Defending your DNS in a post-Kaminsky world

Paul Wouters
<paul@xelerance.com>

Overview

- History of DNS and the Kaminsky attack
- Various DNS problems explained
- Where to address the DNS problem
  - Nameservers, Network, Client and Data
- Bad ideas and why we won't do them
- Proposed and implemented ideas
- The future of DNS(SEC?)

xelerance

DNS resilie

June, 2008

Longer TTL's are much safe

The calculations above indicate the relative ease with which DNS data can be spoofed. For example, using the formula derived earlier on a domain with a 3600 second TTL, an attacker sending 7000 fake response packets/s (a rate of 4.5Mb/s), stands a 10% chance of spoofing a record in the first 24 hours, which rises to 50% after a week.

For a domain with a TTL of 60 seconds, the 10% level is hit after 24 minutes, 50% after less than 3 hours, 90% after around 9 hours.

Note that the attacks mentioned above can be detected by watchful server operators - an unexpected incoming stream of 4.5mbit/s of packets might be noticed.

An important assumption however in these calculations is a known or static destination port of the authentic response.

The IETF

Sunday, July 14, 2008

DNSSEC must happen NOW

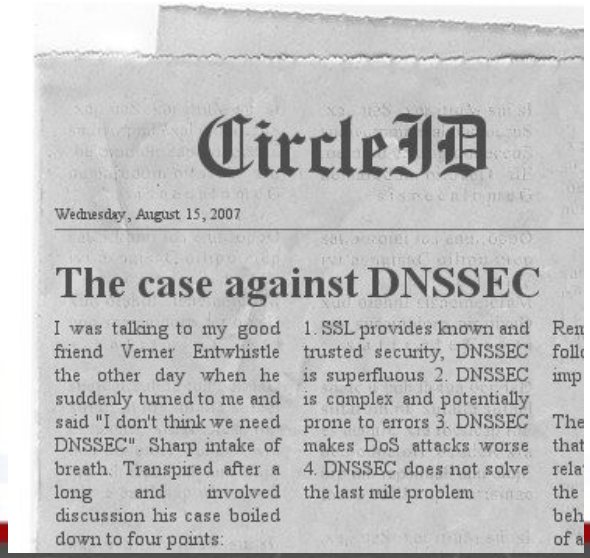

CircleID

Wednesday, August 15, 2007

The case against DNSSEC

I was talking to my good friend Verner Entwhistle the other day when he suddenly turned to me and said "I don't think we need DNSSEC". Sharp intake of breath. Transpired after a long and involved discussion his case boiled down to four points:

1. SSL provides known and trusted security, DNSSEC is superfluous 2. DNSSEC is complex and potentially prone to errors 3. DNSSEC makes DoS attacks worse 4. DNSSEC does not solve the last mile problem

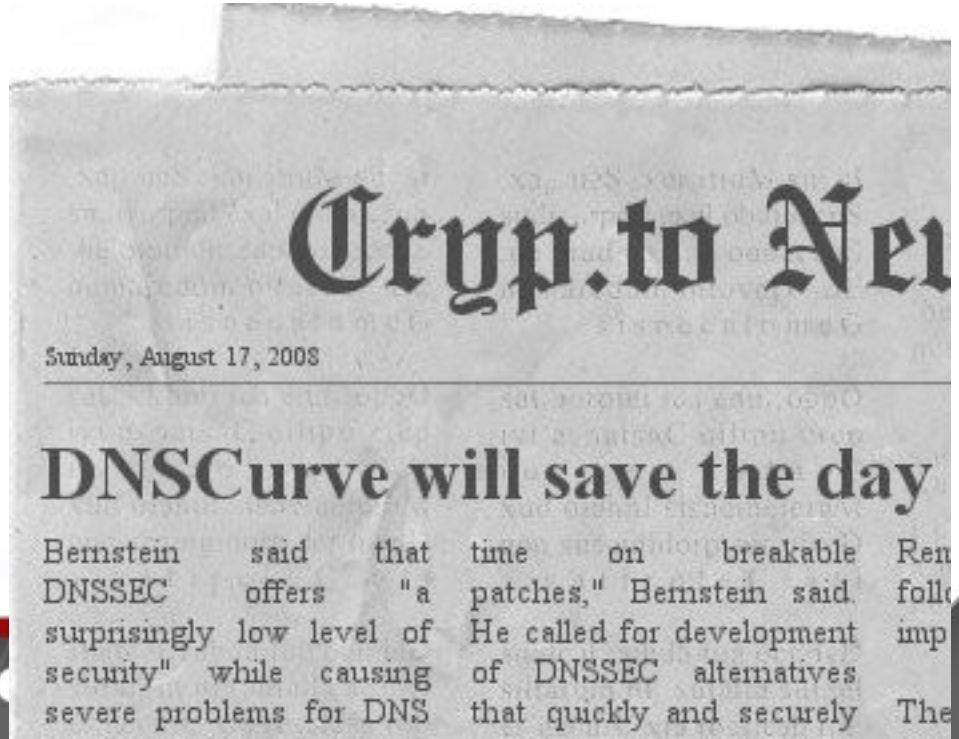

Cryp.to Ne

Sunday, August 17, 2008

DNSCurve will save the day

Bernstein said that DNSSEC offers "a surprisingly low level of security" while causing severe problems for DNS

time on breakable patches," Bernstein said. He called for development of DNSSEC alternatives that quickly and securely

Black Hat Bri

Vendor and NGO's involved

xelerance

Nominum
Connecting Networks and People

NEUSTAR
Trusted to Bring Networks Together

IETF

Microsoft

CERT

IOActive™
COMPREHENSIVE COMPUTER SECURITY SERVICES

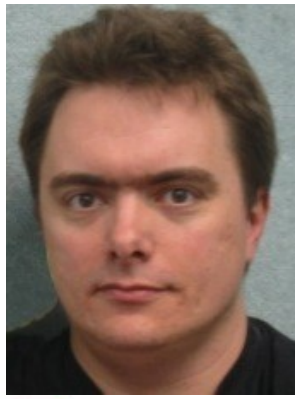

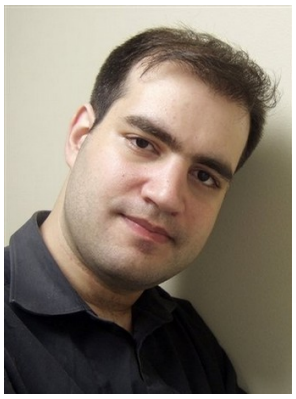

NLnet Labs

OpenDNS

ISC Internet Systems Consortium

Two phase deployment

- First release a generic fix for the Kaminsky attack that does not leak information to the bad guys (source port randomization)
- Then release the bug and patches specifically against the Kaminsky attack

DNS query packet

IP header containing Source IP and Dest IP

UDP or TCP Header containing
Source Port and Dest Port
(if TCP, also random Sequence Number)

DNS Query ID
DNS Query
Option flags

DNS query example

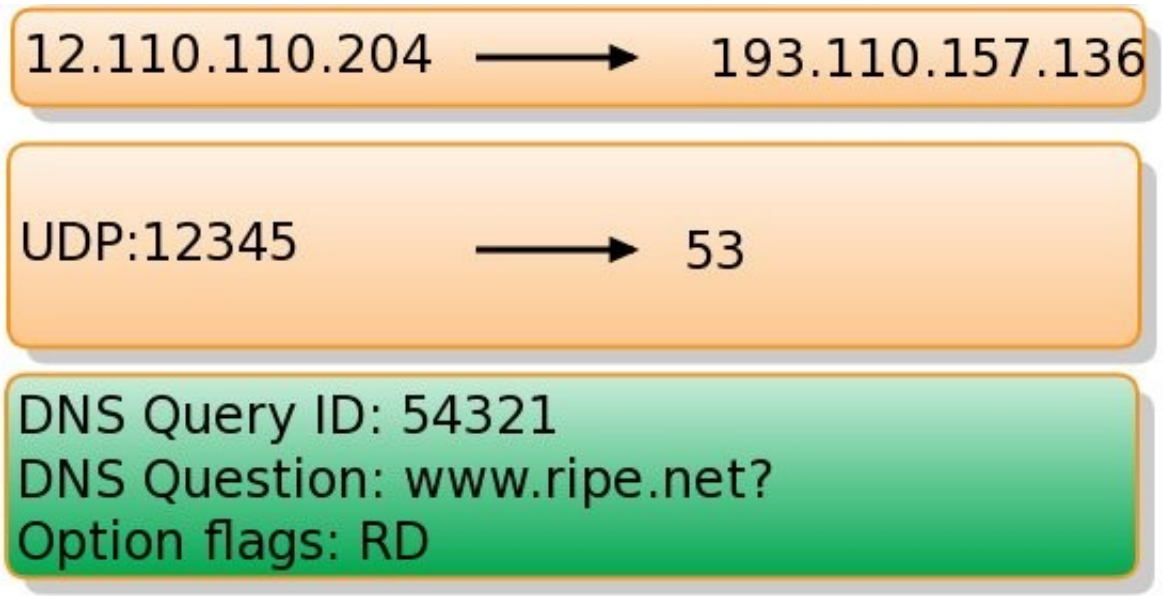

DNS Answer packet

193.110.157.136 → 12.110.110.204

UDP:53 → 12345

QUESTION SECTION
Query ID: 54321
Question: www.ripe.net?

ANSWER SECTION
www.ripe.net = 193.0.0.195 (ttl=172800)

AUTHORITY SECTION
ripe.net NS ns-pri.ripe.net. (ttl=172800)
ripe.net NS ns-ext.isc.org. (ttl=172800)

ADDITIONAL SECTION
ns-pri.ripe.net A 193.0.0.195 (ttl=...)
ns-pri.ripe.net AAAA 2001:610:240:0:53:3

TXID is not enough anymore

- Bellowin's (theoretical) attack (1995)

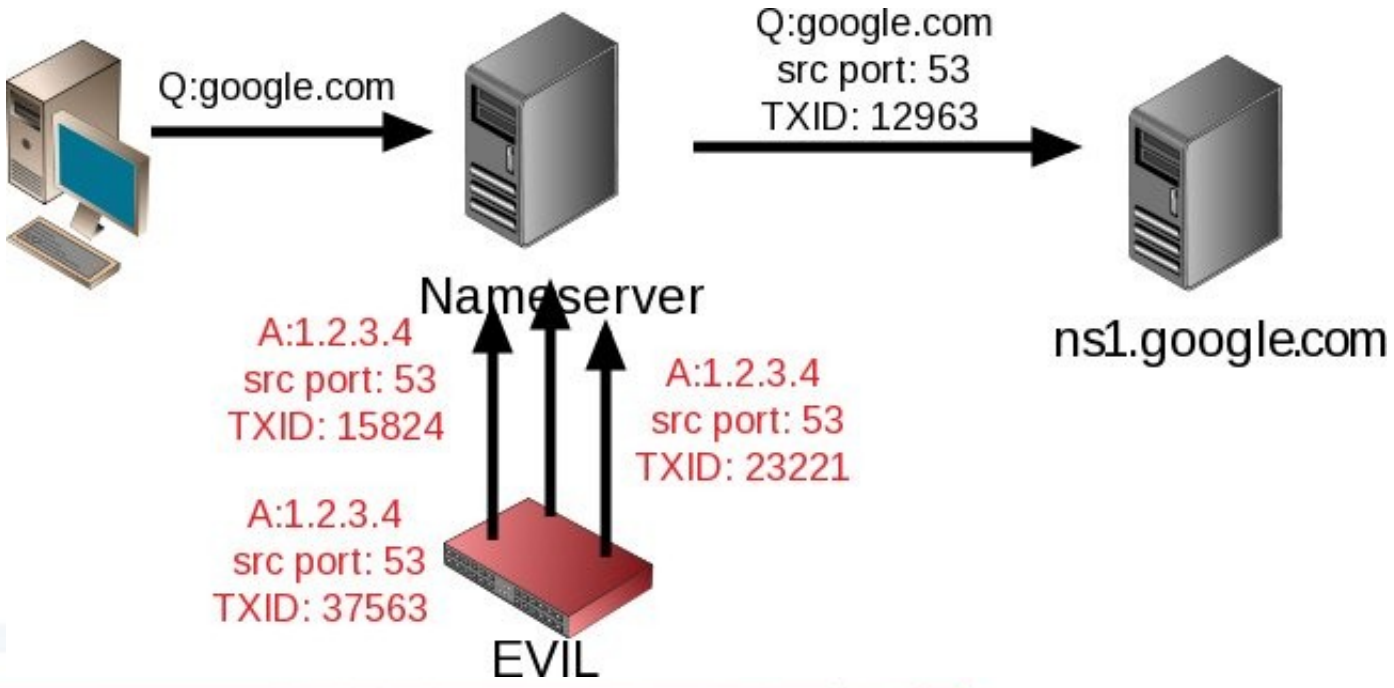

Losing the race

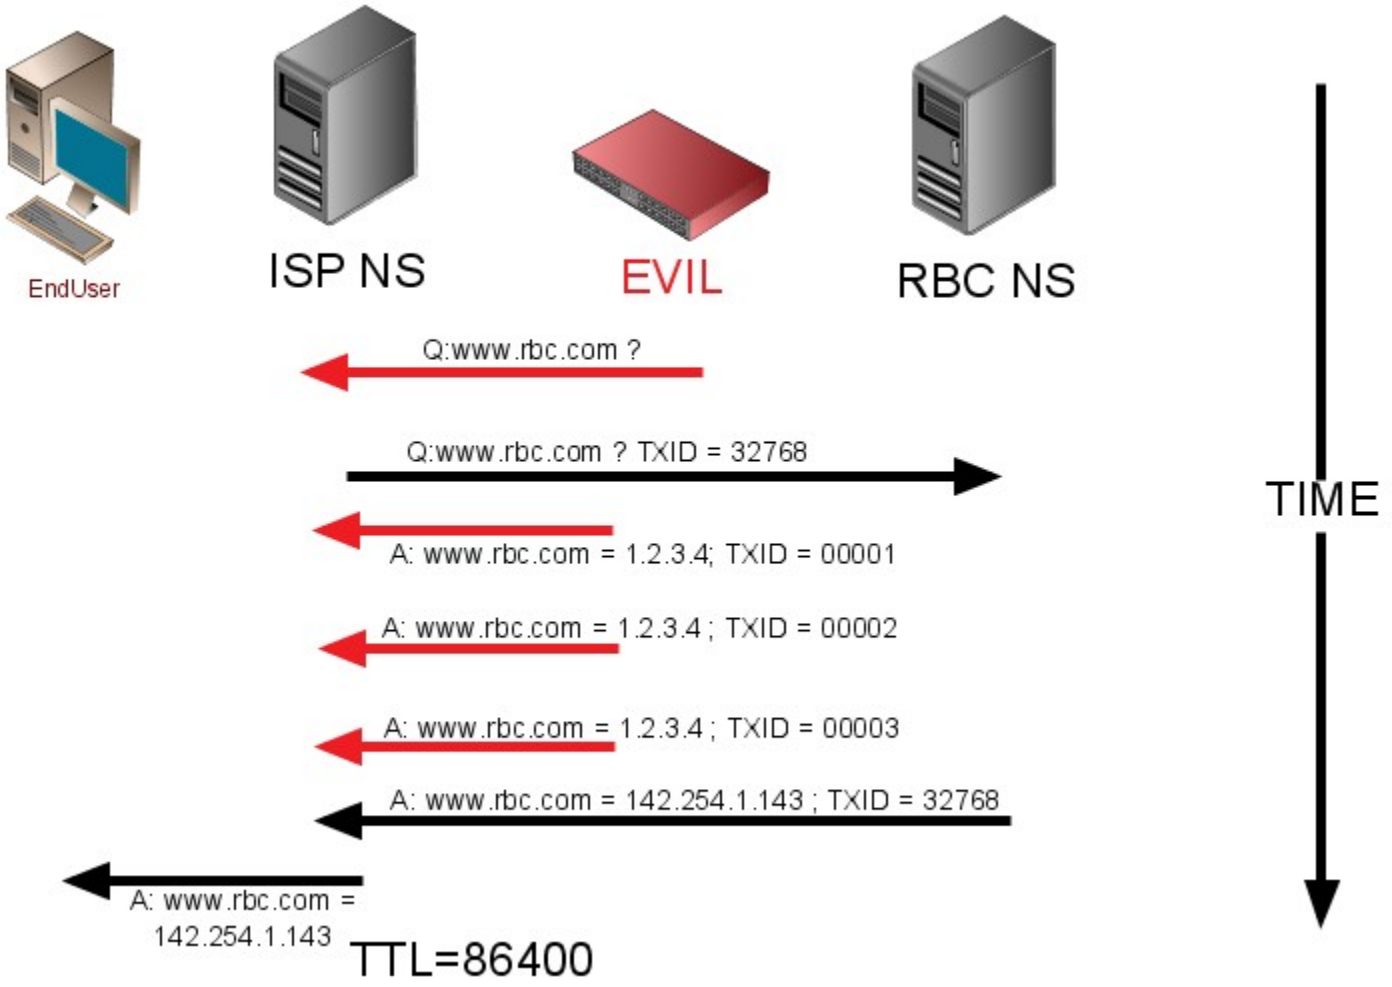

Winning the race

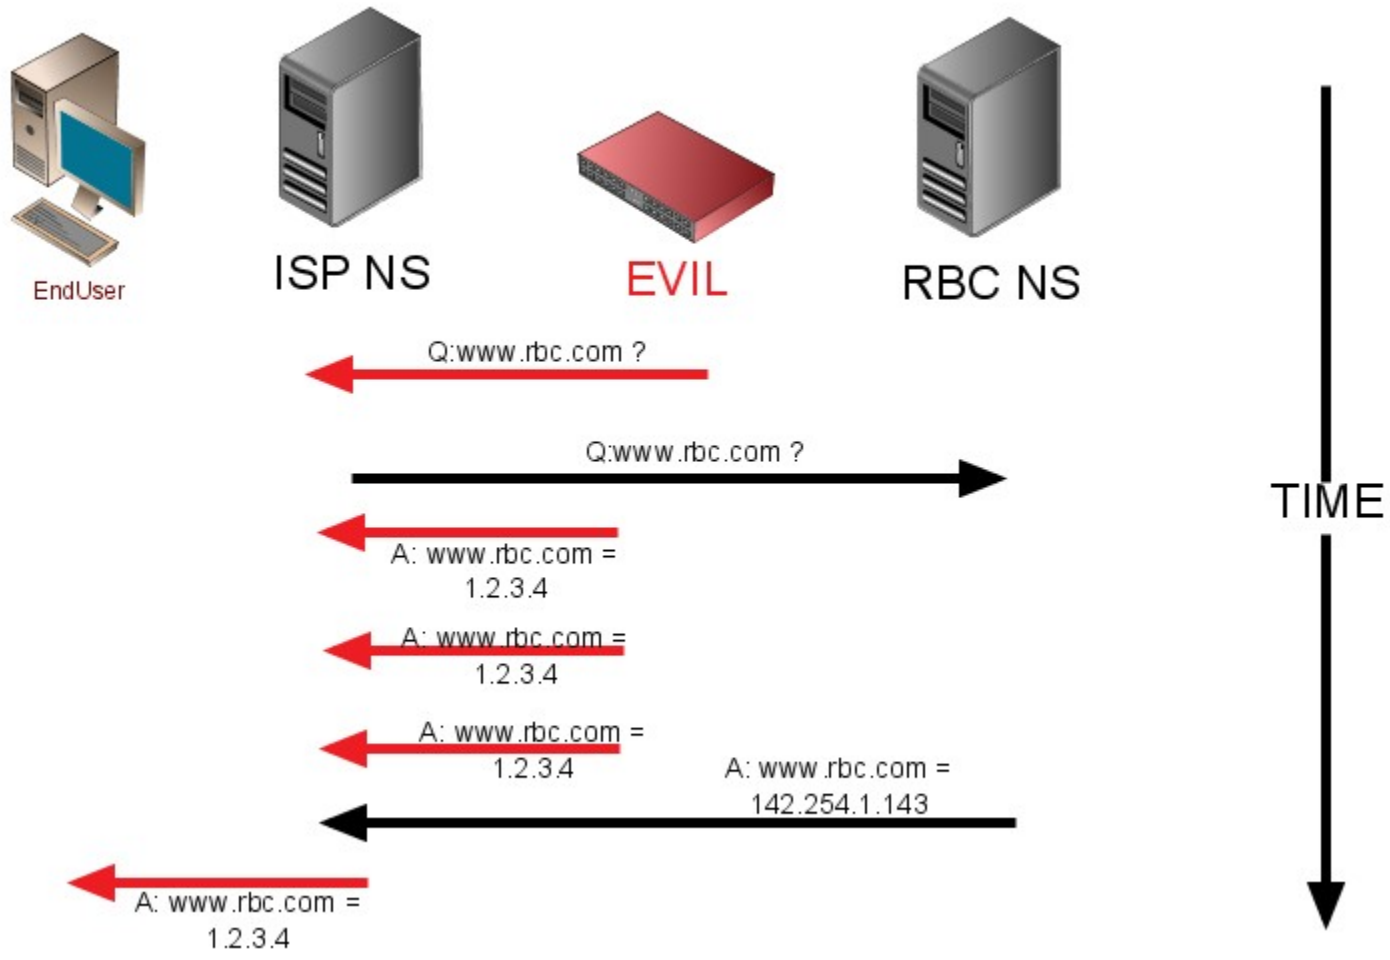

Random source ports

- Bernstein:Use random src ports as entropy

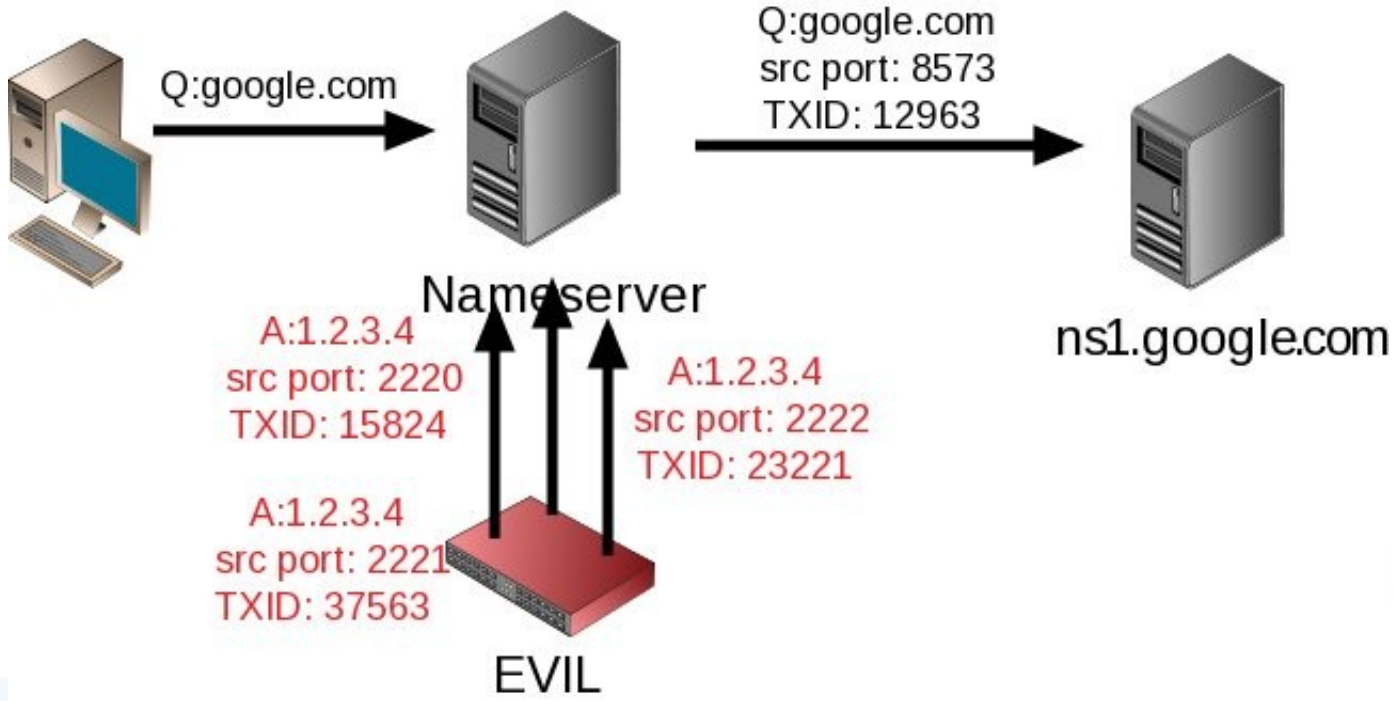


xelerance

DJB's hack is still just a hack

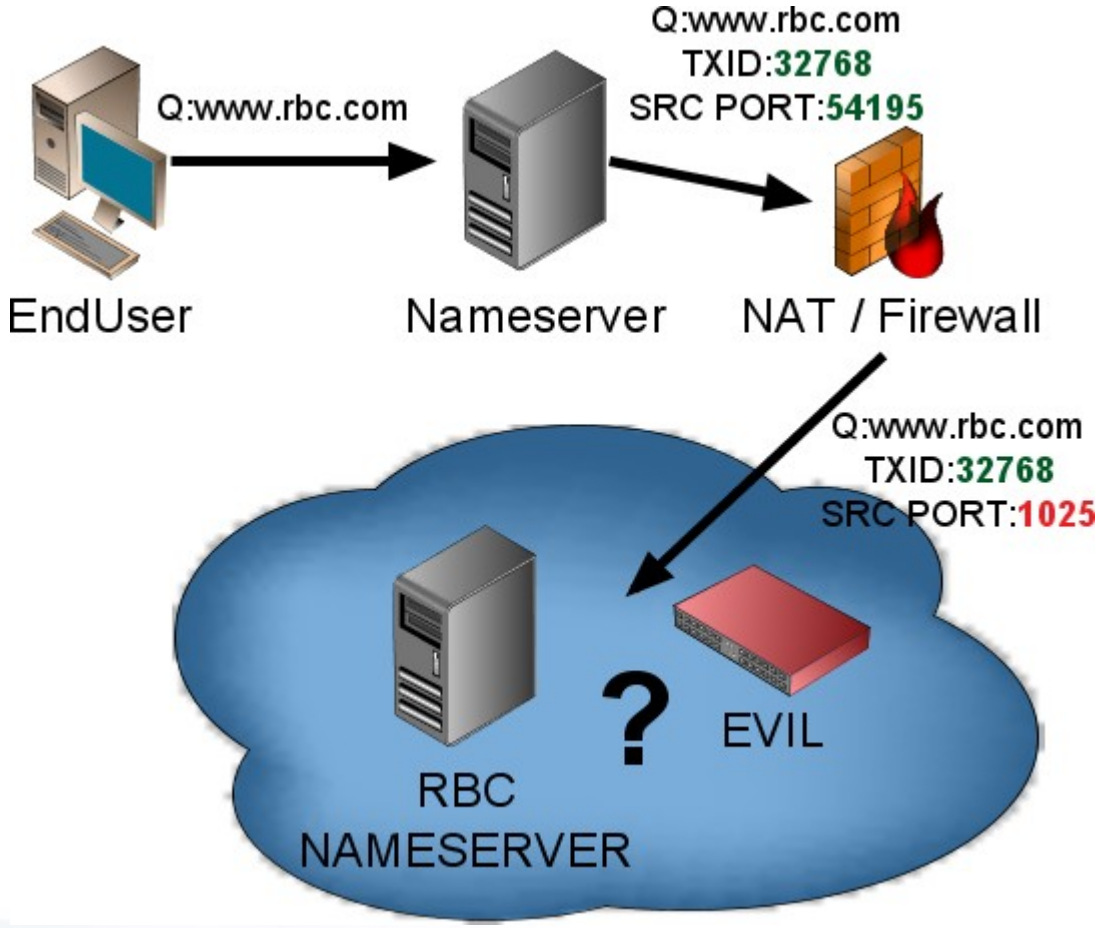

NAT and DNS rebinding

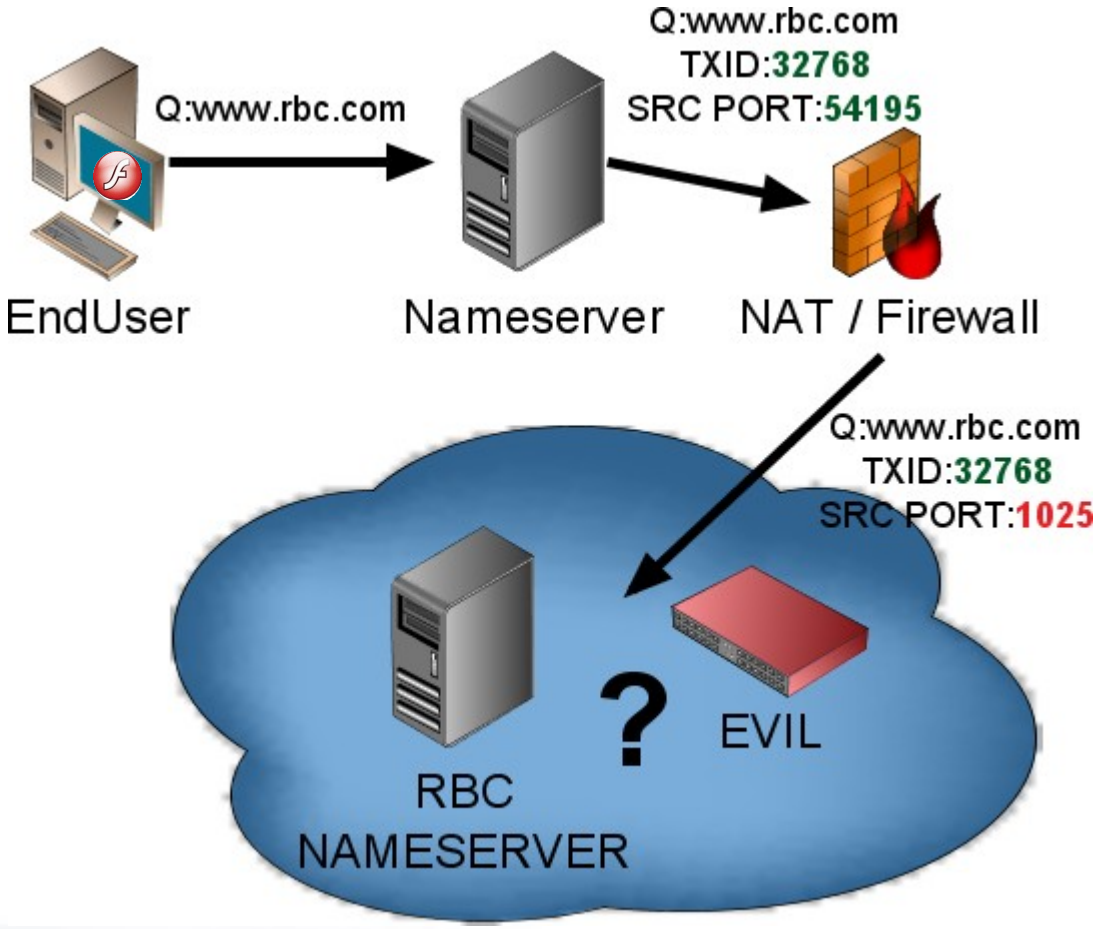

NAT and DNS rebinding (2)

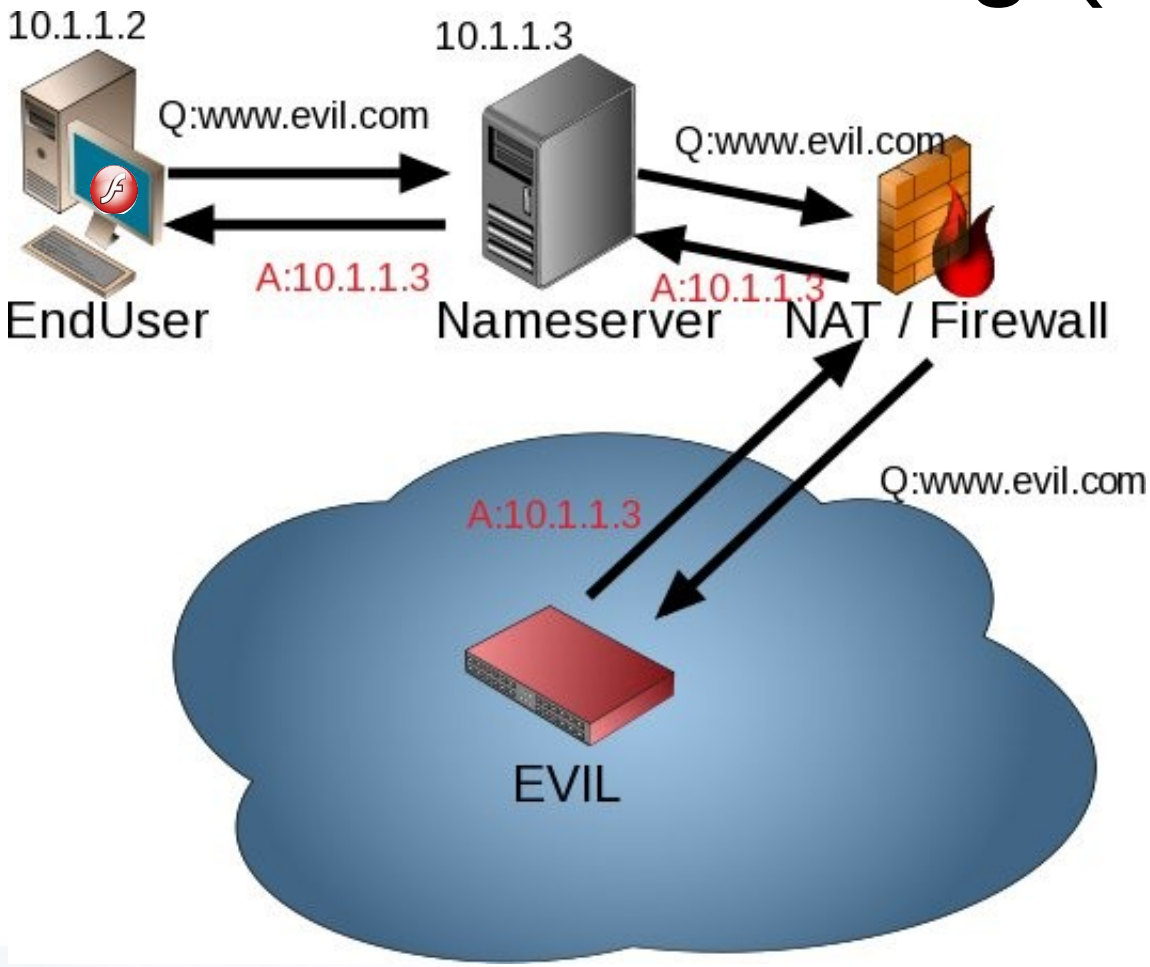

Birthday Attack on src ports

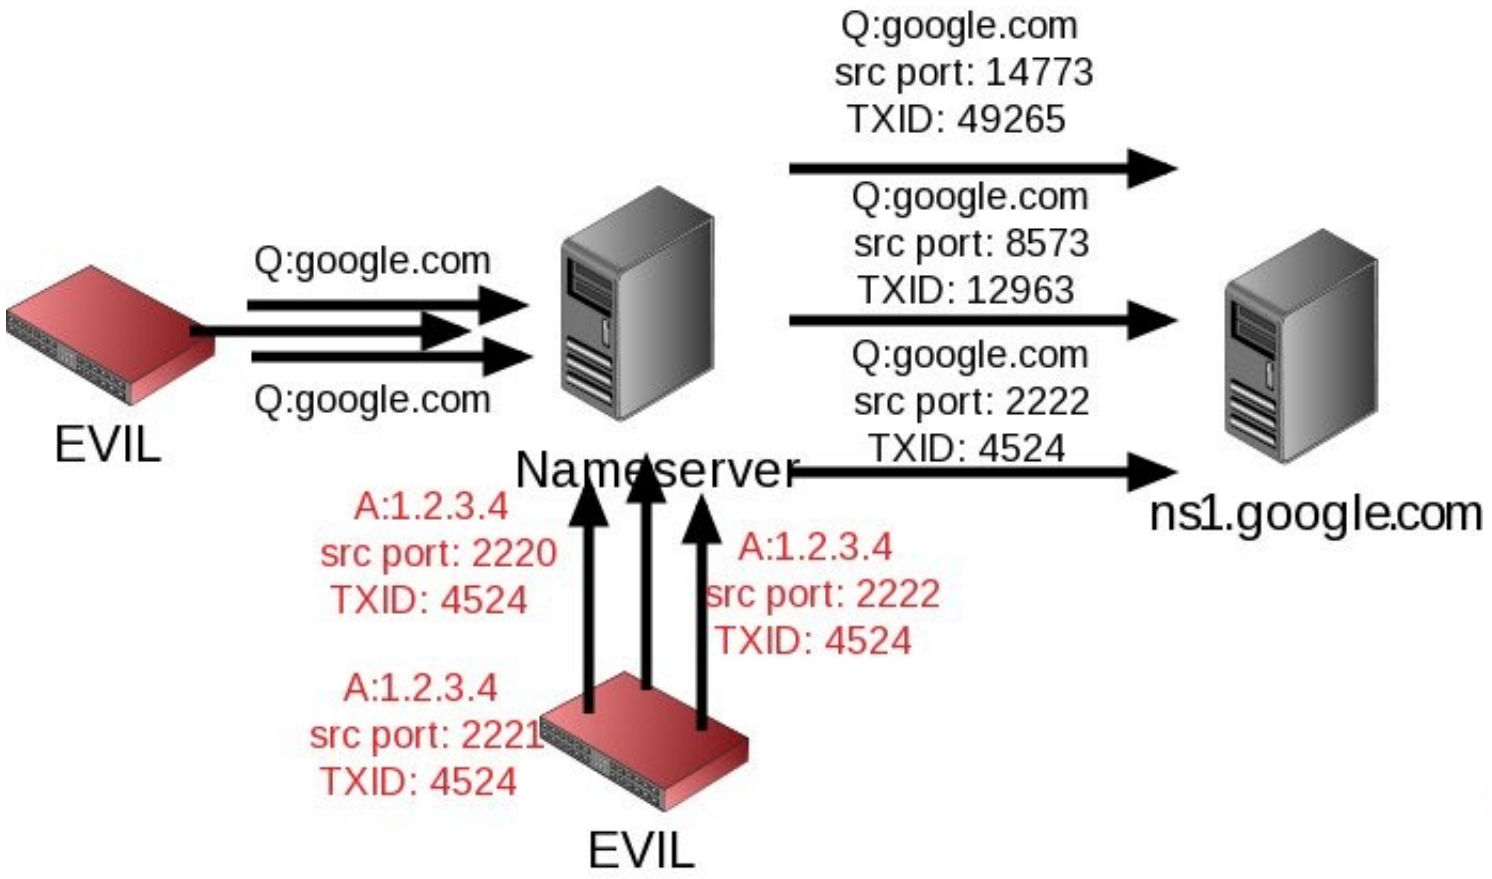

Kasphureff's attack (1997) caused Bailywick restrictions

QUESTION SECTION
Query ID: 54321
Question: www.ripe.net?

ANSWER SECTION
www.ripe.net = 193.0.0.195 (ttl=172800)

AUTHORITY SECTION
ripe.net NS ns-pri.ripe.net.
ripe.net NS ns-ext.isc.org.

ADDITIONAL SECTION
www.paypal.com A 1.2.3.4 (ttl=FOREVER)
google.com NS ns.myevildomain.com.

What protects our DNS?

- Transaction ID (TXID)
- Time To Live (TTL)
- Bailywick

The Kaminsky Attack

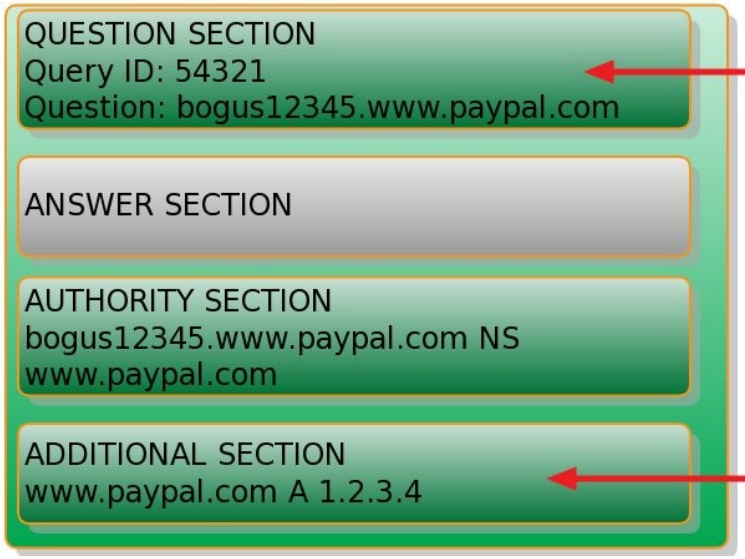

If you lose the race, try bogus12346

Overrides cache

Without source port randomization, this only takes about 65535 packets

DNS related issues: Double Fast Flux

- Botnets use domains with NS and A records with low (eg 3 minute) TTL's
- Change NS records via Registrar very quickly too (hours)
- This makes them next to impossible to shutdown.
-

DNS related issues: The Wifi hotspot

- Captive portals using DNS with mini DNS “server”
- This is so they can serve fake DNS
- This can cause client to cache wrong DNS
- Bad implementations break on EDNS and DNSSEC (hardcoded bits checking)
- Use transparent IP proxy instead

DNS related issues: Double Fast Flux

- Botnets use domains with NS and A records with low (eg 3 minute) TTL's
- Change NS records via Registrar very quickly too (hours)
- This makes them next to impossible to shutdown.
-

Where to fix the DNS ?

- Authorative nameservers
- Recursive nameservers
- Network firewalls and IDS
- Applications

- Protect the data or transport ?

DNS is critical infrastructure

- Backwards compatible (opt-in)
- Non-invasive or intrusive (drop-in)
- Non-disruptive (no CPU/Bandwidth hog)
- No Protocol changes(we have DNSSEC)
- Preferably no TYPE overloading
- No magic such as untested crypto
- Patent / Royalty free

Authorative nameservers

- Upgrade server to allow DNSSEC
- Diversify your infrastructure

```
; <<>> DiG 9.6.0a1 <<>> -t ns xelerance.com
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 57177
;; flags: qr rd ra; QUERY: 1, ANSWER: 3, AUTHORITY: 0, ADDITIONAL: 2

;; QUESTION SECTION:
;xelerance.com.                 IN      NS

;; ANSWER SECTION:
xelerance.com.          844     IN      NS      ns2.xelerance.org.
xelerance.com.          844     IN      NS      ns0.xelerance.nl.
xelerance.com.          844     IN      NS      ns1.xelerance.net.

;; ADDITIONAL SECTION:
ns0.xelerance.nl.       972     IN      A       193.110.157.135
ns1.xelerance.net.      98036   IN      A       209.237.247.134

;; Query time: 118 msec
;; SERVER: 193.110.157.2#53(193.110.157.2)
;; WHEN: Sat Jan 31 12:05:29 2009
;; MSG SIZE  rcvd: 142
```

Network IDS / Firewall

- It's patch work (pun intended)
- Does not address the problems
- Cannot make a decision when an attack is detected. What to do? Blocking is bad (denial of service to yourself)

- Monitor, log and warn. Do not interfere
- Be very careful with DNS load balancers

Monitor Unix based DNS

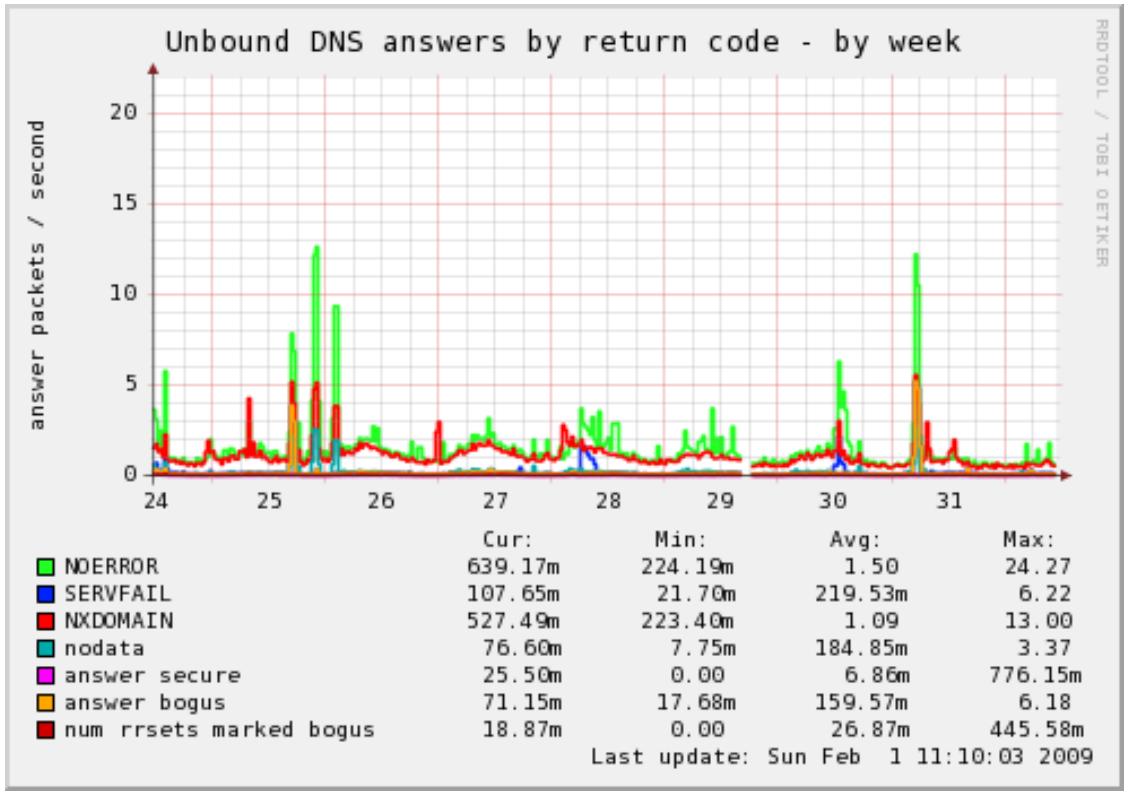

Monitoring using Cisco

Application fixes

- So many different applications to fix
- DNS API for applications is poor
- Easy to fool: DNS Rebinding or Fast Flux
- But let's not build DNS recursive nameservers in every application

  (however a good recursive dns server on each host is a good solution)

The inevitable:
Fix recursive nameservers

- Port randomization
- Sanitize TTL's
- Use more IP addresses per DNS server
- Harden against bogus size packets
- Harden glue
- Additional queries for infrastructure data
- 0x20

Birthday Attack protection

- Do not allow multiple queries for the same question to be outstanding (AKA query chaining)

  (Unbound and PowerDNS support this)

Rebinding protection

- Allow to specify IP addresses that may never appear in “external” domain names

  This way you can ensure 10.1.1.0/24 would never come in through DNS rebinding.
  (supported in Unbound and PowerDNS)

The inevitable:
Fix recursive nameservers

- RFC 5452 “Measures for Making DNS More Resilient against Forged Answers”
- draft-wijngaards-dnsext-resolver-side-mitigation
- draft-vixie-dnsext-0x20

Attacks can be detected

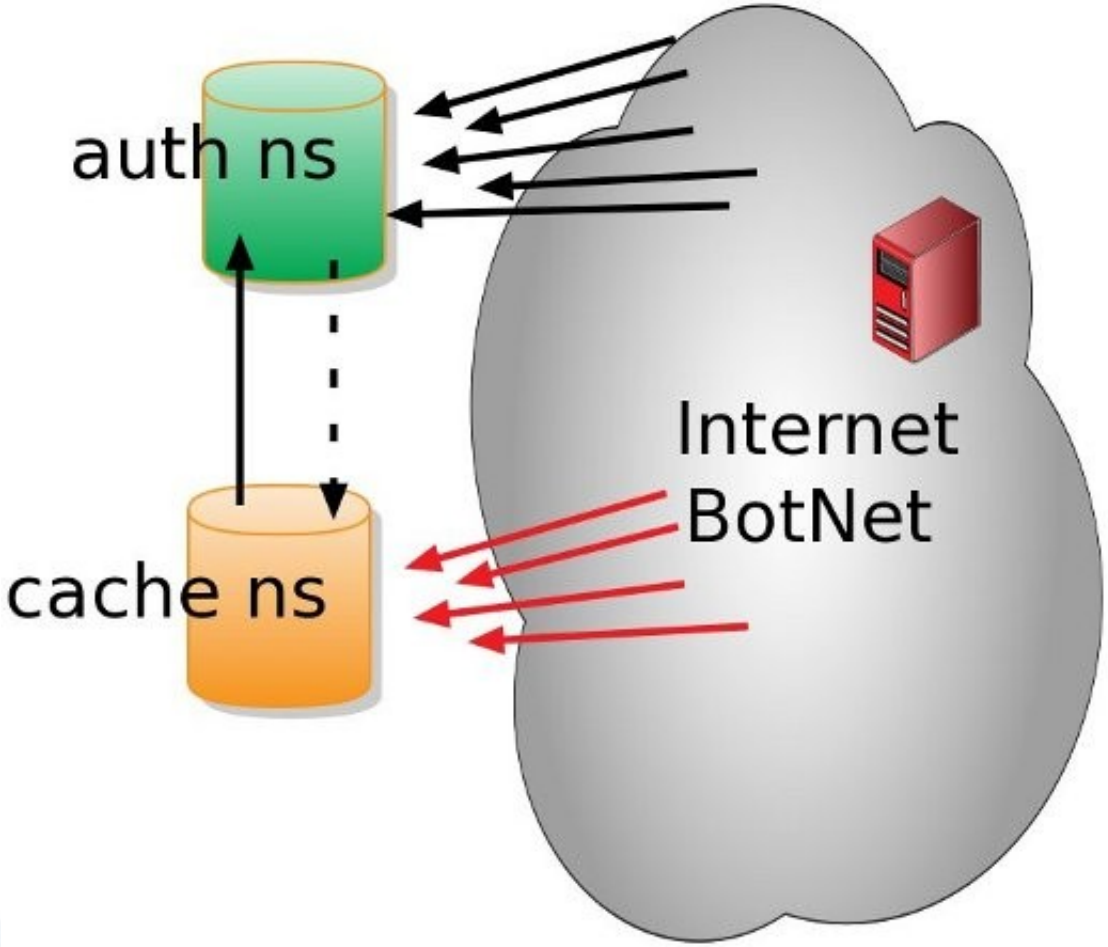

Attack response #1

- At a spoof detection threshold, ignore all answers for that query
- Prevents accepting the right forged answer
- Also prevents accepting the **real** answer spoofmax=?
- Small value : easy DOS
- Large value: might be too late (PowerDNS has spoofmax=20)

Attack response #2

- At a spoof detection threshold throw away the **entire** cache and start from scratch
- Prevents using an accepted forged answer
- Small value : easy DOS on the cache
- Large value: might be too late (Unbound has spoofmax=10M)

Add more NS records?

- If you already have at least two or three, this does not buy you much
- Only makes an attack marginally harder

- Excessive NS records cause other problems (and adds more potentially outdated / vulnerable nameservers)

Chain your caches (esp. the ones behind NAT)

Blacklist IP ranges

- Do not allow certain IP ranges (used internally) to be part of an answer from a public DNS zone not under our control
- This prevents DNS rebinding attacks
- Example: only allow 10.0.0.0/8 in ourdomain.com, nowhere else.

Hardening infrastructure queries

- Before accepting NS records or A records of nameservers, ask **at least** two different nameservers.
- Before accepting glue records or additional data, indepedantly verify these with new queries.

(extra work is only needed once, then we use caching – minimum impact)

Double Fast Flux protection

- Draft-bambenek-doubleflux suggests:

- Replacing the TTL's of NS and A records of NS records with TTL=72 hours.
- Llimit Registrar changes to once per 72h
- Recursors and clients should drop NS or A of NS with TTL < 12

The 0x20 defense (Paul Vixie)

- You don't need “Td-CaNAdaTRuSt.cOm” when you can get “.CoM”
- Fails completely for the root (“.”)

The 0x20 defense (Paul Vixie)

DNS Question: bogus12345.www.paypal.com?
Option flags: RD

The 0x20 defense (Paul Vixie)

DNS Question: bogus12345.www.paypal.com?
Option flags: RD

DNS Query ID: **54321**
DNS Question: bOGus12345.WwW.pAYpaL.Com

The 0x20 defense (Paul Vixie)

DNS Query ID: **54321**
DNS Question: bOGus12345.WwW.pAYpaL.Com

QUESTION SECTION
Query ID: **54321**
Question: BoGUs12345.wWW.pAYPal.cOM

ANSWER SECTION

AUTHORITY SECTION
bogus12345.www.paypal.com NS
www.paypal.com

ADDITIONAL SECTION
www.paypal.com A 1.2.3.4

DNSSEC in a nutshell

- Show DNSSEC signed zone

DNSSEC Lookaside Verification

DNSSEC Bonus

- Offline secure authenticated wireless communication with rendezvous / zeroconf / bluetooth

xelerance

www.xelerance.com/dnssec/

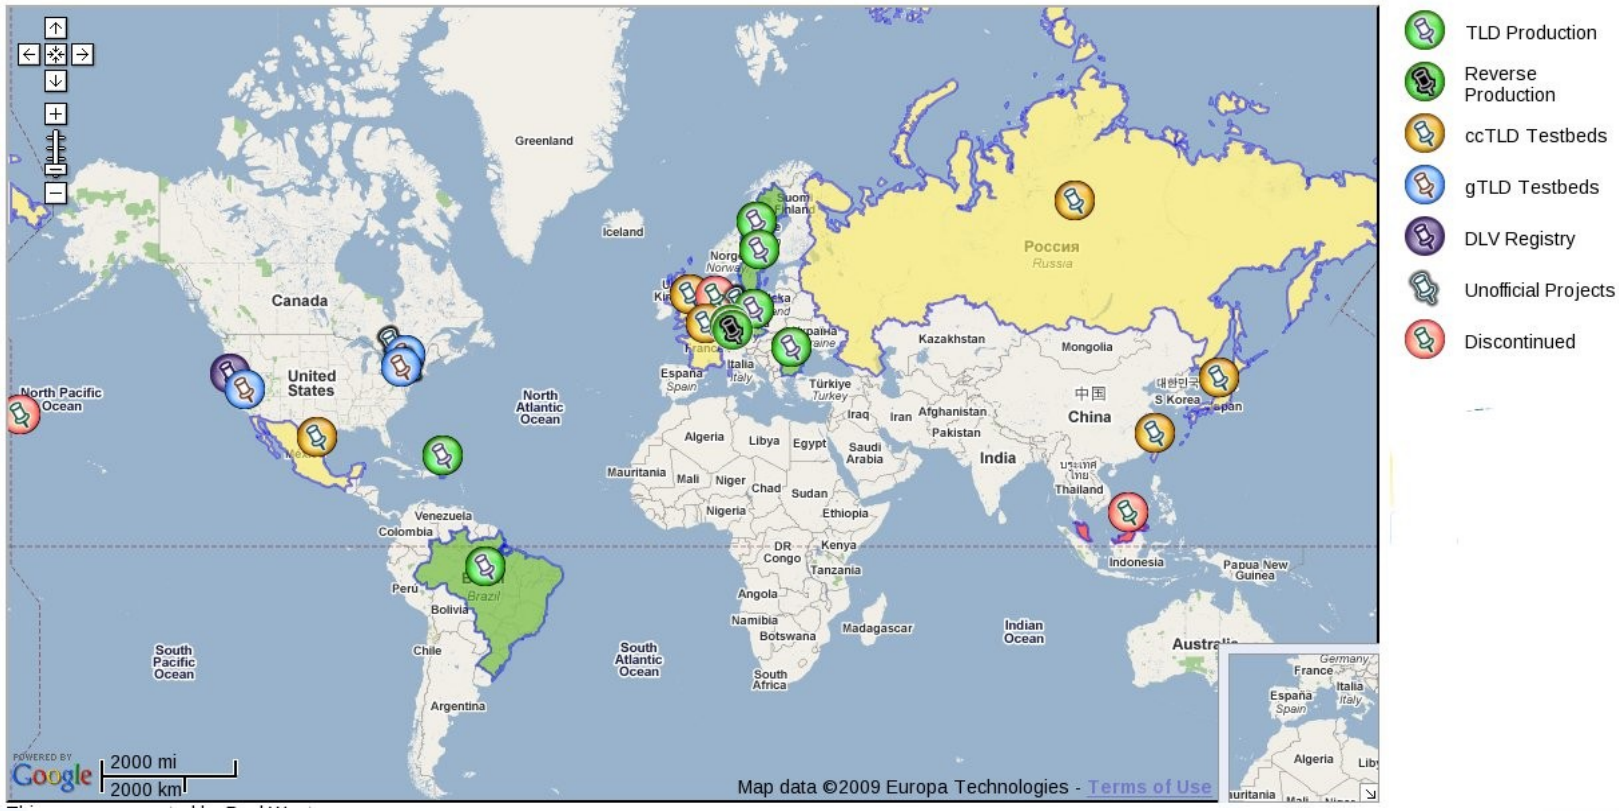

ccNSO survey Nov 2007

- If you have not implemented DNSSEC, are you planning to implement it?

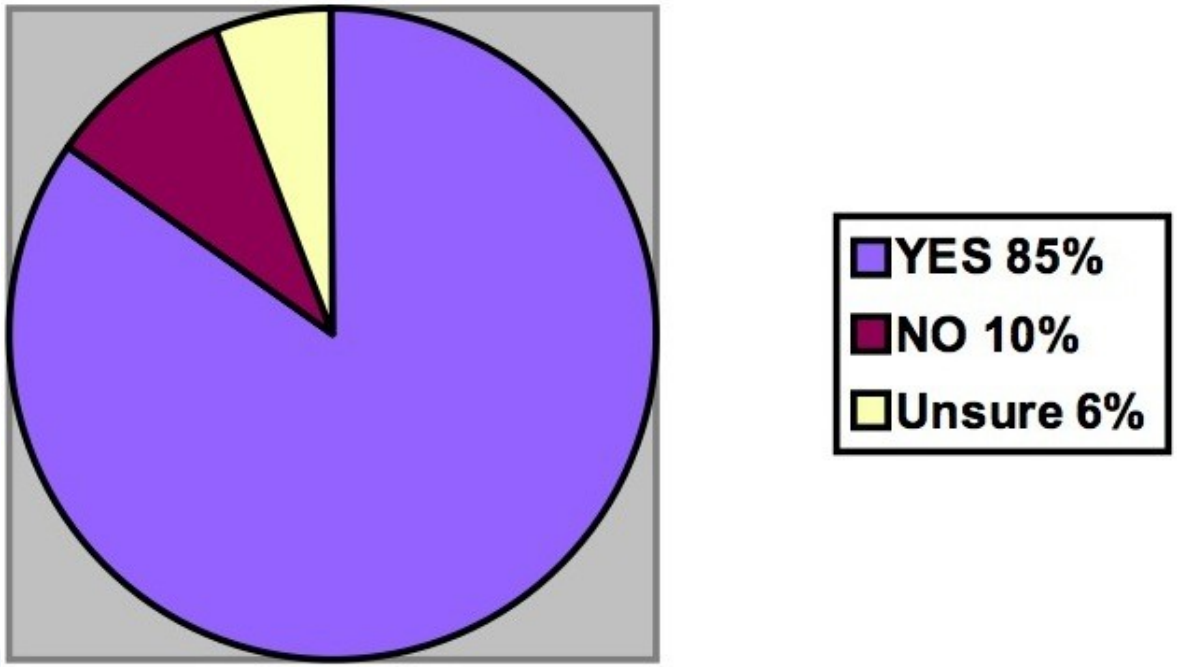

ccNSO survey Nov 2007

- If you have not implemented DNSSEC, when are you planning to implement it?

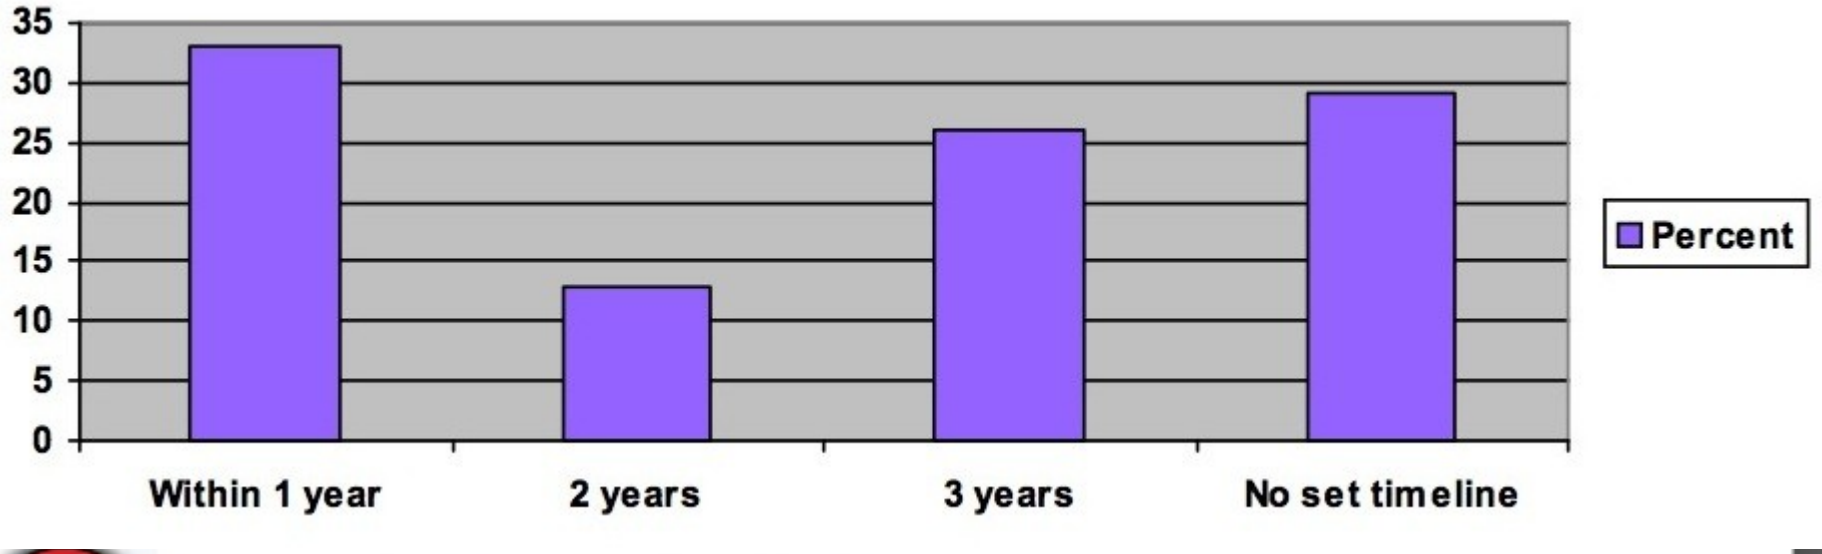

.gov is signed!

well, when I made these slides, it was not,
but now you read it, it should be

DNSSEC for All Top Level .GOV Domains

Published: August 29th, 2008 | Category: Security Vulnerabilities

Last week the Office of Management and Budget released memoranda M-08-23, titled Securing the Federal Government's Domain Name System Infrastructure. The document states that all US government top level .gov domains will use DNSSEC starting in January 2009. This is in response to the DNS cache poisoning attack that Dan Kaminsky made public a few months ago.

New Policy

This memorandum addresses two important issues in following through with the existing policy and expanding its scope to address all USG information systems.

A. The Federal Government will deploy DNSSEC to the top level .gov domain by January 2009. The top level .gov domain includes the registrar, registry, and DNS server operations. This policy requires that the top level .gov domain will be DNSSEC signed and processes to enable secure delegated sub-domains will be developed. Signing the to level .gov domain is a critical procedure necessary for broad deployment of DNSSEC, increases the utility of DNSSEC, and simplifies lower level deployment by agencies.

B. Your agency must now develop a plan of action and milestones for the deployment of DNSSEC to all applicable information systems. Appropriate DNSSEC capabilities must be deployed and operational by December 2009. The plan should follow recommendations in NIST Special Publication 800-81 "Secure Domain Name System

DNS-OARC Domain Name System Operations, Analysis, and Research Center

[Feb 2 2009 meeting information here]

February 3-4, 2009

Global DNS Security, Stability, and Resiliency Symposium

DNS
OARC

[Feb 3-4 meeting information here]

www.govsecinfo.com

The Keys to Deploying DNSSEC: Managing and Meeting Your OMB Domain Name

Thursday, March 12, 2009

Session: 8:30AM - 4:30PM

Presented by:

DNSSEC Development Coordination Initiative

The DNSSEC Deployment Initiative works to encourage all sectors to voluntarily adopt security measures that will improve security of the internet's naming infrastructure, as part of a global, cooperative effort that involves many nations and organizations in the public and private sectors.

Conclusions (1)

- **Update** your nameservers, or place them behind new nameservers.
- Look into more software then just Bind
  - **Unbound**, PowerDNS recursor
- Take a fresh look at your deployment, even when using firewalls and NAT. DNS **will** go through those.

- Ditch DNS captive portals and broken DSL routers

Conclusions (2)

- **Prepare for DNSSEC**

- **Tell your vendor you require DNSSEC on the endnode that uses a dhcp obtained DNS forwarder.**

xelerance

Questions?