


Discovering Path MTU black holes in the Internet using RIPE Atlas

Maikel de Boer
Jeffrey Bosma

25 September 2012

Introduction

- Black holes
 - “A *sphere of influence into which or from which communication or similar activity is precluded.*”
 - ~ Wiktionary.org
 - In layman’s terms: what goes in is forever lost
 - The Internet is full of black holes
 - Many possible causes
 - E.g., misconfiguration, bugs in software, etc.
 - We focus on Path MTU black holes




Research Questions

Where on the Internet do Path MTU black holes occur?

Do Path MTU black holes occur more often on the IPv6-Internet compared to IPv4?

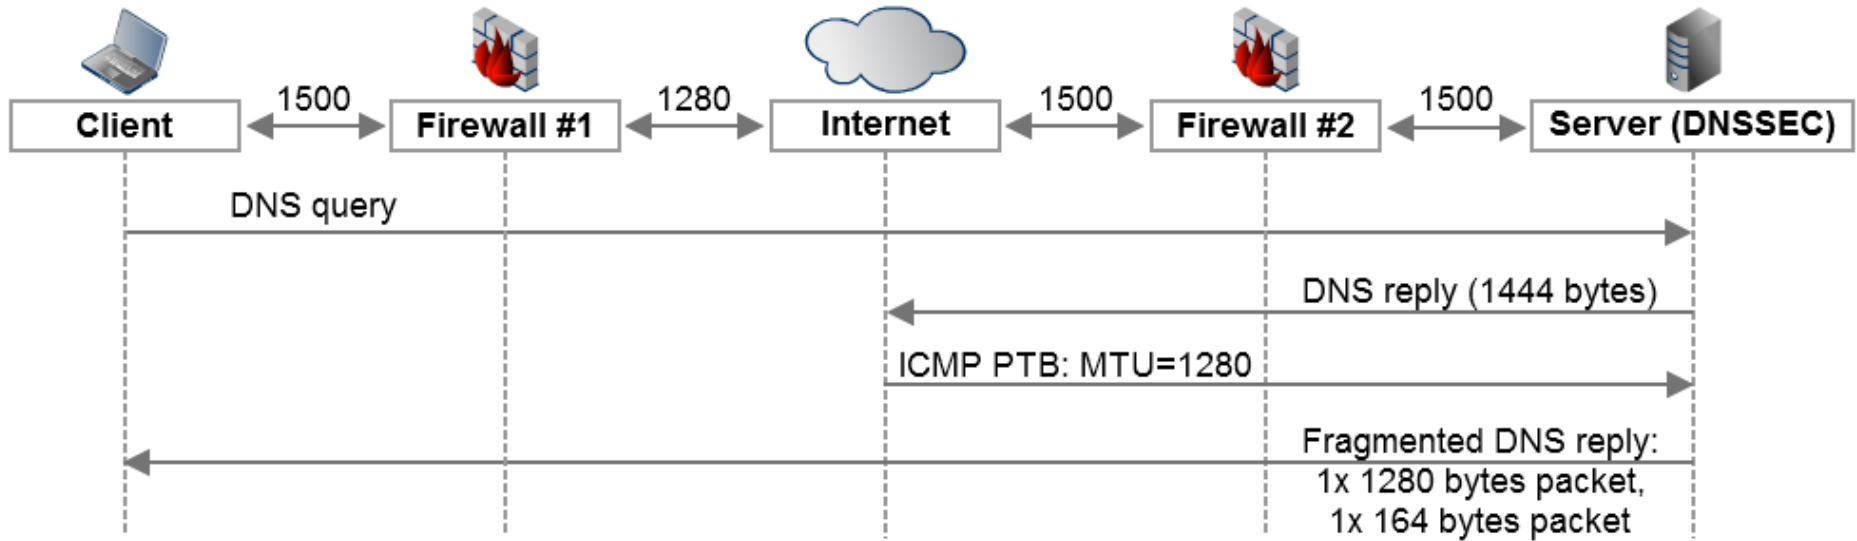
Concepts

- The Internet: enormous collection of links
- Maximum Transmission Unit (MTUs) on network interface
 - Limits the amount of data in packets
 - Two-way limit: sending and receiving
- Path MTU
 - Highest possible MTU for entire path
 - Determined by link with smallest MTU
- Internet Path MTU is commonly 1500 bytes
 - Not always the case
 - Requires Path MTU detection mechanism



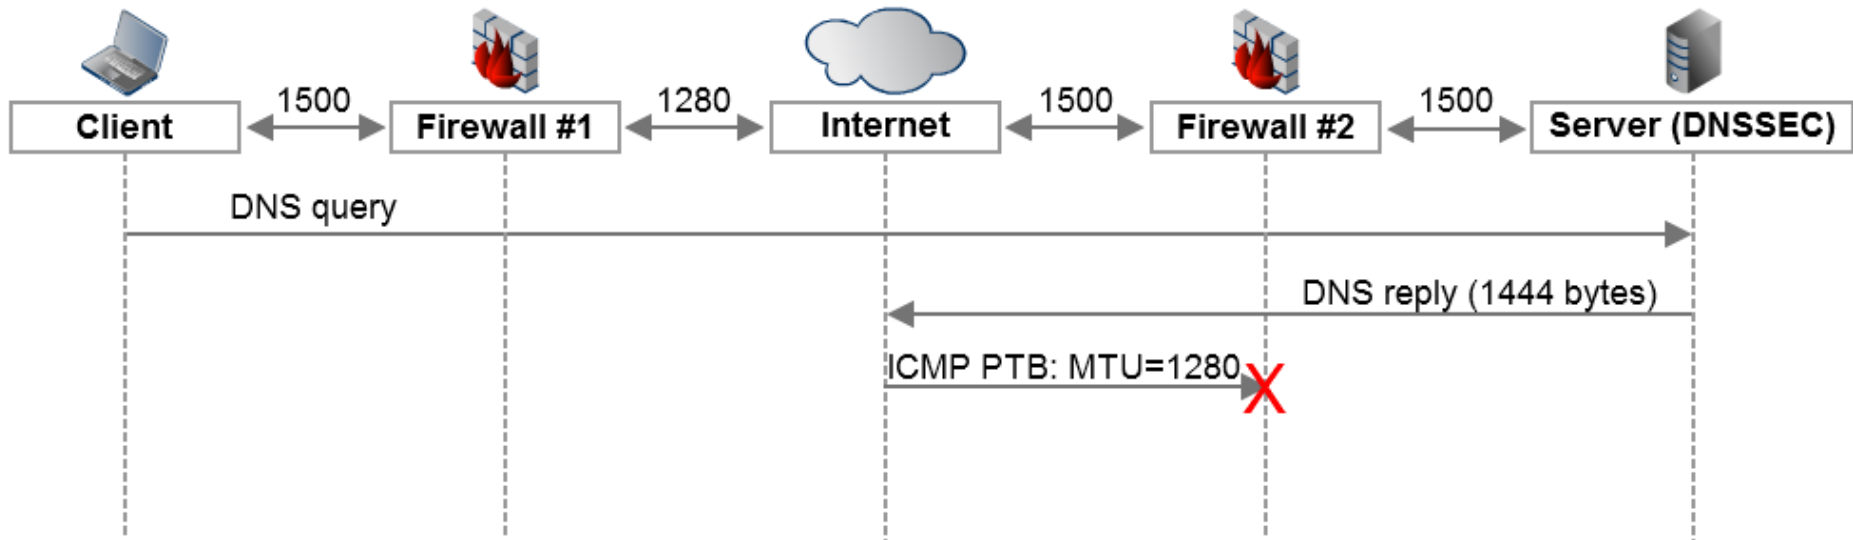
Concepts

Path MTU Discovery (PMTUD)



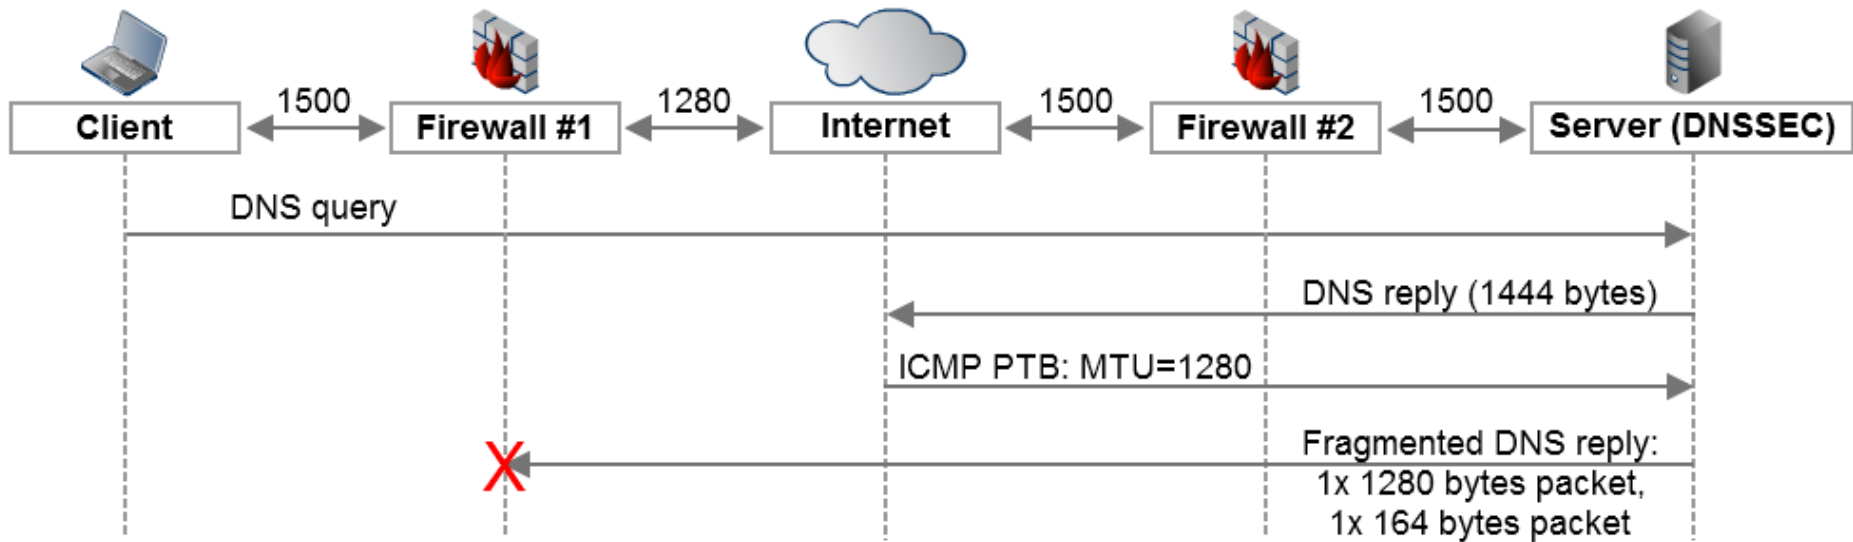
Concepts

Problem #1: ICMP PTB packet filtering




Concepts

Problem #2: IP fragment filtering




Concepts


Experimental setup



Option #1



Option #2



Option #3

Concepts

Experimental setup (cont.)

- Combination of option 1 and 2
 - Capture packets on interface of measurement node
 - Other nodes only need to log success/failure information of measurements
 - Con: no triangulation as possible with option 3
- RIPE Atlas infrastructure probes to actively measure ICMP PTB packet and IP fragment filtering

RIPE Atlas

- Internet measurement system
- Driven by probes
 - USB-powered embedded devices
 - Low volume, non-intrusive measurements
- Default measurement functionality:
 - Ping, Traceroute, DNS, SSL
- Currently around 1800 probes up and running
 - Located primarily in the RIPE NCC service region
 - But also other regions around the globe



RIPE Atlas

Worldwide network of probes



Research Questions

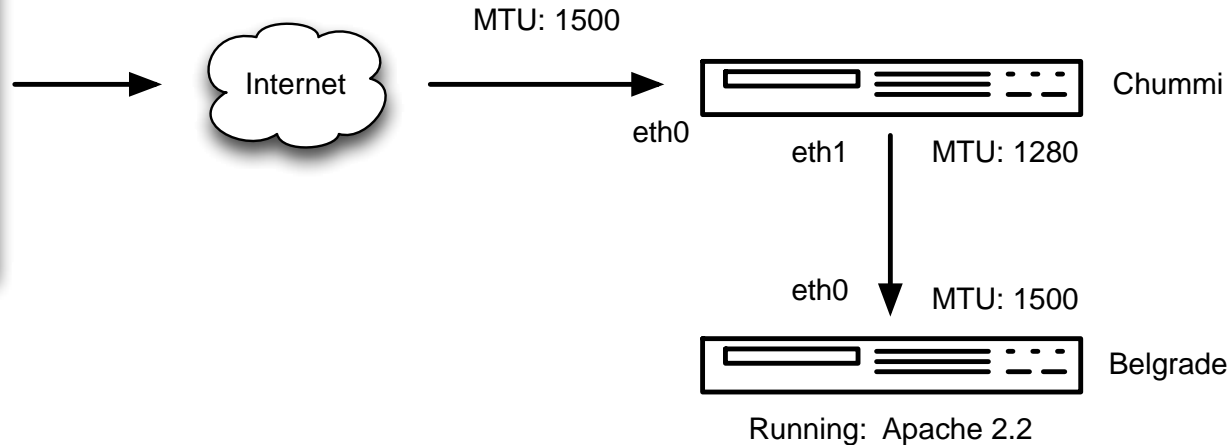
Where on the Internet do Path MTU black holes occur?

Do Path MTU black holes occur more often on the IPv6-Internet compared to IPv4?

Experimental Setup

ICMP PTB packet filtering

```
POST / HTTP/1.1
Host: httppost6.uranus.nlnetlabs.nl
Connection: close
User-Agent: httpget for
atlas.ripe.net
Content-Type: application/x-www-
form-urlencoded
Content-Length: 65528
```

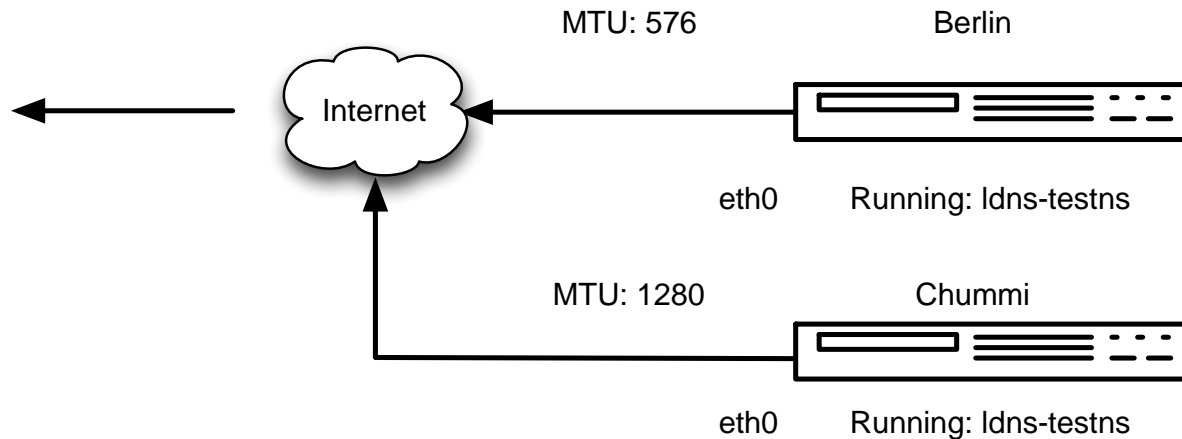


Experimental Setup

IP fragment filtering


```
version.bind. 60 CH TXT
1,002,003,004,005,006,007,008,009,
010,011,012,013,014,015,016,017,01
8,019,020,021,022,023,024,025,
[truncated]
347,348,349,350,351,352,353,354,35
5,356,357,358,359,360,361,362,363,
364,365,366,367,368,369,370,371,37
2,373,374,375,376,377,378,379,380,
381,382,383
```

MSG SIZE snd: 1590




Results

ICMP PTB filtering IPv4 – MTU 1500




Results

ICMP PTB filtering IPv6 – MTU 1500



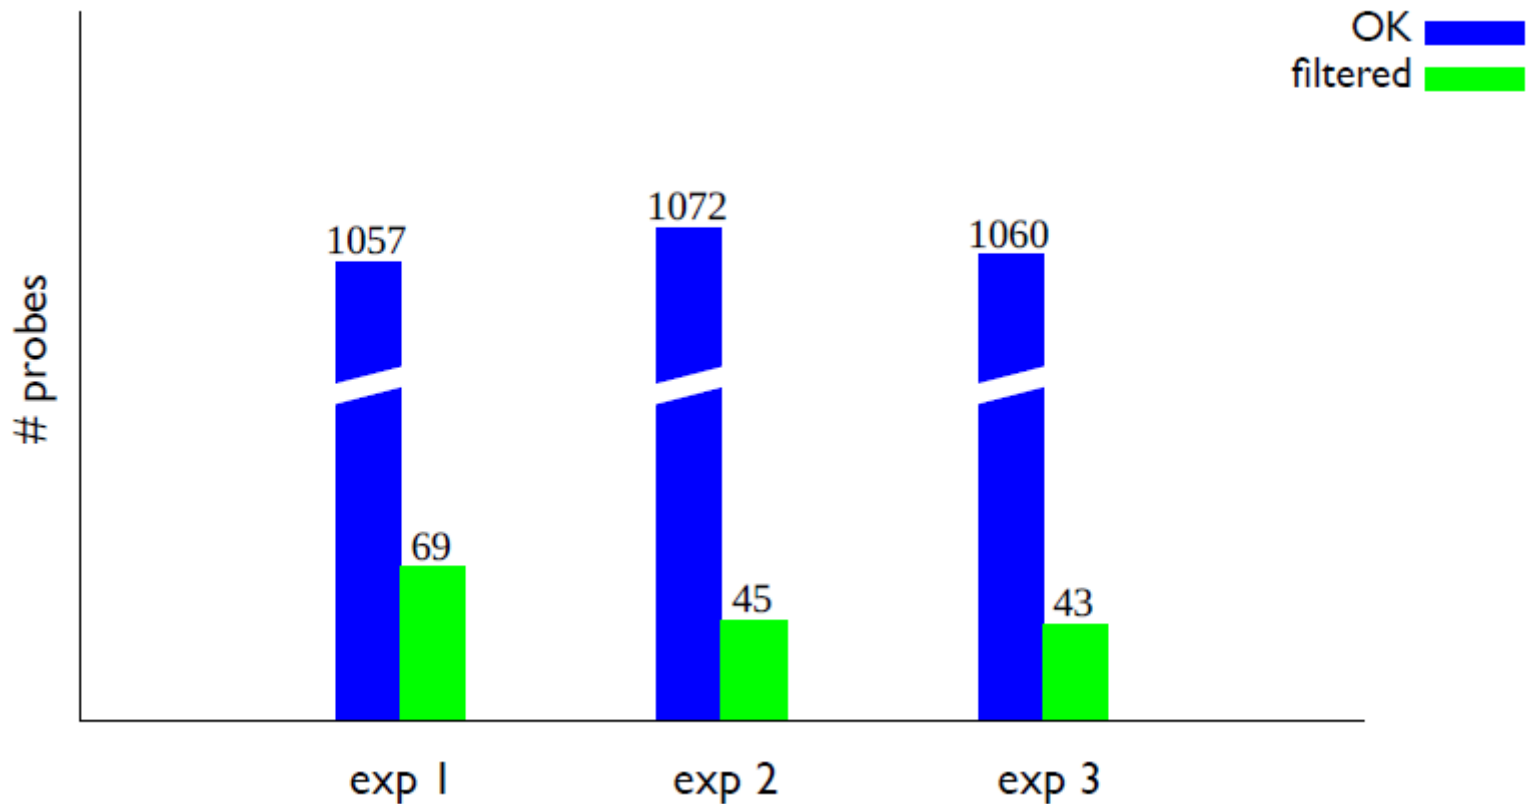
Results

ICMP PTB filtering percentages – MTU 1500




Results

ICMP PTB filtering IPv4 – MTU 1280



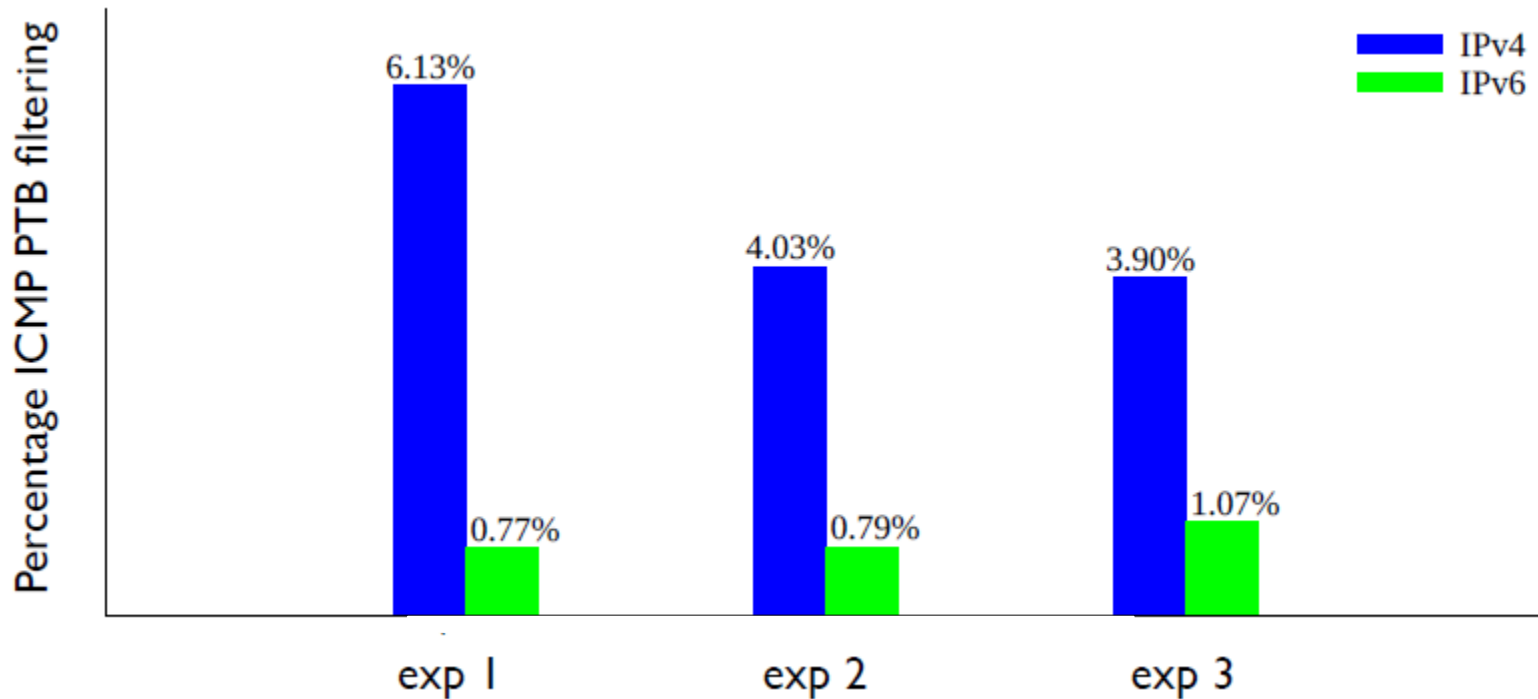
Results

ICMP PTB filtering IPv6 – MTU 1280



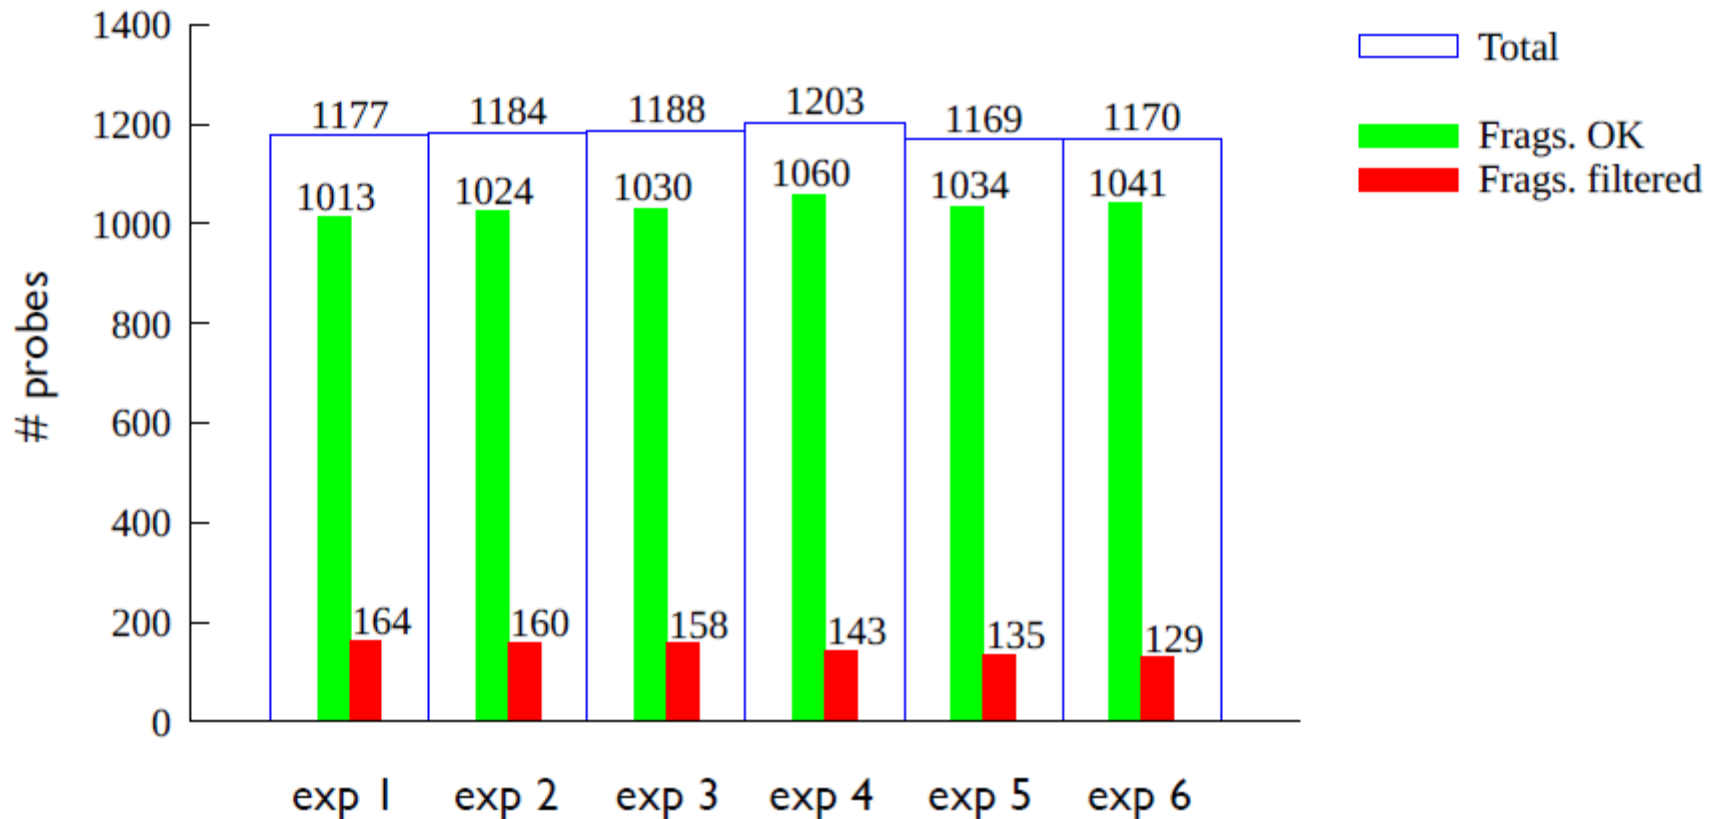
Results

ICMP PTB filtering percentages – MTU 1280



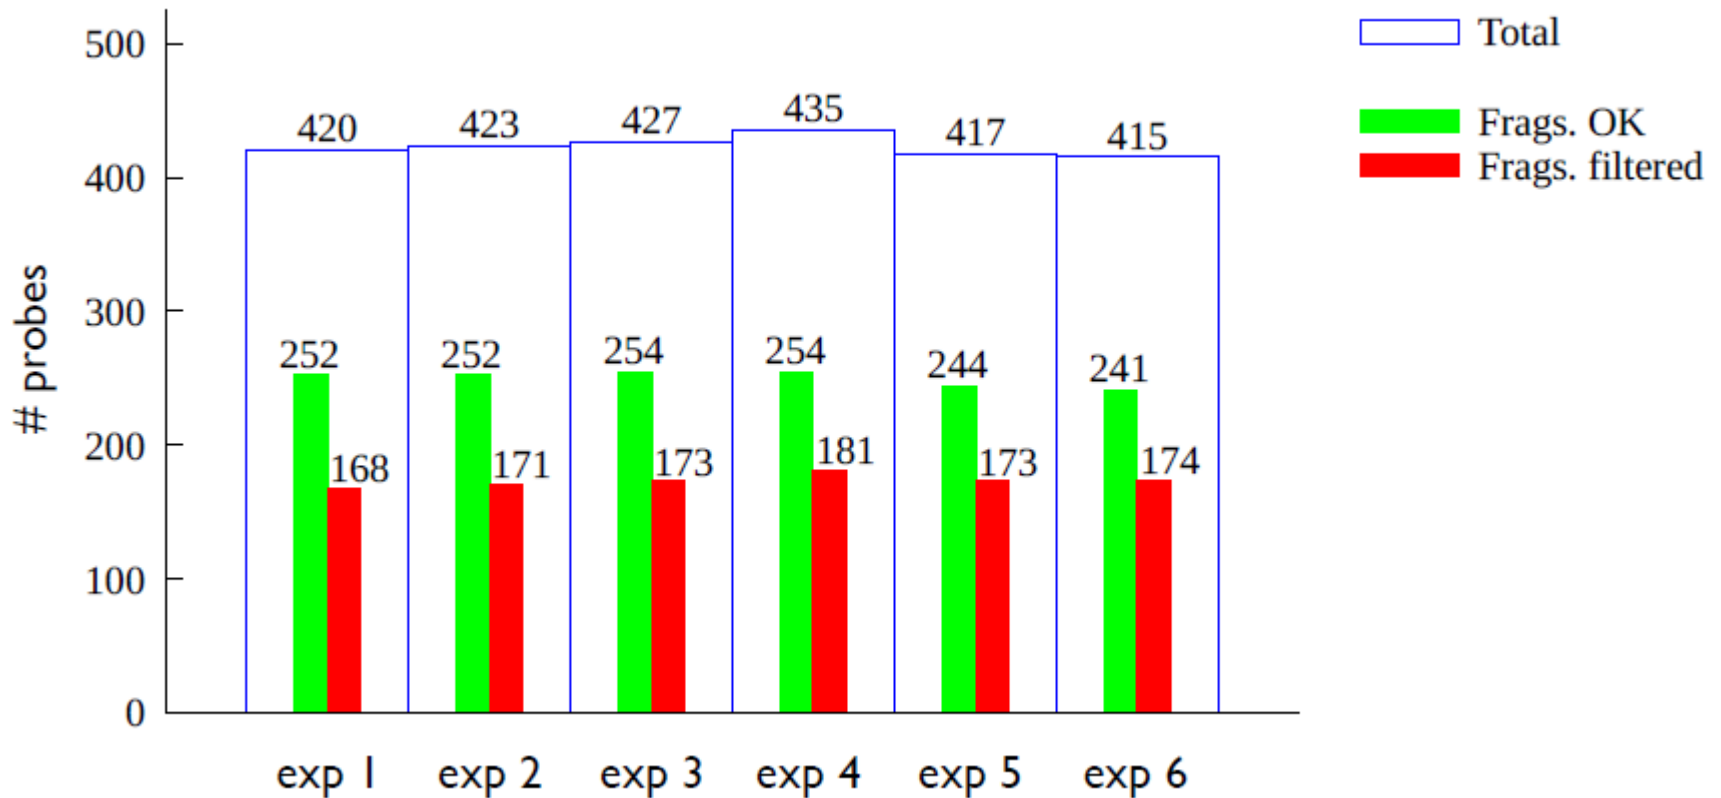
Results

Fragment filtering IPv4 – MTU 1500



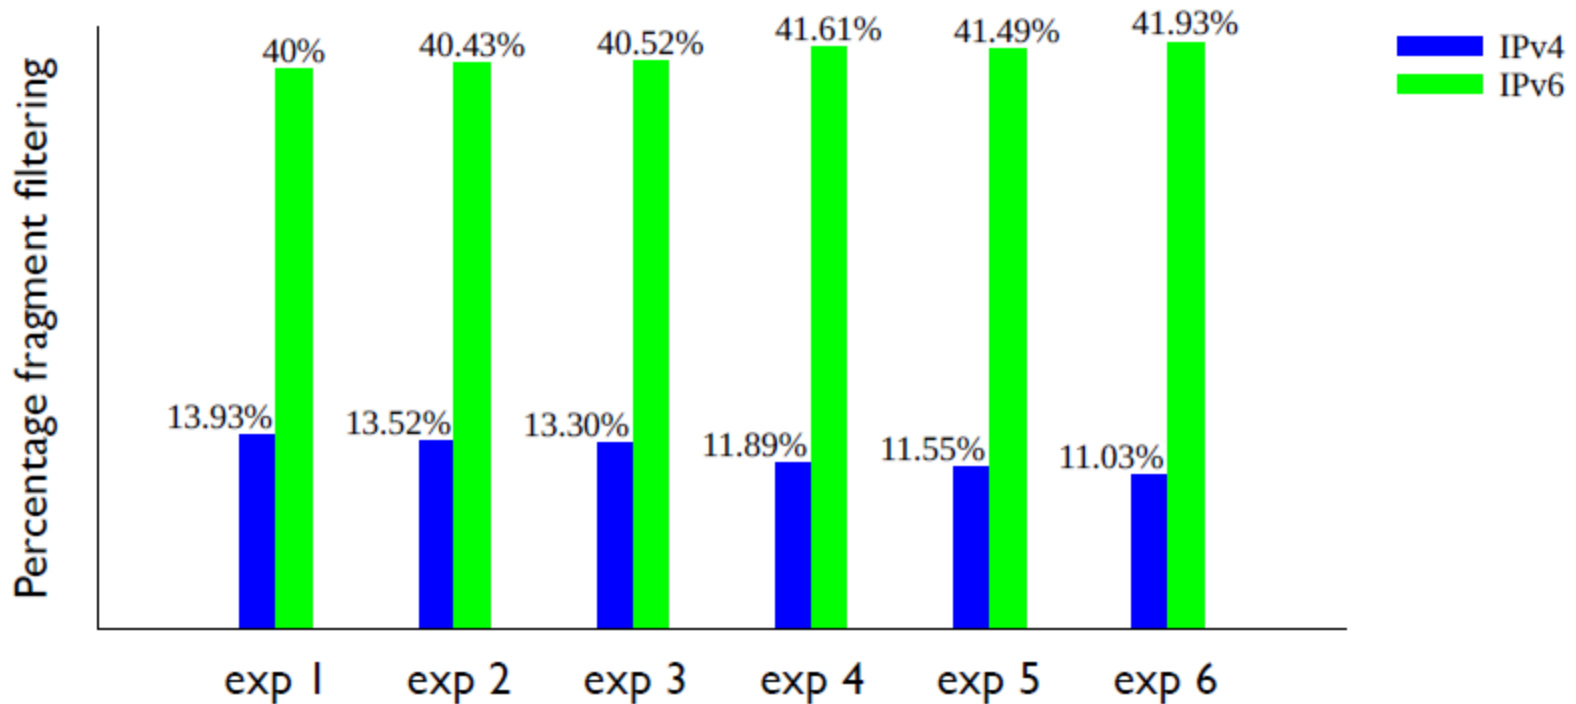
Results

Fragment filtering IPv6 – MTU 1500



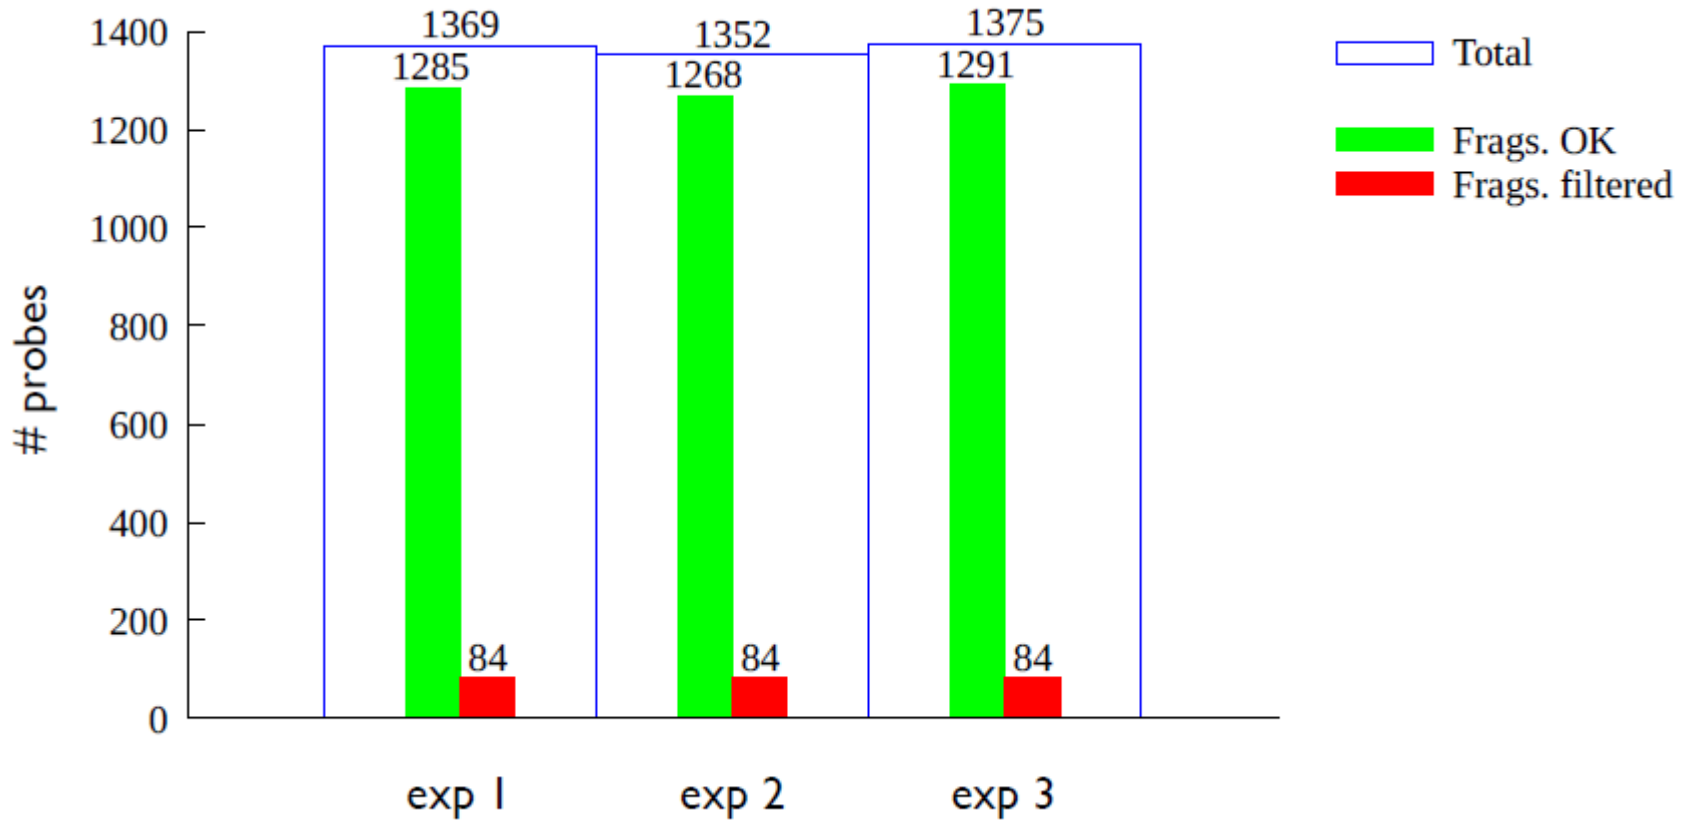
Results

Fragment filtering percentages – MTU 1500




Results

Fragment filtering IPv4 – MTU 576




Results

Fragment filtering IPv6 – MTU 1280

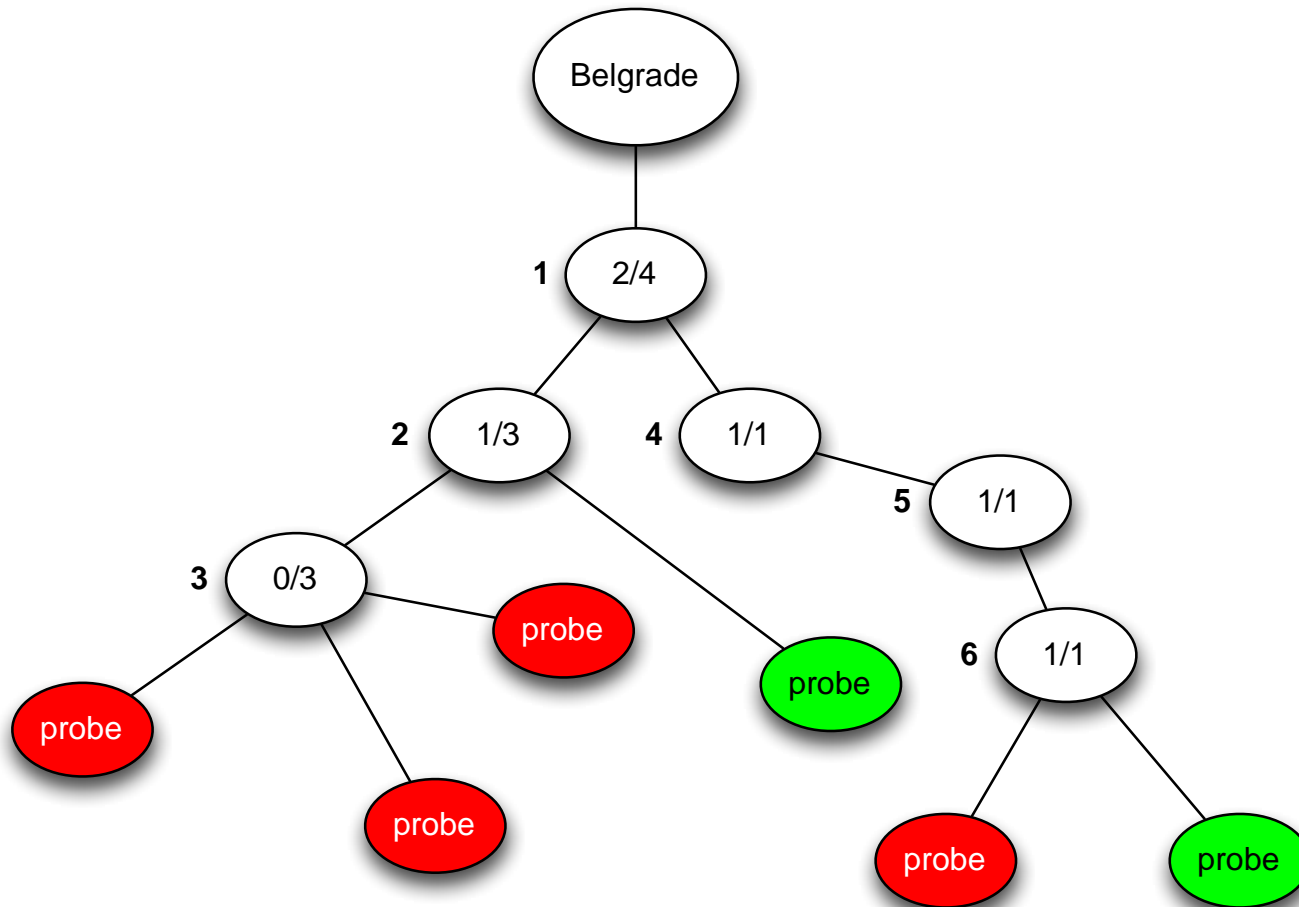


Results

Fragment filtering percentages – MTU 576/1280



Hop counting



Results

Where do IPv4 ICMP PTB packets get filtered?

Bad	Total	Error percentage	IP address
69	1126	6.1%	145.145.19.190
53	810	6.5%	145.145.80.65
16	311	5.1%	145.145.80.73
13	214	6.1%	77.67.72.109
7	199	3.5%	109.105.98.33
2	60	3.3%	62.40.124.157
...			
2	2	100.0%	203.50.6.78
2	2	100.0%	203.50.6.89
2	2	100.0%	61.10.0.118
2	2	100.0%	80.231.159.10
2	2	100.0%	84.116.238.49

Results

Where do IPv6 ICMP PTB packets get filtered?

Bad	Total	Error percentage	IP address
3	391	0.8%	2001:610:158:1916:145:100:99:17
2	292	0.7%	2001:610:e08:64::65
2	131	1.5%	2001:7f8:1::a500:6939:1
1	9	11.1%	2001:470:0:217::2
1	6	16.7%	2001:470:0:67::2
1	46	2.2%	2001:470:0:3f::1
...			
No routers with 100% failure rate			

Results

Where do IPv4 fragments get filtered?

Bad	Total	Error percentage	IP address
84	1369	6.1%	145.145.19.190
56	983	5.7%	145.145.80.65
28	381	7.3%	145.145.80.73
14	256	5.5%	109.105.98.33
21	247	8.5%	77.67.72.109
9	62	14.5%	62.40.124.157
...			
3	3	100.0%	212.188.22.158
2	2	100.0%	146.97.33.137
2	2	100.0%	158.64.16.189
2	2	100.0%	174.35.131.38
2	2	100.0%	188.230.128.10

Results

Where do IPv6 fragments get filtered?

Bad	Total	Error percentage	IP address
181	435	41.6%	2001:610:158:1916:145:100:99:17
138	322	42.9%	2001:610:e08:64::65
74	146	50.7%	2001:7f8:1::a500:6939:1
28	53	52.8%	2001:470:0:3f::1
27	91	29.7%	2001:610:e08:72::73
21	53	39.6%	2001:948:2:6::1
...			
6	6	100.0%	2001:610:f01:9012::14
4	4	100.0%	2001:16d8:aaaa:5::2
4	4	100.0%	2001:7f8:1::a503:9326:1
4	4	100.0%	2a01:348::10:0:1
4	4	100.0%	2a01:348::27:0:1

Conclusion

- ICMP PTB packets get dropped
 - More for IPv4 but nobody notices
 - Complete path is 1500 bytes and DF bit helps
- IP fragments get dropped
 - More in IPv6
 - DNS servers do not respond to ICMP PTB
- Path MTU black holes
 - Occur on the edges of the Internet, not in the core

Recommendations

- Recommendations for Filtering ICMPv6 Messages in Firewalls – RFC4890
- Don't filter IPv4 ICMP Type 3, Code 4
- Packetization Layer Path MTU Discovery – RFC4821
- Don't filter fragments (problematic for DNSSEC)
- Don't reduce MTU on interface
- No MSS clamping

Acknowledgements

- NLnet Labs
 - Benno Overeinder
 - Willem Toorop
- RIPE NCC
 - Philip Homburg
 - Andreas Strikos
 - Vesna Manojlovic
 - Emile Aben

Questions?

Maikel de Boer – maikel.deboer@os3.nl

Jeffrey Bosma – jeffrey.bosma@os3.nl

Read our full report @ www.nlnetlabs.nl/publications/