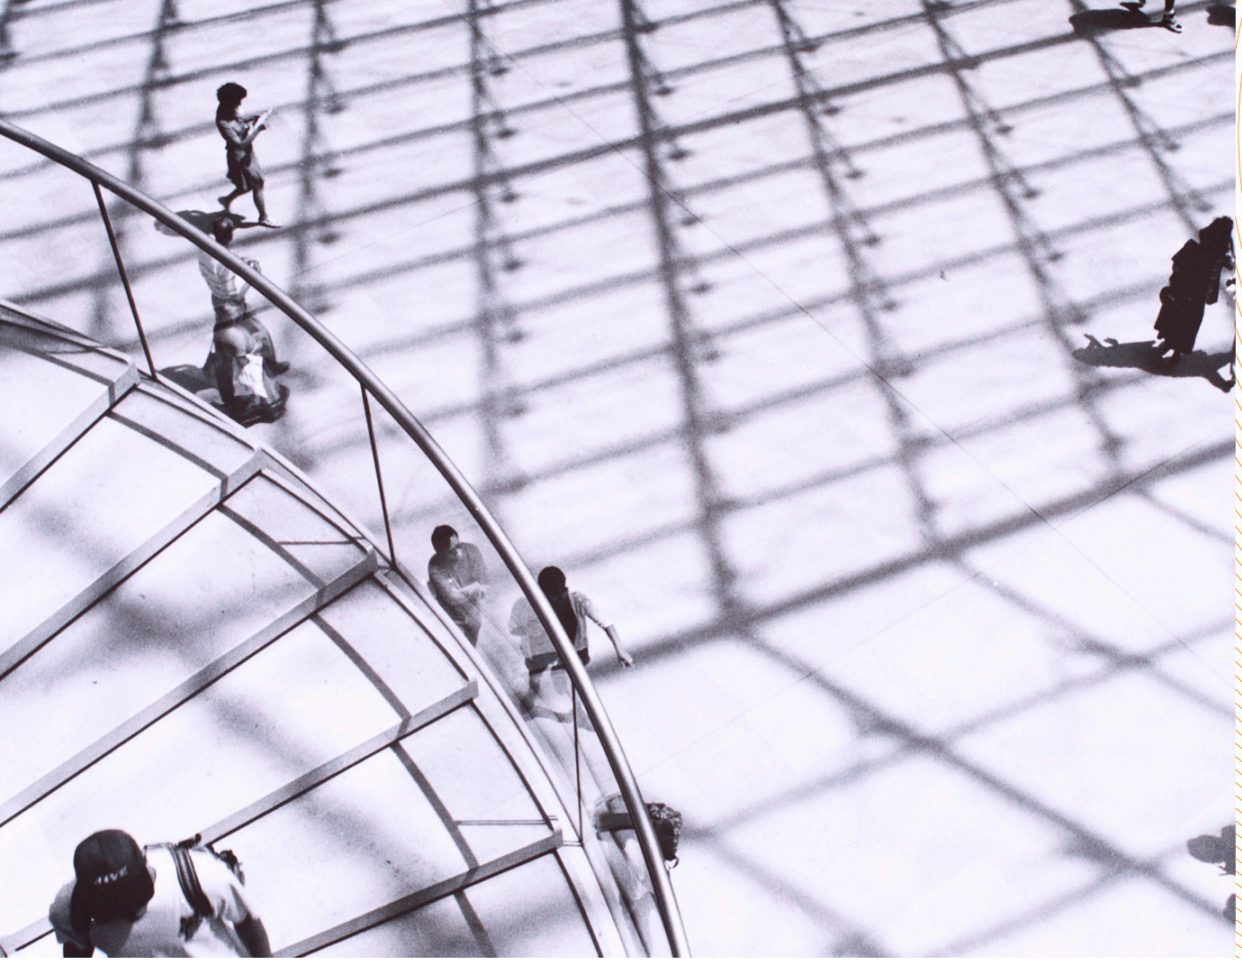

# DNSSEC research at SURFnet

*ICANN 41, Singapore*

June 22nd 2011

Roland van Rijswijk
roland.vanrijswijk [at] surfnet.nl

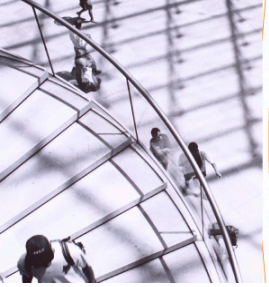

# About SURFnet


National Research and Educational Network

11000+ km ultra-high bandwidth fibre-optic network

‘Shared ICT innovation centre’

≥ 160 connected institutions
±1 million end users

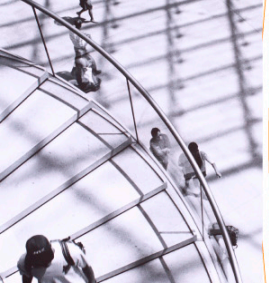

# Measuring validation

- We have a pretty good insight in DNSSEC deployment on the signing side
- Little data is available about the uptake of validation
- A Security Week article triggered us to delve into this
  - http://bit.ly/sw-dnssec-enterprise
    quote: "*There are few if any rewards for an enterprise to actually run DNSSEC live on the Internet today, especially since most ISPs aren't validating yet*"

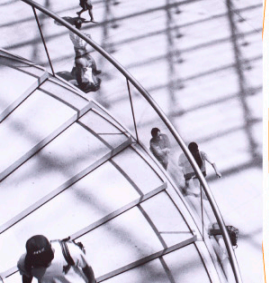

# A starting point

- JPRS presented on “How to count validators” at the DNS-OARC workshop in March 2011 (http://bit.ly/jprs-validators)

- They performed analyses on packet captures

- We had already started a similar effort but instead of analysing offline data we focus on live data

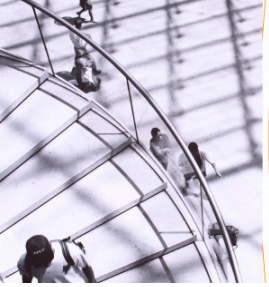

# Strategy

- Assumption:

  Only validating resolvers will send queries for **DS** and **DNSKEY** records

- We implemented simple tooling based on libpcap to capture and parse DNS packets

- We filter out queries for our signed domains (surfnet.nl & gigaport.nl)

- Aggregate queries and send them off to a database server

# Early results

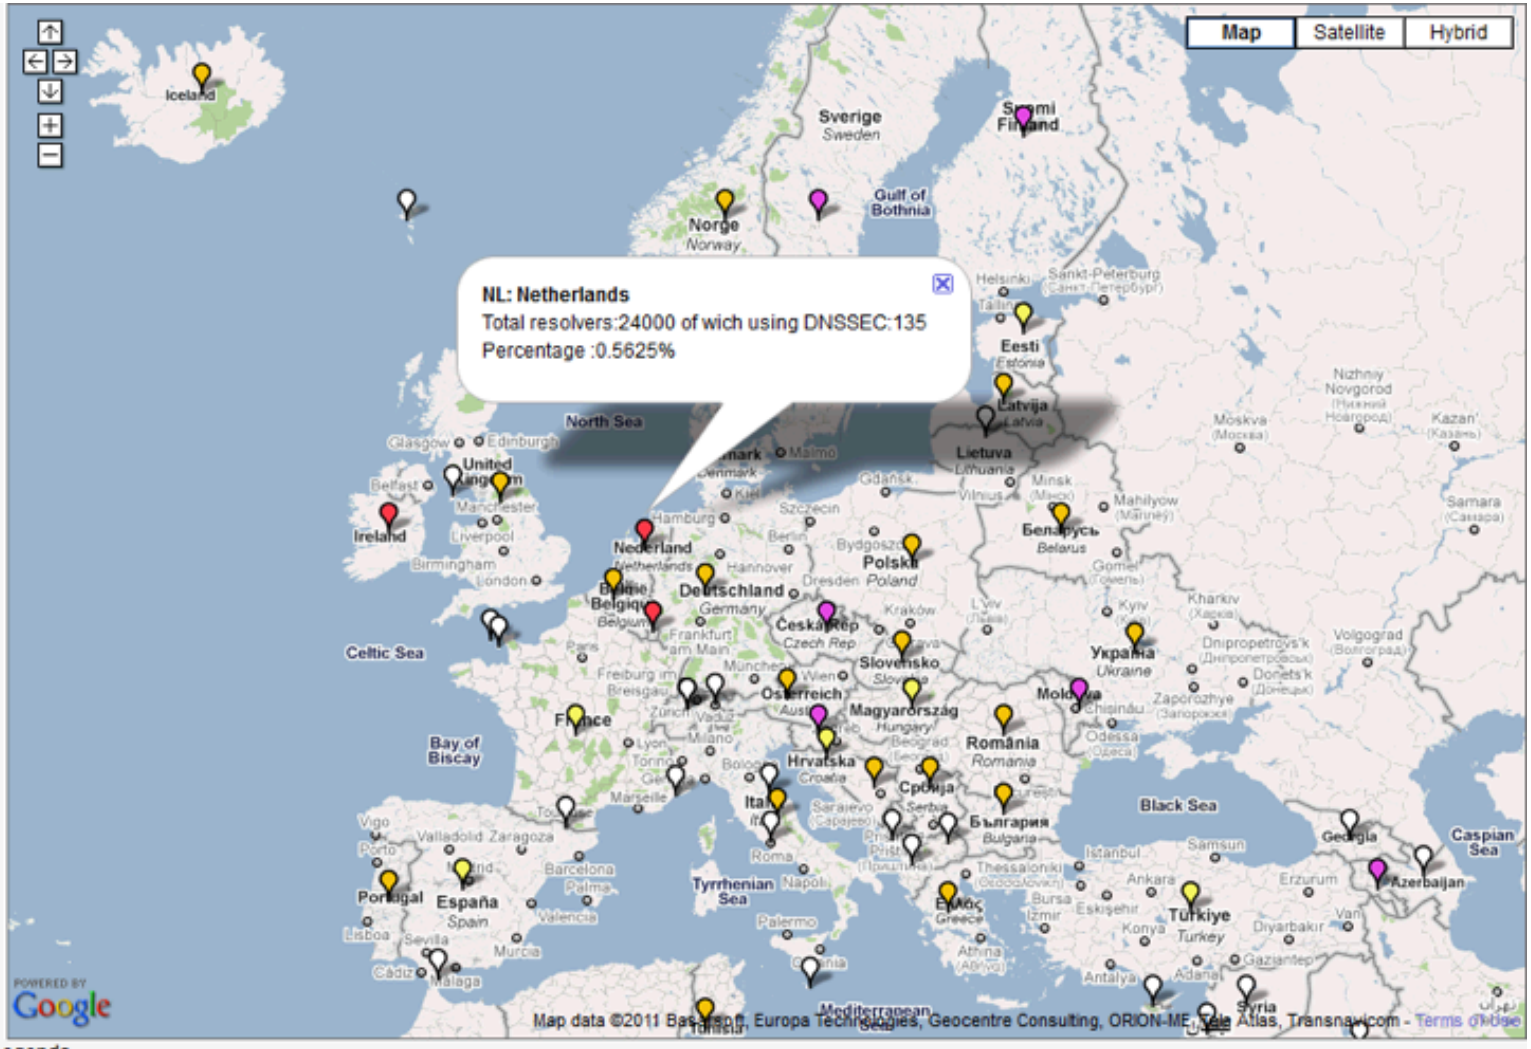

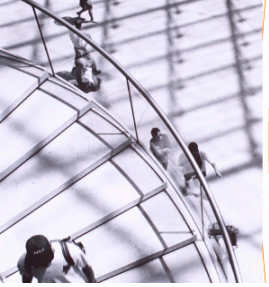

# Early results

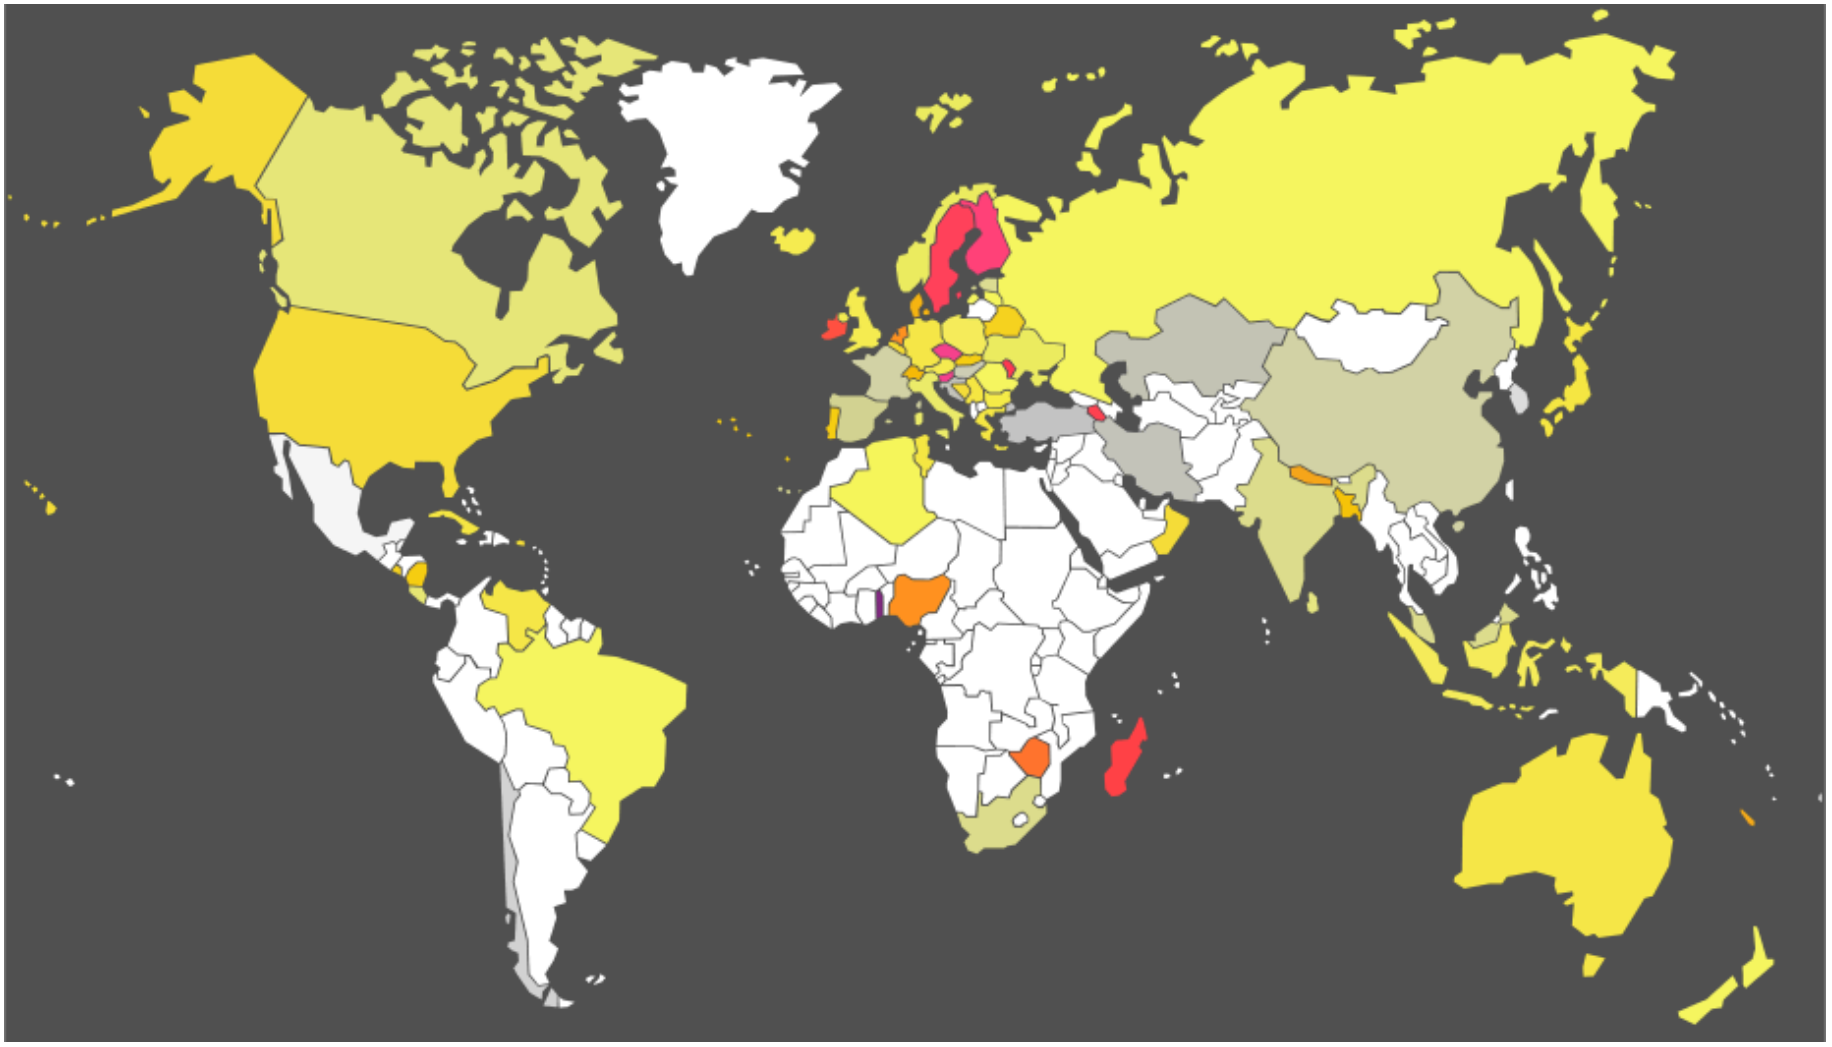

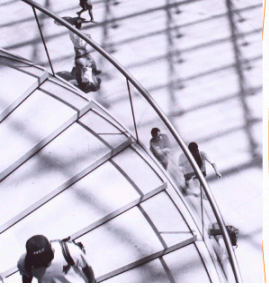

# Early results

| | | | | | |
|---|---|---|---|---|---|
| NC | New Caledonia | 446 | 2 | 0.4484% | 99.5516% |
| NE | Niger | 44 | 0 | 0% | 100% |
| NG | Nigeria | 545 | 3 | 0.5505% | 99.4495% |
| NI | Nicaragua | 297 | 1 | 0.3367% | 99.6633% |
| NL | Netherlands | 24000 | 135 | 0.5625% | 99.4375% |
| NO | Norway | 3129 | 5 | 0.1598% | 99.8402% |
| NP | Nepal | 210 | 1 | 0.4762% | 99.5238% |
| NR | Nauru | 4 | 0 | 0% | 100% |
| NU | Niue | 3 | 0 | 0% | 100% |
| NZ | New Zealand | 2994 | 5 | 0.167% | 99.833% |

| | | | | | | |
|---|---|---|---|---|---|---|
| 192.87.36.36 | SURFnet bv | 53311 | 1320 | 0 | 2011-06-10 17:09:32.914083 | 2011-05-30 18:02:57.866841 |
| 192.87.106.99 | SURFnet bv | 4197 | 1034 | 0 | 2011-06-10 17:06:15.573689 | 2011-05-30 18:10:41.945889 |
| 195.169.124.124 | SURFnet bv | 34037 | 1282 | 0 | 2011-06-10 17:09:20.182531 | 2011-05-30 18:03:49.117727 |
| 194.171.9.20 | SURFnet bv | 106 | 14 | 0 | 2011-06-10 15:31:10.787898 | 2011-05-31 09:20:46.611621 |
| 192.87.106.106 | SURFnet bv | 80516 | 1455 | 0 | 2011-06-10 17:09:46.970471 | 2011-05-30 18:02:54.089563 |
| 131.155.140.130 | Technische Universiteit Eindhoven | 389 | 67 | 0 | 2011-06-10 17:07:26.989673 | 2011-05-30 19:23:46.910177 |
| 84.241.226.7 | T-mobile Netherlands bv. | 2988 | 259 | 0 | 2011-06-10 16:40:33.647419 | 2011-05-30 18:13:08.180237 |
| 84.241.226.137 | T-mobile Netherlands bv. | 3302 | 263 | 0 | 2011-06-10 16:53:39.508226 | 2011-05-30 18:44:06.486269 |

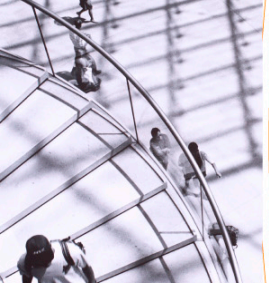

# Plans

- We plan to make this information available to interested parties (no public site planned for the moment)

- We are talking to SIDN to see if we can run similar experiments on the .nl infrastructure

- We will release the tools in open source under a BSD licence

- Please contact me if you are interested or wish to contribute

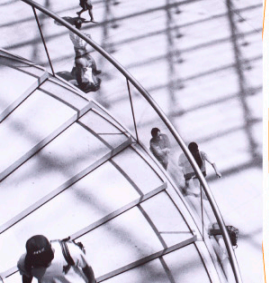

# UDP fragmentation issues

- Late last year we experienced problems with a large ISP in The Netherlands
- surfnet.nl had just gotten a DS in .nl
- Colleagues started complaining that they could not log on to their mail from home
- It turned out to be a firewall at the ISP that discarded UDP fragments
- Even though they did not do validation, they could not resolve our records **(!)**

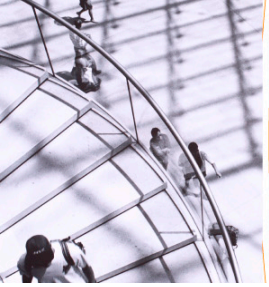

# All is well that ends well?

- We talked to their engineers
- They could not replace the firewall
- In the end, they lowered the EDNS0 buffer size on their resolver to 512 bytes
- Problem solved, right?

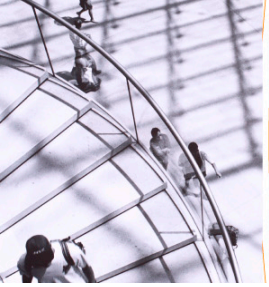

# The saga continues

- Everything worked well until in March 2011 we suddenly started getting complaints from some companies trying to e-mail us

- Lo and behold, they were customers of this same ISP

```
Unable to deliver message to the following recipients, due to being
unable to connect successfully to the destination mail server.
Reporting-MTA: dns;********************
Received-From-MTA: dns;macpro.lan
Arrival-Date: Thu, 17 Mar 2011 14:54:18 +0100
Final-Recipient: rfc822;*************@surfnet.nl
Action: failed
Status: 4.4.7
From:  ***********************************
To:  ******************@surfnet.nl
```

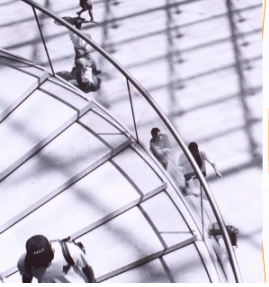

# The firewall strikes back

- It turned out that only customers using the hosted MS Exchange service had issues
- After talking to engineers at the ISP we discovered the problem
- They had upgraded the dedicated resolvers in their hosted exchange environment to Windows 2008R2 which does EDNS0 and sets DO=1
- Solution: tweak some arcane registry setting

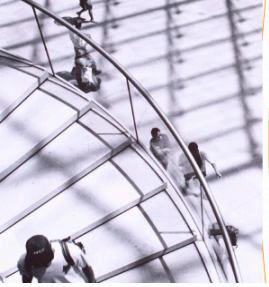

# Co-discovery

- While investigating this issue we discovered something interesting: the resolvers behind the firewall received the first fragment of the UDP packet

- The protocol stack detects that fragments are missing and sends back an ICMP message which we can detect:

```
11:01:59.849643 IP *.*.*.* > ns3.surfnet.nl: ICMP ip reassembly time exceeded, length
92
11:01:59.849655 IP *.*.*.* > ns3.surfnet.nl: ICMP ip reassembly time exceeded, length
92
```

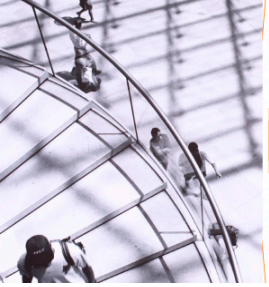

# Research

- We are extending our monitoring tools to detect this issue and log it in our database

- Some initial packet dumping showed scary results

- People even seem to think that UDP fragments are an attack (we have had abuse complaints sent to our CERT team!)

- We have a student who is creating a lab setup to test our theory and write a paper on the results

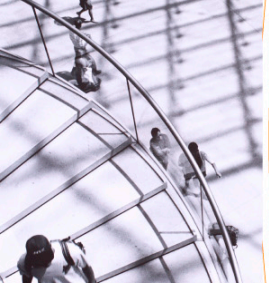

# Conclusion

- This issue requires some serious attention
- It affects owners of signed domains and they can do very little about it
- I have some ideas about making authoritative servers somehow detect this and react to it (but some people are not going to like these ideas)
- If you operate a signed zone you may wish to look into this...

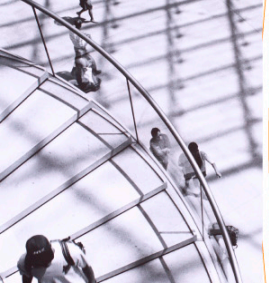

# That's all folks! Questions?

**If you have any questions about this presentation, please feel free to contact me by e-mail**


Roland van Rijswijk

**roland.vanrijswijk [at] surfnet.nl**

**@reseauxsansfil**