A deep dive into heaps of Spoofed DNS Traffic

Lars Nøring
PowerTech

Introduction

- Lars Nøring (Noring)

Introduction

- Lars Nøring (Noring)
- PowerTech Information Systems AS
- ISP established in 1993
- DSL/fiber, SMB, Oslo

Preface

- ddos attack against ftp.powertech.no 195.159.0.153
- 195.159.0.153 is a secondary ip on that box, only for the ftp-service
- Observed rather large dns replies packets from it's secondary name server; 195.159.0.200.
- Believed it to be part of the attack. But it was slow, and didn't stop like the rest of the attack

Example DNS reply

```
UDP 195.159.0.200:53 → 195.159.0.153


;; QUESTION SECTION:
;szcourse214.edufe.cn.          IN      A

;; ANSWER SECTION:
szcourse214.edufe.cn.   3508    IN      CNAME   szcourse214.edufe.cn.cdn20.com.
szcourse214.edufe.cn.cdn20.com. 509 IN  CNAME   1st.ddwscache.glb0.lxdns.com.
1st.ddwscache.glb0.lxdns.com. 29 IN     A       121.11.151.70

;; AUTHORITY SECTION:
ddwscache.glb0.lxdns.com. 120   IN      NS      ns1.glb0.lxdns.com.
ddwscache.glb0.lxdns.com. 120   IN      NS      ns5.glb0.lxdns.com.
ddwscache.glb0.lxdns.com. 120   IN      NS      ns3.glb0.lxdns.com.
ddwscache.glb0.lxdns.com. 120   IN      NS      ns4.glb0.lxdns.com.
ddwscache.glb0.lxdns.com. 120   IN      NS      ns2.glb0.lxdns.com.
```

Investigating 195.159.0.200 - ns2

- Several similar requests from nearby IPs. The requests came from the range 195.159.0.101 to 195.159.1.42.
- But the requests did not actually originate from those IPs, some of them were not even in use.
- Checked flow on the border routers, and found the packets on a transit link.
- Investigated our other DNS servers. Found one other affected as well.

Characteristics

```
02:11:21.351623 IP (ttl 3, id 1083, proto UDP (17), length 66)
195.159.0.214.41918 > 195.159.0.200.53: 1083+ A? szcourse214.edufe.cn. (38)
```

- Spoofed src ip is +-100 of the dns-server (66 different ips in one case and 97 in the other)
- Same ip id and dns query id
- Same id for a long time. Changes in intervals between about 1 and 10 hours
- Same id towards both of our name servers at the same time
- Low ip TTL. Between 1 and 5 (Reply > 40)
- Src port between 32769 (2^15+1 or 1000000000000001) and 42767 (+9998)

More characteristics

```
szcourse214.edufe.cn.   3508    IN      CNAME   szcourse214.edufe.cn.cdn20.com.
szcourse214.edufe.cn.cdn20.com. 509 IN  CNAME   1st.ddwscache.glb0.lxdns.com.
1st.ddwscache.glb0.lxdns.com. 29 IN     A       121.11.151.70
```

- Related to cdn20 and/or lxdns.com (Both registered to ChinaNetCenter)
- 500 000 requests per day
- Requests for 2280 different domains in one day.

Requested domains

214	hk
216	biz
1024	org
1068	tv
2681	cc
21068	net
98363	cn
348198	com

404	perfect5.chinanetcenter.com.
406	perfect3.chinanetcenter.com.
415	perfect2.chinanetcenter.com.
418	perfect1.chinanetcenter.com.
419	perfect4.chinanetcenter.com

Packet Dump

Easy to catch the spoofed packets with tcpdump. For instance on 195.159.0.200:

```
tcpdump –n dst host 195.159.0.200 and dst port 53 and src net \
195.159.0.0/23 and ’ip[8] < 8’ and src portrange 32768-42768
```

A bit more complex with the response:

```
tcpdump –n src host 195.159.0.200 and src port 53 and dst net \
195.159.0.0/23 and dst portrange 32768-42768
```

This will catch too many packets. Filter packets containing lxdns.com or cdn20.net. Could also pay attention to udp query id or match towards incoming requests.

Block

Easy to block the packets, with as a good as no chance of blocking legit packets:

```
iptables -I INPUT -j ACCEPT -s 195.159.0.0/23 -p udp -m ttl \ --
ttl-lt 8 --source-port 32768:42768 --destination-port 53
```

Other affected name servers in our network?

- Collected information about all udp flows towards port 53 within our network on our border routers.
- Made a quick perl script that found all flows within the collected flows that had source ip within the same /24 as the destination ip (not accurate, but good enough).
- Got a lot of hits, but only regarding the two name servers we already knew about.

Other mentions of this online?

- One support issue in a forum of an Australian ISP. But they continued the conversation privately. They have not answered any of our emails.
- The NANOG list. But it was misinterpreted as a NAT-leak. They went on to argue about reverse path filtering.
- Ja.net CSIRT. We've been speaking with them a bit, and see a lot of the same and draw similar conclusions.

How JANET solved it

They contacted ChinaNetCenter and asked what this spoofed traffic was.

The reply was:

```
These are all normal. We are the CDN
service provider.
```

JANET asked them to stop the traffic. The following day, it ceased.

Normal CDN traffic?

Is it normal for CDNs to spoof DNS requests like that? Hardly. But is it «harmless» traffic from a CDN that tries to deliver a good service to its customers? Plausible.

The fear is that this is some sort of malicious attack, involving DNS poisoning or at least trying to map our randomness.

In a country far far away

- China is far away. More than 300 ms of round-trip-time from Norway, sometimes 500ms.
- Some resolvers give up its first request after 500 ms.
- With a little bit of back and forth, especially through several cname layers, a single A request easily exceeds several seconds.
- 5% packet loss towards some Chinese dns-server (2% is not unusual) (In addition there is the 500ms responses that might get dropped)

Somebody set up us the bomb

We contacted ChinaNetCenter and explained that we saw spoofed dns-requests we believed originated from them, and asked them to stop doing it. We also asked why they did it, and was open to discuss more proper ways of achieving what they wanted. They answered:

```
Could you send us the proof. We don't
spoof DNS requests.
```

I sent them some documentation. They answered

```
We will do some work to see how to
reduce the traffic.
```

Soon thereafter, the traffic stopped.

How widespread is it?

We tested the name servers of 19 Norwegian ISPs. We did 3 tests:

- Does the name server respond to queries where src ip is spoofed to be in the same network?
- Is the name server an open recursive resolver?
- Is the name server populated with cached entries from ChinaNetCenter?

ISP	IP	Spoof	Open	Populated
BKK	62.97.193.3			
Broadnet	82.196.201.43			
Dataguard	213.158.233.130			
DirectConnect	82.148.160.2			
Fasthost	80.65.49.14			
Get	84.208.20.110			
Homebase	84.38.159.242			
Lyse	213.167.96.50 (55)			
Lyse	81.167.36.3			
Netcom	212.169.123.67			
Netpower	212.33.131.67			
Nextgentel	217.13.4.24			
Safety Computing	82.199.2.201			
Signal/IT connect	80.89.32.10			
Tafjord Mimer	213.184.200.1			
TDC	62.65.30.10			
Tele2	193.216.1.10			
Telenor	148.122.161.3			
Ventelo	193.75.75.75			
Google	8.8.8.8			

ISP	IP	Spoof	Open	Populated
BKK	62.97.193.3	Yes		
Broadnet	82.196.201.43	Yes		
Dataguard	213.158.233.130	Yes		
DirectConnect	82.148.160.2	Yes		
Fasthost	80.65.49.14	No		
Get	84.208.20.110	No		
Homebase	84.38.159.242	Yes		
Lyse	213.167.96.50 (55)	Yes		
Lyse	81.167.36.3	Yes		
Netcom	212.169.123.67	No		
Netpower	212.33.131.67	Yes		
Nextgentel	217.13.4.24	No		
Safety Computing	82.199.2.201	Yes		
Signal/IT connect	80.89.32.10	No		
Tafjord Mimer	213.184.200.1	No		
TDC	62.65.30.10	No		
Tele2	193.216.1.10	Yes		
Telenor	148.122.161.3	Yes		
Ventelo	193.75.75.75	No		
Google	8.8.8.8	No		
		11 of 20		

ISP	IP	Spoof	Open	Populated
BKK	62.97.193.3	Yes	No	
Broadnet	82.196.201.43	Yes	No	
Dataguard	213.158.233.130	Yes	No	
DirectConnect	82.148.160.2	Yes	No	
Fasthost	80.65.49.14	No	No	
Get	84.208.20.110	No	No	
Homebase	84.38.159.242	Yes	Yes	
Lyse	213.167.96.50 (55)	Yes	No	
Lyse	81.167.36.3	Yes	No	
Netcom	212.169.123.67	No	No	
Netpower	212.33.131.67	Yes	No	
Nextgentel	217.13.4.24	No	No	
Safety Computing	82.199.2.201	Yes	No	
Signal/IT connect	80.89.32.10	No	Yes	
Tafjord Mimer	213.184.200.1	No	No	
TDC	62.65.30.10	No	No	
Tele2	193.216.1.10	Yes	No	
Telenor	148.122.161.3	Yes	No	
Ventelo	193.75.75.75	No	No	
Google	8.8.8.8	No	Yes	
		11 of 20	3 of 20	

ISP	IP	Spoof	Open	Populated
BKK	62.97.193.3	Yes	No	N/A
Broadnet	82.196.201.43	Yes	No	N/A
Dataguard	213.158.233.130	Yes	No	Yes
DirectConnect	82.148.160.2	Yes	No	No
Fasthost	80.65.49.14	No	No	No
Get	84.208.20.110	No	No	N/A
Homebase	84.38.159.242	Yes	Yes	Yes
Lyse	213.167.96.50 (55)	Yes	No	No
Lyse	81.167.36.3	Yes	No	Yes
Netcom	212.169.123.67	No	No	No
Netpower	212.33.131.67	Yes	No	Yes
Nextgentel	217.13.4.24	No	No	No
Safety Computing	82.199.2.201	Yes	No	N/A
Signal/IT connect	80.89.32.10	No	Yes	No
Tafjord Mimer	213.184.200.1	No	No	N/A
TDC	62.65.30.10	No	No	N/A
Tele2	193.216.1.10	Yes	No	No
Telenor	148.122.161.3	Yes	No	No
Ventelo	193.75.75.75	No	No	N/A
Google	8.8.8.8	No	Yes	No
		11 of 20	3 of 20	4 of 13

Spoofing details

Create a dns zone and entry

```
labs.noring.no.                IN NS 195.159.11.12

*.loopback.labs.noring.no. IN A  127.0.0.1
```

Send a spoofed packet requesting a unique host name

```
UDP 8.8.8.9 -> 8.8.8.8:53 A?
unique8888.loopback.labs.noring.no.
```

See if the request reaches the authoritative name server

Details around populated cache

We sent A requests for ChinaNetCenter hosts from within the ISPs network.

```
UDP -> 8.8.8.8:53 A? szcourse214.edufe.cn.
```

A cached response is received in less than 1s:

```
szcourse214.edufe.cn.   3508    IN      CNAME   szcourse214.edufe.cn.cdn20.com.
szcourse214.edufe.cn.cdn20.com. 509 IN  CNAME   1st.ddwscache.glb0.lxdns.com.
1st.ddwscache.glb0.lxdns.com. 29 IN     A       121.11.151.70
```

The TTL indicates that it was already cached.
A non-cached response will take several seconds and have a typical initial TTL value:

```
szcourse214.edufe.cn.   3600    IN      CNAME   szcourse214.edufe.cn.cdn20.com.
szcourse214.edufe.cn.cdn20.com. 600 IN  CNAME   1st.ddwscache.glb0.lxdns.com.
1st.ddwscache.glb0.lxdns.com. 120 IN     A      121.11.151.70
```

Conclusion

- It is probable that ChinaNetCenter are behind it and doing it to improve the speed of the content of their customers, without any malicious intent
- It works to some extent, as many ISPs are open to spoofed dns requests
- But they could do better
- If it was malicious traffic, it is no longer attacking us

Aftermath

- A couple of months later, the traffic ceased altogether

end

Lars Nøring
lars@powertech.no