DNSSEC & fragmentation
a prickly combination

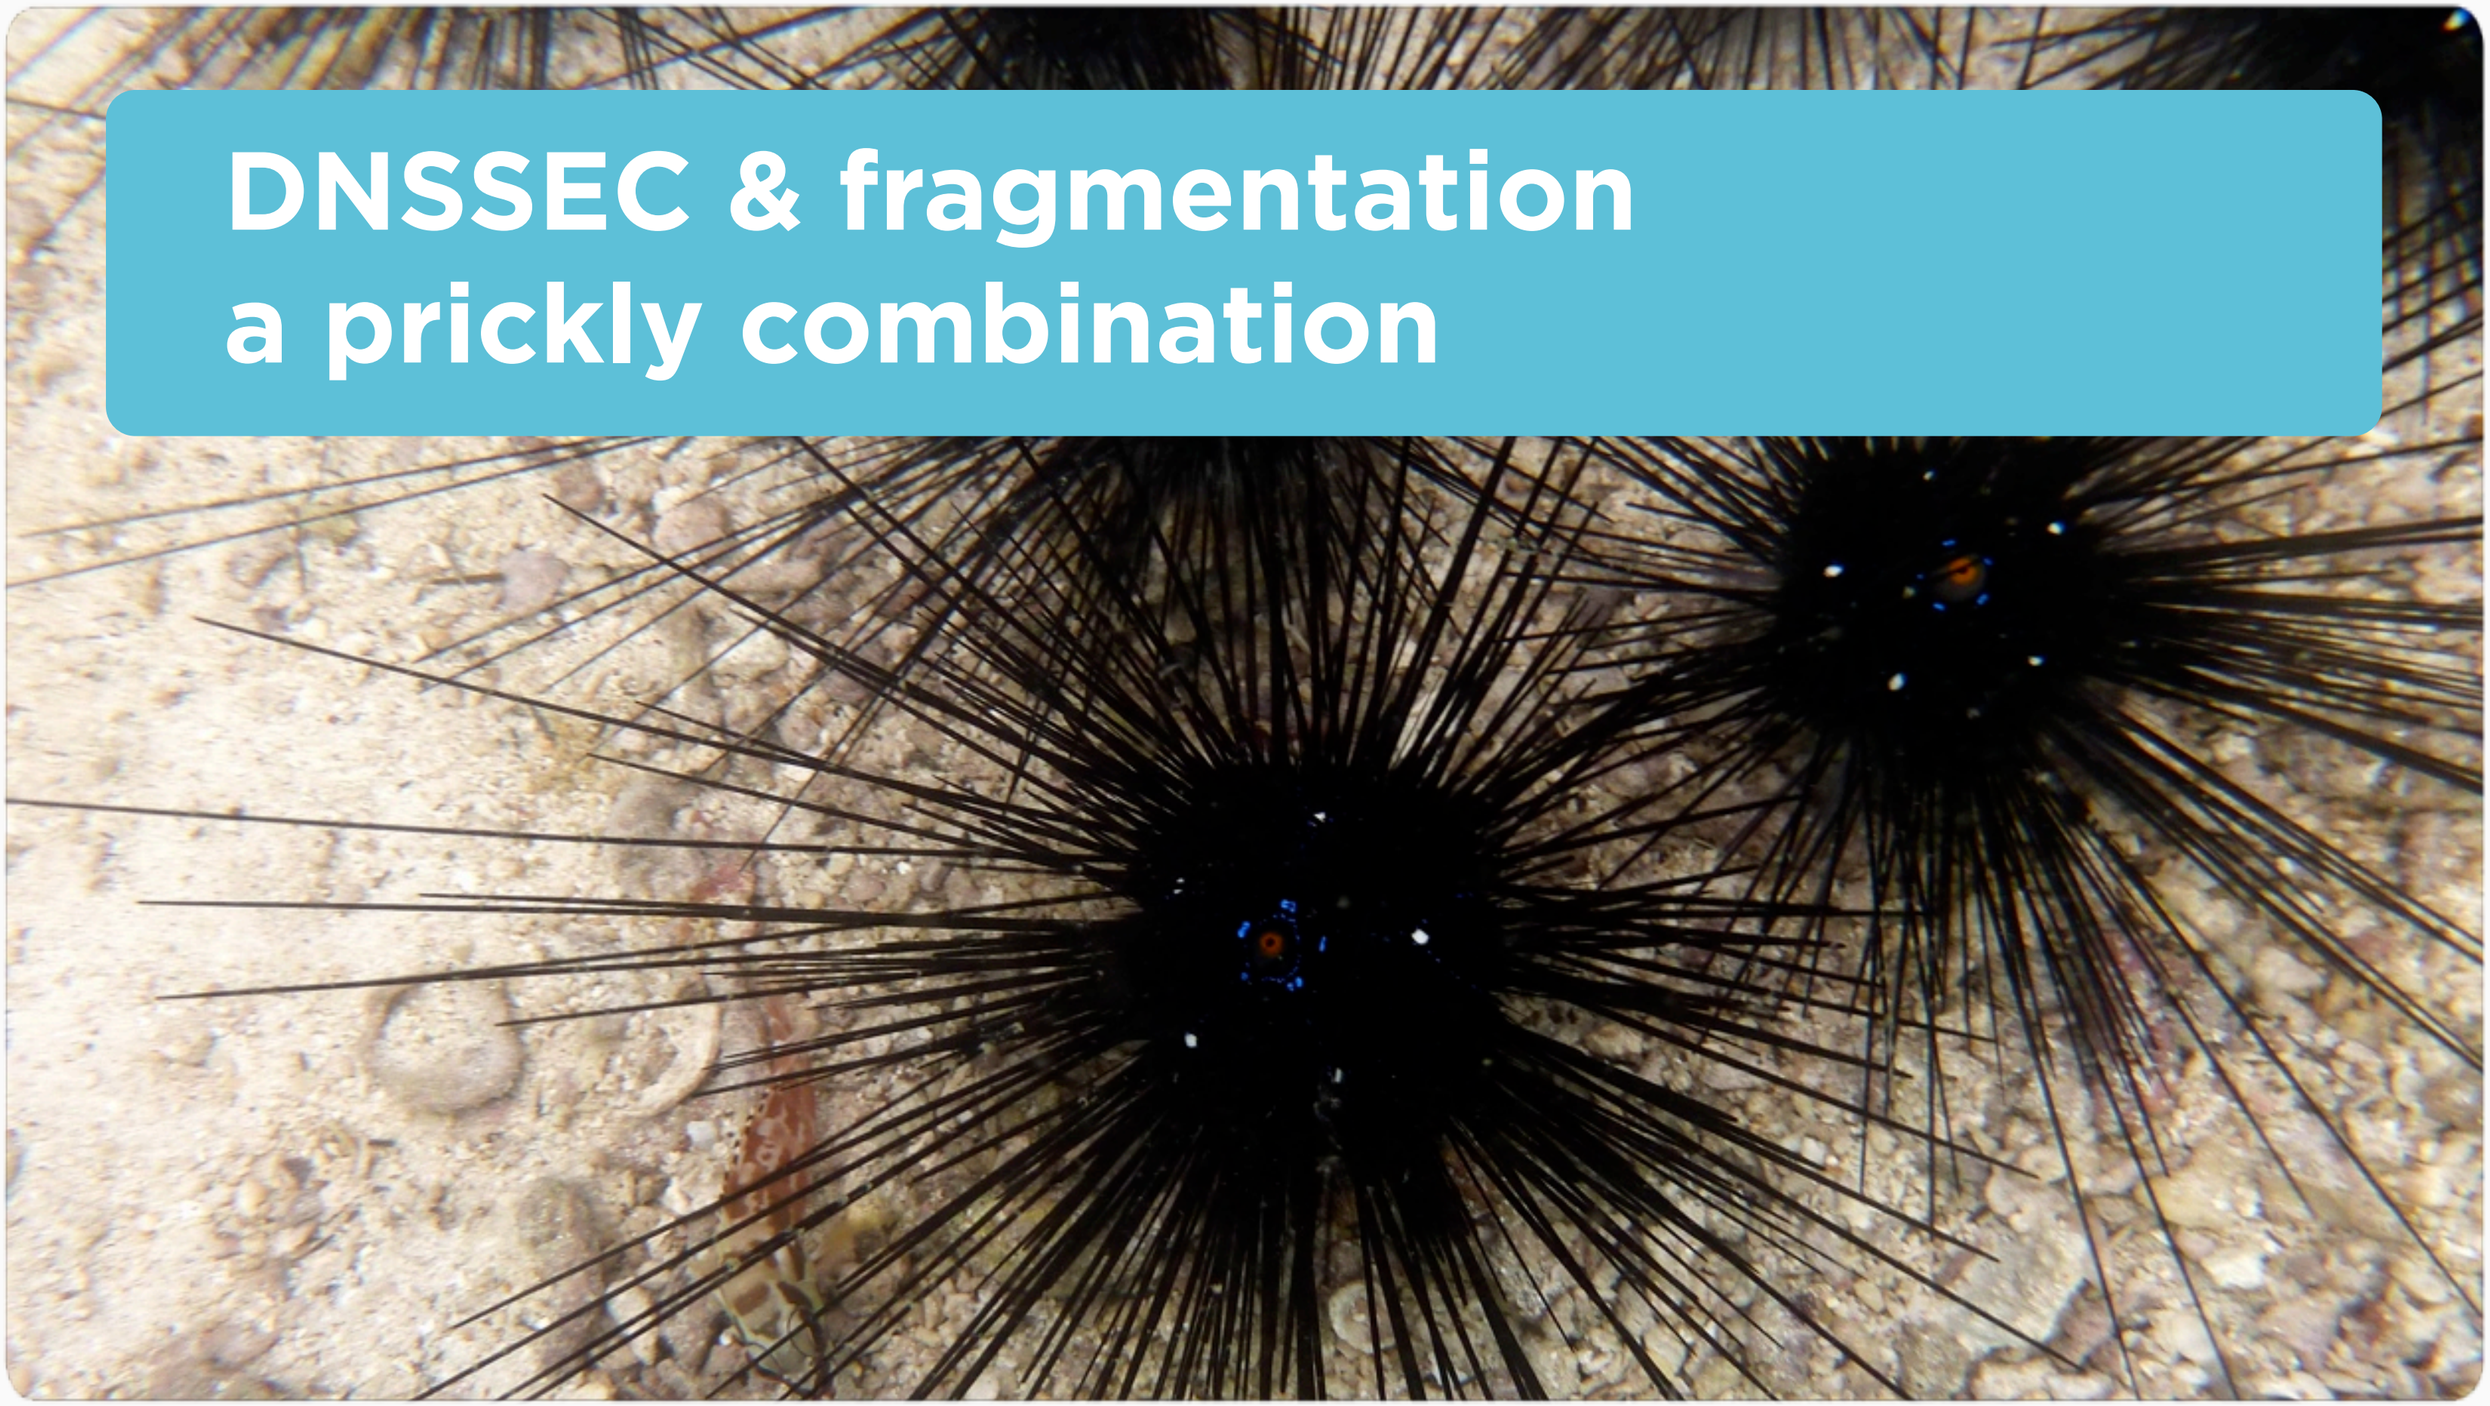

Roland van Rijswijk - Deij

roland.vanrijswijk@surfnet.nl

SURFNET

The problem in 1 slide

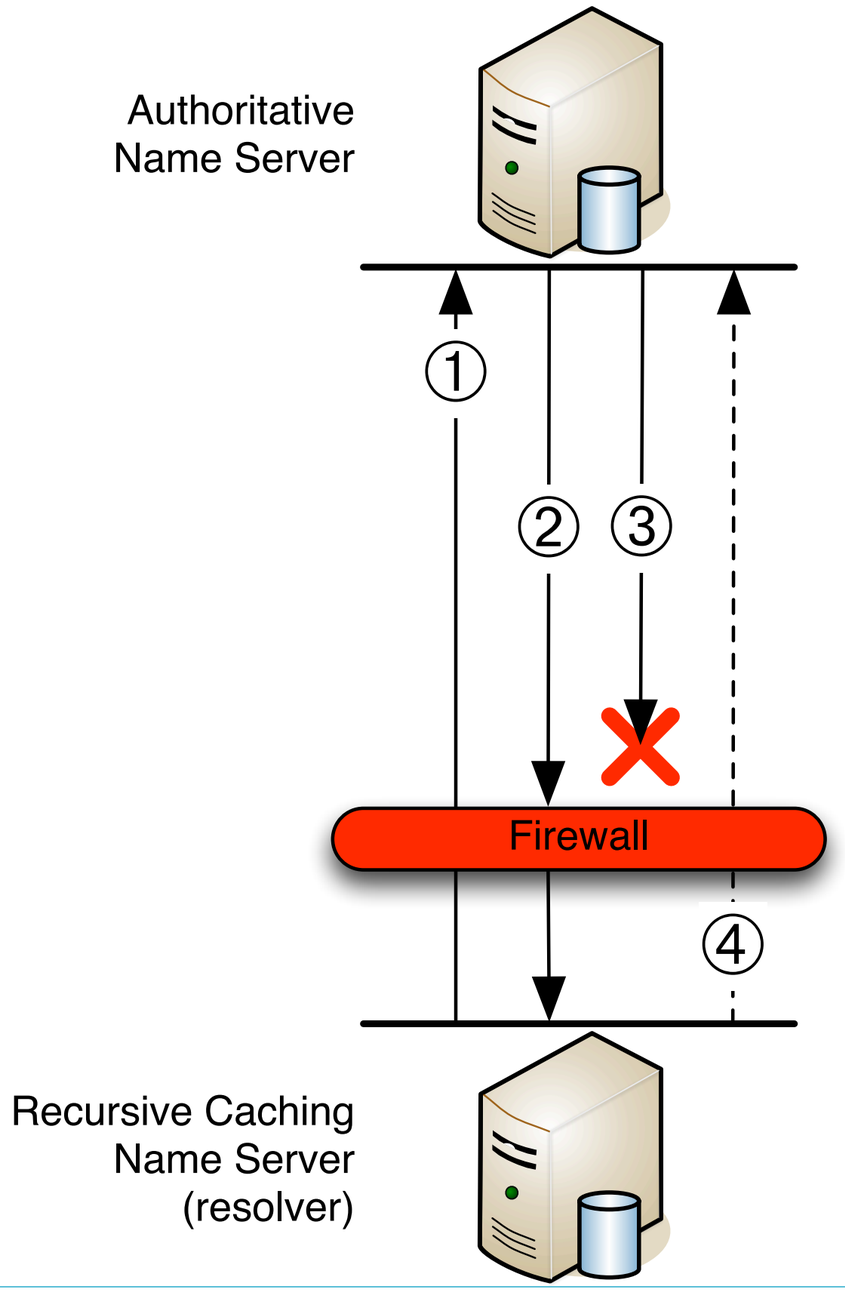

Extent of the problem

- 9% of all internet hosts may have problems receiving fragmented UDP messages [1];

- 2% – 10% of all resolving name servers experience problems receiving fragmented DNS responses [2]

[1] Weaver, N., Kreibich, C., Nechaev, B., and Paxson, V.: Implications of Netalyzr's DNS Measurements. In: Proceedings of the First Workshop on Securing and Trusting Internet Names (SATIN), Teddington, United Kingdom, (2011).

[2] Van den Broek, J., Van Rijswijk, R., Pras, A., Sperotto, A., "DNSSEC and firewalls - Deployment problems and solutions", Private Communication, Pending Publication, (2012).

The problem biting us for real

- SURFnet deployed DNSSEC for surfnet.nl in 2010 (first secure delegation in .nl)
- Within a week we had problems
- Cause: largest ISP (2.5M users) in the country blocks fragments on service network edge
- Helpdesk:
  "*SURFnet is doing something wrong*" :-(

Solutions

- Resolving name servers **SHOULD** advertise a proper max. response size to avoid fragmentation issues [RFC 2671BIS (DRAFT)];

Not explicitly stated in standards yet, nor widely implemented;

- Until then: set maximum response size at some authoritative name servers

Resolver experiments (1)
Normal operations

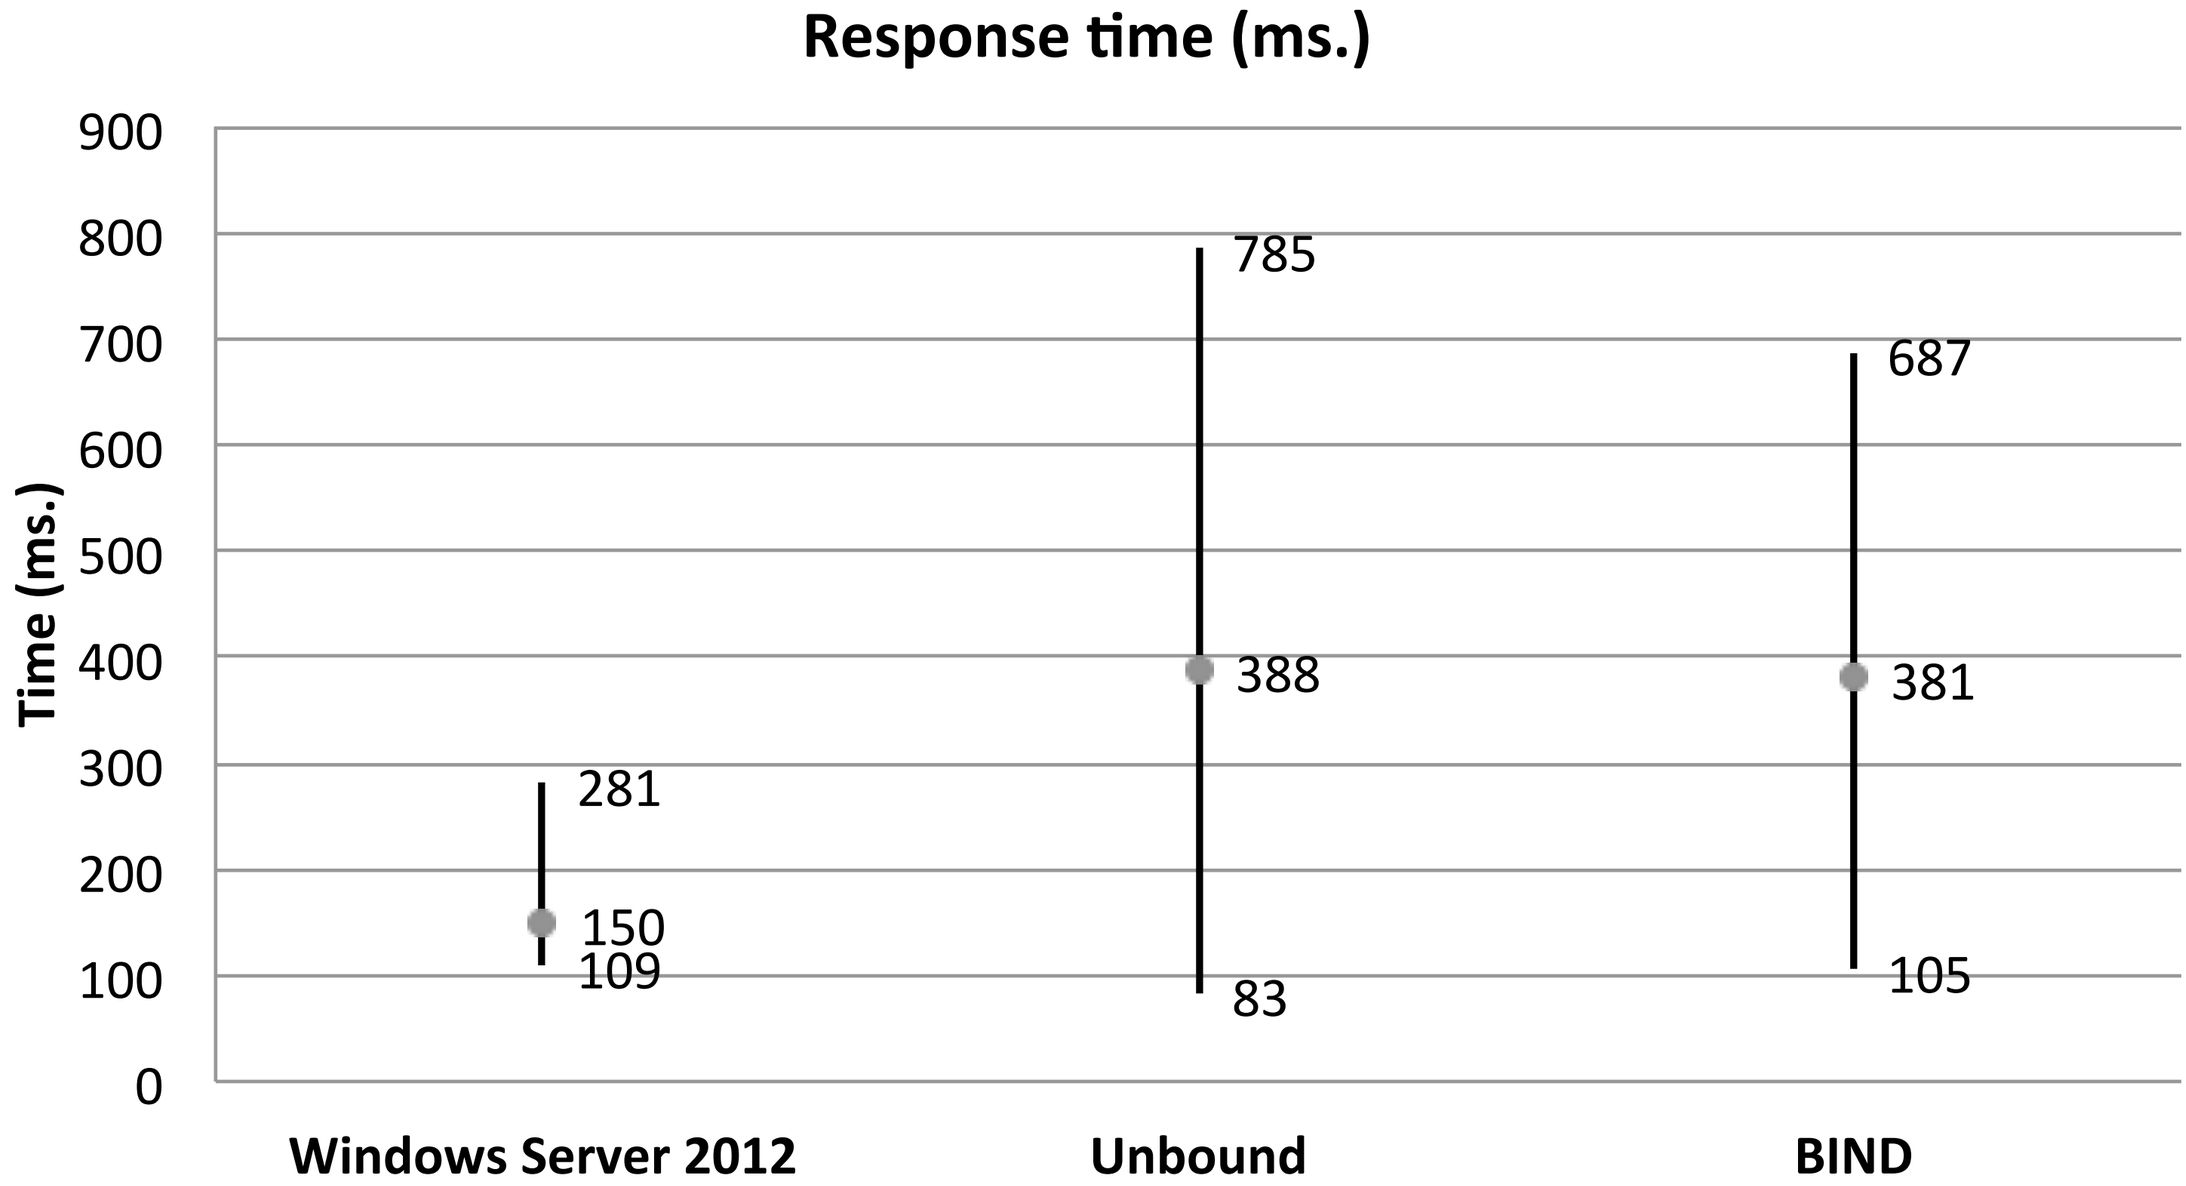

Resolver experiments (2)
Blocking fragments

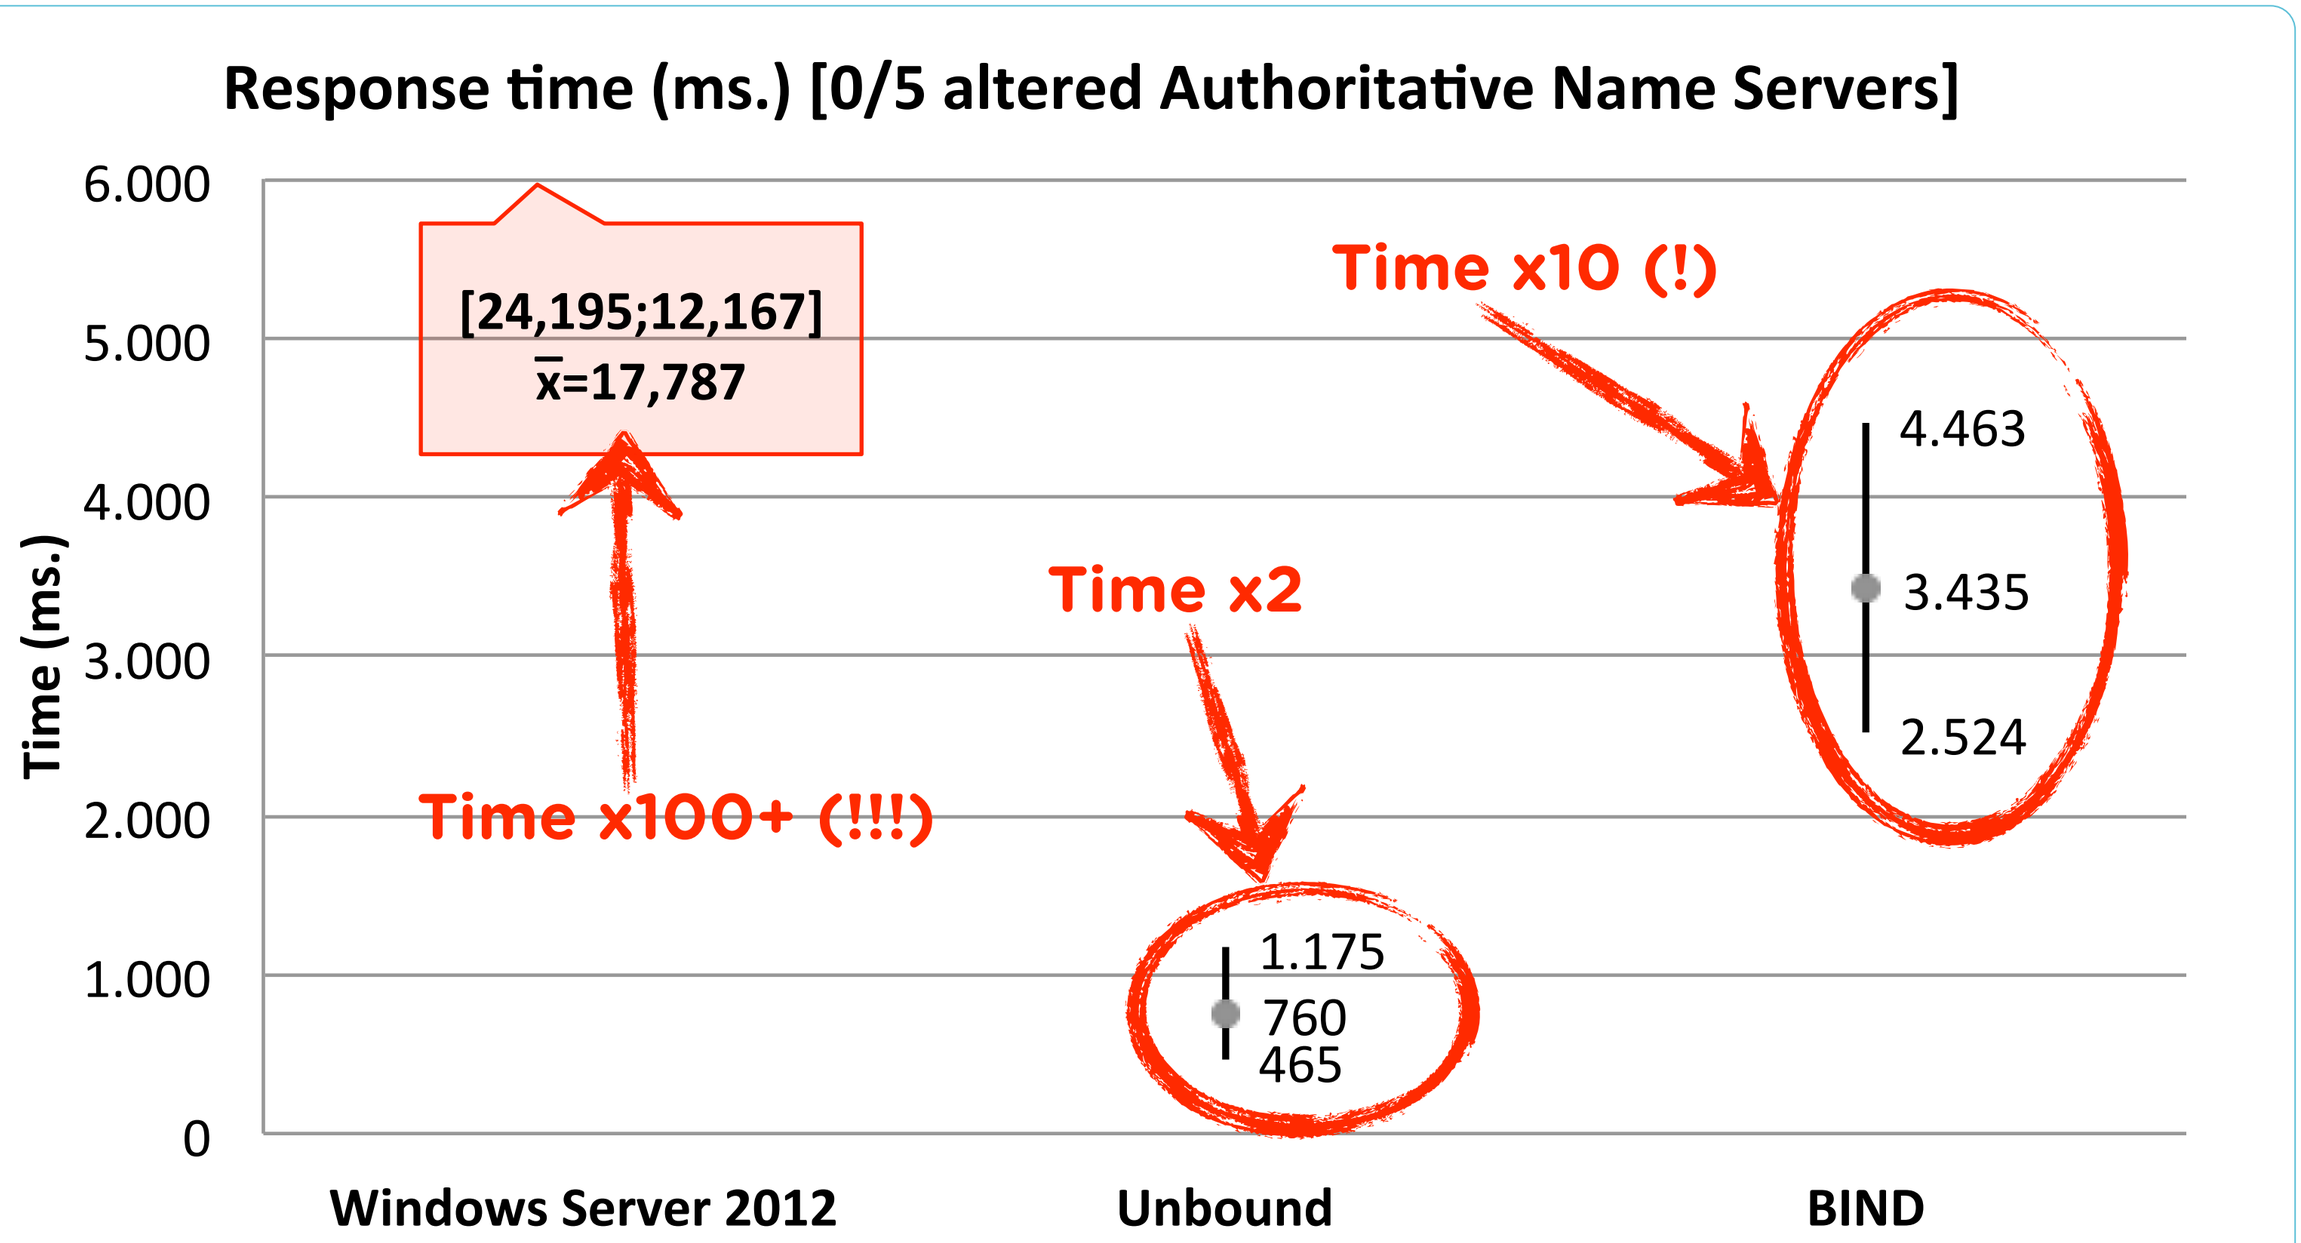

Resolver experiments (3)
Max. resp. size on 1 authNS

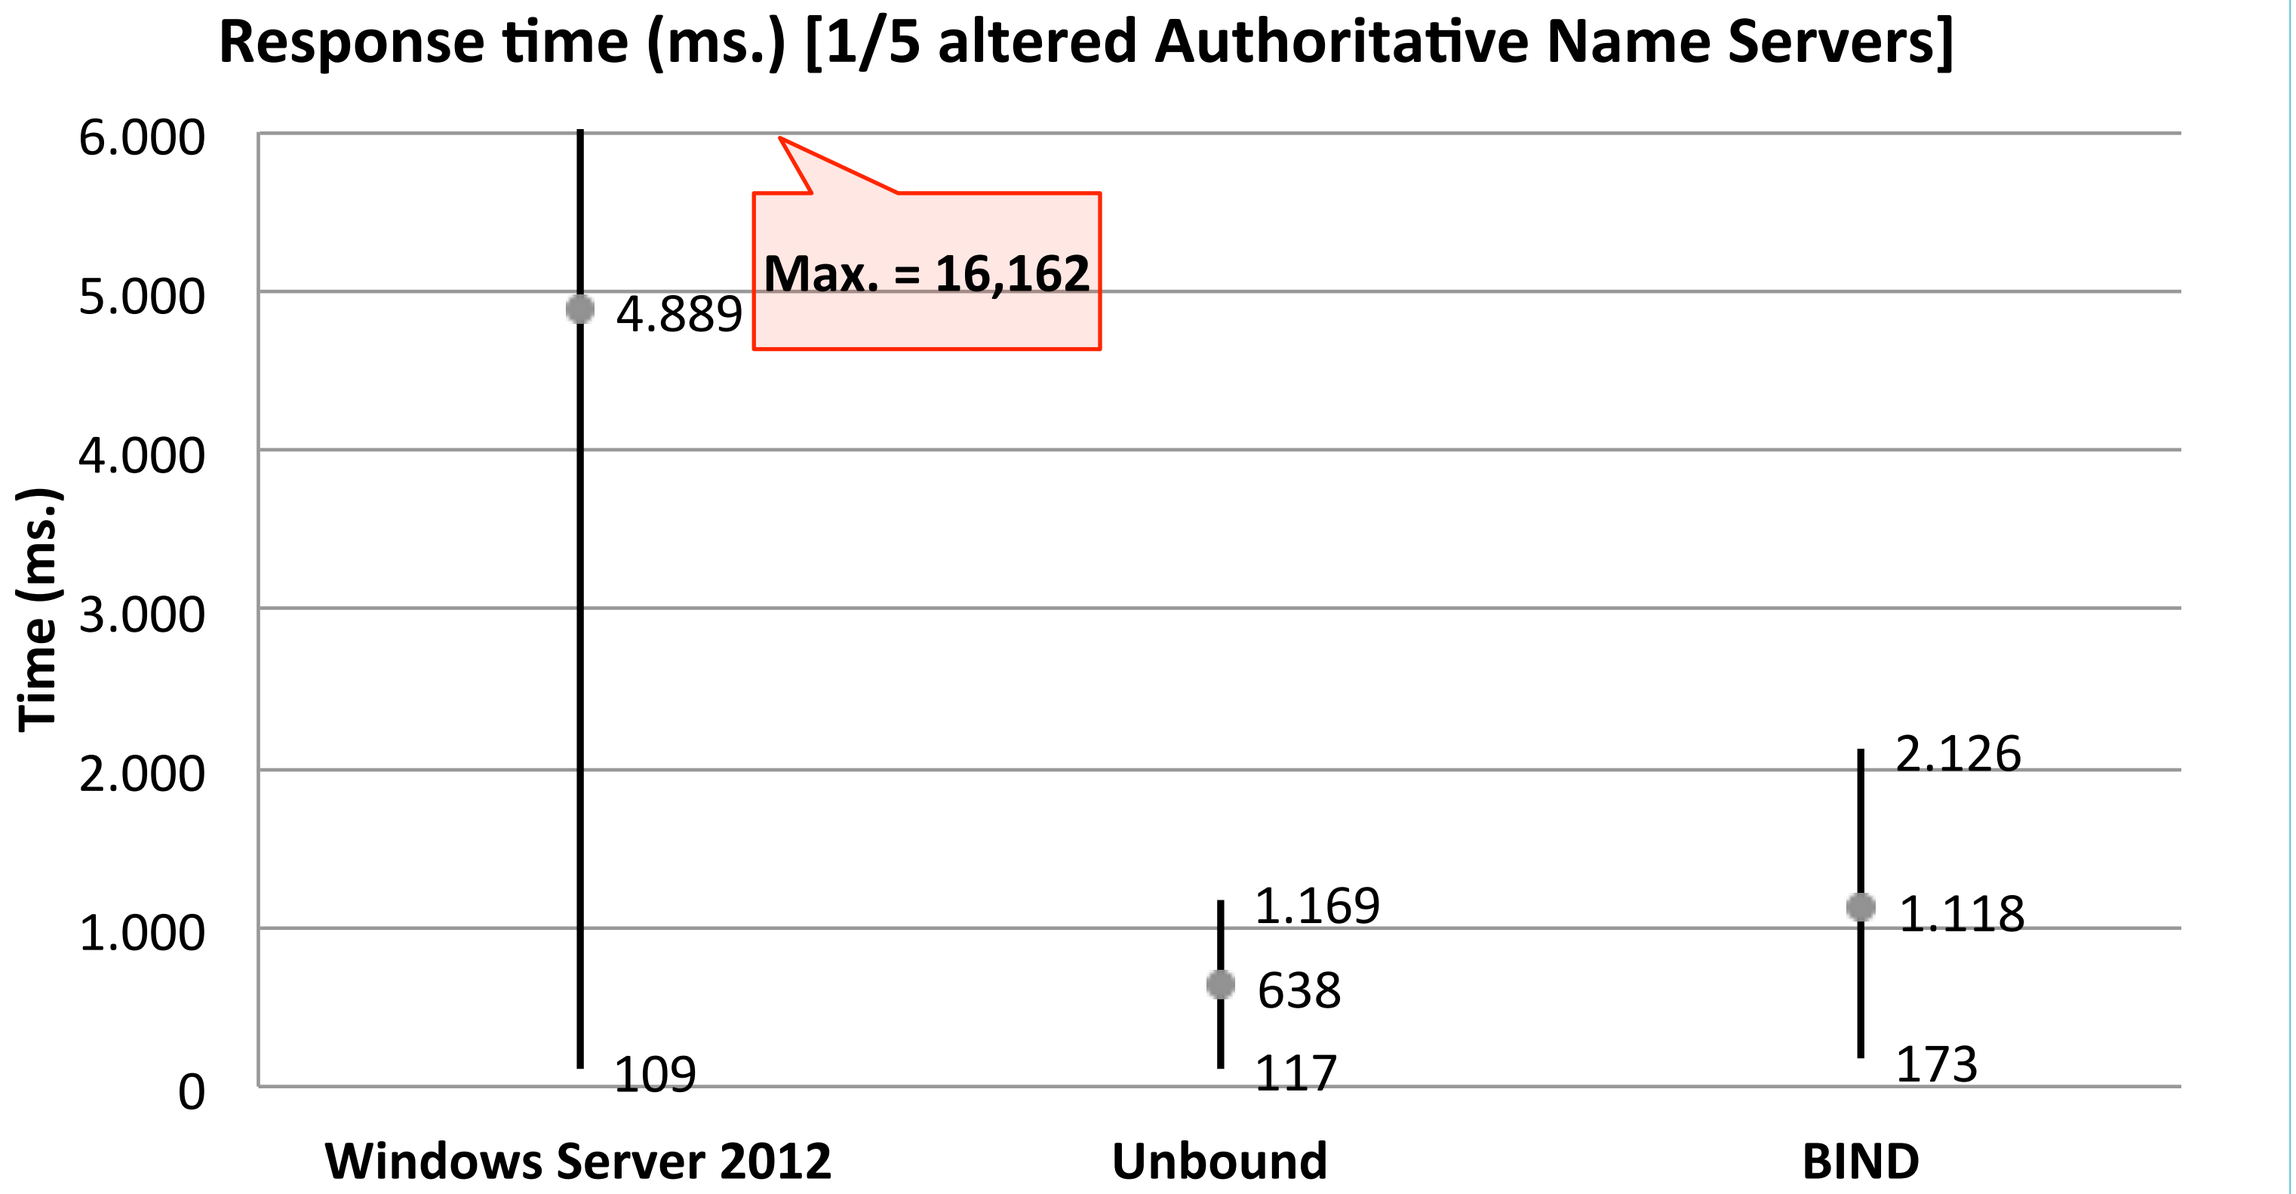

Resolver experiments (4)
Max. resp. size on 2 authNS

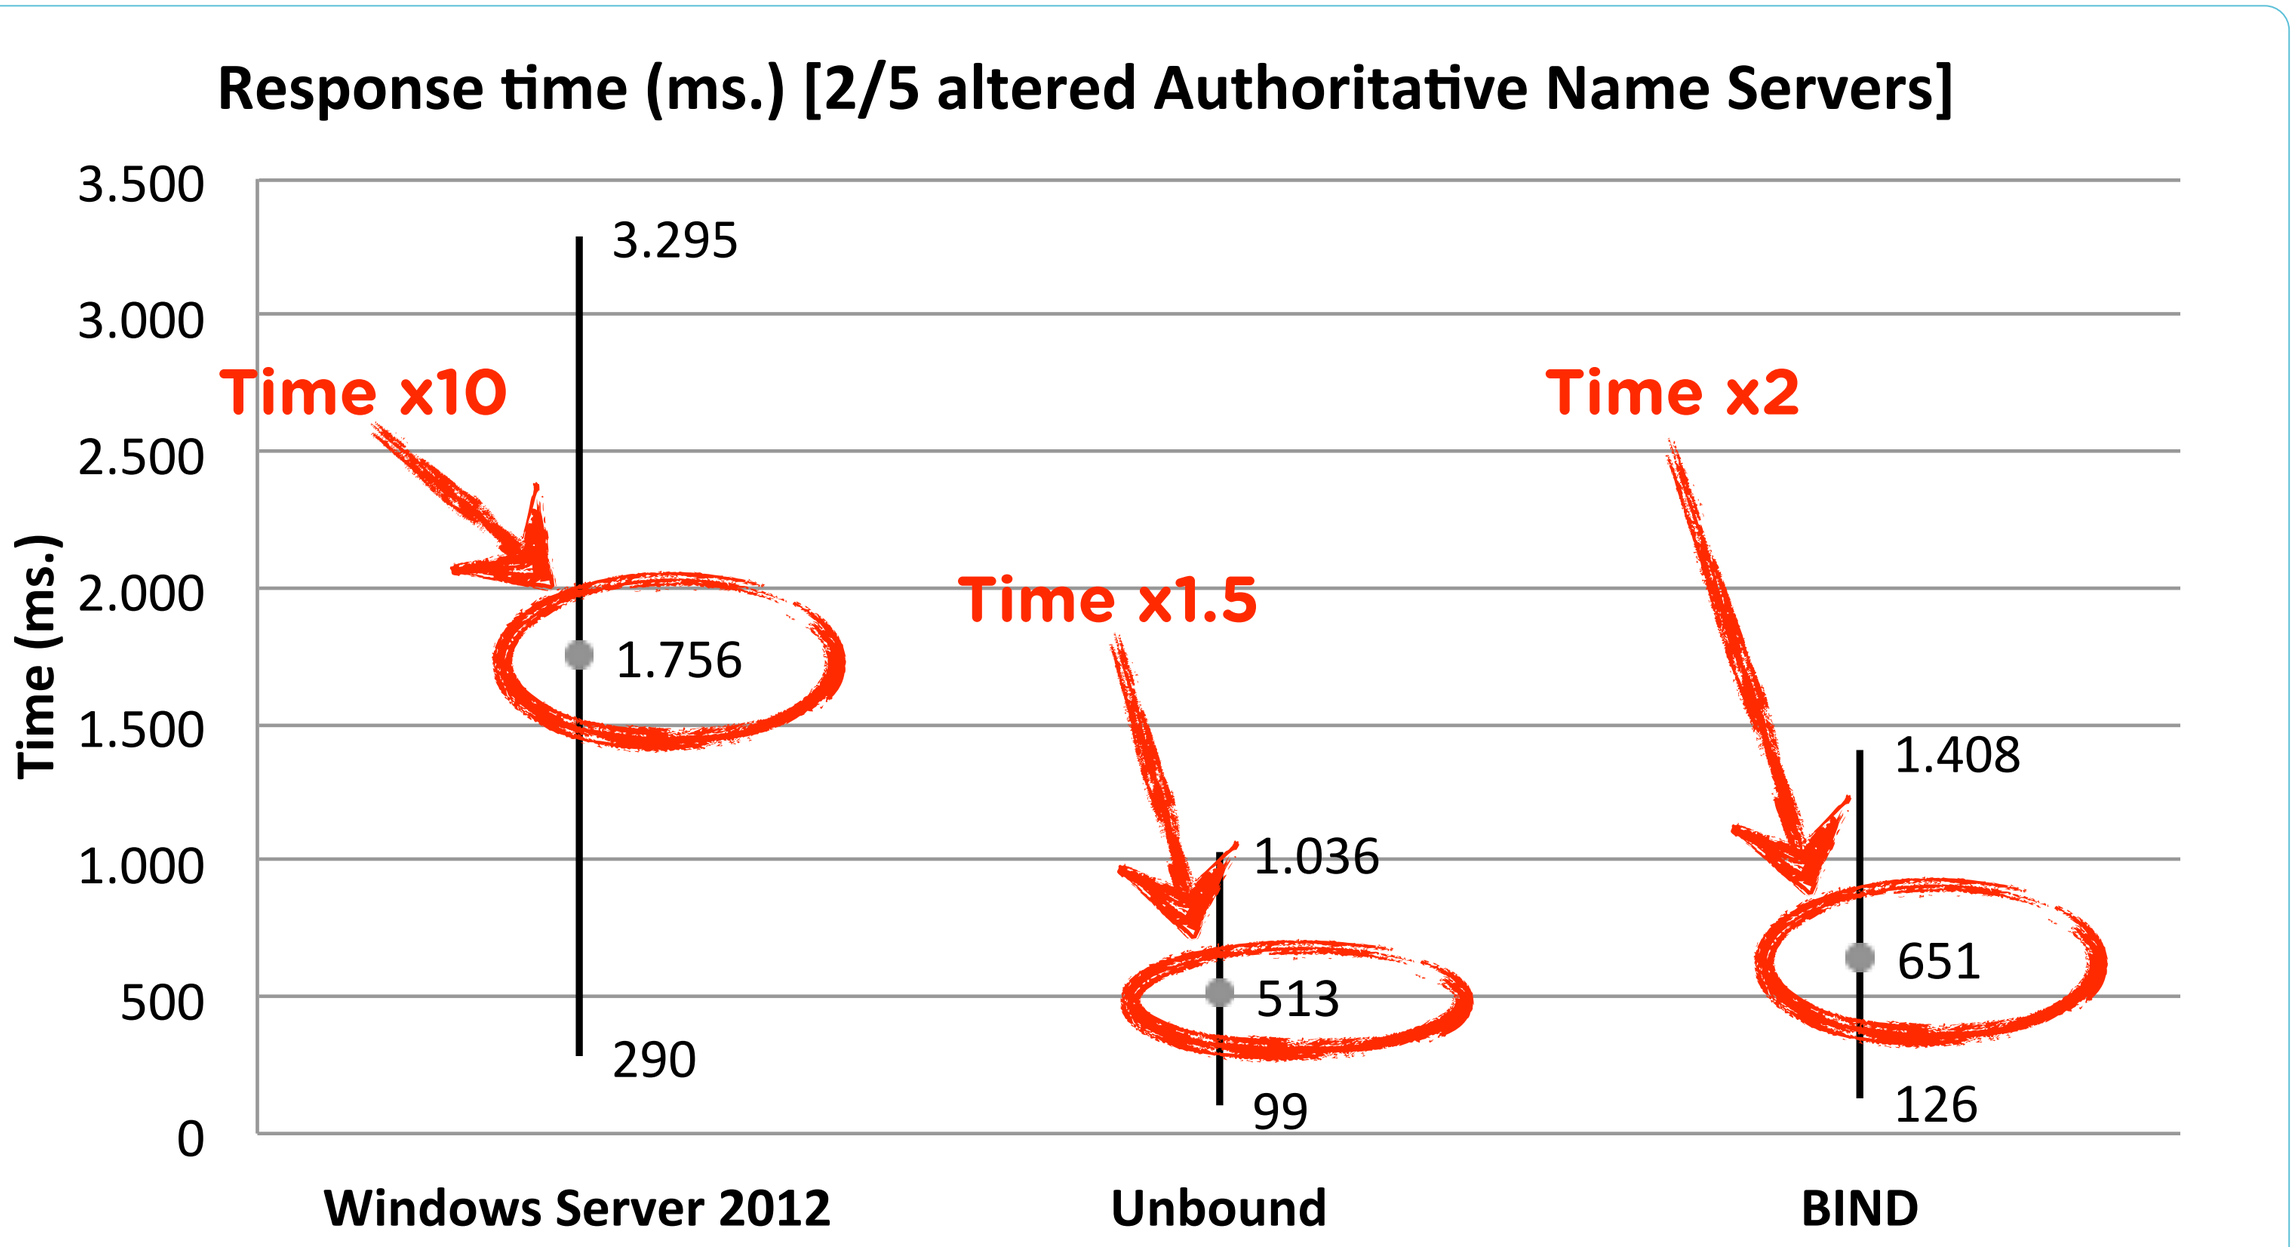

Experiment on live authNS

Traffic (IPv4 + IPv6)	Normal Operations	Max. response size 1232 bytes
Fragmented responses	28.9%	0.0%*
Fragment receiving resolvers	57.3%	0.0%*
Truncated UDP responses	0.8%	0.9%
ICMP FRTE messages	5649/h	< 1/h*
ICMP FRTE sending resolvers	1.3%	0.0%*
Total retries	25.8%	25.5%

*Statistically significant difference between experiments

Rise in truncated answers

- Experiment:
  - Querying 995 zones in .com, .edu, .mil, .net and .nl
  - All zones are signed and have a www-node
  - Results:

Max. response	A for www	AAAA for www	DNSKEY
4096	0.0%	0.0%	0.0%
1472	1.8%	1.8%	8.1%
1232	2.9%	3.5%	**40.0%**

  - 30% truncations were expected for a maximum response size of 1232 bytes by Rikitake, K., Nogawa, H., Tanaka, T., Nakao, K. and Shimojo, S. “An Analysis of DNSSEC Transport Overhead Increase”, IPSJ SIG Technical Reports 2005-CSEC-28, Vol. 2005, No. 33, pp. 345-350, ISSN 0919-6072, 2005

How to move forward?

- Working on a recommendation in the RIPE DNS working group (http://bit.ly/ripe-draft-frag)

- Make sure your resolver(s) set the maximum response size to something that actually works!

Learn how:
http://bit.ly/sn-dnssec-vali

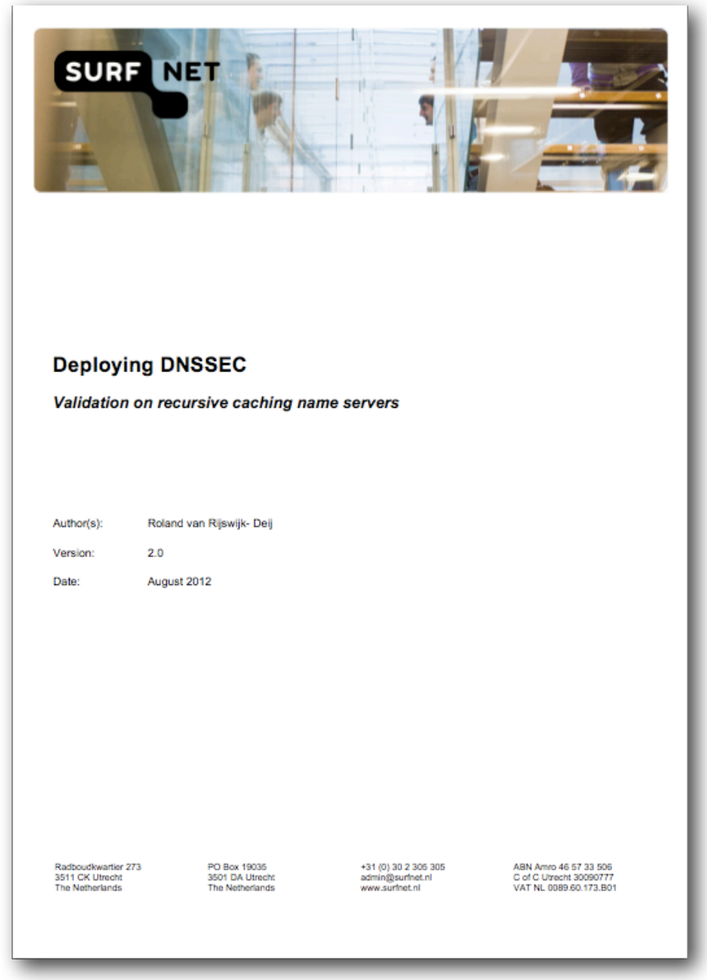

SURF NET

roland.vanrijswijk@surfnet.nl

nl.linkedin.com/in/rolandvanrijswijk

@reseauxsansfil

Questions? Remarks?

Read our blog: https://dnssec.surfnet.nl/