Compact Denial of Existence in DNSSEC

Shumon Huque & Christian Elmerot

March 30th 2023
DNS Operations Working Group
Internet Engineering Task Force (IETF) 116 Meeting
Yokohama, Japan

Compact DNSSEC denial of existence

- Or “Compact Answers” (formerly “Black Lies”)
- Originally described in an expired Internet draft ~ March 2016:
  - https://datatracker.ietf.org/doc/html/draft-valsorda-dnsop-black-lies
- Never proposed for RFC publication in any category.
- But is widely deployed amongst commercial online signers:
  - Cloudflare, NS1, and Amazon Route53
- Eliminates NXDOMAIN responses.
- This has some operational implications.

NXDOMAIN considered unnecessary?

- For names that don’t exist, it pretends that they do actually exist, but don’t have any data associated with the queried type.
- i.e. they return NODATA answers (NOERROR response code, an empty ANSWER section)
- Rationale:
    - **More compact answers**. A signed NODATA response requires just 1 NSEC record (and corresponding signature).
    - **Higher performance**: only 1 online cryptographic signing operation is needed.
    - By contrast, an NXDOMAIN response requires up to 2 NSEC or up to 3 NSEC3 records, and their corresponding signatures.
    - Authoritative servers don’t need to know the detailed **structure of the zone**.

Operational Implications?

- For typical end users, probably nothing; a NODATA response is treated mostly identically to NXDOMAIN.
  - However NODATA may result in additional follow-on queries, which NXDOMAIN would have suppressed (e.g. other types at same name)
- But a variety of diagnostic, troubleshooting, traffic characterization, & provisioning tools may need adaptations to correctly deal with this protocol.
  - Especially tools that rely on the correctness of the DNS Response Code field.
  - Arguably, the RCODE should not be relied on, because it is unauthenticated.
  - But then we must infer non-existence of a name from signed data in the response (namely, NSEC records)
  - Can this inference be reliably drawn with Compact Answers?

Distinguish NXDOMAIN from ENT

- Empty Non-Terminals (ENT) are names that have no resource record type associated with them, but have descendant names that do.
- In the described Compact Answers spec, they are indistinguishable from non-existent names, because they have the same type bitmap (“NSEC RRSIG”) in the NSEC record.
- Other hacks (not in the written spec):
    - Some implementations work around this, by returning all supported RRtypes in the type bitmap *except* for the qtype in responses to ENTs.
    - Can’t distinguish ENT any more!
    - May bloat the size of the type bitmap
    - May confuse tools that perform type inference from the bitmap

Distinguish ENT from NXDOMAIN (cont)

- Either we have to insert a distinguisher in the response for NXDOMAIN or for ENT (could do both, but unnecessary)
- A “pseudo” RR type in the NSEC Type Bitmaps field.
- In “Empty Non-Terminal Sentinel for Black Lies”, an ENT distinguisher was proposed:
  - https://www.ietf.org/archive/id/draft-huque-dnsop-blacklies-ent-01.html
  - This is deployed in the field today by NS1, using private use RR type# 65281
  - Allowed the same NSEC type bitmap pattern to identify NXDOMAIN in other implementations that used the ‘all types except qtype’ bitmap hack for ENTs.
  - Also, less work for the authority server since ENTs are far less common than NXDOMAIN (not a practical concern - work required here is extremely minimal).

Empty Non-Terminal Response - enhanced

```
$ dig +nostats +dnssec ent1.sfdcsd.net. AAAA

;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 3727
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 0, AUTHORITY: 4, ADDITIONAL: 1

;; QUESTION SECTION:
;ent1.sfdcsd.net.         IN    AAAA

;; AUTHORITY SECTION:
sfdcsd.net.     1799    IN    SOA    dns1.p08.nsone.net. hostmaster.nsone.net.
1619363158 43200 7200 1209600 3600
sfdcsd.net.     1799    IN    RRSIG    SOA 13 2 3600 20210728150036 20210726150036 44688
sfdcsd.net. xSv1lHZIPbKJ5f8pJf0Es0vSg+mr0SFk37Nh1OabvD96UdncINFGxYWG
vDNDcK7jXqRw8cwOK5jjCI8PWsx50w==
ent1.sfdcsd.net.    3599    IN    NSEC    \000.ent1.sfdcsd.net. RRSIG NSEC TYPE65281
ent1.sfdcsd.net.    3599    IN    RRSIG    NSEC 13 3 3600 20210728150036 20210726150036
44688 sfdcsd.net. UElkkdTBMg00mu6v0HFMkEc89IjNNbMg6C4zsBv2RaFsHJFI455oHhaA
3L0rxhuiKW0//pXWHjOx9iwVaIeTcA==
```

New proposal: NXDOMAIN distinguisher

- Implement the “pseudo” type distinguisher for NXDOMAIN (instead of for ENT).
- Mnemonic “**NXNAME**” (for “Non Existent Name”)
- Use the lowest available RR type number from the “meta” range (128?) to minimize the type bitmap size
- *Cloudflare has recently implemented this with private RR type code 65283 (while we wait for a formal IANA allocation).*

NXDOMAIN Response (new)

```
$ dig +dnssec nxd.example.com. AAAA

;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 3913
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 0, AUTHORITY: 4, ADDITIONAL: 1

;; QUESTION SECTION:
;nxd.example.com.         IN    AAAA

;; AUTHORITY SECTION:
example.com.         1800    IN    SOA    dns1.example.com. hostmaster.example.com.
1619363158 43200 7200 1209600 3600
example.com.         1800    IN    RRSIG    SOA 13 2 3600 20210728145830 20210726145830
5986 sfdcsd.net. 8yFF++j9XBPARG+4jcZ/w0IvkVgPeS0eU5n3jS7d6RSFPQcO2k+9oU5V
3H3aev8Qcj0+7m5ht1Z4oaXkZLFclA==
nxd.example.com.    3600    IN    NSEC    \000.nxd.example.com. RRSIG NSEC NXNAME
nxd.example.com.    3600    IN    RRSIG    NSEC 13 4 3600 20230228145830 20230226145830
5986 example.com. TK5ccSxJ8Dt5oHmLi/6cykmglsjT2dMwZAnlbCfdsdN8DxXpu4wULBy9
k/ws0sECMh7AQcs54VJAR1W/XZCFwA==
```

NXDOMAIN restoration in RCODE?

- Lots of security, diagnostic, and traffic characterization tools only examine the RCODE field today.
- In theory they could be enhanced to recognize NXNAME in the NSEC type bitmap, but will they ever be?
- For non-validating queriers, an NXNAME recognizing resolver could re-instate NXDOMAIN in the RCODE field of responses it sends back.
- What about validating (DO=1) queriers though?
- And is there a way to put NXDOMAIN RCODE in the response from the authority server itself? (probably only via EDNS signaling)

New draft

- https://datatracker.ietf.org/doc/html/draft-huque-dnsop-compact-lies-01
- Describes the current Compact Answers scheme.
- The new NXDOMAIN distinguisher RR type.
- Some discussion of RCODE substitution schemes.

Can we adopt this draft?

- Compact Answers is widely deployed today.
- And yet has no formally published specification. We need to fix this.
- And make it work better (by definitively restoring the NXDOMAIN signal).

- Mailing list threads (dnsop@ietf.org, March 2023)
  - https://mailarchive.ietf.org/arch/msg/dnsop/oic5BUS9Ae2vrMM5ltEe1MDe-l0/
  - https://mailarchive.ietf.org/arch/msg/dnsop/86xo4YZfThXm43CxZ9Tk-woPILs/