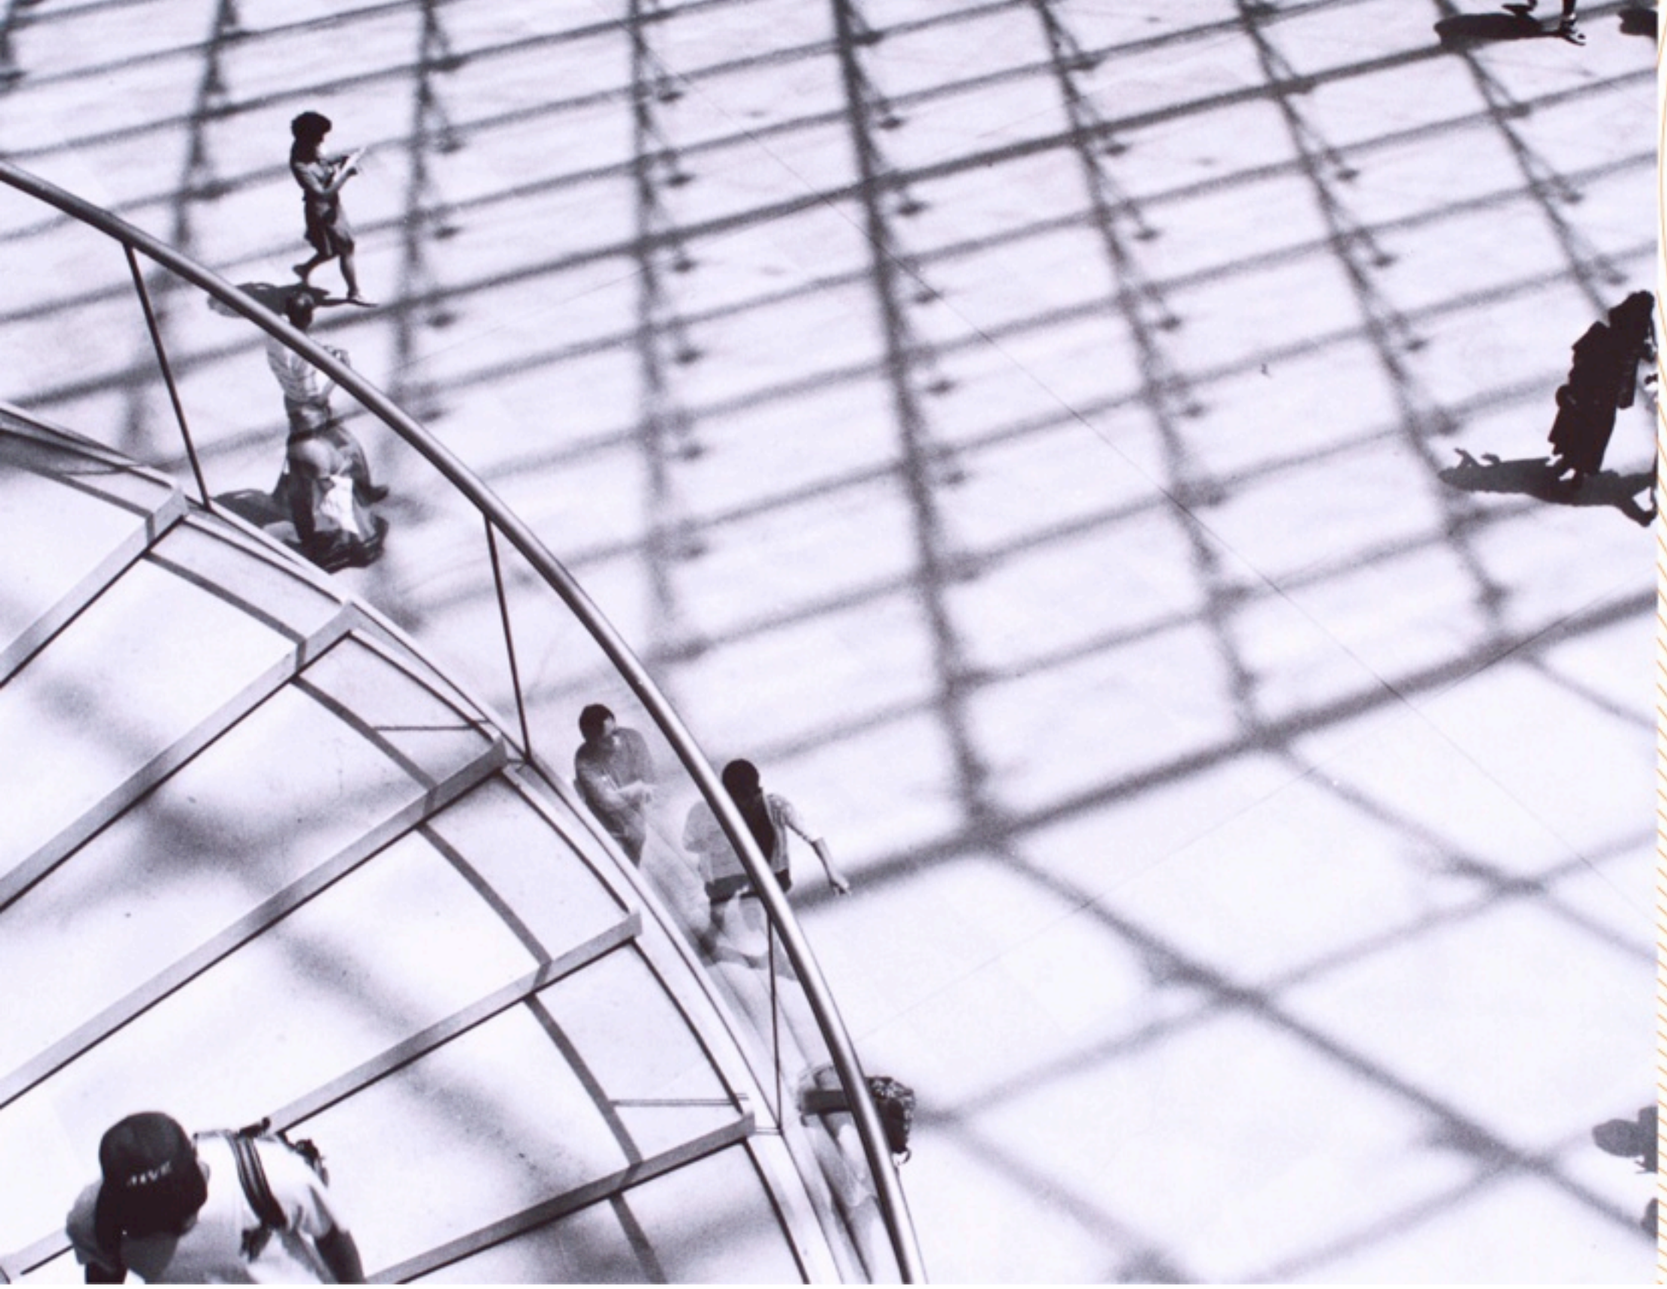

DNSSEC: dealing with hosts that don't get fragments

RIPE 64 DNS wg, Ljubljana, April 18th 2012

Roland van Rijswijk
roland.vanrijswijk@surfnet.nl

Introducing the issue

- In 2010 and in March last year we had “issues” with a very large ISP in The Netherlands

- Customers of the ISP were unable to resolve names in surfnet.nl

- The cause turned out to be an issue with the ISP’s firewall

A picture to make it clearer ;-)

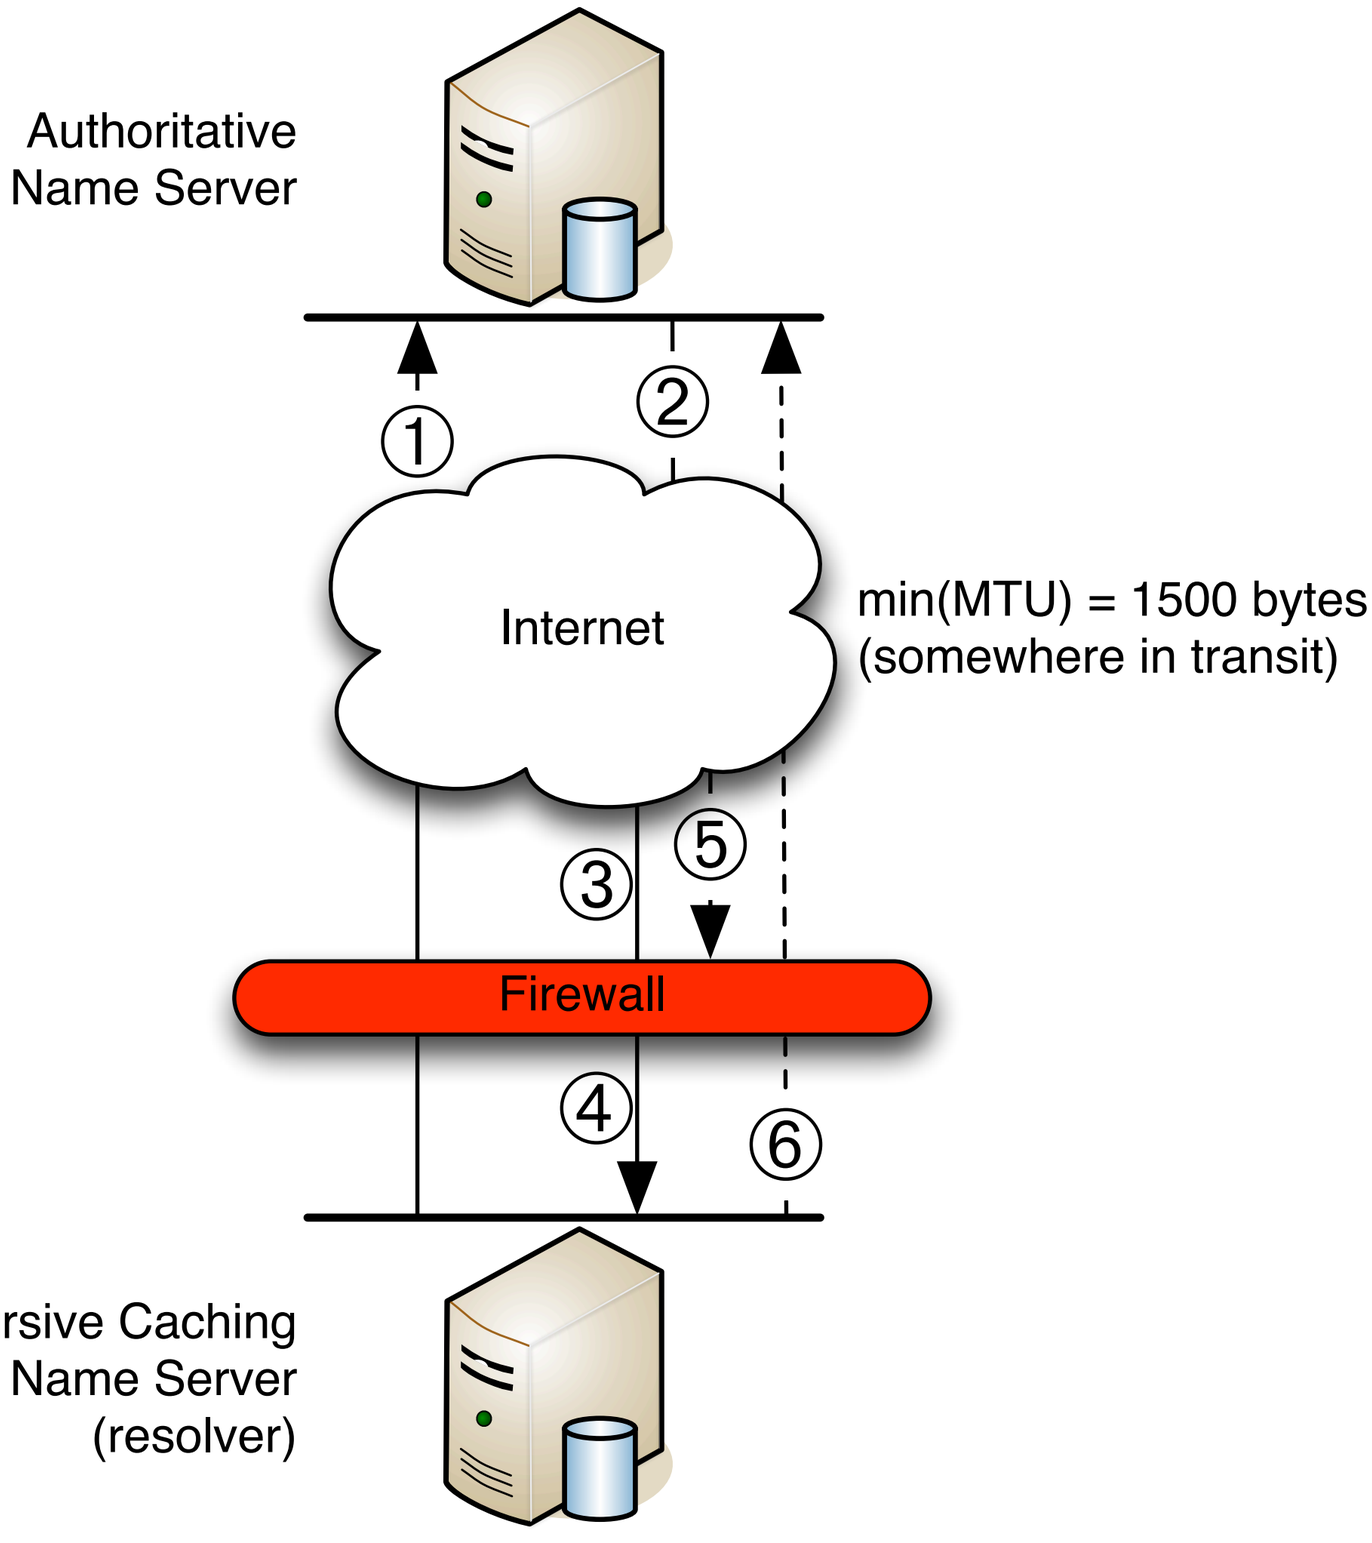

Serious business

- Even though we do everything by the book w.r.t. DNSSEC, and even if people don't validate they still have trouble resolving host names in our zone

- We are a research network, so a few bumps in the road don't scare us

- But think of the big enterprises we are trying to convince to start deploying DNSSEC!

- Also: the ISP was unable/unwilling to change the firewall setting ("It's almost Christmas")

Research at SURFnet

- Short student assignment to confirm the problem http://bit.ly/dnssec-frags

- Student research confirmed: FRTE messages show up when UDP fragments are dropped

- Currently: M.Sc. student working on problem mitigation options and better detection

How big is the problem?

#1 -- EDNS0 use:

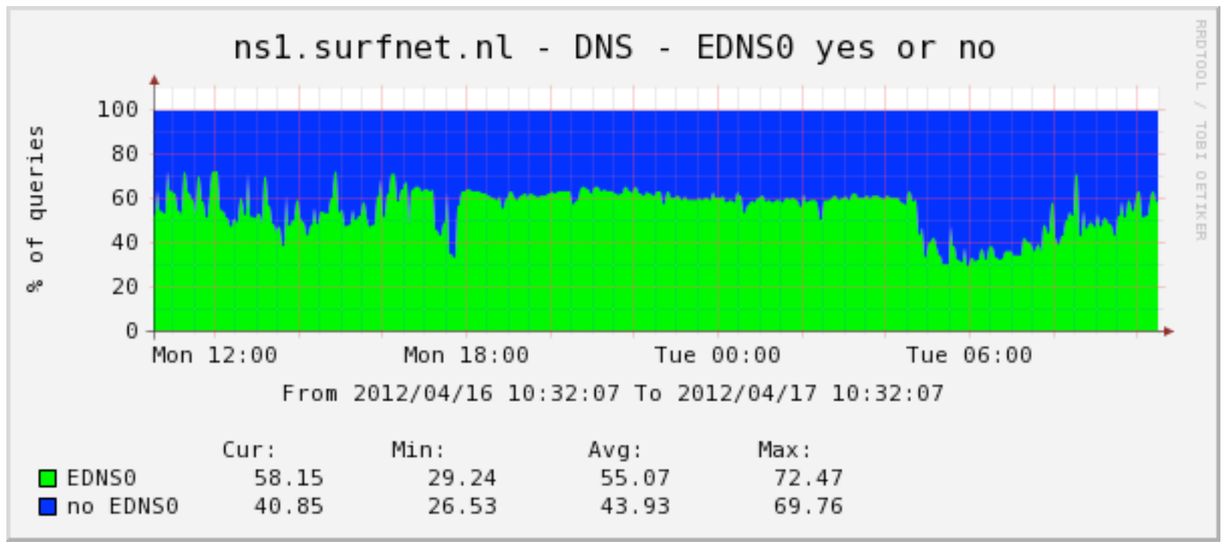

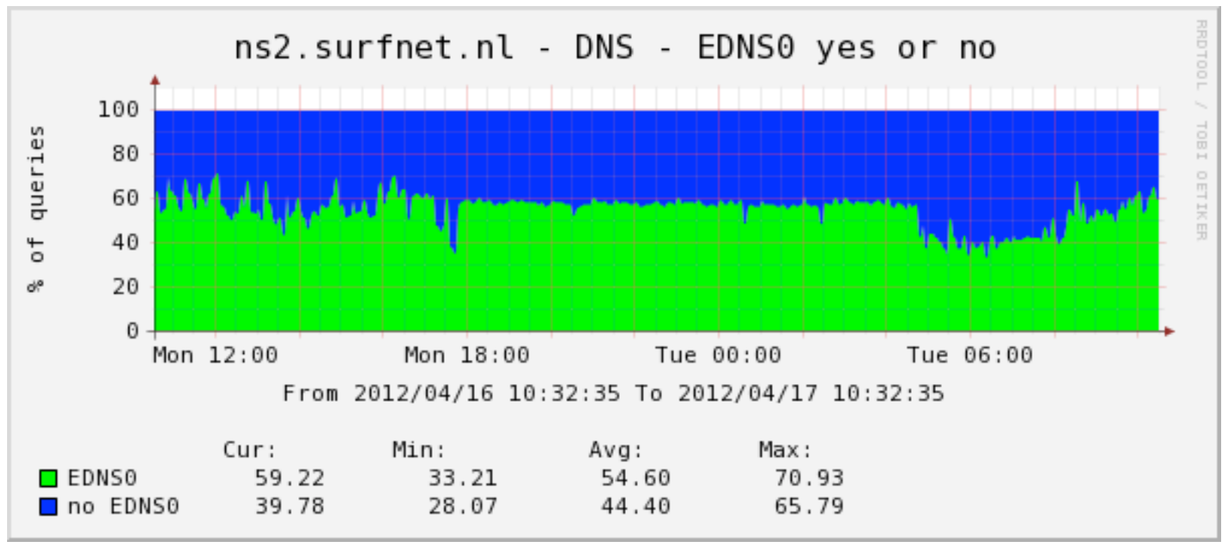

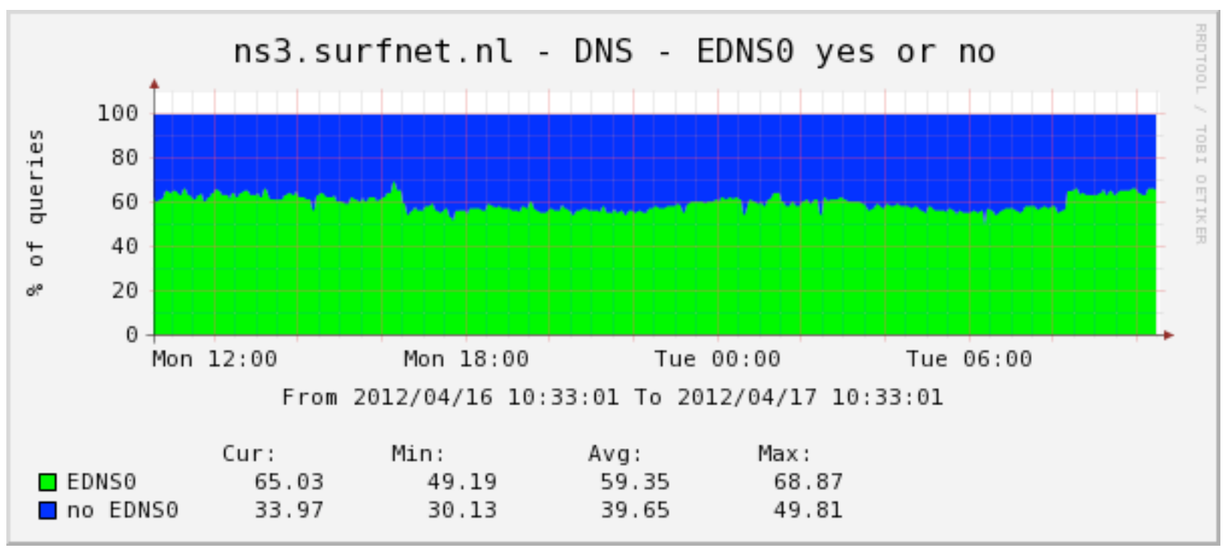

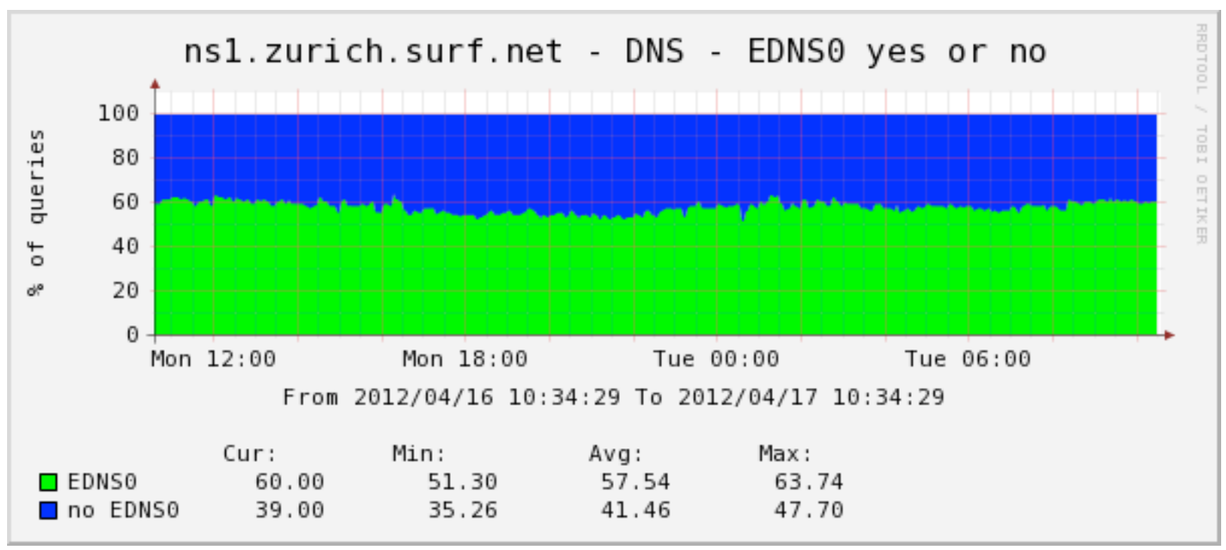

Well over 50% of querying hosts use EDNS0

How big is the problem?

#2 -- EDNS0 advertised buffer size

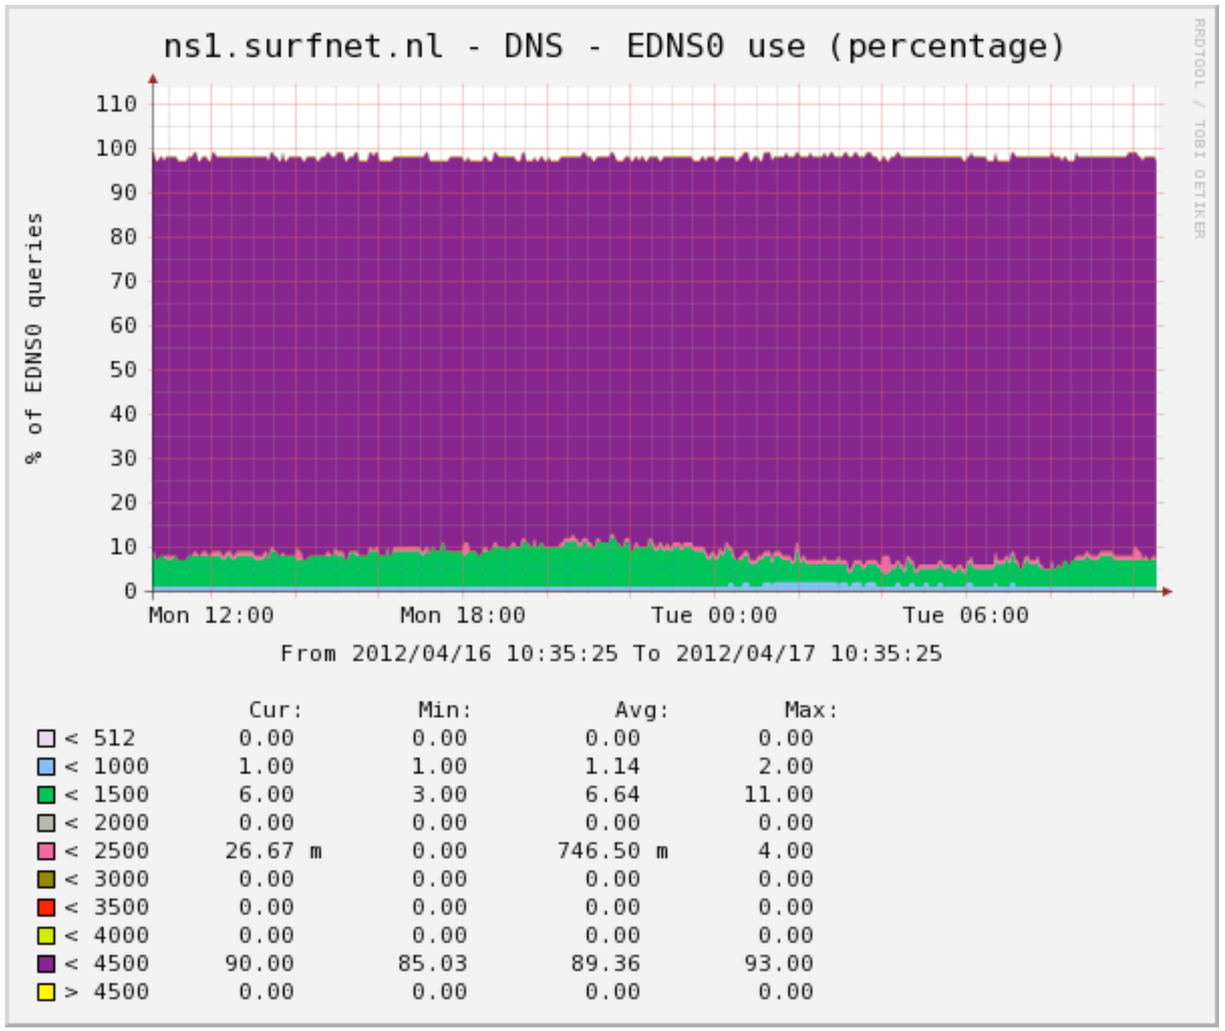

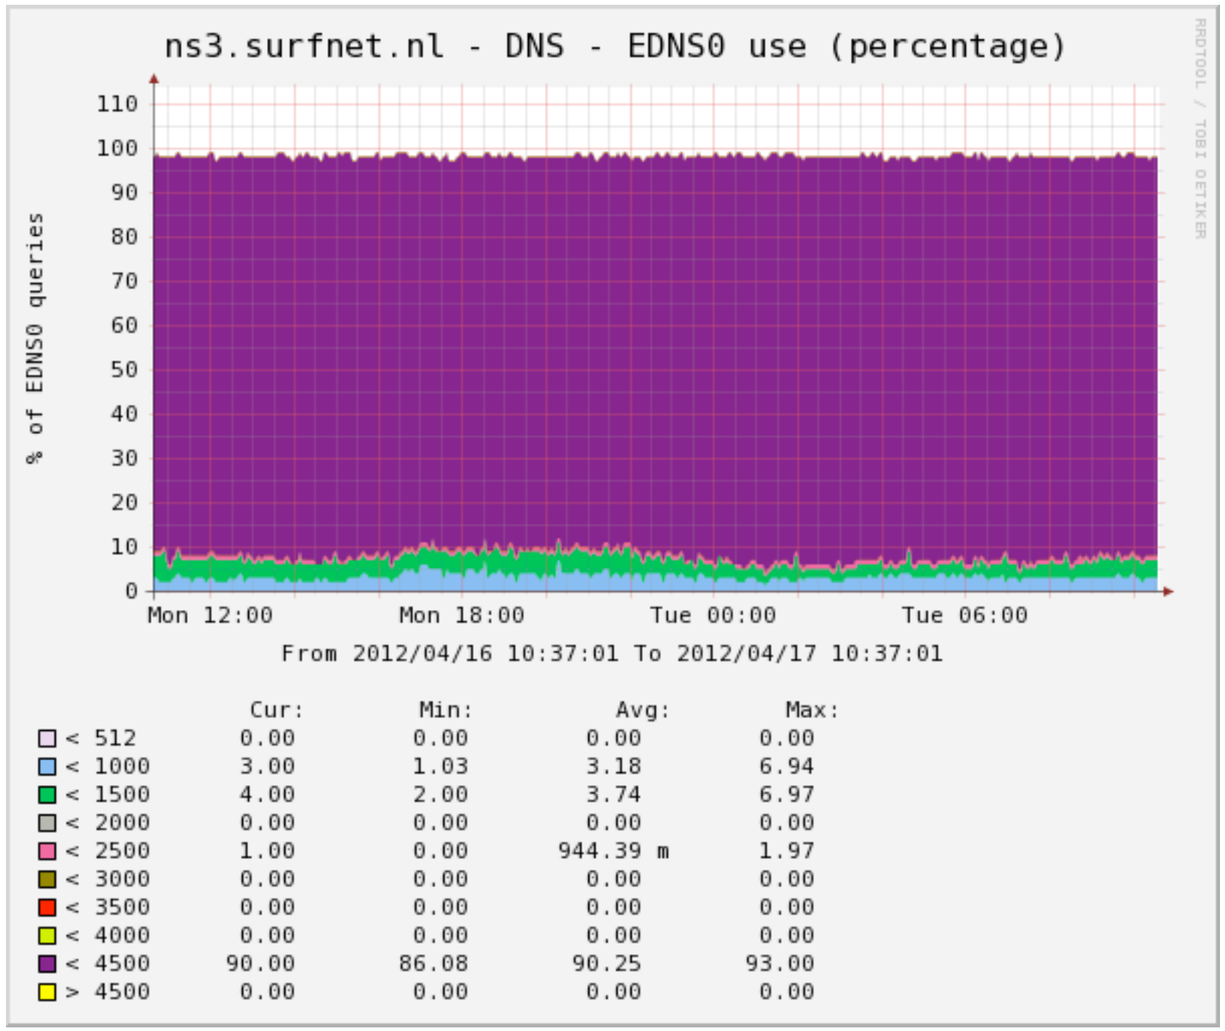

About 90% advertise (default) 4K buffer size

How big is the problem?

#3 -- DNSSEC OK bit set:

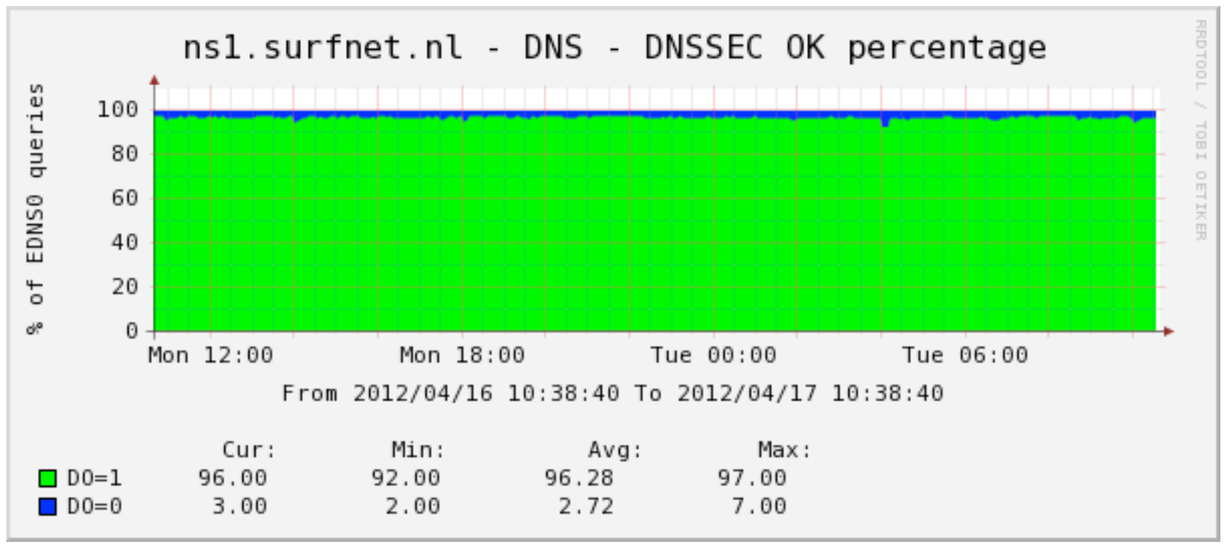

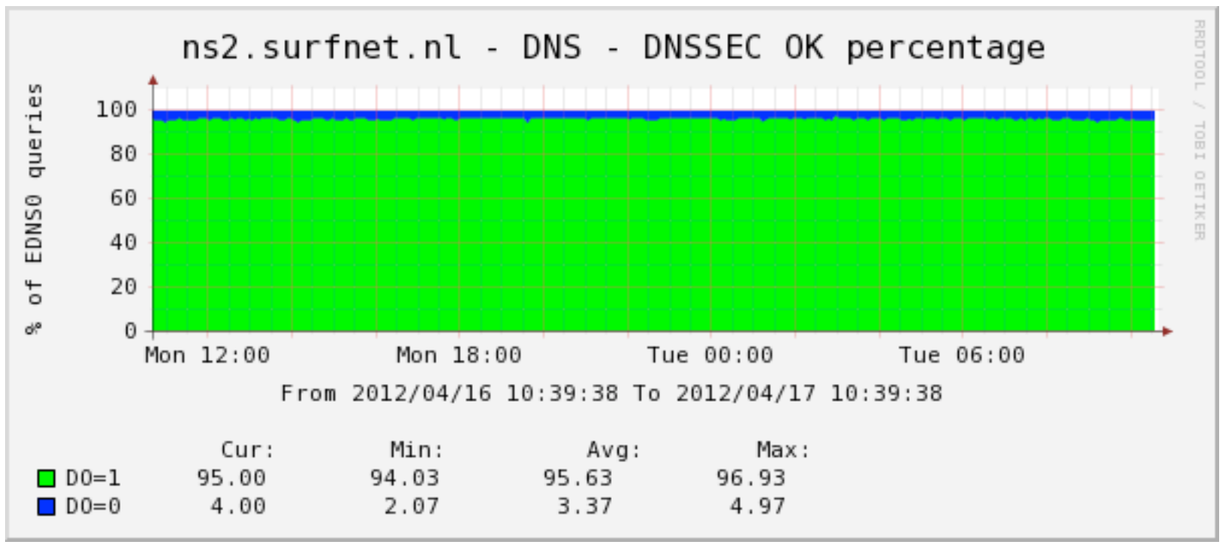

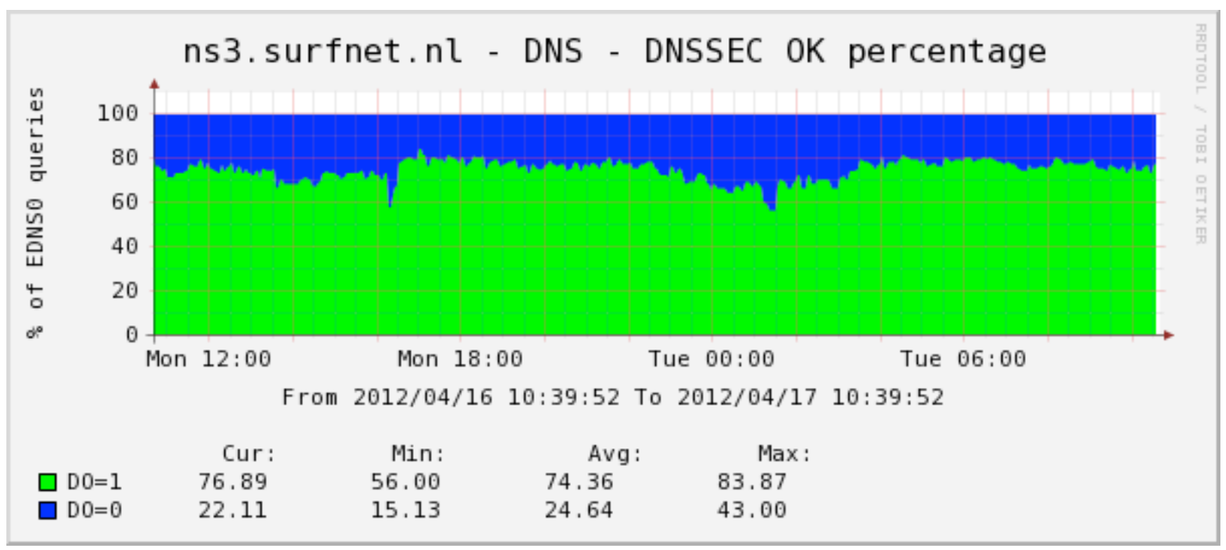

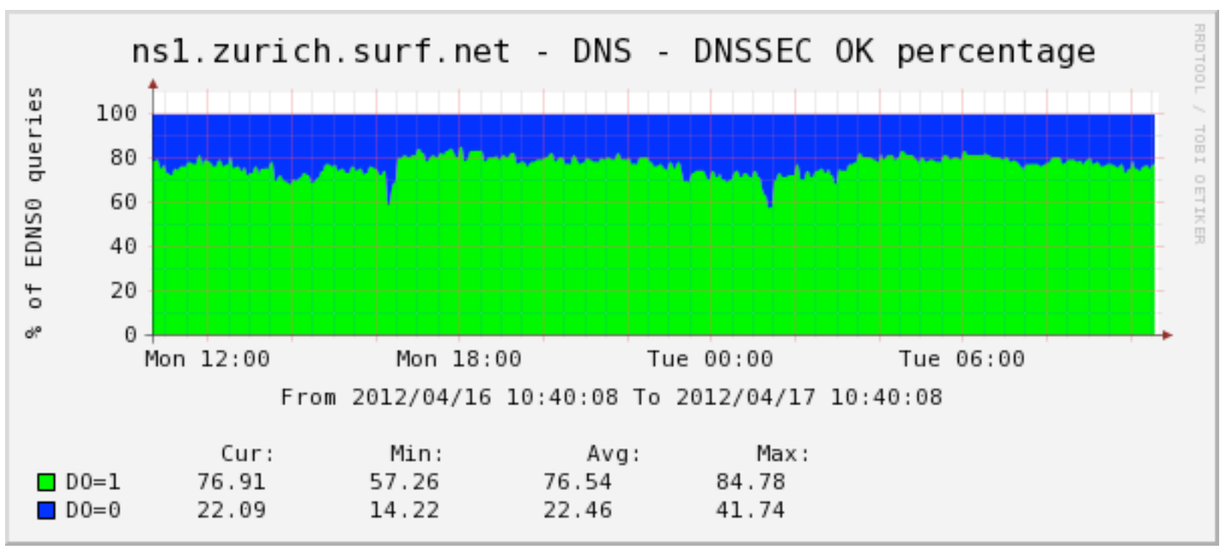

The vast majority sets DO=1

Mitigation approaches

- Two approaches to mitigation

- One: lowering the EDNS0 buffer size on one of the authoritative name servers in the NS set of a domain

- Two: detecting problem hosts with a sensor and adapting name server behaviour (dynamically adjusting EDNS0 buffer size)

Real detection

- ICMP may be blocked by a firewall

- How to detect problem hosts that aren't allowing ICMP through?

- Heuristic approach, 5 rules

#1	ICMP FRTE is seen
#2	EDNS0 header toggled on/off by querying host
#3	(Excessive) retries within TTL of record
#4	Changing EDNS0 buffer size in queries
#5	Fallback to TCP without truncation

Experiments

- Experiment #1:

  Lowering the EDNS0 buffer size on one authoritative name server to 1232 bytes, so below IPv6 minimum MTU

- Experiment #2:

  Selectively modify advertised EDNS0 buffer size in queries originating from “problem” hosts before they reach the name server

Problem hosts detected

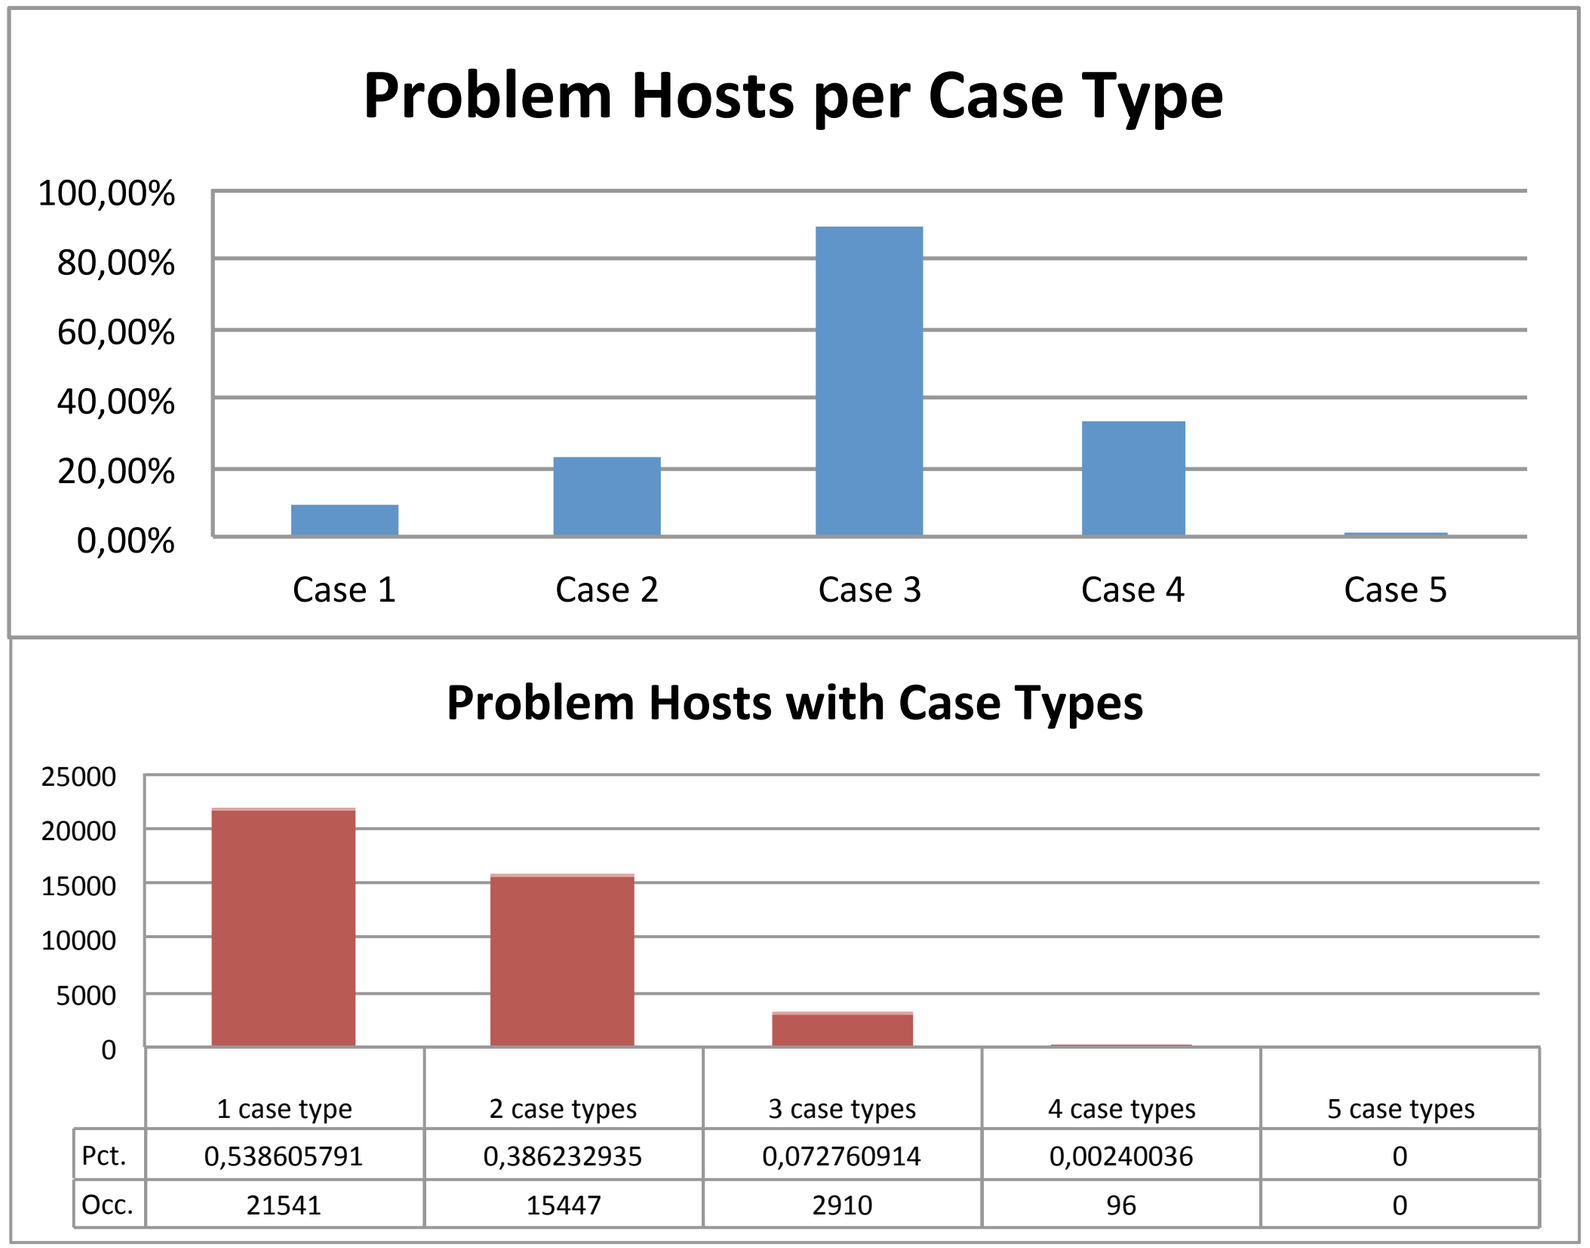

Analysis shows: ≥2% confirmed problem host

ICMP FRTE behaviour

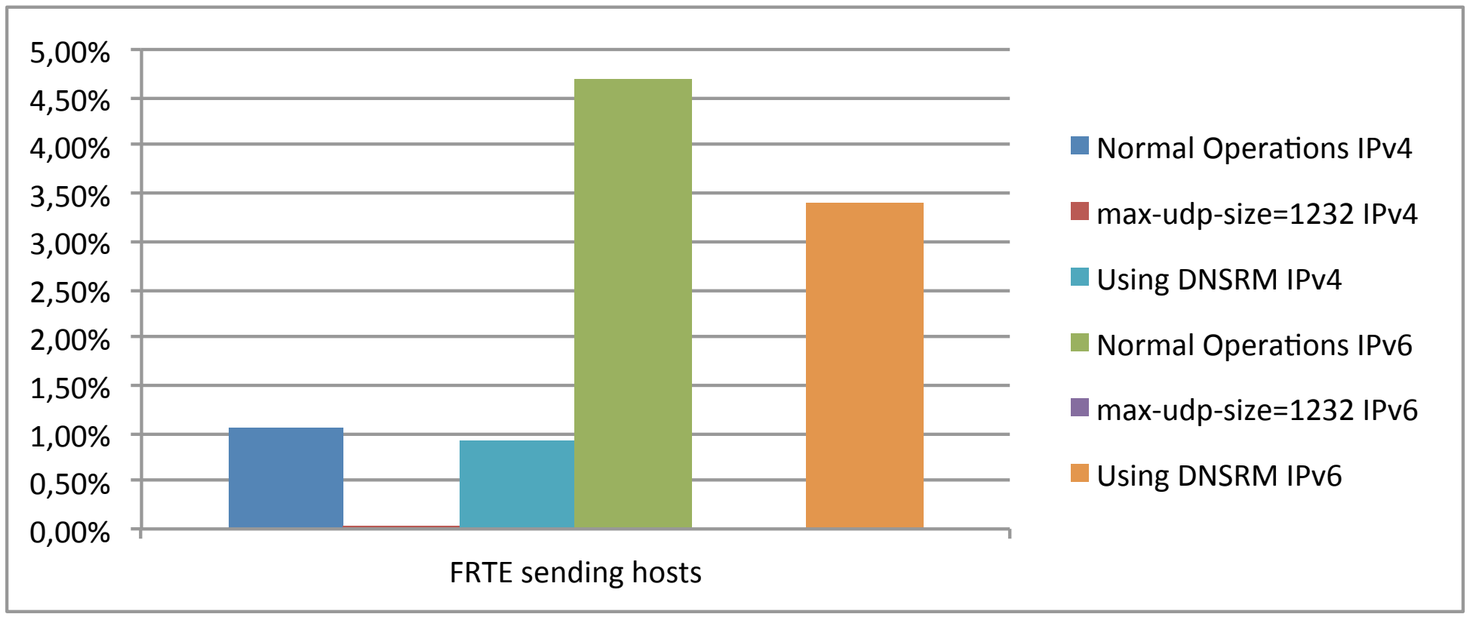

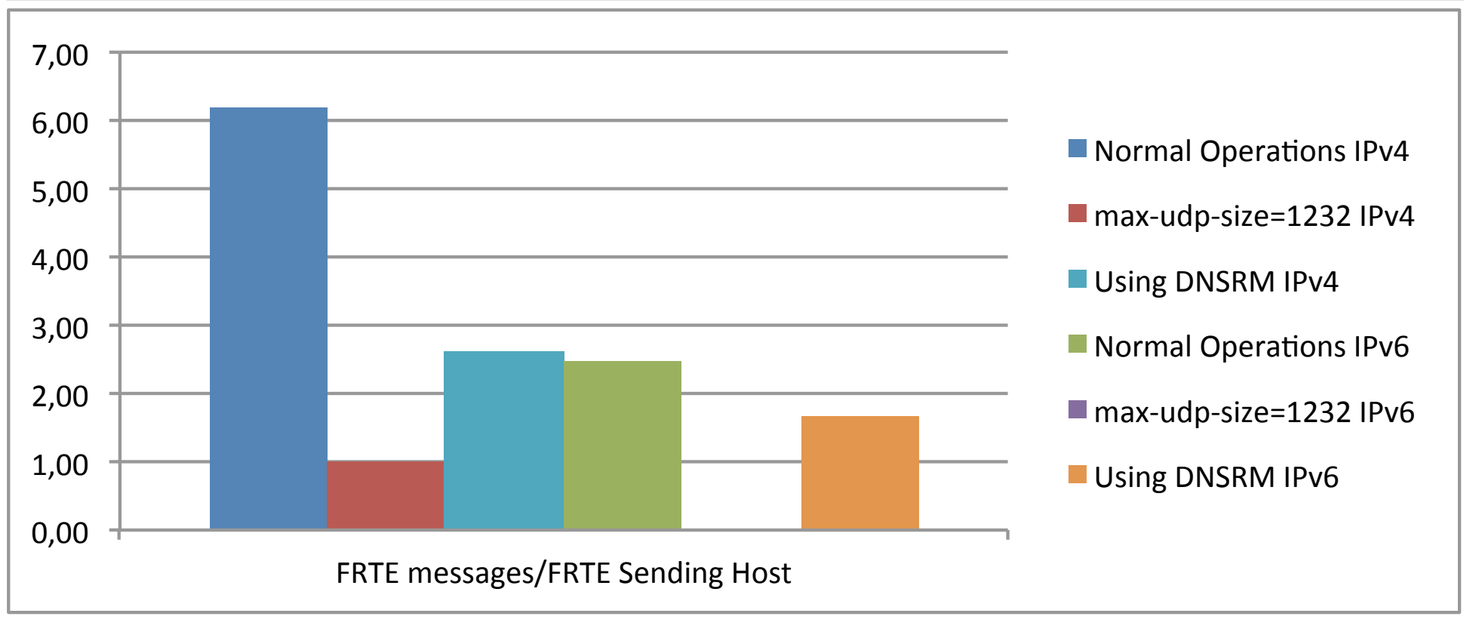

Bottom line: both approaches tackle the problem

Some side-effects

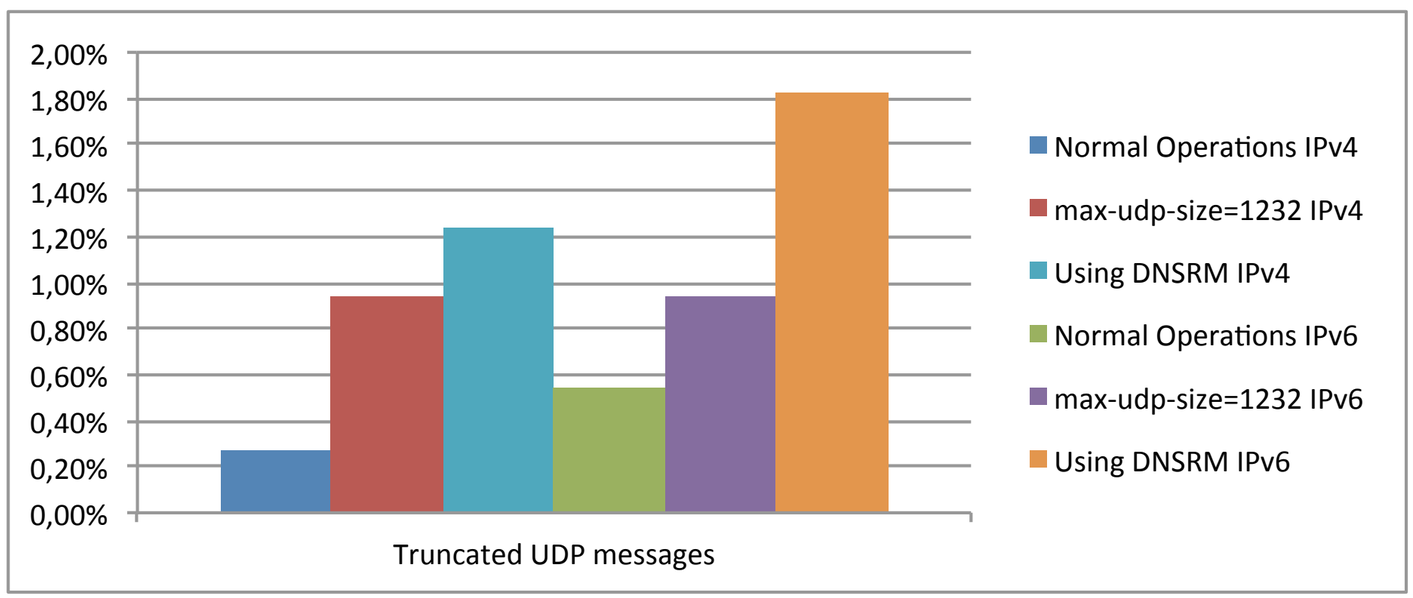

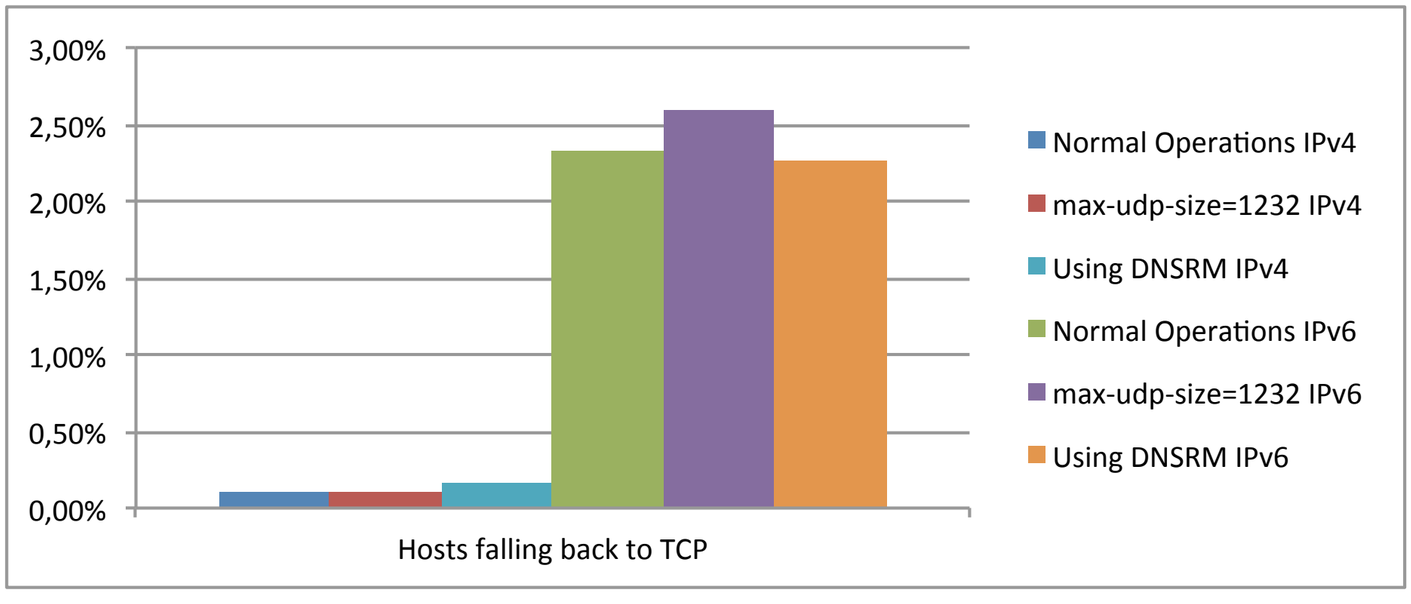

Note: long bars, but very low percentages

Conclusion

- This seems to be a serious issue for DNSSEC-signed zones

- There are ways to ameliorate the problem

- We are considering writing a best-practice paper (or even an informational RFC)

- Expect a paper in IEEE CC Review or ACM Transactions on Networking

- Check your firewall settings if you start doing DNSSEC validation on your resolvers!

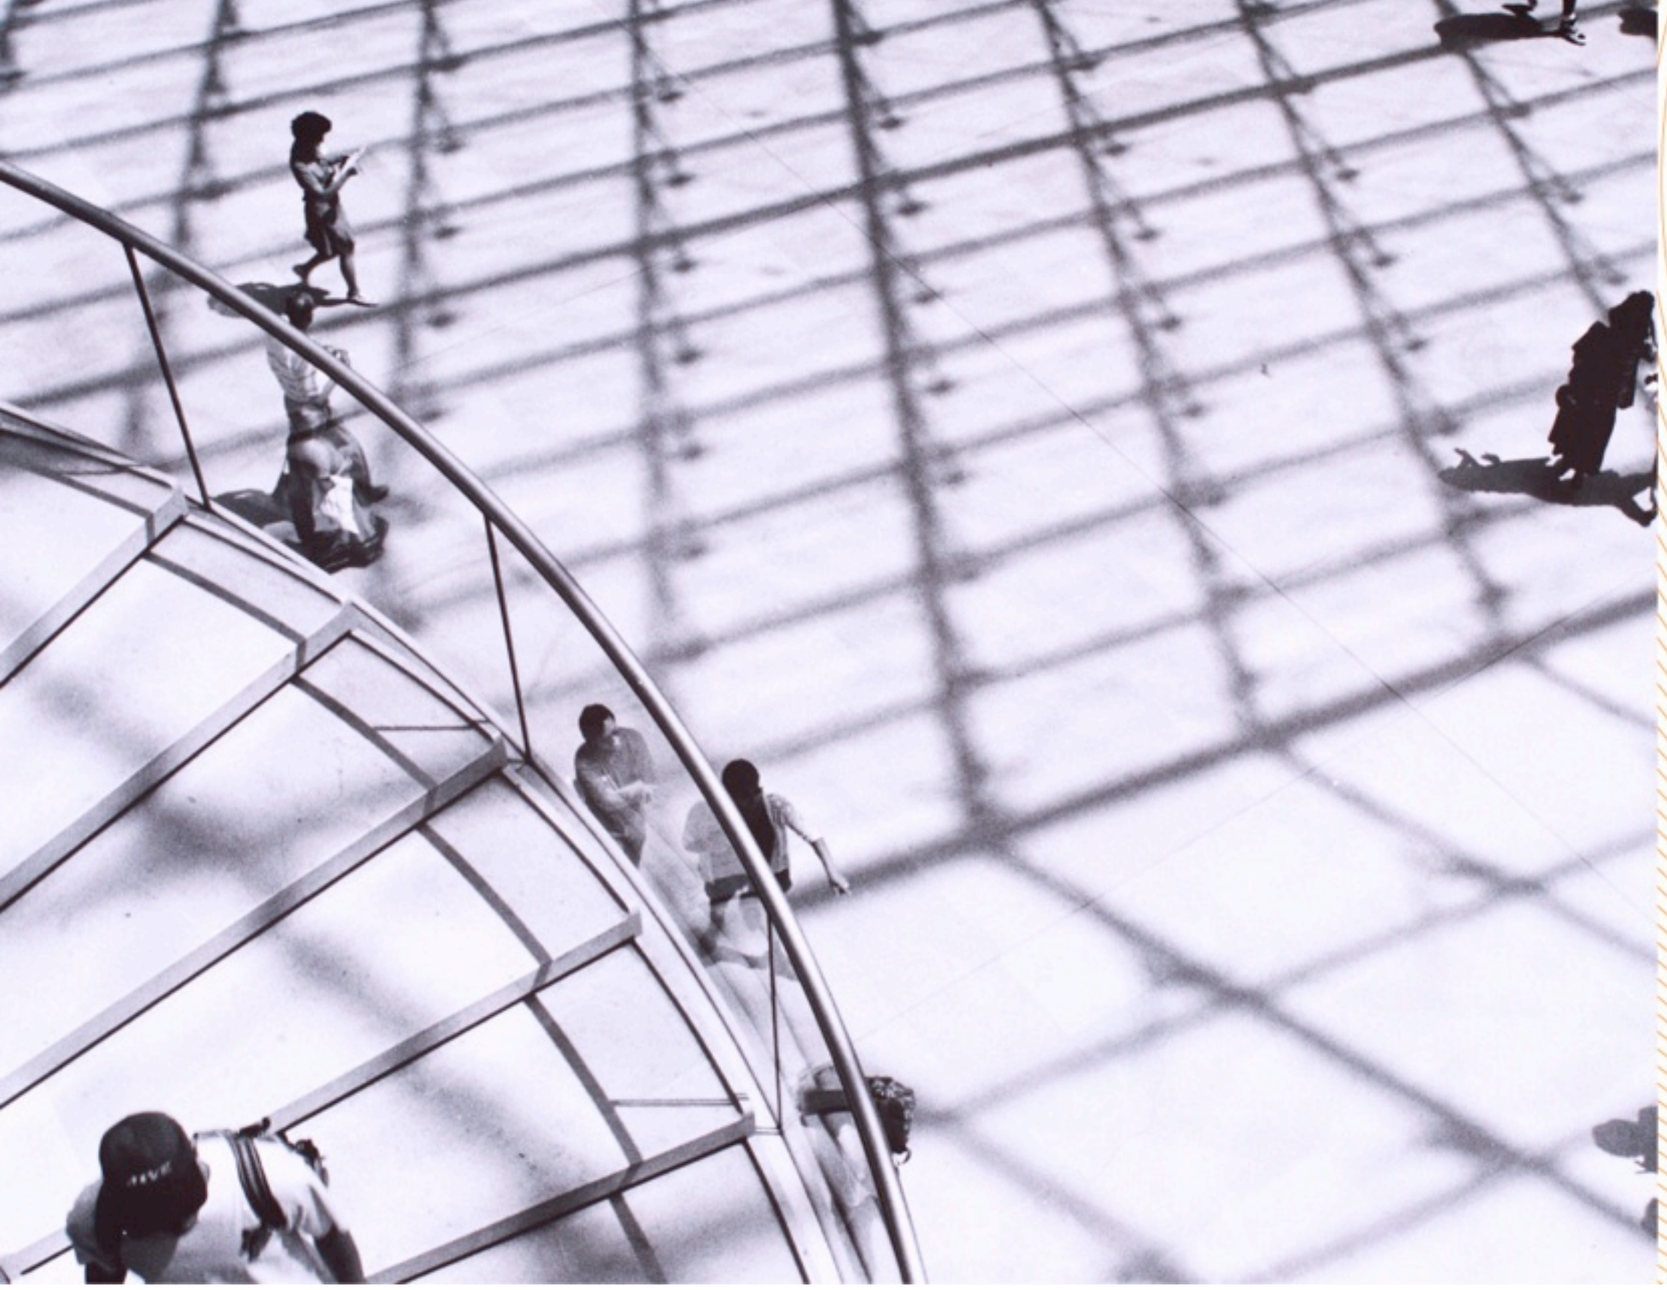

SURF NET

Questions? Comments?

Please contact me!

roland.vanrijswijk@surfnet.nl

nl.linkedin.com/in/rolandvanrijswijk

@reseauxsansfil