

IP FRAGMENTATION AND MEASURES AGAINST DNS- CACHE-POISONING

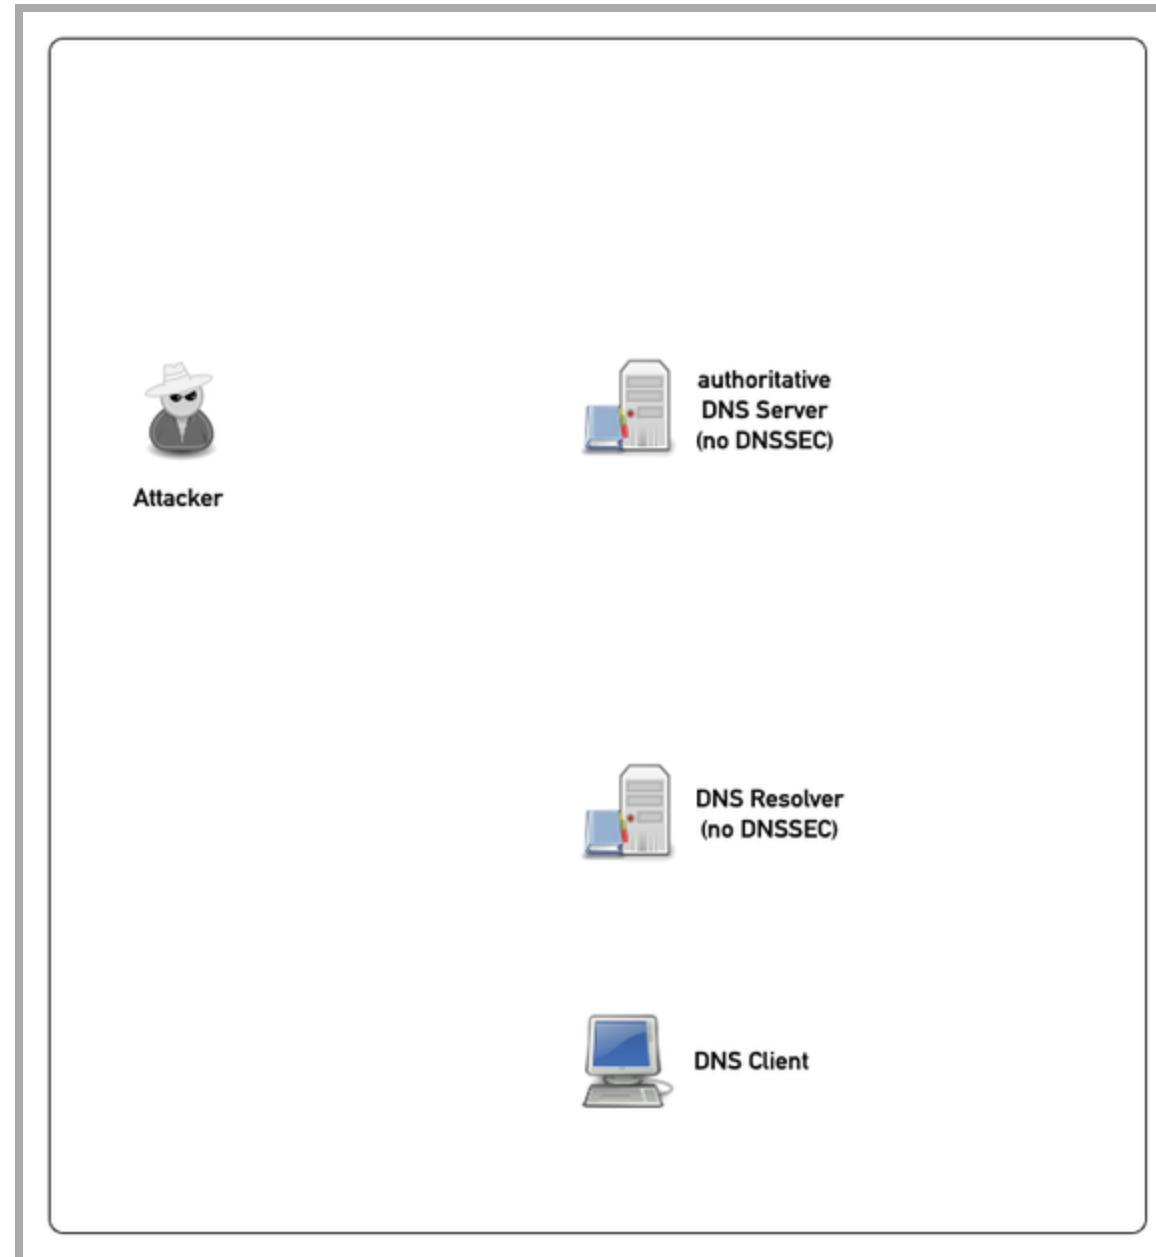
Roland van Rijswijk-Deij (NLnetLabs), Patrick Koetter (sys4), Carsten Strotmann (sys4)

About the project

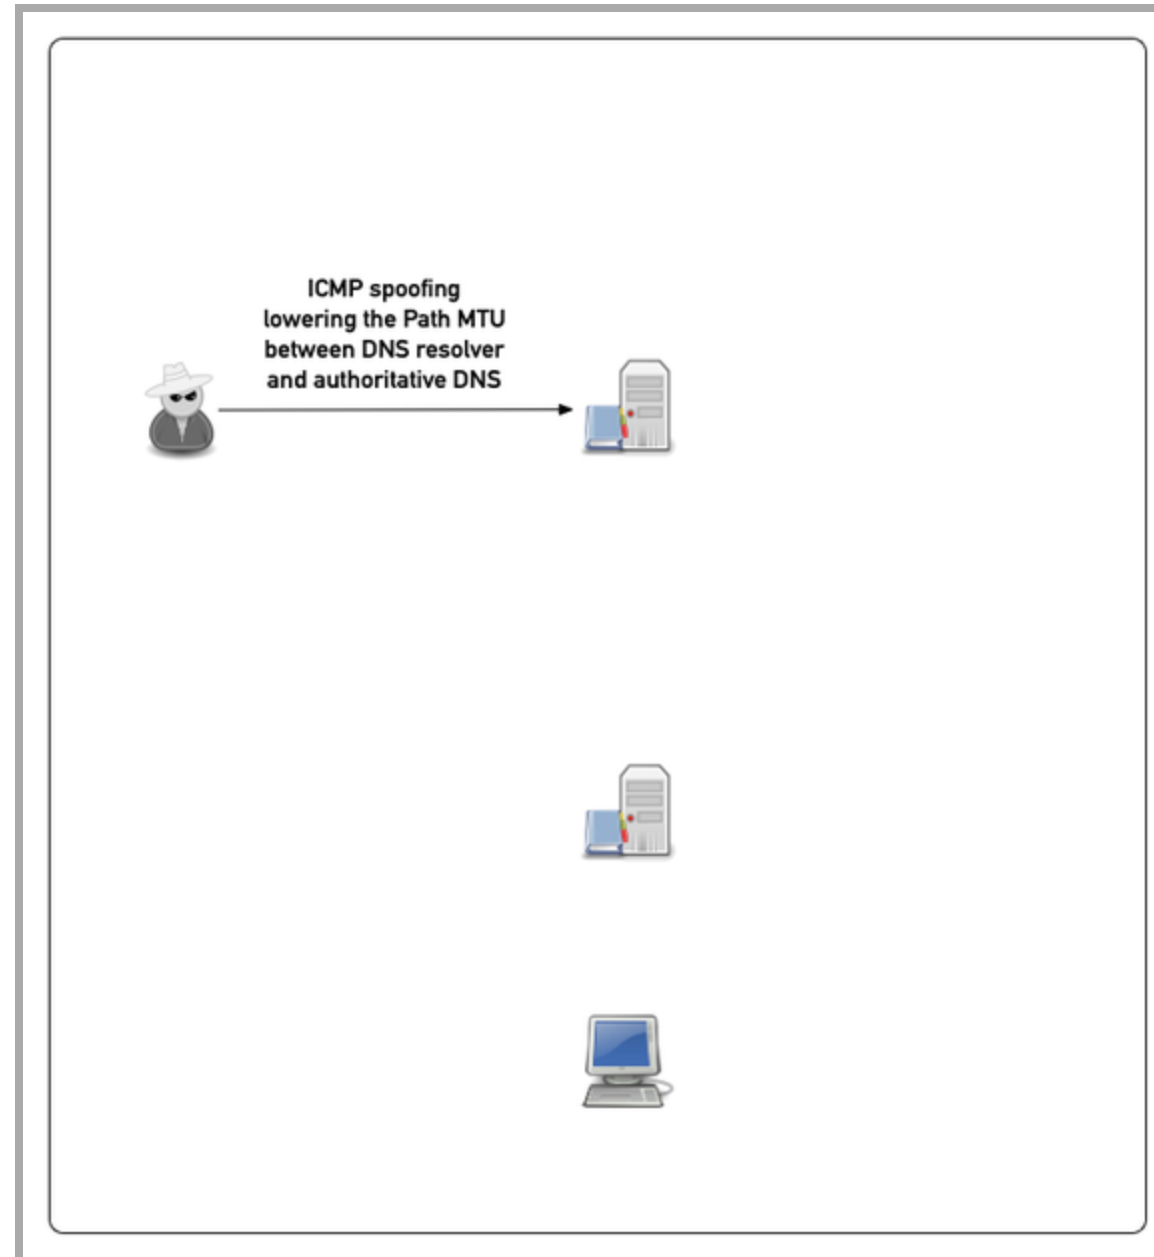
- Study under contract from BSI (German Federal Office for Information Security) between December 2019 and September 2021
- Roland van Rijswijk-Deij (NLnetLabs), Patrick Koetter (sys4), Carsten Strotmann (sys4)
- Questions:
 - Do DNS cache poisoning attacks via fragmentation impose a real threat?
 - Is it possible to mitigate such attacks?
 - How effective are these mitigations?

DNS cache poisoning via DNS fragmentation

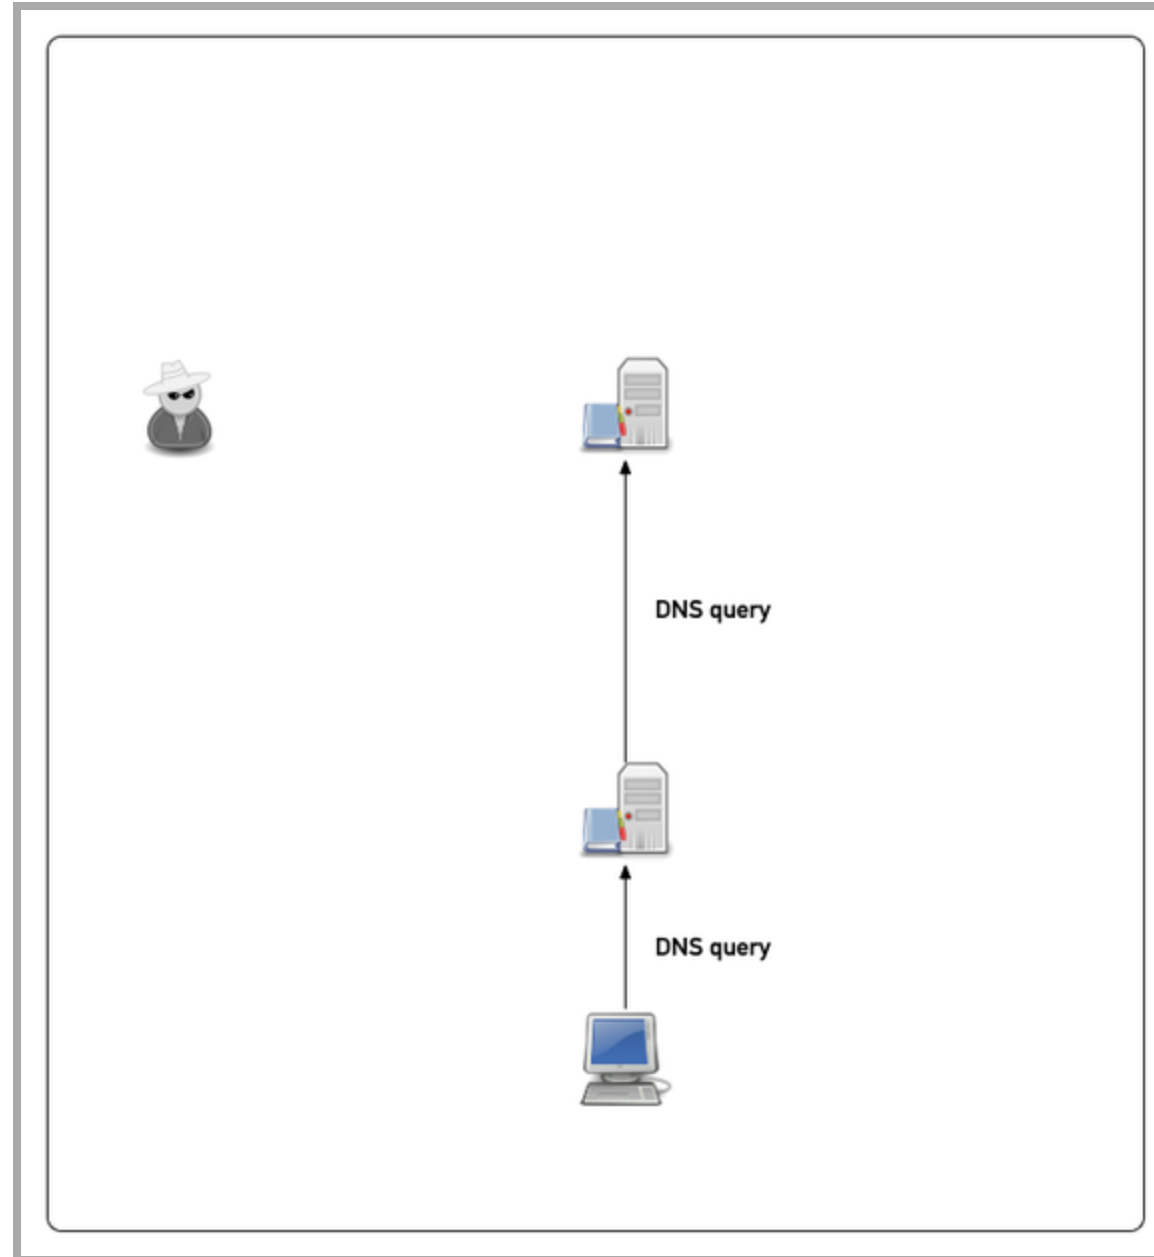
Example attack (simplified)



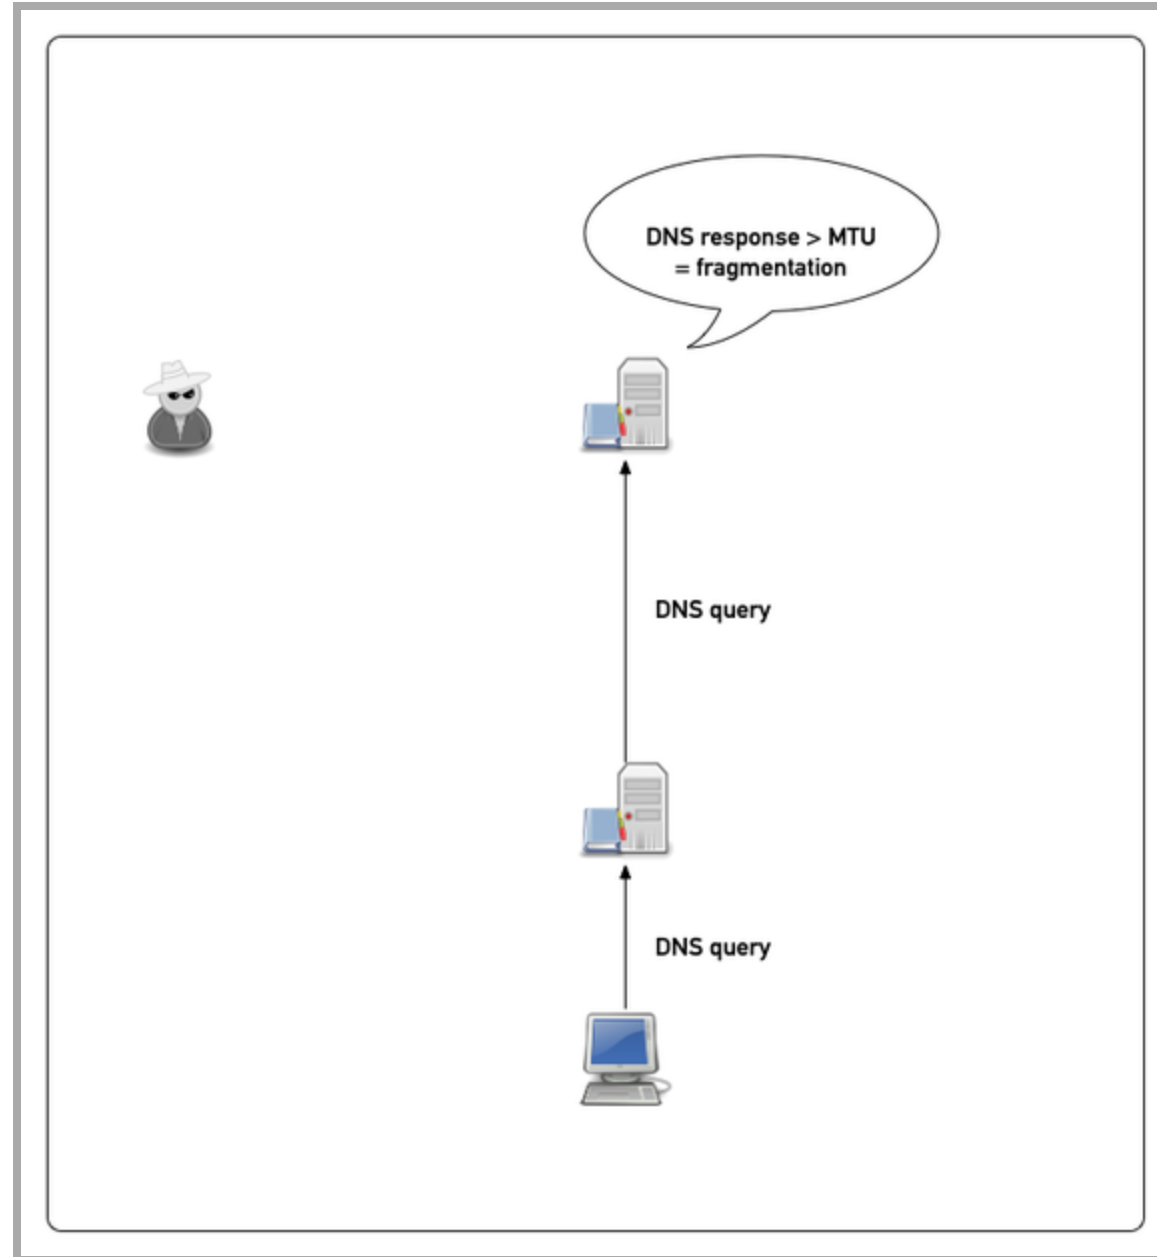
Example attack (simplified)



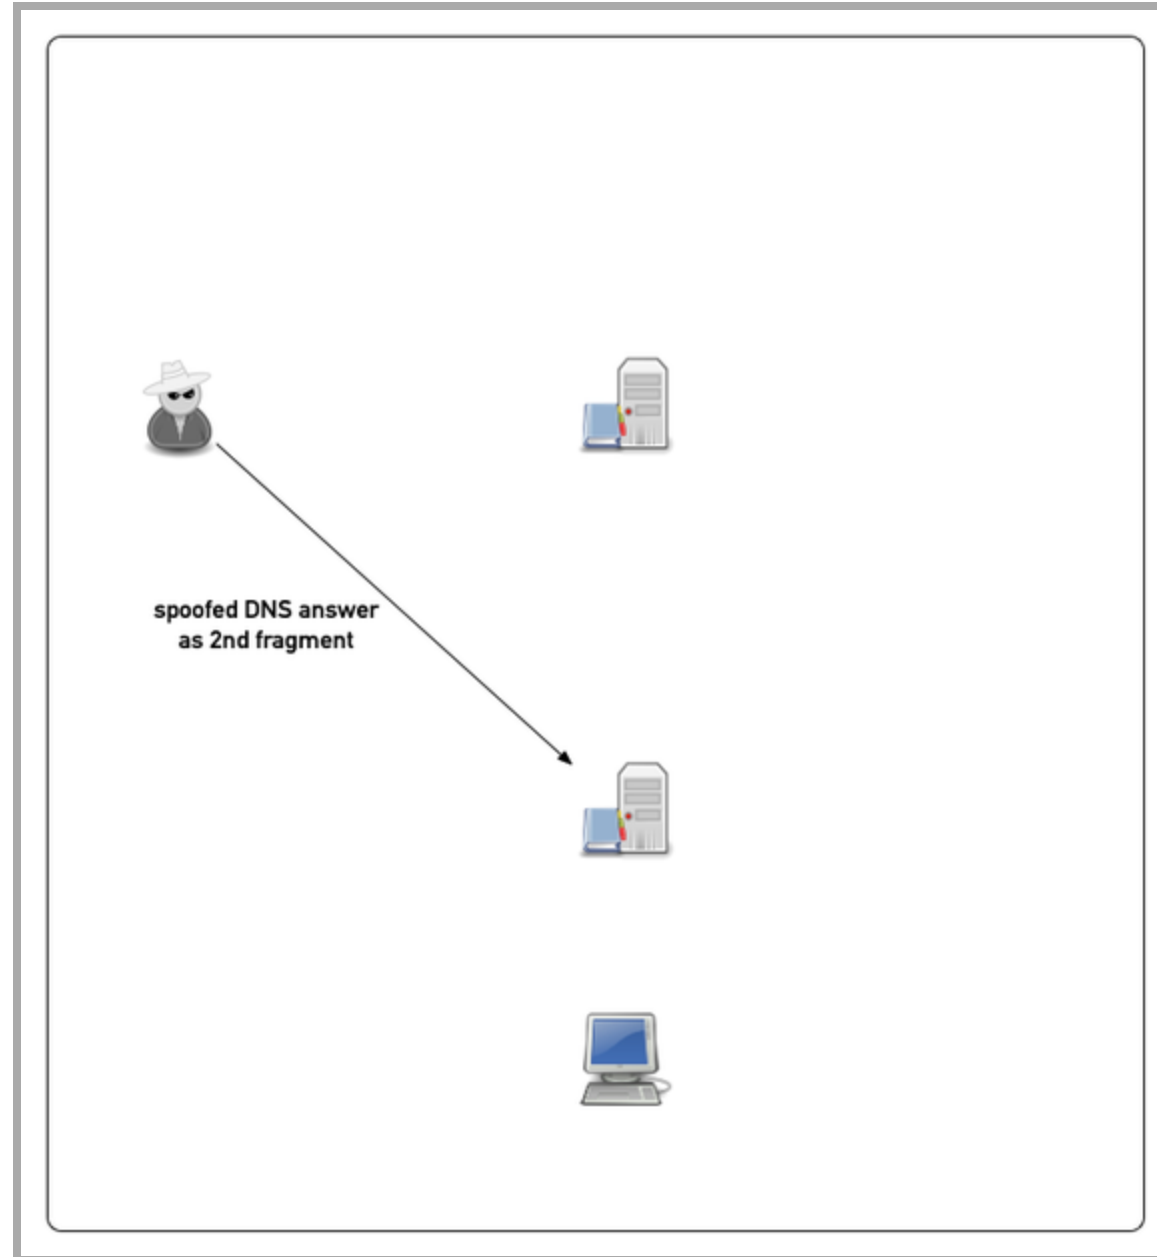
Example attack (simplified)



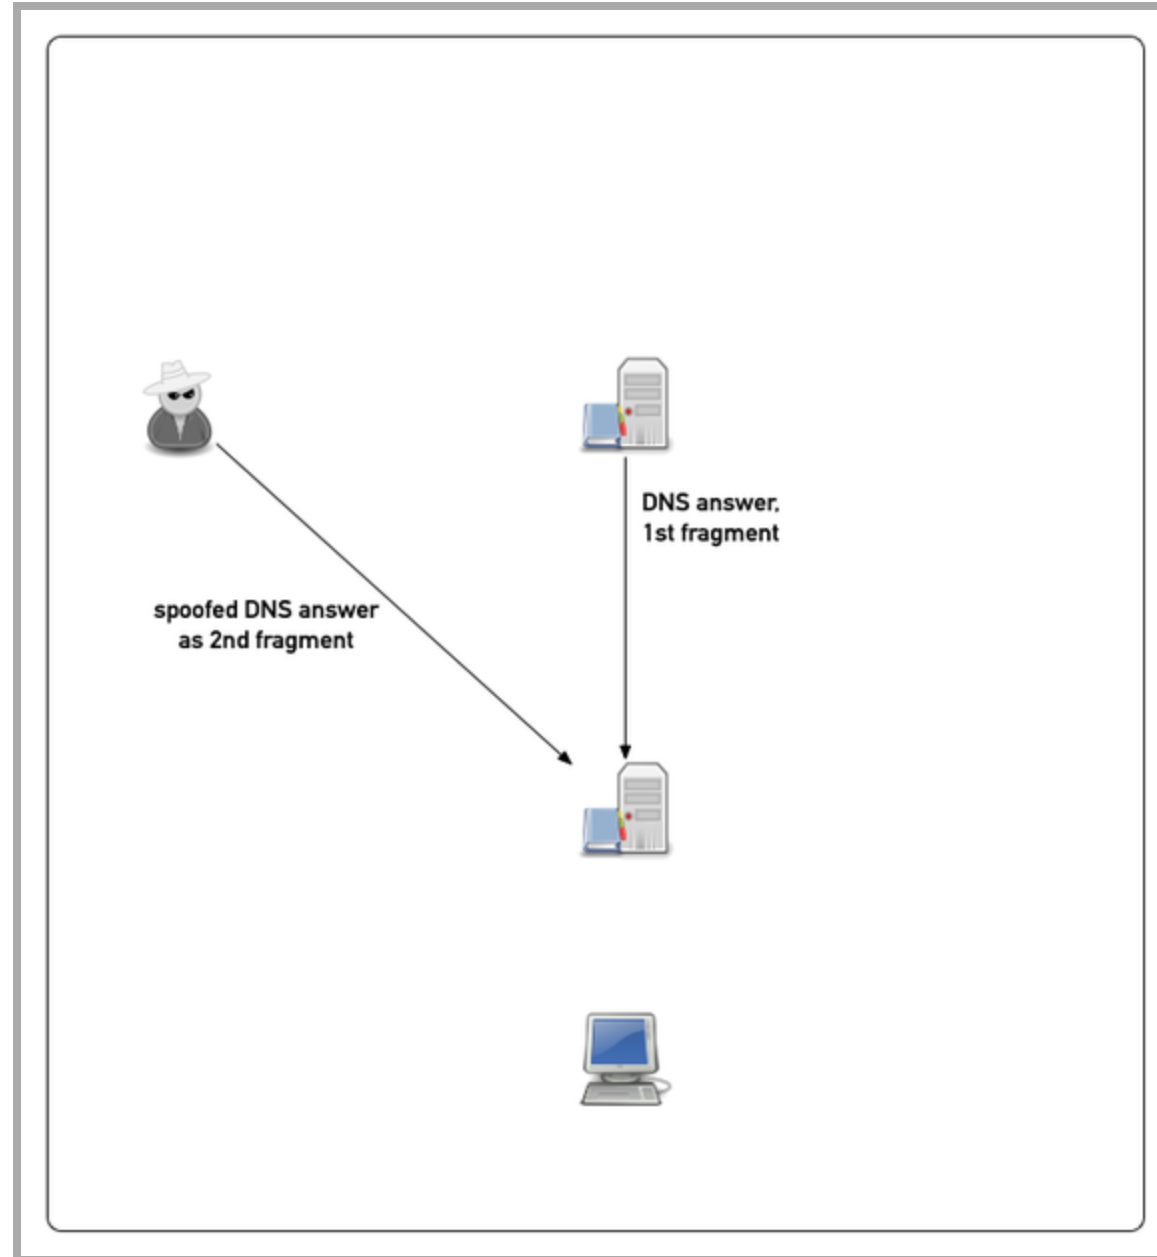
Example attack (simplified)



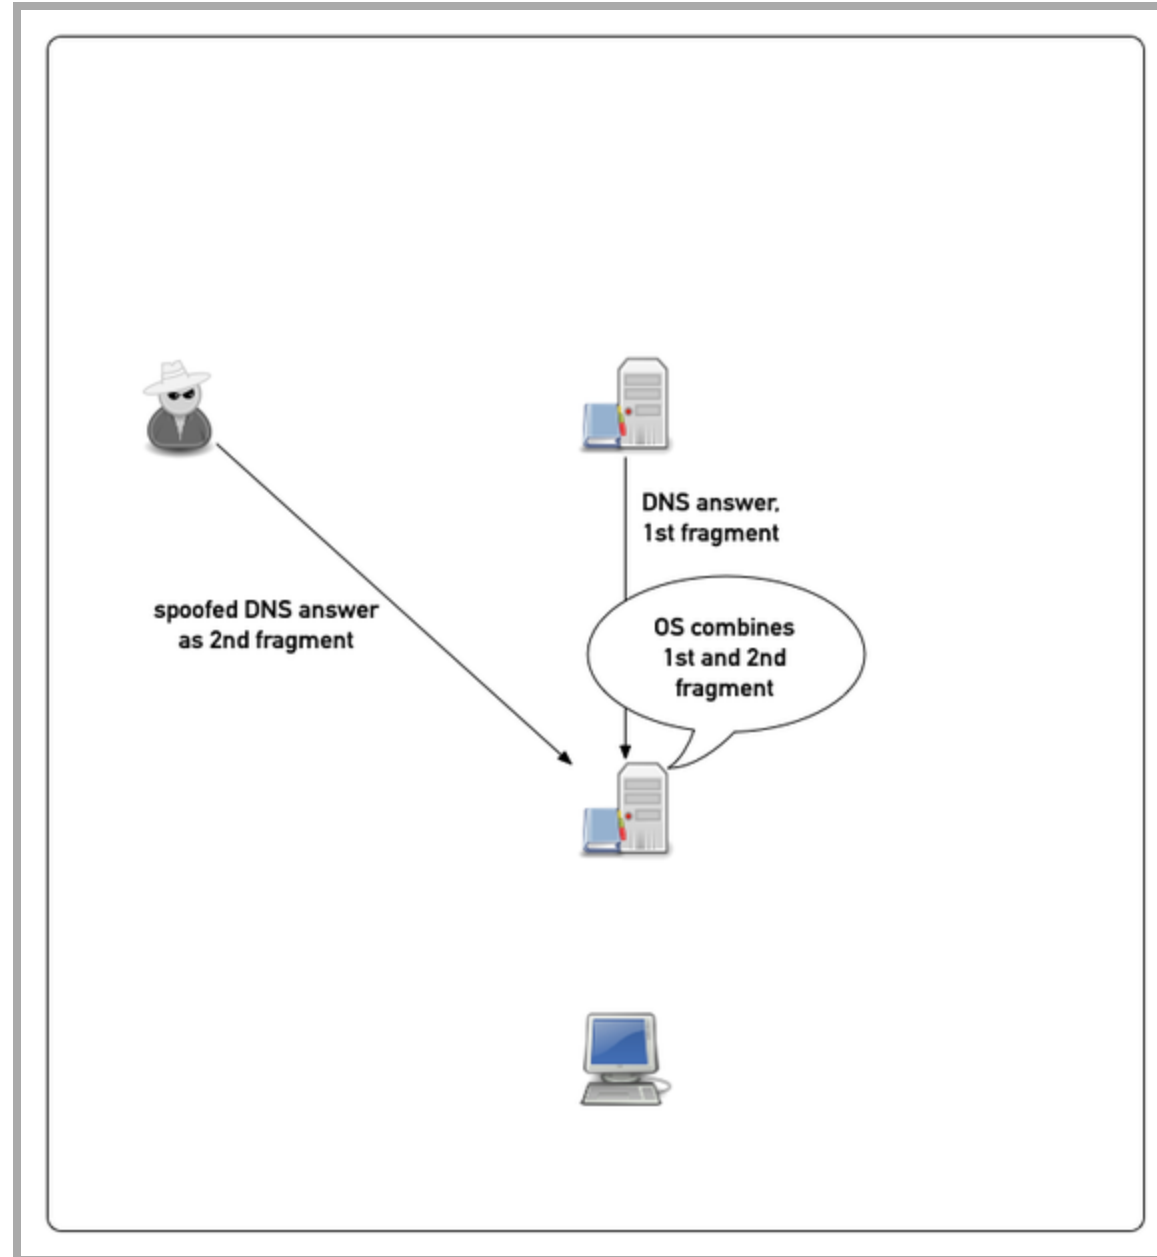
Example attack (simplified)



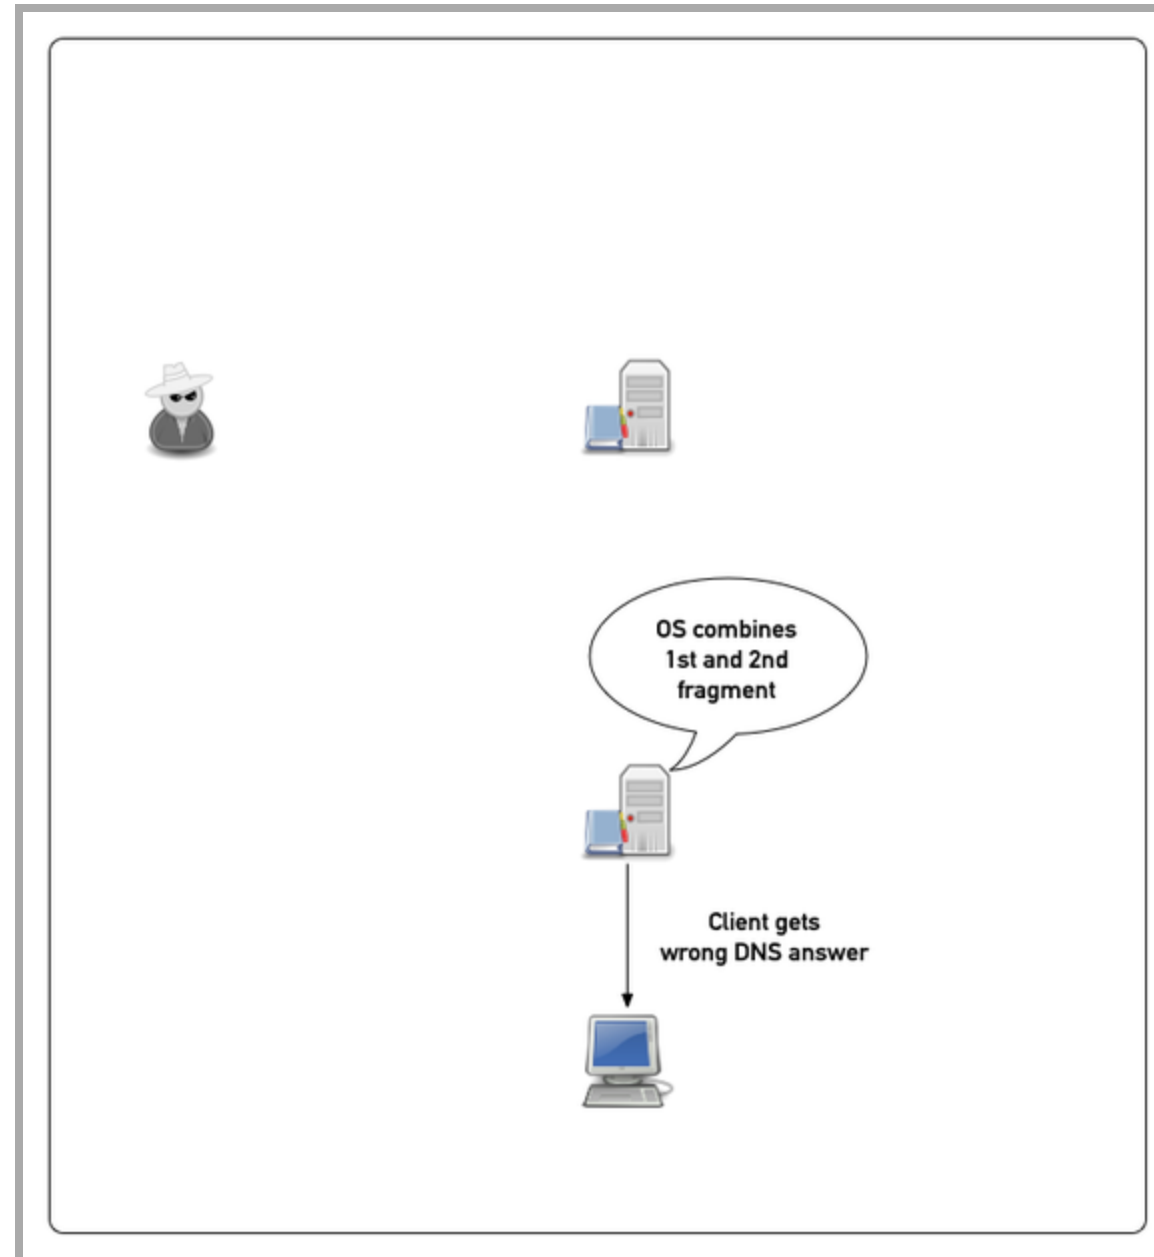
Example attack (simplified)



Example attack (simplified)



Example attack (simplified)



DNS fragmentation on an ISP DNS resolver

DNS fragmentation as perceived on an ISP DNS resolver


- Measurement of number, size and source of fragmented DNS responses on a DNS resolver
- Conducted in July 2020 at a large German ISP with about 4 million home and business Internet access customers

DNS fragmentation on an ISP DNS resolver

- IPv4: 55 064 DNS responses from a total of 54 023 478 have been fragmented (0.10 %)
- IPv6: 104 129 DNS responses from a total of 96 620 298 have been fragmented (0.11 %)
- DNSSEC: 93% (IPv6) and 97% (IPv4) of fragmented DNS answers came from DNSSEC signed zones


Fragmented DNS responses distribution per 24 hours

Number of fragmented DNS responses seen over 24 hours

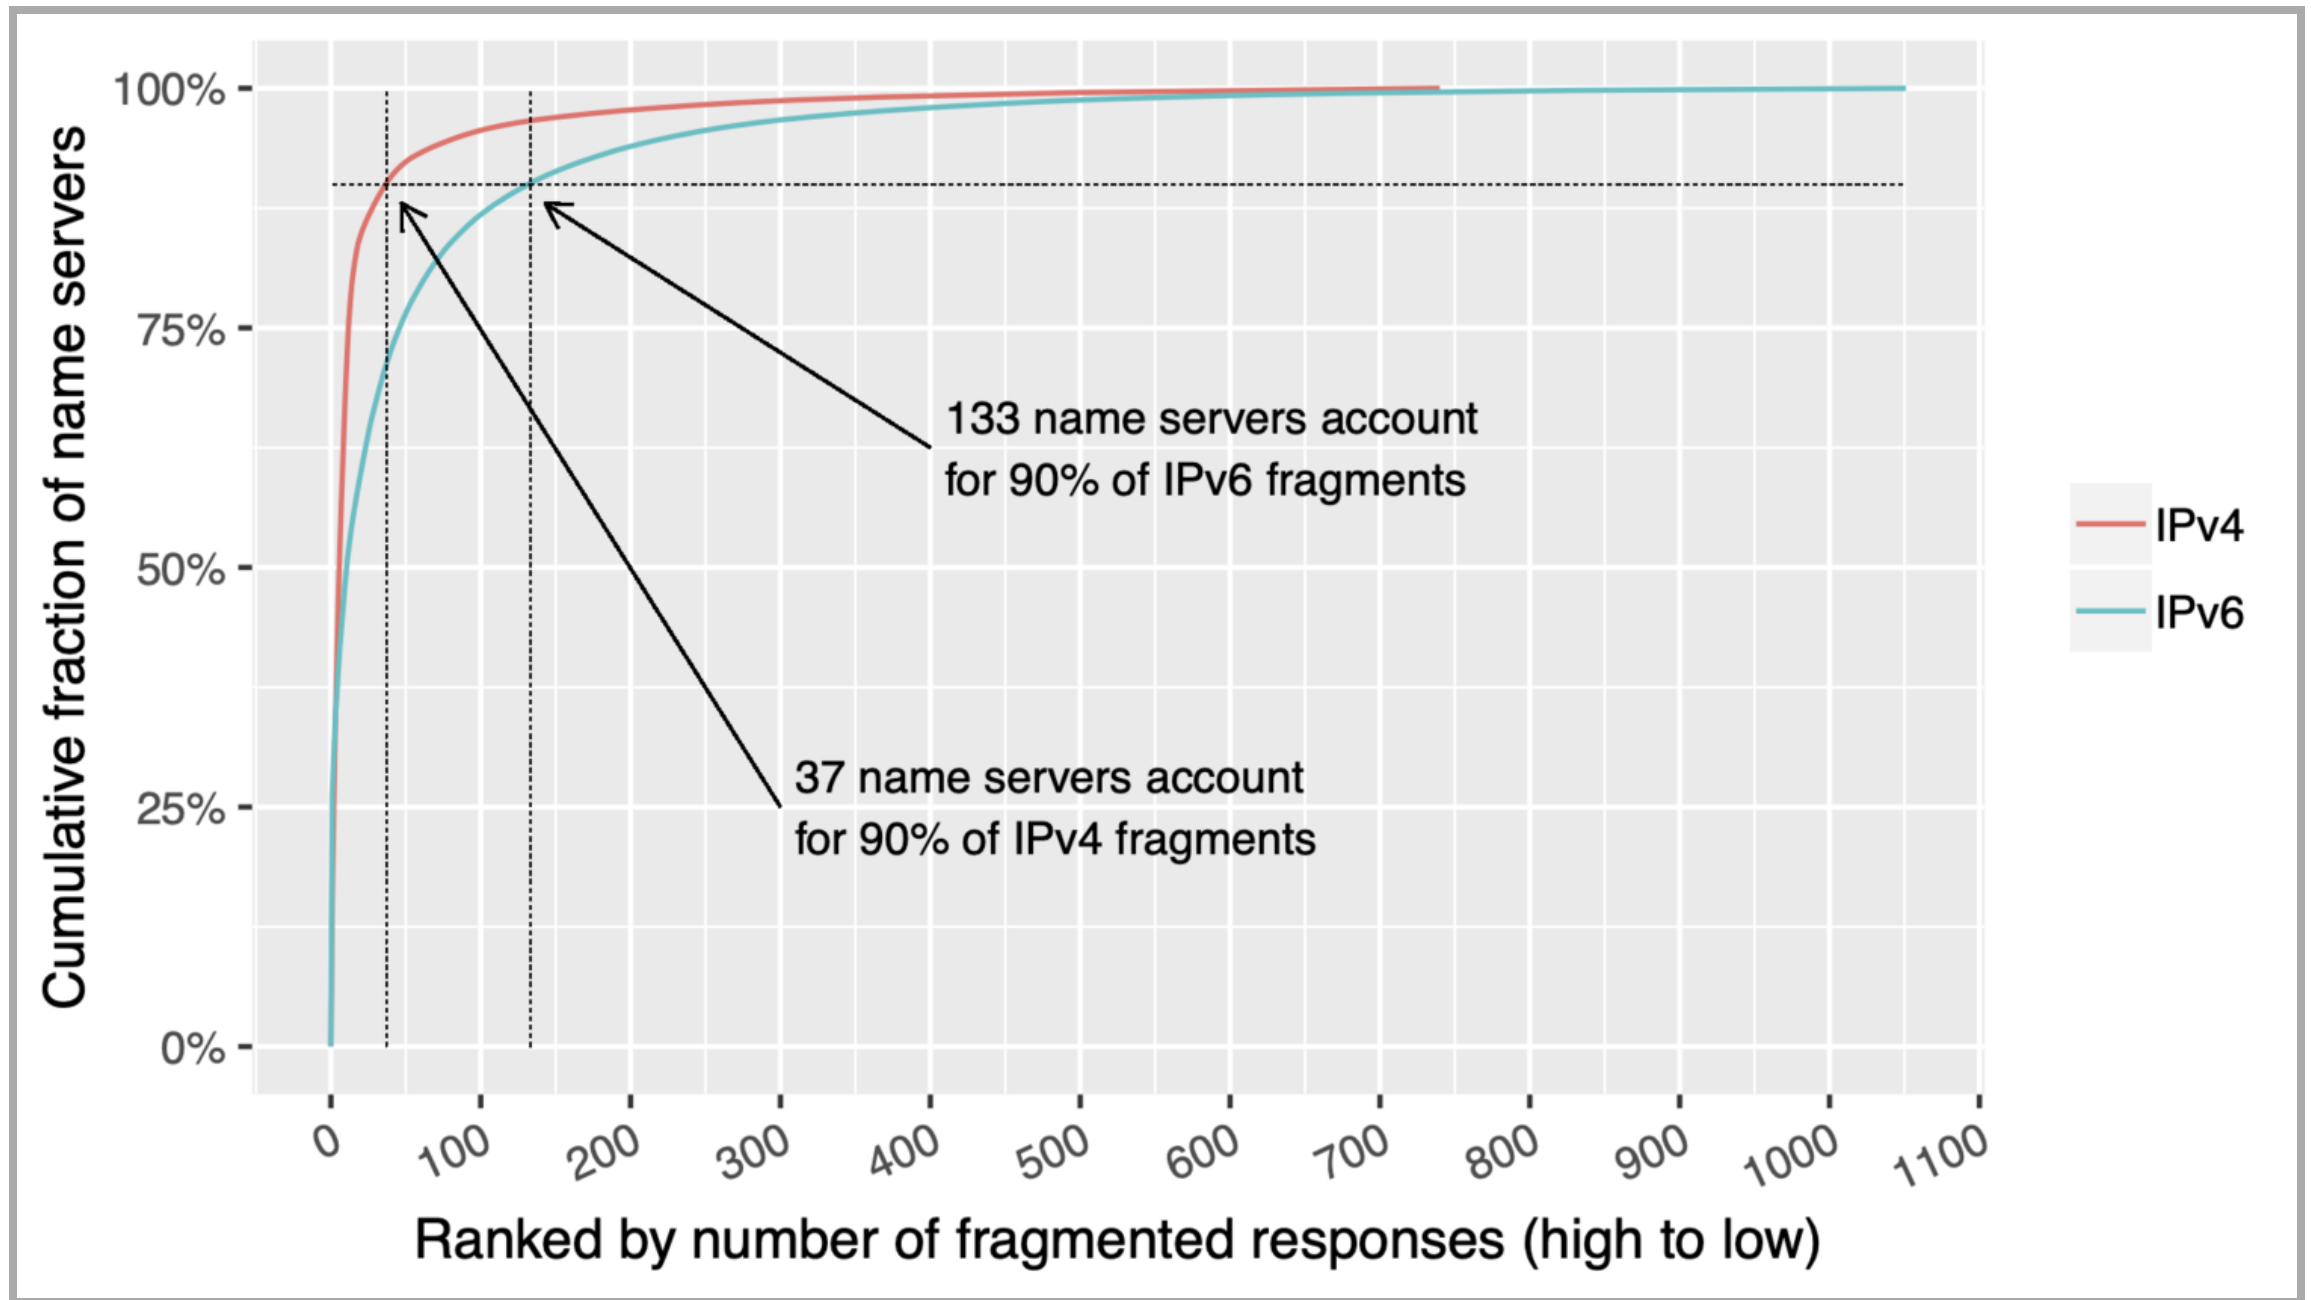


Fragmented DNS responses distribution per 24 hours

Percentage of fragmented DNS responses from the total number of responses over 24 hours



DNS Server sending fragmented DNS responses



Notable domains with fragmented DNS responses

- Domains from where fragmented DNS responses have been seen
 - office.com (Microsoft)
 - army.mil (US Army)
 - fnfis.com (Fidelity National Information Services)
 - ekom21.de (kommunales Gebietsrechenzentrum Hessen)
 - fraunhofer.de (Fraunhofer Gesellschaft)
 - rwe.de (RWE Aktiengesellschaft)
 - agilent.com (Agilent, Research)
 - checkpoint.com (Check Point Security - Firewall and VPN products)
 - salesforce.com and force.com (Salesforce.com, Inc - Cloud based customer relationship management solutions)
 - fedex.com (FexEx Corporation - multi national delivery services company)
 - gnome.org (Gnome Desktop Software - open source GUI desktop for Linux and Unix)


OpenINTEL

- OpenINTEL is an Internet research platform that collects DNS responses from 227 000 000 DNS domains
- OpenINTEL observes around 60% of the public Internet
- OpenINTEL processes 2.4 billion DNS datasets per day
- This study looked into the fragmentation seen in NS, A and AAAA DNS responses

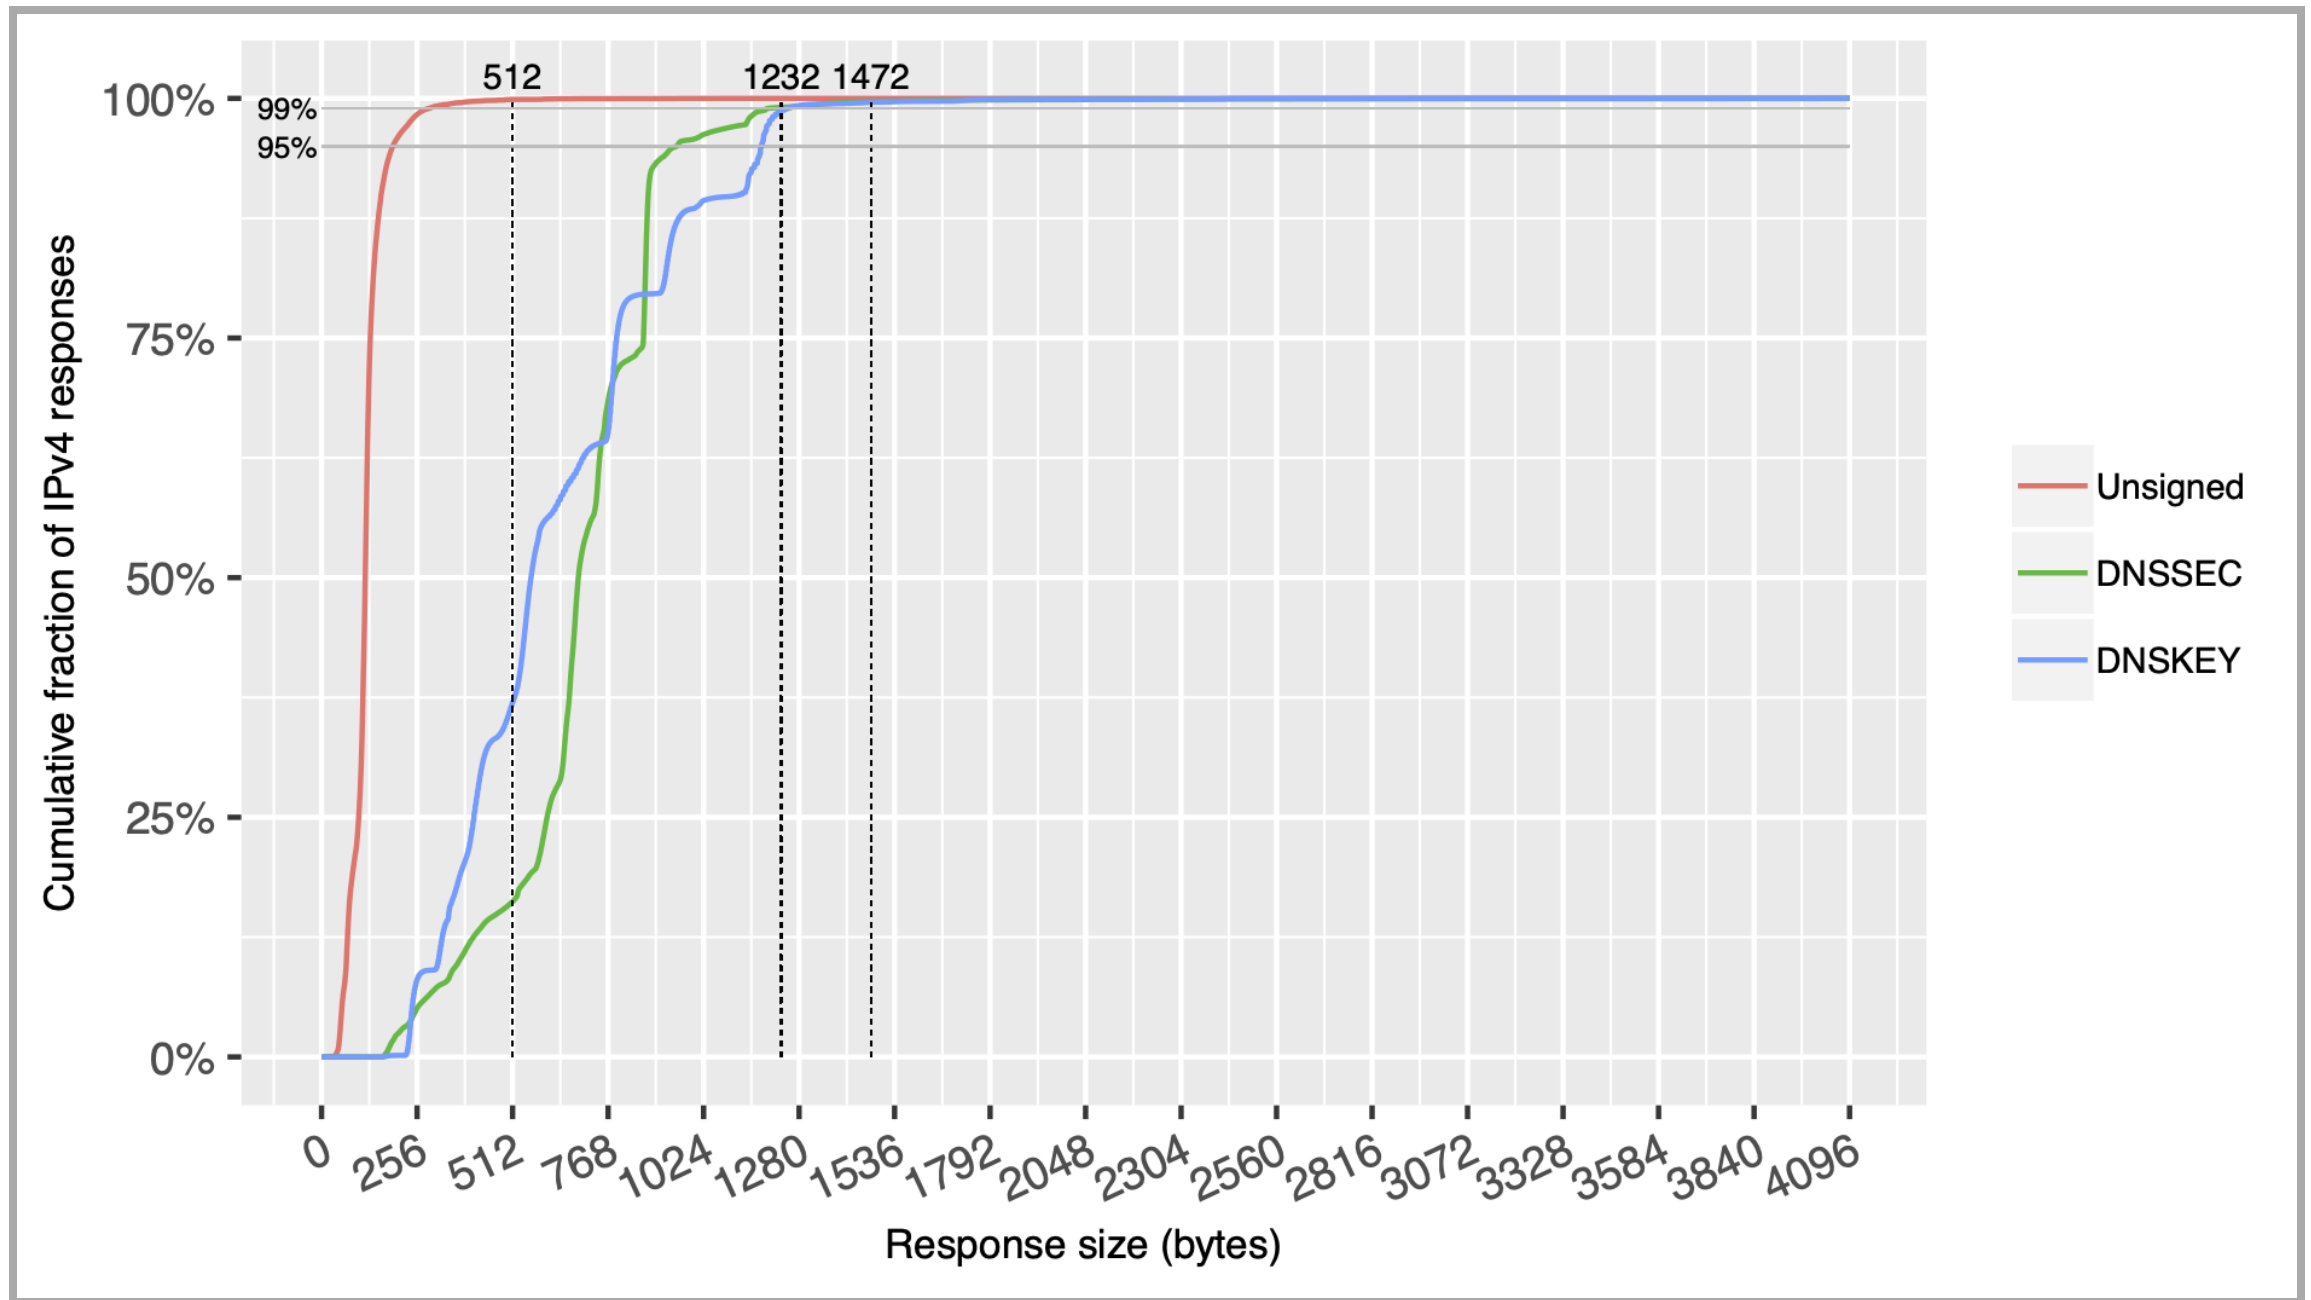
How much DNS fragmentation is seen by OpenINTEL?

- 3 893 453 582 DNS responses
 - IPv4: 2 837 177 438 [72.870%]
 - IPv6: 1 056 276 144 [27.130%]
- fragmented responses
 - IPv4: 1 334 549 [0.047%]
 - IPv6: 1 008 894 [0.096%]

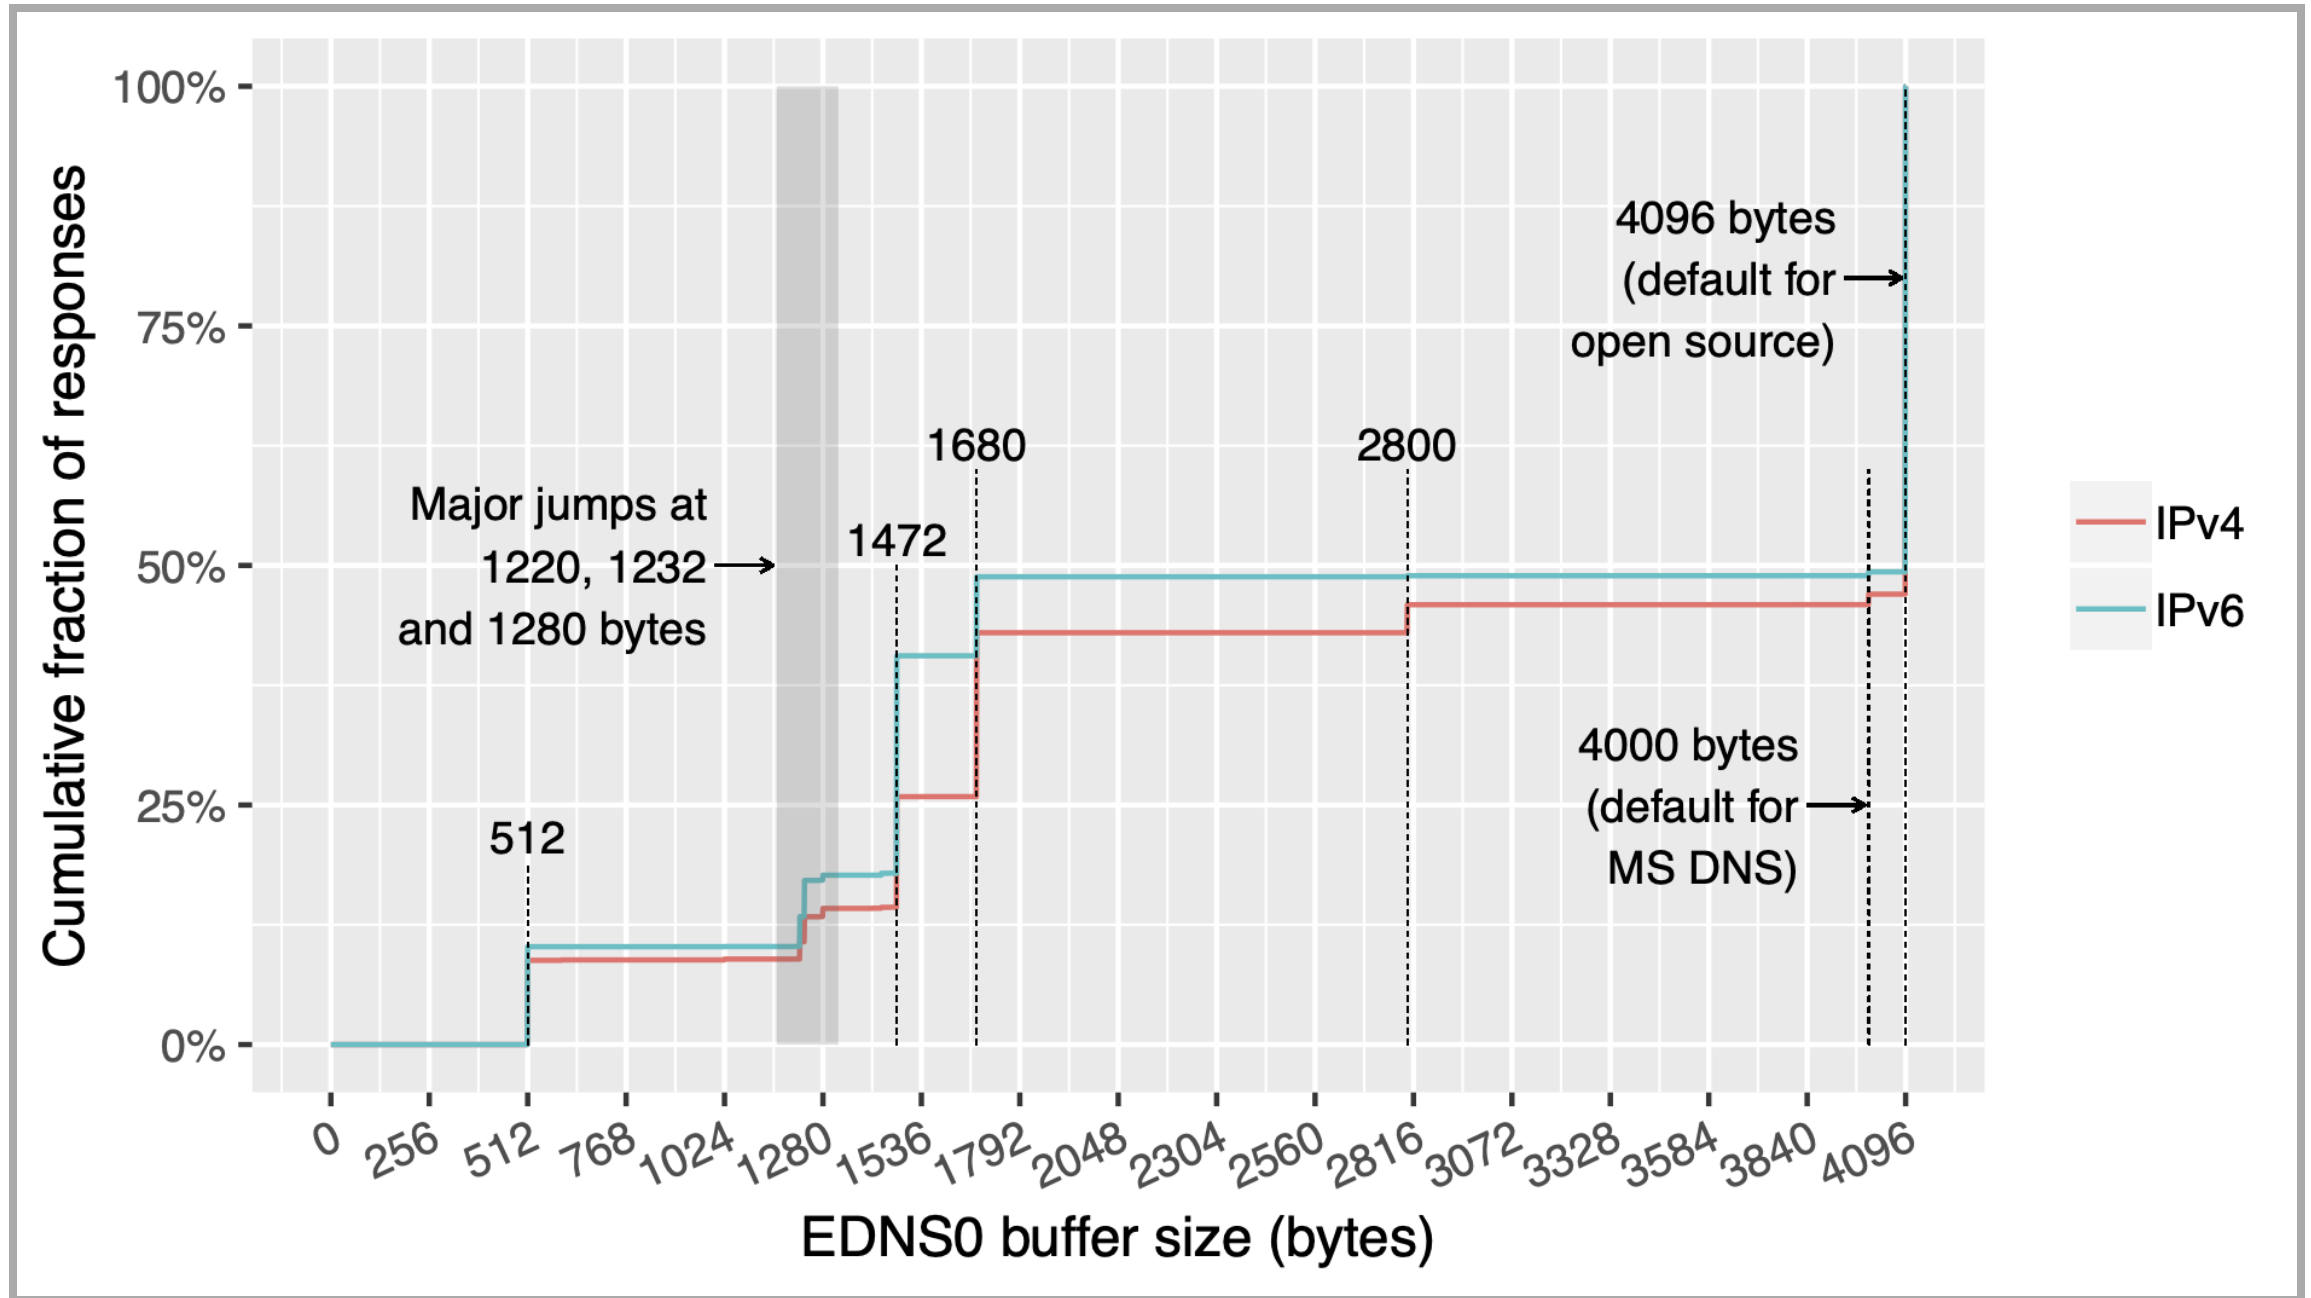
OpenINTEL: Size of DNS datagrams over IPv6



OpenINTEL: Size of DNS datagrams over IPv4



OpenINTEL: Size of the advertised EDNS buffer



Authoritative DNS servers supporting TCP

- A DNS response that does not fit into an UDP response must be sent over TCP
- Response size limits of DNS UDP messages:
 - 512 Byte: classic DNS RFC 1034/1035 (1987)
 - 4096 Byte: EDNS RFC 2671 (1999)
 - 1232 Byte: popular recommendation to prevent DNS fragmentation
- Question: How many authoritative DNS servers support DNS/TCP?
 - How *popular* are the domains that are hosted on DNS sever that **do not** support DNS/TCP?


TCP Support

- 879 345 IPv4/IPv6 addresses of authoritative DNS server
 - This DNS server are authoritative for 202 765 149 domains
 - 197 773 383 (97.57%) of these domains have at least one DNS server offering DNS/TCP
 - From 183 549 827 (90.55%) domains all announced DNS servers (NS Record) offer DNS/TCP
 - 4 925 715 (2.43%) of the surveyed domains have **no** DNS server supporting DNS/TCP

TCP Support

- Domains where at least one DNS server does not support DNS/TCP contain popular Internet destinations such as `live.com`, `office.com` (Microsoft) and `yahoo.com` (Yahoo)
- 1.5% of all domains of the Tranco 1M list (list of the 1 million most popular Internet domains) have **no** DNS server with TCP support

Rank of Tranco 1M domains lacking TCP support



DNS/TCP Support - Conclusion

- Few, but also some *popular* domains do not support DNS over TCP
- Usage of DNS over TCP to mitigation DNS fragmentation attacks is therefore not recommended

Which Operating Systems are vulnerable to ICMP PathMTU Spoofing?

- To increase the success of a DNS attack via fragmentation, a attacker would try to lower the Path-MTU between the DNS resolver and the authoritative DNS server
 - This can be done by sending spoofed ICMP error messages
- Question: Which popular operating systems are vulnerable to ICMP Path-MTU spoofing?

Operating-Systems and ICMP Path-MTU Spoofing

- Tested the vulnerability of popular operating systems for ICMP Spoofing in a lab environment
- Question: Would an authoritative DNS server send fragmented DNS responses after a successful Path-MTU spoofing attack?

Operating Systems and ICMP PathMTU Spoofing

| Operating System | minMTU IPv4 | minMTU IPv6 | success IPv4 | success IPv6 |
|--|-------------|-------------|--------------|--------------|
| Debian 6 / Kernel 2.6.32-5-amd64 | 552 | 1.280 | X | X |
| Ubuntu 14.04.1 / Kernel 3.13.0-45-generic (12/2014) | 552 | 1.280 | X | X |
| Ubuntu 14.04.1 LTS / Kernel 3.13.0-170-generic (05/2019) | 552 | 1.280 | X | X |
| Ubuntu 16.04.6 LTS / Kernel 4.4.0-184-generic (06/2020) | 552 | 1.280 | X | X |
| Ubuntu 18.04.4 LTS / Kernel 4.15.0-106-generic (06/2020) | 1.500 | 1.280 | - | X |
| CentOS 6 / Kernel 2.6.32-504.3.3.el6.x86_64 (12/2014) | 552 | 1.280 | X | X |
| CentOS 7 / Kernel 3.10.0-1127.10.1.el7.x86_64 (06/2020) | 1.500 | 1.280 | - | X |
| CentOS 8 / Kernel 4.18.0-147.8.1.el8_1.x86_64 (04/2020) | 1.500 | 1.280 | - | X |
| SUSE EL 15SP1 / Kernel 4.12.14-197.45-default (06/2020) | 1.500 | 1.280 | - | X |
| FreeBSD 12.1 / Kernel 12.1-RELEASE r354233 GENERIC amd64 | 1.500 | 1.280 | - | X |
| OpenBSD 6.7 / Kernel 6.7 GENERIC#234 i386 | 1.500 | 1.280 | - | X |
| Windows Server 2008R2 | 1.500 | 1.280 | - | X |
| Windows Server 2012R2 | 1.500 | 1.280 | - | X |
| Windows Server 2016 | 1.500 | 1.280 | - | X |
| Windows Server 2019 | 1.500 | 1.280 | - | X |

Operating Systems used for DNS Server

- The popular BIND 9 DNS server software responds with it's version number over DNS on request
 - This version number often contains the version of the Linux-Kernel and the version of the Linux distribution
 - We've used OpenINTEL to query for the versions used on authoritative DNS server

Operating Systems used for DNS Server (Summer 2020)



| Linux OS | Number of Server | Percent from total |
|--------------|------------------|--------------------|
| RedHat Linux | 240.481 | 28.2% |
| RedHat EL5 | 7.876 | 0.9% |
| Redhat EL6 | 98.443 | 11.5% |
| RedHat EL7 | 121.103 | 14.2% |
| RedHat EL8 | 1.594 | 0.2% |
| Ubuntu Linux | 25.034 | 2.9% |
| Ubuntu 14.04 | 5.110 | 0.6% |
| Ubuntu 16.04 | 9.314 | 1.1% |
| Ubuntu 18.04 | 9.467 | 1.1% |

Operating-Systems and ICMP Path-MTU spoofing - conclusion

- Windows operating systems are not vulnerable (to Path-MTU spoofing)
- Older Linux-Kernel are vulnerable
 - These older Linux-Kernel are still in use in long-term support Enterprise-Linux systems!
 - The vulnerable Linux versions are used for authoritative DNS server on the Internet

IP Fragmentation, DNS and TCP

- A joint study of Tianxiang Dai, Dr Haya Shulman and Professor Michael Waidner (Summer 2021) found that DNS communication over TCP can also be vulnerable to IP fragmentation attacks
 - Attacking DNS-over-TCP with IP fragmentation is very similar to attacks on DNS-over-UDP traffic

The recommendations to move to TCP to avoid the fragmentation attacks that apply to UDP, do not solve the problem.

- <https://blog.apnic.net/2021/11/09/dns-over-tcp-considered-vulnerable/>
- <https://www.youtube.com/watch?v=MAfvUbwktOM&t=2285s>

RECOMMENDATIONS FOR DNS RESOLVER OPERATORS



- Use the default Ethernet MTU in DNS server networks
- Restrict the DNS response size over UDP (1 232 Byte)
- Support DNS over TCP
- Evaluate the security risks of running long term supported OS
- Drop Fragmented DNS Responses
- Monitor DNS Traffic for DNS fragmentation

RECOMMENDATIONS FOR OPERATORS OF AUTHORITATIVE DNS SERVERS



- Use the default Ethernet MTU in DNS server networks
- Restrict the DNS response size over UDP (1 232 Byte)
- Support DNS over TCP
- Evaluate the security risks of running long term supported OS
- Deploy DNSSEC Signed Zones
- Avoid Large DNS Resource Record Sets
- Minimize "ANY" Responses
- Enable Minimal Responses

Conclusion (1/2)

- It is possible to attack DNS content by means of DNS fragmentation
- The amount of *natural* (non attack) DNS fragmentation in the Internet is minimal yet still significant
- Popular domains are vulnerable
- Fragmentation of DNS responses should be avoided
 - Among the tested mitigations, lowering the EDNS buffer is the most effective one
 - once the EDNS buffer is lowered, no *natural* fragmentation should occur. All remaining fragmentation can be dropped at (host-)firewall level

Conclusion (2/2)

- The mitigations against DNS fragmentation focus on the effect and do not eliminate the cause
 - DNS cache poisoning and many other attacks on DNS infrastructure would cease to exist if operators began to DNSSEC-sign their DNS zones and DNS resolvers would DNSSEC-verify DNS responses by default

The future

- The DNS protocol is seeing significant changes these days:
 - new encrypted transport protocols such as DNS-over-TLS and DNS-over-HTTPS are being deployed
 - QUIC is a new transport protocol being standardized in the IETF, DNS-over-QUIC might replace the classic DNS-over-UDP in the future
- These new transport protocols (might) solve the problems with fragmentation for DNS
 - But it will be likely more than a decade until widespread deployment is being seen in the Internet