WikiLeaks DNS mirrors and the limits of the DNS

Stéphane Bortzmeyer
<stephane+blog@bortzmeyer.org>

First publication of this article on 12 December 2010. Last update on of 13 December 2010

https://www.bortzmeyer.org/size-dns-wikileaks.html

Following the big crackdown against WikiLeaks and the many attempts to censor the independent information Web site, many people, eager to do something for the freedom of information, have set up mirrors <http://www.wikileaks.ch/mirrors.html> of the WikiLeaks content. Some of these mirrors do not actually store content, they are just DNS mirrors, which store the IP addresses of the actual Web sites. The goal is to allow these Web sites to be found even if many domain names are deleted, following brutal takedowns like the one performed by the US customs <http://www.justice.gov/iso/opa/ag/speeches/2010/ag-speech-101129.html> in another case. But these DNS mirrors must be careful to test that they work even if they go beyond the traditional limits of the DNS.

The biggest of these limits is the size : at its beginning, DNS accepted only answers of up to 512 bytes (RFC 1035 [1], section 2.3.4). This limit was lifted ten years ago, in RFC 2671. But ten years is short in an environment as ossified as the Internet and many DNS resolvers still cannot handle properly larger answers.

Let's take as an example the DNS mirror all-wikileaks.bortzmeyer.fr. It now stores 157 IP addresses (it may change in the future since mirrors appear and disappear all the time and you have to test them often <https://www.bortzmeyer.org/testing-wikileaks.html>). It allows a big resilience since only one working IP address in the set is sufficient to reach WikiLeaks. The total size of the DNS answer is 2753 bytes (a bit less if you query only one IP address family, for instance only IPv4). It is much larger than the traditional limit and it is even larger than the Ethernet MTU, the most common maximum packet size today. Does it work ?

It depends on the resolver. With a standard, out-of-the-box DNS resolver like BIND or Unbound, it works fine, the name can be resolved and a Web browser can visit the site. That's because these programs correctly implement the DNS as it is today. They use the EDNS0 option of RFC 2671 and have a default buffer size of 4096 bytes, which is sufficient today (the largest theoretical size now is 65536 bytes).

If you want to test by hand, with the common DNS testing tool dig, be careful that, unlike a proper resolver, it does not use EDNS0 by default. You have to indicate it on the command line :

1. Pour voir le RFC de numéro NNN, https://www.ietf.org/rfc/rfcNNN.txt, par exemple https://www.ietf.org/rfc/rfc1035.txt

```
% dig +bufsize=4096 ANY all-wikileaks.bortzmeyer.fr
...
;; Query time: 4 msec
;; SERVER: ::1#53(::1)
;; WHEN: Fri Dec 10 17:16:03 2010
;; MSG SIZE  rcvd: 2753
```

or to put it once for all in the `/.digrc` configuration file :

```
% cat ~/.digrc
+bufsize=4096

% dig ANY all-wikileaks.bortzmeyer.fr
...
;; Query time: 4 msec
;; SERVER: ::1#53(::1)
;; WHEN: Fri Dec 10 17:17:14 2010
;; MSG SIZE  rcvd: 2753
```

But many ISP or local networks provide broken DNS resolvers to their users. Even if the DNS resolver is OK, it is at the mercy of a firewall or another middlebox which mangles either the DNS query (by deleting the EDNS0 option) or the answer (by being unable of transmitting the larger-than-512-bytes reply). (For details on middleboxes, see the SSAC 035 document <https://www.icann.org/en/committees/security/sac035.pdf> or RFC 5625.)

Possible issues are many : for instance, if the query does not use EDNS0, the answer, too large, won't be entirely sent and the authoritative name server will set the TC bit which means "Data truncated"). The resolver is then supposed to retry over TCP and some cannot, either because they don't support TCP (a big mistake in today's Internet) or because they are blocked by a misconfigured firewall which allows DNS only on UDP.

Here is an example on the Free network (Free being the second largest ISP in France). While EDNS0 queries seem to go just fine, the resolvers do not accept TCP. When the dig testing tool queries them, it receives a truncated response, retries over TCP and is blocked :

```
% dig  A all-wikileaks.bortzmeyer.fr
;; Truncated, retrying in TCP mode.
;; Connection to 212.27.40.241#53(212.27.40.241) for \
           all-wikileaks.bortzmeyer.fr failed: connection refused.
```

Among the programs which may have problems, the PowerDNS recursor, which does not have EDNS0 by default (for the reasons explained by its author <http://mailman.powerdns.com/pipermail/pdns-users/2010-December/007247.html>). As a result, PowerDNS Recursor always falls back to TCP, which may be a problem with some broken firewalls.

Another interesting test (thanks to Marco Davids) was done through a AVM Fritz!Box 7170 Annex A with firmware version 58.04.82. Like many cheap boxes, it has an internal DNS proxy which does not do TCP and does not accept answers ¿512 bytes. TCP connections are here rejected :

```
% dig +bufsize=4096 ANY all-wikileaks.bortzmeyer.fr @192.168.68.7
;; Truncated, retrying in TCP mode.
;; Connection to 192.168.68.7#53(192.168.68.7) for
all-wikileaks.bortzmeyer.fr failed: connection refused.
```

With EDNS0 on the dig side (option `+bufsize`), and the interdiction to fall back to TCP (option `+ignore`):

```
% dig +ignore +bufsize=4096 ANY all-wikileaks.bortzmeyer.fr @192.168.68.7
...
;; flags: qr tc rd ra; QUERY: 1, ANSWER: 23, AUTHORITY: 0, ADDITIONAL: 0
...
;; Query time: 4 msec
;; SERVER: 192.168.68.7#53(192.168.68.7)
;; WHEN: Mon Dec 13 09:35:10 2010
;; MSG SIZE  rcvd: 509
```

The box has stuffed as many answers it could, to stay below 512 bytes.

If you want to test yourself, I'll be happy to receive (at `bortzmeyer+testdnssize@bortzmeyer.org`) reports of issues but please, do not forget to indicate the resolver used (make and model), as well as existing middleboxes such as firewalls) or the network used (name of the ISP and city) if you do not control the DNS resolver yourself. Complete output of dig would be a big plus. Another big DNS mirror, to have several tests, is `wklk.eu.org`.

Thanks to Niall O'Reilly for his style and technical checking. All the opinions are mine.