CISPA
HELMHOLTZ CENTER FOR
INFORMATION SECURITY

Evaluating DNS Resiliency with Truncation, Fragmentation and DoTCP Fallback

Pratyush Dikshit, Mike Kosek^, Nils Faulhaber^, Jayasree Sengupta*, and Vaibhav Bajpai**

**CISPA Helmholtz Center for Information Security, Germany*
^Technical University of Munich, Germany

IFIP Networking Conference 2023 | Volkswagen Foundation | June 2023

—Paper—

Preliminaries

Truncation

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

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”*

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

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

Research Questions and Findings

To investigate how DNS service providers around the world have adopted the DNS Flag Day 2020 recommendations in their software

I. How resilient is DoTCP over DoUDP (for IPv4 and IPv6) while comparing the failure rate?
II. What is the scale of usage and performance of DoTCP?
III. Which buffer sizes are currently used in DNS traffic around the globe (EDNS Configuration)?

Findings:

I. The observed resilience of DoTCP over IPv4 is higher than over IPv6
II. 6/10 Public resolvers announced either very small (512B) or very large (4096B) EDNS(0) Buffer Sizes
III. Several Public DNS resolvers still lack adoption to the DNS Flag Day 2020 recommendations

Methodology

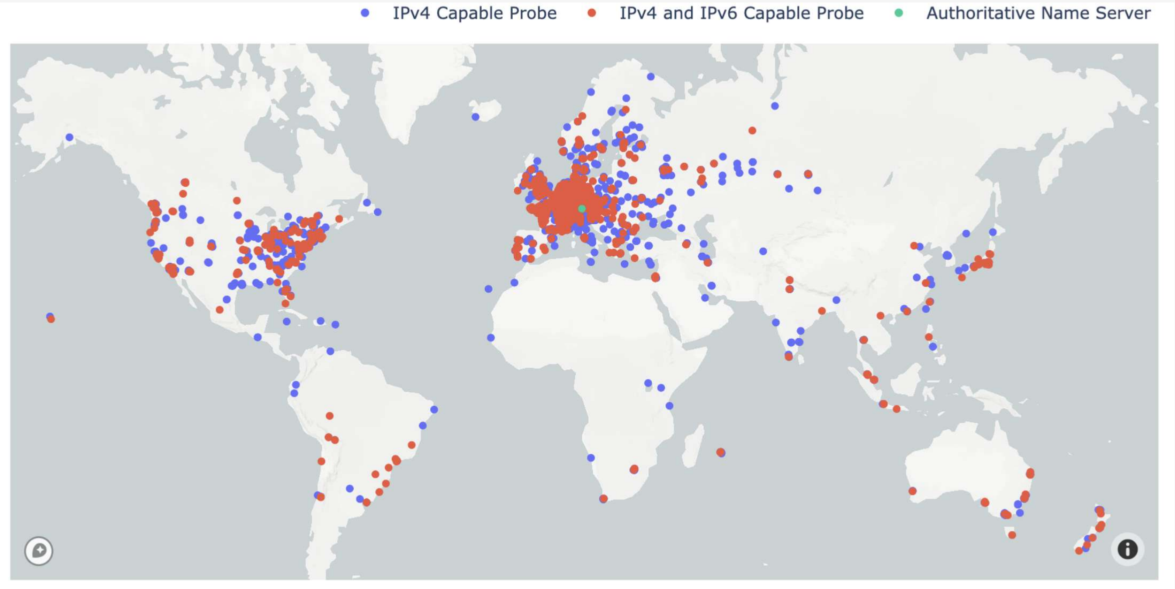

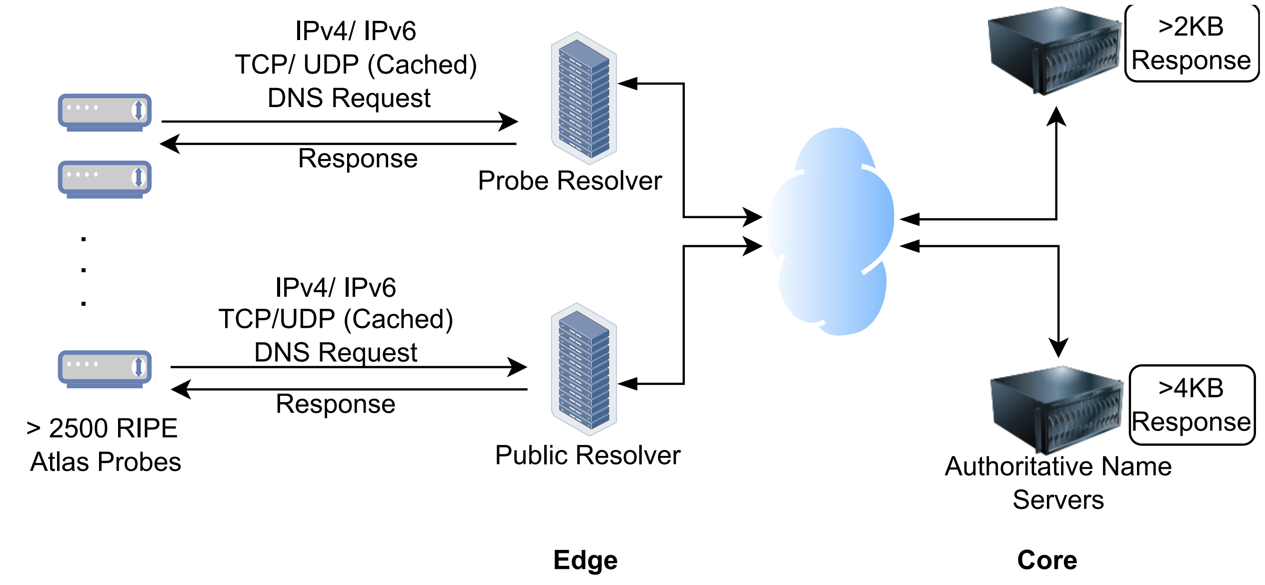

- 2527 globally distributed RIPE Atlas probes
- 88% of the probes are located in North America and Europe.
- RIPE Altas probes are hardware devices that volunteers can host by connecting them to their local router via Ethernet

- 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

Methodology (Evaluation from the Edge)

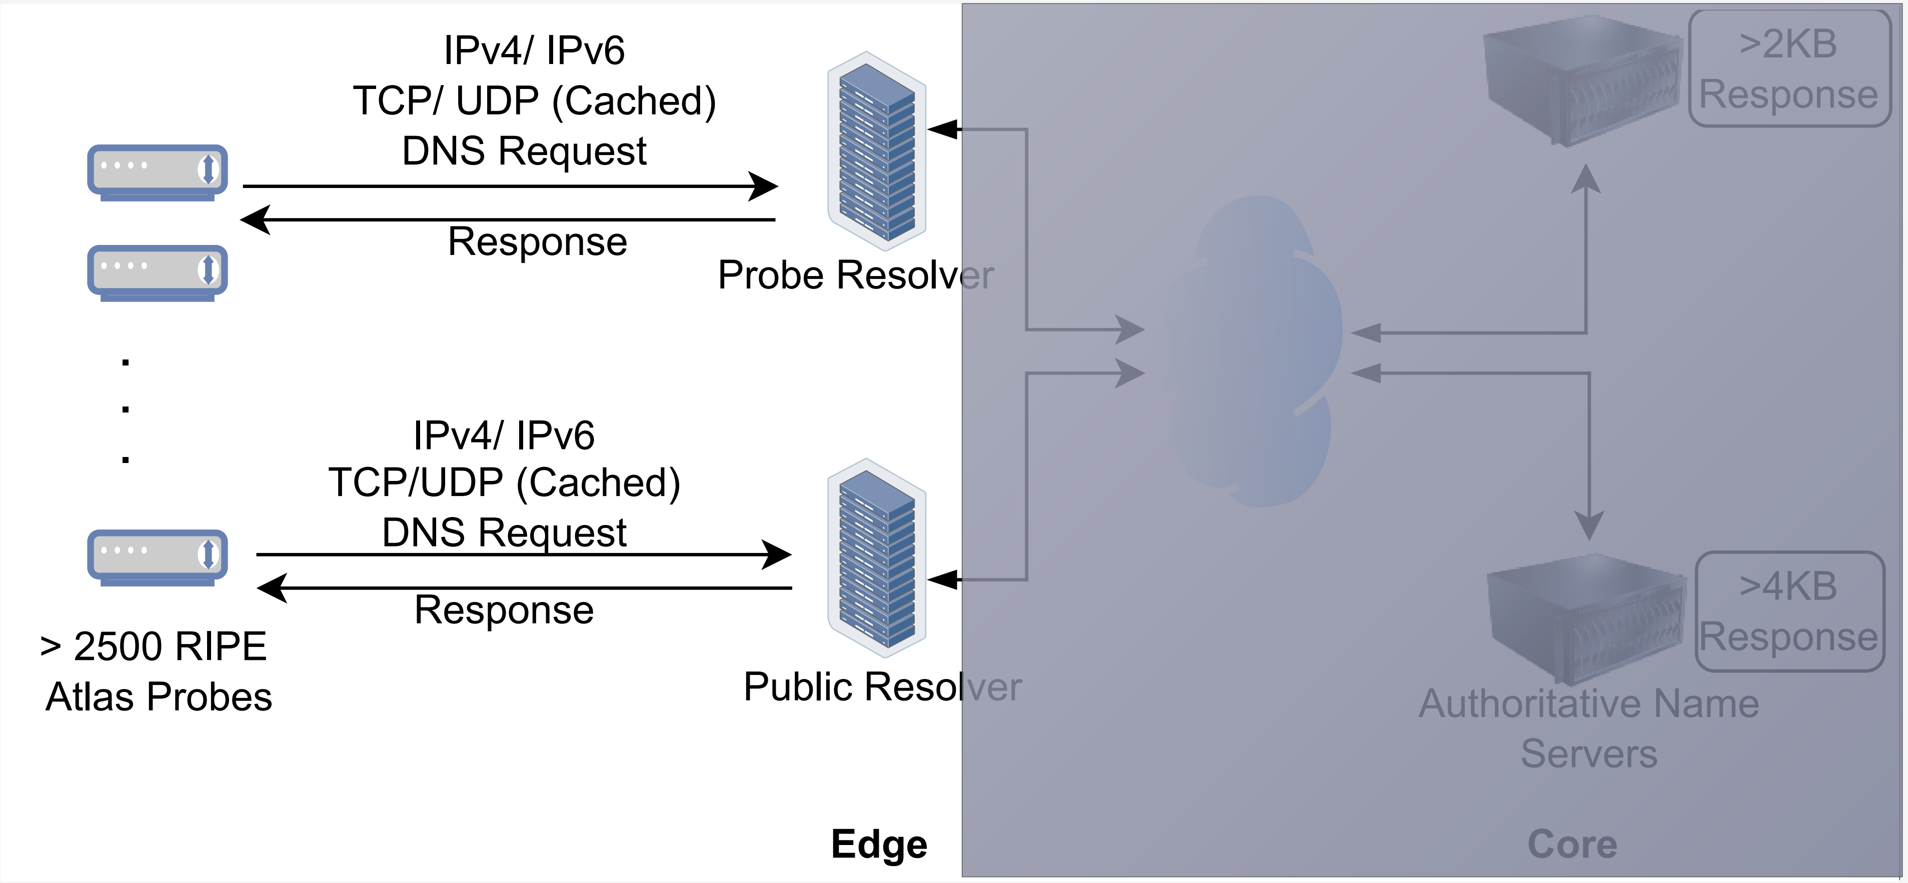

- **One week, 10 blocks every day, 10 requests per block per resolver**

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

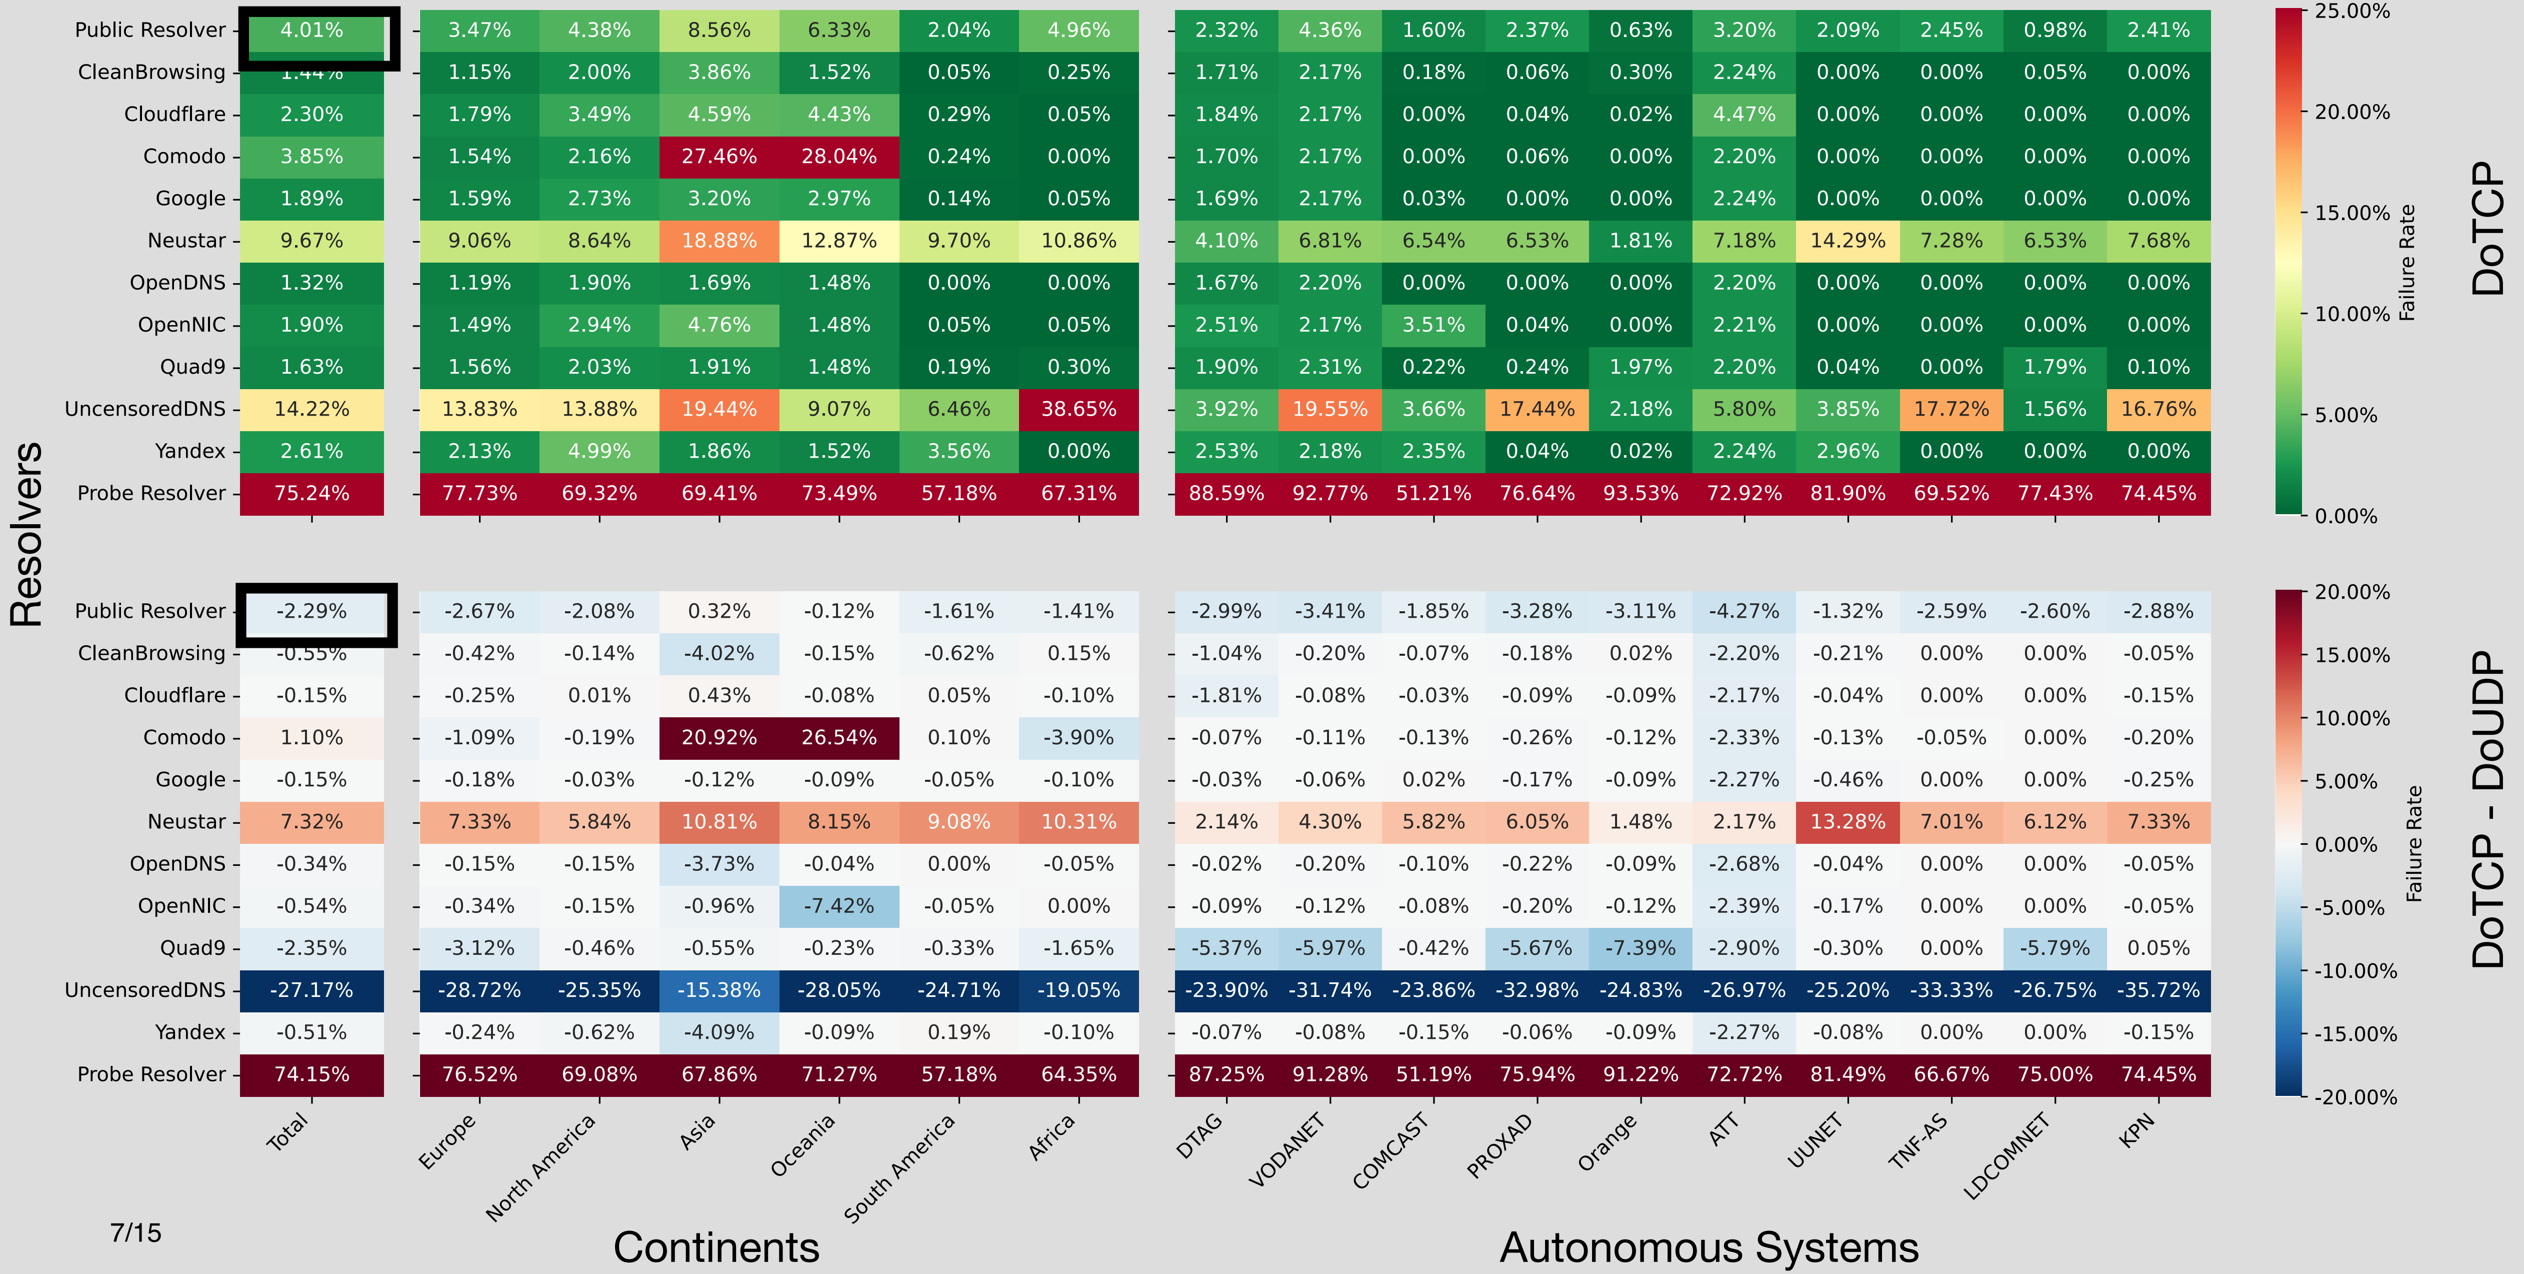

DoTCP

Resolvers	Total	Europe	North America	Asia	Oceania	South America	Africa	DTAG	VODANET	COMCAST	PROXAD	Orange	ATT	UUNET	TNF-AS	LDCOMNET	KPN
Public Resolver	4.01%	3.47%	4.38%	8.56%	6.33%	2.04%	4.96%	2.32%	4.36%	1.60%	2.37%	0.63%	3.20%	2.09%	2.45%	0.98%	2.41%
CleanBrowsing	1.44%	1.15%	2.00%	3.86%	1.52%	0.05%	0.25%	1.71%	2.17%	0.18%	0.06%	0.30%	2.24%	0.00%	0.00%	0.05%	0.00%
Cloudflare	2.30%	1.79%	3.49%	4.59%	4.43%	0.29%	0.05%	1.84%	2.17%	0.00%	0.04%	0.02%	4.47%	0.00%	0.00%	0.00%	0.00%
Comodo	3.85%	1.54%	2.16%	27.46%	28.04%	0.24%	0.00%	1.70%	2.17%	0.00%	0.06%	0.00%	2.20%	0.00%	0.00%	0.00%	0.00%
Google	1.89%	1.59%	2.73%	3.20%	2.97%	0.14%	0.05%	1.69%	2.17%	0.03%	0.00%	0.00%	2.24%	0.00%	0.00%	0.00%	0.00%
Neustar	9.67%	9.06%	8.64%	18.88%	12.87%	9.70%	10.86%	4.10%	6.81%	6.54%	6.53%	1.81%	7.18%	14.29%	7.28%	6.53%	7.68%
OpenDNS	1.32%	1.19%	1.90%	1.69%	1.48%	0.00%	0.00%	1.67%	2.20%	0.00%	0.00%	0.00%	2.20%	0.00%	0.00%	0.00%	0.00%
OpenNIC	1.90%	1.49%	2.94%	4.76%	1.48%	0.05%	0.05%	2.51%	2.17%	3.51%	0.04%	0.00%	2.21%	0.00%	0.00%	0.00%	0.00%
Quad9	1.63%	1.56%	2.03%	1.91%	1.48%	0.19%	0.30%	1.90%	2.31%	0.22%	0.24%	1.97%	2.20%	0.04%	0.00%	1.79%	0.10%
UncensoredDNS	14.22%	13.83%	13.88%	19.44%	9.07%	6.46%	38.65%	3.92%	19.55%	3.66%	17.44%	2.18%	5.80%	3.85%	17.72%	1.56%	16.76%
Yandex	2.61%	2.13%	4.99%	1.86%	1.52%	3.56%	0.00%	2.53%	2.18%	2.35%	0.04%	0.02%	2.24%	2.96%	0.00%	0.00%	0.00%
Probe Resolver	75.24%	77.73%	69.32%	69.41%	73.49%	57.18%	67.31%	88.59%	92.77%	51.21%	76.64%	93.53%	72.92%	81.90%	69.52%	77.43%	74.45%

DoTCP - DoUDP

Resolvers	Total	Europe	North America	Asia	Oceania	South America	Africa	DTAG	VODANET	COMCAST	PROXAD	Orange	ATT	UUNET	TNF-AS	LDCOMNET	KPN
Public Resolver	-2.29%	-2.67%	-2.08%	0.32%	-0.12%	-1.61%	-1.41%	-2.99%	-3.41%	-1.85%	-3.28%	-3.11%	-4.27%	-1.32%	-2.59%	-2.60%	-2.88%
CleanBrowsing	-0.55%	-0.42%	-0.14%	-4.02%	-0.15%	-0.62%	0.15%	-1.04%	-0.20%	-0.07%	-0.18%	0.02%	-2.20%	-0.21%	0.00%	0.00%	-0.05%
Cloudflare	-0.15%	-0.25%	0.01%	0.43%	-0.08%	0.05%	-0.10%	-1.81%	-0.08%	-0.03%	-0.09%	-0.09%	-2.17%	-0.04%	0.00%	0.00%	-0.15%
Comodo	1.10%	-1.09%	-0.19%	20.92%	26.54%	0.10%	-3.90%	-0.07%	-0.11%	-0.13%	-0.26%	-0.12%	-2.33%	-0.13%	-0.05%	0.00%	-0.20%
Google	-0.15%	-0.18%	-0.03%	-0.12%	-0.09%	-0.05%	-0.10%	-0.03%	-0.06%	0.02%	-0.17%	-0.09%	-2.27%	-0.46%	0.00%	0.00%	-0.25%
Neustar	7.32%	7.33%	5.84%	10.81%	8.15%	9.08%	10.31%	2.14%	4.30%	5.82%	6.05%	1.48%	2.17%	13.28%	7.01%	6.12%	7.33%
OpenDNS	-0.34%	-0.15%	-0.15%	-3.73%	-0.04%	0.00%	-0.05%	-0.02%	-0.20%	-0.10%	-0.22%	-0.09%	-2.68%	-0.04%	0.00%	0.00%	-0.05%
OpenNIC	-0.54%	-0.34%	-0.15%	-0.96%	-7.42%	-0.05%	0.00%	-0.09%	-0.12%	-0.08%	-0.20%	-0.12%	-2.39%	-0.17%	0.00%	0.00%	-0.05%
Quad9	-2.35%	-3.12%	-0.46%	-0.55%	-0.23%	-0.33%	-1.65%	-5.37%	-5.97%	-0.42%	-5.67%	-7.39%	-2.90%	-0.30%	0.00%	-5.79%	0.05%
UncensoredDNS	-27.17%	-28.72%	-25.35%	-15.38%	-28.05%	-24.71%	-19.05%	-23.90%	-31.74%	-23.86%	-32.98%	-24.83%	-26.97%	-25.20%	-33.33%	-26.75%	-35.72%
Yandex	-0.51%	-0.24%	-0.62%	-4.09%	-0.09%	0.19%	-0.10%	-0.07%	-0.08%	-0.15%	-0.06%	-0.09%	-2.27%	-0.08%	0.00%	0.00%	-0.15%
Probe Resolver	74.15%	76.52%	69.08%	67.86%	71.27%	57.18%	64.35%	87.25%	91.28%	51.19%	75.94%	91.22%	72.72%	81.49%	66.67%	75.00%	74.45%

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

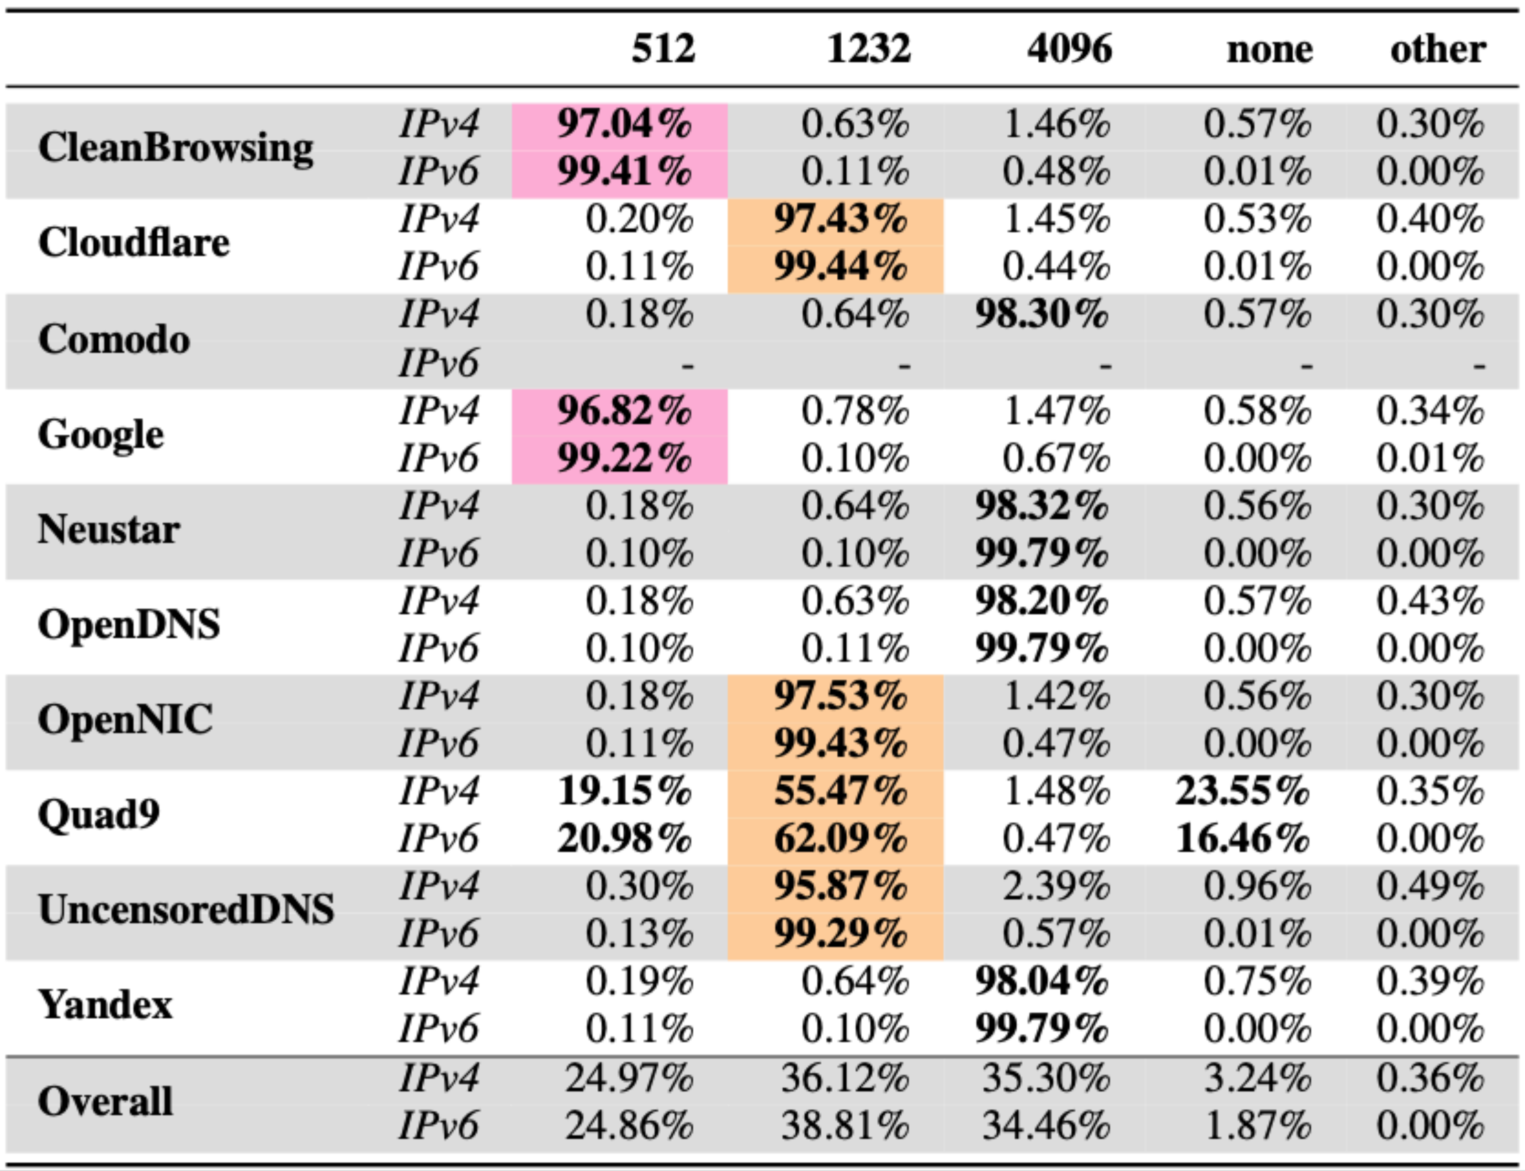

		512	1232	4096	none	other
CleanBrowsing	*IPv4*	**97.04%**	0.63%	1.46%	0.57%	0.30%
	IPv6	**99.41%**	0.11%	0.48%	0.01%	0.00%
Cloudflare	*IPv4*	0.20%	**97.43%**	1.45%	0.53%	0.40%
	IPv6	0.11%	**99.44%**	0.44%	0.01%	0.00%
Comodo	*IPv4*	0.18%	0.64%	**98.30%**	0.57%	0.30%
	IPv6	-	-	-	-	-
Google	*IPv4*	**96.82%**	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
- 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

Methodology (Evaluation from the Core)

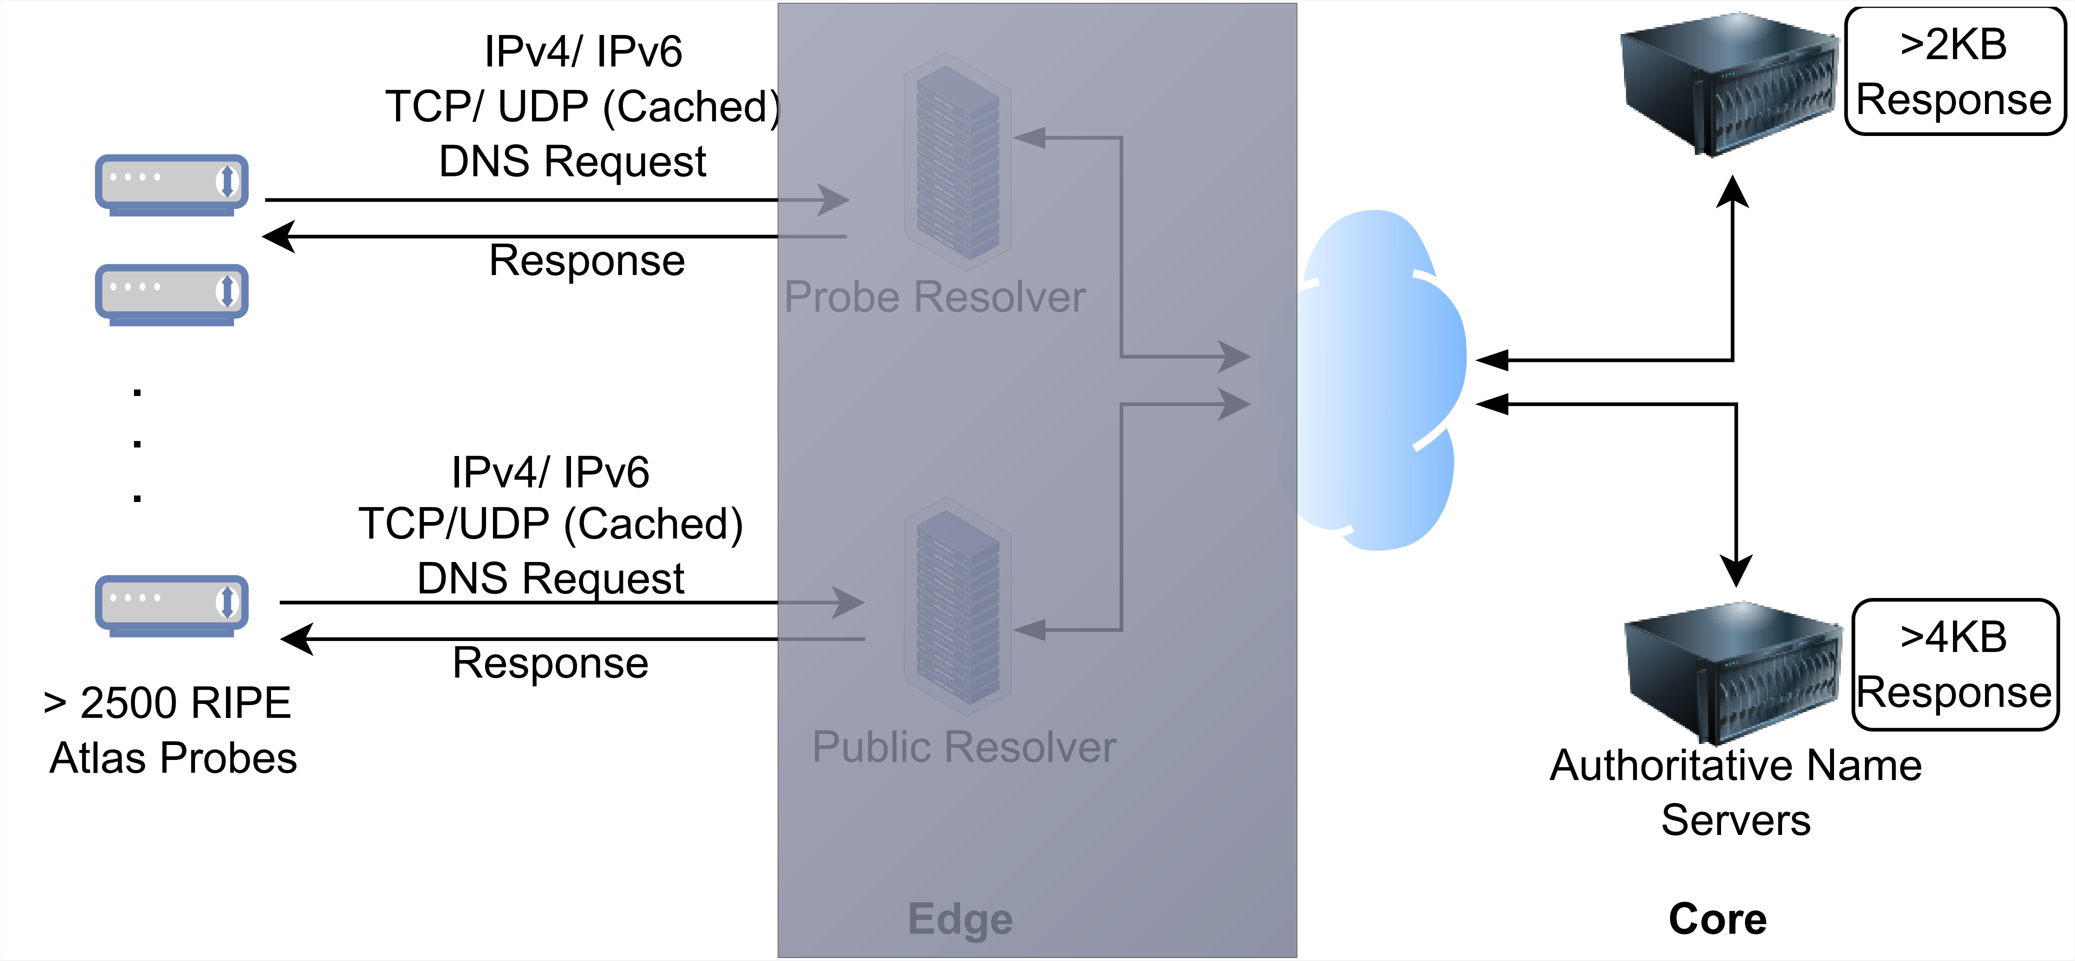

- **One week, 10 blocks every day, 10 requests per block per resolver**

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

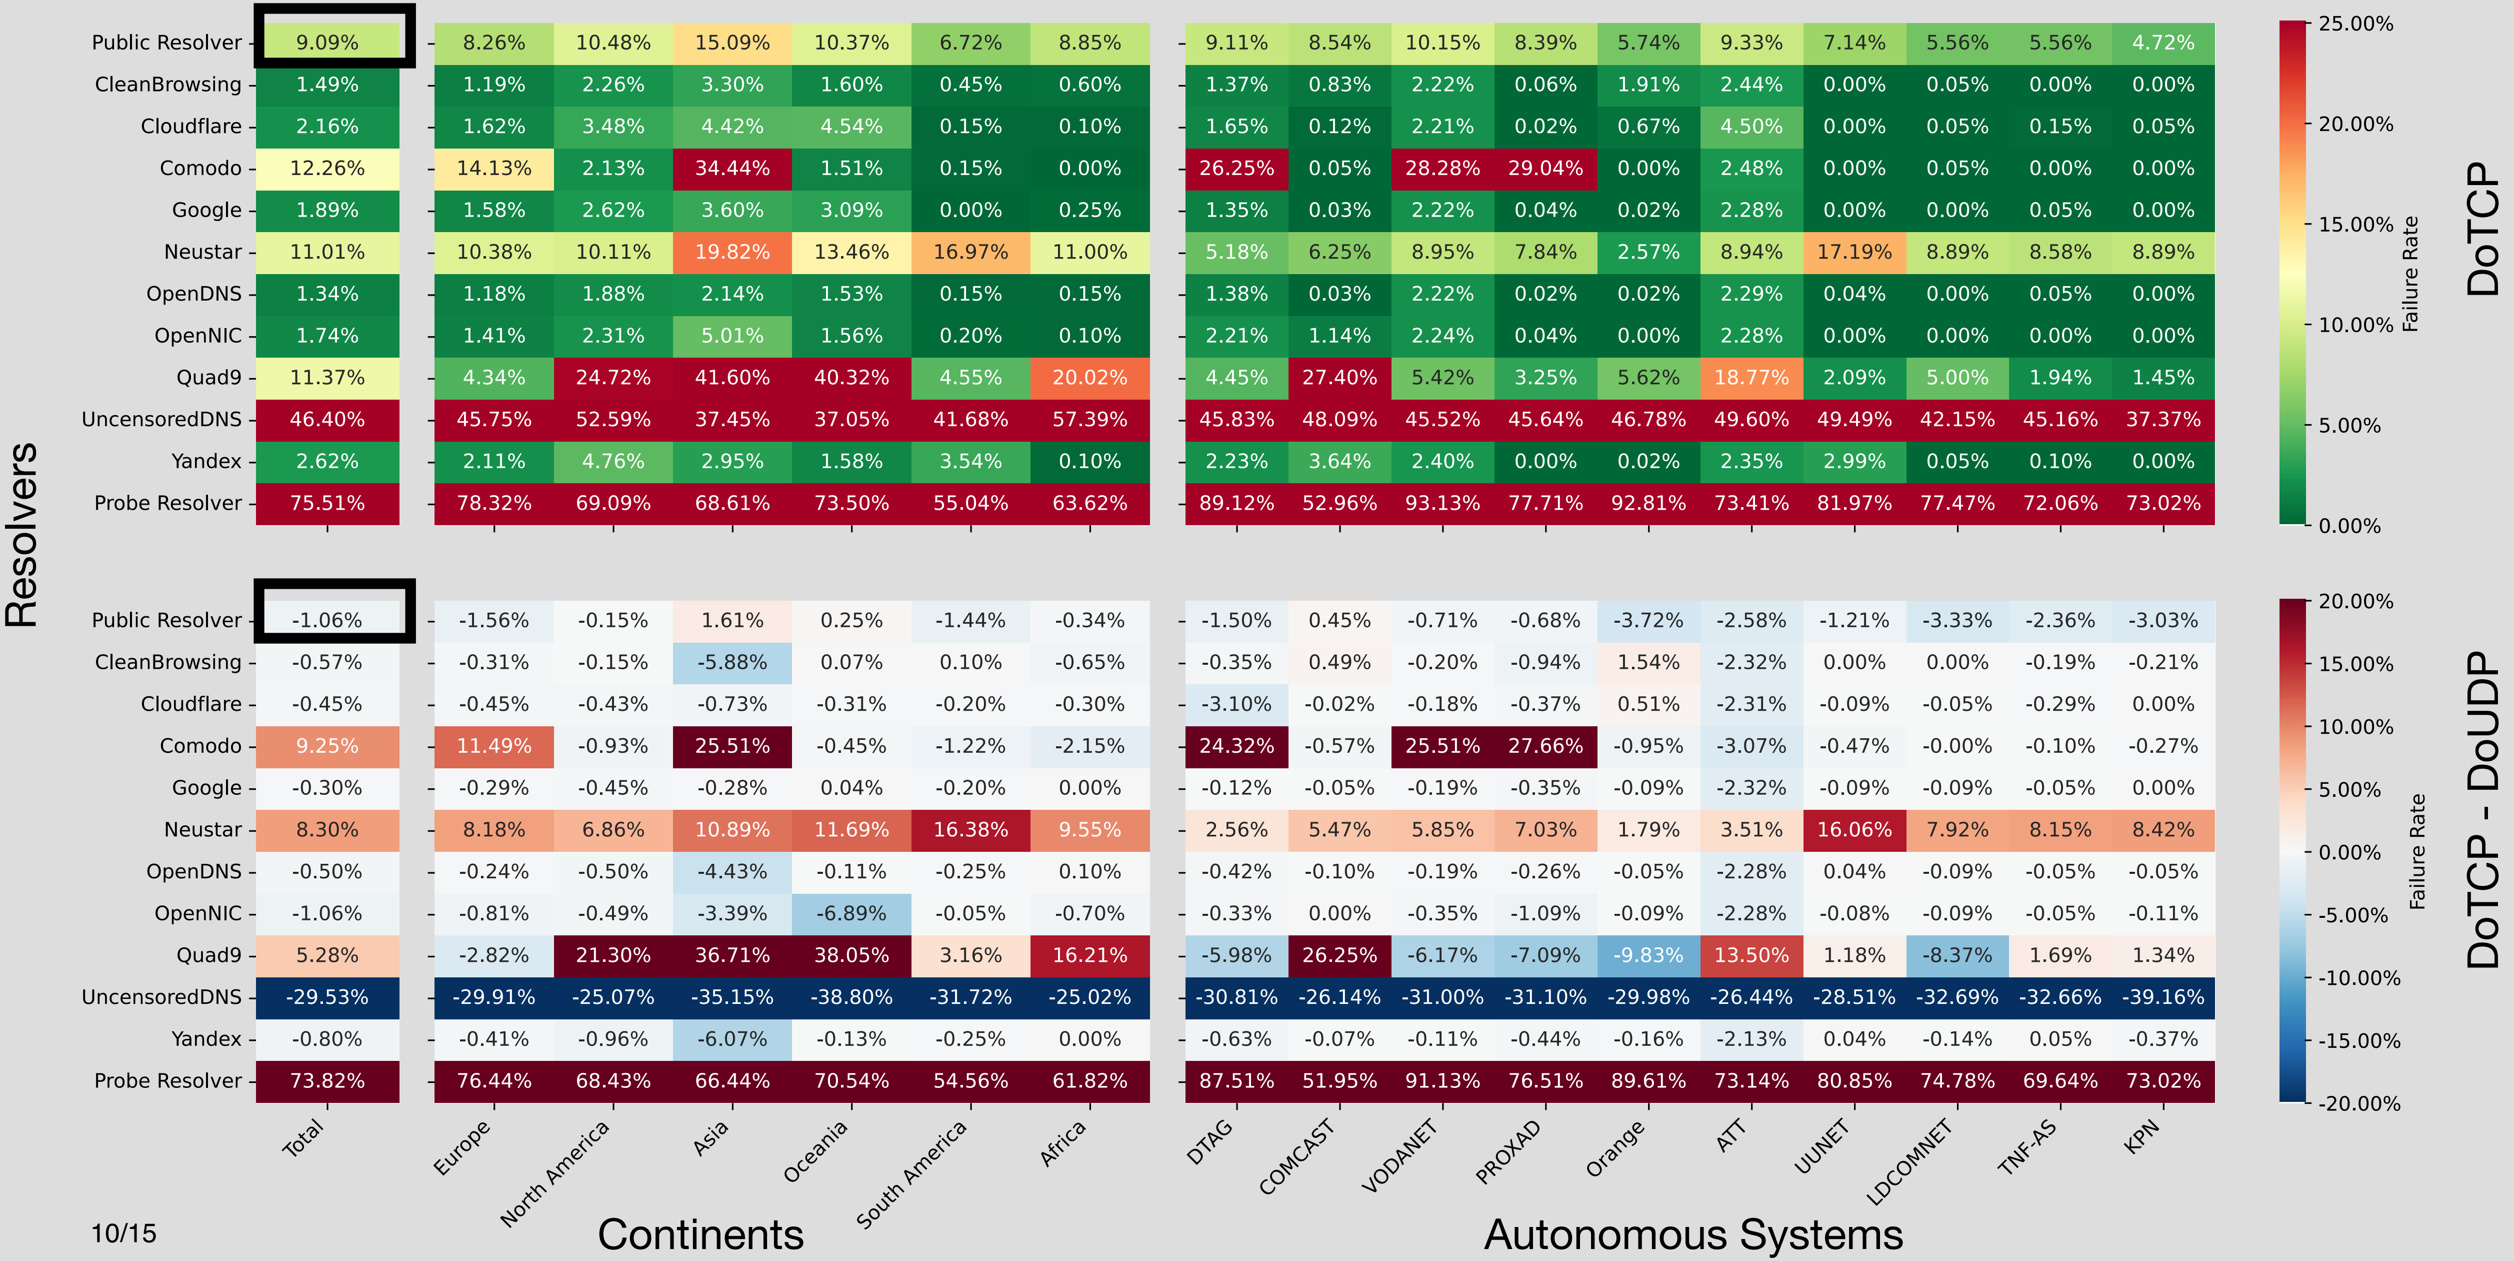

	Total	Europe	North America	Asia	Oceania	South America	Africa	DTAG	COMCAST	VODANET	PROXAD	Orange	ATT	UUNET	LDCOMNET	TNF-AS	KPN
Public Resolver	9.09%	8.26%	10.48%	15.09%	10.37%	6.72%	8.85%	9.11%	8.54%	10.15%	8.39%	5.74%	9.33%	7.14%	5.56%	5.56%	4.72%
CleanBrowsing	1.49%	1.19%	2.26%	3.30%	1.60%	0.45%	0.60%	1.37%	0.83%	2.22%	0.06%	1.91%	2.44%	0.00%	0.05%	0.00%	0.00%
Cloudflare	2.16%	1.62%	3.48%	4.42%	4.54%	0.15%	0.10%	1.65%	0.12%	2.21%	0.02%	0.67%	4.50%	0.00%	0.05%	0.15%	0.05%
Comodo	12.26%	14.13%	2.13%	34.44%	1.51%	0.15%	0.00%	26.25%	0.05%	28.28%	29.04%	0.00%	2.48%	0.00%	0.05%	0.00%	0.00%
Google	1.89%	1.58%	2.62%	3.60%	3.09%	0.00%	0.25%	1.35%	0.03%	2.22%	0.04%	0.02%	2.28%	0.00%	0.00%	0.05%	0.00%
Neustar	11.01%	10.38%	10.11%	19.82%	13.46%	16.97%	11.00%	5.18%	6.25%	8.95%	7.84%	2.57%	8.94%	17.19%	8.89%	8.58%	8.89%
OpenDNS	1.34%	1.18%	1.88%	2.14%	1.53%	0.15%	0.15%	1.38%	0.03%	2.22%	0.02%	0.02%	2.29%	0.04%	0.00%	0.05%	0.00%
OpenNIC	1.74%	1.41%	2.31%	5.01%	1.56%	0.20%	0.10%	2.21%	1.14%	2.24%	0.04%	0.00%	2.28%	0.00%	0.00%	0.00%	0.00%
Quad9	11.37%	4.34%	24.72%	41.60%	40.32%	4.55%	20.02%	4.45%	27.40%	5.42%	3.25%	5.62%	18.77%	2.09%	5.00%	1.94%	1.45%
UncensoredDNS	46.40%	45.75%	52.59%	37.45%	37.05%	41.68%	57.39%	45.83%	48.09%	45.52%	45.64%	46.78%	49.60%	49.49%	42.15%	45.16%	37.37%
Yandex	2.62%	2.11%	4.76%	2.95%	1.58%	3.54%	0.10%	2.23%	3.64%	2.40%	0.00%	0.02%	2.35%	2.99%	0.05%	0.10%	0.00%
Probe Resolver	75.51%	78.32%	69.09%	68.61%	73.50%	55.04%	63.62%	89.12%	52.96%	93.13%	77.71%	92.81%	73.41%	81.97%	77.47%	72.06%	73.02%

	Total	Europe	North America	Asia	Oceania	South America	Africa	DTAG	COMCAST	VODANET	PROXAD	Orange	ATT	UUNET	LDCOMNET	TNF-AS	KPN
Public Resolver	-1.06%	-1.56%	-0.15%	1.61%	0.25%	-1.44%	-0.34%	-1.50%	0.45%	-0.71%	-0.68%	-3.72%	-2.58%	-1.21%	-3.33%	-2.36%	-3.03%
CleanBrowsing	-0.57%	-0.31%	-0.15%	-5.88%	0.07%	0.10%	-0.65%	-0.35%	0.49%	-0.20%	-0.94%	1.54%	-2.32%	0.00%	0.00%	-0.19%	-0.21%
Cloudflare	-0.45%	-0.45%	-0.43%	-0.73%	-0.31%	-0.20%	-0.30%	-3.10%	-0.02%	-0.18%	-0.37%	0.51%	-2.31%	-0.09%	-0.05%	-0.29%	0.00%
Comodo	9.25%	11.49%	-0.93%	25.51%	-0.45%	-1.22%	-2.15%	24.32%	-0.57%	25.51%	27.66%	-0.95%	-3.07%	-0.47%	-0.00%	-0.10%	-0.27%
Google	-0.30%	-0.29%	-0.45%	-0.28%	0.04%	-0.20%	0.00%	-0.12%	-0.05%	-0.19%	-0.35%	-0.09%	-2.32%	-0.09%	-0.09%	-0.05%	0.00%
Neustar	8.30%	8.18%	6.86%	10.89%	11.69%	16.38%	9.55%	2.56%	5.47%	5.85%	7.03%	1.79%	3.51%	16.06%	7.92%	8.15%	8.42%
OpenDNS	-0.50%	-0.24%	-0.50%	-4.43%	-0.11%	-0.25%	0.10%	-0.42%	-0.10%	-0.19%	-0.26%	-0.05%	-2.28%	0.04%	-0.09%	-0.05%	-0.05%
OpenNIC	-1.06%	-0.81%	-0.49%	-3.39%	-6.89%	-0.05%	-0.70%	-0.33%	0.00%	-0.35%	-1.09%	-0.09%	-2.28%	-0.08%	-0.09%	-0.05%	-0.11%
Quad9	5.28%	-2.82%	21.30%	36.71%	38.05%	3.16%	16.21%	-5.98%	26.25%	-6.17%	-7.09%	-9.83%	13.50%	1.18%	-8.37%	1.69%	1.34%
UncensoredDNS	-29.53%	-29.91%	-25.07%	-35.15%	-38.80%	-31.72%	-25.02%	-30.81%	-26.14%	-31.00%	-31.10%	-29.98%	-26.44%	-28.51%	-32.69%	-32.66%	-39.16%
Yandex	-0.80%	-0.41%	-0.96%	-6.07%	-0.13%	-0.25%	0.00%	-0.63%	-0.07%	-0.11%	-0.44%	-0.16%	-2.13%	0.04%	-0.14%	0.05%	-0.37%
Probe Resolver	73.82%	76.44%	68.43%	66.44%	70.54%	54.56%	61.82%	87.51%	51.95%	91.13%	76.51%	89.61%	73.14%	80.85%	74.78%	69.64%	73.02%

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

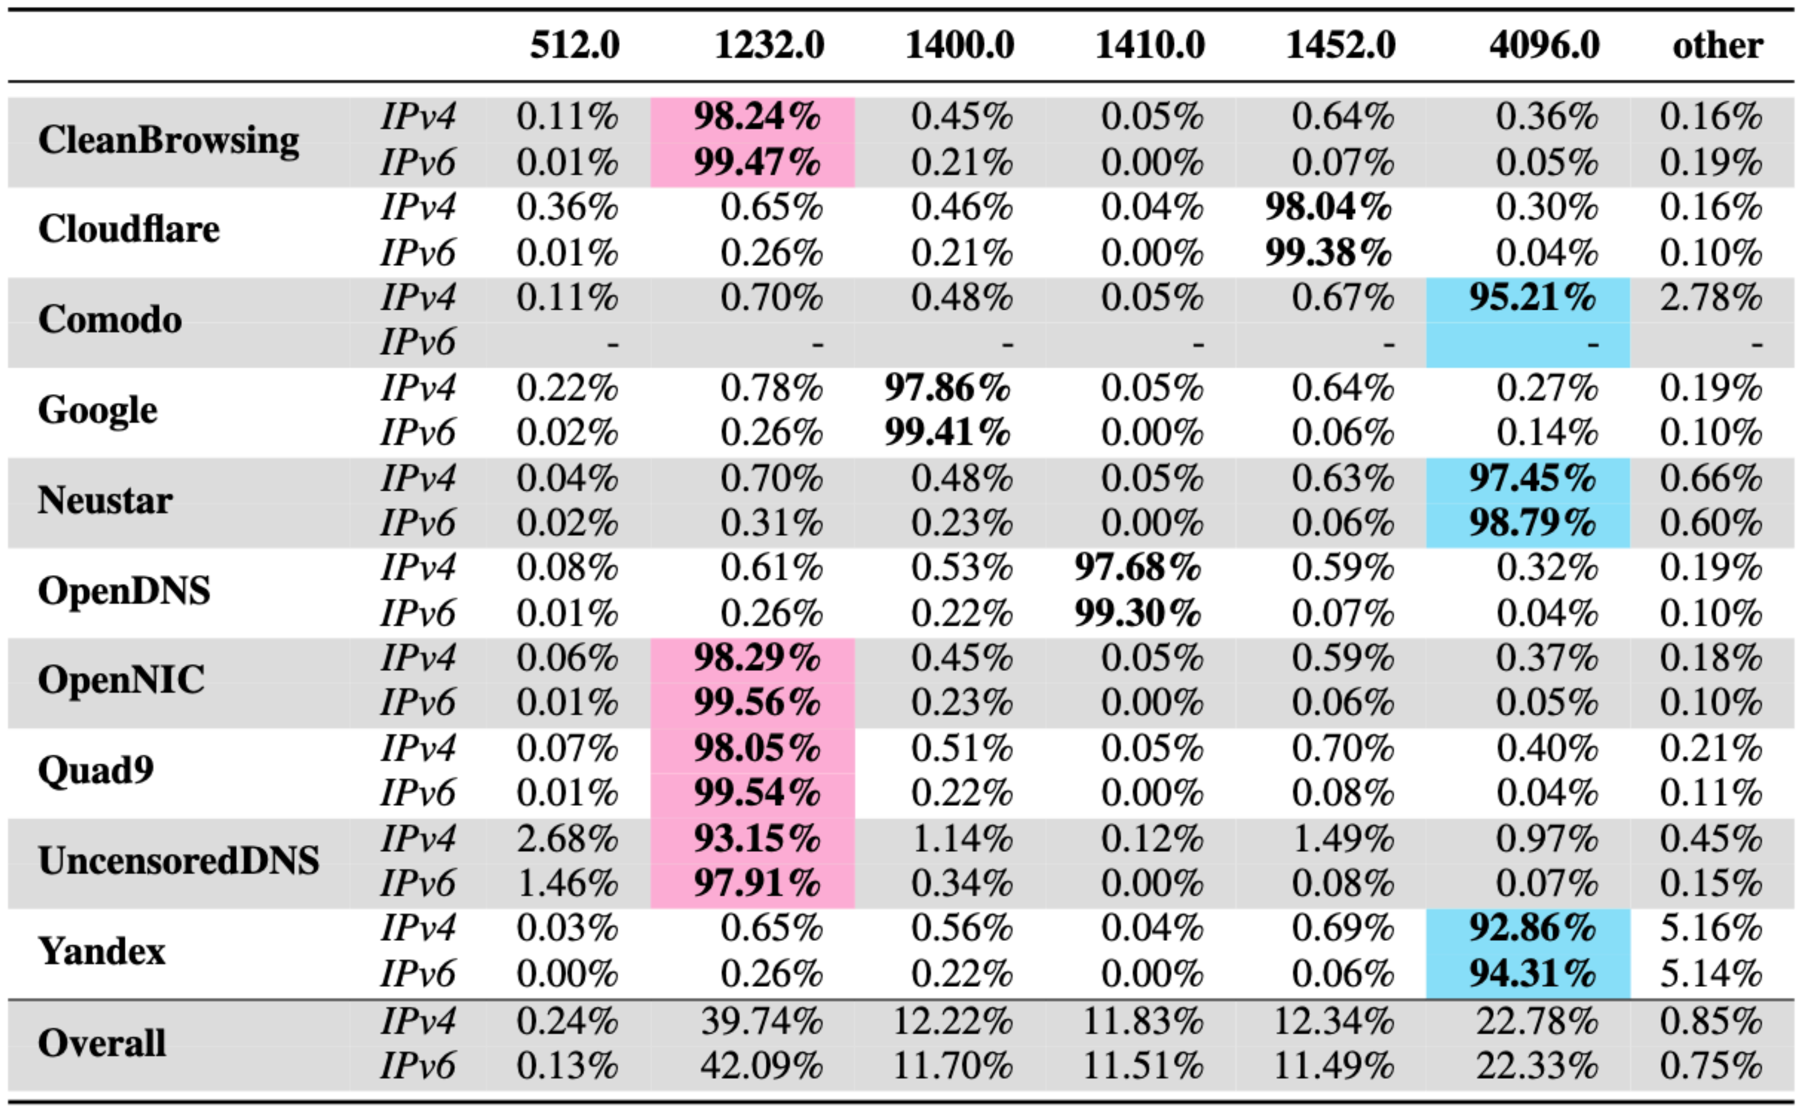

		512.0	1232.0	1400.0	1410.0	1452.0	4096.0	other
CleanBrowsing	*IPv4*	0.11%	**98.24%**	0.45%	0.05%	0.64%	0.36%	0.16%
	IPv6	0.01%	**99.47%**	0.21%	0.00%	0.07%	0.05%	0.19%
Cloudflare	*IPv4*	0.36%	0.65%	0.46%	0.04%	**98.04%**	0.30%	0.16%
	IPv6	0.01%	0.26%	0.21%	0.00%	**99.38%**	0.04%	0.10%
Comodo	*IPv4*	0.11%	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.
- CleanBrowsing, OpenNIC, Quad9, and UncensoredDNS conforming to the DNS Flag Day 2020 recommendations

Findings (Evaluation from the Core - EDNS options)

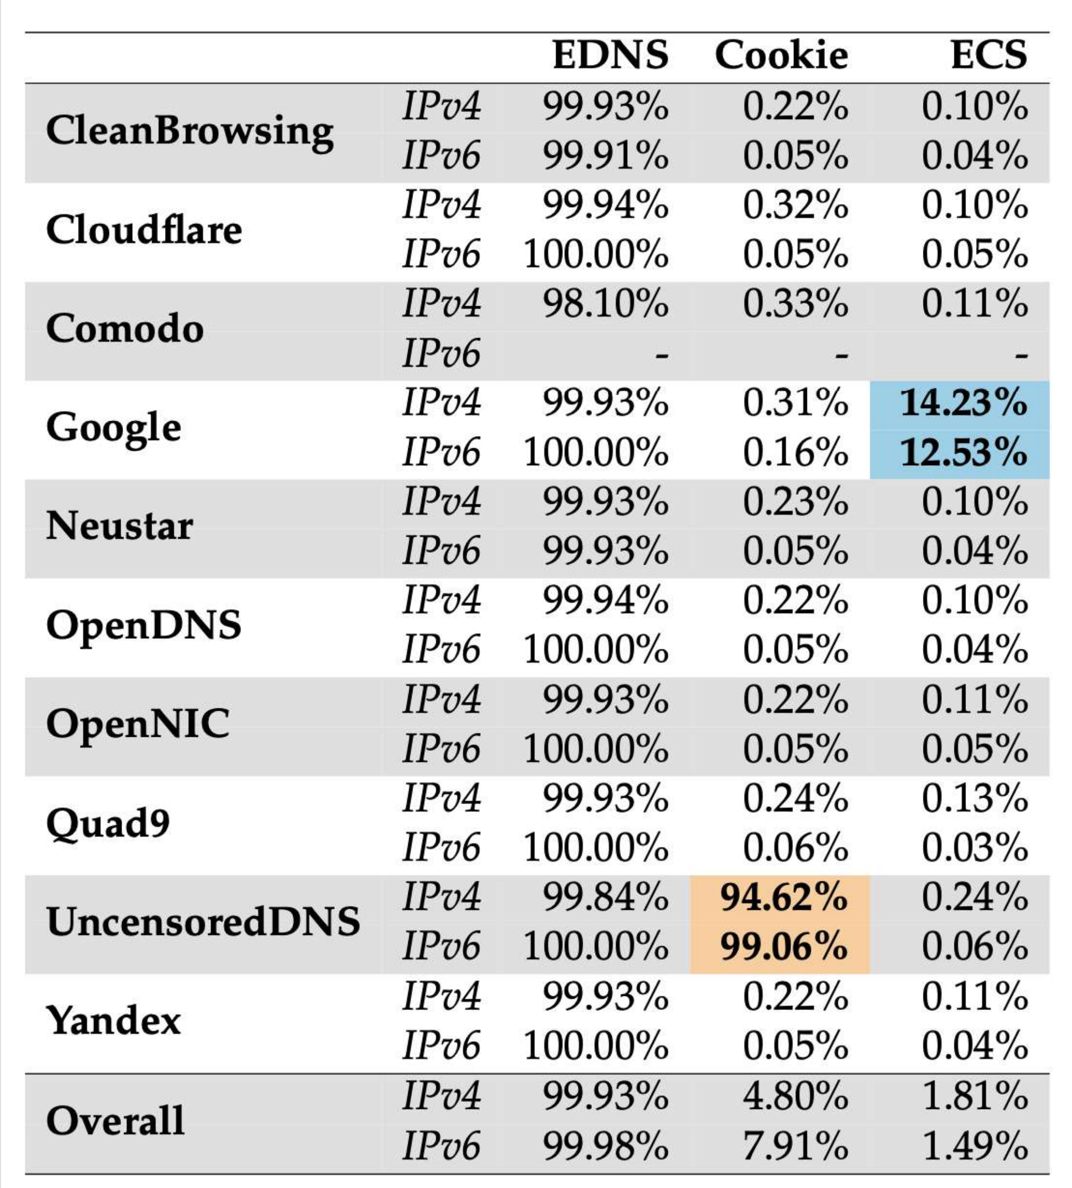

		EDNS	Cookie	ECS
CleanBrowsing	*IPv4*	99.93%	0.22%	0.10%
	IPv6	99.91%	0.05%	0.04%
Cloudflare	*IPv4*	99.94%	0.32%	0.10%
	IPv6	100.00%	0.05%	0.05%
Comodo	*IPv4*	98.10%	0.33%	0.11%
	IPv6	-	-	-
Google	*IPv4*	99.93%	0.31%	**14.23%**
	IPv6	100.00%	0.16%	**12.53%**
Neustar	*IPv4*	99.93%	0.23%	0.10%
	IPv6	99.93%	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:

- 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.
- Most Public resolvers at least in a small percentage of requests use the backend of other resolvers like Google without having any control over the EDNS configuration.

Findings (Evaluation from the Core: TCP Usage - 2KB)

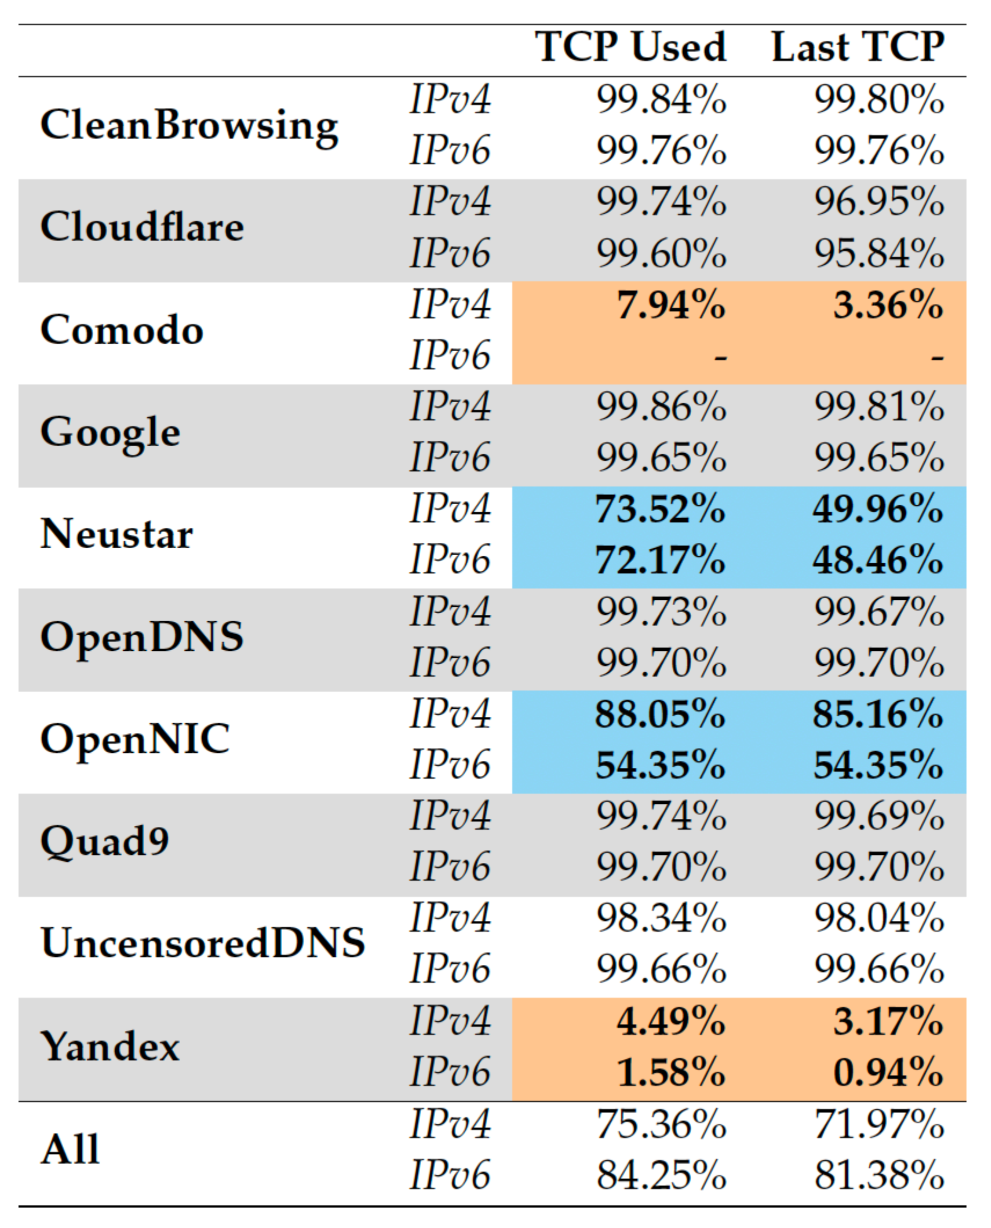

		TCP Used	Last TCP
CleanBrowsing	*IPv4*	99.84%	99.80%
	IPv6	99.76%	99.76%
Cloudflare	*IPv4*	99.74%	96.95%
	IPv6	99.60%	95.84%
Comodo	*IPv4*	**7.94%**	**3.36%**
	IPv6	-	-
Google	*IPv4*	99.86%	99.81%
	IPv6	99.65%	99.65%
Neustar	*IPv4*	**73.52%**	**49.96%**
	IPv6	**72.17%**	**48.46%**
OpenDNS	*IPv4*	99.73%	99.67%
	IPv6	99.70%	99.70%
OpenNIC	*IPv4*	**88.05%**	**85.16%**
	IPv6	**54.35%**	**54.35%**
Quad9	*IPv4*	99.74%	99.69%
	IPv6	99.70%	99.70%
UncensoredDNS	*IPv4*	98.34%	98.04%
	IPv6	99.66%	99.66%
Yandex	*IPv4*	**4.49%**	**3.17%**
	IPv6	**1.58%**	**0.94%**
All	*IPv4*	75.36%	71.97%
	IPv6	84.25%	81.38%

About:

- TCP Used represents all scenarios in which it is used at any point in the request sequence

- The last TCP considers only those sequences ending with a DoTCP request.

Takeaways:

- Neustar shows high differences between the usage rates of TCP in general and in the last request.

- OpenNIC shows high differences between IPv4 and IPv6.

- Yandex and Comodo rarely use DoTCP when they deal with responses of 2KB. This causes fragmentation of the DNS response on its path to the respective resolver

Findings (Evaluation from the Core: TCP Usage - 4KB)

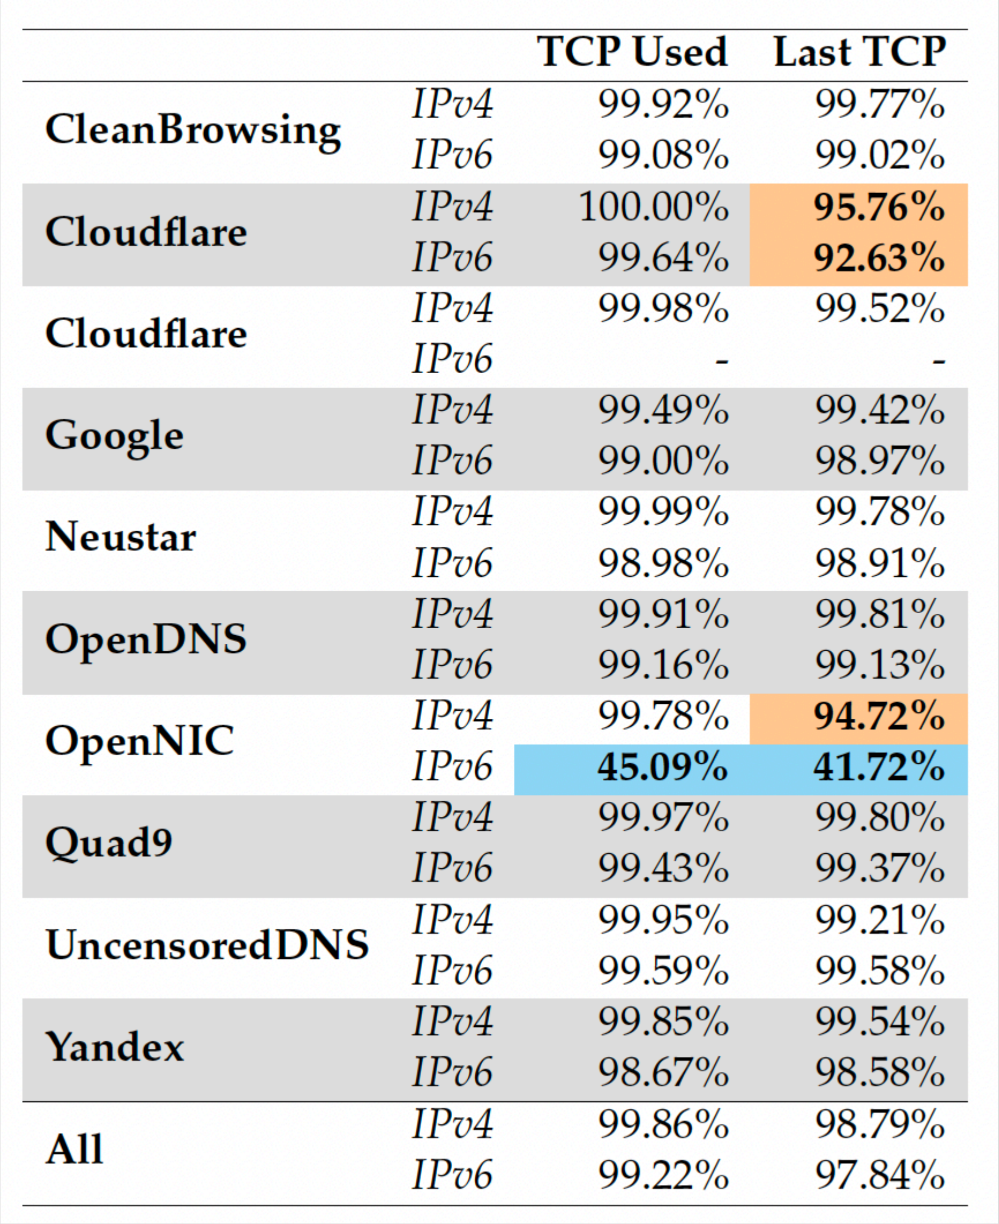

		TCP Used	Last TCP
CleanBrowsing	*IPv4*	99.92%	99.77%
	IPv6	99.08%	99.02%
Cloudflare	*IPv4*	100.00%	**95.76%**
	IPv6	99.64%	**92.63%**
Cloudflare	*IPv4*	99.98%	99.52%
	IPv6	-	-
Google	*IPv4*	99.49%	99.42%
	IPv6	99.00%	98.97%
Neustar	*IPv4*	99.99%	99.78%
	IPv6	98.98%	98.91%
OpenDNS	*IPv4*	99.91%	99.81%
	IPv6	99.16%	99.13%
OpenNIC	*IPv4*	99.78%	**94.72%**
	IPv6	**45.09%**	**41.72%**
Quad9	*IPv4*	99.97%	99.80%
	IPv6	99.43%	99.37%
UncensoredDNS	*IPv4*	99.95%	99.21%
	IPv6	99.59%	99.58%
Yandex	*IPv4*	99.85%	99.54%
	IPv6	98.67%	98.58%
All	*IPv4*	99.86%	98.79%
	IPv6	99.22%	97.84%

Takeaways:

- Over IPv6, OpenNIC uses TCP in less than half of the sequences that reach the name server
- A non-negligible number of measurements for that no TCP is used at all by Comodo, IPv6. (May cause truncation)
- Cloudflare and OpenNIC tend to use DoTCP in their last requests less often than the other resolvers.

Key Takeaways

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

3/10 resolvers announce very large EDNS(0) buffer sizes (4096B) both from the Edge as well as from the Core, which potentially causes fragmentation.

The **resolvers exhibit one preferred buffer size** which is advertised to the name servers in more than 90% of the cases.

In response to responses (2KB and 4KB) from ANSes, some resolvers do not fall back to DoTCP in many cases. This bears the risk of fragmented responses.

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

-paper-