Exploring EDNS-Client-Subnet Adopters in Your Free Time

IMC 2013, Barcelona

Florian Streibelt

<florian@inet.tu-berlin.de>

TU-Berlin, Germany - FG INET

www.inet.tu-berlin.de

October 24th 2013

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

With special thanks to Walter Willinger.

Non-ISP (aka 'public') DNS usage increases

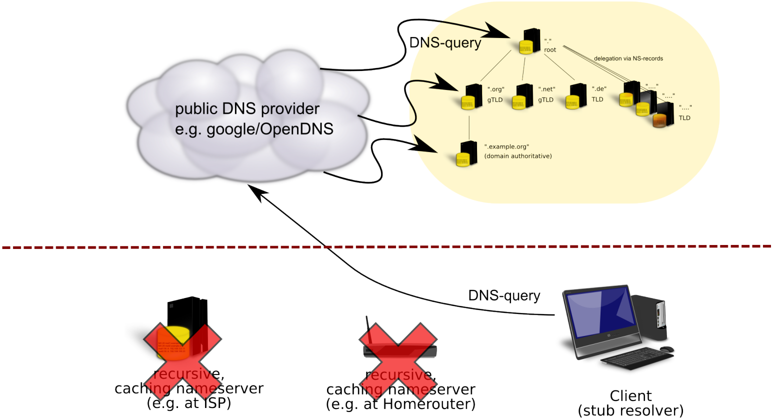

Usage at 8.6% in December 2011

According to Otto et al. in "Content delivery and the natural evolution of DNS: remote DNS trends, performance issues and alternative solutions" (IMC 2012)

Challenge for CDNs/CPs

- Non-ISP resolvers are gaining momentum
- Clients are far away from resolvers
- CDNs often make heavy use of DNS for client location
- Using the DNS request origin for client-location now leads to (more) wrong results
- Mis-location of clients gives end-users bad performance

Introducing: Client IP information in EDNS (ECS)

- Recursive nameserver adds client subnet information (network prefix) to the query directed at the authoritative nameserver
- EDNS0 extension is introduced to transport this data
- Proposal by Google, OpenDNS and others (A faster Internet consortium)
- Performance gain can be observed, again see Otto et al. (IMC 2012)
- We find roughly 13% of the top 1M Alexa list seem to support this extension already

(Ab)using ECS for Measurements

Intended use of ECS:

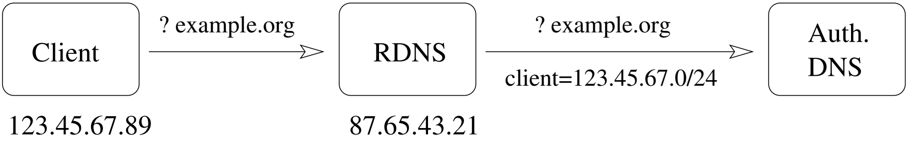

(Ab)using ECS for Measurements

Intended use of ECS:

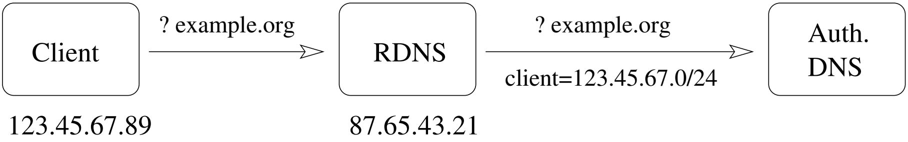

Doing our measurements:

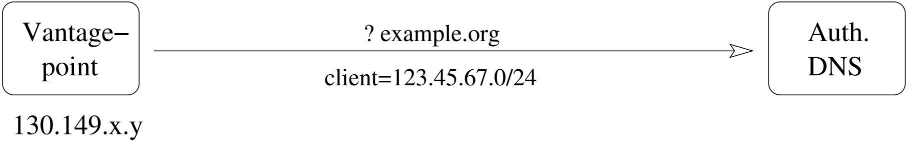

(Ab)using ECS for Measurements

Intended use of ECS:

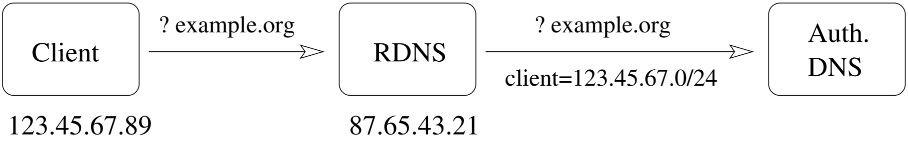

Doing our measurements:

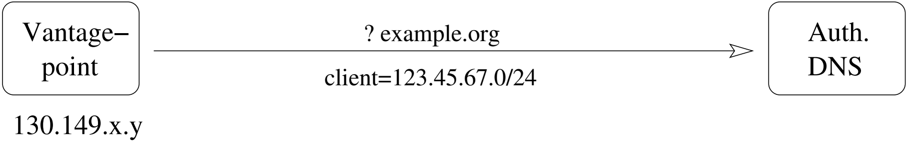

⇒ We can impose every client 'location'.

Protocol: Client IP information in EDNS (ECS)

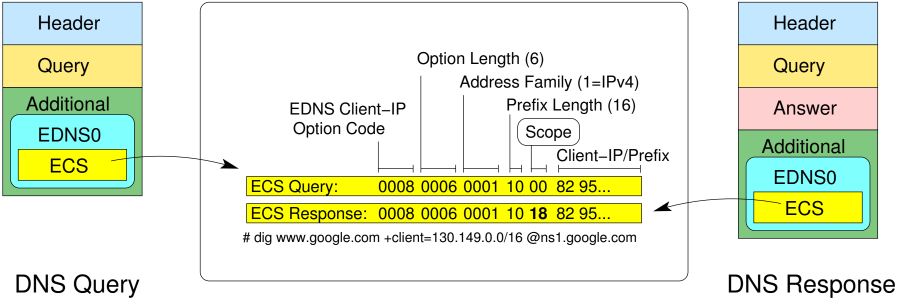

- The scope returned allows for caching (applied as netmask)
- The client IP information cannot be checked

ECS as a Measurement Tool

- Using arbitrary client subnet information, we can impose every client 'location'
- This gives us the opportunity to
  - find the location of CDN caches within ISPs,
  - observe the growth of CDN footprints,
  - infer client-to-server mappings (to some extend),
  - analyze dynamic changes by repeated measurements.
- As demonstration we present a subset of our experiments, using Google as example.

Measurements

- Single vantage point[1] is sufficient to use *arbitrary* Client IP/prefix
- As Client Subnets we use all network prefixes from RIPE RIS (sanity check using Routeviews)
- We compare with Client Subnets derived from:
  popular resolvers, subnets of an ISP, educational networks
- Measurements are done for: Google/YouTube, MySqueezebox, Edgecast and others
- Data to look at: A-records (servers) and scope (caching) returned

[1] we checked from four different locations

Comparing sources for Client Subnets

	Prefix set	Server	Sub	AS	Countries
Google (03/26/13)	RIPE	6,340	329	166	47
	RV	6,308	328	166	47
	PRES	6,088	313	159	46
	ISP	207	28	1	1
	ISP24	535	44	2	2
	UNI	123	13	1	1

- RIPE RIS and Routeviews give nearly identical results
- The 280k most popular resolvers, as seen by a CDN, yield similar results – but dataset is not freely available
- Mapping to GGCs is working, as can been seen at the UNI and ISP datasets

Looking at the A-Records of Google

- Resolving www.google.com via `ns1.google.com`
- Using all network prefixes from RIPE RIS as client subnets
- Different synchronized vantage points (plausibility check)

Date (RIPE)	IPs	Sub nets	ASes	Countries
2013-03-26	6340	329	166	47
2013-03-30	6495	332	167	47
2013-04-13	6821	331	167	46
2013-04-21	7162	346	169	46
2013-05-16	9762	485	287	55
2013-05-26	9465	471	281	52
2013-06-18	14418	703	454	91
2013-07-13	21321	1040	714	91
2013-08-08	21862	1083	761	123

see also the next presentation:

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

Looking at the A-Records of Google

Selected results from combined experiments:

- We see GGC (Google Global Cache edge servers) in various ISP networks
- These ISPs are not allowed to advertise the GGC, but we are
- Huge increase in the footprint can be observed, also for YouTube
- Comparing results from different vantage points we observe redirection of clients and prefixes, probably due to load balancing the GGCs
- We see that most of the time clients indeed are served from caches in their respective AS
- We see large overlap in the returned A records in the results from the different vantage points, both for Google and YouTube

Comparing Google and Edgecast Scopes

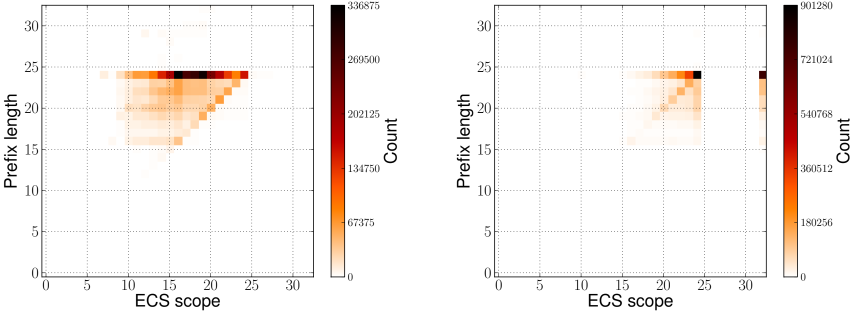

Edgecast (left) aggregates while Google (right) returns more specific scopes.

Conclusion

- Enabling ECS gives better performance for clients
- This comes with a tradeoff for DNS providers and CDNs: it also reveals internal information
- It enables researchers (and competitors) to investigate e.g. global footprint, growth-rate, user-to-server mapping, etc.
- No filtering e.g. based on number of client prefixes was yet observed
- We show that this extension offers interesting opportunities for measurements

Contact:

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

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

http://projects.inet.tu-berlin.de/projects/ecs-adopters/wiki

Image sources:

own work and http://openclipart.org/

A Textbook DNS-Lookup

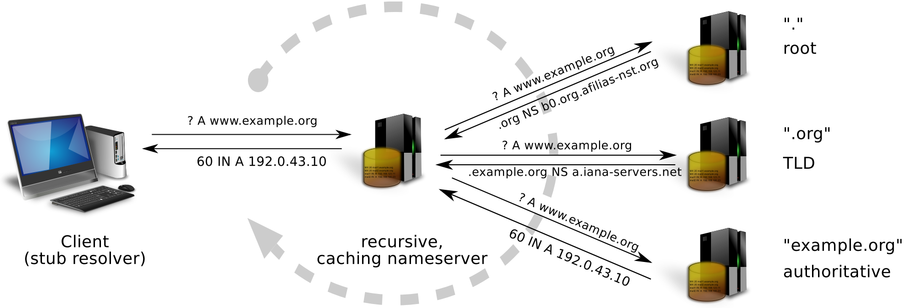

- Client asks a recursive nameserver (e.g., at the ISP)
- This nameserver follows the delegation, contacts the authoritative server
- Assumption: Client located near the recursive nameserver

How to enable ECS?

- Primary nameservers must be ECS enabled (Supported by PowerDNS: yes, Bind: no)
- If there are other systems in front: these as well
- Primary nameservers need to be whitelisted (manually) by e.g., OpenDNS, Google
- Note: We find that roughly 13% of the top 1 million domains (Alexa) may be already ECS enabled.

Framework used

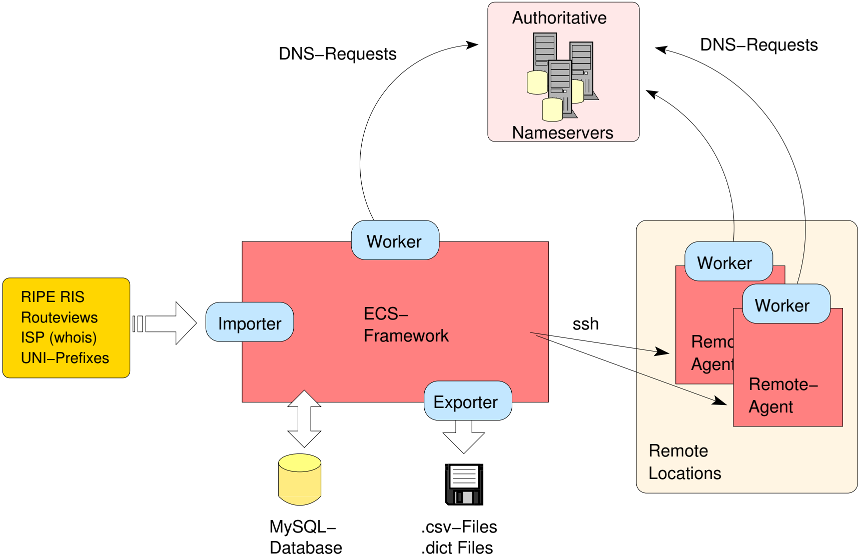

RIPE RIS prefix length vs. ECS-scopes

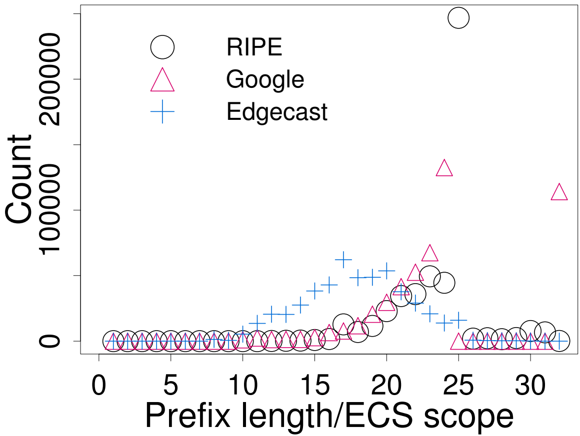

Prefix length and scope distribution do not match and differ between adopters, also note the /32s!

Client and AS mappings

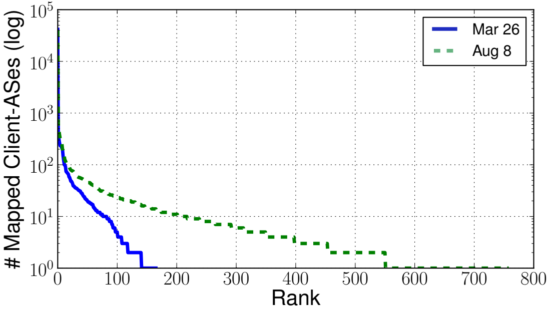

In August we see more ASes served from more than one 'server-AS'.