

IETF 121 hackathon: greasing DNS answers

Stéphane Bortzmeyer

<stephane+blog@bortzmeyer.org>

First publication of this article on 10 November 2024

<https://www.bortzmeyer.org/hackathon-ietf-121.html>

On November 2 and 3 was the IETF hackathon in Dublin. I worked on the greasing of DNS answers from an authoritative name server. What is greasing? Continue reading.

One of the big technical problems of the Internet is its ossification : software is written by people who did not read the technical standards, or did not understand them, specially software in the middleboxes (load balancers, firewalls, etc). As a result, some things that are possible according to the technical specification are de facto forbidden by broken software. This makes difficult to deploy new things. For instance, TLS 1.3 had to pretend to be 1.2 (and add an extension to say "I am actually 1.3") because too many middleboxes prevented the establishment of TLS sessions if the version was 1.3 (see RFC 8446¹, section 4.1.2). This problem is widespread in the Internet, specially since there is typically no way to talk to the middlebox software authors and these boxes are popular among managers.

A way to fight ossification is **greasing**. Basically, the idea is to exercise all the features and options of a protocol from day one, not waiting that you really need them. This way, broken software will be detected immediately, not many years after, when it is entrenched. TLS was the first protocol to go that way (see RFC 8701) and it proved effective. QUIC also uses greasing (RFC 9287).

The DNS could benefit from greasing as well, since it is often difficult to deploy new features, because they sometimes break bad software (it was the case with the cookies of RFC 7873). Hence the current Internet Draft `draft-ietf-dnsop-grease`.

OK, so, let's grease the rusted parts of the Internet but where exactly, and how? DNS servers are basically of two kinds : resolver and authoritative servers. The version -00 of the draft only mentions resolvers because they are in the best place to test greasing and to report what broke. The general idea is that the resolver sends its queries with "unexpected" values (unallocated EDNS options, unallocated EDNS flags, etc, all of them "legal" according to the RFC). If it receives no reply from the authoritative

1. Pour voir le RFC de numéro NNN, <https://www.ietf.org/rfc/rfcNNN.txt>, par exemple <https://www.ietf.org/rfc/rfc8446.txt>

server (or a bad one such as FORMERR), and, if retrying without greasing work, the resolver knows there is a problem in the path to this authoritative name server and can log it and/or report it (for instance through RFC 9567). The remaining question is : what we can grease? We need options that are legal to send but new and unexpected. For plain DNS, there is no hope : there is only one remaining (unallocated) bit <https://www.iana.org/assignments/dns-parameters/dns-parameters.xml#dns-parameters-12> in the flags (RFC 1035, section 4.1.1) of the DNS query. So, it means we can grease only with EDNS stuff : EDNS version number <https://www.iana.org/assignments/dns-parameters/dns-parameters.xml#dns-parameters-14>, EDNS options <https://www.iana.org/assignments/dns-parameters/dns-parameters.xml#dns-parameters-11>, and EDNS flags <https://www.iana.org/assignments/dns-parameters/dns-parameters.xml#dns-parameters-13>. For instance, unknown EDNS options are supposed to be ignored (otherwise, it would never be possible to deploy new options). Now, how to choose the unallocated values to send? TLS decided to reserve ranges of values for which to choose randomly. The risk is that some bad software will treat this range in a special way but, at least, it guarantees there will be no collision with a future allocation.

This is the current version (-00) of the draft. Now, the work at the hackathon. First, I decided to work on an authoritative server. A priori, it is less useful than a resolver, because, unlike the resolver, the authoritative name server cannot know if its reply was accepted or not, or created problems. But it could be useful on test zones, to see (for instance through the use of RIPE Atlas probes <https://atlas.ripe.net>) if they have resolution issues. The work was done on the software Drink <https://framagit.org/bortzmeyer/drink/>.

First test, sending back in the reply two EDNS records. Sending two OPT records in a response does not seem forbidden by the RFC (which prohibit it only in a query, RFC 6891, section 6.1.1) but dig does not like it :

```
% dig +norec grease.courbu.re SOA @31.133.134.59
;; Warning: Message parser reports malformed message packet.
```

It creates problems with many other programs and it is not clear if it is legitimate so let's stop here.

Second test, sending an EDNS reply with a version number which is higher than the one requested. This is legal, the last paragraph of Section 6.1.3 of RFC 6891 says that a responder can respond with a higher EDNS version than what was requested by the requestor. (And it explains why, and the limits, for instance to keep the same format.) I tried that for DNS greasing and typical resolvers seem to be happy with it. But DNS testing tools (very useful tools, do not forget to test your zones with them!) disagree. ednscmp <https://ednscmp.isc.org/> says "expect : OPT record with version set to 0" (not greater-or-equal, strictly equal). DNSviz <https://dnsviz.net/> says "The server responded with EDNS version 1 when a request with EDNS version 0 was sent, instead of responding with RCODE BADVERS. See RFC 6891, Sec. 6.1.3." (We obviously do not read this section in the same way. To me, it mentions BADVERS only in a different context.) And Zonemaster <https://zonemaster.net/> also disagrees with me. So, there is a debate : when a responder knows both version 0 and some higher version (say, version 1), can it reply to a EDNS=0 query with a EDNS=1 response? Can we use that for greasing?

Less controversial, adding EDNS options and flags. You can see the result here :

<https://www.bortzmeyer.org/hackathon-ietf-121.html>

