurement Pitfalls

IPv6 Extension Header Edition

Ana Custura*
Gorry Fairhurst*

IRTF MAPRG, Yokohama March 2023

*** University of Aberdeen**

with funding from RIPE NCC

Wide-Scale Measurements

- **Measurements - useful to guide protocol standardisation**
- but… the Internet is heterogeneous
  - 120K registered ASes (~25% in US) = billions of paths
  - Lots of diversity: mobile, CDNs, data center networks
  - Wide-scale measurements needed to target as many (diverse) Internet paths as possible

Measurement Approach

- By technique: *active or passive*, depending on whether measurement traffic is observed or is generated

- By vantage point: *endpoint or in-network*, e.g., where traffic is observed/generated under the control of researcher

- By traffic and aggregation level: *per-packet, per-flow,* etc

- By metric: *performance measurements* (packet loss, throughput), *functional measurements* (transparency to protocols), and more!

Example: IPv6 Extension Headers

IPv6 Extension Headers

- IPv6 was standardised in RFC2474 in the 1990s
  - Designed to be extensible, EHs enable new functionality
  - EHs had a rocky start - some networks drop EH packets
  - Let's look at measuring end-to-end EH traversal…

	Destination Option EH	Hop-by-Hop Option EH
RFC 7872 [1]	80-90%	45-60%
My own data [2]	70-75%	15-20%
APNIC [3]	30-80%	0%
JAMES [4]	94-97%	8-9%

  - …apparently conflicting results?

EH Measurement is *Hard*

- Some devices might not support EH to begin with
- Network can be configured to read beyond the EHs
  - Brokeness can be subtle, for network devices that inspect upper layer protocol information
- Network can be configured to filter EH
  - Edge network devices, transit networks

Active measurements

- Traffic is generated, one or more endpoints controlled by researcher
- Vantage point -> EH traffic -> Destination
- Can measure end-to-end to determine traversal
- ... does it matter ***where*** "problems" or "bottlenecks" occur?

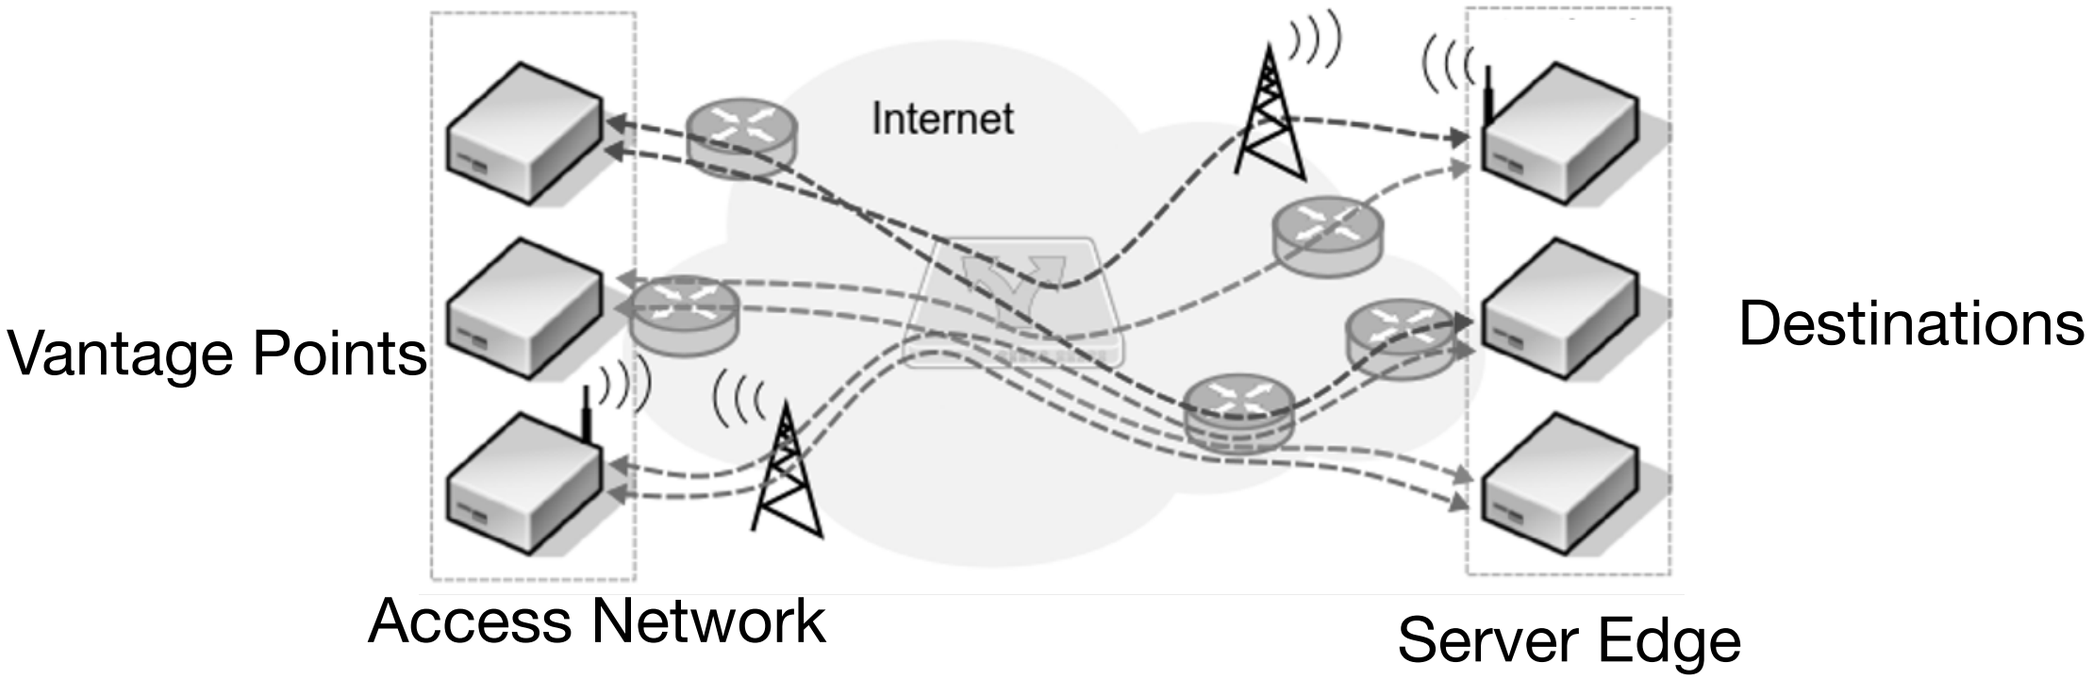

Examples of measurements

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

Where we measure from: Vantage Points

- 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

How we measure: Methodology

- Example 7: different protocols can reveal path info
- Example 8: the same path can reveal unexpected results

Examples of measurements

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

Where we measure from: Vantage Points

- 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: Destinations

- Example 7: different protocols can reveal path info
- Example 8: the same path can reveal unexpected results

***How* we measure: Methodology**

How Not to Ruin your Measurement Campaign:
Choosing the Vantage Points

- Use many vantage points, in multiple AS-es
  - Know your cloud providers
  - Ensure transparency to what you will measure
  - Active measurement platforms: RIPE Atlas, CAIDA Ark etc
- Might need a split between edge/core paths!

- Avoid the Sampling Bias Pitfall

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

End-to-End support percentage for an 8 Byte HBH Options EH - measured in 2022

- Digital Ocean, AWS, Linode - did not support HbH options
  - Still a valid measurement point!
  - But unable to do wide scale measurements from here

Diverse vantage points tell better stories!

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

- Core often more transparent than edge
  - Edge networks can also differ: e.g., mobile, satellite, …

Understanding core/edge helps pinpoint brokenness

How Not to Ruin your Measurement Campaign:
Choosing the Destinations

- Top 1M list of choice:
    - Multiple web, mail, DNS server targets
        - Not diverse, always should include a per-AS split!
        - List needs to be resolved and filtered
- Crowd sourcing: great for clients/edge, harder to reproduce

Results may look different for different types of destinations

Example 3: Destinations

	Per-Host	Per-AS
UK (JANET) - Destination Options	71%	92%
UK (JANET) - Hop-by-Hop Options EH	12%	38%
Canada (OVH) - Destination Options	72%	94%
Canada (OVH) - Hop-by-Hop Options EH	19%	59%

Per-AS vs per-host comparison of the same dataset

End-to-End percentage traversal for an 8 Byte Destination/Hop-by-Hop Option EH, to the authoritative DNS servers for n=20082 destinations in 2867 different ASes.

- One third of destinations = hosted by a few major players

Top 1X lists: considering hosts only can make things look better or worse that they are!

Examples 4-6: Destinations

Dataset	DO8	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%)

- Web vs DNS server data in RFC 7872: per-server split
  - Crowd sourced measurements (APNIC): different story

Infrastructure may look different for different server types

Examples 4-6: Destinations

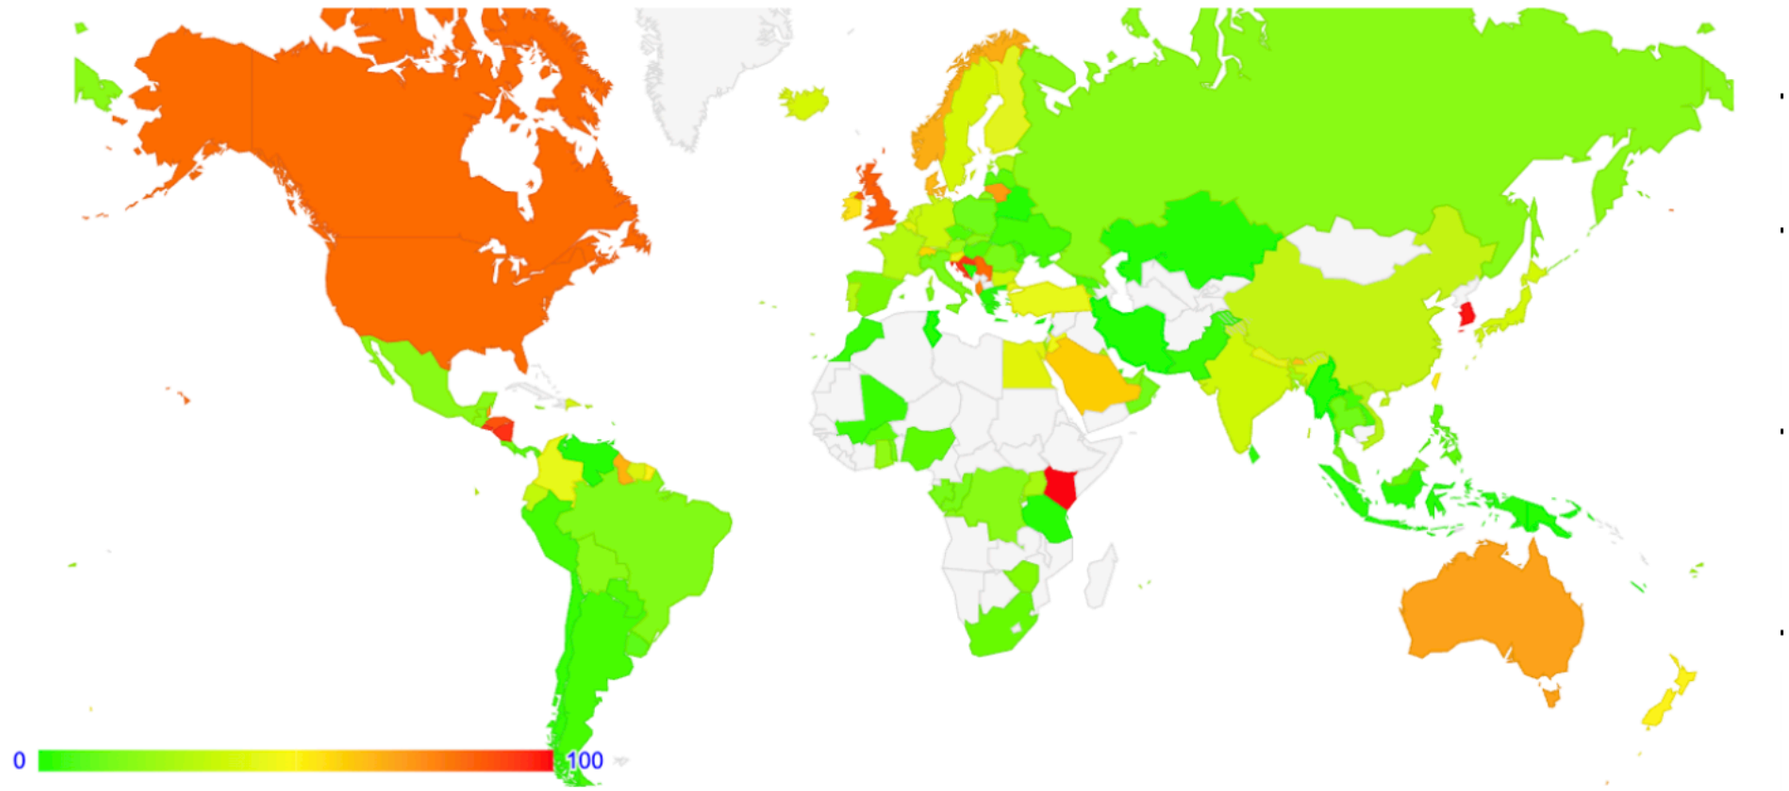

Figure 8 — DST option drop rate, October 2022.

- Web vs DNS server data in RFC 7872: per-server split
  - Crowd sourced measurements (APNIC): different story

Infrastructure may look different for different server types

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

0

Figure 8 — DST option drop rate, October 2022.

- Web vs DNS server data in RFC 7872: per-server split
  - Crowd sourced measurements (APNIC): different story

Infrastructure may look different for different server types

How Not to Ruin your Measurement Campaign:
Choosing the Measurement

- Combine measurement approaches:
  - passive measurements, crowd sourcing
  - gather path info - traceroute-based tools
  - end-to-end-testing; PATHSpider, Scamper
- Measure longitudinally, open source your data
- Compare your methodology and results

How Not to Ruin your Measurement Campaign:
Choosing the Traffic/Protocol

- Combine measurement approaches

  .... and measure multiple upper layer protocols

- Because of:
  - load balancing in the network ;)
  - load balancing at the server edge ;)
  - firewalls and other configured policies
  - .... and more

Example 8: Protocol Differences

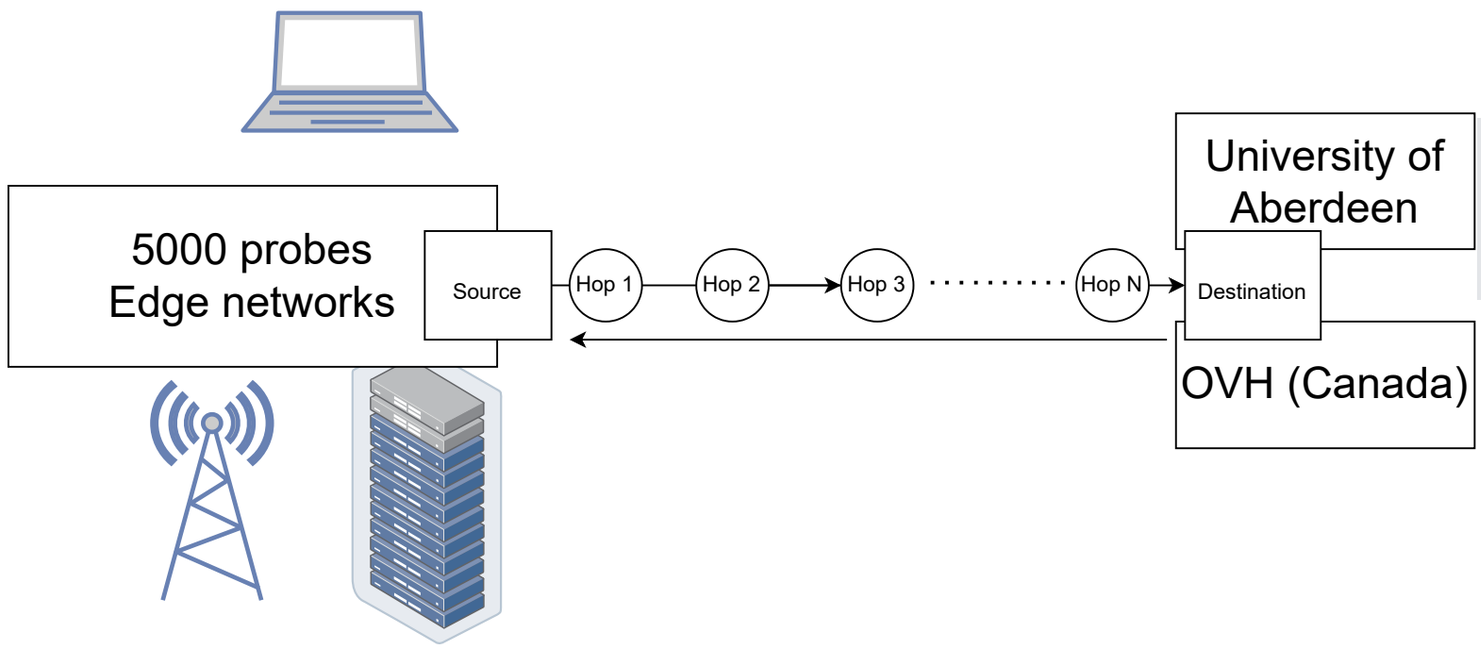

	DOPT	HBHOPT	
University of Aberdeen	~92%	~11%	UDP
OVH (Canada)	~68%	~9%	TCP

- 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?

Traversal of IP features can depend on the transport protocol

Example 9: Load Balancing

Path measured with packets without EH

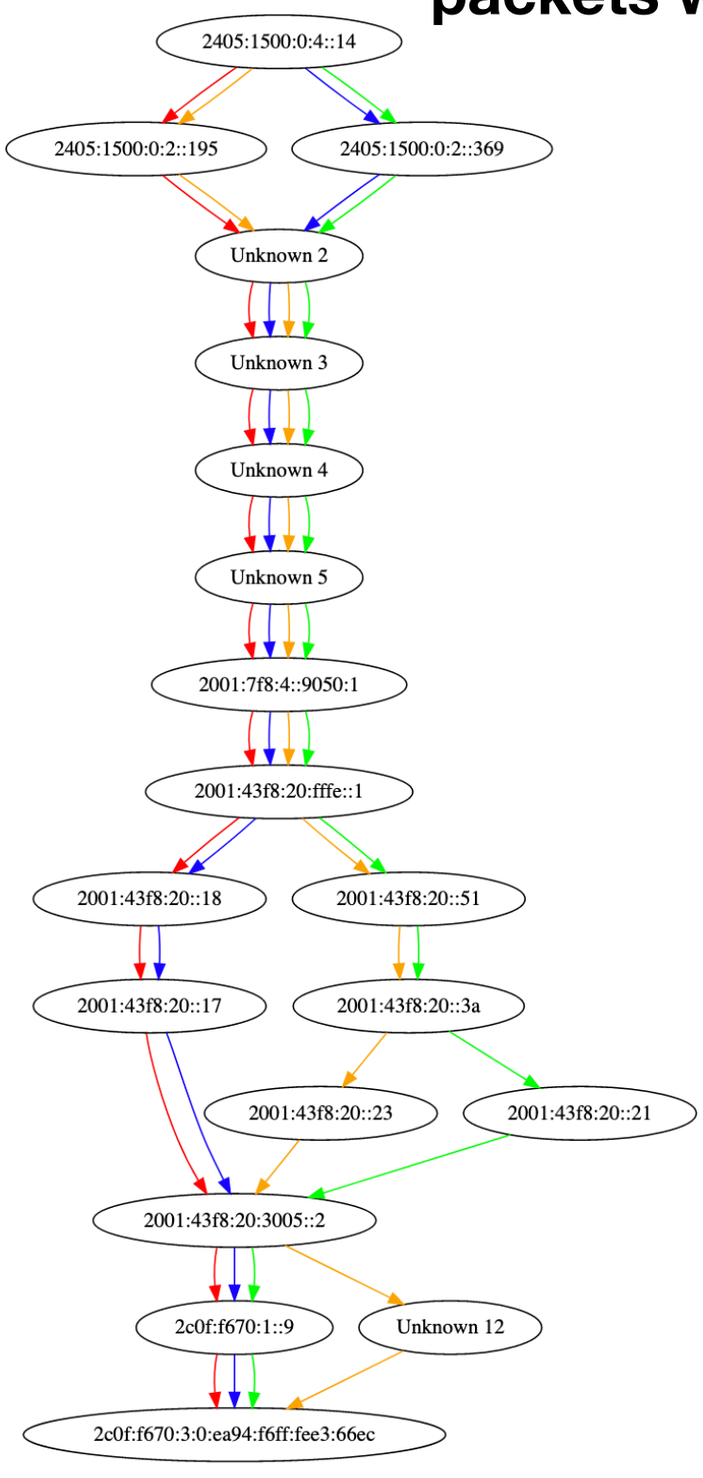

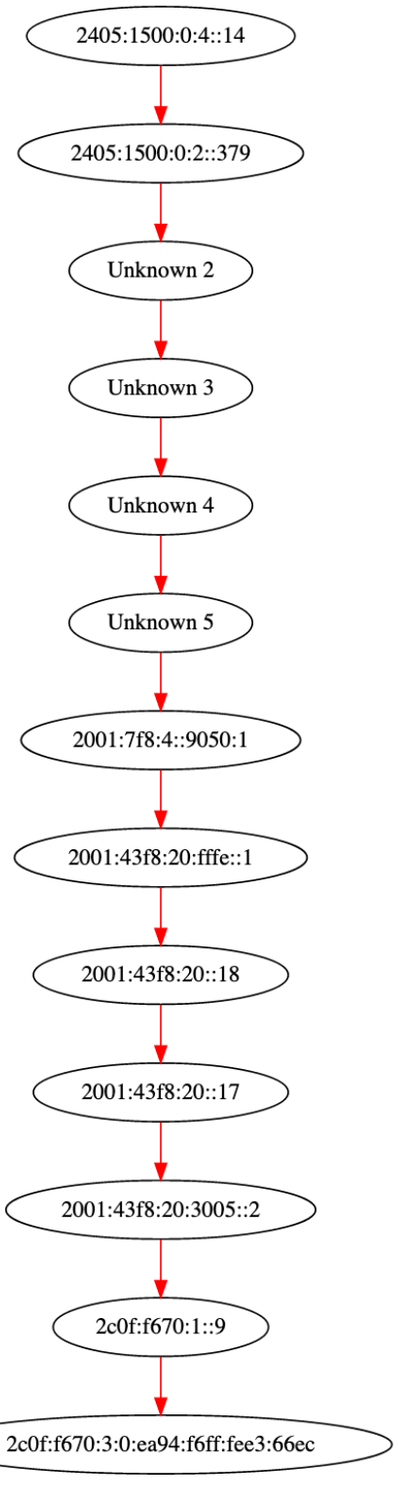

Same path measured with a Destination Options EH

Example 9: Load Balancing

Path measured with packets without EH

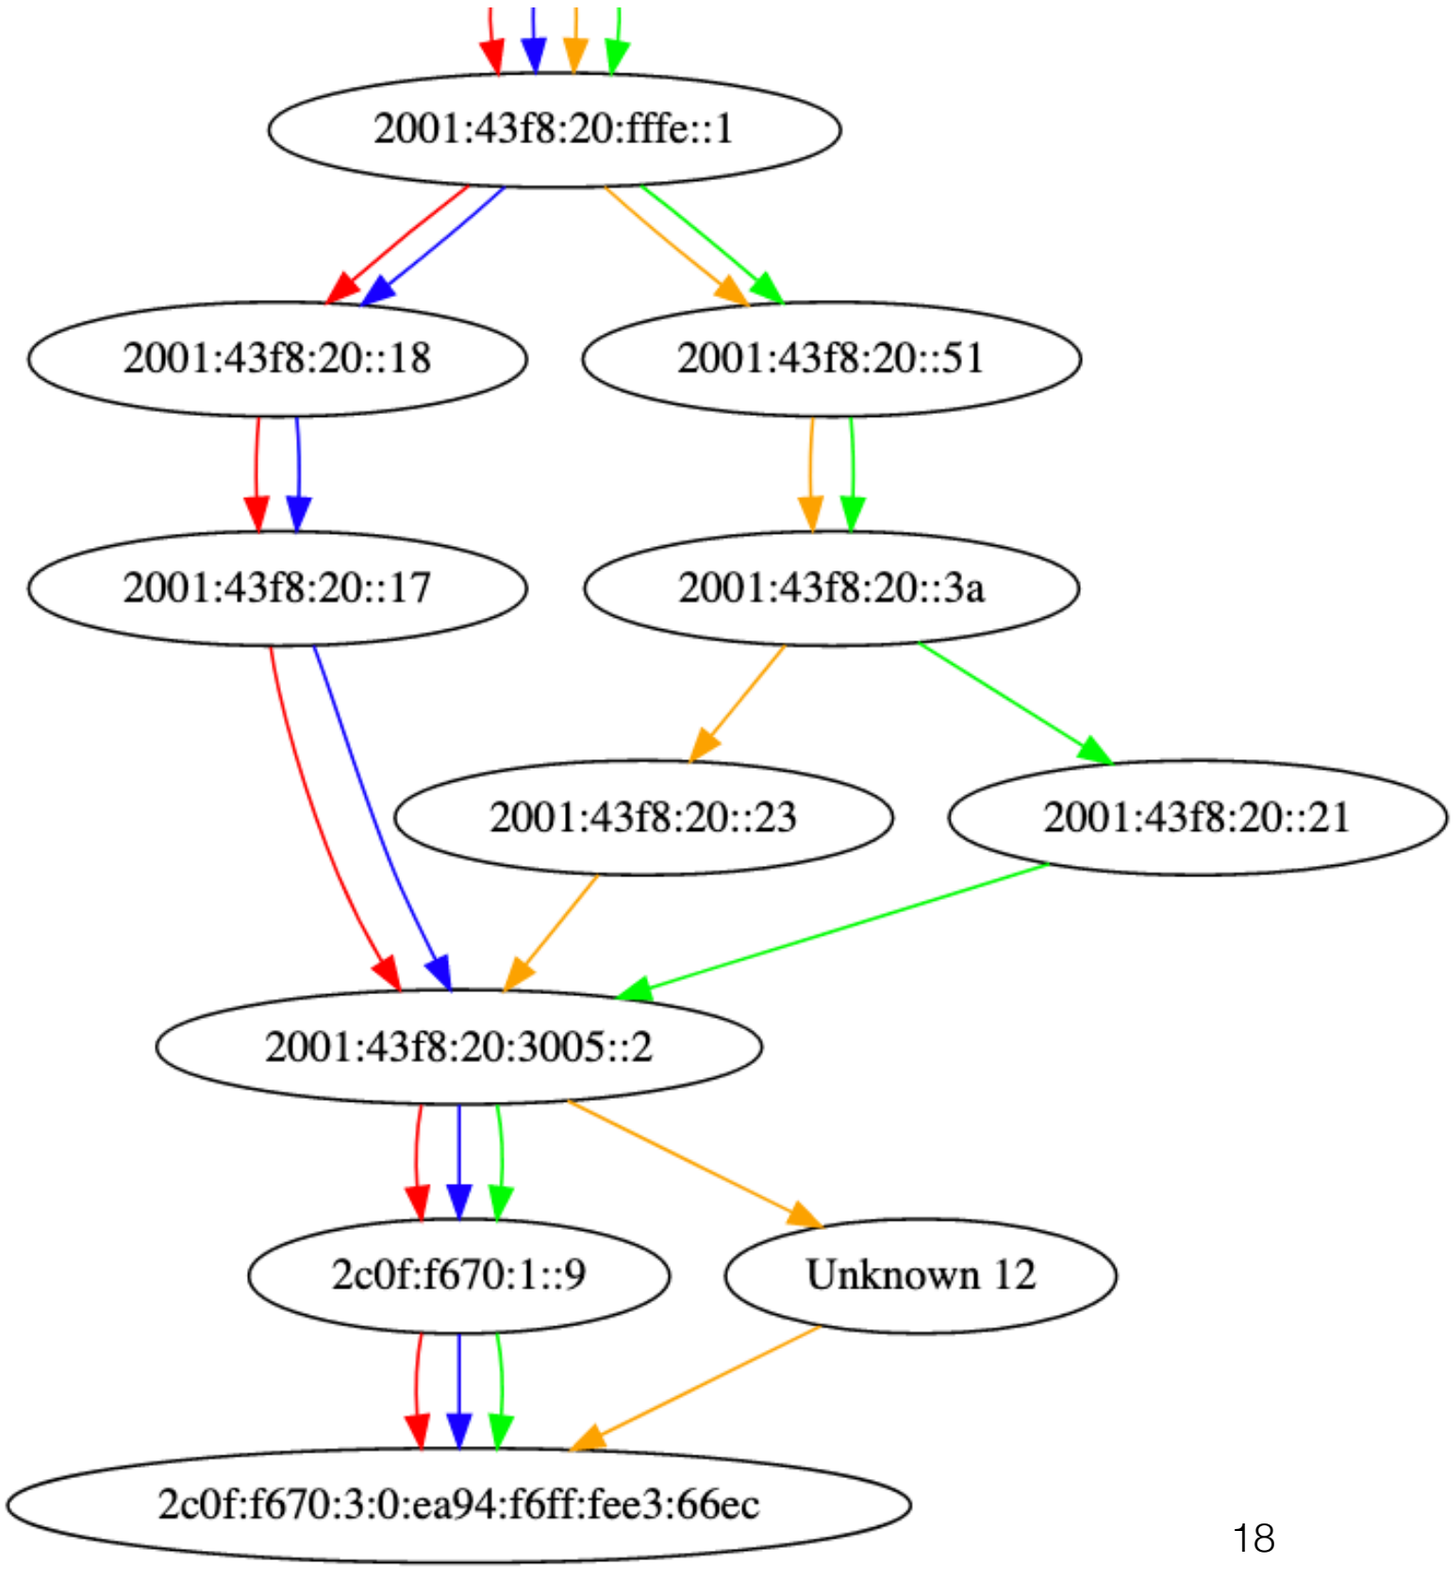

Same path measured with a Destination Options EH

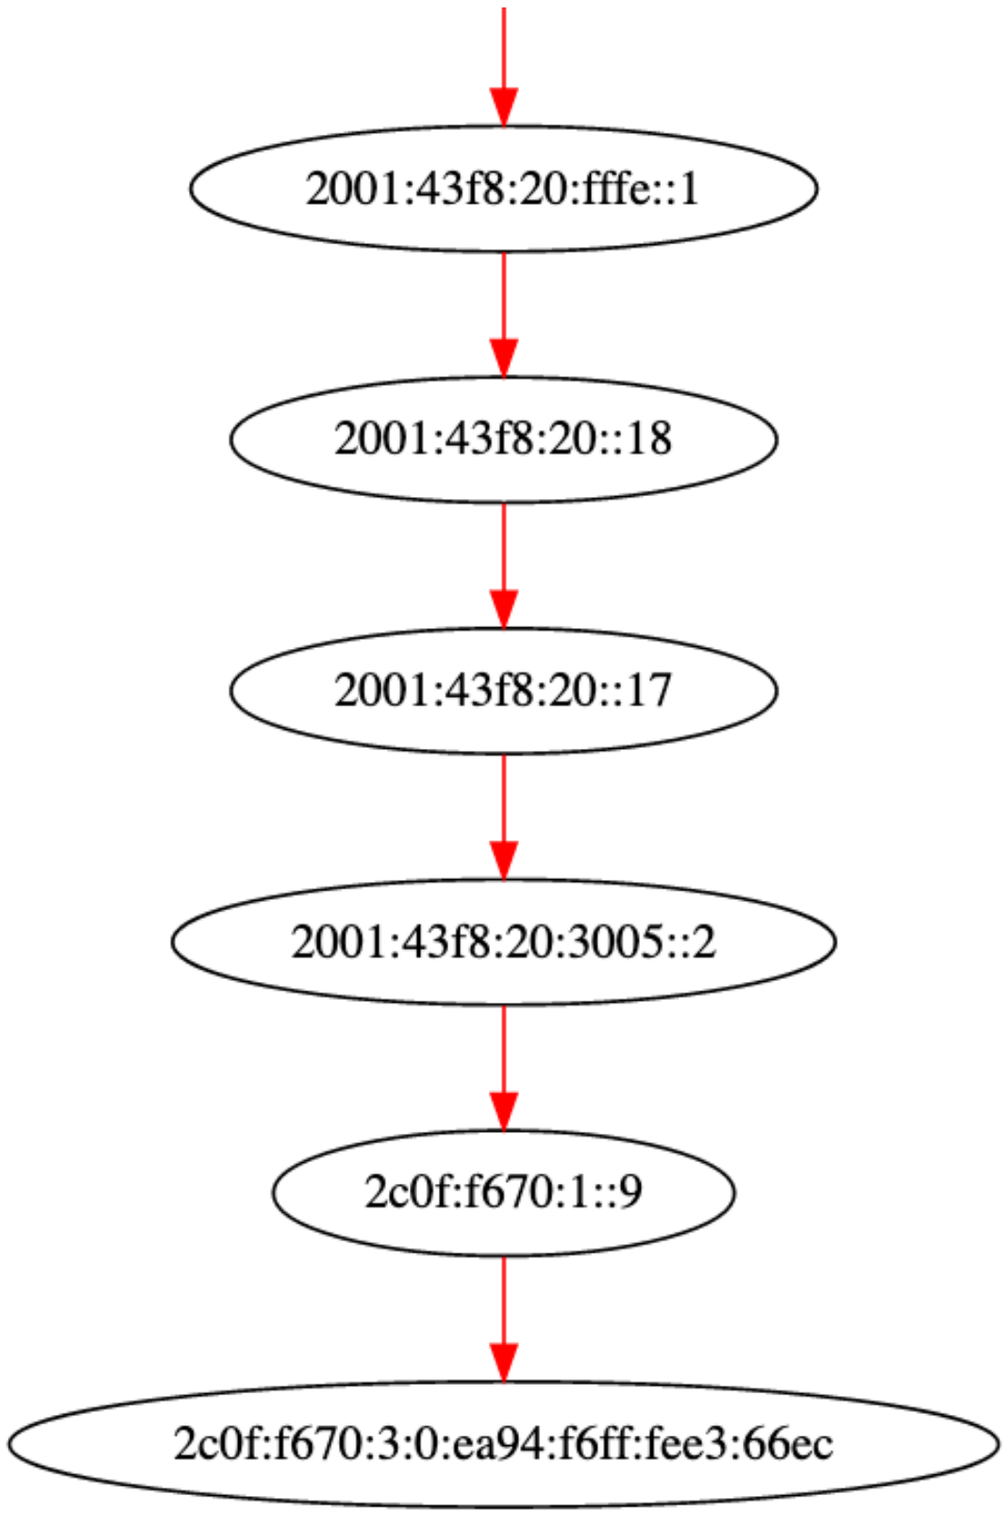

How Not to Ruin your Measurement Campaign

Recommendations:

- Try both *active or passive* technique(s)
- Use many vantage points and destinations
- Consider measurement aggregation level and metric
  - Cross-check data when possible
  - Open-source your data
- Expect the unexpected!

End-to-end + path measurements + diverse categories of targets, destinations and protocols mitigates limitations of each way to measure!

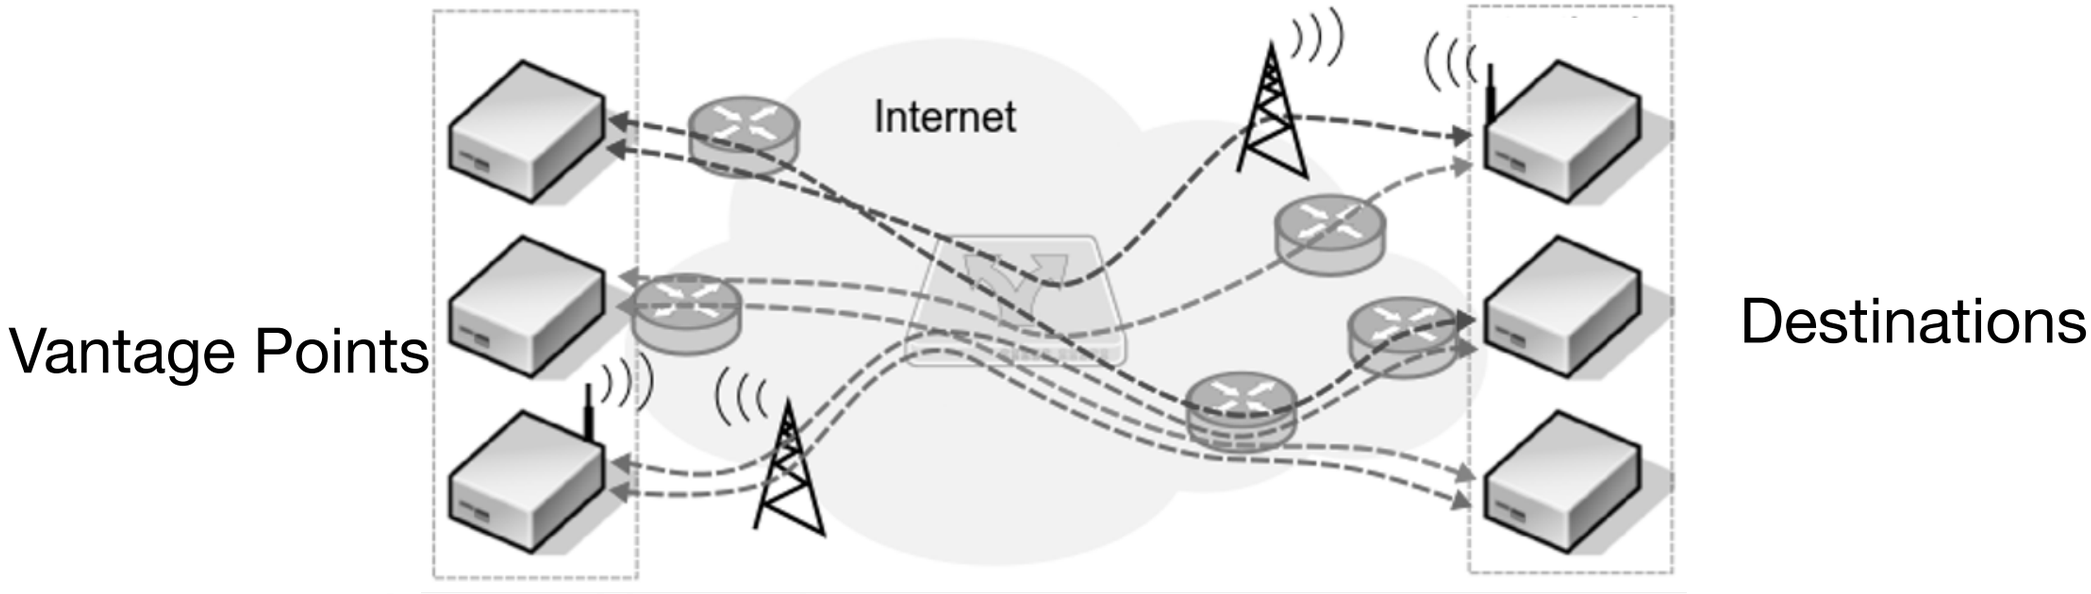

Questions?

- [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/

IPv6 Extension Headers

- IPv6 was standardised in RFC2474 in the 1990s
  - Designed to be extensible, EHs enable this new functionality
  - Defined arbitrary number of EHs following base IPv6 header
  - First routers did not support IPv6 EH processing in hardware
    - Packets processed in software, vulnerable to DoS attacks
    - Many networks drop packets with EH.
  - Bugs in less-used IPv6 code also remain

IPv6 Measurement is Difficult

- Many edge networks still do not support IPv6
  - e.g., mobile networks, broadband in Europe, ...
- IPv6 servers - hosting companies, e.g., Cloudflare, do not always proxy IPv6 request to an IPv6 origin server
  - The IPv6 top domains lists are not very diverse
  - Hard to scan, but there are IPv6 hitlists
- Measurements should take load-balancing into account

Existing EH measurements

	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