nominet

Rollover and Die?

George Michaelson, APNIC

Geoff Huston, APNIC

Patrik Wallström, IIS

Roy Arends, Nominet UK

We’re under attack!!!

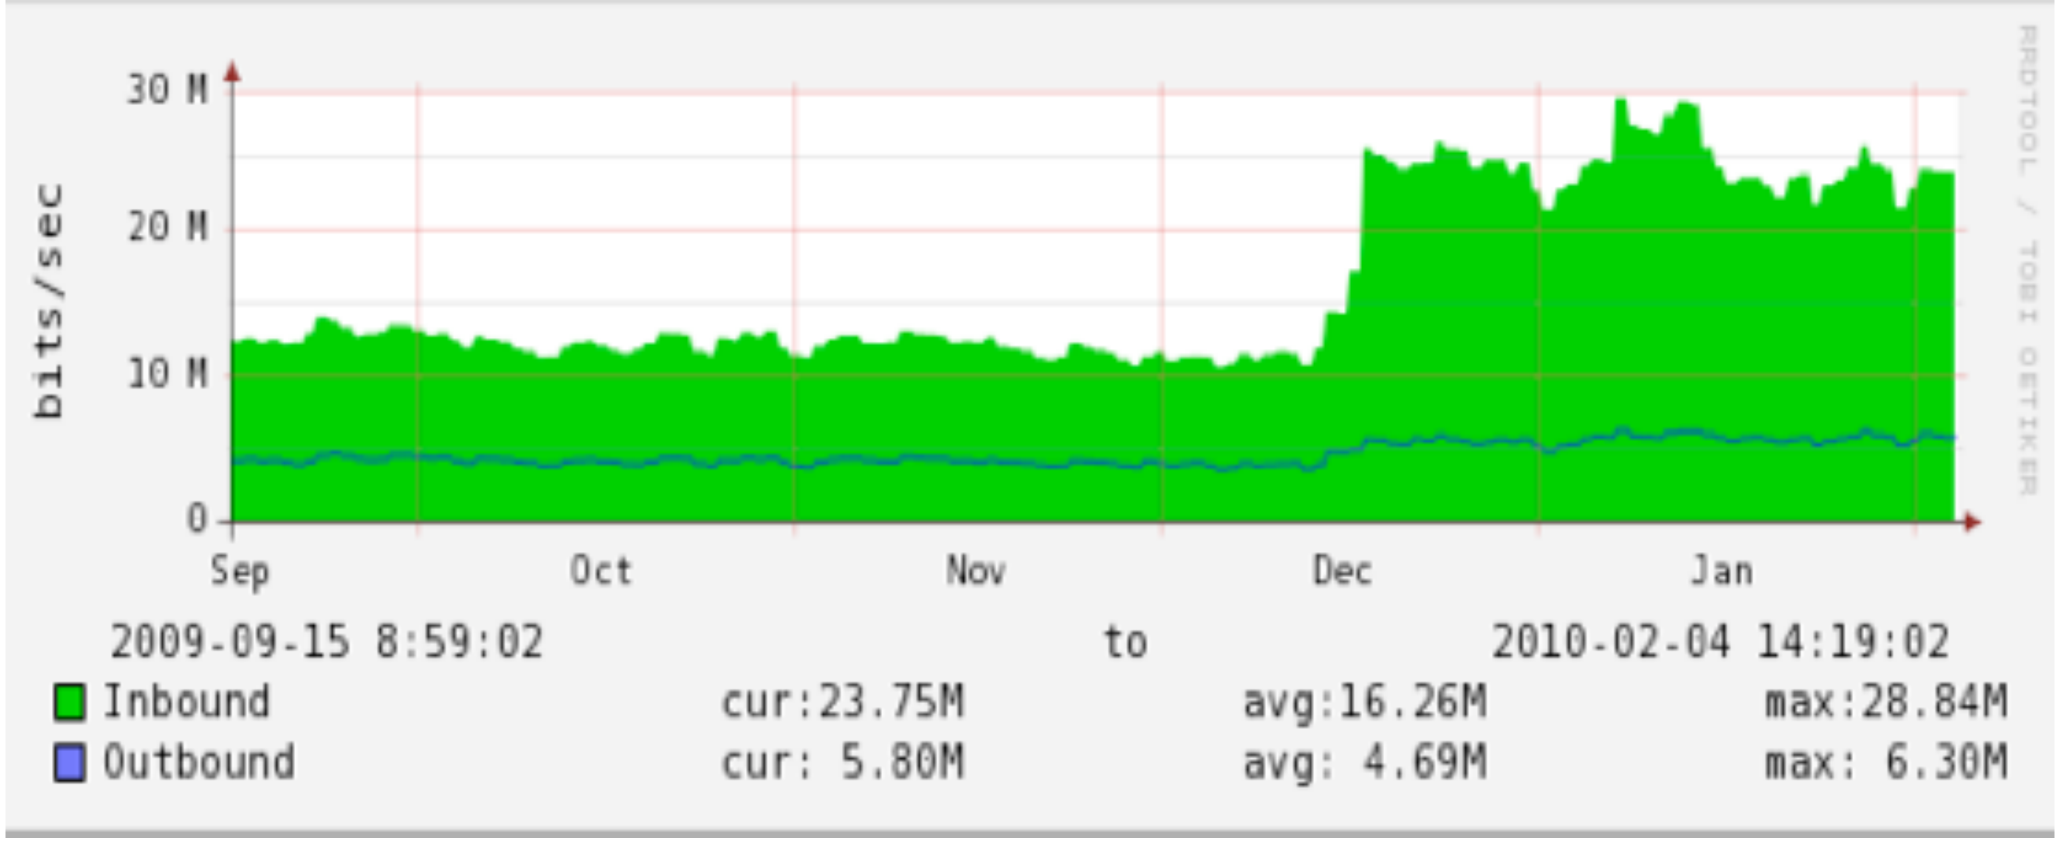

On the 16th of december, traffic more than doubled

DNSKEY amplification attack

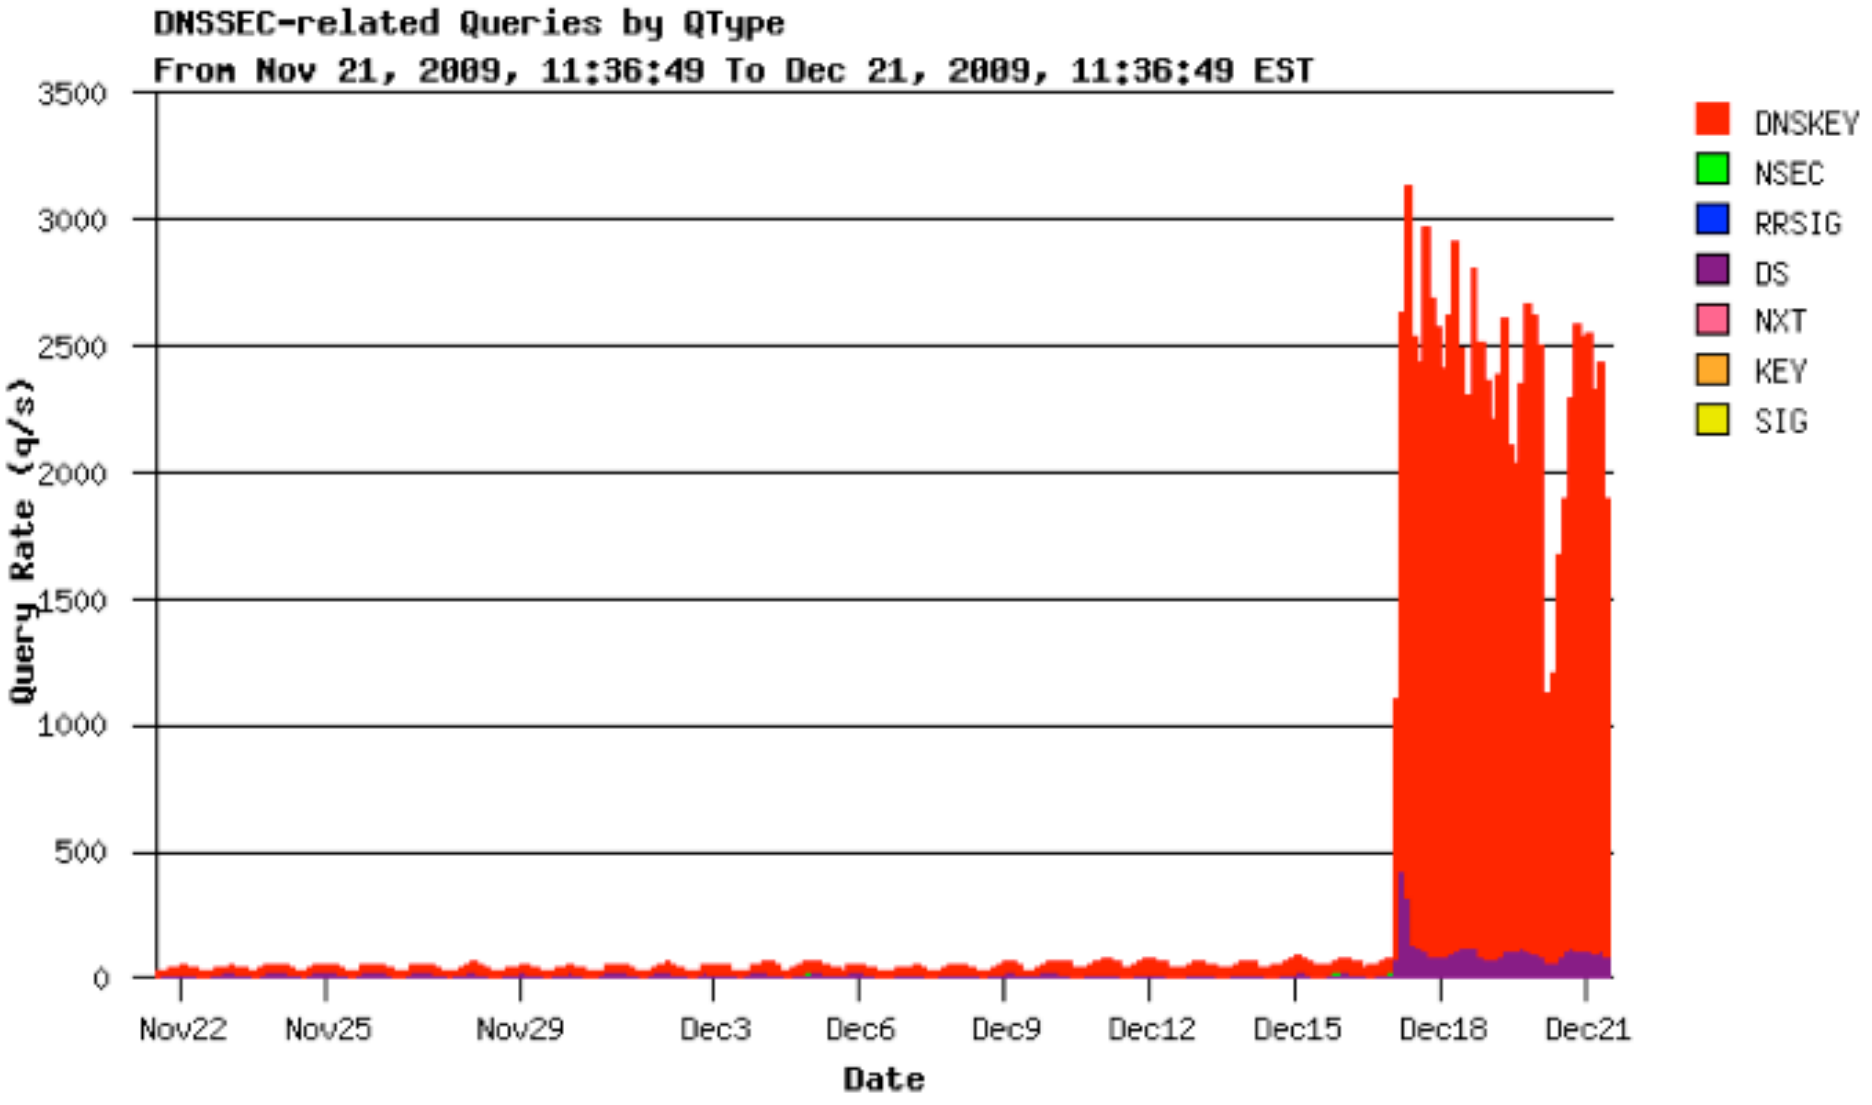

DNSKEY response size

Response size:	990 Bytes
Query rate:	2000 qps

15.8 Mbps

Additional load

Who does this?

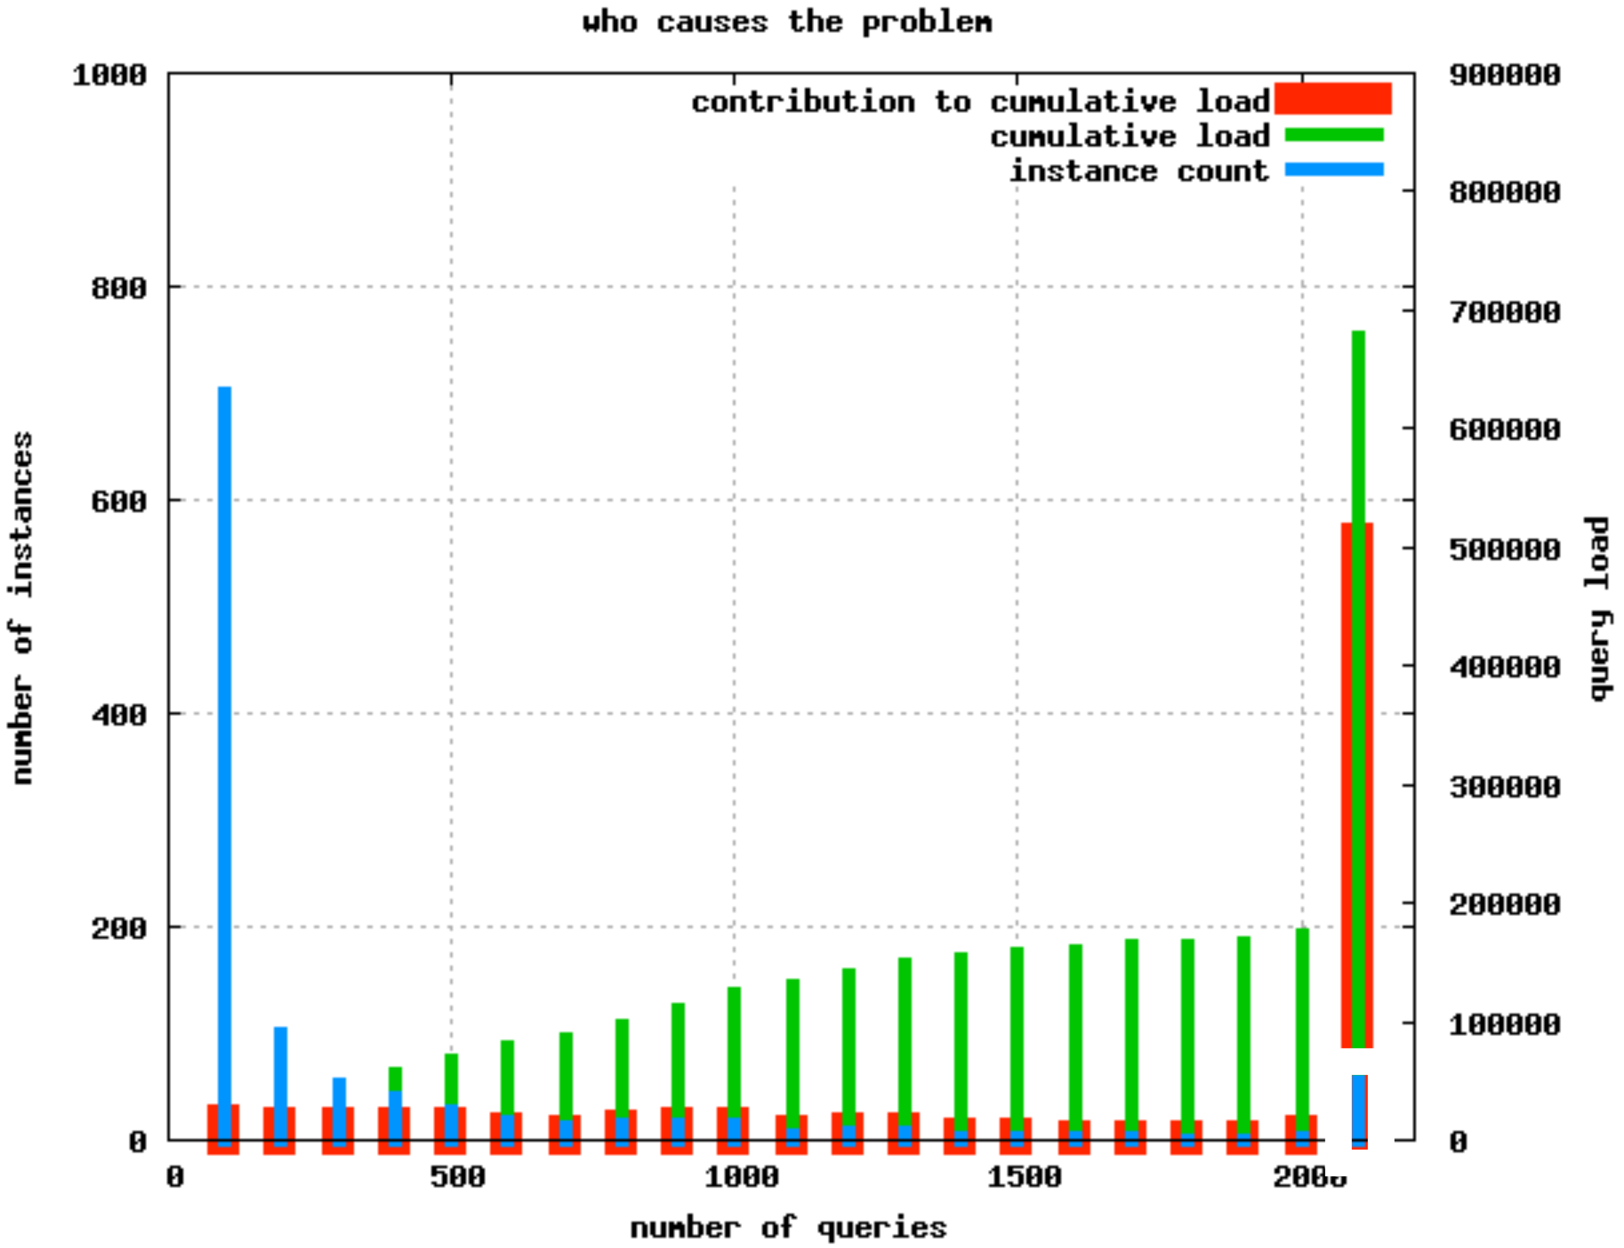

nominet

What was special about the 16th?

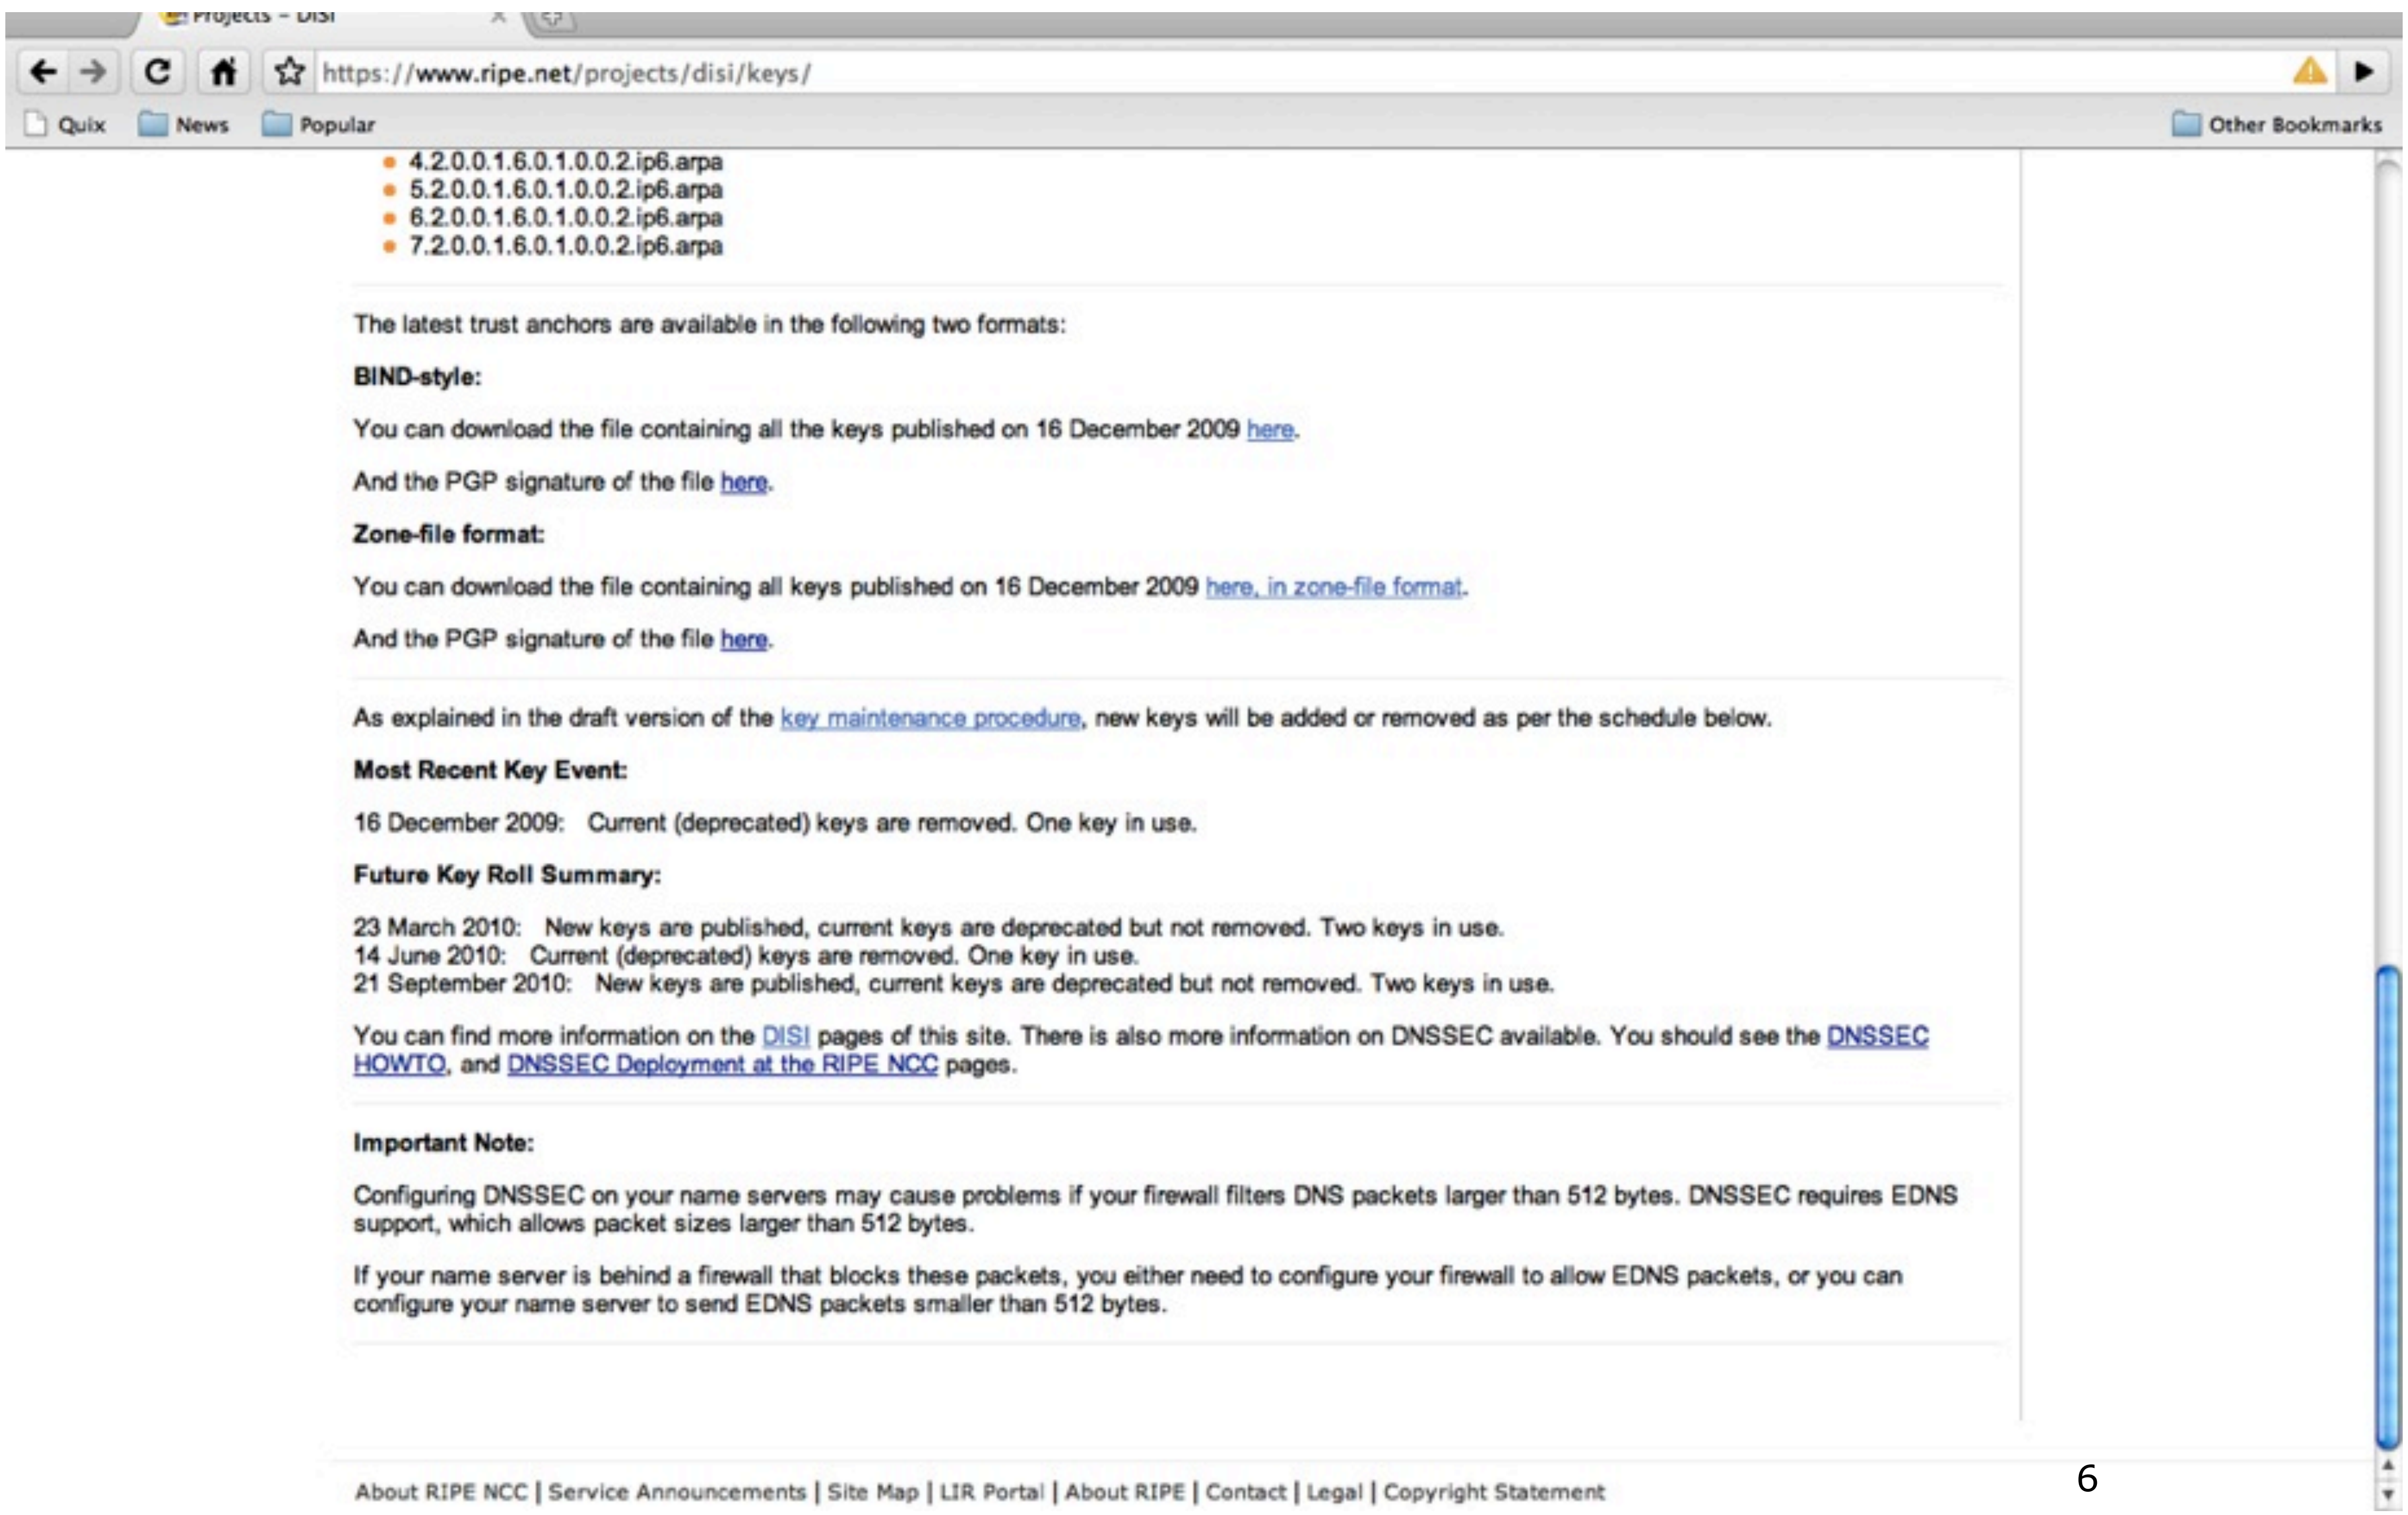

- 4.2.0.0.1.6.0.1.0.0.2.ip6.arpa
- 5.2.0.0.1.6.0.1.0.0.2.ip6.arpa
- 6.2.0.0.1.6.0.1.0.0.2.ip6.arpa
- 7.2.0.0.1.6.0.1.0.0.2.ip6.arpa

The latest trust anchors are available in the following two formats:

BIND-style:

You can download the file containing all the keys published on 16 December 2009 here.

And the PGP signature of the file here.

Zone-file format:

You can download the file containing all keys published on 16 December 2009 here, in zone-file format.

And the PGP signature of the file here.

As explained in the draft version of the key maintenance procedure, new keys will be added or removed as per the schedule below.

Most Recent Key Event:

16 December 2009: Current (deprecated) keys are removed. One key in use.

Future Key Roll Summary:

23 March 2010: New keys are published, current keys are deprecated but not removed. Two keys in use.
14 June 2010: Current (deprecated) keys are removed. One key in use.
21 September 2010: New keys are published, current keys are deprecated but not removed. Two keys in use.

You can find more information on the DISI pages of this site. There is also more information on DNSSEC available. You should see the DNSSEC HOWTO, and DNSSEC Deployment at the RIPE NCC pages.

Important Note:

Configuring DNSSEC on your name servers may cause problems if your firewall filters DNS packets larger than 512 bytes. DNSSEC requires EDNS support, which allows packet sizes larger than 512 bytes.

If your name server is behind a firewall that blocks these packets, you either need to configure your firewall to allow EDNS packets, or you can configure your name server to send EDNS packets smaller than 512 bytes.

About RIPE NCC | Service Announcements | Site Map | LIR Portal | About RIPE | Contact | Legal | Copyright Statement

nominet

What was special about the 16th?

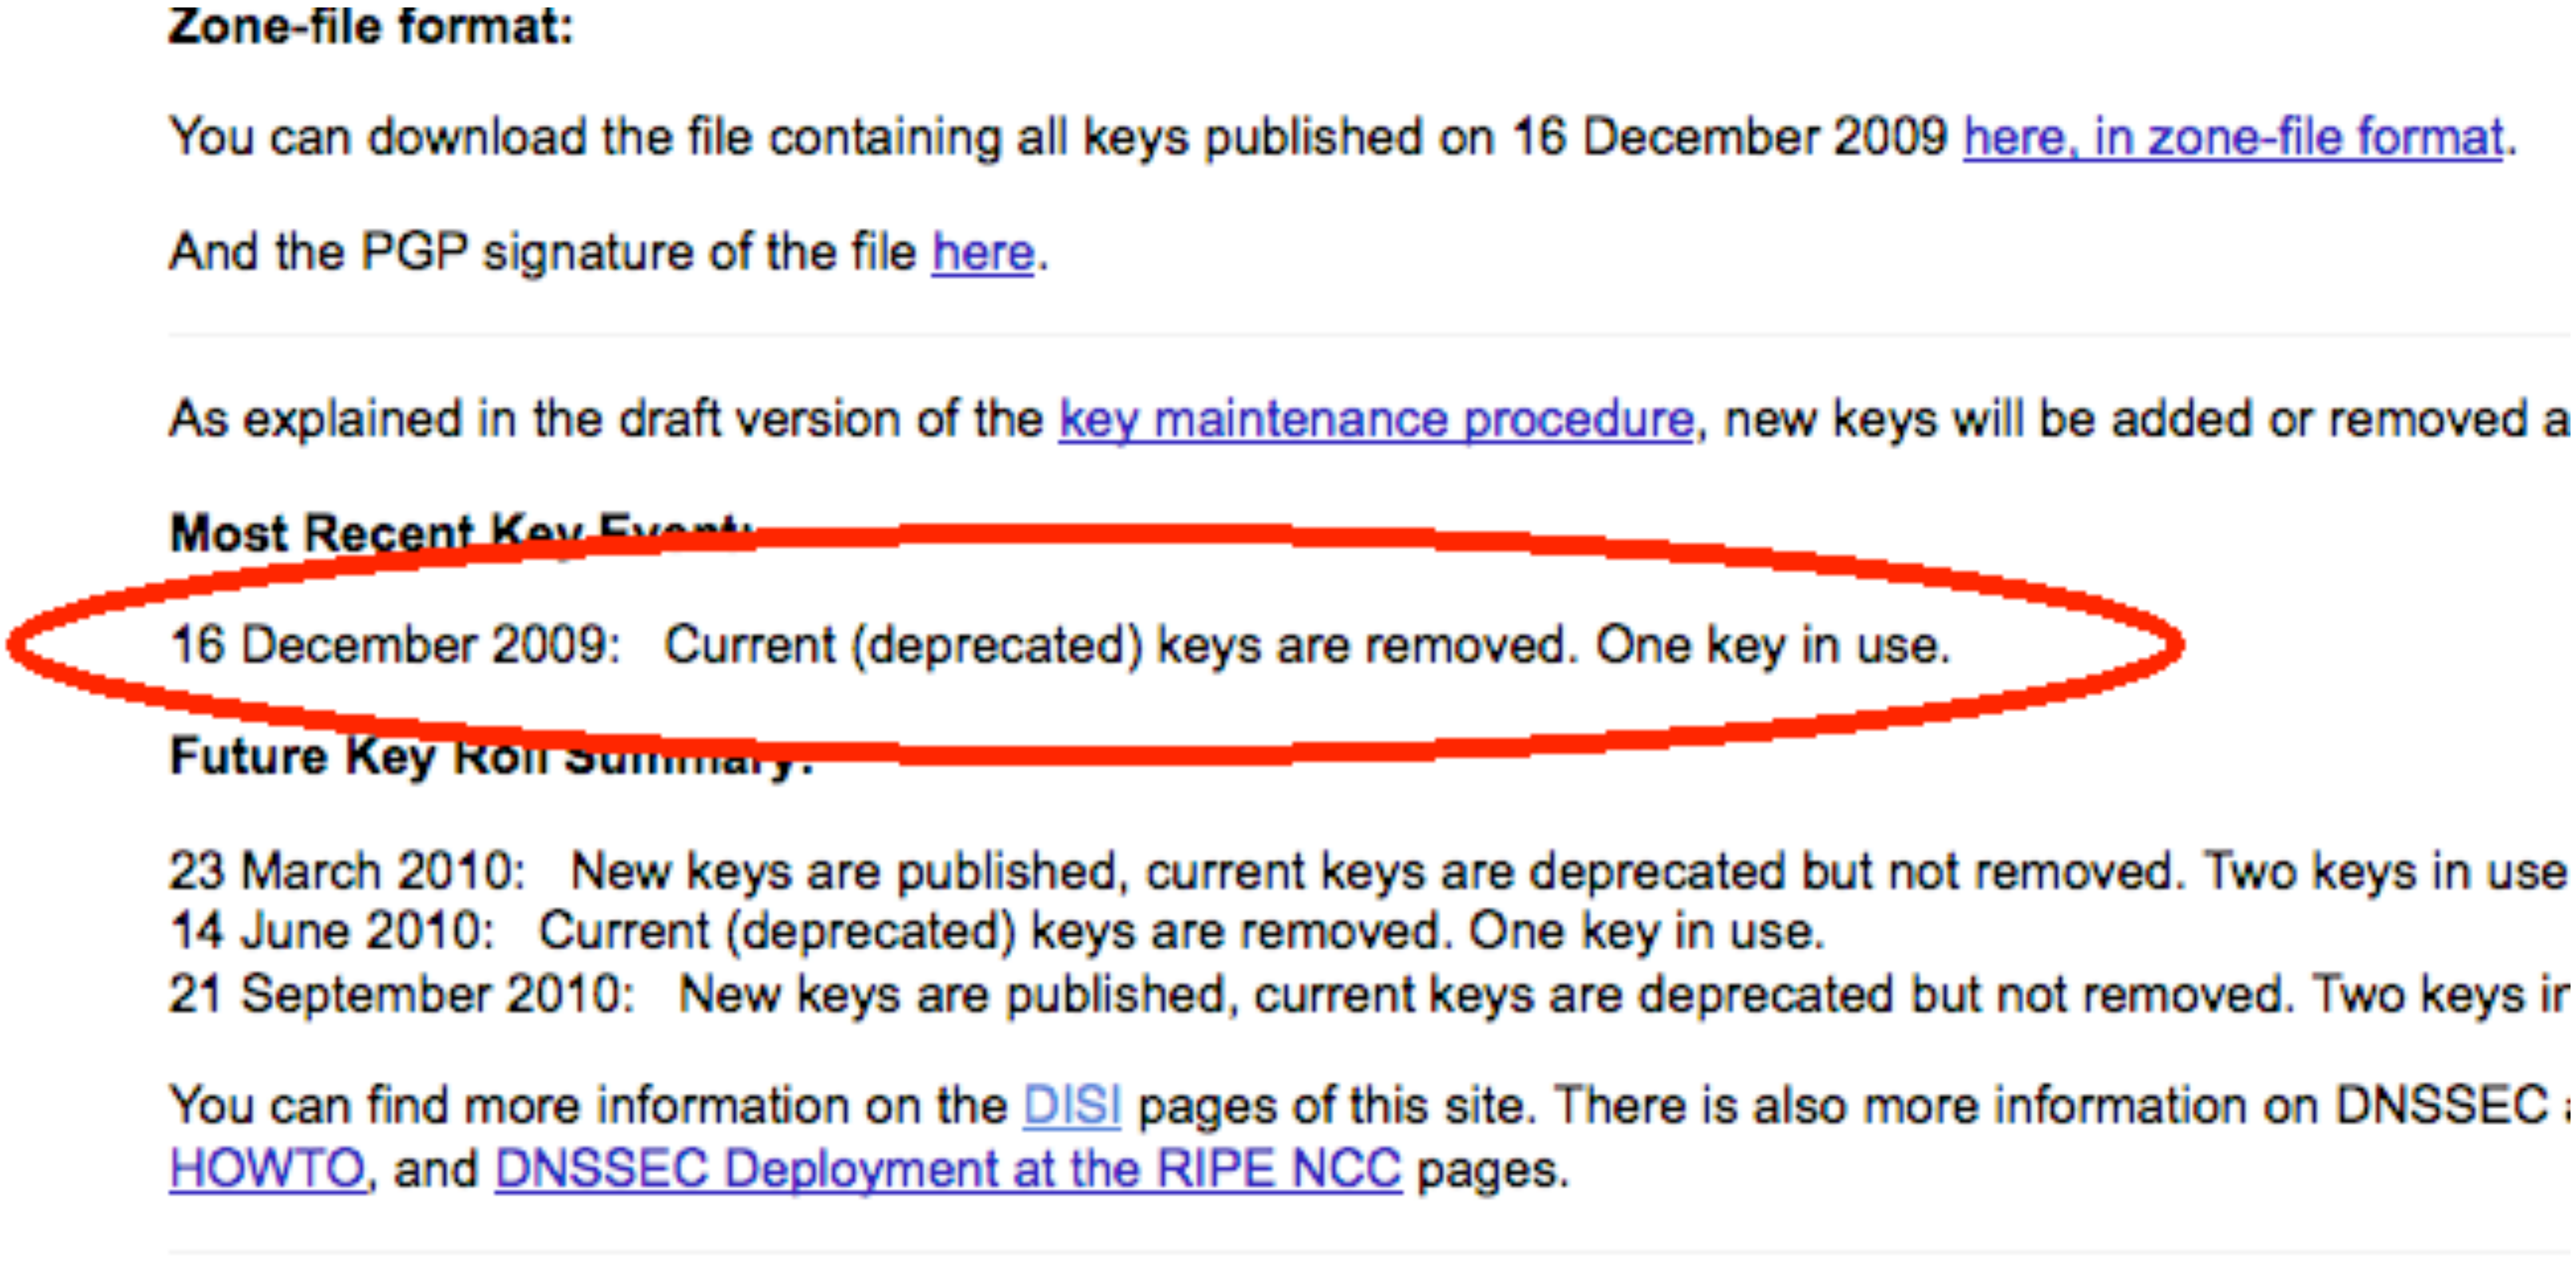

Zone-file format:

You can download the file containing all keys published on 16 December 2009 here, in zone-file format.

And the PGP signature of the file here.

As explained in the draft version of the key maintenance procedure, new keys will be added or removed a

Most Recent Key Event:

16 December 2009: Current (deprecated) keys are removed. One key in use.

Future Key Roll Summary:

23 March 2010: New keys are published, current keys are deprecated but not removed. Two keys in use
14 June 2010: Current (deprecated) keys are removed. One key in use.
21 September 2010: New keys are published, current keys are deprecated but not removed. Two keys ir

You can find more information on the DISI pages of this site. There is also more information on DNSSEC a HOWTO, and DNSSEC Deployment at the RIPE NCC pages.

Hanlon's razor

Never attribute to malice that which can be explained by stupidity.

Why so many clients?

- Fedora bug report 17th Jan 2010
  - (1 month after the roll)
- operator reports getting 240.000 log entries in 24hr
  - “no valid key”
- dnssec-conf tool contained a hard-configured trust anchor file
  - obsolete after the 16th.

What was special about the 16th?

what a great lesson

Randy Bush’s response

Current load for in-addr.arpa

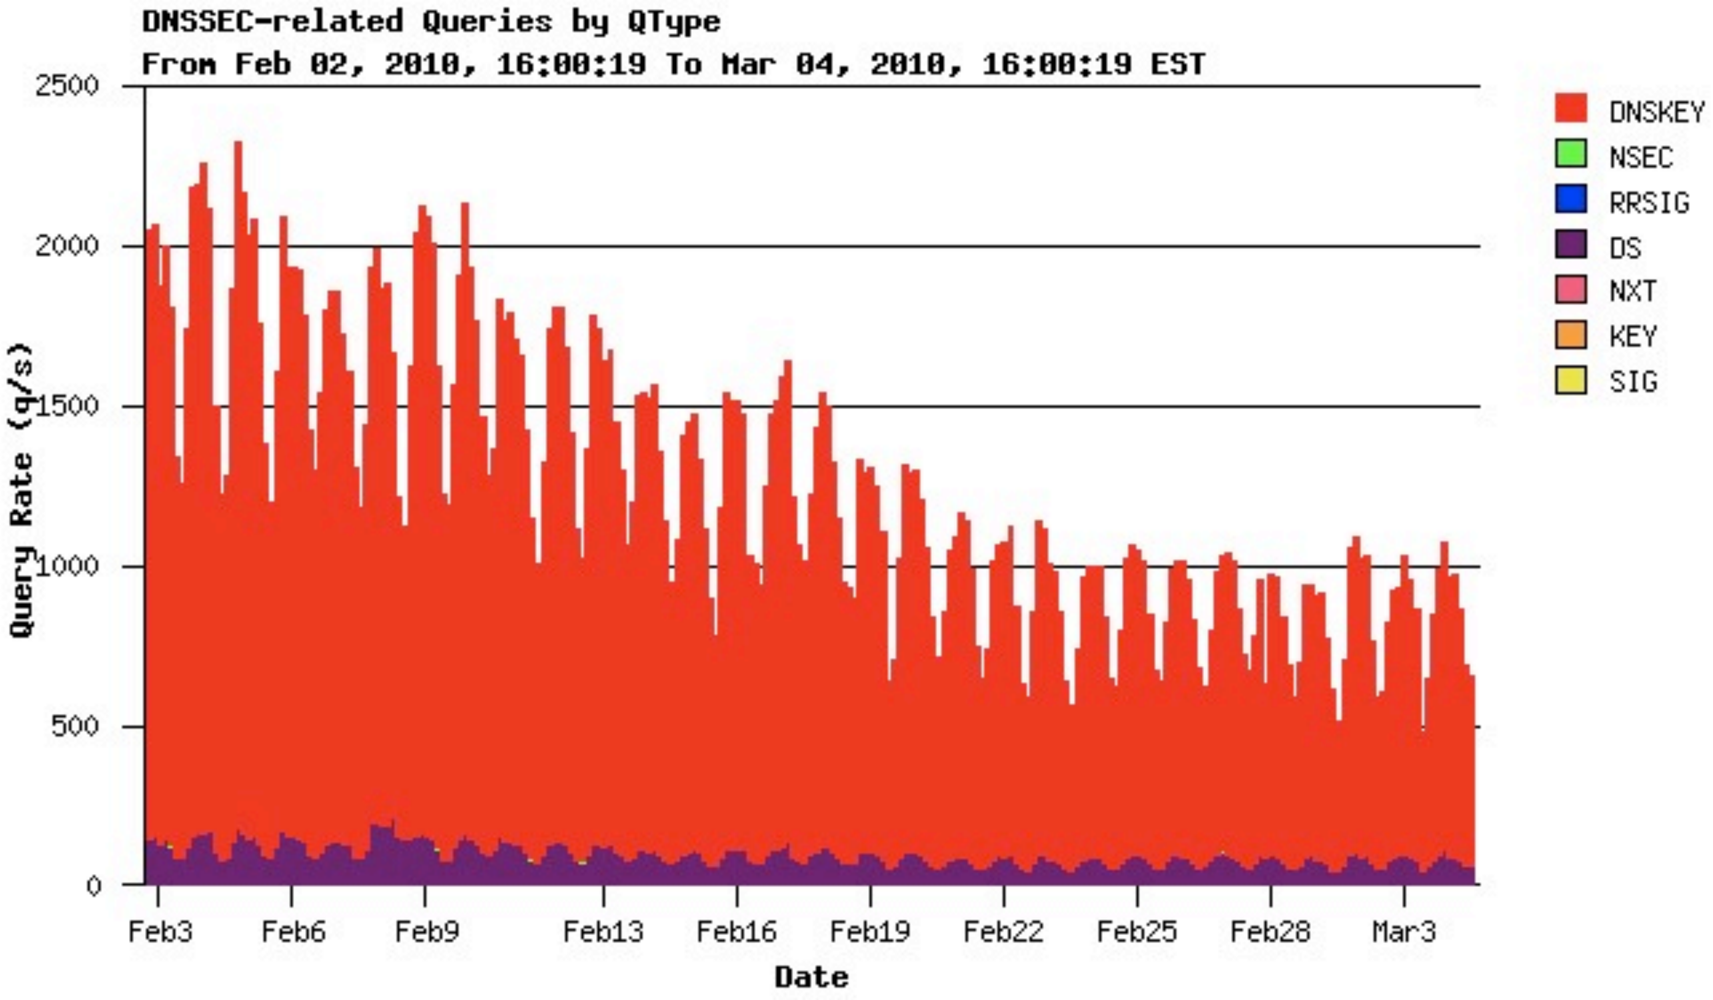

getting better, below 1000 qps right now

But decline not fast enough before new roll

The Load Projection

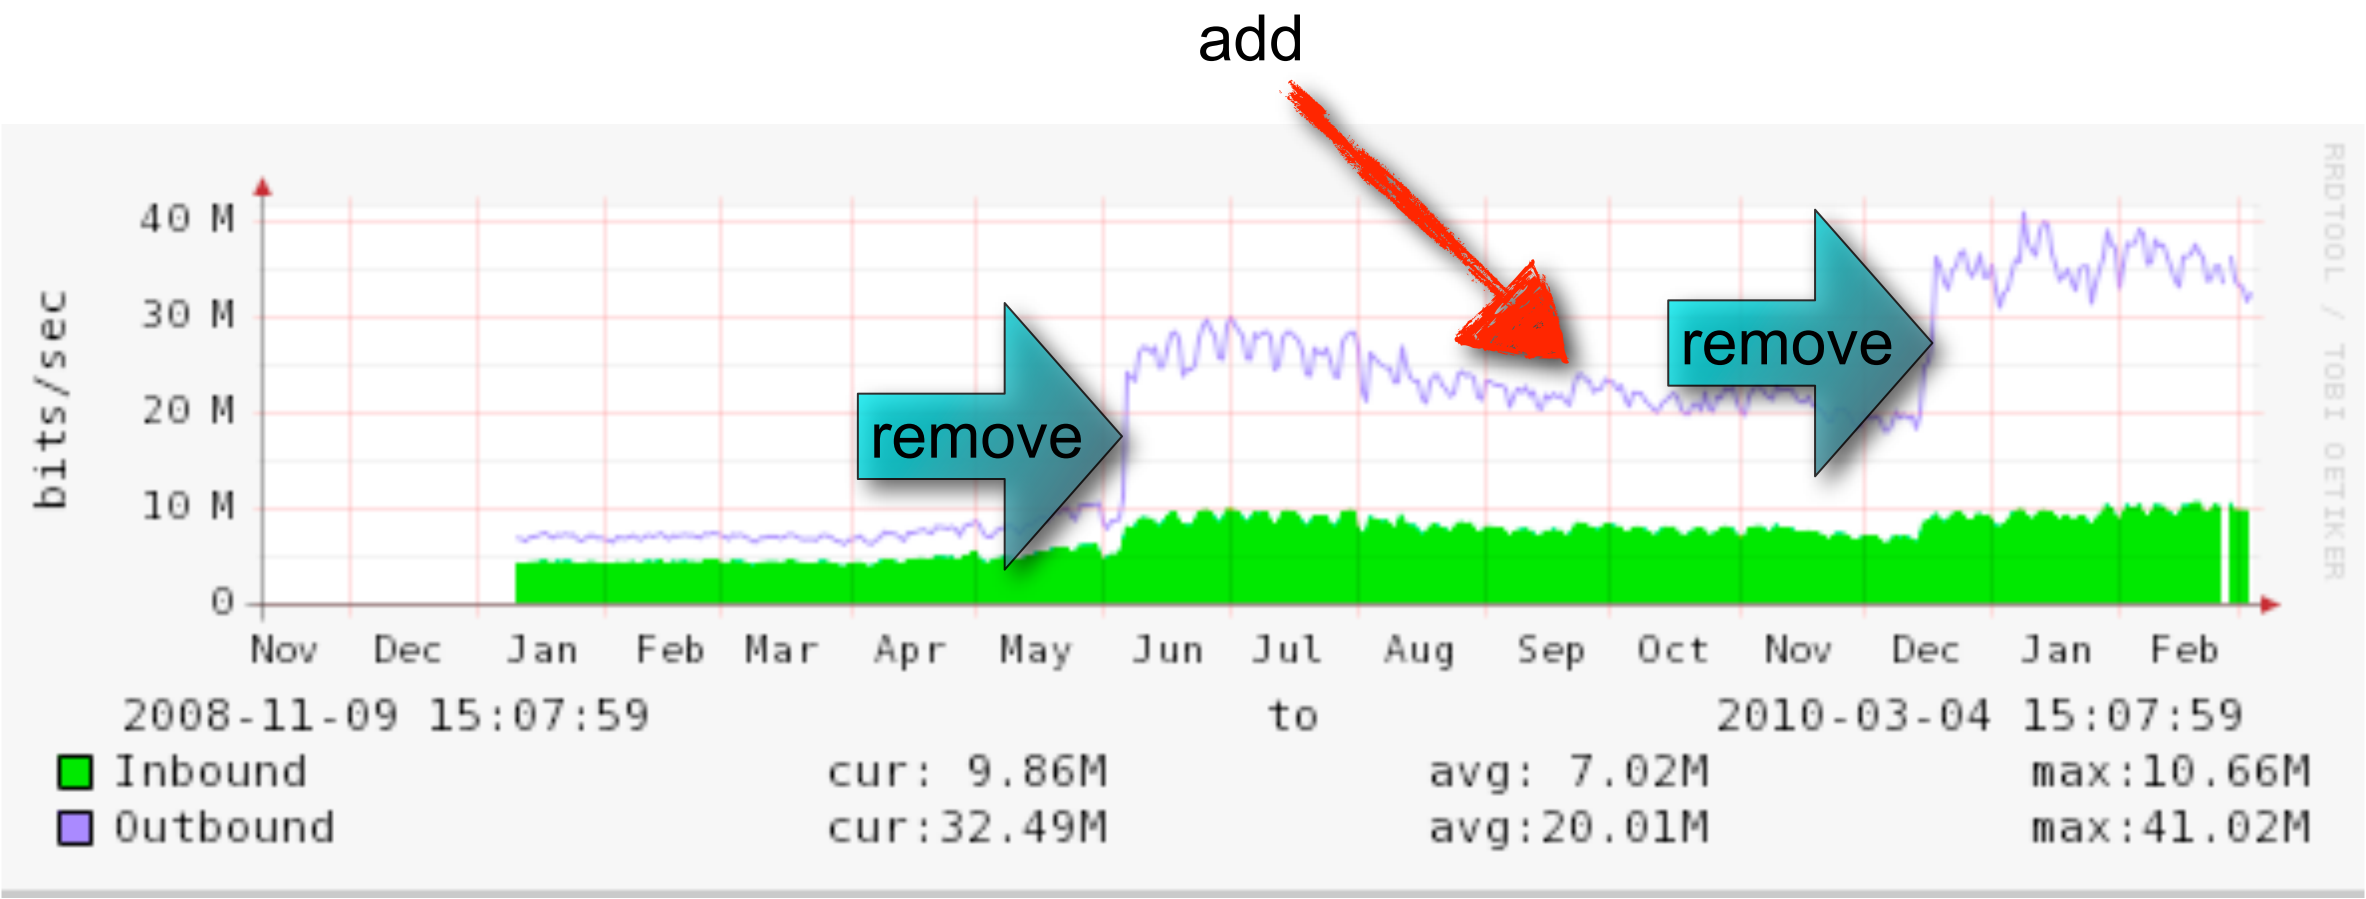
add

remove

remove

Was this a one off event?

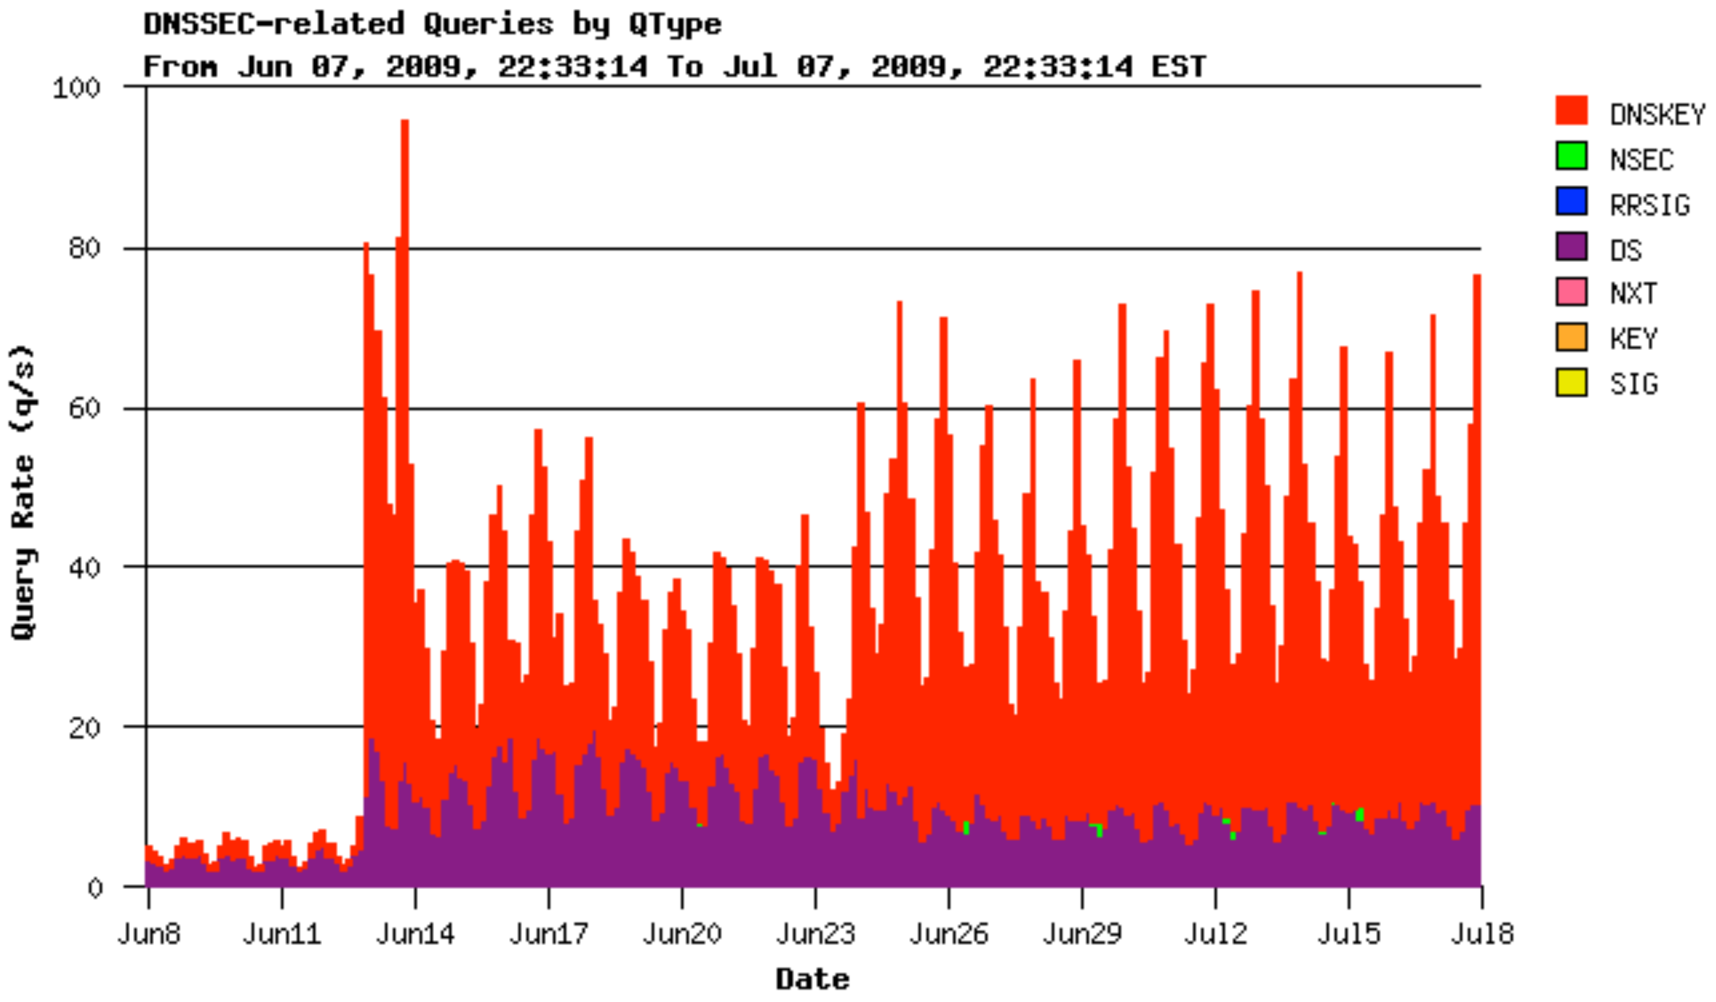

Sweden, june 2009

Was this a one off event?

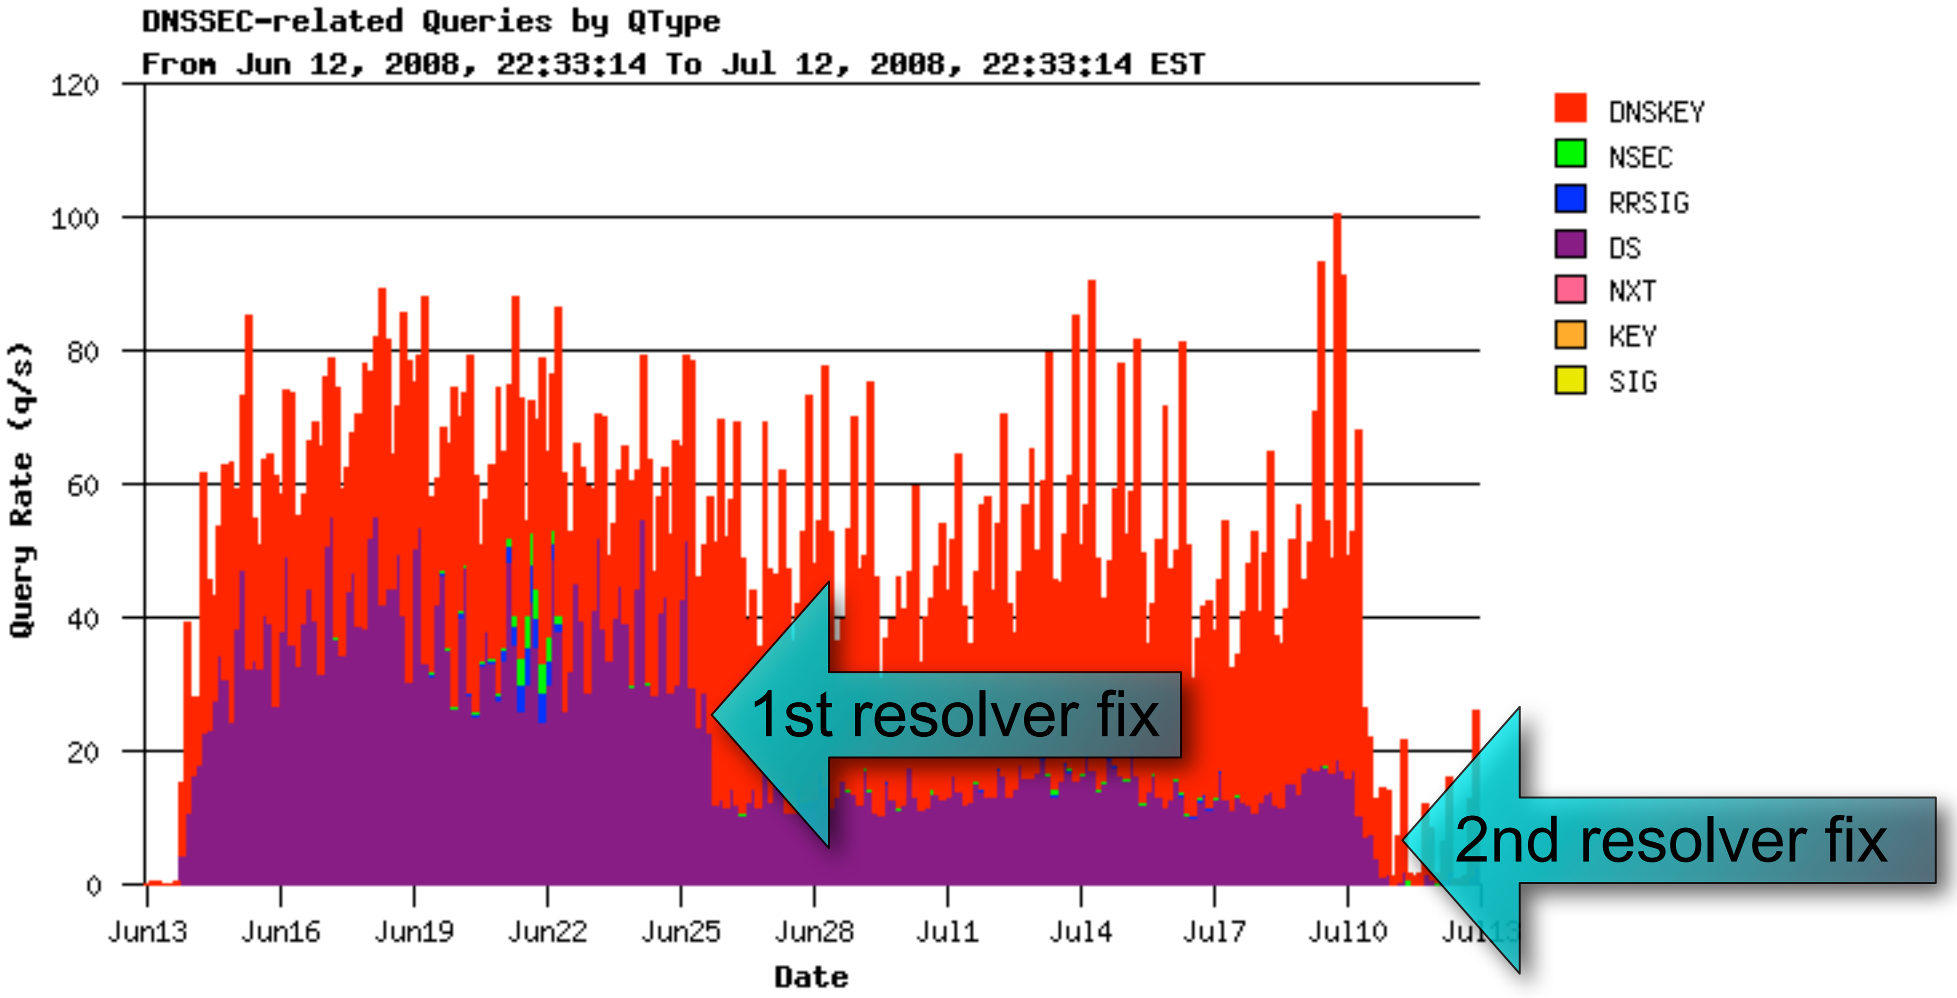

1st resolver fix

2nd resolver fix

Sweden, june 2008

Why so many Queries?

- Resolvers are supposed to cache dnskey
- Even when those are bad
- Resolvers should be nice, not aggressive
- So many resolvers, so few servers

Why so many Queries?

- Bind bug in all versions
- Depth First Search (DFS) problem
- Chain of trust validation:

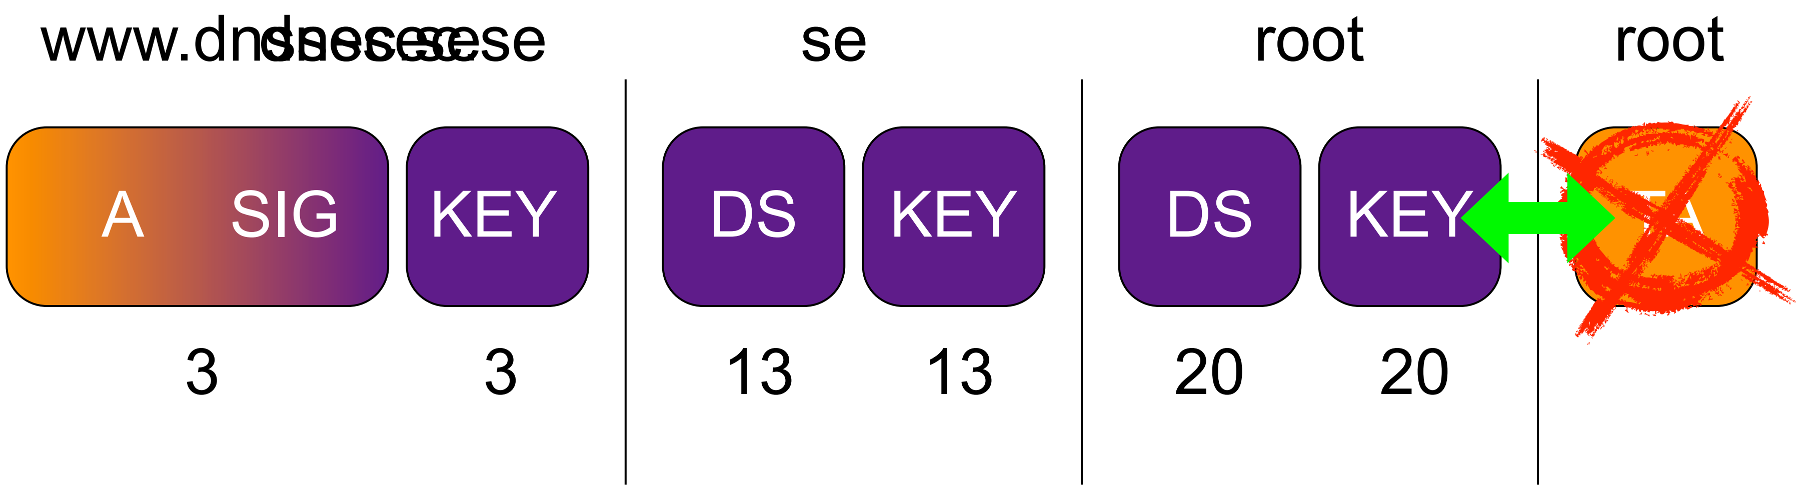

3 * 3 * 13 * 13 * 20 * 20 = 608400 queries

ISC

- Reported the depth first search bug on februari 8th
- Acknowledged the problem
  - fundamental fix, needs thorough testing.
- released BIND 9.7.0 & 9.6.2
  - first version that can validate the root
  - Exercise caution
- ISC released the patched versions 15th march.

The Perfect Storm

- DNSSEC deployment at root (DURZ)
  - guess what: lame trust-anchor, don’t configure

```
02:23:13.904447 IP 94.254.84.99.29484 > 192.112.36.4.53: 5784% [1au] DNSKEY? . (28)
02:23:14.063617 IP 94.254.84.99.56185 > 202.12.27.33.53: 58470% [1au] DNSKEY? . (28)
02:23:14.096800 IP 94.254.84.99.19540 > 192.33.4.12.53: 63411% [1au] DNSKEY? . (28)
02:23:14.202476 IP 94.254.84.99.23210 > 128.63.2.53.53: 43288% [1au] DNSKEY? . (28)
02:23:14.302964 IP 94.254.84.99.61614 > 193.0.14.129.53: 60641% [1au] DNSKEY? . (28)
02:23:14.443820 IP 94.254.84.99.39117 > 128.8.10.90.53: 52235% [1au] DNSKEY? . (28)
02:23:14.580610 IP 94.254.84.99.1832 > 192.228.79.201.53: 41792% [1au] DNSKEY? . (28)
02:23:14.749730 IP 94.254.84.99.42450 > 192.203.230.10.53: 52903% [1au] DNSKEY? . (28)
02:23:14.934376 IP 94.254.84.99.32392 > 199.7.83.42.53: 48480% [1au] DNSKEY? . (28)
02:23:15.073805 IP 94.254.84.99.18993 > 192.5.5.241.53: 53794% [1au] DNSKEY? . (28)
02:23:15.083405 IP 94.254.84.99.18362 > 192.58.128.30.53: 32638% [1au] DNSKEY? . (28)
02:23:15.536684 IP 94.254.84.99.40824 > 198.41.0.4.53: 63668% [1au] DNSKEY? . (28)
02:23:17.237648 IP 94.254.84.99.43118 > 192.36.148.17.53: 20348% [1au] DNSKEY? . (28)
02:23:17.497613 IP 94.254.84.99.26253 > 192.112.36.4.53: 27565% [1au] DNSKEY? . (28)
02:23:17.541230 IP 94.254.84.99.13293 > 128.8.10.90.53: 14401% [1au] DNSKEY? . (28)
02:23:17.677963 IP 94.254.84.99.12985 > 192.58.128.30.53: 21457% [1au] DNSKEY? . (28)
02:23:17.686715 IP 94.254.84.99.47565 > 202.12.27.33.53: 11950% [1au] DNSKEY? . (28)
02:23:17.719576 IP 94.254.84.99.52505 > 193.0.14.129.53: 27749% [1au] DNSKEY? . (28)
02:23:17.744421 IP 94.254.84.99.12667 > 192.203.230.10.53: 10018% [1au] DNSKEY? . (28)
02:23:17.929291 IP 94.254.84.99.4109 > 128.63.2.53.53: 46561% [1au] DNSKEY? . (28)
```

The Perfect Storm

- No automatic trust anchor roll (5011)
  - 9.6.2 not planned
- DLV mishaps:
  - DLV registry promiscuously scrapes TLD keys
    - Just another chain of trust
  - .PR rolled its key
    - was unavailable to DLV users for days
    - caused a major packet storm

The Perfect Storm

- Multiple trust anchor problem
  - TLD Trust Anchors trump Root Trust Anchor
    - stale TLD Trust Anchor trumps valid Root Trust Anchor

- Doom scenario:
  - TLD registers DS in root
  - new policy: don’t announce rolls, depend on root
    - That is the way NS records works as well
  - Operators won’t update TLD trust anchor anymore
    - Why would they, they’ve configured root trust-anchor

A Series Of Unfortunate Events

- buggy “dnssec provisioning” software
- DNSSEC @ root
- multiple trust anchor problem
- no 5011 deployment
- Frequent Rollover Syndrome
  - rolling rolling rolling, keep them DNSKEYs rolling.

Frequent Rollover Syndrome

- Advice seems to be:
  - roll the key as often as you can
  - Some roll twice a year, some roll monthly
- Advice is misguided:
  - too many sigs do not leak the key.
  - Intention is to mitigate a compromised key fallout
  - no perfect forward security
- If a key can be compromised in 1 year, it can be compromised in 6 months for twice the cost
- Other reasons: educate operators, exercise procedures
  - all irrelevant, never mess with a critical production system

Solution

- Stop and test DNSSEC provisioning software.
- Don’t roll keys (too often)
  - be practical
- Do not endorse configuration of trust-anchors when parent is signed.
  - no 5011, no web-page with listed keys, no DLV, no ITAR
  - Manage all through a signed parent.
- When parent is not signed:
  - Use proper 5011. Use ISC’s DLV.

- Help fund development of ISC’s BIND-10.

Questions ? Remarks ? Observations ?

http://www.potaroo.net/ispcol/2010-02/rollover.html

Thanks to
Anand Buddhdev
Patrik Wallström
George Michaelson
Geoff Huston
David Conrad
Folks at ISC

Question: If you’ve deployed DNSSEC and rolled your (ksk) key, look at the stats around that period, and (pretty) please report them back to us.