

Accidentally Importing Censorship

The I-root instance in China

NANOG49: San Francisco, California 2010

Doug Madory
Alin Popescu
Earl Zmijewski
Renesisys Corporation

Chinese Censorship



- The Great Firewall (GFW) is reported to ...
 - Block access to certain IPs and entire prefixes
 - Intercept and return incorrect DNS responses
 - Intercept and reset TCP connections
- DNS queries routed through the GFW can ...
 - Return bogus answers
 - Impact users outside of China

*Note: GFW is a term of convenience for the strange non-point-source effects observed; no evidence of responsibility, state or otherwise. **"It's complicated."***

Try the Chinese firewall for yourself ...

- Repeatedly ...

dig @dns1.chinatelecom.com.cn. www.facebook.com. A

- Answers vary ...

www.facebook.com.	11556	IN	A	37.61.54.158
www.facebook.com.	24055	IN	A	78.16.49.15
www.facebook.com.	38730	IN	A	203.98.7.65

- Results are all over the place.
 - 37/8 is currently unallocated by IANA
 - 78.16/14 is announced by AS 2110 (BT Ireland)
 - 203.98/18 is announced by AS 4768 (TelstraClear, NZ)
- Note: Queries are to “China Telecom” (but may not ever get there).

I-root: Just the facts



- IP address: 192.36.148.17
- Prefixes: 192.36.148.0/23 & 192.36.148.0/24
- Origin: AS 29216 (Dedicated to I-root)
- Single Upstream: AS 8674 (Netnod)
 - AS 8674 has ~80 BGP adjacencies
 - I-root is run by Autonomica
 - Subsidiary of Sweden's Netnod
 - I-root is *anycast* from around the world
 - 14 instances in EMEA
 - 14 in Asia Pacific
 - 6 in North America

DNS-Operations Report (24 March 2010)

Hi there! A local ISP has told us that there's some strange behavior with at least one node in i.root-servers.net (traceroute shows mostly China) It seems that when you ask A records for facebook, youtube or twitter, you get an IP and not the referral for .com

It doesn't happen every time, but we have confirmed this on 4 different connectivity places (3 in Chile, one in California)

This problem has been reported to Autonomica/Netnod but I don't know if anyone else is seeing this issue.

This is an example of what are we seeing:

```
$ dig @i.root-servers.net www.facebook.com A ;
```

```
....
```

```
ANSWER SECTION: www.facebook.com. 86400 IN A 8.7.198.45
```

Mauricio Vergara Ereche
Santiago CHILE

Explanation

- In March, we saw the AS path for 192.36.148.0/24 traverse a Chinese AS before arriving at the I-root:
[...] 10026 **7497 7497 24151** 8674 29216
- By crossing Chinese infrastructure before arriving at I-root, DNS queries were subject to tampering from GFW.
- *29216 is the I-root ASN*
8674 is the Netnod ASN
24151 is the China Internet Network Information Center
7497 is the Chinese Academy of Science

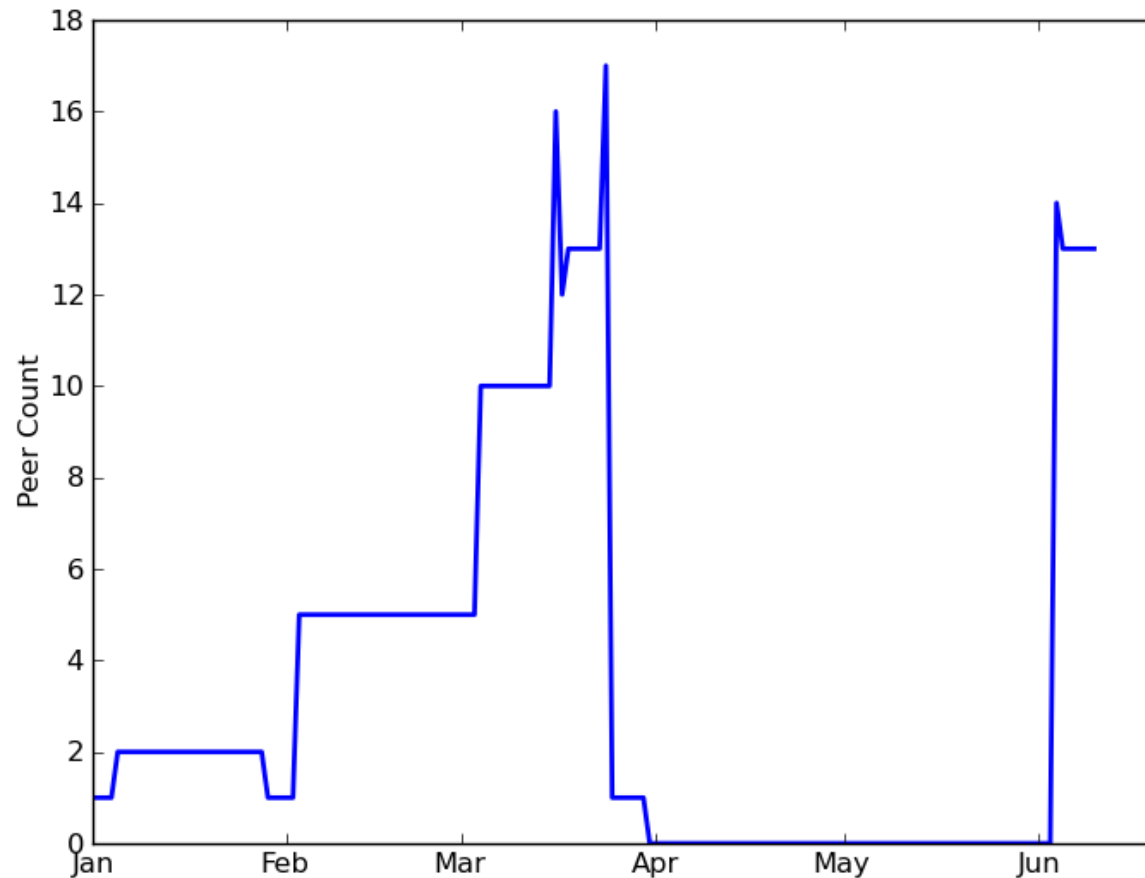
But for the “potential” to be realized, you need an unlikely series of events ...

- Query `www.facebook.com` (or other blocked domain)
- Not cached (locally or by your server)
- And `.com` is not cached either (has a 48 hr TTL)
- Ask the I-root (rather than A, B, C, ... roots)
- Get directed to China's I-root instance
- Game over!
 - Query should return the `.com` servers
 - Instead returns incorrect A record for Facebook
 - Your DNS cache is now poisoned

Timeline



- Peer count for 192.36.148.0/24 (Since Jan 1, 2010)



- Nothing for 192.36.148.0/23

Timeline



- January – March: I-root visible outside of China.
- March 24: Bogus DNS results from I-root first reported. (Here is the accidental importation of censorship.)
- March 25: Netnod withdraws routes.
- June 3: Netnod routes are leaked again via PacNet and PCCW – a larger footprint of potential impact.
- June 14: The leak continues. Answers seem legitimate for now.

Chinese Client – Bad Result (10 June 2010)

```
dig @i.root-servers.net. www.facebook.com. A
; <<>> DiG X.X.X <<>> @i.root-servers.net. www.facebook.com. A
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 26148
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

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

;; ANSWER SECTION:
www.facebook.com.          300    IN      A      59.24.3.173      ← Korea Telecom

;; Query time: 4 msec
;; SERVER: 192.36.148.17#53(i.root-servers.net.)
;; WHEN: Thu Jun 10 18:41:40 2010
;; MSG SIZE rcvd: 50
```

Chinese Client Packet Capture – Bad Result

18:06:17.581240 IP SRC-IP.57520 > 192.36.148.17.53: 54947+ A?
www.facebook.com. (34)

18:06:17.585669 IP 192.36.148.17.53 > SRC-IP.57520: 54947 1/0/0 A
59.24.3.173 (50) ← Bad answer

18:06:17.600736 IP 192.36.148.17.53 > SRC-IP.57520: 54947* 1/0/0 A
243.185.187.39 (66) ← Another bad answer (for good measure)

18:06:17.600778 IP SRC-IP.128 > 192.36.148.17: icmp 102: SRC-IP.128
udp port 57520 unreachable ← 2nd bad answer is not accepted, so 1st was

- This is completely expected behavior.
 - The GFW is known to tamper with DNS packets.
 - The client is inside of China.
 - You can see the exact same behavior querying *any other* root name server from inside China.

US client with PCCW transit (12 June 2010)

```
# traceroute i.root-servers.net
```

```
traceroute to i.root-servers.net (192.36.148.17), 30 hops max, 40 byte packets
```

```
1 sc-smv1494.servint.net (206.214.212.60) 0.044 ms 0.025 ms 0.017 ms
```

```
2 ge9-18.br01.lax05.pccwbtn.net (63.218.42.201) 0.617 ms 0.758 ms 0.777 ms
```

```
3 cni.ge9-1.br02.hkg04.pccwbt.net (63.218.2.146) 154.789 ms 154.795 ms 154.865ms
```

```
4 8.198 (159.226.254.253) 242.666 ms 242.650 ms 242.629 ms
```

```
5 * * *
```

```
6 218.241.96.193 (218.241.96.193) 244.294 ms 244.280 ms 244.498 ms
```

```
7 i.root-servers.net (192.36.148.17) 240.107 ms 240.306 ms 240.247 ms
```

- Second to last hop originated as ...
 - 218.241.96.0/20
 - AS 24151 (China Network Information Center)

US Client – Good Result (12 June 2010)

```
# dig @i.root-servers.net. www.facebook.com. A

; <<>> DiG X.X.X <<>> @i.root-servers.net. www.facebook.com. A
...
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 15872
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 13, ADDITIONAL: 14

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

;; AUTHORITY SECTION:
com.          172800 IN      NS      a.gtld-servers.net.
com.          172800 IN      NS      c.gtld-servers.net.

... more of the same ...

;; ADDITIONAL SECTION:
a.gtld-servers.net.  172800 IN      A       192.5.6.30
a.gtld-servers.net.  172800 IN      AAAA    2001:503:a83e::2:30

... more of the same ...
```

Same Chinese Client – Good Result! (12 June 2010)

```
$ dig @i.root-servers.net. www.facebook.com. A
```

```
; <<>> DiG X.X.X <<>> @i.root-servers.net. www.facebook.com. A
```

```
;; global options: printcmd
```

```
;; Got answer:
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 57990
```

```
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 13, ADDITIONAL: 14
```

```
;; QUESTION SECTION:
```

```
;;www.facebook.com.      IN      A
```

```
;; AUTHORITY SECTION:
```

```
com.      172800 IN      NS      l.gtld-servers.net.  
com.      172800 IN      NS      i.gtld-servers.net.  
com.      172800 IN      NS      m.gtld-servers.net.  
com.      172800 IN      NS      g.gtld-servers.net.  
com.      172800 IN      NS      j.gtld-servers.net.  
com.      172800 IN      NS      a.gtld-servers.net.  
com.      172800 IN      NS      b.gtld-servers.net.  
com.      172800 IN      NS      c.gtld-servers.net.  
com.      172800 IN      NS      h.gtld-servers.net.  
com.      172800 IN      NS      e.gtld-servers.net.  
com.      172800 IN      NS      d.gtld-servers.net.  
com.      172800 IN      NS      k.gtld-servers.net.  
com.      172800 IN      NS      f.gtld-servers.net.
```

Same Chinese Client – Good Result Continued (12 June 2010)

...

:: ADDITIONAL SECTION:

a.gtld-servers.net.	172800	IN	A	192.5.6.30
a.gtld-servers.net.	172800	IN	AAAA	2001:503:a83e::2:30
b.gtld-servers.net.	172800	IN	A	192.33.14.30
b.gtld-servers.net.	172800	IN	AAAA	2001:503:231d::2:30
c.gtld-servers.net.	172800	IN	A	192.26.92.30
d.gtld-servers.net.	172800	IN	A	192.31.80.30
e.gtld-servers.net.	172800	IN	A	192.12.94.30
f.gtld-servers.net.	172800	IN	A	192.35.51.30
g.gtld-servers.net.	172800	IN	A	192.42.93.30
h.gtld-servers.net.	172800	IN	A	192.54.112.30
i.gtld-servers.net.	172800	IN	A	192.43.172.30
j.gtld-servers.net.	172800	IN	A	192.48.79.30
k.gtld-servers.net.	172800	IN	A	192.52.178.30
l.gtld-servers.net.	172800	IN	A	192.41.162.30

:: Query time: 69 msec

:: SERVER: 192.36.148.17#53(i.root-servers.net.)

:: WHEN: Sat Jun 12 16:20:24 2010

:: MSG SIZE rcvd: 506

Same Chinese Client – Bad Result (12 June 2010 – 45 minutes later)

```
$ dig @i.root-servers.net. www.facebook.com. A
```

```
; <<>> DiG X.X.X <<>> @i.root-servers.net. www.facebook.com. A
```

```
;; global options: printcmd
```

```
;; Got answer:
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 52408
```

```
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0
```

```
;; QUESTION SECTION:
```

```
www.facebook.com.      IN      A
```

```
;; ANSWER SECTION:
```

```
www.facebook.com.     39245  IN      A      203.98.7.65 ← TelstraClear NZ
```

```
;; Query time: 15 msec
```

```
;; SERVER: 192.36.148.17#53(i.root-servers.net.)
```

```
;; WHEN: Sat Jun 12 17:05:10 2010
```

```
;; MSG SIZE rcvd: 66
```


Global problem .. or harmless local politics?

- Non-Chinese clients of Beijing I-root continue to receive correct answers.
- We see no recent evidence of bogus I-root responses outside of China.
- Netnod is doing everything cleanly: serving correct data, routing properly

Global problem .. or harmless local politics?

BUT ...

- The DNS injections observed in March outside of China are typical of what a Chinese client might see today within China – not just across these particular ASNs that leaked the I-root's domestic route. This is not a point source problem.
- F- and J-roots also have anycast Beijing instances, but they are not visible outside China for weeks and months at a time.
- As long as the route leak stands, I-root clients are at increased risk.

What does Netnod have to say?



<https://lists.dns-oarc.net/pipermail/dns-operations/2010-June/005724.html>

"What we understand from these discussions, the occurrence of these incorrect responses for queries sent to i.root-servers.net was a mistake. I have no insight into why or how the mistake happened, but we have been assured it won't be possible for it to happen again."

– Kurt Erik Lindqvist, CEO Netnod

Our Recommendation: Trust But Verify

- Root server operators should keep a very close eye on the routes people are using to reach their instances. Especially in “challenging networking environments.”
- It's great to operate domestic/local instances of global services. If that's your intent, though, you have an affirmative responsibility to keep them domestic/local.
- The NANOG community can take some responsibility here. Leaks happen several relationships from the source.
- Use your connections and clue to plug them!

Thank You

Doug Madory dmadory@renesys.com
Alin Popescu alin@renesys.com
Earl Zmijewski earl@renesys.com