IPv6 Extension Headers in the Real World v2.0

Jen Linkova <furry@google.com>
Fernando Gont <fgont@si6networks.com>

IEPG 90
Toronto, Canada. July 19, 2014

Motivation

- You can love ext. headers or hate them but can we really use them?
- Measurements have been done (mostly for fragment header) - but what is the trend?
- Where and why packets are dropped?

Past IEPG presentations

- Fernando Gont @ IEPG 88:
  - > 50% drop rate for small EHs (e.g. DOH of 8 bytes)
  - > 40% drop rate for Fragmented traffic
  - > 90% drop rate for large EHs (e.g. DOH of 1K)
- Tim Chown & Fernando @ IETF 89:
  - > 60% of packet drops >= 7 hops from destination

Some unanswered questions

- Where are packets with EHs dropped?
  - same AS as destination?
  - different AS?

Jen's Measurements

Topology

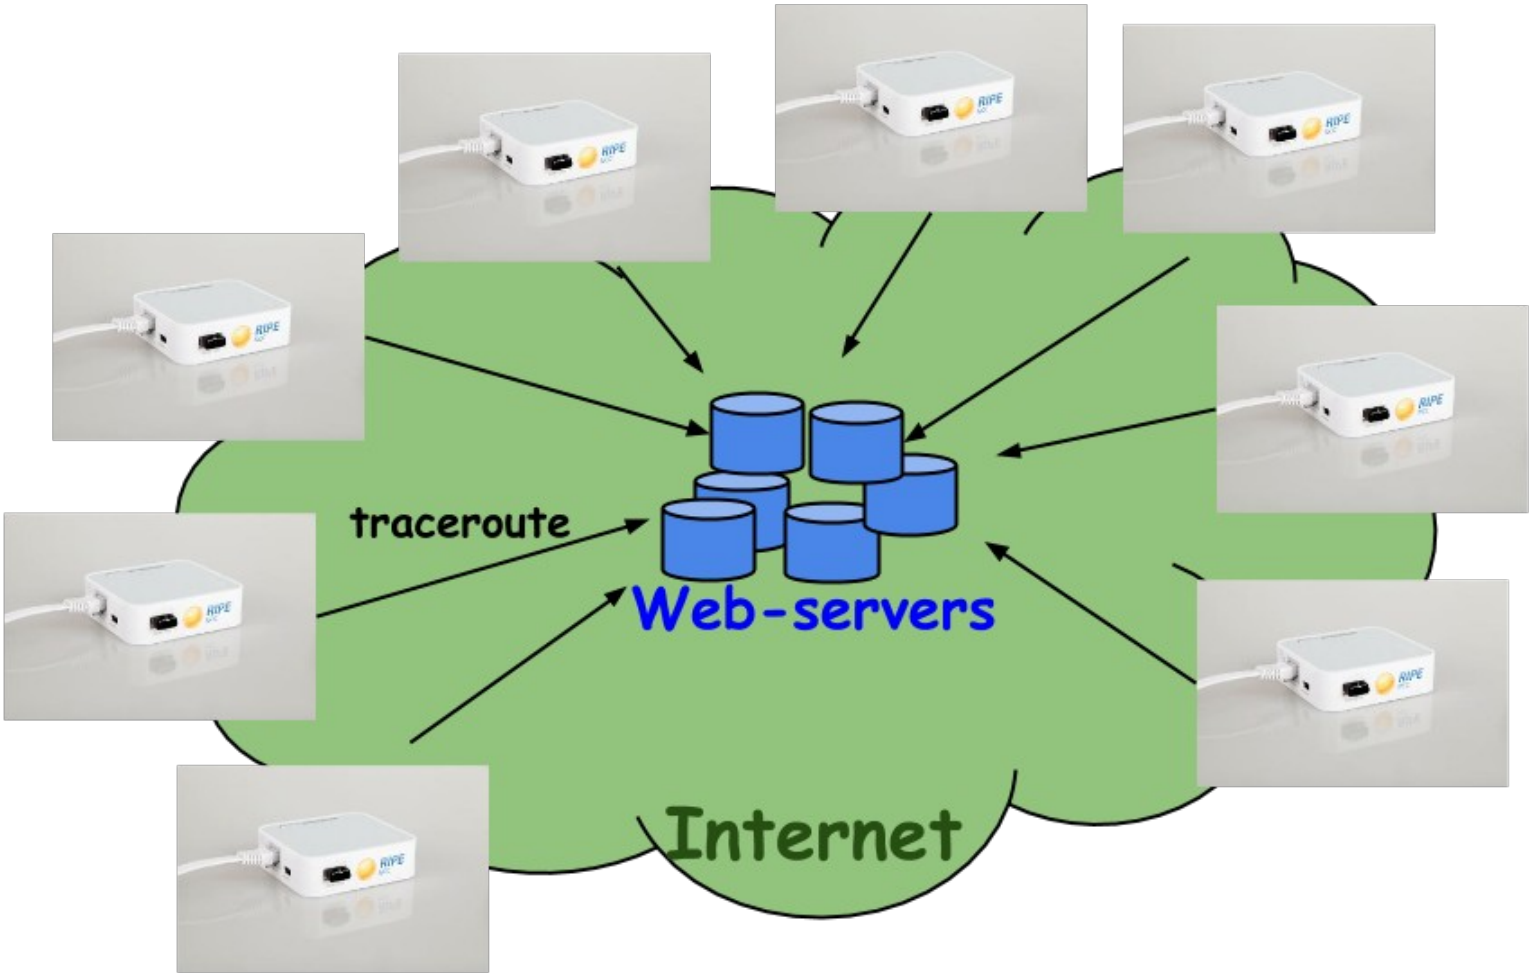

Choosing targets

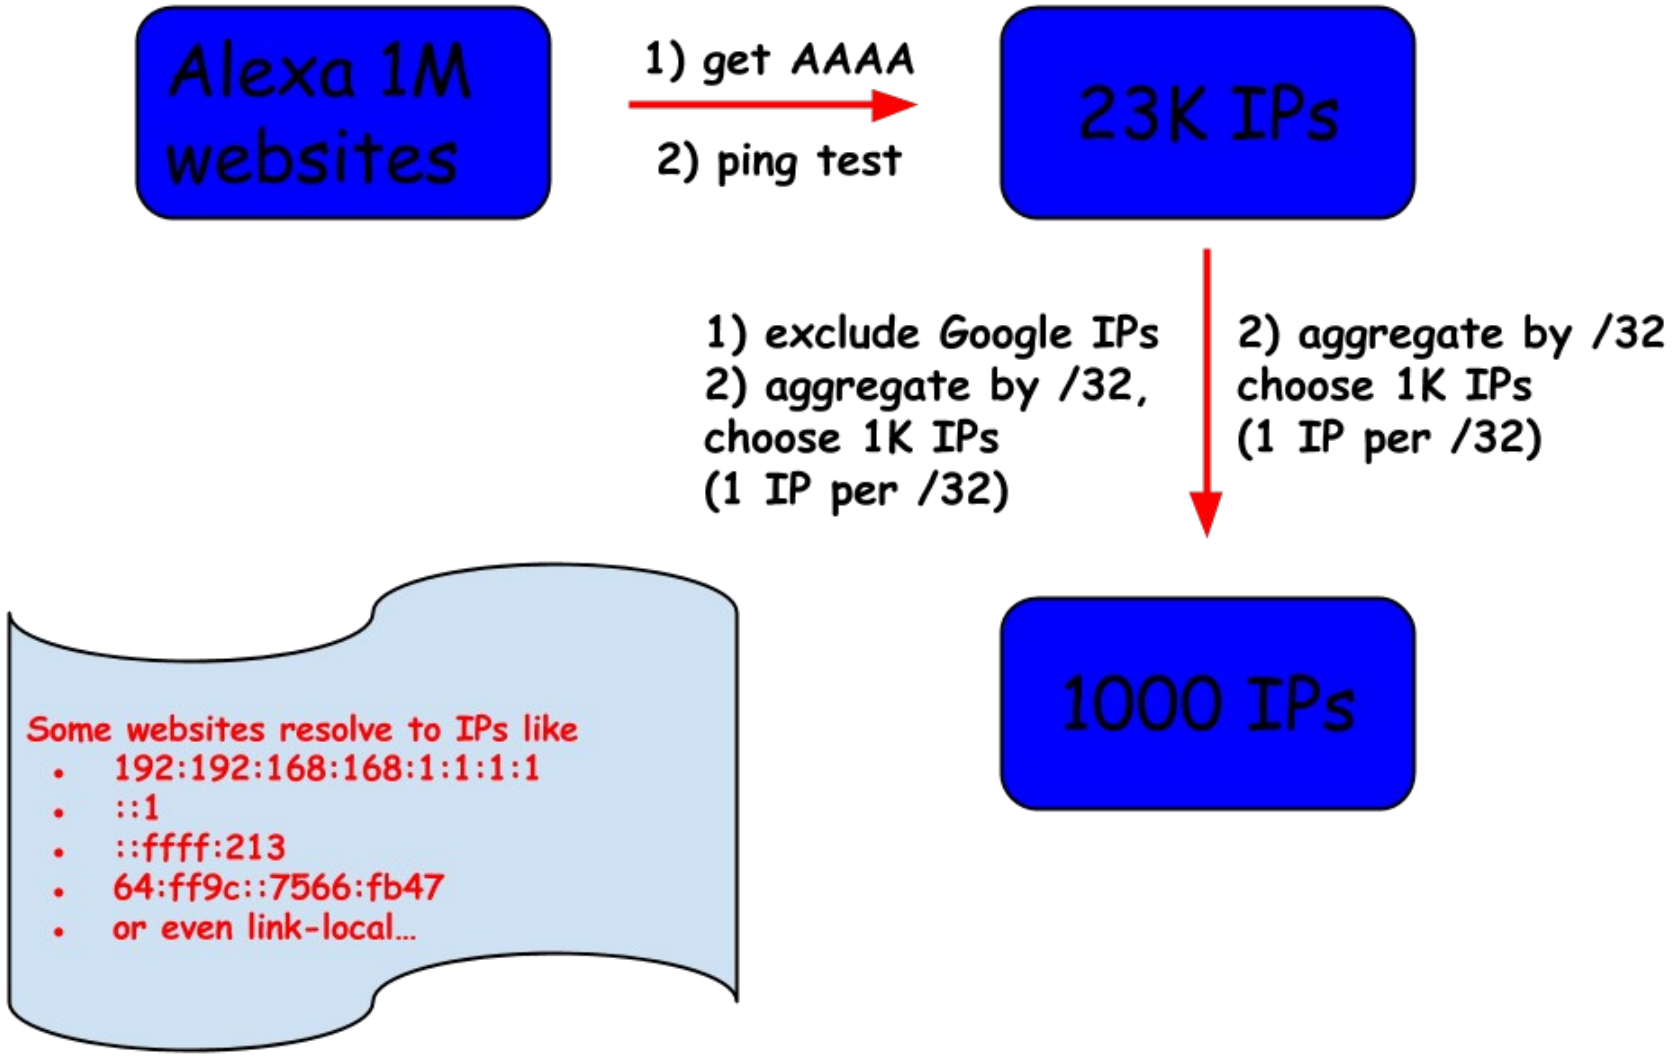

Choosing RIPE Atlas probes

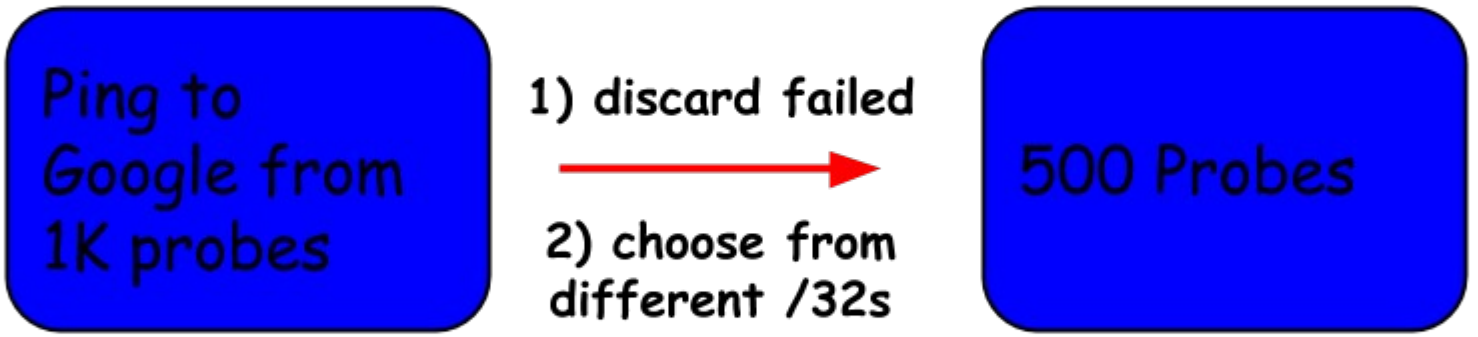

Methodology

- To each destination from each probe the full test consists of:
  - control measurement (ICMP traceroute)
  - 9 ICMP traceroute tests:
    - Hop-by-Hop Options:
      8 bytes, 512 bytes, 1024 bytes
    - Destination Options
      8 bytes, 512 bytes, 1024 bytes
    - Hop-by-Hop + Destination Options
      4 bytes + 4 bytes
      128 bytes + 128 bytes
      512 bytes + 512 bytes

Processing the results

- For each (probe; destination) test:
  - discard the test if the control test failed
  - discard the test if not all 10 sub-tests were run on the probe (you don't necessarily get all probes you requested)

- Final data set:

  238179 tests in total (10 traceroutes) to 746 destinations

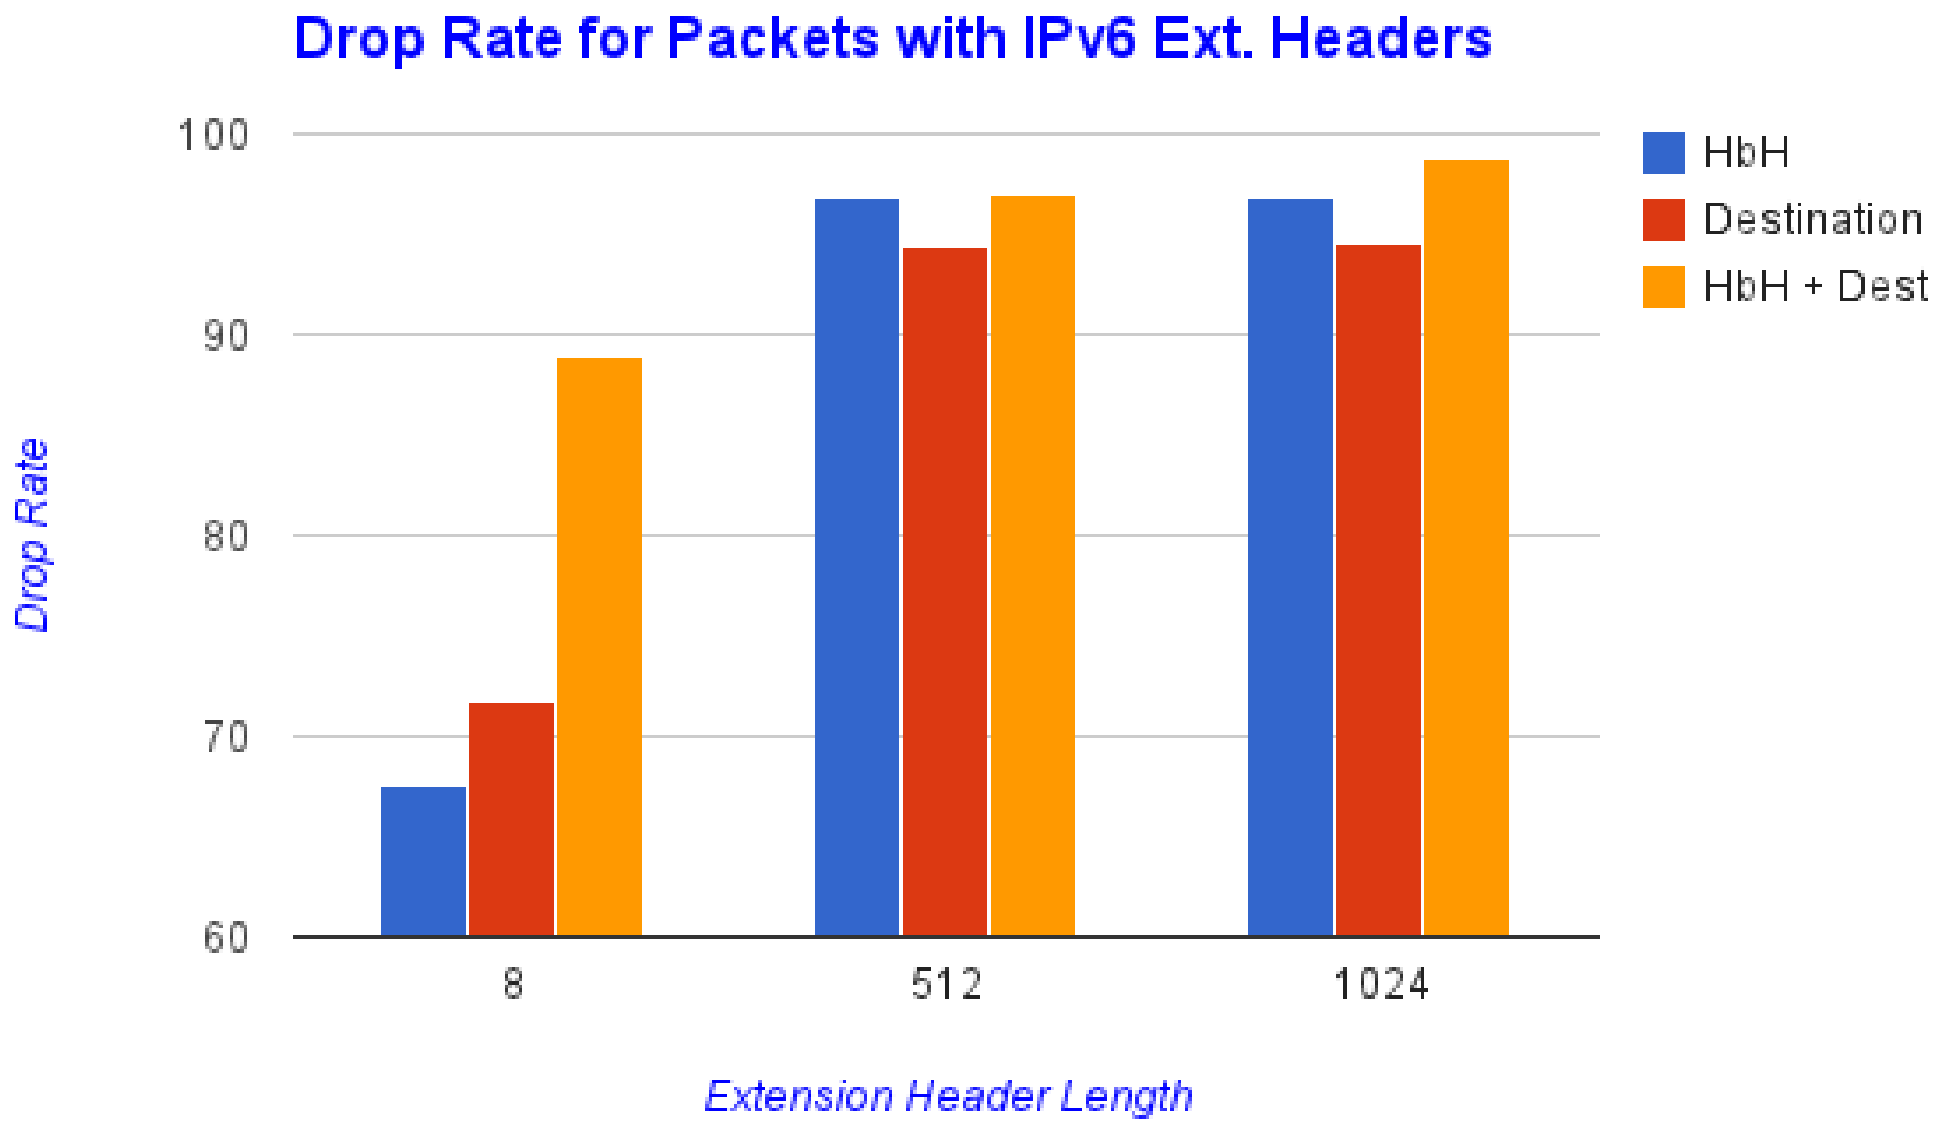
Drop Rate for Packets with IPv6 Ext. Headers
100
90
80
70
60
Drop Rate
8
512
1024
Extension Header Length
HbH
Destination
HbH + Dest

Where packets are dropped?

- Finding origin AS for each traceroute hops
- Ignoring invalid IPs/link-local/ULAs/etc
- Comparing ‘AS_PATH’ for control test and the measurement;
  - If AS_PATH for failed test has length 0 or 1:

    packet could not leave the origin network
  - If last AS in AS_PATH for failed test is destination AS or PHP AS from the control test:

    packet was dropped in the destination network or on its edge

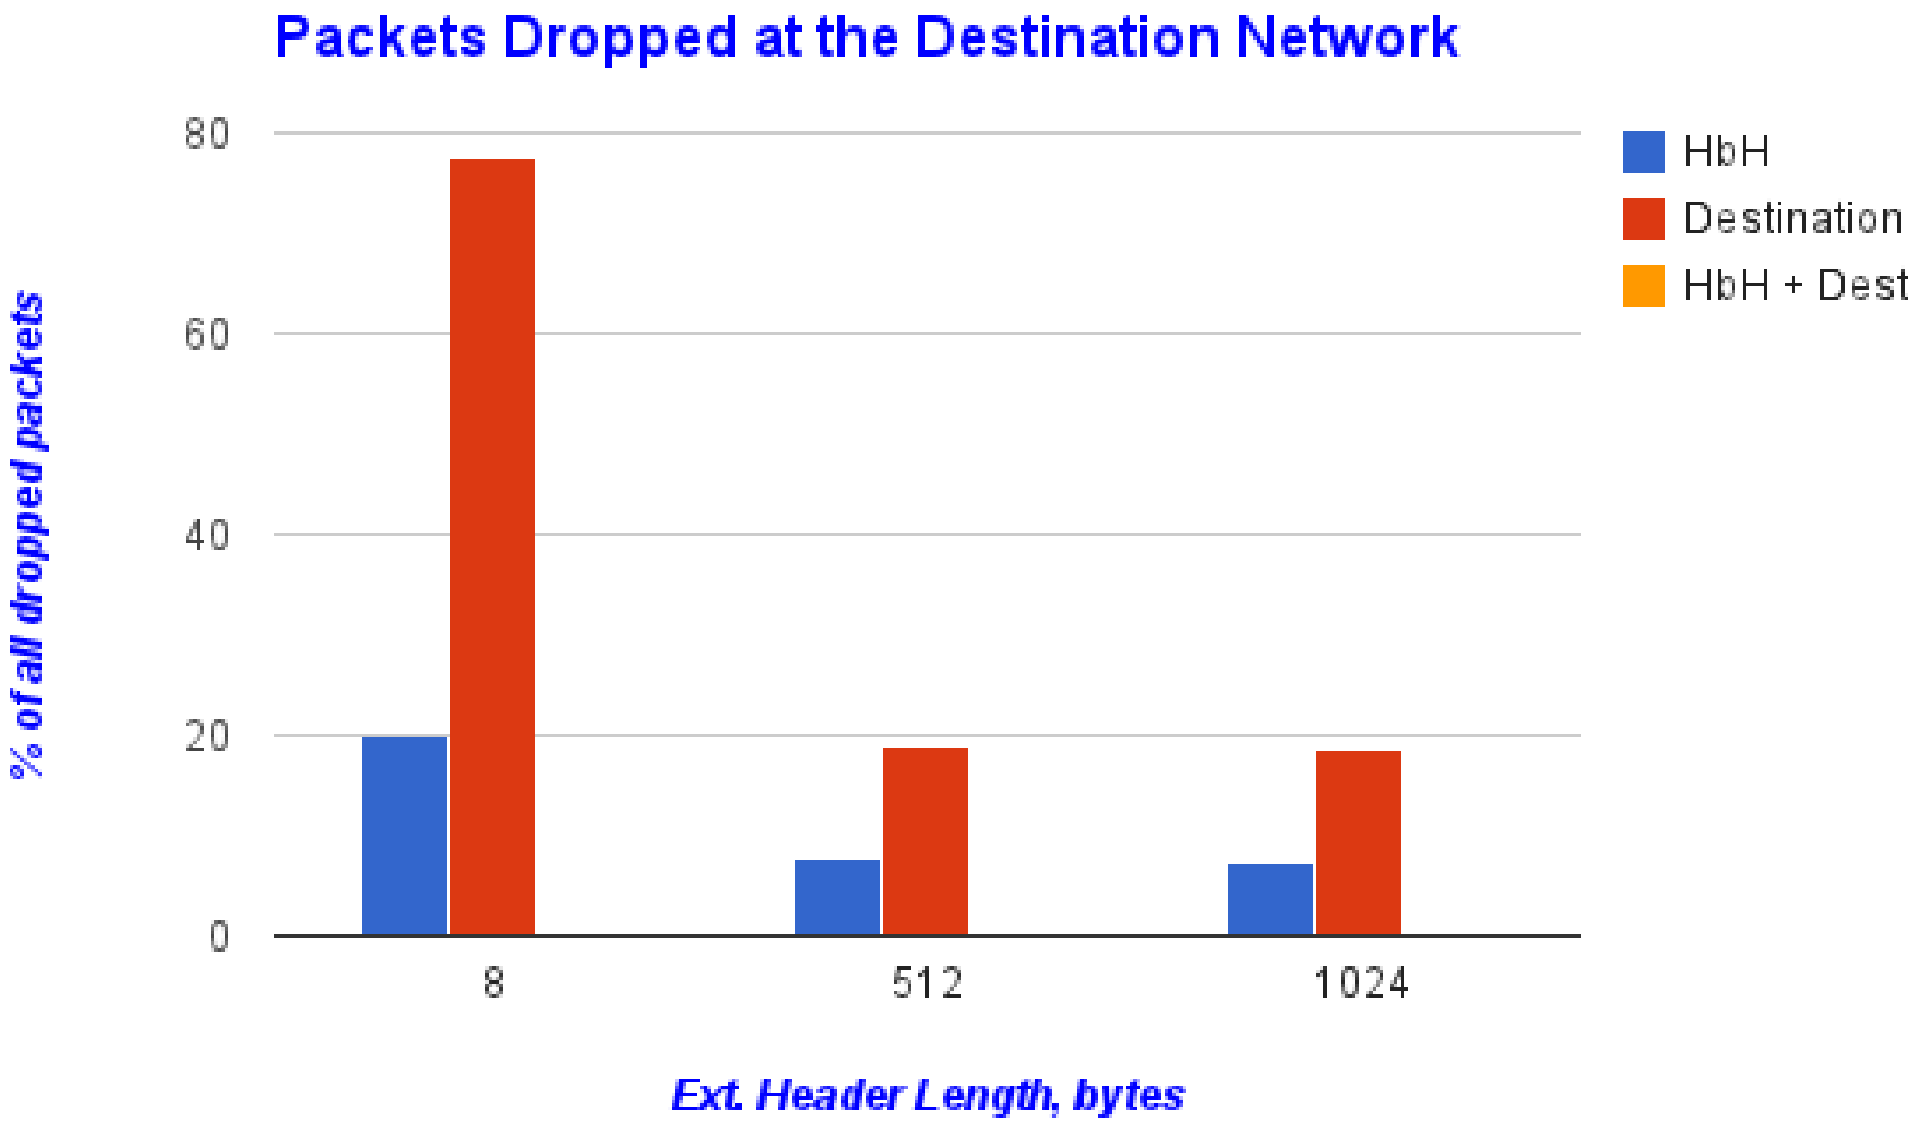
Packets Dropped at the Destination Network
80
60
40
20
0
% of all dropped packets
8
512
1024
Ext. Header Length, bytes
HbH
Destination
HbH + Dest

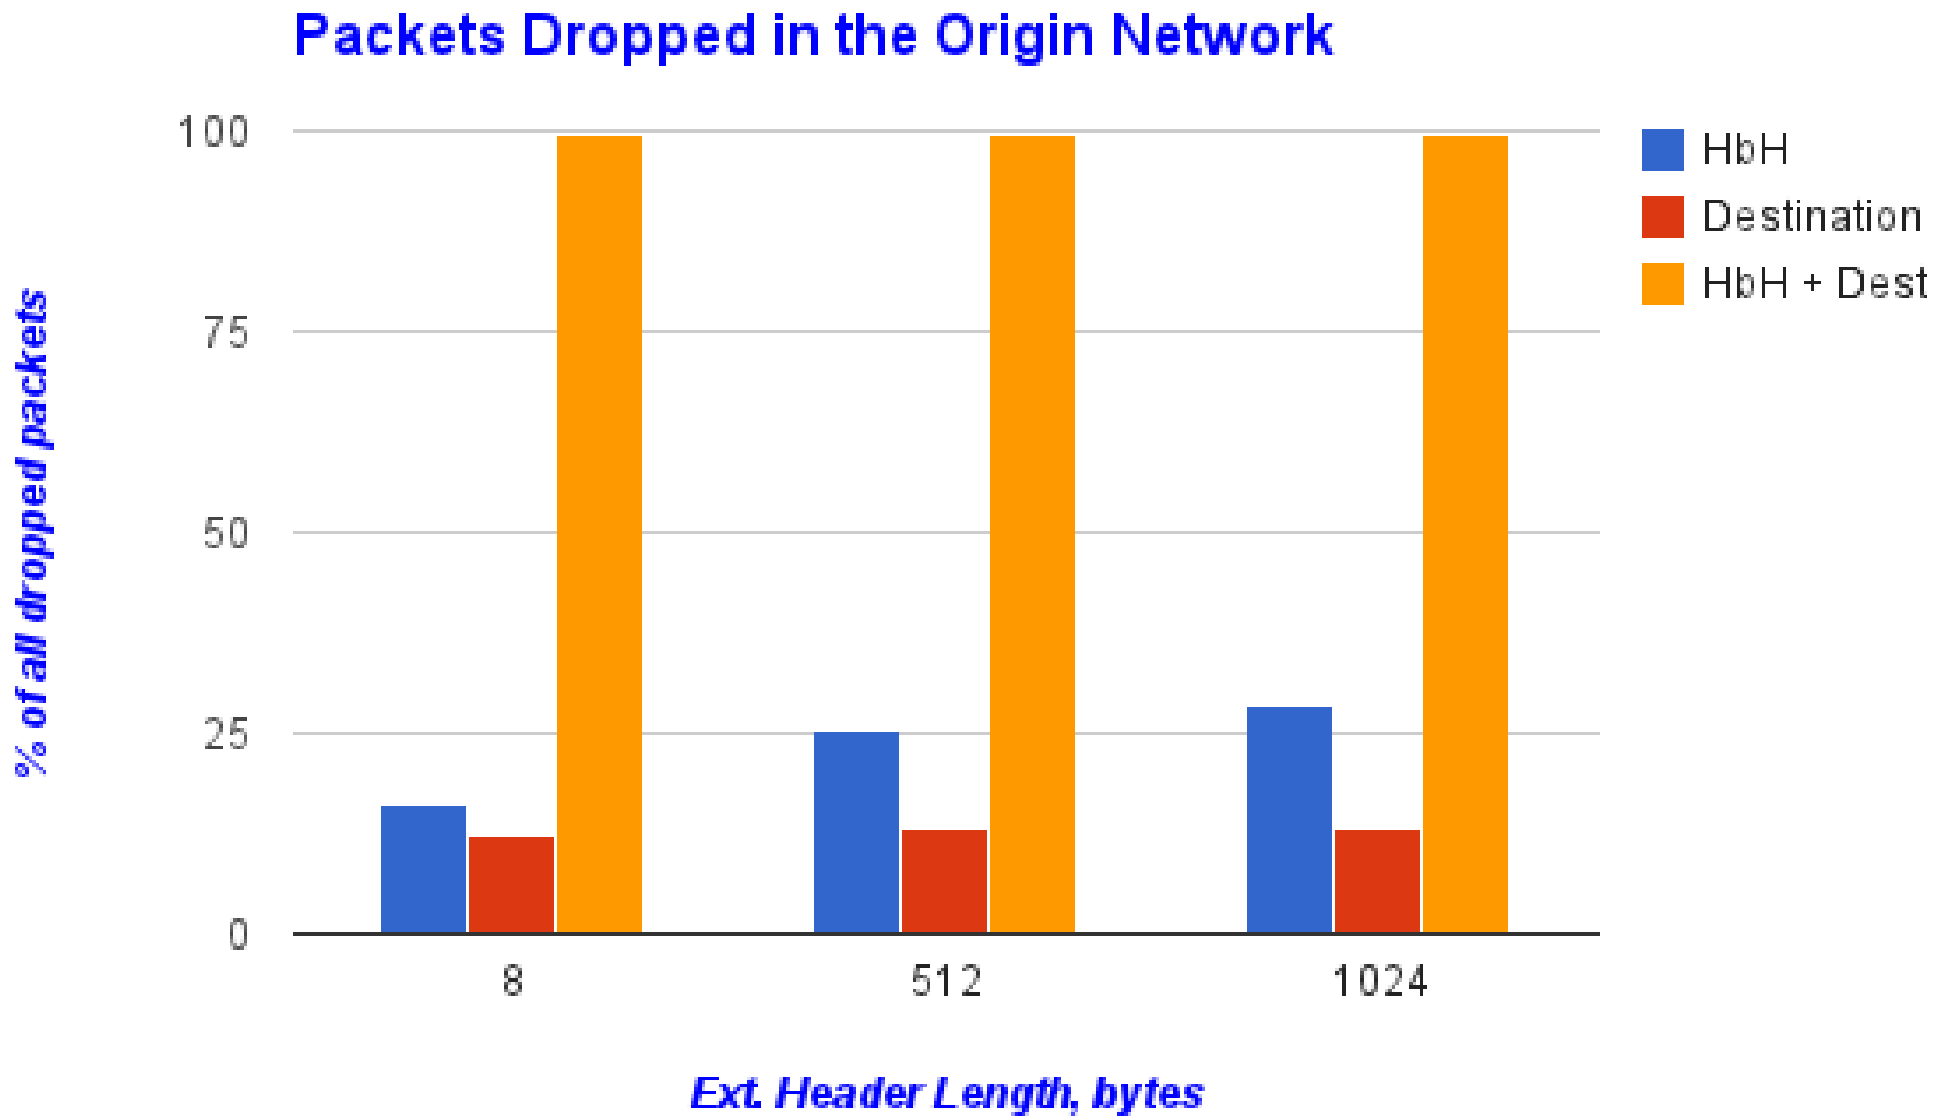
Packets Dropped in the Origin Network
100
75
50
25
0
% of all dropped packets
8
512
1024
Ext. Header Length, bytes
HbH
Destination
HbH + Dest

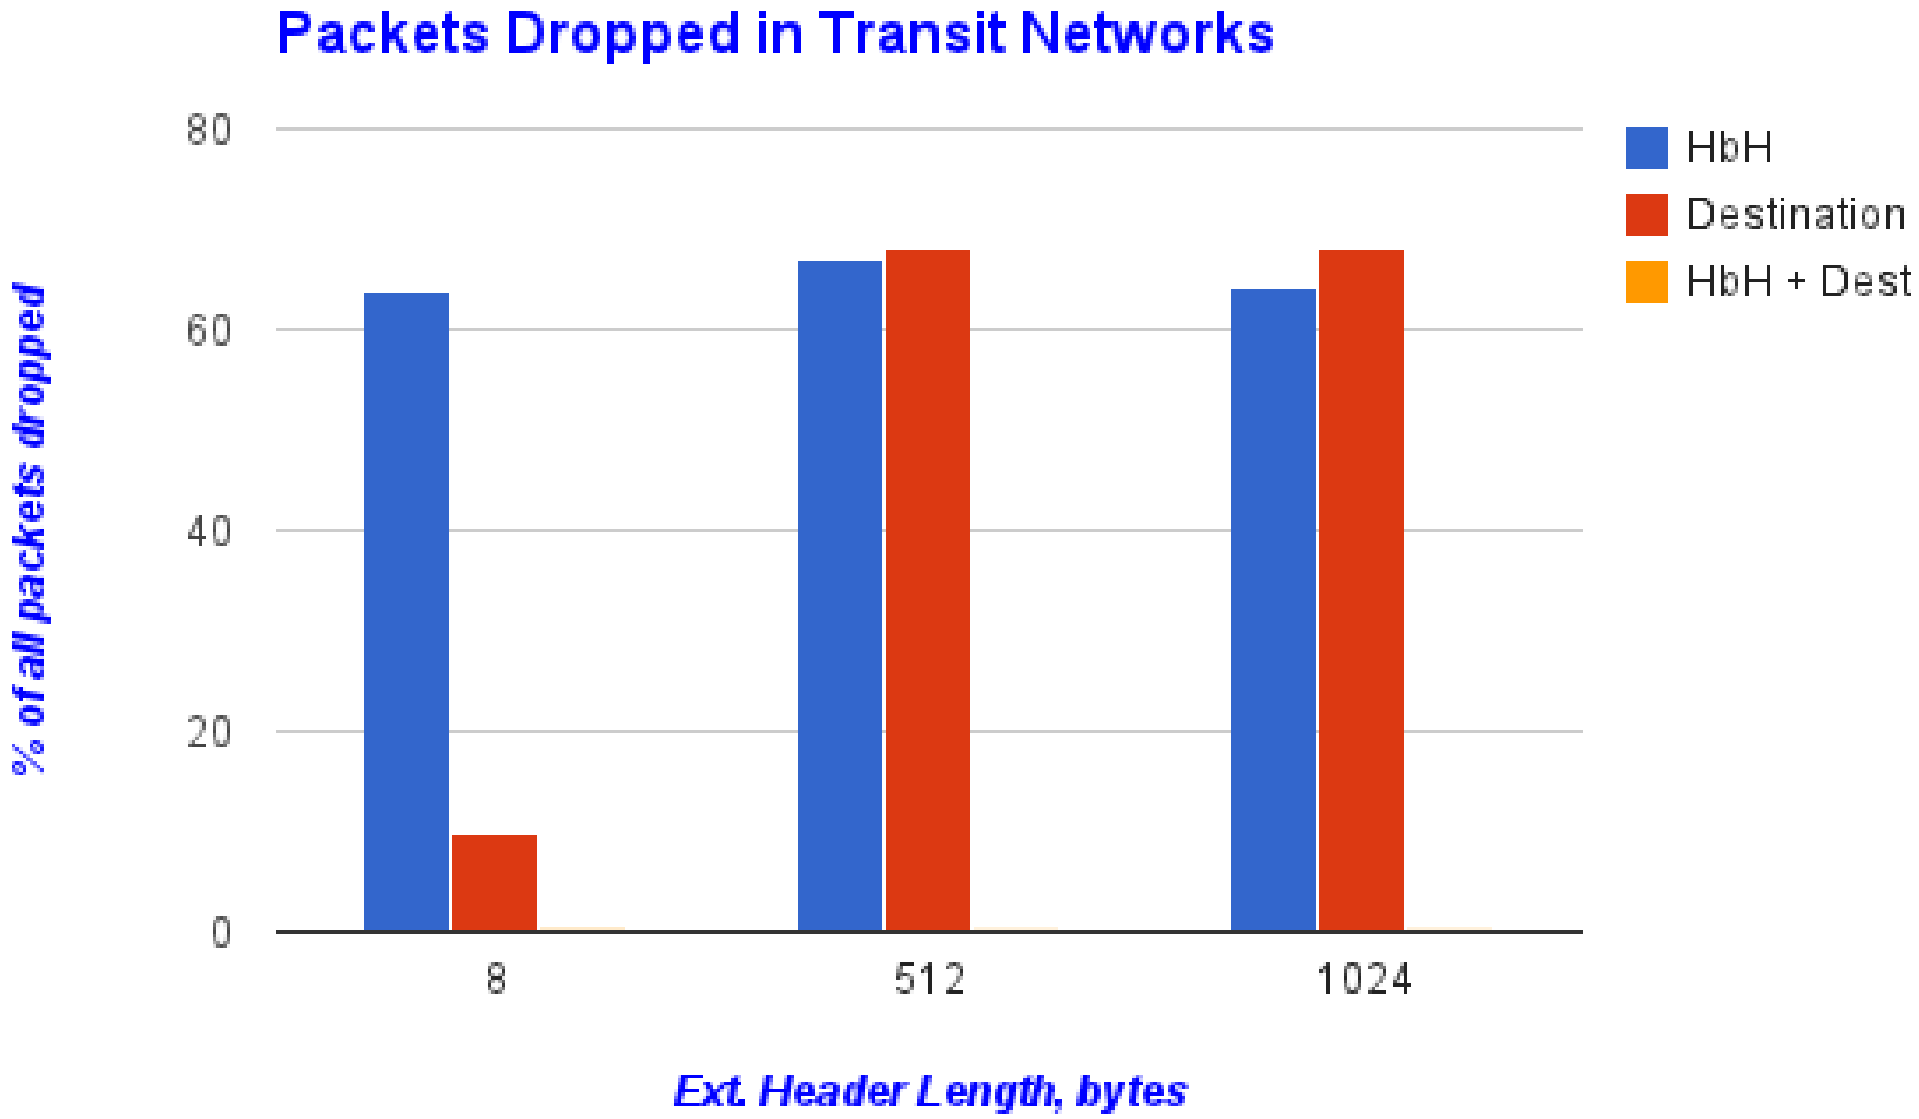
Packets Dropped in Transit Networks
80
60
40
20
0
% of all packets dropped
8
512
1024
Ext. Header Length, bytes
HbH
Destination
HbH + Dest

Conclusions

- Packets with EHs ARE DROPPED ;(
- Short EHs have lower drop rate
  - most chips could not look deeper than first 64-128-256 bytes?
  - 78 % of packets with short Destination Options EH could reach the destination network even
- Adding second EH increases the drop rate to 99% packets could not leave the origin network
  - routers/firewall can not look deeper than the first EH?

Fernando's Measurements

Methodology/Caveats

Finding the dropping node (I)

- Given the output of “traceroute6” for no-EH and EHs:

	No EHs	With EHs
	1. fc00:1:1:1000::1	1. fc00:1:1:1000::1
	2. fc00:1:1:2000::4	2. fc00:1:1:2000::4
	3. fc00:1:2:4000::1	3. fc00:1:2:4000::1
	4. fc00:2:1:4000::1	4. fc00:2:1:4000::1
	5. fc00:a:2:1000::1	5. fc00:a:2:1000::1
	6. fc00:a:4:4000::1	6. fc00:a:4:4000::1
DROP	7. fc00:b:1:1000::1	
	8. fc00:b:2:5000::1	
	9. fc00:b:4:5000::1	
	10. fc00:d::1	

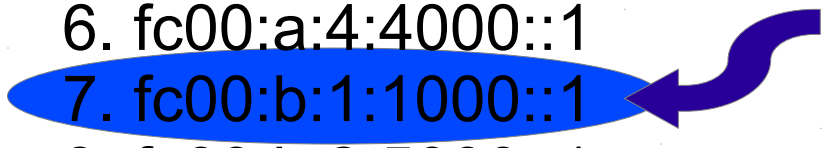

Finding the dropping node (II)

- We assume ingress filtering...
- Otherwise dropping node actually is M rather than M+1

Finding the dropping AS

- Lookup ASN of dropping node, but...
- There may be ambiguity when finding the AS of the dropping node:
  - who provides the address space for the peering?

Peering address space (case 1)

- Case 1: Address space provided by AS Y

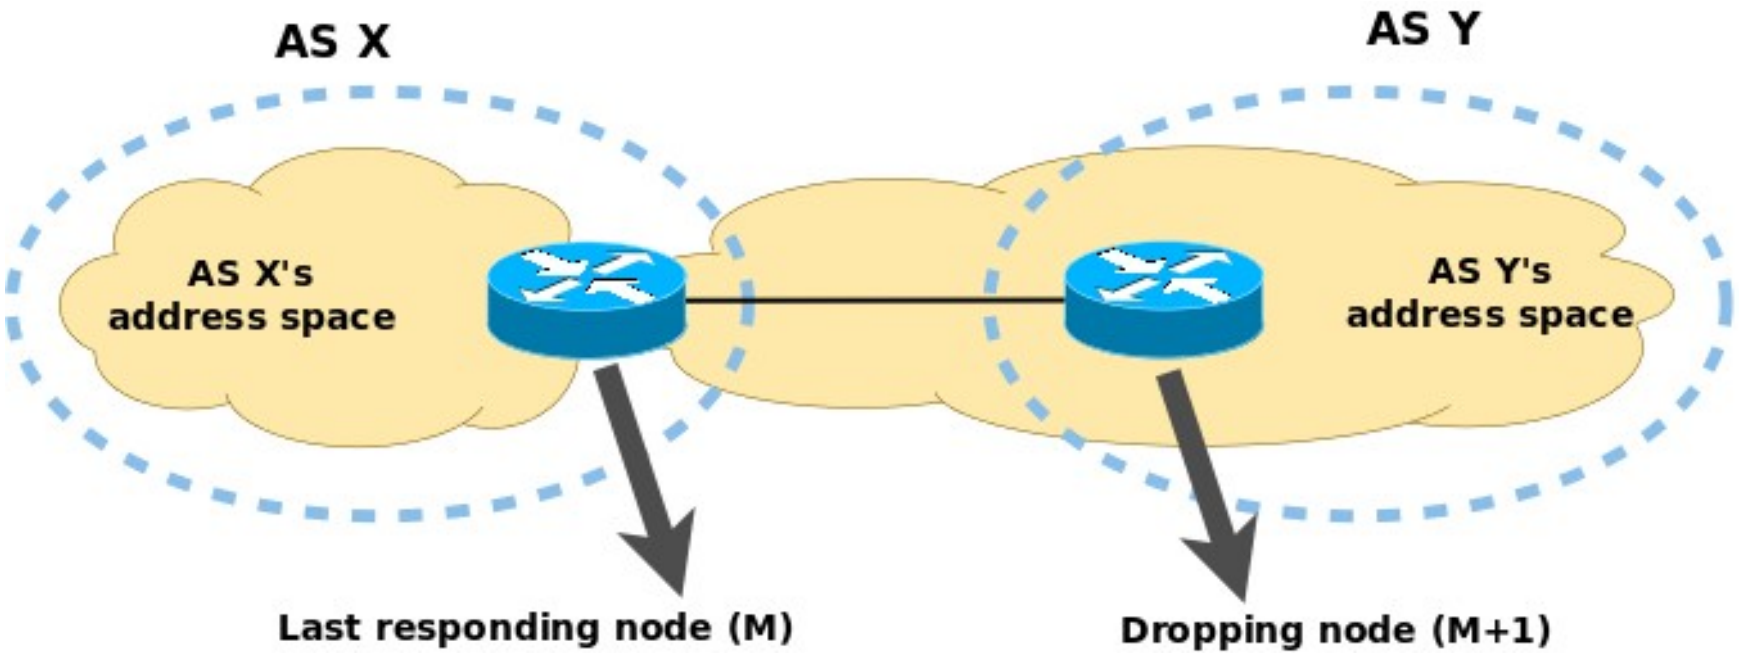

Peering address space (case 2)

- Case 1: Address space provided by AS X

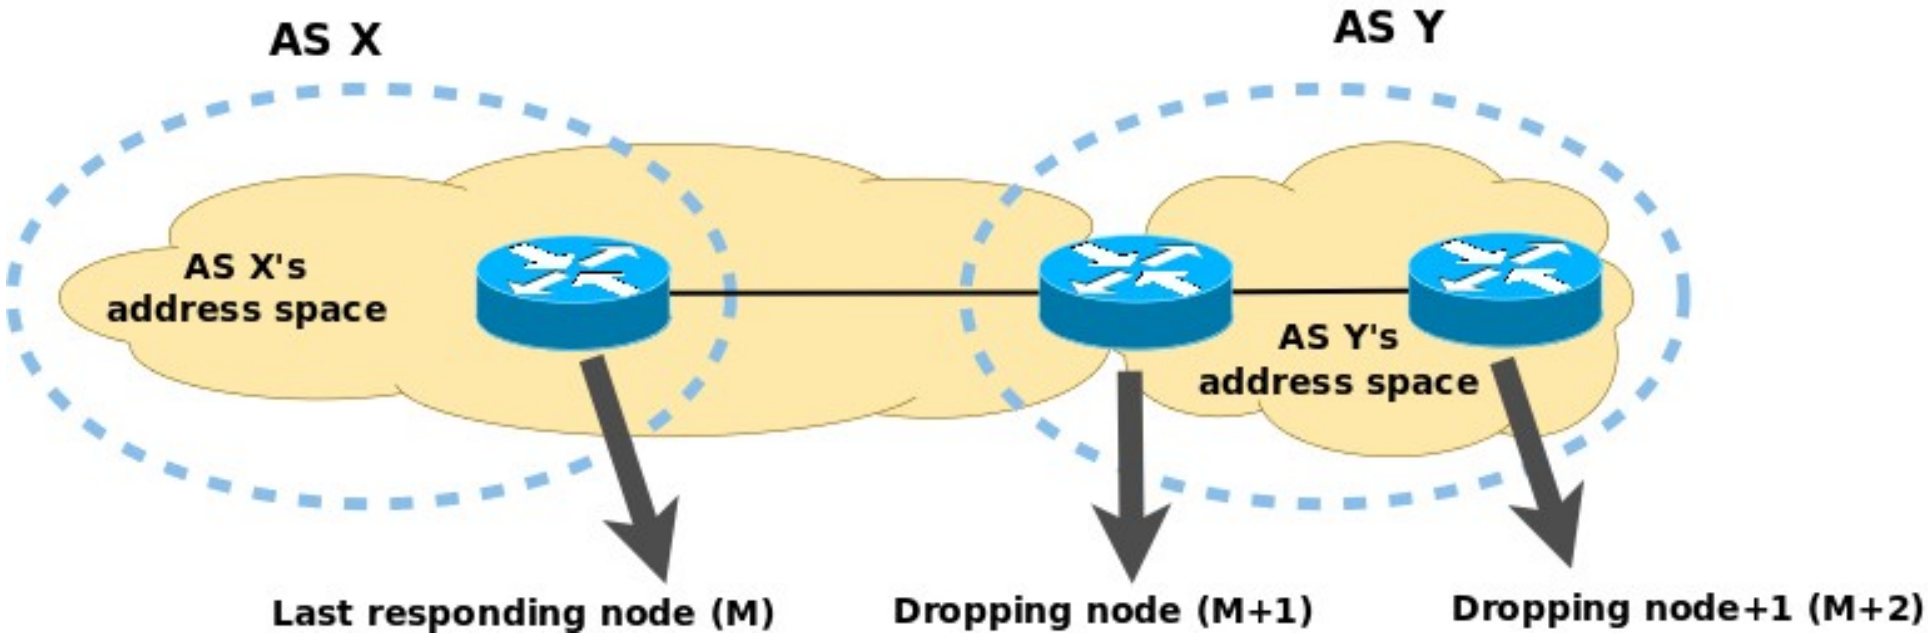

Measurement Results

WIPv6LD dataset: Packet Drop rate

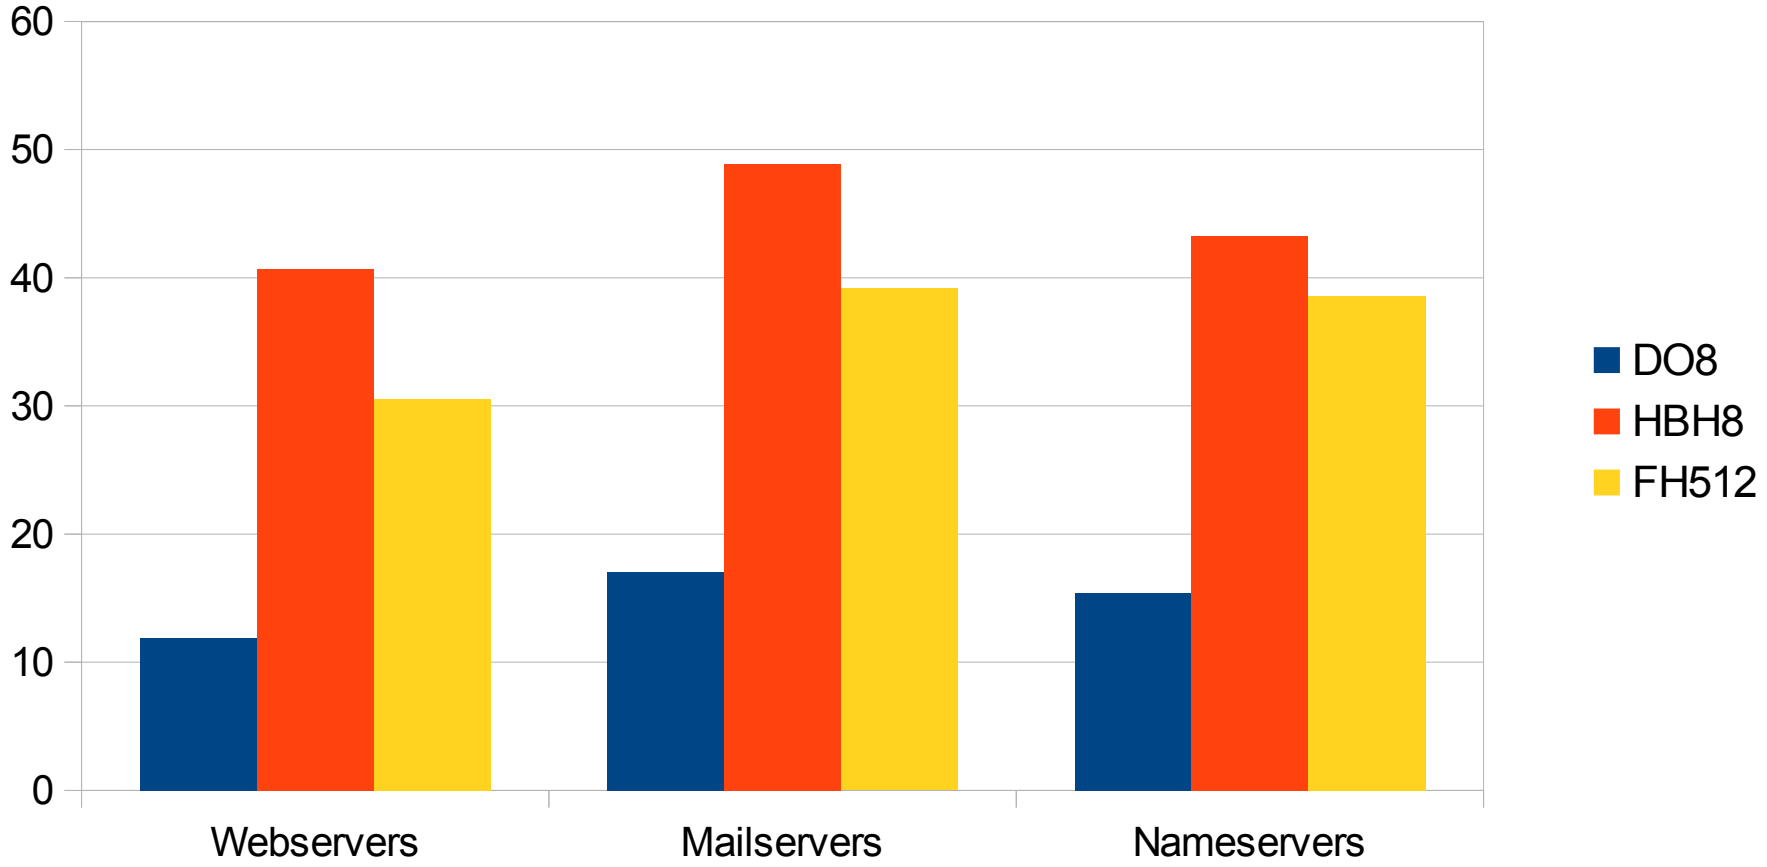

WIPv6LD dataset: Drops by diff. AS

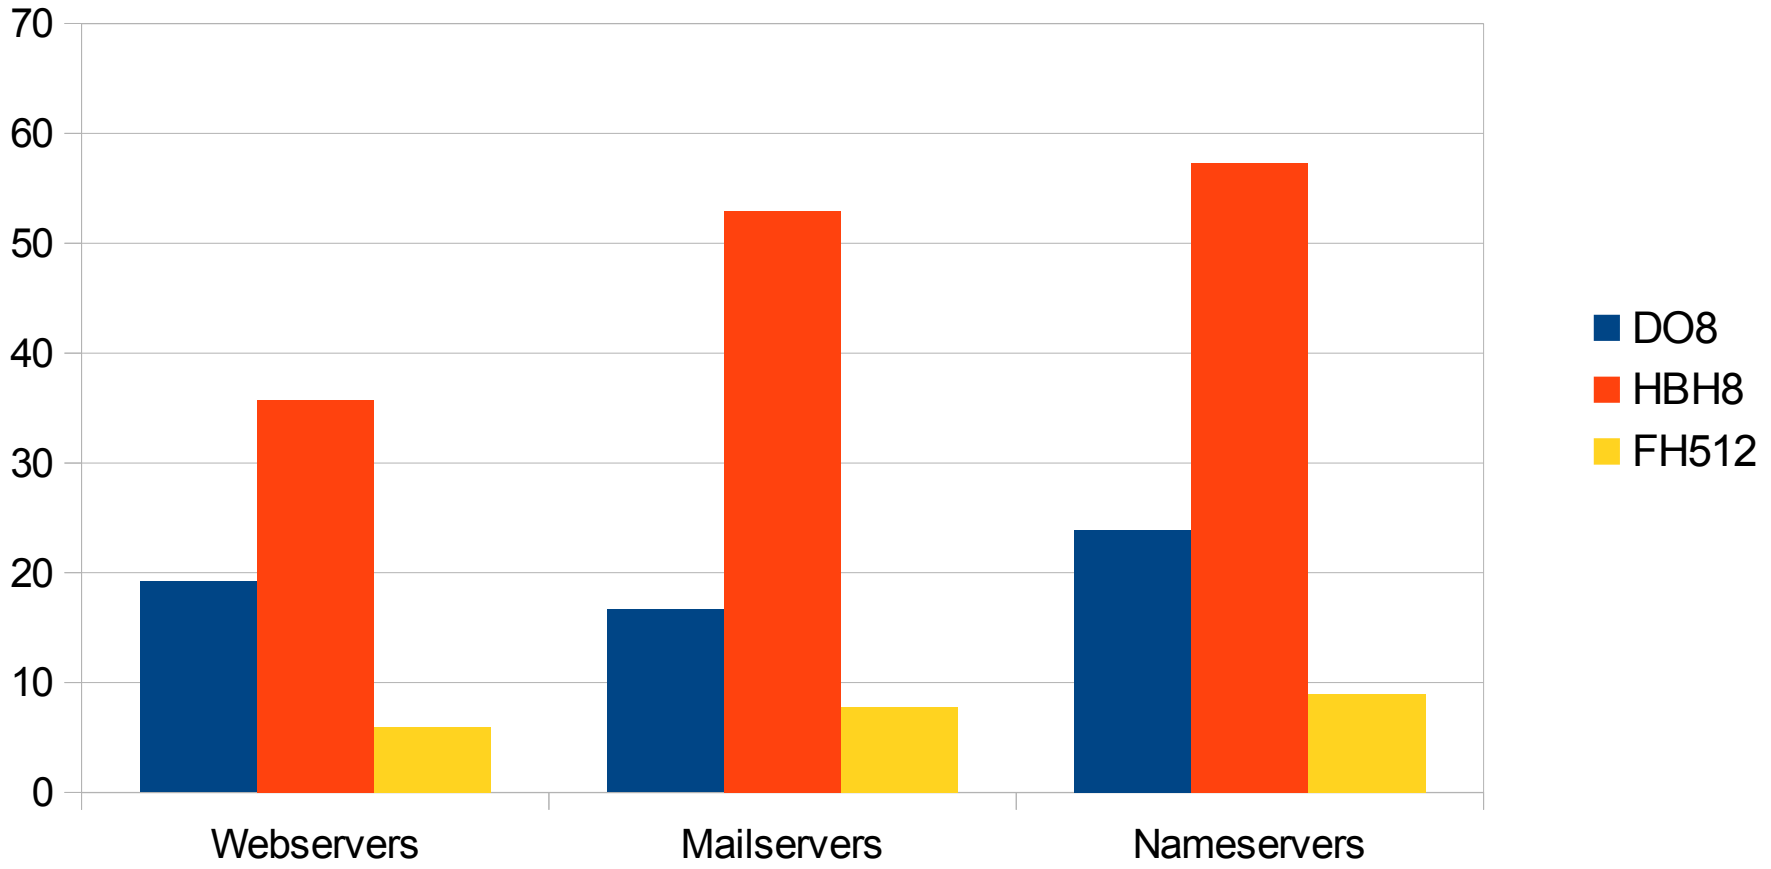

Alexa dataset: Packet Drop rate

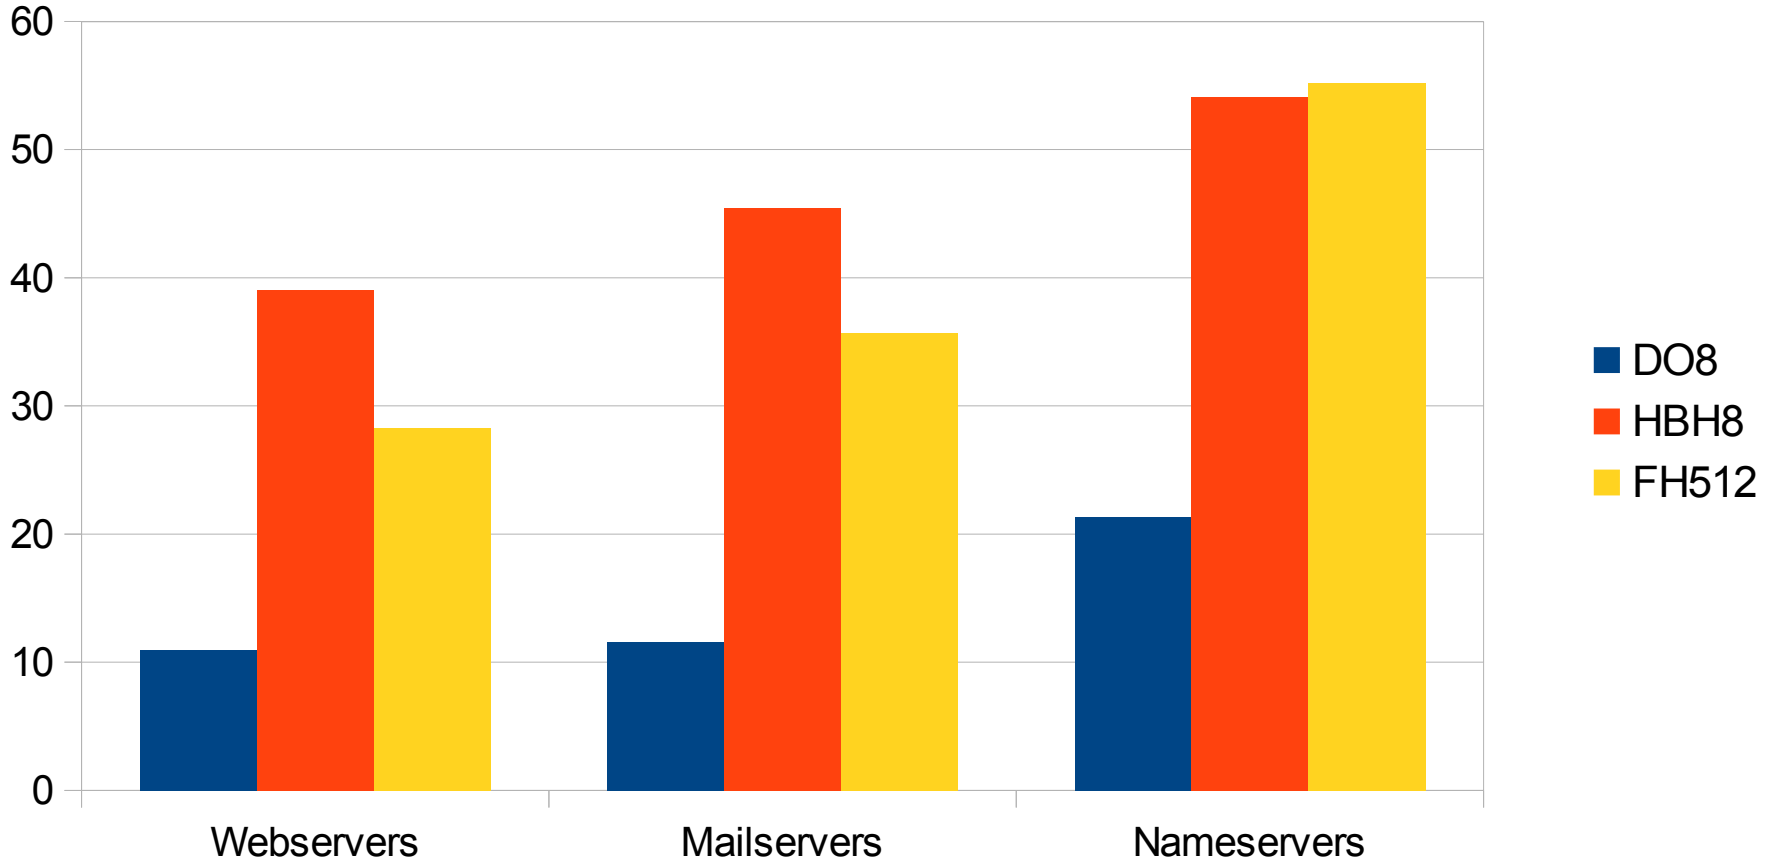

Alexa dataset: Drops by diff. AS

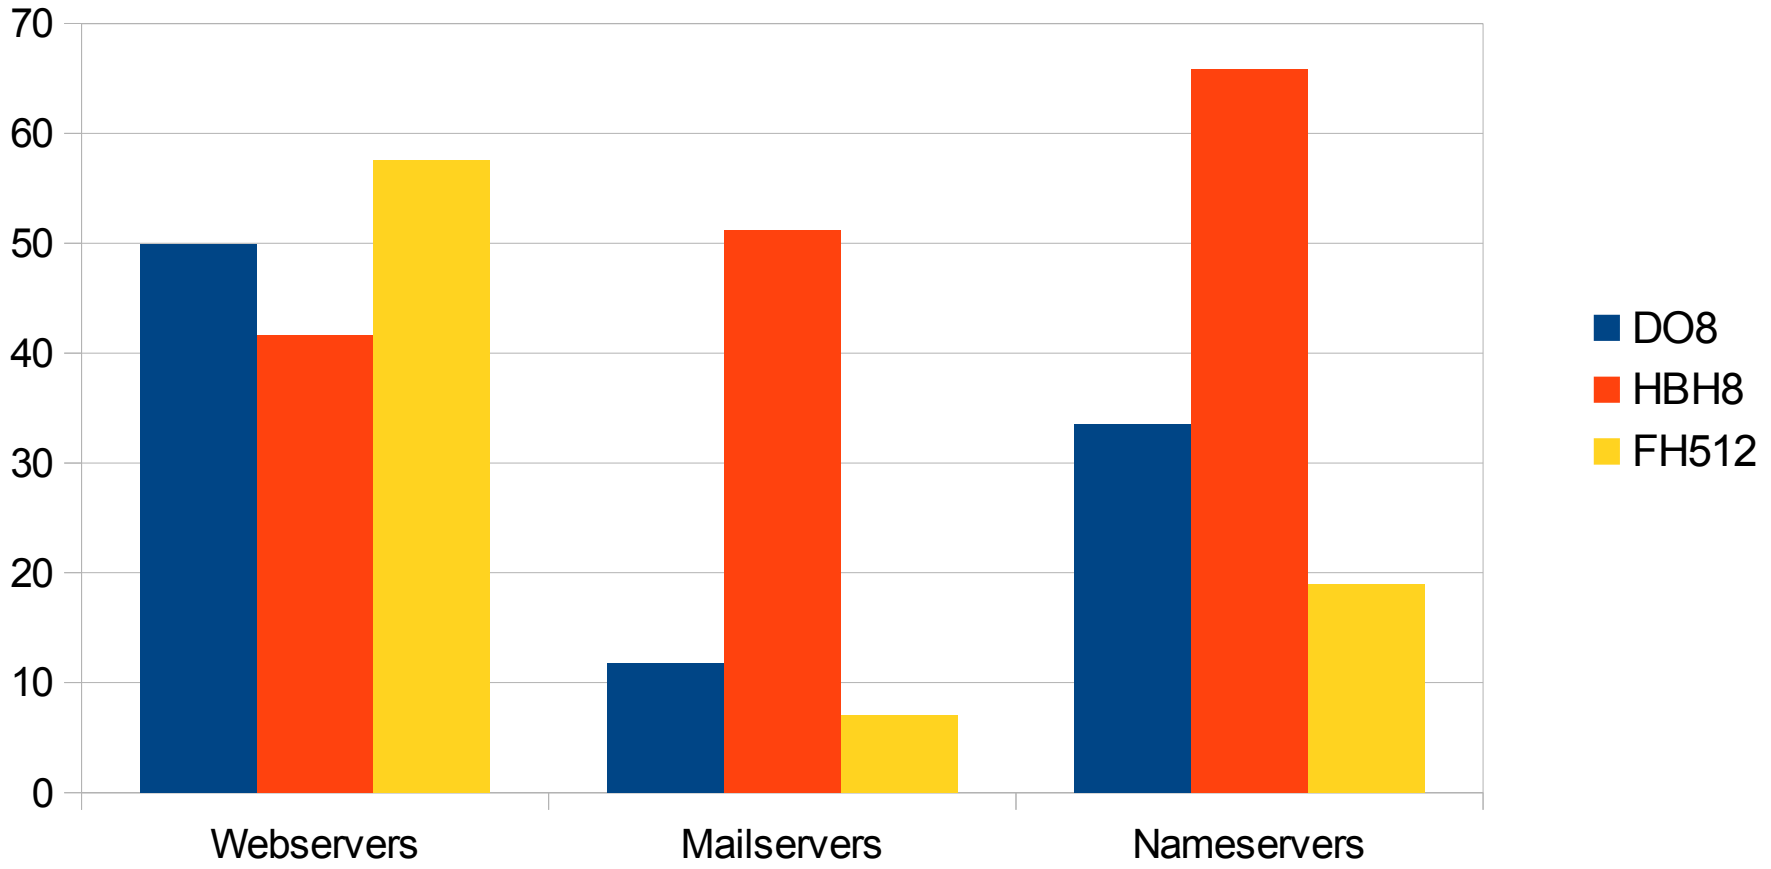

Future work?

- Fragment Header from servers to clients
- More details analysis of where packets are dropped
  - how many dropped by the host?
  - TCP vs UDP vs ICMP
- Repeat the measurement in 1 year
- ...any other ideas?

Future Work

Questions?