The DNS Violations Project

I've seen things you people wouldn't believe.

Ondřej Surý • ondrej.sury@nic.cz • 11. 5. 2017

CZ.NIC | SPRÁVCE DOMÉNY CZ

DNS Violations Project

- DNS is complex!
- There's a lot of DNS protocol violations "out there"
- DNS Resolvers just "cope" with that
- We started doing internal list
- Idea: Let's make this community effort
- DNS Violations project was born

Purpose

- To better understand the breakages of the DNS protocol
- To make DNS better
- To share knowledge
- For Authoritative DNS implementors, to avoid common pitfalls
- For DNS Resolver implementors, to verify they can handle it
- Nothing sensitive (only public DNS information)

Common Violations

- CDNs have a special department in hell
- Garbage at the end of packet
- Case-sensitive DNS servers
- QNAME Minimization
  - Empty Non-Terminals: NXDOMAIN vs NODATA
- EDNS-related breakages
- General protocol ignorance
- DNSSEC-related brokeness
- Unknown RRType

Current Violations

Garbage at the end of packet

- DVE-2017-0001: KMPMedia nameservers returns garbage at the end of the packet
  - Looks like a static size buffer that's returned as whole (69 bytes)

```
;; Got bad packet: FORMERR
69 bytes
b2 38 85 a0 00 01 00 01 00 00 00 01 03 6c 6f 67            .8...........log
08 6b 6d 70 6c 61 79 65 72 03 63 6f 6d 00 00 01            .kmplayer.com…
00 01 c0 0c 00 01 00 01 00 00 00 1e 00 04 6e 2d            ..............n-
c3 d1 00 00 00 00 00 00 00 00 00 00 00 00 00 00            ................
00 00 00 00 00                                             .....
```

“Generated” NS records

- DVE-2017-0002: KMPMedia generate invalid nameservers
  - NS records generated on the fly
  - This might break strict QNAME minimization

```
;; QUESTION SECTION:
;; 8107.kmpmedia.net.          IN    NS

;; ANSWER SECTION:
8107.kmpmedia.net.      300   IN    NS    n1.8107.kmpmedia.net.
8107.kmpmedia.net.      300   IN    NS    n2.8107.kmpmedia.net.
8107.kmpmedia.net.      300   IN    NS    n5.8107.kmpmedia.net.
8107.kmpmedia.net.      300   IN    NS    n6.8107.kmpmedia.net.
8107.kmpmedia.net.      300   IN    NS    n7.8107.kmpmedia.net.
```

Case sensitive nameservers

- DVE-2017-0003: McAfee nameservers silently break when 0x20 randomization is used
  - So called 0x20 technique might be used to add more entropy to outgoing queries (make the case random per-letter)
  - local.cloud.mcaffee.com fails to provide correct answer if upper-case is used in QNAME

```
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 40072
;; Flags: qr aa rd; QUERY: 1; ANSWER: 1; AUTHORITY: 1; ADDITIONAL: 0
;; QUESTION SECTION:
;; b-0.19-23003008.1481.1518.19cf.3ea1.410.0.ekzijnekvvvg7gb38qcwur561b.avqs.mcafee.com.        IN    A

;; ->>HEADER<<- opcode: QUERY; status: NXDOMAIN; id: 49080
;; Flags: qr aa rd; QUERY: 1; ANSWER: 0; AUTHORITY: 1; ADDITIONAL: 0
;; QUESTION SECTION:
;; b-0.19-23003008.1481.1518.19cf.3ea1.410.0.ekzijnekvvvg7gb38qcwur561B.avqs.mcafee.com.        IN    A
```

Case sensitive nameservers

- DVE-2017-0006: CDNetworks Co. nameserver fails to honour 0x20
  - {ns1,ns2,ns3,ns4}.panthercdn.com. return RCODE=REFUSED when there's upper-case in QNAME

```
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 51999
;; Flags: qr aa rd; QUERY: 1; ANSWER: 0; AUTHORITY: 1; ADDITIONAL: 0
;; QUESTION SECTION:
;; p7677.cdngc.net.           IN    A

;; ->>HEADER<<- opcode: QUERY; status: REFUSED; id: 11762
;; Flags: qr aa rd; QUERY: 1; ANSWER: 0; AUTHORITY: 0; ADDITIONAL: 0
;; QUESTION SECTION:
;; p7677.CDNgc.net.           IN    A
```

Malformed packets on EDNS queries

- DVE-2017-0004: *.webcfs00.com nameservers produces malformed packets with EDNS
  - NOTE: EDNS was standardized 18 years ago!

```
$ dig +edns=0 @ns01.webcfs00.com. IN A
726170696473736c2d63726c.67656f7472757374.636f6d.80hc70747be.webcfs00.com.
;; Got bad packet: FORMERR
118 bytes
e4 1b 84 00 00 01 00 01 00 00 00 00 18 37 32 36       .............726
31 37 30 36 39 36 34 37 33 37 33 36 63 32 64 36       170696473736c2d6
33 37 32 36 63 10 36 37 36 35 36 66 37 34 37 32       3726c.67656f7472
37 35 37 33 37 34 06 36 33 36 66 36 64 0b 38 30       757374.636f6d.80
68 63 37 30 37 34 37 62 65 08 77 65 62 63 66 73       hc70747be.webcfs
30 30 03 63 6f 6d 00 00 01 00 01 00 00 00 00 00       00.com..........
00 00 00 00 00 00 c0 0c 00 01 00 01 00 00 38 40       ..............8@
00 04 cc d4 aa 69                                     .....i
```

Malformed packets on EDNS queries

- DVE-2017-0012: České Radiokomunikace a.s. CDN nameservers provide malformed response when EDNS is used
  - NOTE: EDNS was standardized 18 years ago!

```
$ dig +norec +edns=0 IN NS se04.se.prima-vod-prep-sec.service.cdn.cra.cz @sr01.cdn.cra.cz.
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOTIMP, id: 31825
;; flags: qr ad; QUERY: 1, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 0

;; WARNING: EDNS query returned status NOTIMP - retry with '+noedns'
;; WARNING: Message has 11 extra bytes at end

;; QUESTION SECTION:
;se04.se.prima-vod-prep-sec.service.cdn.cra.cz. IN NS

;; Query time: 0 msec
;; SERVER: 82.99.164.132#53(82.99.164.132)
;; WHEN: Thu Jan 26 10:38:50 CET 2017
;; MSG SIZE  rcvd: 74
```

CDN: Who cares about standards…

- DVE-2017-0007: České Radiokomunikace a.s. CDN nameservers fails to provide answer to NS queries
- DVE-2017-0011: GitBook CDN nameservers return A records on any QTYPE query
- DVE-2017-0007: České Radiokomunikace a.s. CDN nameservers fails to provide answer to NS queries
- DVE-2017-0007: (fixed) České Radiokomunikace a.s. CDN nameservers sets AD bit on every response

Breaks QNAME Minimization

- Empty Non-Terminals
  - www.**ent**.dns.rocks. IN AAAA → 2001:1488:0:3::5
    - ent.dns.rocks. → RCODE=NOERROR + SOA
- DVE-2017-0007: České Radiokomunikace a.s. CDN nameservers fails to provide answer to NS queries
  - dig IN NS se04.se.prima-vod-prep-sec.service.cdn.cra.cz @sr01.cdn.cra.cz. → RCODE=**NOTIMPL**
- DVE-2017-XXXX: Akamai nameservers return NXDOMAIN for ENTs
  - This was fixed, but the fix had to be reverted.
  - Now the fix is deployed in stages.
  - dig IN A gov.edgekey.net. @a12-65.akam.net. → RCODE=**NXDOMAIN**
  - RFC8020: NXDOMAIN: There Really Is Nothing Underneath

DNSSEC related: broken signing(?)

- DVE-2017-0009: axc.nl bogus DNSSEC denial of existence
  - Returns NODATA instead of NXDOMAIN
    - IN TLSA _25._tcp.cameras-kopen.nl @nszero1.axc.nl. → RCODE=**NOERROR**
  - Invalid NSEC chain excluding wildcard

```
mwalet.nl.              86400    IN    NSEC    mwalet.nl. A NS SOA MX TXT RRSIG NSEC DNSKEY
mwalet.nl.              86400    IN    RRSIG   NSEC 8 2 86400 20170316000000 20170223000000 53476 mwalet.nl.
```

- DVE-2017-0010: infracom.nl bogus DNSSEC denial of existence
  - Cycle in the wildcard NSEC record:
    - kdig +dnssec IN TXT foo.fivelpoort.nu. @ns3.infracom.nl.

```
*.fivelpoort.nu.        NSEC  *.fivelpoort.nu. CNAME RRSIG NSEC
```

“DNS must be easy” class of errors

- DVE-2017-0011: GitBook CDN nameservers return A records on any QTYPE
  - and failed to set **AA** bit in the responses (some of it has been fixed)

```
$ kdig +dnssec IN AAAA cdn.gitbook.com. @ns1.gitbook.me.
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 40264
;; Flags: qr rd; QUERY: 1; ANSWER: 1; AUTHORITY: 0; ADDITIONAL: 0

;; QUESTION SECTION:
;; cdn.gitbook.com.           IN    AAAA

;; ANSWER SECTION:
cdn.gitbook.com.        3600  IN    A     95.85.1.232
```

Duplicate RRs

- DVE-2017-0016: Google nameservers return duplicate RRs
  - The duplicate RRs should be suppressed
  - Hard to catch (a very dynamic service)

```
$ kdig +norec @ns1.google.com dns.google.com
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 8531
;; Flags: qr aa; QUERY: 1; ANSWER: 3; AUTHORITY: 0; ADDITIONAL: 0

;; QUESTION SECTION:
;; dns.google.com.            IN    A

;; ANSWER SECTION:
dns.google.com.         300    IN    A     216.58.214.206
dns.google.com.         300    IN    A     216.58.214.206
dns.google.com.         300    IN    A     216.58.214.206

;; Received 80 B
;; Time 2017-05-09 14:33:25 CEST
;; From 216.239.32.10@53(UDP) in 27.8 ms
```

Fresh from the oven...

Run, run, NXDOMAIN is everywhere...

- DVE-2017-0017: Raiffeisen Bank nameservers for rb.cz return NXDOMAIN instead of NOERROR on non-existing RRTYPEs
  - and doesn't bother to add SOA for negative caching (which saves them from total failure)

```
$ kdig +norec IN SOA lb-dns-live-01.rb.cz @lb-dns-live-01.rb.cz.
;; ->>HEADER<<- opcode: QUERY; status: NXDOMAIN; id: 2847
;; Flags: qr aa; QUERY: 1; ANSWER: 0; AUTHORITY: 0; ADDITIONAL: 0

;; QUESTION SECTION:
;; lb-dns-live-01.rb.cz.            IN      SOA
```

- DVE-2017-0018: Raiffeisen Bank nameservers drops DNS packets with EDNS version > 0

```
$ kdig +edns=1 +norec IN TXT www.rb.cz @lb-dns-live-03.rb.cz.
    ;; WARNING: response timeout for 89.233.149.41@53(UDP)
```

Fixed Issues!

Fixed: No Response Issue on unknown RR Types

- DVE-2017-0014: domaincontrol.com nameservers filter TLSA queries

```
$ kdig IN TLSA _25._tcp.svr-zeta.uspta.org @216.69.185.50
;; connection timed out; no servers could be reached

$ kdig IN A _25._tcp.svr-zeta.uspta.org @216.69.185.50
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 18411
;; flags: qr aa; QUERY: 1, ANSWER: 0, AUTHORITY: 8, ADDITIONAL: 1
```

Fixed: EDNS(1) Google Public DNS handling

- DVE-2017-0005: Google DNS returns incorrect response to EDNS(1)
  - Google DNS returns a packet with no Question section
  - The minimal response MUST be the DNS header, question section, and an OPT record.

```
$ dig +edns=1 @8.8.8.8 IN A www.google.com
;; ->>HEADER<<- opcode: QUERY, status: BADVERS, id: 54428
;; flags: qr rd ra; QUERY: 0, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 512
```

  - After the fix

```
$ kdig +edns=1 @8.8.8.8 IN A www.google.com
;; ->>HEADER<<- opcode: QUERY; status: BADVERS; id: 18482
;; Flags: qr rd ra; QUERY: 1; ANSWER: 0; AUTHORITY: 0; ADDITIONAL: 1

;; EDNS PSEUDOSECTION:
;; Version: 0; flags: ; UDP size: 512 B; ext-rcode: BADVERS

;; QUESTION SECTION:
;; www.google.com.            IN    A
```

Fixed: Misc Protocol Errors – AD bit w/o DNSSEC

- DVE-2017-0007: České Radiokomunikace a.s. CDN nameservers sets AD bit on every response
  - (DNSSEC) AD flag doesn't make sense for authoritative servers
  - But the answer for non-auth is still wrong

```
; <<>> DiG 9.10.3-P4-Debian <<>> +norec +noedns IN A example.com. @sr01.cdn.cra.cz.
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 16639
;; flags: qr ad; QUERY: 1, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 0
```

  - After the fix:

```
$ kdig +norec +noedns IN A example.com. @master.dns.rocks
;; ->>HEADER<<- opcode: QUERY; status: REFUSED; id: 64833
;; Flags: qr; QUERY: 1; ANSWER: 0; AUTHORITY: 0; ADDITIONAL: 0
```

Fixed: Invalid unsigned DNSSEC delegation

- DVE-2017-0008: Cloudflare omits SOA from NOERROR/NODATA response to DS query
  - info.nominet.uk. is insecure delegation from secure nominet.uk. zone
  - Fails to include NS + SOA in negative answer

```
$ kdig +dnssec +norec @173.245.58.93 info.nominet.uk IN DS
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 24442
;; Flags: qr aa; QUERY: 1; ANSWER: 0; AUTHORITY: 2; ADDITIONAL: 1
```

  - After the fix:

```
$ kdig +dnssec +norec @curt.ns.cloudflare.com. info.nominet.uk IN DS
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 50665
;; Flags: qr aa; QUERY: 1; ANSWER: 0; AUTHORITY: 4; ADDITIONAL: 1
```

Other related work

- EDNS Compliance Project by Mark Andrews @ ISC
  - https://ednscomp.isc.org/

Violation Submission

- Community Effort
- Anyone can submit
- But it needs a review from somebody else
- Consists of:
  - DVE-<year>-####: short description
  - Description
  - Evidence (textual + dnstap)
  - Proposed fix/Workaround
  - DNS Operator/Vendor Response
  - Metadata
    - Submit-Date
    - Report-Date
    - Fixed-Date
    - ...

What can we do about it?

- DNS Resolver Vendors
  - Be patient, educate
  - Coordinate
  - Make a plan to remove the workarounds
  - Stop resolving the most blatant violations
  - Add (small) time penalty?
- DNS Community
  - Be inclusive
  - Invite the DNS operator people
  - Invite the CDN people
  - Promote use of existing solutions
  - DNS-OARC as neutral platform for this kind of discussion

How can I help?

- The DNS Violations DVE Repository: https://github.com/dns-violations/dns-violations
- Mailing list: https://lists.dns-oarc.net/mailman/listinfo/dns-violations
- A decent website would be nice (we have domain, but no web): https://dns-violations.org generated from DVEs...
- Join the team of the reviewers!
- Report violations!
- And report violations to the DNS operators!
- Invite DNS operators to DNS-OARC
- Write blogposts about correct behaviour of DNS
- ...

Thanks for listening!

Questions?

DNS Violations Project
https://github.com/dns-violations

Ondřej Surý • ondrej.sury@nic.cz • 11. 5. 2017