



DNS Fast-Flux



John Kristoff
Research Analyst
jtk@cymru.com



First... a primer





DPU netseminar
2009-10-16

John Kristoff – Team Cymru


2

I need an IPv4 address for
www.cdm.depaul.edu.

Please get it (recursion desired) for me?




Check cache.
If empty, ask a parent.
Follow delegation if necessary.



Local Caching Server
(full resolver)

parent zones: cdm.depaul.edu.
depaul.edu.
edu.

.




**Let's assume cache is empty,
and all it knows about is (.) root.***

A.root-servers.net

...


M.root-servers.net

***Do you see why a reliable and trustworthy root is so important?**




I need an IPv4 address for
www.cdm.depaul.edu.


Can you tell me or refer me to someone?



Local Caching Server
(full resolver)




root (.) server




Don't know.
Try one of these .edu servers:

A.GTLD-SERVERS.NET
C.GTLD-SERVERS.NET
D.GTLD-SERVERS.NET
E.GTLD-SERVERS.NET
F.GTLD-SERVERS.NET
G.GTLD-SERVERS.NET
L.GTLD-SERVERS.NET



Local Caching Server
(full resolver)




root (.) server




**Does the caching server have
something in its cache now?**

Raise your hand for yes.




Ultimately we should get here...




You've come to the right place.
The authoritative answer is:

140.192.32.142


and that answer is valid
for 3600 seconds


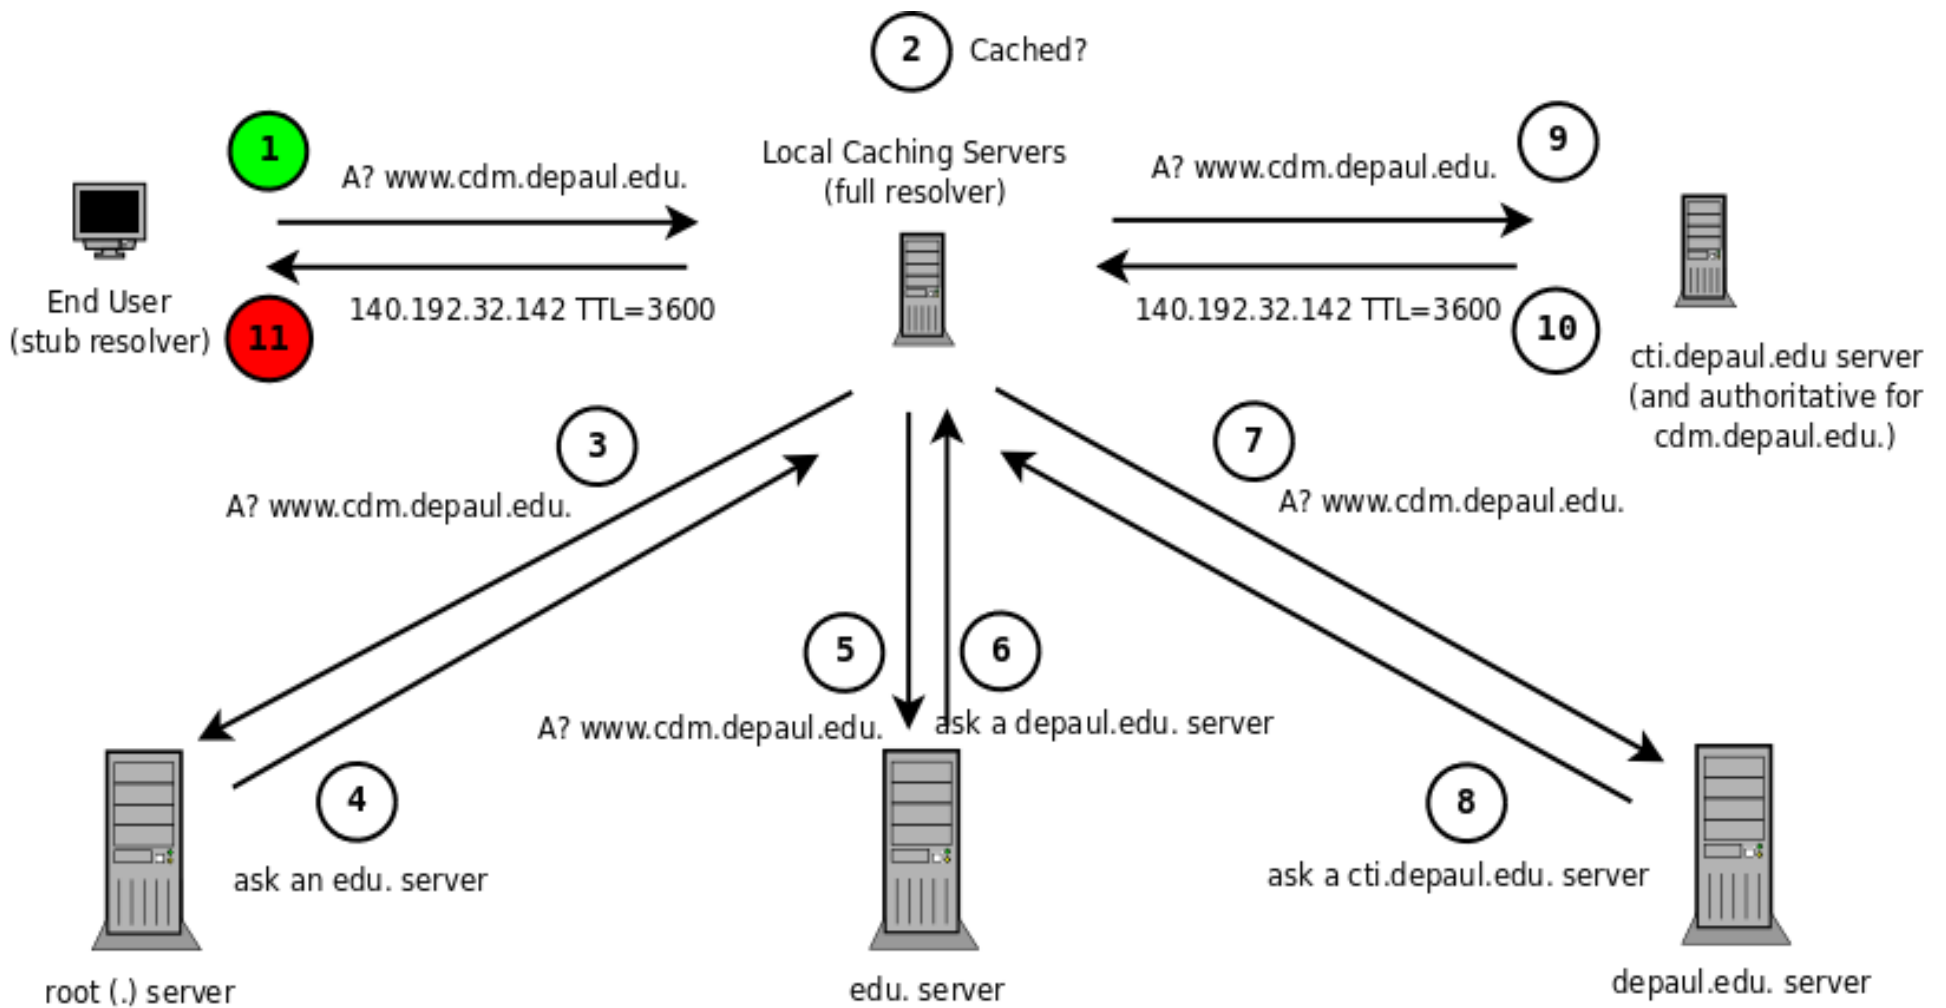


Local Caching Server
(full resolver)



ns1.cti.depaul.edu.
or
ns2.cti.depaul.edu.
or
ns3.cti.depaul.edu.





Some DNS terminology review

- RR or RRset
 - Resource record or RR set, one or more records containing information about a domain name
- A
 - DNS RR of type A, for IPv4 address record(s)
- NS
 - DNS RR of type NS, for name servers names authoritative for a particular zone



Dig output of an “A” query

```
$ dig www.depaul.edu @ns1.depaul.edu
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 12345  
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 3
```

```
;; QUESTION SECTION:  
www.depaul.edu.          IN      A
```

```
;; ANSWER SECTION:  
www.depaul.edu.          300     IN      A       140.192.23.180
```

```
;; AUTHORITY SECTION:  
depaul.edu.              86400   IN      NS      ns1.depaul.edu.  
depaul.edu.              86400   IN      NS      ns2.depaul.edu.  
depaul.edu.              86400   IN      NS      ns4.depaul.edu.
```

```
;; ADDITIONAL SECTION:  
ns1.depaul.edu.          86400   IN      A       140.192.0.2  
ns2.depaul.edu.          86400   IN      A       140.192.239.2  
ns4.depaul.edu.          86400   IN      A       64.30.240.250
```



Dig output of an “NS” query

```
$ dig ns depaul.edu @ns1.depaul.edu
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 12345  
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 3, AUTHORITY: 0, ADDITIONAL: 3
```

```
;; QUESTION SECTION:
```


```
;depaul.edu.                IN      NS
```

```
;; ANSWER SECTION:
```

```
depaul.edu.                86400  IN      NS      ns1.depaul.edu.  
depaul.edu.                86400  IN      NS      ns2.depaul.edu.  
depaul.edu.                86400  IN      NS      ns4.depaul.edu.
```


```
;; ADDITIONAL SECTION:
```

```
ns1.depaul.edu.          86400  IN      A       140.192.0.2  
ns2.depaul.edu.          86400  IN      A       140.192.239.2  
ns4.depaul.edu.          86400  IN      A       64.30.240.250
```



Affecting availability with DNS

- The RRs in an answer or NS RRset
- RRset TTL
- Unique answer based on origin (geoloc/views)
- Unique answer based on time
- Wildcards, answering authoritatively



Different origin, different answer

```
$ dig www.google.com
```


```
;; ANSWER SECTION:
```

www.google.com.	604800	IN	CNAME	www.l.google.com.
www.l.google.com.	300	IN	A	74.125.95.147
www.l.google.com.	300	IN	A	74.125.95.99
www.l.google.com.	300	IN	A	74.125.95.103
www.l.google.com.	300	IN	A	74.125.95.104

```
$ dig www.google.com @4.2.2.2
```

```
;; ANSWER SECTION:
```

www.google.com.	43190	IN	CNAME	www.l.google.com.
www.l.google.com.	300	IN	A	209.85.171.103
www.l.google.com.	300	IN	A	209.85.171.104
www.l.google.com.	300	IN	A	209.85.171.147
www.l.google.com.	300	IN	A	209.85.171.99



A RRset fast-fluxing

```
$ dig vqthe.cn
```

```
;; ANSWER SECTION:
```

vqthe.cn.	180	IN	A	89.46.127.47
vqthe.cn.	180	IN	A	123.237.100.126
vqthe.cn.	180	IN	A	123.237.108.142
vqthe.cn.	180	IN	A	190.191.142.122
vqthe.cn.	180	IN	A	196.202.6.66
vqthe.cn.	180	IN	A	71.239.64.226
vqthe.cn.	180	IN	A	78.92.180.208
vqthe.cn.	180	IN	A	85.254.64.153

```
;; AUTHORITY SECTION:
```

vqthe.cn.	180	IN	NS	ns2.kr crab.com.
vqthe.cn.	180	IN	NS	ns1.czwill.com.
vqthe.cn.	180	IN	NS	ns4.kr crab.com.
vqthe.cn.	180	IN	NS	ns2.czwill.com.

```
;; ADDITIONAL SECTION:
```

ns1.czwill.com.	172799	IN	A	78.92.180.208
ns2.czwill.com.	172799	IN	A	85.67.171.146
ns2.kr crab.com.	172799	IN	A	61.61.61.61
ns4.kr crab.com.	172799	IN	A	138.16.6.201



Re-query, notice changes

```
$ dig vqthe.cn
```

```
;; ANSWER SECTION:
```


vqthe.cn.	180	IN	A	85.67.171.146
vqthe.cn.	180	IN	A	89.44.56.76
vqthe.cn.	180	IN	A	89.102.112.60
vqthe.cn.	180	IN	A	116.72.241.170
vqthe.cn.	180	IN	A	124.125.245.32
vqthe.cn.	180	IN	A	190.245.216.89
vqthe.cn.	180	IN	A	79.140.228.27
vqthe.cn.	180	IN	A	85.29.210.207

```
;; AUTHORITY SECTION:
```

vqthe.cn.	180	IN	NS	ns1.czwill.com.
vqthe.cn.	180	IN	NS	ns2.krcrab.com.
vqthe.cn.	180	IN	NS	ns2.czwill.com.
vqthe.cn.	180	IN	NS	ns4.krcrab.com.


```
;; ADDITIONAL SECTION:
```

ns1.czwill.com.	172585	IN	A	78.92.180.208
ns2.czwill.com.	172585	IN	A	85.67.171.146
ns2.krcrab.com.	172585	IN	A	61.61.61.61
ns4.krcrab.com.	172585	IN	A	138.16.6.201




Single-flux vs double-flux

- If you know two more buzzwords than the other guy, you're an expert
- Single-flux: the A RRs in the answer flux
- Double-flux: the authoritative name servers flux too
 - authoritative name server IP addresses are changing (perhaps as far up as a parent zone into the TLD where glue records are changing, possibly automated through a registrar's API).




Why is this cool? or... Why is this bad?

- Its a great way to ensure availability
- Taking away any single host has almost no impact
- How do you take down potentially dozens, if not hundreds of hosts participating in the A RRset?
- Take down the name?
 - Not all registrars or registries are willing and/or able to support this whack-a-mole process




Law enforcement problem

Fast-flux makes it literally impossible for them to bust bad guys. They are just not equipped to deal with these challenges and infrastructure. :-)



How to address the problem?

- Debatable, everything is duct tape and bubble gum
- ICANN/GNSO talking a lot about this
 - I'm not convinced they'll ultimately have much of an impact
- Registrars could potentially do a lot
 - For that to really work/happen, that probably means stricter procedures – this is something ICANN can have an impact on




DNS fast-flux detection ideas

- A/ASN RR diversity
- A/NS count
- A/NS history
- A/ASN reputation
- A/NS fingerprinting
- TTLs
- fwd/rev mapping
- whois - registrar
- whois – dates
- name structure
- name reputation
- NS query probing
- parent/child NS consistency
- A RR HTTP probing
- A/NS fingerprints
- Passive DNS history
- query history



Unobtrusive fast-flux detection using variance

pretty catchy title eh?




pcapff

- From DNS pcap, for each answer RRset examine:
 - qname, answer A RRset, TTL
- If current A RRset \neq past A RRset, mark as ff
 - *note: this alone results in some some FPs
- Maybe augment with...
 - History, IP/name reputation, white/black list, etc.
 - Obtrusive analysis




abbeynational97.com|10|300|410393465|9698
bcneoise.cn|8|180|273412252|14806|m
best-rx-onlinestore.com|6|300|945584867|7200
casinotreasure.com|8|300|25116914|7626
testnameserver.com|10|300|673397541|8357|m
thinkaboutus.cn|6|300|945584867|7200
zxjoiwgc.cn|8|180|775935496|15065|m

abbeynational97.com|10|300|452615284|12837|f
bcneoise.cn|8|180|355949343|8127|f
best-rx-onlinestore.com|6|300|945584867|7200
casinotreasure.com|8|300|21206785|7626
testnameserver.com|10|300|673397541|8357|m
thinkaboutus.cn|6|300|945584867|7200
zxjoiwgc.cn|8|180|355949343|8127|f




TODO

- Evaluate FPs and reduce/remove them
 - e.g. NTP pools, IRC nets, content distribution providers
- Retain unobtrusive analysis
 - No querying of the fast-flux names directly
 - No active probing of A RRs
- Avoid hacks
 - White/black lists
 - IP reputation
 - “scoring” algorithms



References

- <http://forum.icann.org/lists/gnso-ff-pdp-may08/>
- <http://www.icann.org/en/committees/security/sac025.pdf>
- <http://www.uoregon.edu/~joe/fastflux/simple.cgi>
- <http://atlas.arbor.net/summary/fastflux>
- <http://honeyblog.org/junkyard/paper/fastflux-malware08.pdf>
- <http://www.honeynet.org/papers/ff>



Contact us

- jtk@cymru.com
- Team PGP key 0x79B109F9

<http://www.team-cymru.org/About/teamcymru-pgp.txt>

