

Measuring EDNS Client Subnet Extension

IEPG, IETF 87, Berlin

Florian Streibelt

<florian@inet.tu-berlin.de>

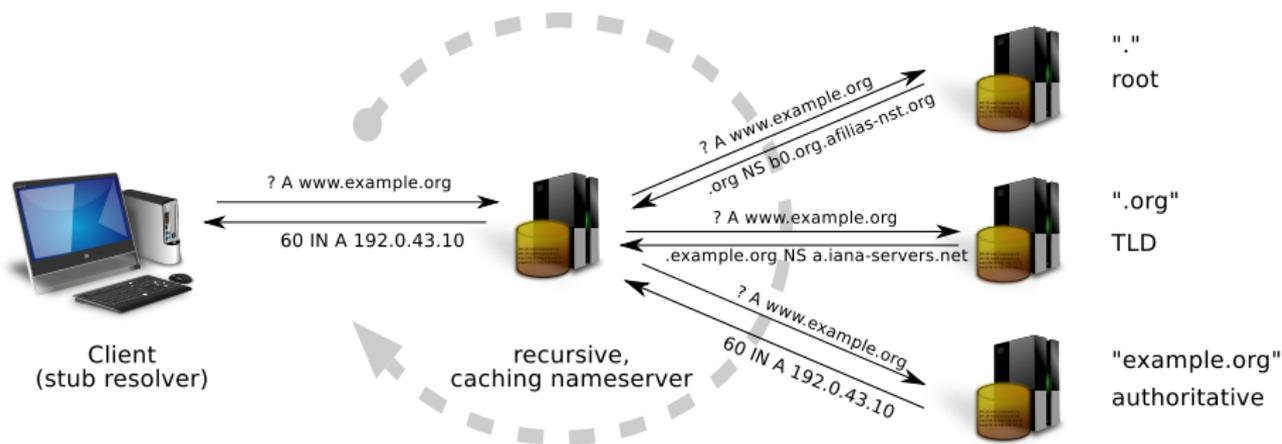
TU-Berlin, FG INET

www.inet.tu-berlin.de

July 28th 2013

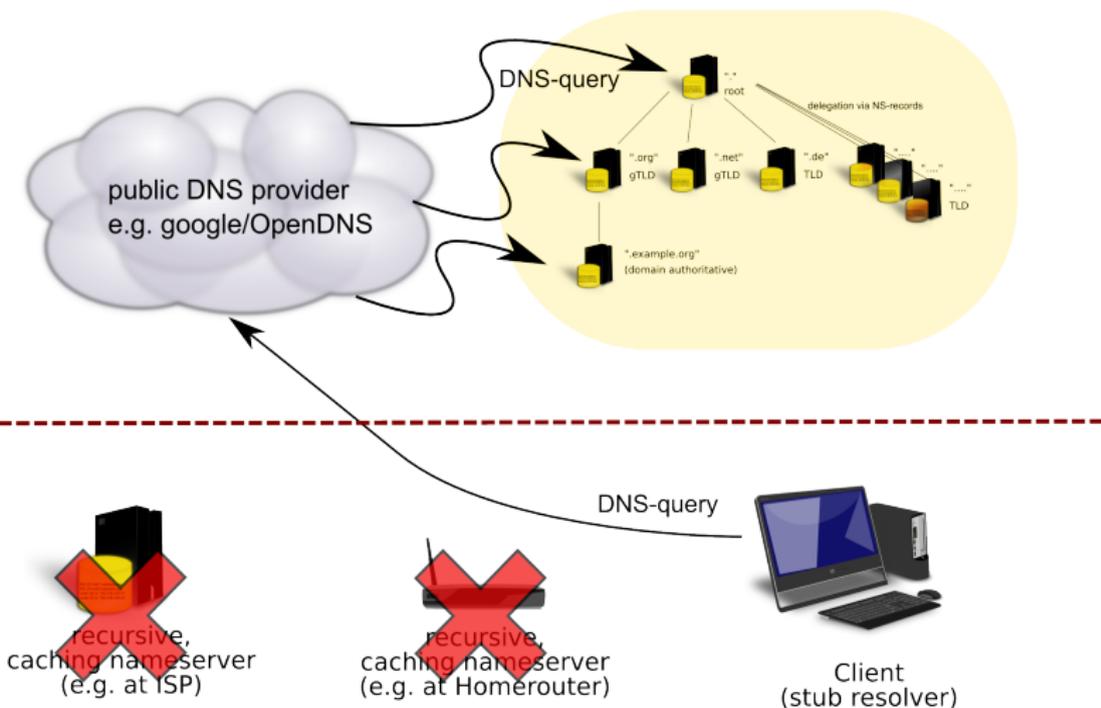
Preliminary results, full results at IMC 2013 (see last slide)

Textbook DNS-Lookup



- Stub resolver on the client asks a recursor (e.g., at the ISP)
- Recursor follows the delegation

Today: Public DNS usage increases



Otto et al. [2]: usage at 8.6% in December 2011

CDNs/CPs loose control

- Non-ISP Resolvers are gaining momentum
- CDNs using the DNS request origin for client-location thus are *blinded*
- Workarounds exist but don't scale well - e.g. check against known list of google NS IPs and their geolocation¹

¹<https://developers.google.com/speed/public-dns/faq#locations>

Introducing: Client IP information in EDNS (CIP)

- proposal by google, OpenDNS and others:
<http://afasterinternet.com/>
- EDNS0 extension to transport Client IP information:
<http://tools.ietf.org/html/draft-vandergaast-edns-client-subnet-02>
- Recursor adds client IP-information (usually a netmask) to the query directed at the authoritative NS
- Client-specific answers, e.g. based on geolocation, are again possible
- Scope to allow caching is returned in the answer

⇒ We can impose every 'location' using arbitrary Client IP information

Protocol: Client IP information (CIP) in EDNS

```
# dig www.google.com +client="100.100.100.101"
/-----\
|Contents of|          option length (8)
|Additional |          |
|Section for|          |      adress family (1=IPv4)
|EDNS, CIP  |          |      |
\-----~/          |      |      source netmask(=32)
                    |      |      |
                    EDNS-CIP    |      |      scope netmask
                    Option code |      |      |
                    |           |      |      |      |CLIENT-IP
                    |---\    |---\    |---\    |      |      |-----\
query:      00 08 00 08 00 01 20 00 64 64 64 65
response:   00 08 00 08 00 01 20 20 64 64 64 65
old code:   50 fa                                ~~
```

Measurements

- Single vantage point is sufficient to *arbitrary* Client IP/mask
- We can use all network prefixes collected by RIPE/Routeviews
- Subset of our experiments:
 - Compare scopes to original prefix lengths
 - How do scopes differ between DNS-providers?
 - Relation between A-Records and Client-IPs?
 - Find datacenters/global footprint of adopters

Looking at the A-Records

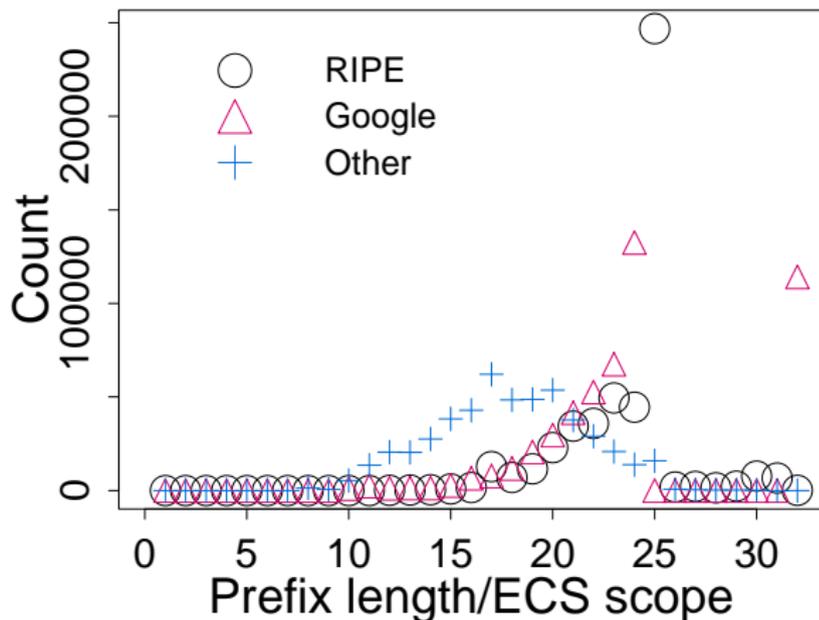
Using google as example:

- resolving `www.google.com` via `ns1.google.com`
- using all network prefixes from RIPE route collection
- repeated after 3 months
- 8,735 (6,284) frontend IP addresses (not servers)
- 282 (163) ASes
- 52 (47) countries
- within 3 months: 40% increase of IPs
- also: in various non-google datacenters

see also:

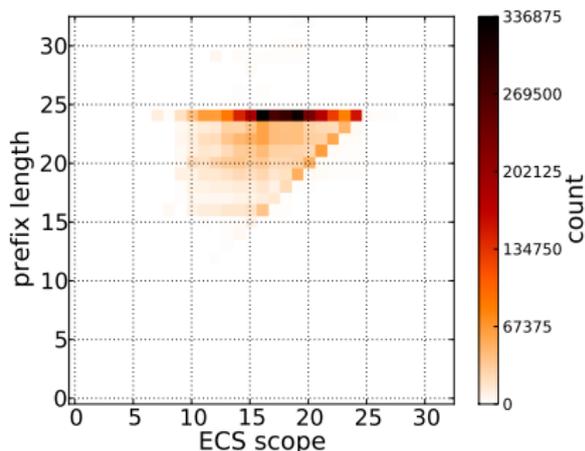
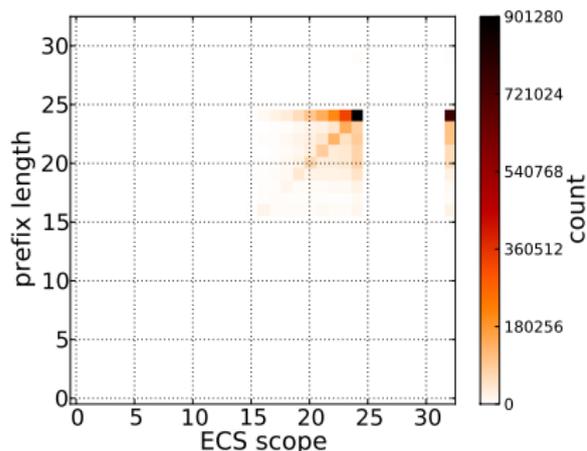
Calder et al.: Mapping the Expansion of Google's Serving Infrastructure [1]

RIPE prefix length vs. CIP-scopes



Prefix length and scope distribution do not match and differ between adopters

Comparing google and another adopter



The smaller ECS-adopter (right) aggregates while google (left) returns more specific scopes.

Conclusion

- Enabling Client IP Information gives better performance for clients
- This comes with a tradeoff: it also reveals internal information
- It enables researchers (and competitors) to investigate e.g. global footprint, user-to-server mapping
- By chance it reveals more information than desired (server and service distribution)
- No filtering e.g. based on number of client prefixes was observed
- Future Adopters should be aware of these facts

Bibliography I

- [1] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John Heidemann, and Ramesh Govindan.
Mapping the expansion of Google's serving infrastructure.
Technical Report TR 13-935, University of Southern California
Computer Science Department, June 2013.
- [2] John S. Otto, Mario A. Sánchez, John P. Rula, and Fabián E. Bustamante.
Content delivery and the natural evolution of dns: remote dns
trends, performance issues and alternative solutions.
*In Proceedings of the 2012 ACM conference on Internet
measurement conference, IMC '12*, pages 523–536, New York,
NY, USA, 2012. ACM.

Contact:

Florian Streibelt <florian@inet.tu-berlin.de>

Related publication:

Unintended Consequences: Exploring EDNS-Client-Subnet Adopters in your Free Time

Internet Measurement Conference, October 2013

<http://conferences.sigcomm.org/imc/2013/>

Authors:

Florian Streibelt, Jan Böttger, Nikolaos Chatzis, Georgios Smaragdakis, Anja Feldmann, Walter Willinger

The paper, software and raw data will be published in October 2013.

Image sources:

own work and <http://openclipart.org/>