Improving Response Deliverability in DNS(SEC)

Gijs van den Broek[*‡], Roland van Rijswijk[‡], Aiko Pras*, Anna Sperotto*

*University of Twente, Enschede, The Netherlands (j.g.vandenbroek@student.utwente.nl, a.pras@utwente.nl, a.sperotto@utwente.nl)

[‡]SURFnet bv, Utrecht, The Netherlands (roland.vanrijswijk@surfnet.nl)

The Domain Name System provides a critical service on the Internet, where it allows host names to be translated to IP addresses. However, it does not provide any guarantees about authenticity and origin integrity of resolution data. DNSSEC attempts to solve this through the application of cryptographic signatures to DNS records. These signatures generally result in larger responses compared to plain DNS responses. Some of these larger responses experience fragmentation, which in turn might be partially blocked by some firewalls. Apparently unresolvable zones may in those cases be a consequence. Analysis of DNS traffic suggests that at least one per cent of all resolvers experience this problem with our signed zones. However, we suspect this number to be much larger. In our presentation we will elaborate on the potential extent of this problem and propose to test two solutions. We intend to test both solutions in our production environment.

Problem Overview

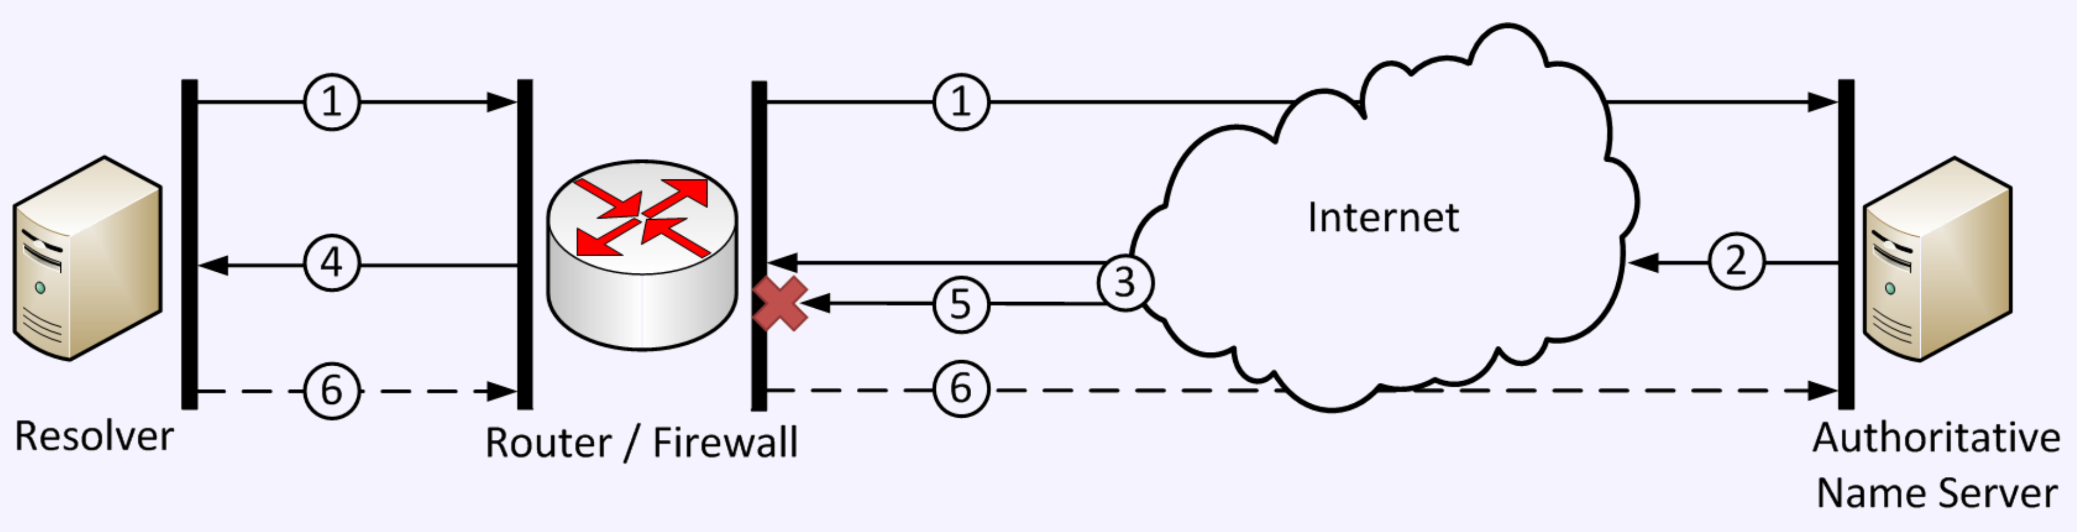

	Description
1	DNS query sent to authoritative name server
2	DNS response returned
3	DNS response fragmented into IP fragments due to lower MTU
4	First IP fragment of DNS response arrives at resolver
5	Second IP fragment of DNS response is blocked at firewall
6	An ICMP Fragment Reassembly Time Exceeded message is sent 30 seconds later

Table 1: Problem resolver not able to obtain a fragmented DNS response.

DNS Traffic Analysis

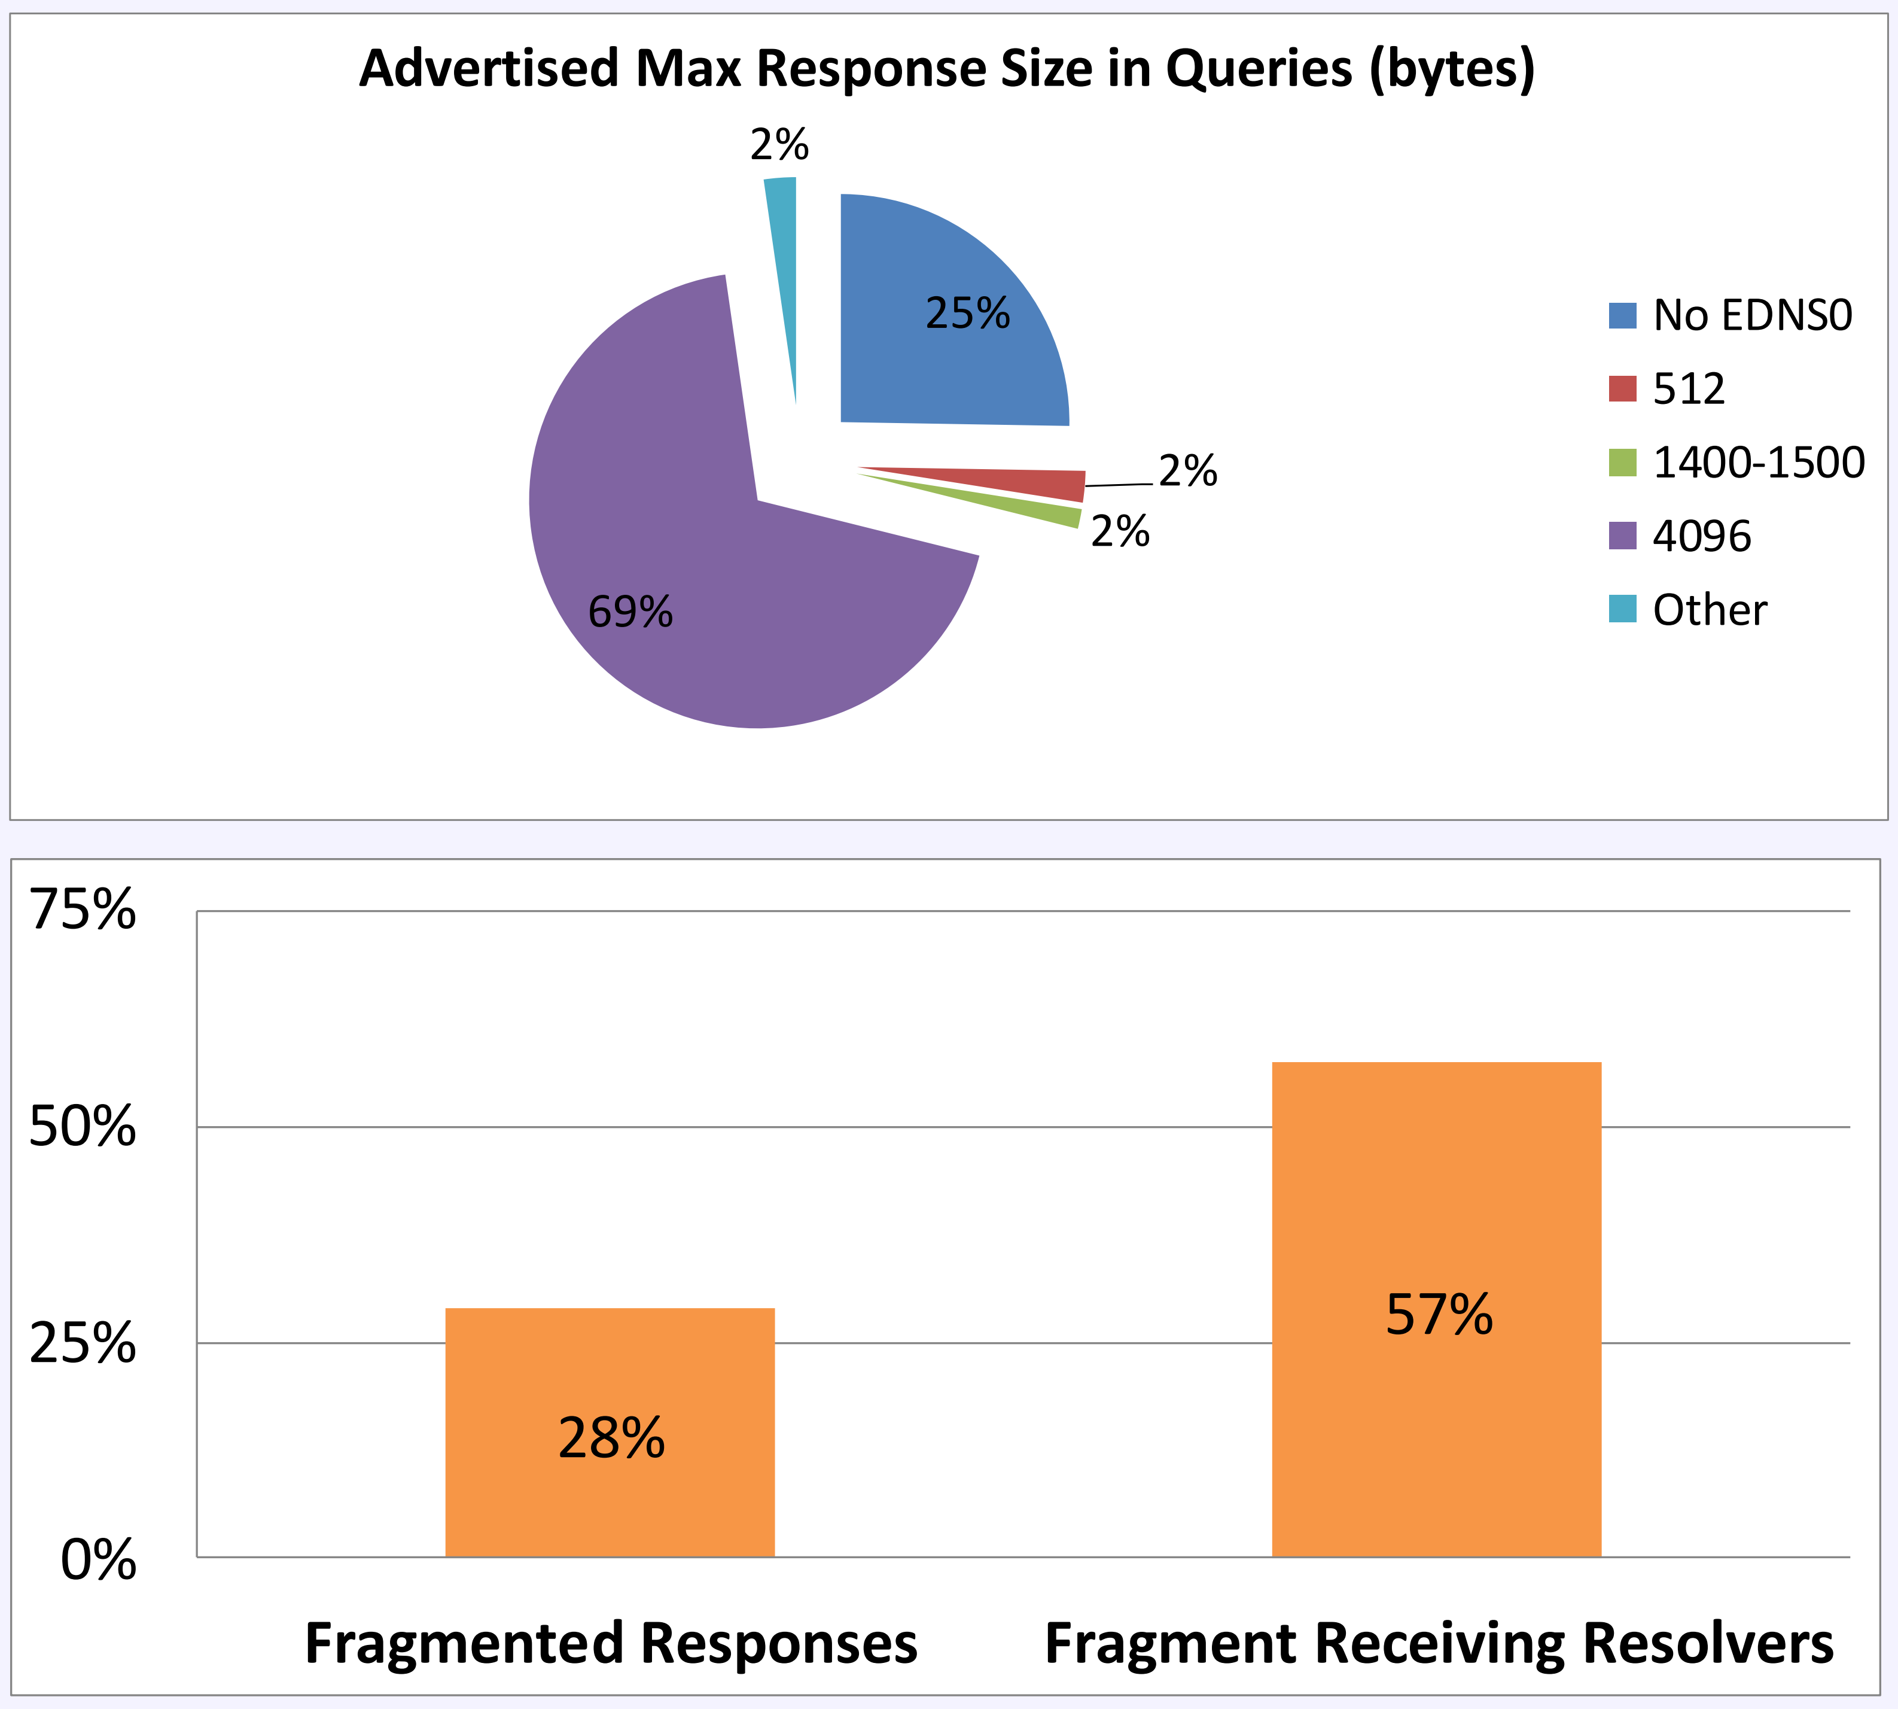

Figure 1, 2: Distribution of max. DNS response size in queries, percentage of all UDP DNS responses being fragmented and percentage of all resolvers receiving fragments (at ns3.surfnet.nl, sample 8.4 mln. messages).

Resolver Behavioural Characteristics	Unique Resolvers
Case 1: Sending ICMP Fragment Reassembly Time Exceeded (FRTE)	1.1%
Case 2: Removal of EDNS0 header in retries	2.4%
Case 3: Retries for large (>512 bytes) responses exceed 4%	9.7%
Case 4: Reduced advertised buffer size in retries	3.5%
Case 5: TCP fallback w/o truncated UDP response preceding it	<0.1%

Table 2: Resolver characteristics detected at ns1, ns2 and ns3.surfnet.nl. This involved 386,648 unique IPv4 and IPv6 resolvers and 40.5 mln messages.

Experiment #1

We will use three authoritative name servers part of the surfnet.nl zone for both experiments. All of these name servers have the same configuration, unless stated otherwise.

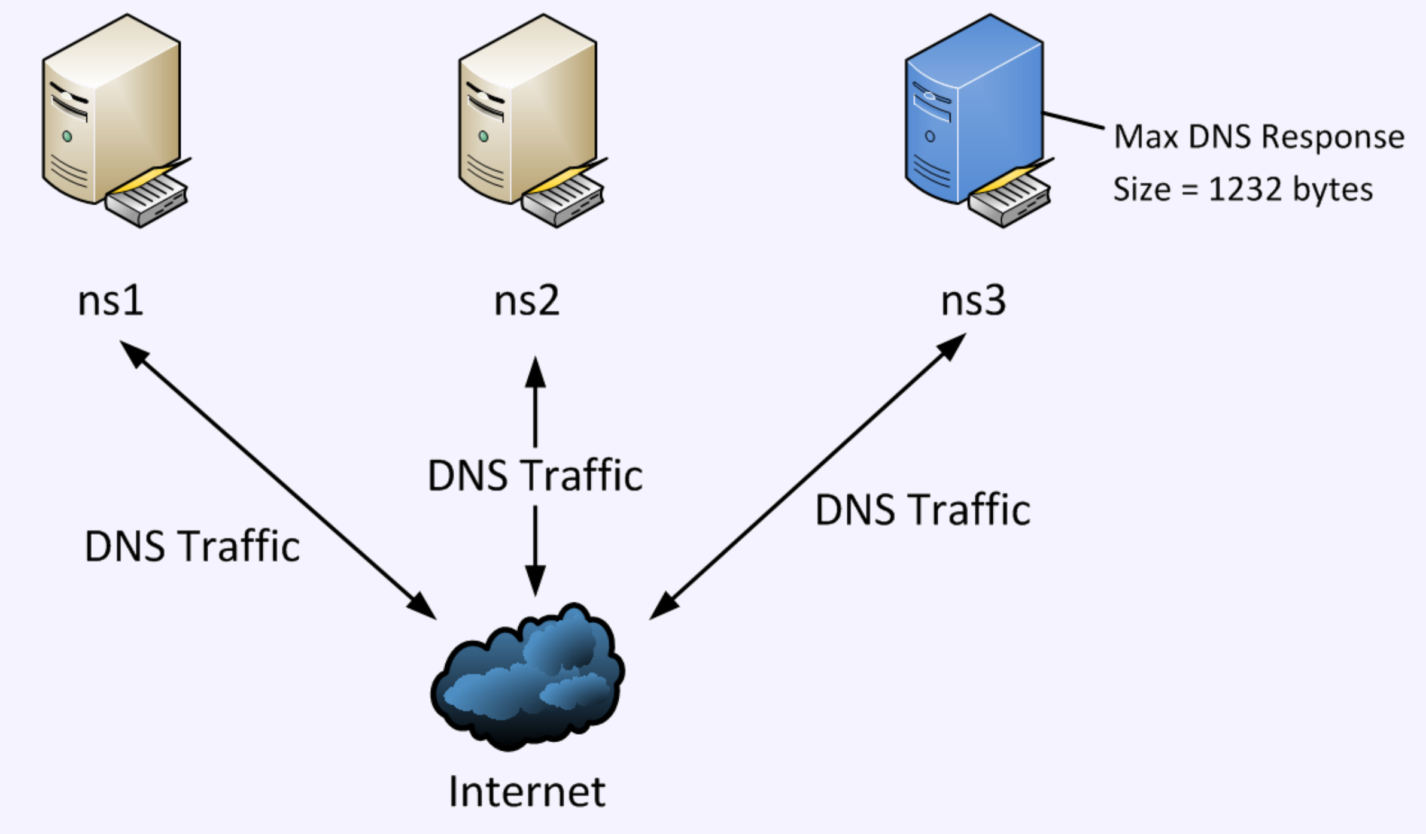

- Name server ns3 returns responses with a maximum size of: *min*('max. size advertised in query', 1232) bytes.

Experiment #2

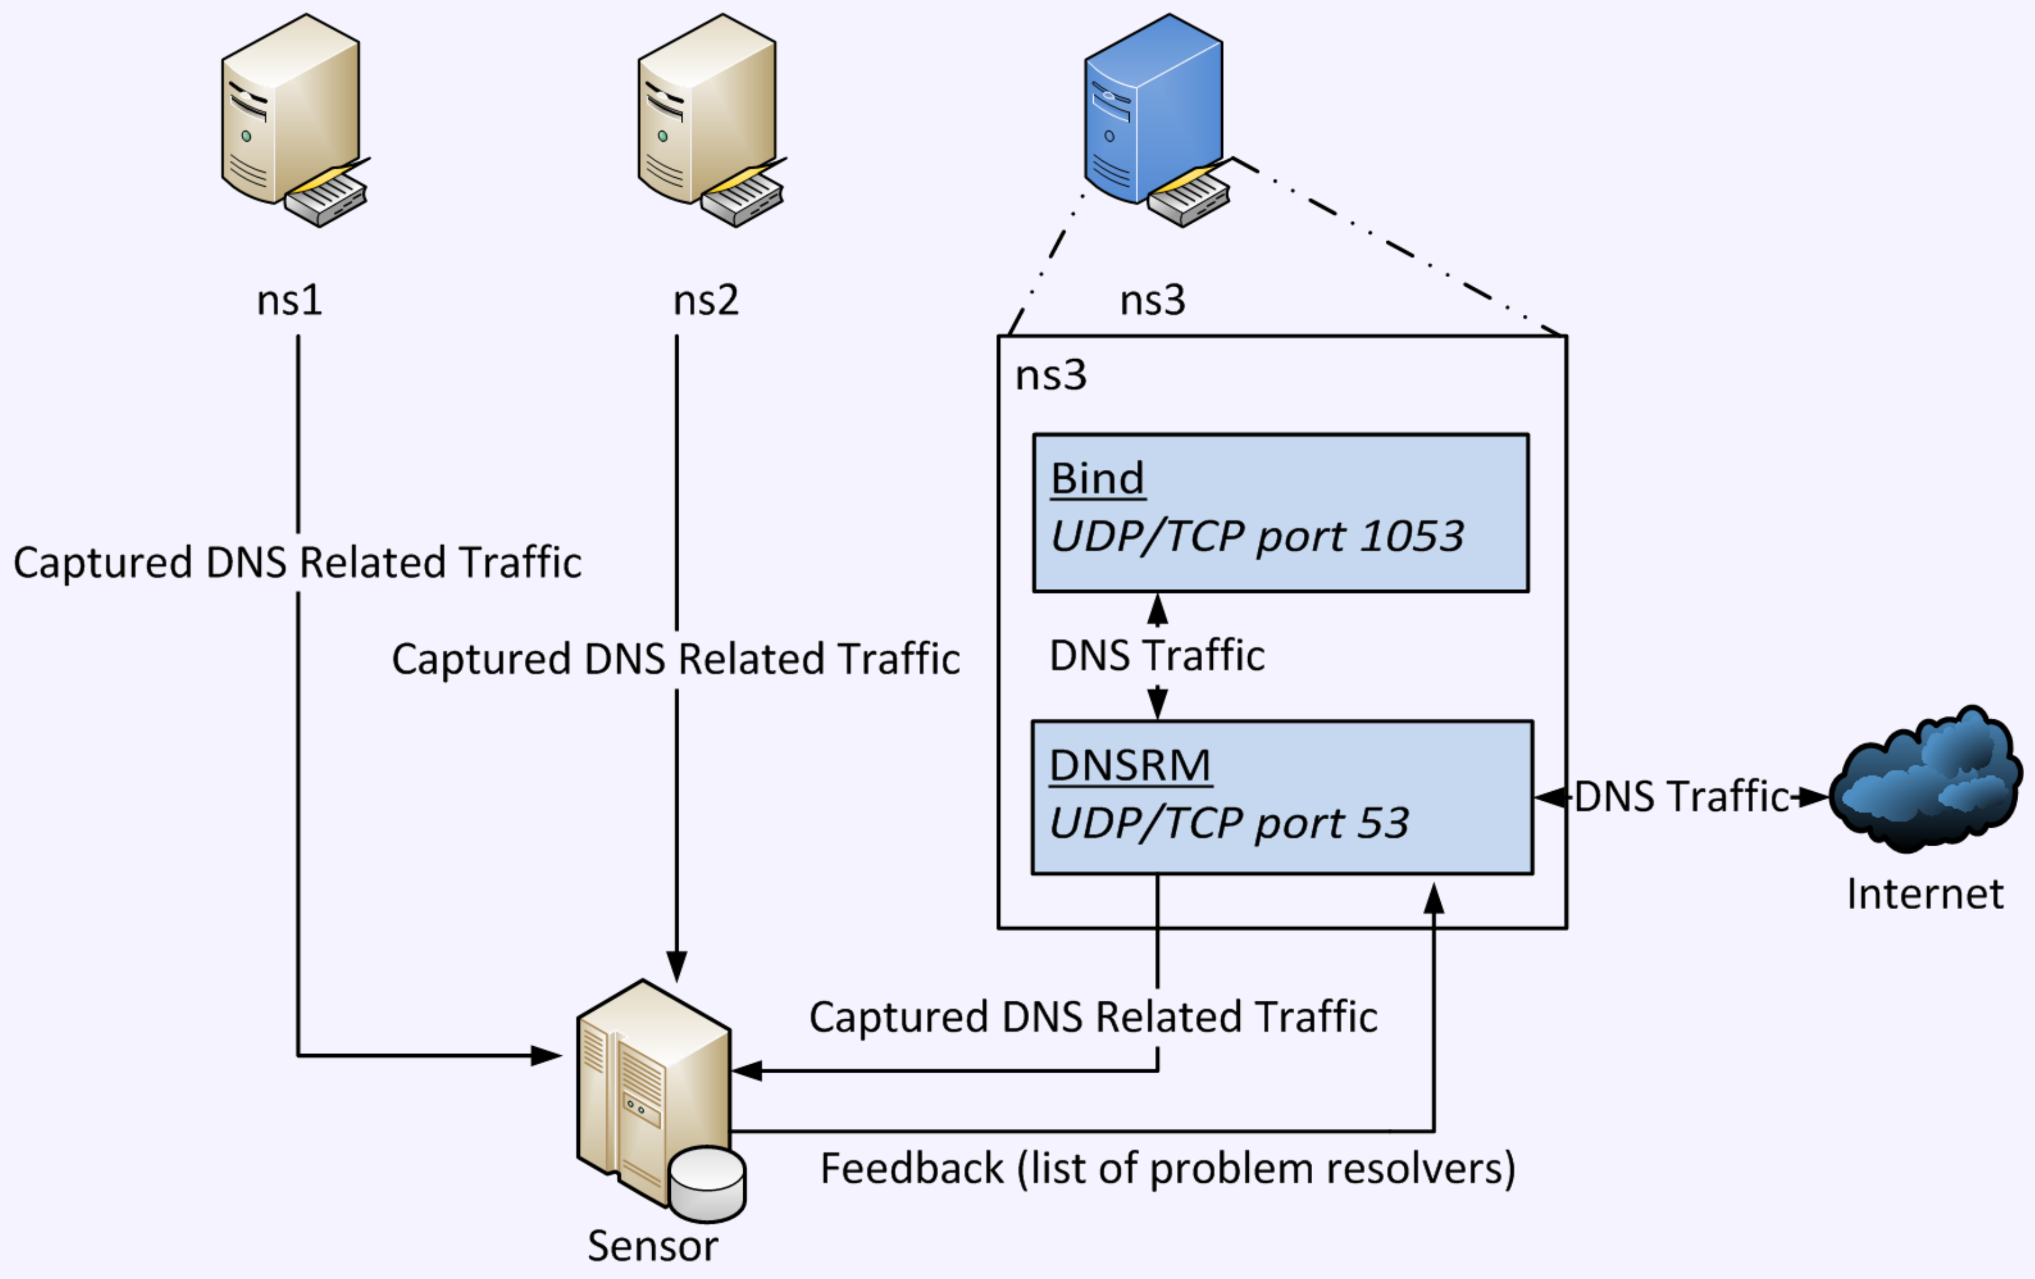
ns1
ns2
ns3
ns3
Bind
UDP/TCP port 1053
DNS Traffic
DNSRM
UDP/TCP port 53
DNS Traffic
Internet
Captured DNS Related Traffic
Captured DNS Related Traffic
Captured DNS Related Traffic
Feedback (list of problem resolvers)
Sensor

- Name server ns3 runs DNSRM on port 53, which forwards inbound DNS traffic to a local Bind process on port 1053. It also sends a copy of the live traffic feed to the sensor.
- Name servers ns1 and ns2 merely forward a copy of their live traffic feed to the sensor to improve the detection of problem resolvers.
- The sensor analyses this traffic in order to determine problem resolvers. It returns a list of detected problem resolvers to the DNSRM process on ns3.
- DNSRM alters advertised maximum response size advertisements for problem resolvers only to: *min*('max. size advertised in query', 'detected maximum for this resolver').

Comparison

Experiment 1 affects every resolver querying authoritative name server ns3, while experiment 2 involves the detection of problem resolvers and manipulating only those queries from such problem resolvers. Experiment 1 is very simple compared to the elaborate setup required for experiment 2, but it may be considered bad netizenship. It is questionable if the results of experiment 2 outweigh the results and simplicity of experiment 1.

SURF NET

Presentation Wednesday 23/05
14:00 - 15:30 Hall 3

UNIVERSITY OF TWENTE.