

The Sichuan pepper "attack": from DNS censorship to DNS redirection

Stéphane Bortzmeyer
<stephane+blog@bortzmeyer.org>

First publication of this article on 5 February 2015

<https://www.bortzmeyer.org/sichuan-pepper.html>

It is well-known, for many years, that the Chinese government censors the Internet via several means, including DNS lies. For a long time, these DNS lies have been generated by the network itself : when a DNS query for a censored name is seen, an active censorship device generates a lie and sends a reply with a wrong IP address. A few weeks ago, there have been a change in this system : the IP addresses returned by the Great FireWall are more often actual addresses used by real machines, which suddently have to sustain a big traffic from China.

This technique have been studied and documented in several papers such as "The Great DNS Wall of China <<http://cs.nyu.edu/~pcw216/work/nds/final.pdf>>" or "Source code to identify the fake DNS packets from China <<https://lists.dns-oarc.net/pipermail/dns-operations/2010-March/005340.html>>". Basically, every DNS request in a chinese network, whatever the destination IP address, is examined and, if the qname (Query Name) in it matches a predefined list of censored domains, a fake reply is generated and sent to the requester. The bad consequences of this technique outside of China have been described in articles like "Accidentally Importing Censorship <<http://research.dyn.com/2010/03/fouling-the-global-nest/>>" or "China censorship leaks outside Great Firewall via root server <<http://arstechnica.com/tech-policy/2010/03/china-censorship-leaks-outside-great-firewall-via-root-server/>>" : since the censorship system acts whatever the destination IP address is, **if one of your DNS packets happen to goes through China**, you will be subjected to chinese censorship <<https://labs.ripe.net/Members/pk/denic-case-study-using-ripe-atlas>>.

To see this DNS tampering, one just has to query any IP address in China (it does not need to be an existing machine, since the fake DNS reply is made by the network, not by an evil DNS server, the address here was choosen at random and tested to be sure it does not reply to any other packet) :

```
% dig @218.76.74.42 A www.ssi.gouv.fr
; <<>> DiG 9.9.5-8-Debian <<>> @218.76.74.42 A www.ssi.gouv.fr
; (1 server found)
;; global options: +cmd
;; connection timed out; no servers could be reached
```

As expected, this non-existing machine does not reply. But if we try with a censored name, here Facebook:

```
% dig @218.76.74.42 A www.facebook.com

; <<>> DiG 9.9.5-8-Debian <<>> @218.76.74.42 A www.facebook.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 30344
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;www.facebook.com.      IN A

;; ANSWER SECTION:
www.facebook.com.      2655 IN A 67.205.10.141

;; Query time: 313 msec
;; SERVER: 218.76.74.42#53 (218.76.74.42)
;; WHEN: Wed Feb 04 12:06:43 CET 2015
;; MSG SIZE rcvd: 50
```

This time, we get an IP address, and completely unrelated to Facebook (it is a USAn hoster). It is not just a match of the string that triggers the lying answer, the system actually understands DNS:

```
% dig @218.76.74.42 A www.facebook.com.example.net

; <<>> DiG 9.9.5-8-Debian <<>> @218.76.74.42 A www.facebook.com.example.net
; (1 server found)
;; global options: +cmd
;; connection timed out; no servers could be reached
```

Besides going to China and testing from your laptop, there are other ways to see this DNS tampering: one is to use the RIPE Atlas probes <<https://atlas.ripe.net/>>. Few of them are in China and many seem immune to the DNS tampering, probably because they are located on a network with a safe VPN connection to the outside.

```
% python resolve-name.py --country=CN --requested=30 www.facebook.com
Measurement #1854647 for www.facebook.com/A uses 15 probes
[66.220.158.19] : 4 occurrences
[179.60.192.3] : 2 occurrences
[31.13.79.246] : 3 occurrences
[31.13.68.84] : 3 occurrences
[173.252.74.22] : 1 occurrences
[153.122.20.47] : 2 occurrences
[31.13.68.70] : 3 occurrences
[67.205.10.141] : 1 occurrences
[173.252.73.52] : 1 occurrences
[114.200.196.34] : 1 occurrences
[31.13.76.102] : 1 occurrences
Test done at 2015-02-04T11:11:19Z
```

(The program `resolve-name.py` is on Github <<https://github.com/RIPE-Atlas-Community/ripe-atlas-community-contrib>>.) Most answers are actually Facebook's but not all (for instance, 114.200.196.34 is a Korean access provider).

And a last solution is to use the cloud, actually a Web site hosted in China which allows you to perform DNS requests, for instance, <<http://viewdns.info/chinesefirewall>>.

Until the beginning of 2015, the returned IP addresses were apparently non-reachable addresses, unallocated, or even "class D" ("*multicast*", 224.0.0.0/4) or "class E" (unused, 240.0.0.0/4) addresses. When the unsuspected chinese user tried to reach Facebook, he got one of these unroutable addresses and it ended in a timeout. But a change was done at the beginning of 2015. Now, the returned IP addresses are, much more often, actually assigned to a real machine. Suddenly, several system administrators reported a lot of traffic coming from China, asking for sites like Facebook, something that never happened before.

The first report, seen from the chinese site (chinese users sent to unexpected Web sites) was "Visitors to blocked sites redirected to porn <<https://en.greatfire.org/blog/2015/jan/gfw-upgrade-fail-visitors>>". Many other reports documented the other side, the point of view of the site suddenly receiving chinese traffic. See "Fear China <<http://furbo.org/2015/01/22/fear-china/>>", "DDOS on La Quadrature du Net, analysis <<https://benjamin.sonntag.fr/DDOS-on-La-Quadrature-du-Net-analysis>>" or "DDoS from China [Caractère Unicode non montré ¹] Facebook, WordPress and Twitter Users Receiving Sucuri Error Pages <<http://blog.sucuri.net/2015/01/ddos-from-china-facebook-wordpress-and-t.html>>".

Let's look at this traffic, as seen by one of the Web servers of CGT. The HTTP server log contains :

```
x.y.z.t - - [27/Jan/2015:07:48:29 +0100] "GET /plugins/like.php?href=https://www.facebook.com/yvesrocher.nederla
```

(The source IP address was from "Graduate University of Chinese Academy of Sciences <<http://english.gu.cas.cn/>>") The request is for `/plugins/like.php`, Facebook's "Like" button. It does not exist on the server and the HTTP status code is therefore 404 (not found). What is interesting is the Referer : HTTP field, here <https://secureorder.yves-rocher.nl/control/main/>. It shows that the chinese client did not want explicitly to visit Facebook (he probably knows that this would be hopeless from China) but he visited a site (<<https://secureorder.yves-rocher.nl/control/main/>>) which includes a Facebook "Like" button, therefore triggering a HTTP request to Facebook and, because of the DNS tampering, actually going to the CGT server. Here is a part of the HTML source code of the Referer :

```
<iframe src="https://www.facebook.com/plugins/like.php?href=https://www.facebook.com/yvesrocher.nederland&width=
```

Unfortunately, the HTTP server, like most HTTP servers, did not log the Host : field in the HTTP request. This field indicates which host was requested by the client. Here, we can guess it was `www.facebook.com`, from the requested path (`/plugins/like.php`). But it would be better if all HTTP servers were to log the Host : field (in Apache, it is the `%v` format directive). On another HTTP server, which was victim of the same Sichuan pepper issue, and had this logging activated, we see :

1. Car trop difficile à faire afficher par \LaTeX

```
x.y.z.t - - [21/Jan/2015:00:53:33 +0100] "GET /plugins/like.php?href= [...] "Mozilla/5.0 (Linux; U; Android
```

The `www.facebook.com` clearly indicates that the original user really wanted to go to Facebook, and was distracted by the DNS tampering.

Another very common HTTP request is :

```
x.y.z.t - - [27/Jan/2015:07:48:41 +0100] "GET /announce?info_hash=%0eo%40%e7.u%f7%a3%3e%e6%e9%a9%5e%e45%8bK
```

(The original source IP address was "China Mobile Communications Corporation".) The requested path, `/announce?info_hash` is typical of BitTorrent clients going to a BitTorrent tracker. QQdownload is a popular BitTorrent client in China. On another machine, where Host : logging is activated, we see that the requested host was indeed a tracker, `tracker.thepiratebay.org`, also censored in China.

OK, so we see a lot of HTTP traffic, coming almost only from chinese IP addresses, and we see that the requested names are indeed censored in China. Can we prove that the Great FireWall redirected to the IP addresses of the victims? We can do it with `<http://PassiveDNS.cn/>`, a passive DNS (a system which records DNS answers observed on the network) database located in China. First, we can check (like we did with `dig` to a chinese IP address) that names like `tracker.thepiratebay.org` are indeed tampered with, using the API client of PassiveDNS.cn, `flint` :

```
% flint rrset tracker.thepiratebay.org A
[api error]: http://www.passivedns.cn/api/rrset/keyword/tracker.thepiratebay.org/rtype/1/
```

:-(This request works for non-censored names. I suspect that censored names, being redirected to many IP addresses, exceed some limit of PassiveDNS.cn, leading to this bug. But the Web interface of PassiveDNS.cn still works so we can see indeed that we have many IP addresses for `tracker.thepiratebay.org`. It is not a trick specific to The Pirate Bay, all the other censored names show the same behaviour. But what is more interesting is how many names point to the IP address of the victim, `212.234.228.143`?

```
% flint rdata 212.234.228.143 A | more
212.234.228.143 A In rdata
-----
50congres.cgt.fr          212.234.228.143 2014-11-26 10:10:13
accounts.youtube.com     212.234.228.143 2015-01-28 10:26:25
adecco.cgt.fr            212.234.228.143 2014-11-29 12:56:14
adm-salaries.cgt.fr     212.234.228.143 2015-01-04 06:39:26
aful.cgt.fr              212.234.228.143 2015-02-04 06:04:32
platform.twitter.com     212.234.228.143 2015-01-15 23:20:37
plus.google.com          212.234.228.143 2015-02-03 03:23:34
tracker.thepiratebay.org 212.234.228.143 2015-01-31 23:58:02
www.bloomberg.com        212.234.228.143 2015-02-01 10:00:13
www.facebook.com         212.234.228.143 2015-02-01 21:37:59
www.kanzhongguo.com      212.234.228.143 2015-01-15 23:18:35
...
```

<https://www.bortzmeyer.org/sichuan-pepper.html>

So, we can see that the original interpretation is correct : the Great Firewall, through DNS tampering, redirects many very popular names to innocent servers.

How many sites are used as "sinkholes" by the chinese censorship system? Let's query repeatedly `www.facebook.com` to an IP address in China (123.123.123.123) during 17 hours (on February 2nd). We retrieve 5559 answers. This is 1856 distinct IP addresses because some addresses are sent several times. So, it does not look random. Here are the most popular addresses (the owner name has been retrieved through the very useful `cymruwhois` package <<https://pypi.python.org/pypi/cymruwhois>>):

```
205.186.162.167 (MEDIATEMPLE - Media Temple, Inc.,US): 26
77.66.57.6 (NGDC NetGroup A/S,DK): 24
205.157.169.156 (ASN-PENNWELL - PennWell corporation,US): 24
216.201.83.226 (NATIONALNET-1 - NationalNet, Inc.,US): 24
64.20.49.2 (NJIIX-AS-1 - NEW JERSEY INTERNATIONAL INTERNET EXCHANGE LLC,US): 24
193.188.112.80 (AS6453 - TATA COMMUNICATIONS (AMERICA) INC,US): 23
114.130.54.22 (BCC-MANGOCIENT-AS-AP Bangladesh Computer Council,BD): 22
216.57.200.175 (WHIDBEY1 - Whidbey Internet Services,US): 22
74.121.192.250 (BLACKMESH-RST - BlackMesh Inc.,US): 21
137.117.70.70 (MICROSOFT-CORP-MSN-AS-BLOCK - Microsoft Corporation,US): 21
70.32.110.223 (MEDIATEMPLE - Media Temple, Inc.,US): 21
184.173.133.194 (SOFTLAYER - SoftLayer Technologies Inc.,US): 21
5.9.26.245 (HETZNER-AS Hetzner Online AG,DE): 20
195.205.239.197 (TPNET Orange Polska Spolka Akcyjna,PL): 20
222.230.141.241 (VECTANT VECTANT Ltd.,JP): 20
91.213.100.50 (BRACK-AS Brack.ch AG,CH): 20
14.139.212.165 (NKN-CORE-NW NKN Core Network,IN): 20
54.235.118.83 (AMAZON-AES - Amazon.com, Inc.,US): 20
200.57.151.168 (Triara.com, S.A. de C.V.,MX): 20
62.109.134.70 (IGNUM-AS Igunum s.r.o.,CZ): 20
...
```

As you can see, many IP addresses are used in the Great FireWall lies.

Now, let's indulge in some speculation. How are the IP addresses of the victim chosen? At random, and the Great FireWall administrators do not care of the consequences? On purpose, to turn every chinese Internet user into an involuntary accomplice of the dDoS? We must admit that we don't know.

This sort of "attack by referral" is a scourge of the Internet, because there is a very little to do against it. A famous example a few years ago, not involving the DNS, was D-Link NTP "attack" <<http://slashdot.org/story/06/04/07/130209/d-link-firmware-abuses-open-ntp-servers>>.

Thanks to Benjamin Sonntag and Éric Duval for the data.