Evaluating DNS Resiliency and Responsiveness with Truncation, Fragmentation and DoTCP Fallback

Pratyush Dikshit | CISPA Helmholtz Center for Information Security, DE
Mike Kosek, Nils Faulhaber | Technical University of Munich, DE
***Jayasree Sengupta**, Vaibhav Bajpai | CISPA Helmholtz Center for Information Security, DE*

Published at IFIP Networking Conference, June 2023

RIPE87, Rome
Plenary (Onsite)| 28 November 2023

Overview

Truncation

When DNS responses over UDP exceed the buffer size limit (due to DNSSec/IPv6), **truncation bit (TC) is set**. This signals the resolvers and clients that the message could not be transferred correctly.

Extension Mechanisms of DNS (EDNS)

Buffer **sizes ranging from (512-4096)B** over DoUDP.

EDNS is also used for sending general information **from resolvers to name servers** about clients' geographic location in the form of the **EDNS Client Subnet (ECS)** option

DNS Flag Day, 2020

This is an event connecting important DNS providers to react to current research and new developments in the ecosystem. It is supported by the DNS Operations Analysis and Research Center (DNS-OARC).

"default in the DNS software should reflect *the minimum safe size -1232B"*

Fragmentation

IPv4 allows fragmentation, which **divides the datagram into pieces**. Each piece is small enough to pass over the link it is fragmented for, using the MTU parameter configured for that interface.

The IPv6 sender performs fragmentation at the source.

Motivation

DNS-over-UDP (DoUDP)	DNS-over-TCP (DoTCP)
Limited Payload Size (512B) -> Truncation	Unlimited Payload Size -> No Truncation
Introduction of EDNS -> Larger Buffer Sizes	Path MTU Discovery -> Fragmentation avoidance
Fragmentation -> Default Buffer Size: 1232B	Fallback option

- DoTCP is mandatory for hosts
- Introduction of EDNS (to 4096 bytes)
- DNS Flag Day 2020 recommended 1232 bytes of UDP buffer size

Jayasree Sengupta | jayasree.sengupta@cispa.de | CISPA Helmholtz Center for Information Security

Research Questions and Findings

Goals:

- Investigate DNS Resiliency
- Whether DNS service providers have adopted the DNS Flag Day 2020 recommendations

Questions:

- How resilient is DoTCP and DoUDP over IPv4 and IPv6?
- What is the scale of usage and performance of both DoTCP and DoUDP?
- Which buffer sizes are currently in use in DNS traffic around the globe?

Findings:

- DoTCP tends to fail less often than DoUDP over IPv4
- Several Public DNS resolvers still lack adoption to the DNS Flag Day 2020 recommendations

Jayasree Sengupta | jayasree.sengupta@cispa.de | CISPA Helmholtz Center for Information Security

Methodology

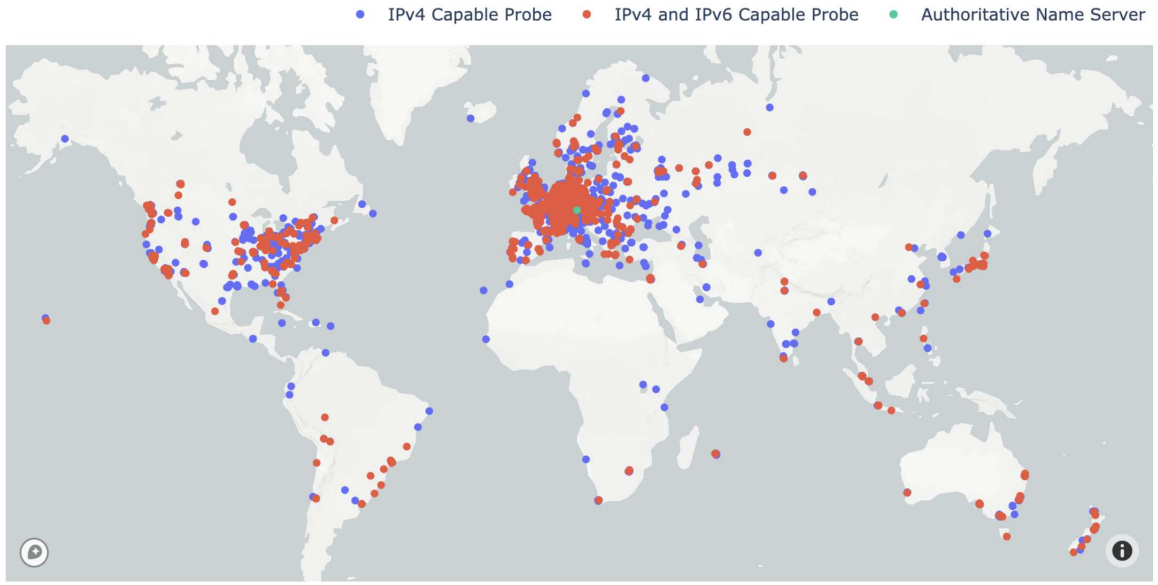

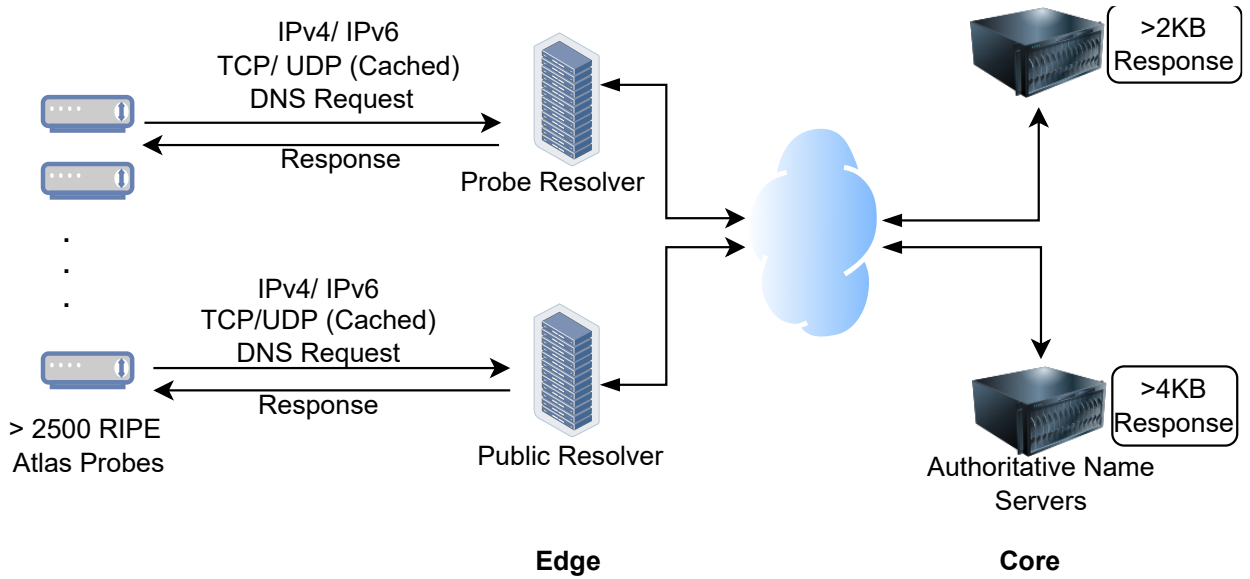

- Probe (versions 3 and 4) with home-tag
- Probes supporting IPv4/IPv6 or both
- 2527 globally distributed RIPE Atlas probes
- Probe Density:
  - 70% in Europe
  - 12% in North America
  - 6% in Asia and 3% in Oceania

- RIPE Atlas Probes communicate the DNS requests with the Edge (Probe and Public Resolvers) and with the Core (authoritative NSes) using IPv4 and IPv6.

- Cached DNS responses are sent by the Edge, while uncached DNS responses (2KB and 4KB) are sent by the Core

Jayasree Sengupta | jayasree.sengupta@cispa.de | CISPA Helmholtz Center for Information Security

Methodology (Evaluation from the Edge)

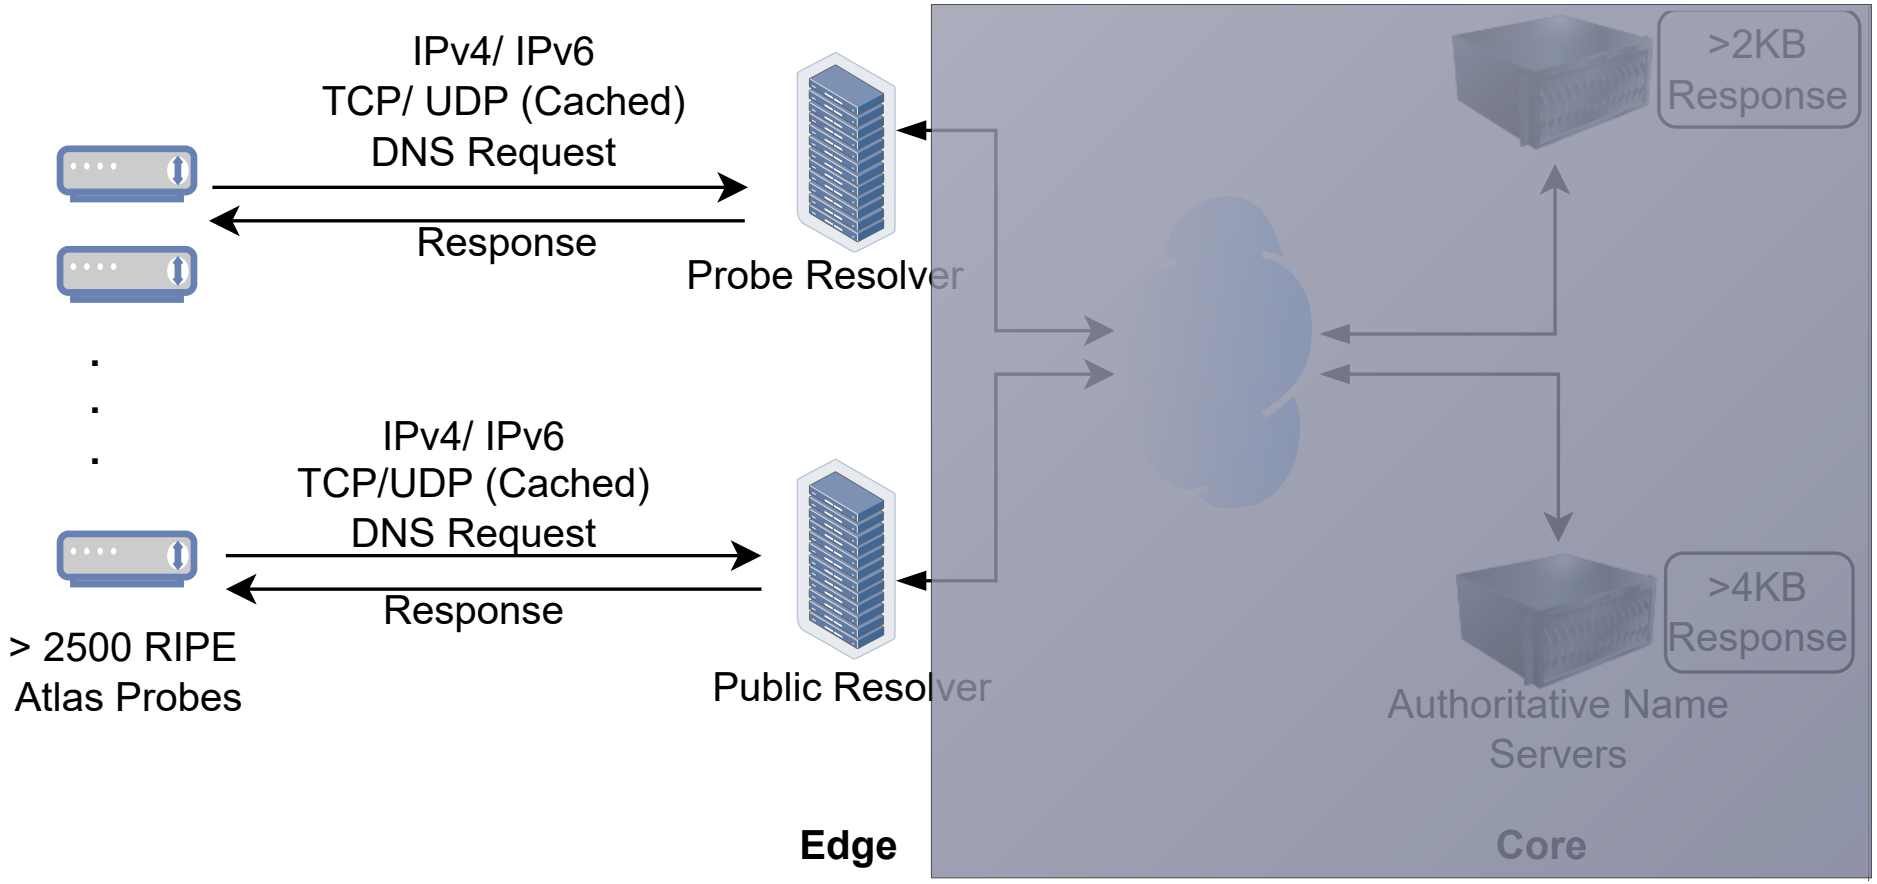

- **One week, 10 blocks every day, 10 requests per block per resolver**

Jayasree Sengupta | jayasree.sengupta@cispa.de | CISPA Helmholtz Center for Information Security

Findings (Evaluation from the Edge - Failure Rate over IPv4)

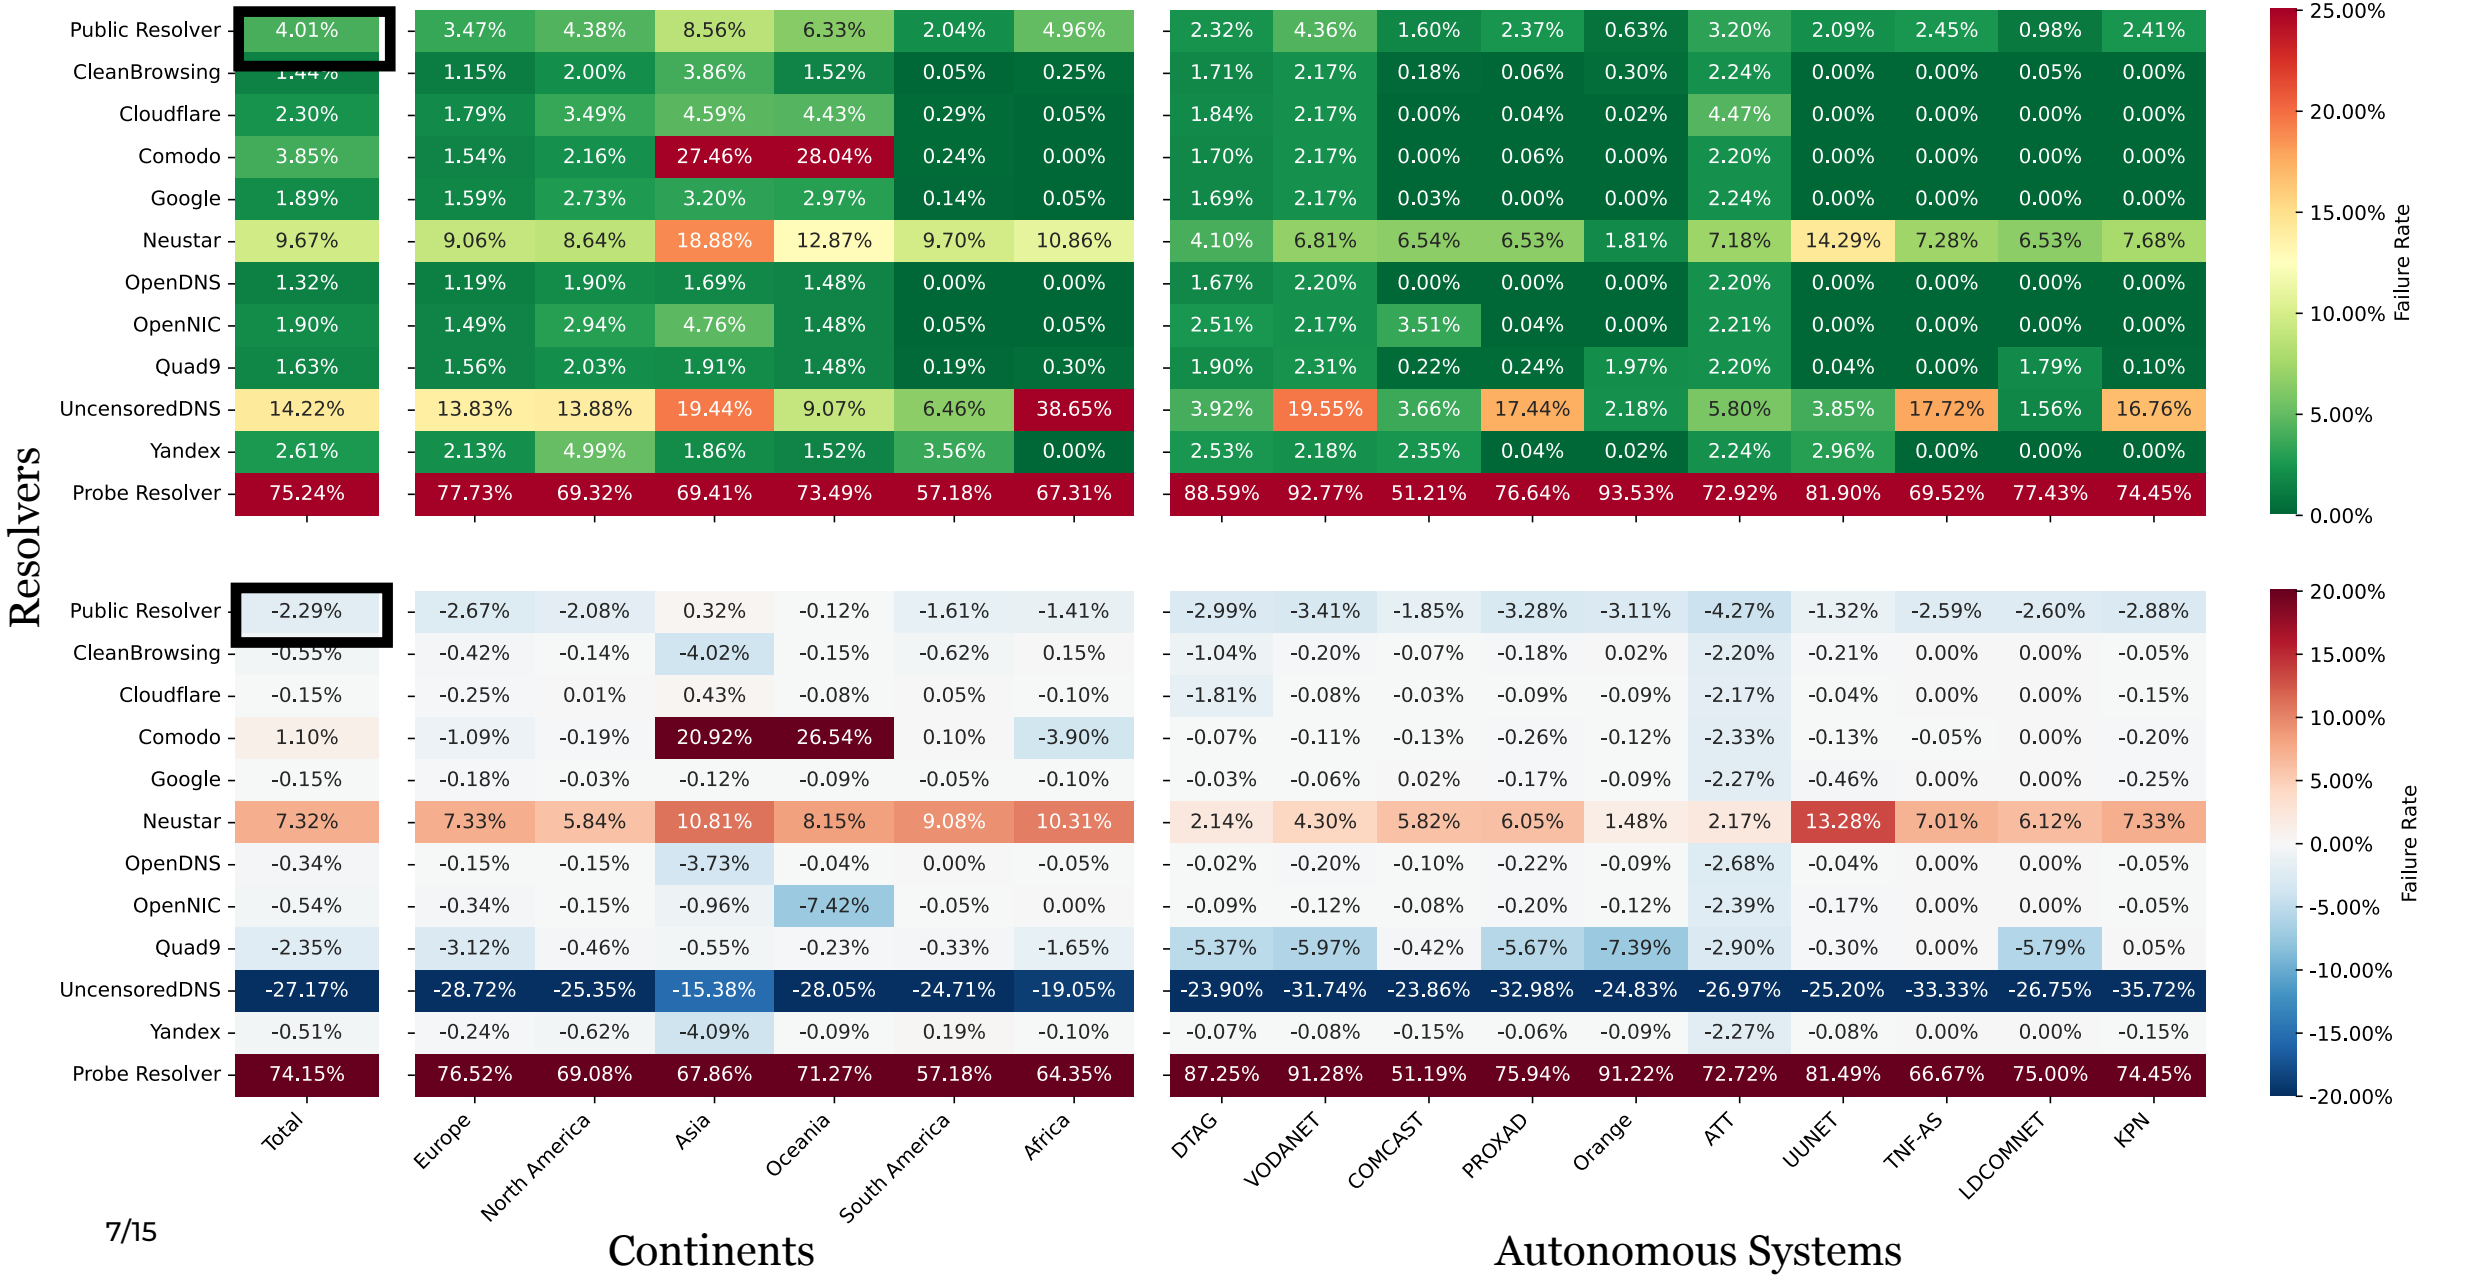

Findings (Evaluation from the Edge - EDNS(0))

		512	1232	4096	none	other
CleanBrowsing	*IPv4*	**97.04%**	0.63%	1.46%	0.57%	0.30%
	IPv6	**99.41%**	0.11%	0.48%	0.01%	0.00%
Cloudflare	*IPv4*	0.20%	**97.43%**	1.45%	0.53%	0.40%
	IPv6	0.11%	**99.44%**	0.44%	0.01%	0.00%
Comodo	*IPv4*	0.18%	0.64%	**98.30%**	0.57%	0.30%
	IPv6	-	-	-	-	-
Google	*IPv4*	**96.82%**	0.78%	1.47%	0.58%	0.34%
	IPv6	**99.22%**	0.10%	0.67%	0.00%	0.01%
Neustar	*IPv4*	0.18%	0.64%	**98.32%**	0.56%	0.30%
	IPv6	0.10%	0.10%	**99.79%**	0.00%	0.00%
OpenDNS	*IPv4*	0.18%	0.63%	**98.20%**	0.57%	0.43%
	IPv6	0.10%	0.11%	**99.79%**	0.00%	0.00%
OpenNIC	*IPv4*	0.18%	**97.53%**	1.42%	0.56%	0.30%
	IPv6	0.11%	**99.43%**	0.47%	0.00%	0.00%
Quad9	*IPv4*	**19.15%**	**55.47%**	1.48%	**23.55%**	0.35%
	IPv6	**20.98%**	**62.09%**	0.47%	**16.46%**	0.00%
UncensoredDNS	*IPv4*	0.30%	**95.87%**	2.39%	0.96%	0.49%
	IPv6	0.13%	**99.29%**	0.57%	0.01%	0.00%
Yandex	*IPv4*	0.19%	0.64%	**98.04%**	0.75%	0.39%
	IPv6	0.11%	0.10%	**99.79%**	0.00%	0.00%
Overall	*IPv4*	24.97%	36.12%	35.30%	3.24%	0.36%
	IPv6	24.86%	38.81%	34.46%	1.87%	0.00%

Takeaways:

- 4/10 resolvers use 1232B buffer size
- 4/10 resolvers use very large buffer size (4096B)
- The differences in the buffer sizes of IPv4 and IPv6 are low (<3.5%) except Quad9
- 1/4 - 1/5 times Quad9 does not use EDNS

Jayasree Sengupta | jayasree.sengupta@cispa.de | CISPA Helmholtz Center for Information Security

Methodology (Evaluation from the Core)

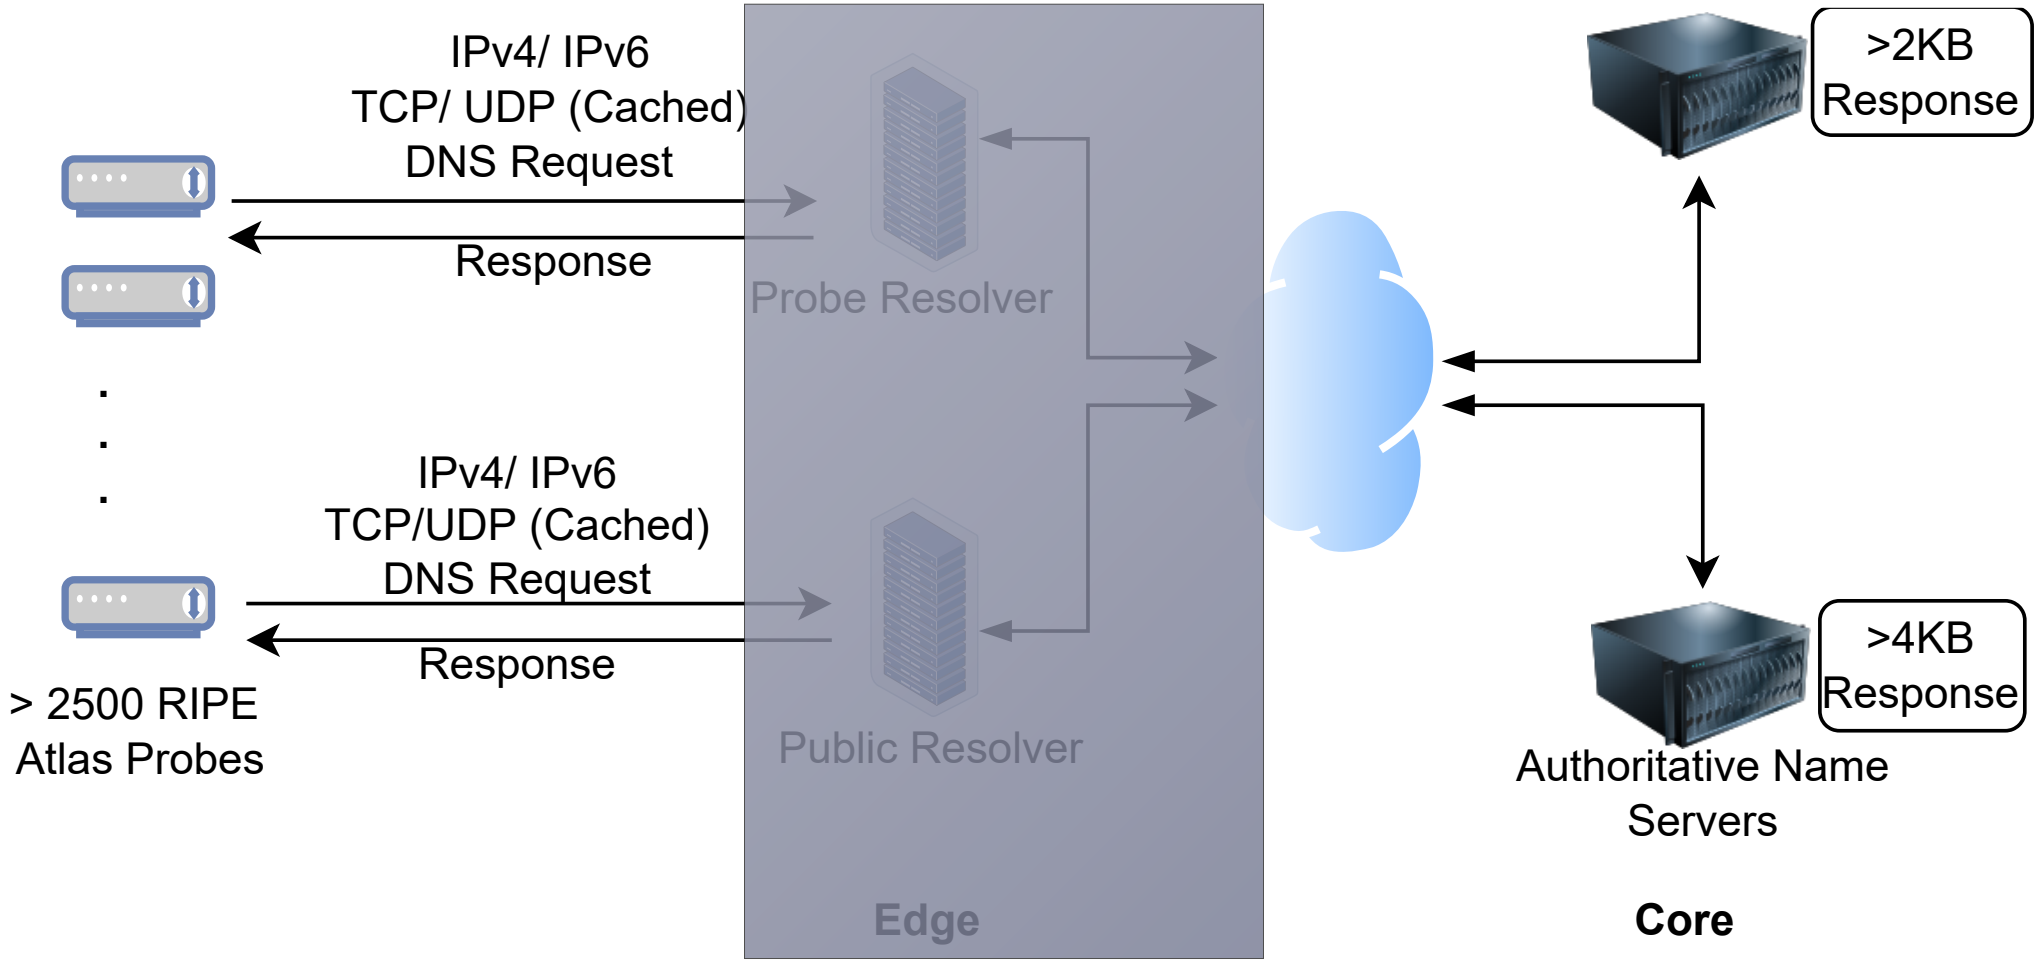

- **One week, 10 blocks every day, 10 requests per block per resolver**

Jayasree Sengupta | jayasree.sengupta@cispa.de | CISPA Helmholtz Center for Information Security

Findings (Evaluation from the Core - Failure Rate over IPv4)

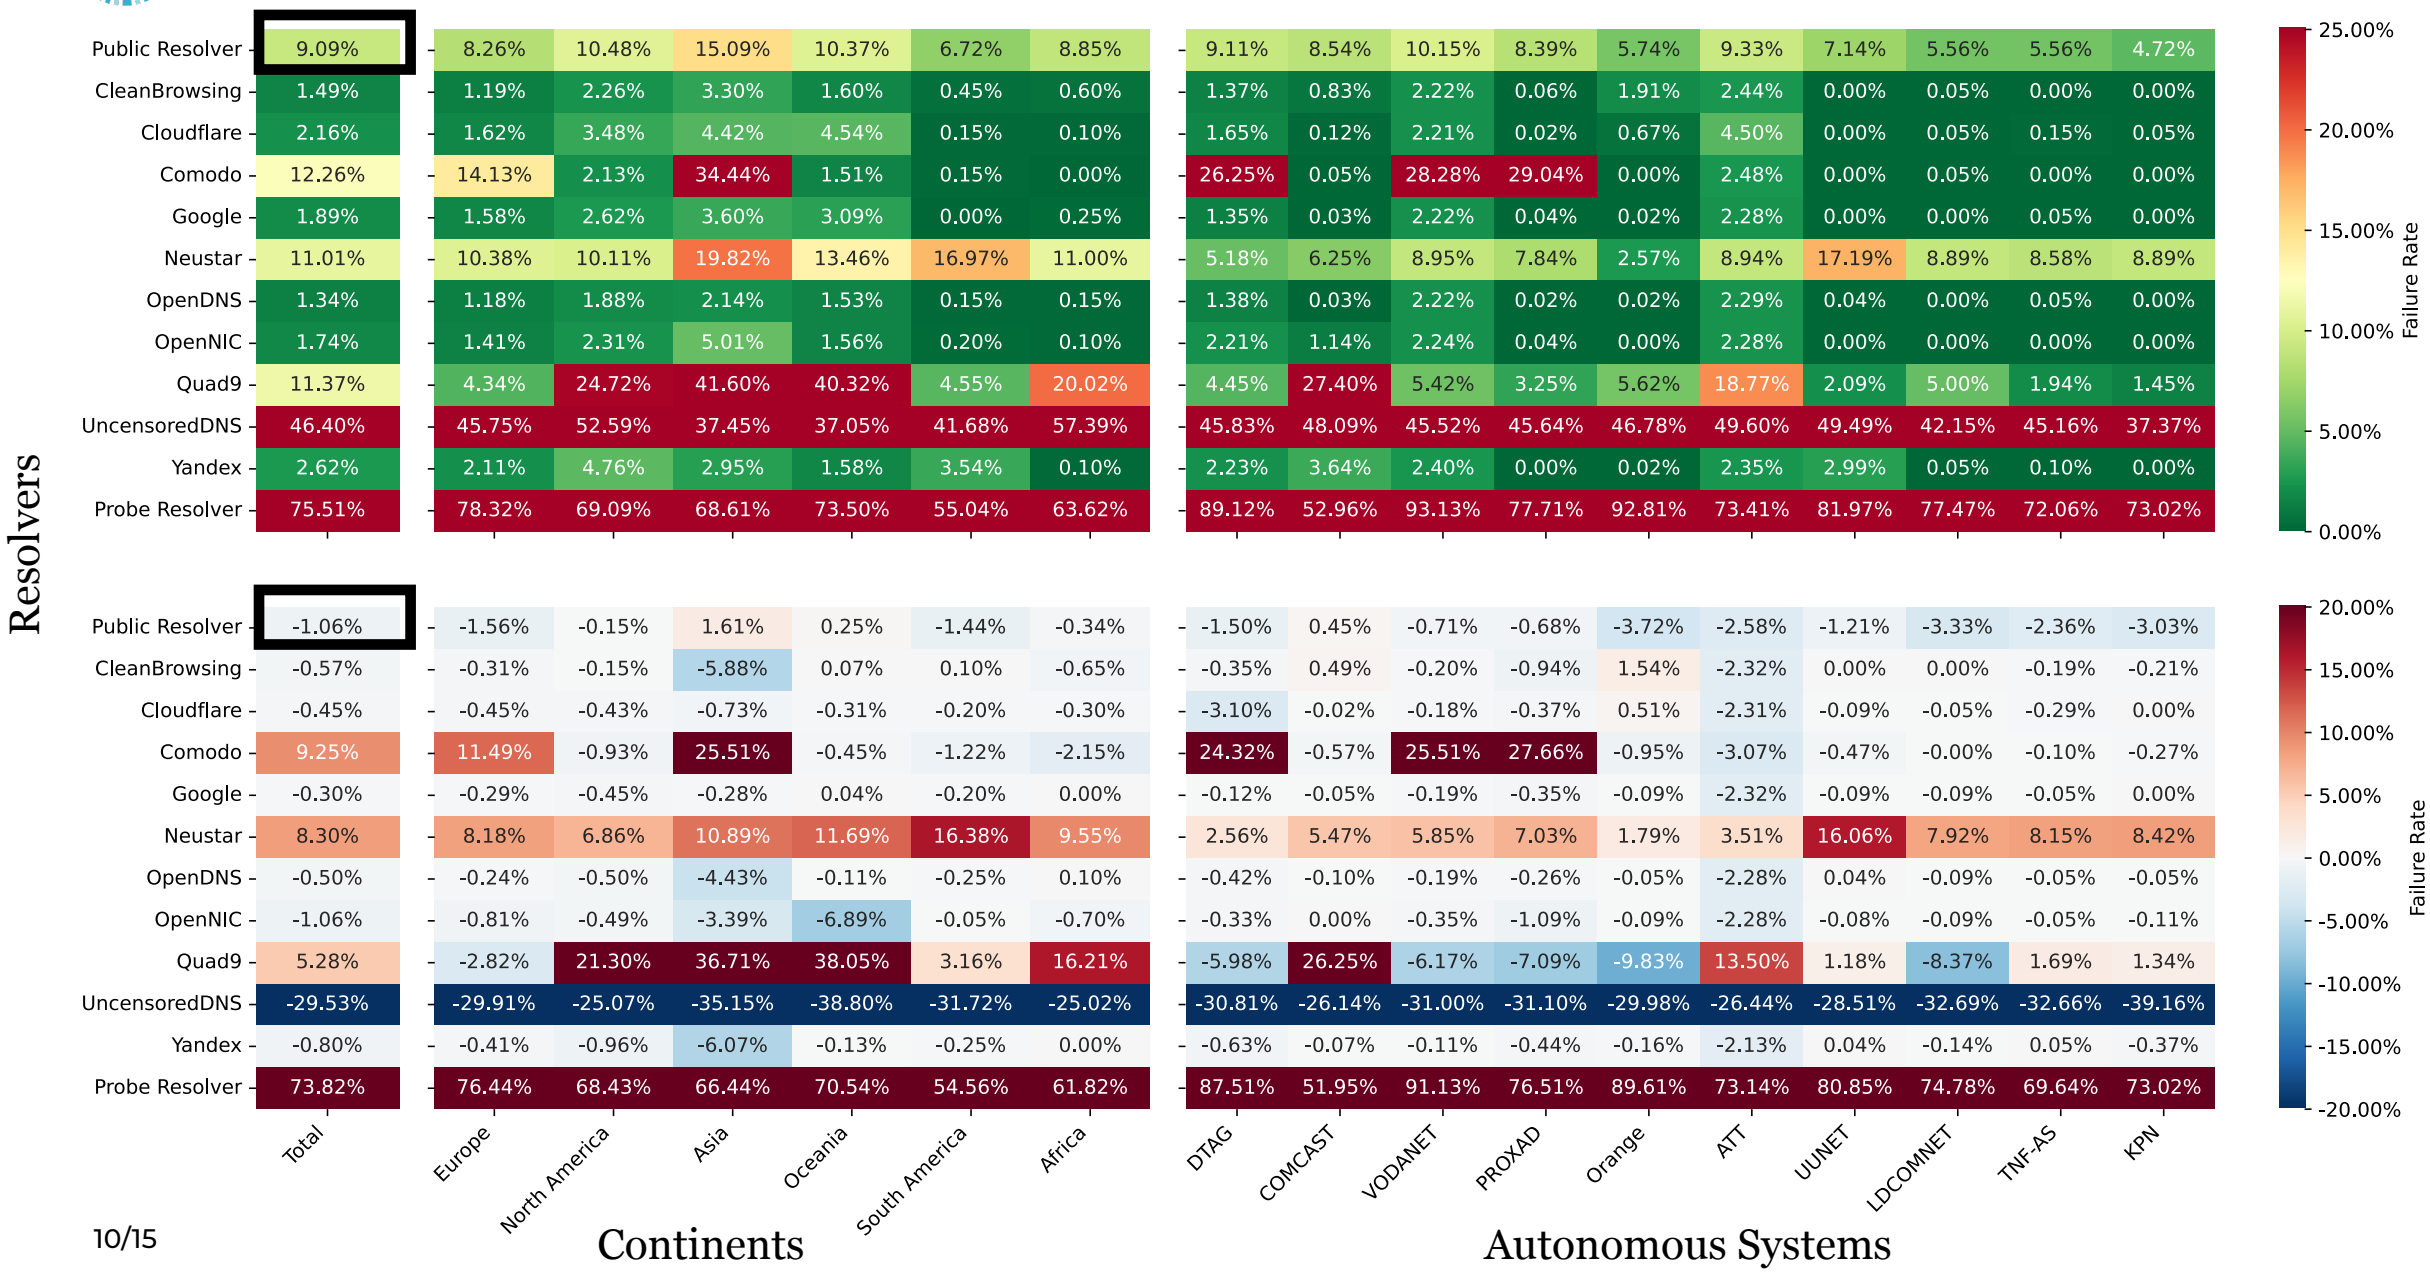

Findings (Evaluation from the Core - EDNS(0))

		512.0	1232.0	1400.0	1410.0	1452.0	4096.0	other
CleanBrowsing	*IPv4*	0.11%	**98.24%**	0.45%	0.05%	0.64%	0.36%	0.16%
	IPv6	0.01%	**99.47%**	0.21%	0.00%	0.07%	0.05%	0.19%
Cloudflare	*IPv4*	0.36%	0.65%	0.46%	0.04%	**98.04%**	0.30%	0.16%
	IPv6	0.01%	0.26%	0.21%	0.00%	**99.38%**	0.04%	0.10%
Comodo	*IPv4*	0.11%	0.70%	0.48%	0.05%	0.67%	**95.21%**	2.78%
	IPv6	-	-	-	-	-	-	-
Google	*IPv4*	0.22%	0.78%	**97.86%**	0.05%	0.64%	0.27%	0.19%
	IPv6	0.02%	0.26%	**99.41%**	0.00%	0.06%	0.14%	0.10%
Neustar	*IPv4*	0.04%	0.70%	0.48%	0.05%	0.63%	**97.45%**	0.66%
	IPv6	0.02%	0.31%	0.23%	0.00%	0.06%	**98.79%**	0.60%
OpenDNS	*IPv4*	0.08%	0.61%	0.53%	**97.68%**	0.59%	0.32%	0.19%
	IPv6	0.01%	0.26%	0.22%	**99.30%**	0.07%	0.04%	0.10%
OpenNIC	*IPv4*	0.06%	**98.29%**	0.45%	0.05%	0.59%	0.37%	0.18%
	IPv6	0.01%	**99.56%**	0.23%	0.00%	0.06%	0.05%	0.10%
Quad9	*IPv4*	0.07%	**98.05%**	0.51%	0.05%	0.70%	0.40%	0.21%
	IPv6	0.01%	**99.54%**	0.22%	0.00%	0.08%	0.04%	0.11%
UncensoredDNS	*IPv4*	2.68%	**93.15%**	1.14%	0.12%	1.49%	0.97%	0.45%
	IPv6	1.46%	**97.91%**	0.34%	0.00%	0.08%	0.07%	0.15%
Yandex	*IPv4*	0.03%	0.65%	0.56%	0.04%	0.69%	**92.86%**	5.16%
	IPv6	0.00%	0.26%	0.22%	0.00%	0.06%	**94.31%**	5.14%
Overall	*IPv4*	0.24%	39.74%	12.22%	11.83%	12.34%	22.78%	0.85%
	IPv6	0.13%	42.09%	11.70%	11.51%	11.49%	22.33%	0.75%

Takeaways:

- 3/10 resolvers show large buffer sizes which may lead to fragmentation attacks
- 4/10 resolvers use 1232B buffer size
- Resolvers exhibit preferred buffer sizes >90% of cases

Jayasree Sengupta | jayasree.sengupta@cispa.de | CISPA Helmholtz Center for Information Security

Findings (Evaluation from the Core - EDNS options)

		EDNS	Cookie	ECS
CleanBrowsing	*IPv4*	99.93%	0.22%	0.10%
	IPv6	99.91%	0.05%	0.04%
Cloudflare	*IPv4*	99.94%	0.32%	0.10%
	IPv6	100.00%	0.05%	0.05%
Comodo	*IPv4*	98.10%	0.33%	0.11%
	IPv6	-	-	-
Google	*IPv4*	99.93%	0.31%	**14.23%**
	IPv6	100.00%	0.16%	**12.53%**
Neustar	*IPv4*	99.93%	0.23%	0.10%
	IPv6	99.93%	0.05%	0.04%
OpenDNS	*IPv4*	99.94%	0.22%	0.10%
	IPv6	100.00%	0.05%	0.04%
OpenNIC	*IPv4*	99.93%	0.22%	0.11%
	IPv6	100.00%	0.05%	0.05%
Quad9	*IPv4*	99.93%	0.24%	0.13%
	IPv6	100.00%	0.06%	0.03%
UncensoredDNS	*IPv4*	99.84%	**94.62%**	0.24%
	IPv6	100.00%	**99.06%**	0.06%
Yandex	*IPv4*	99.93%	0.22%	0.11%
	IPv6	100.00%	0.05%	0.04%
Overall	*IPv4*	99.93%	4.80%	1.81%
	IPv6	99.98%	7.91%	1.49%

About:

- EDNS Client Subnet (ECS) allows clients to pass the network information through the chain of DNS queries from the DNS client to name servers
- The EDNS Cookie option is a lightweight security mechanism for DoUDP. Client and server exchange cookies of a minimum length of 64-bit allowing the communication parties to identify spoofed requests

Takeaways:

- All resolvers use EDNS (0) >99% of cases
- Small difference in usage rates between IPv4 vs IPv6
- UncensoredDNS uses the EDNS Cookie option in the majority while all other resolvers send cookies in <=0.33% of their requests
- Google mostly uses ECS. The other ones send Client Subnet information in <=0.24% of their requests

Key Takeaways

Failure rates of DoUDP > DoTCP over both IPv4 and IPv6 while evaluating from the edge and the core as well.

3/10 and 4/10 resolvers announce very large EDNS(0) buffer sizes (4096B) from the Core and Edge respectively, which may potentially cause fragmentation.

The **resolvers exhibit one preferred buffer size** for >95% of the cases both from the Edge and the Core.

DNS-over-QUIC (DoQ) (RFC 9250) solves fragmentation by means of the QUIC protocol (RFC 9000), thereby supporting increased DNS message sizes.

-paper-

Jayasree Sengupta | jayasree.sengupta@cispa.de | CISPA Helmholtz Center for Information Security