Understanding the Growth and Security Considerations of ECS

A Longitudinal Study Exploring the Deployment of a DNS Extension

Athanasios Kountouras*, Panagiotis Kintis*, Athanasios Avgetidis*,
Thomas Papastergiou*, Chaz Lever*,
Michalis Polychronakis†, Manos Antonakakis*
*Georgia Institute of Technology,†Stony Brook University
{kountouras, kintis, avgetidis, tpapastergiou, chazlever, manos}@gatech.edu, mikepo@cs.stonybrook.edu

***Abstract*—The Domain Name System (DNS) is fundamental to communication on the Internet. Therefore, any proposed changes or extensions to DNS can have profound consequences on network communications. In this paper, we explore the implications of a recent extension to DNS called EDNS Client Subnet (ECS). This extension extends the visibility of client information to more domain operators by providing a prefix of a client's IP address to DNS nameservers above the recursive nameserver. This raises numerous questions about the impact of such changes on network communications that rely on DNS.**

In this paper, we present the results of a longitudinal study that measures the deployment of ECS using several DNS vantage points. We show that, despite being an optional extension, ECS has seen steady adoption over time—even for sites that do not benefit from its use. Additionally, we observe that the client subnet provided by ECS may provide less privacy than originally thought, with most subnets corresponding to a /24 CIDR or smaller. Lastly, we observe several positive and negative consequences resulting from the introduction of DNS. For example, DNS can help aid security efforts when analyzing DNS data above the recursive due to the addition of client network information. However, that same client information has the potential to exacerbate existing security issues like DNS leakage. Ultimately, this paper discusses how small changes to fundamental protocols can result in unintended consequences that can be both positive and negative.

I. Introduction

The Domain Name System (DNS) is a fundamental network protocol on the Internet. Its most visible function is enabling humans to remember simple names to reach network resources. However, it is also used by countless applications and is critical to the operation of many services on the Internet. In fact, attacks targeting DNS infrastructure [8] have been responsible for taking down large portions of the Internet. However, it is not just users, benign applications, and services that rely on DNS for communication. DNS is used by malicious actors to provide resilient communication for malware [23], [24], and security researchers frequently use DNS to track malicious abuse on the Internet [9], [10], [21], [28], [31]. Due to its importance to communication in IP based networks, proposed changes or extensions to DNS should be well understood.

Network and Distributed Systems Security (NDSS) Symposium 2021
21-25 February 2021, Virtual
ISBN 1-891562-66-5
https://dx.doi.org/10.14722/ndss.2021.24343
www.ndss-symposium.org

This paper studies the deployment of a DNS extension called EDNS Client Subnet (ECS) [15], which was originally proposed as an experimental Internet draft in 2011 [14]. ECS changes how client IP information is shared with DNS infrastructure in order to optimize Content Delivery Network (CDN) selection. It provides more granular client information to authoritative DNS servers that reveals information about the underlying client making a request. This entire process is transparent to end users who receive no indication whether ECS will be used by their recursive DNS server. ECS is currently "on by default" for all traffic through many of the largest open DNS recursive servers. To date, this extension has been adopted by many of the largest open DNS providers on the Internet—including Google Public DNS, OpenDNS, Quad9, and NextDNS [2]–[5].

The result of this widespread adoption is that DNS client information is now shared across many networks on the Internet that previously did not have access to such information. This is potentially problematic because DNS sometimes leaks information about user behavior. For example, the automated DNS behavior of some applications (e.g., web browsers) can reveal limited, indirect information about local users [32]. For this reason, it is recommended that DNS prefetching is disabled for sensitive applications. In addition, DNS attacks on Tor anonymity have been demonstrated, but are either patched or are based on the non-trivial capability of monitoring the recursive footprint used by most Tor nodes [20]. Furthermore, information leaks through anonymity networks have long been addressed by SOCKS tunneling of UDP queries [30]. Beyond these potential security issues, ECS changes the privacy expectations traditionally associated with the use of a shared recursive DNS server. Therefore, it is important that potentially negative repercussions of ECS are studied and better understood.

Beyond the negative side effects briefly discussed above, we note that ECS does introduce benefits to both end users and security researchers. There are numerous open DNS providers that offer extra features built on top of DNS such as content filtering, ad blocking, malware protection, and more [3], [5]. These features may be attractive to end users looking to protect their networks. However, prior to ECS, such users may have incurred a performance penalty for switching DNS providers due to CDN optimization based on proximity to a client's recursive DNS server rather than the client's actual location. Thus, ECS allows users greater choice of DNS providers without incurring a performance penalty. Additionally, ECS can allow security researchers to gain greater insights into potential infections above the recursive DNS server. This is particularly

valuable when monitoring DNS queries at the authoritative DNS server, which sees all non-cached queries for a zone. Additionally, ECS provides added benefit for organizations running DNS sinkholes. By providing client level visibility, DNS sinkholes can be used to estimate infected populations and provide different responses to clients in different networks. This extra level of visibility can be used to generate insights that were previously impossible before introduction of ECS.

The goal of this study is to understand the real-world adoption of ECS since its inception. It seeks to provide a discussion of the potential security benefits and pitfalls introduced by ECS. It accomplishes these goals by providing a longitudinal study of ECS deployment using DNS data collected from several DNS vantage points both before and after its official adoption in 2016. The outcomes of our investigation can be summarized in the following contributions:

- We measure the ECS adoption from the perspective of three different DNS authoritative name servers to show how the protocol has grown both before and after its official adoption in 2016. We show that, despite being an optional extension, ECS has seen steady adoption over the years with numerous DNS providers now supporting it.
- We show how ECS reveals more information about the clients making DNS requests and discuss the effects of this increased visibility. We discuss how it may provide more freedom to end users and aid security practitioners. At the same time, we discuss how it potentially exacerbates existing threats (i.e., DNS leaks).
- We examine the practical benefit provided by ECS to end users, and using a combination of Alexa and passive DNS data, we show that the vast majority of highly ranked ECS-enabled domains *do not* benefit from the use of ECS. Thus, most ECS-enabled domains appear to exacerbate existing privacy problems related to DNS without any benefit to the end user.

The following is a blueprint to help navigate our findings. We discuss the necessary background to understand both DNS and the changes introduced by ECS in Section II. This is followed by a description of the datasets and methodology used to perform our study in Section III. In Section IV-A, we discuss how the default configuration of ECS may introduce some unintended privacy consequences—which have been noted, discussed, and partially addressed in various iterations of the ECS proposal. Next, we show that ECS has seen steady adoption over time, despite being an optional extension, in Section IV-B. Then, in Section IV-C, we discuss how the client-level visibility provided by ECS to authoritative DNS networks is extremely granular, with most client prefixes corresponding to a /24 or smaller network CIDR. This effectively enables client level tracking using ECS. In Section IV-D, we study the practical impact of ECS on end users and show that most of the domains in the Alexa top million do not benefit from ECS. Finally, Section IV-E discusses how ECS leaks client level information to every AS on-path between a recursive DNS server and the authoritative DNS server—exacerbating existing DNS related privacy concerns. A summary of our findings and their impact on Internet communications can be found in Section VII.

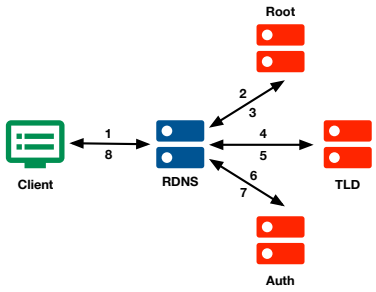

Fig. 1: *Illustration of the iterative name resolution process. In the diagram, the recursive is labeled as RDNS, and the authority is referred to as Auth.*

In summary, DNS is a fundamental protocol for communication on the Internet. This paper studies the deployment of a DNS extension called ECS, which introduces client information into communication above the recursive DNS server. The effects of this introduction extend beyond its original goal of making the Internet faster by helping optimize CDN selection at the DNS level. Through a longitudinal study using several DNS vantage points, we study both the positive and negative impacts of ECS on DNS communication.

II. Background

The Domain Name System (DNS) [35], [36] translates memorable names into IP addresses. Figure 1 shows the steps involved in resolving a domain name. In step (1), a stub resolver, located at the client, sends a request to a recursive DNS server, often simply referred to as the "recursive." If a cached answer is not available, the recursive iteratively queries other servers in the DNS hierarchy until it reaches the *authoritative* DNS server (referred to as authority throughout the paper) that can answer the current request, as seen in steps (2) to (7). Finally, in step (8), the recursive forwards the response from the authority to the stub resolver and caches the response for a period of time dictated by the response's time-to-live field.

The resolution process can be conceptually split into two parts. The first is the communication between the stub resolver (client) and the recursive, seen in steps (1) and (8), which is said to occur *below the recursive*. The second is the iterative resolution process shown in steps (2) to (7), which is said to occur *above the recursive*. Below, we discuss how the adoption of ECS enables entities above the recursive to acquire client-specific information that was not available to them before.

A. Evolution of DNS with ECS

EDNS Client Subnet (ECS) [15], which has been adopted by most large open recursives [19], [38], does not change the resolution process below the recursive but augments the information exchanged between recursives and authorities. Without ECS, only communication below the recursive (step 1) reveals the IP address of the clients. Thus, authorities receive no information about who is performing a query other than the IP address of the recursive. For example, as shown in Figure 2, only the IP address of the recursive DNS (RDNS) server is revealed to entities outside the local autonomous system (AS).

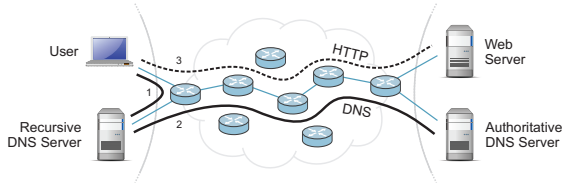

Fig. 2: *Legacy DNS network topology. Typically, recursion took place in the user's own autonomous system, and authorities were often situated in the same AS as the web server. Both DNS and HTTP traffic followed the same network path.*

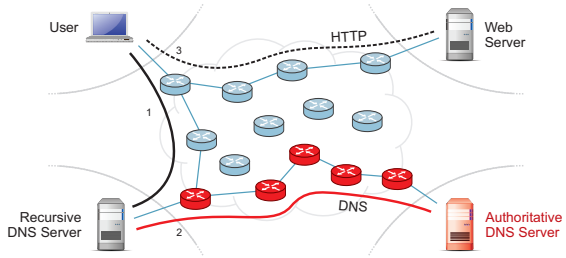

Fig. 3: *Modern DNS network topology. Increasingly, clients query a "cloud DNS" host or open resolver situated at a different autonomous system, and modern web sites frequently outsource DNS management to third-parties. Due to the inclusion of ECS information in DNS requests, a fraction of autonomous systems that would otherwise be unrelated to the path between the user and the actual web server are now in a position to gain client-specific information about browsing (or other) activity.*

What changes with the adoption of ECS is the information contained in the communications above the recursive, which is shown in steps (2) to (7) in figure 1. The steps are still the same; the main difference is that when the recursive resolver and the authoritative DNS servers support ECS, the DNS packet contains the extra information to help the authority identify the user's broad geographic location. This change has come about due to the changing landscape in how DNS resolutions are performed nowadays. The rise of open recursive DNS servers, which are typically situated in separated ASes than the users (as shown in Figure 3), disrupts the optimal delivery of content (e.g., as performed by CDNs), which previously assumed users were proximate to their resolvers. When a user resolves the name of a CDN-enabled web site, the authority DNS server would respond with a web server address close to the recursive, instead of the actual user. Thus, a North American user relying on a European DNS server could be directed to a non-local CDN mirror, slowing the resulting TCP connection.

ECS attempts to address this issue by including a truncated portion of the client's IP address, referred to as the *source netmask*, in all subsequent requests made by the recursive to an authority supporting ECS. An authority usually indicates that it supports ECS by including a scope netmask in reply to an ECS enabled query. On the other hand, some recursive resolvers send ECS enabled queries to all authorities. This added user information allows selecting a mirror that is in close proximity to the actual user, not just their cloud recursive. According to the ECS protocol [15], the source netmask should be determined using the most detailed network information available to the recursive, but by default, it includes the first three octets of a client's IP address. An authority may include in its response a *scope netmask* that can guide a recursive's future choice of source netmask. The inclusion of a scope netmask by the authority is one way to signal that the authority supports ECS. The scope netmask indicates the authority's desired source netmask length, which should correspond to the minimum length that will allow for an optimal answer with respect to network performance. The recursive resolver also uses the scope netmask to help with caching an answer; based on the documentation, the answer from an ECS query with a scope netmask indicates the scope under which the answer is valid, and the recursive can proceed with caching the answer for the clients under the specified netmask. The caching behavior and the potential issues that can arise from it are further discussed in [7]. Finally, the discovery process of ECS authorities by the resolvers varies but usually relies on the recursive resolver sending ECS enabled queries and observing if the authority responds with a scope netmask, in most cases. Some operators repeat the discovery process over time, while other recursives do not keep a list of authorities that support ECS and instead send ECS enabled queries to all the authorities.

B. Implications of ECS Misuse

Although one might think that ECS information does not introduce any additional privacy leakage, as the actual HTTP traffic will eventually reveal a user's IP address to the web server (and all entities along that path). This is not true on the modern web for two main reasons:

1) The recursive DNS server is often situated in a different AS than the user.
2) The authoritative DNS server is often situated in a different AS than the web site.

Consequently, when resolving a domain name, there is no guarantee (and should be no assumption) that the same organization will manage both the DNS server and the web server for example. ECS introduces new ways in which the added user information can be exposed to parties that would normally not have visibility in the traditional DNS resolution case. When using an ECS enabled cloud based recursive, the DNS resolution request could have to follow a different path between the recursive and the authority (red line Figure 3). For example, a user located in a European country using an open cloud-based DNS resolver could have their DNS packet information leave the confines of their ISP and traverse outside third party networks on the way to the DNS resolver. This is a case of below the recursive information leakage and applies to all DNS queries both before and after ECS adoption made to recursive resolvers outside the user's ISP network. On the other hand, after the cloud-based DNS recursive accepts the query if the recursive support ECS all subsequent hops from the recursive to ECS supporting authorities would include the ECS related IP prefix of the client. When, in this case, the authority is on a different path, signified by the red line in Figure 3, all the hops will receive the unencrypted ECS IP prefix of the client. This ties back to point (1) where even when the authoritative is in the same AS as the web server (e.g., as

shown in Figure 2), the network path from the user to the web server will be different than the path from the third-party open DNS resolver (e.g., Google's 8.8.8.8) to the web server. More importantly, regarding (2), due to the increasing reliance on third-party DNS hosting services (e.g., No-IP, EveryDNS, EasyDNS, Afraid, Zoneedit, Cloudflare), the path between the recursive and the authoritative may be completely different than the path between the user and the web server, as shown in Figure 3.

As shown in section IV, it appears there is significant misuse of ECS in the Internet. While, in many cases, this may pose no privacy concerns, in others, the users' anonymity may be seriously jeopardized. For example, Kintis et al. [27] discussed how this information could enable highly stealthy and targeted man-in-the-middle and surveillance attacks against dissidents, minorities, and even entire industry sectors.

In any case, it cannot be denied that users can indeed benefit from ECS, however, its correct and absolutely necessary deployment is paramount in order to minimize the possibility of privacy leakage.

III. Methodology

In this section, we study the adoption of the ECS protocol by recursives from three different vantage points and investigate the client information sharing due to ECS from the perspective of an authority (i.e., what *additional* client-related information ECS shares with authorities) and its applications. We first describe the datasets and provide statistics about the observed legacy and ECS enabled requests throughout our different collection sources. Then, we show that the ECS protocol is widely adopted across our sources and constitutes a significant percentage of the DNS traffic. Next, we demonstrate how the ECS enabled traffic can be utilized to provide a view of the clients behind the ECS enabled recursives that make the DNS queries to the authoritative servers through a sinkhole authority case study.

A. Datasets

Top Level Domain (TLD): This historical dataset consists of queries to popular Top-Level Domain (TLD) zones. The DNS queries for this dataset span one year, from July 2014 to July 2015, beginning just before the wide adoption of ECS. This source gives us a coherent view into the first years of the ECS adoption and shortly before the release of the RFC (RFC 7871 [15]) in 2016.

DNS Zones: This passive DNS dataset consists of DNS traffic to authoritative DNS servers for several popular zones. The data from this authority ranges from March 2017 to June 2019. It contains DNS traffic for 9.8 Million unique IPs. We are utilizing this dataset to get a coherent view of ECS adoption after the official release of the RFC.

Sinkhole: Our sinkhole passive DNS data consists of a total of 24 sinkholed domain names related to targeted attacks from Advanced Persistent Threats (APTs) and typosquatting [6], [44] and combosquatting [28]. When users visit these sinkholed domains, as a result of social engineering or a typographical error when typing the domain name in the browser, our domain names get resolved and we record the resolution process.

TABLE I: *The four types of passive DNS datasets that we utilize in our study. For the first three datasets, the dates span from July 2014 before the official adoption of ECS and then follows the evolution and growth of its deployment using popular DNS Zone authority data that we collected.*

Dataset Type	Size	Time Period
TLD Authority	141.9T	2014/07/01—2015/07/09
Zones Authority	50.9T	2017/03/10—2019/07/17
Sinkhole Authority	455.6G	2017/09/10—2019/06/20
ISP DNS dataset	4.2T	2019/04/01—2019/04/06

ISP DNS dataset: This dataset was collected by a large ISP (top 10) in North America over the first five days of April 2019. This ISP provides services over the entire North American region and provides us with real-world information about the state of ECS and its usage. We use this dataset in section IV-D to provide us with more insight into the benefits that ECS enabled domains obtain by adopting the protocol.

Alexa: Finally, we use the list of the most popular domains compiled by Amazon's Alexa. We use this dataset to help identify popular domains on the internet and examine their support for ECS and how they might benefit, or not, from supporting ECS.

Table I provides a detailed view of the first three aforementioned datasets, the size for each one, and the time period they cover. At this point, we should note that the TLD authority data is approximately 1.5 years older than the Zones and Sinkhole datasets. Even though this could seem inconsistent at first, our results will demonstrate the statistical significance of our measurements, even though time periods might not overlap. Moreover, obtaining contiguous datasets of such large volume and different time periods is particularly difficult. However, we chose to use all three datasets in our study to paint a clear and longitudinal picture of the different ECS uses and changes from the very early adoption days until recently. The fourth dataset, the ISP DNS dataset is used to provide our study insights into the beneficial aspects of the adoption of ECS such as it's utilization from CDN providers.

B. Identifying ECS in Our Datasets

Figure 4 shows the volume of ECS-enabled and legacy DNS resolution requests for the TLD and the DNS Zones authorities. We observe that in both authorities the vast majority of the DNS requests are non ECS-enabled. In the TLD authority dataset, which goes back to July 2014, we observe no ECS-enabled queries until mid-August 2014. We spot the first noticeable volume of ECS-enabled queries on August 20 2014, with a total of 2.6 million queries from 95 different recursives, all of which can be traced back to Google by using the Route Views [45] BGP announcement project database. ECS-enabled traffic constitutes about 0.2% of the daily TLD traffic in 2014-2015, featuring an increasing trend over time. After the adoption of the ECS RFC, we notice that the ECS enabled requests make up 30% of the daily DNS traffic with a mean of 295 million requests per day. This clearly indicates the large growth after ECS was adopted as an RFC.

Our sinkhole dataset contains observations between September 10, 2017, and June 20, 2019. On the first day of our experiment we had 11 domain names sinkholed. We kept

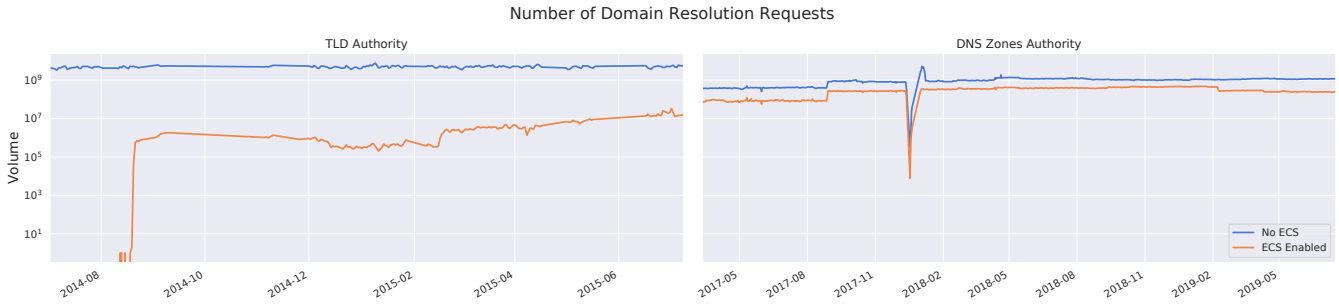

Fig. 4: *The number of daily legacy and ECS-enabled DNS requests to the authorities. The non ECS-enabled requests constitute the majority of the DNS requests. The dip in December 2017 in the DNS Zones authorities is a result of collection issues (missing data) on that period.*

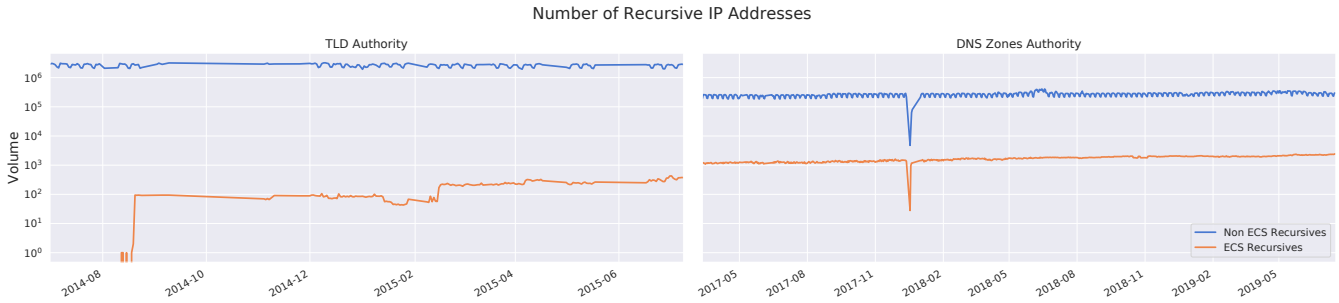

Fig. 5: *The number of different legacy and ECS-enabled recursives that resolved domain names in the authorities. The majority of the recursives do not utilize the ECS protocol while ECS traffic is emanating from a very small number of recursives. The dip in December 2017 is a result of collection issues (missing data) on that period for the DNS Zones authority.*

incorporating more and more domain names in our sinkhole, reaching a maximum of 24 domains. Figure 6 shows the volume of ECS-enabled and normal DNS resolution requests. Contrary to the previous two datasets, we observe that the ECS enabled requests constitute the majority of the traffic to our sinkhole authority while the daily query volume is unsurprisingly orders of magnitude smaller than that of the other two authorities. In total, we saw 11.5 billion DNS requests from which 69% were ECS enabled.

By looking at the IPs of the recursives making the DNS requests at the TLD and DNS Zones authority in Figure 5, we initially observe that the vast majority of the recursives that query the authorities do not use ECS. The ECS requests come from a small number of recursives that increased from less than 100 in the first months of ECS-enabled traffic in 2014 to more than 1,000 in 2017-2019. Figure 7 shows a similar trend for the recursives at the sinkhole authority. The ASes that host ECS-enabled recursives are predominately owned by Google in the tune of 1,579 (85%) , 2,444 (46%) and 500 (78%) of the recursives for the sinkhole, DNS Zones and the TLD authorities. Likewise, Google's recursives handle the majority of the ECS enabled traffic, as shown in Table II, which shows that in all of the authorities Google handles the vast majority of the ECS traffic.

IV. Measuring ECS in the Real World

In the previous section, we presented the state of ECS adoption as can be observed through passive measurements. In this section, we will present a measurement study of the real-world deployment of ECS among popular websites, with the goal of examining whether its use is justified. We will also present a study of sinkholed domains and how the use of ECS can provide us with information about the clients connecting to our sinkhole. To begin, we investigate the privacy preservation claims in the ECS RFC [15] with respect to the length of the source netmask (Section IV-A), and show that the prefixes suggested for ECS do not necessarily reflect the reality, in terms of routing on the Internet, where the vast majority (50%) of prefixes have the same /24 that ECS recommends. Second, in Section IV-B, we revisit prior work by measuring the deployment and distribution of ECS-enabled resolvers and identify steady growth in the adoption rate of ECS across both popular and less popular sites. Third, we examine various properties of the observed ECS speakers. We show that over the years, more and more domains have opted to support ECS, even though many ECS speakers do not appear to represent content delivery networks, the very technology ECS is meant to assist. In fact, the majority of ECS-enabled domains (80%) do not exhibit any kind of CDN behavior. (Section IV-D). Lastly, in Section IV-E, we show that not only ECS-enabled domains do not exhibit CDN behavior , but also utilize outsourced and managed DNS services by commercial providers. These services reside in different autonomous systems (AS), and anyone else on the DNS path is positioned to collect client-specific information (through the ECS client netmask) that would be otherwise unavailable to them if not for ECS.

A. Revisiting the Default ECS Configuration

In Section 11.1 of RFC 7871 for ECS [15], the authors discuss some privacy considerations due to the use of the

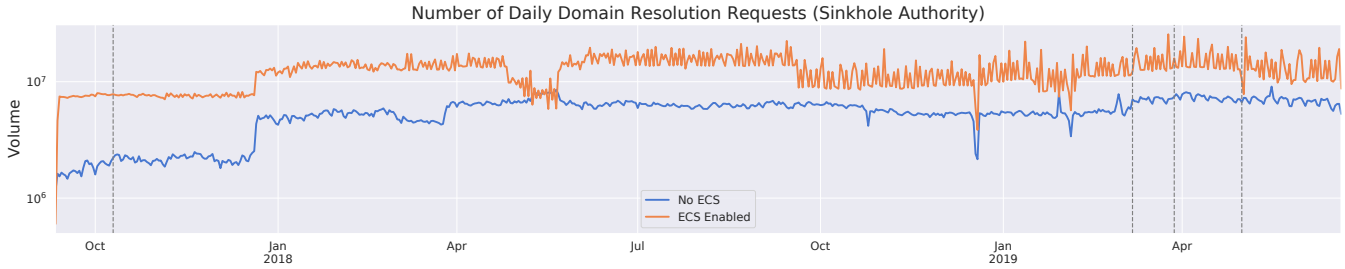

Fig. 6: *The number of daily legacy and ECS-enabled DNS requests to the sinkhole authority. The dashed lines represent the event of the addition of new domain names to the authority. Contrary to the global authority data, the ECS-enabled requests constitute the majority of the overall traffic.*

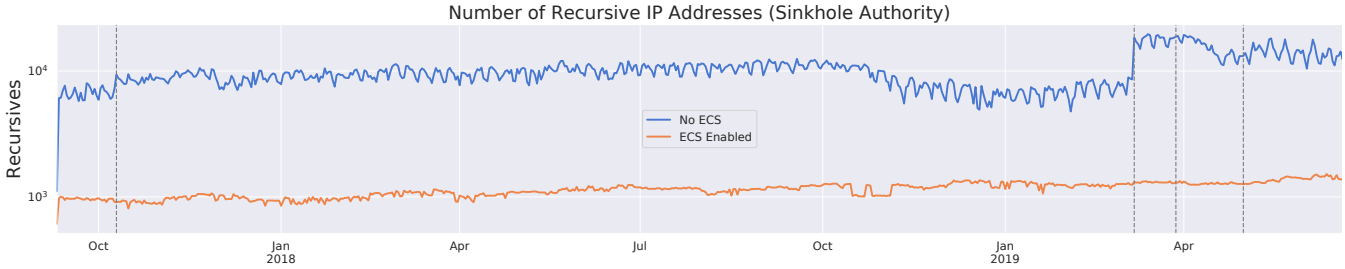

Fig. 7: *The number of different legacy and ECS-enabled recursives that resolved the sinkholed domains. The vertical dashed lines represent the addition of a new sinkhole domain name to the authority. A large number of legacy recursives have submitted resolution requests, whereas the ECS-enabled requests originate from a very small number of recursives.*

TABLE II: *Top 5 Autonomous Systems where the ECS enabled recursives reside. Clearly, the vast majority of the ECS-enabled requests to all of the authorities come from Google's recursives.*

TLD Authority		DNS Zones		Sinkhole Authority	
Recursive IP AS Owner	Queries	Recursive IP AS Owner	Queries	Recursive IP AS Owner	Queries
GOOGLE - Google LLC, US	743M	GOOGLE - Google LLC, US	212B	GOOGLE - Google LLC, US	8B
AS-APPRIVER - APPRIVER LLC, US	3,546	OVH, FR	19B	AS-CHOOPA - Choopa, LLC, US	155,307
COMCAST-7922 - Comcast Cable Communications, LL	2,649	OPENDNS - Cisco OpenDNS, LLC, US	11B	DYNDNS - Oracle Corporation, US	40,754
CHINANET-BACKBONE No.31,Jin-rong Street, CN	365	IPV6 Internet Ltda, BR	844M	CHINANET-IDC-GD China Telecom (Group), CN	38,270
DETEQUE - Deteque LLC, US	142	DYNDNS - Oracle Corporation, US	148M	SKYTEL-AS, GE	36,140

proposed extension. Their suggestion is for recursive servers to truncate the client's IP address into a source network mask, which typically contains the first 24 bits, in an effort to preserve privacy. The authors also suggest that entities responsible for operating ECS-enabled recursives should adjust the source netmask such that it reveals the least information possible about the client, while still providing beneficial information to ECS-enabled authorities, so that they can serve the client the proper IP. To the best of our knowledge, there is no study that demonstrates what portion, if any, of an IP address can be safely revealed while still preserving user privacy. Thus, it is unclear if the RFC's suggestion of using a source netmask is sufficient.

Before evaluating the privacy suggestions in the RFC, we first compared the assumed "/24 default" in the protocol to the general allocation and delegation practices found in IPv4. The RFC suggests that the /24 of the client's IP is sufficient to protect users' privacy while allowing better geolocation identification. However, that brings up the question of whether another choice, for example /16 or some other mask size, would perform as well while leaking less information.

To that end, we show the distribution of the announced prefixes on the Internet for two snapshots, one in April 2015 and one in June 2019. The dataset includes also the organization for each announced prefix. These figures showcase the changes that the Internet has undergone over these few years.

For the 2015 dataset, we downloaded the public delegation files from each of the five Regional Internet Registries (RIR) [25]. Since delegations are updated daily, we picked a single snapshot spanning April 6, 2015, to April 7, 2015, to limit experimental complexity and account for time zone variations across registries. The delegation files only map prefixes to autonomous systems. Therefore we use the Team Cymru IP to ASN tool [1] to associate the prefixes and their associated organizations. On April 9, 2015, there were 325,680 prefixes found for the entire IPv4 address space. While the vast majority of these were /24s, the shortest prefix found was a /8 and the longest was a /32. Figure 8 shows the distribution of these prefixes across the different subnet masks. Note that the vast majority (50%) of the prefixes are allocated as /24, and only 12,496 (4%) are less than or equal to a /16.

We examine the same statistics for June 2019, this time through the Route Views [45] dataset (that already has the ASN to organization mappings), in order to measure how this

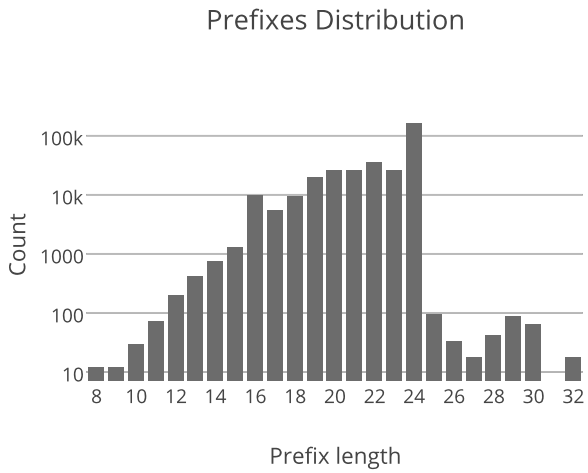

Fig. 8: *The distribution of prefixes (log scale) announced on the Internet for 2015 as reported by Team Cymru's IP to ASN Mapping service.*

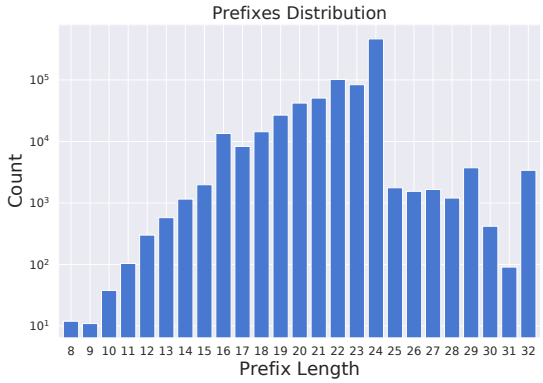

Fig. 9: *The distribution of prefixes (log scale) announced on the Internet as reported by Route Views in 2019.*

aspect of the internet has evolved. As indicated by Figure 9, the total number of prefixes currently is 732,182, more than double compared to four years ago. This is expected since the internet landscape has been constantly evolving over time with more and more companies opting to use it. However, the distribution of the prefixes has remained more or less the same, especially for the lower prefixes (up to /24). Approximately 396,045 (54%) of the announced prefixes, are /24 and 17,588 (2.5%) are less or equal to /16. The utilization of prefixes larger than /24 has also increased considerably, but that is not of immediate interest to our study since the RFC has determined that at most a /24 can be leaked without compromising client's privacy. Overall, by comparing the two figures side by side, it is evident that there is a trend of smaller portions of the IP space being delegated to the organizations, resulting in a higher percentage of longer network prefixes.

Considering that ECS reveals the IP address of the stub resolver and depending on an organization's network infrastructure, ECS might reveal information about the stub's behavior to third parties that might not be necessary for the optimal routing of information back to the client. This applies

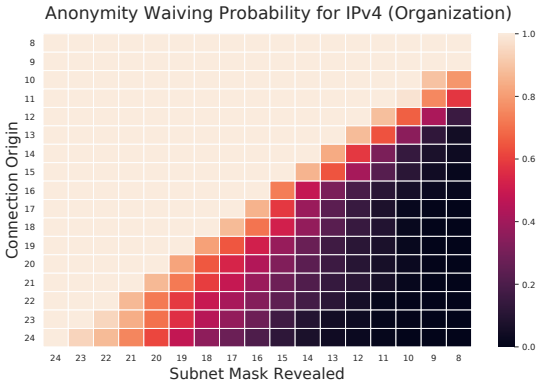

Fig. 10: *The probability that a client's organization can be precisely identified, given its actual network prefix (y-axis) and the revealed source network mask through ECS (x-axis).*

particularly to cases where the client inside a network uses a DNS provider outside the organization network with ECS enabled as the default. In such cases, ECS should ideally use only the prefix length that would provide with enough geographical information for optimal content delivery. With this in mind, we attempt to measure the extent to which this if feasible as we vary the length of the subnet mask. We expect that, by reducing the prefix length, the behavior of clients within an organization will no longer be able to be uniquely identified. This is also a way to theoretically test the ability to set custom ECS network masks as per the RFC. Organizations that want to avoid this type of information leak should augment their IT policies to make sure that any client operating on their network is not using cloud based DNS providers that might expose their organization's IP to outside parties. Considering the prevalence in the use of cloud based DNS resolvers and their use in devices like mobile phones, we believe that this part of the study can help shed some light into the potential issues that can arise in such a scenario.

We set up the measurement as follows: Using the Route Views dataset [45] of prefixes and corresponding organizations, we organized the prefixes as a radix tree, so that it is easy to collect the organizations that are covered by a given prefix. Then, we calculated the probability of a network prefix falling into a single organization for a given source network mask. We chose to use network prefixes instead of IP addresses, because they align more closely with organizational delegations from RIRs. We limit our computations to prefixes between "/8" and "/24", because ECS suggests that a "/24" prefix is sufficient. For each possible network prefix in the Route Views dataset, we sequentially reduced the length of the source network mask being revealed and measured the number of unique candidate organizations after each reduction. The probability computed for a given network prefix and source network mask equals the percentage of prefixes for which only a single organization was a possible match, i.e., the probability of uniquely identifying an organization for a given prefix and source network mask.

For example, assume a stub resolver connects from a given prefix X.Y.Z.W/24 (with subnet mask length 24). We reduce the length of the network mask to 22, and we get all the

organizations that are covered by X.Y.Z.W/22 with prefix lengths 22,23 and 24 (there is no point examining lengths more than 24 since we assumed that the origin subnet is a /24). If only one organization appears, then the target organization can be uniquely identified, even after reducing the mask length by 2 bits.

Figure 10 shows that changing the source network mask does not always increase the number of the candidate organizations a request may originate from. For example, a user connecting to the Internet through a /24 and revealing a /16 source network mask can be linked to a *single* organization about 20% of the time. The likelihood of uniquely identifying the organization jumps to 50% if a user connects from a /16 and reveals a /14 mask. Consequently, it is often quite easy to precisely identify the originating organization, despite changes in the source network mask introduced by ECS. Most organizations, though, have publicly documented network boundaries, so this is only a consideration for not documented organizations.

With these observations in mind, we can observe that the default subnet considerations made by the ECS RFC are proving to have wider implications in terms of identifying behavior from within networks, especially because so much of the internet infrastructure is based around network masks of /24. This is also the observation that the authors of the RFC had made early during draft making progress [13] that there should be a selectable mask length flag due to the potential privacy concerns of the initial proposal.

B. ECS Adoption Over the Years

In 2013, a group of researchers from the Technical University of Berlin measured ECS adoption among the Alexa top million websites [43]. In their findings, it is stated that approximately 13% of the Alexa top million domains were found to provide some support for ECS. With the goal of measuring how the adoption has evolved over time, we set up a custom resolver that complies with RFC 7871 and implements ECS. Then we performed similar experiments by sending ECS enabled requests to authorities and analyzing their responses. The pipeline is simple. First, we collect all the Alexa top million domains. Second, we query for the nameservers that are authoritative for these domains. As a final step, we pull the domain's A record from the corresponding authority using our ECS-enabled resolver with a random client subnet. The detailed results for two randomly selected days, one in April 2015 and one in June 2019, are presented next. These snapshots will reveal useful information regarding the ECS adoption.

In August 2015, there were 5,607 ECS-enabled authorities, which account for approximately 3% of all authorities (187,730), that serve domains in the Alexa top million. Due to network errors and misconfigured authorities, we were able to successfully measure 731,813 (73%) of these domains, out of which 161,302 were ECS-enabled and served by the previously identified ECS-enabled authorities. Almost 22% of the domains are ECS-enabled; this represents a 9% increase in ECS-enabled domains since 2013.

In the June 2019 case, approximately 92% of the total domains were successfully measured, and the number of ECS-enabled authorities are 19,133 (out of 173,905 authorities), a huge increase compared to 2015 data. This accounts for 11% of the unique authorities that serve the top domains in the Alexa dataset. However, the ECS adoption over the years has been strong, and 418,314 out of the measured 922,139 (45%) domains support ECS. This is a significant increase compared to 2015 as well.

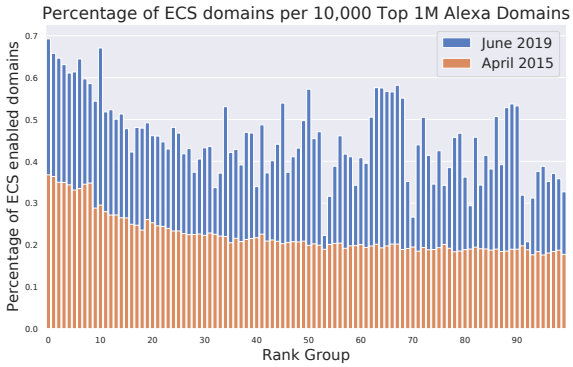

Fig. 11: *The percentage of ECS-enabled domains from the domains that responded, aggregated into buckets of 10,000 elements, for the Alexa top million web sites for 2019 and 2015. As expected, the most popular domain names are also ECS-enabled. In total, we identified 161,302 ECS-enabled domains in April 2015 and 418,314 in June 2019.*

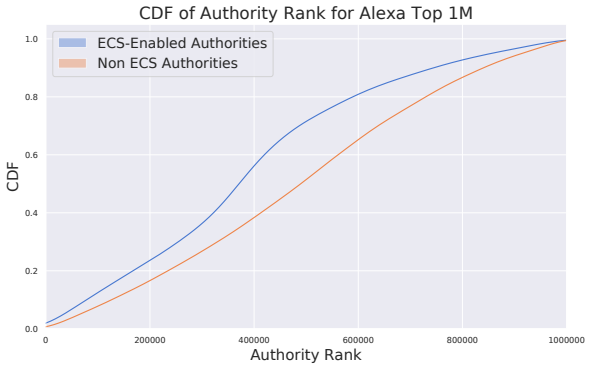

Fig. 12: *CDF of the authority rank for ECS and non-ECS enabled authorities. The authority rank is the average Alexa rank of the domains that this authority is authoritative for. The ECS-enabled domains are served by 19,133 authorities in June 2019, compared to 5,607 authorities in April 2015.*

Given that ECS aims to improve network performance, intuition suggests that popular sites are more likely to benefit from its use, and therefore, they should be more likely to use ECS-enabled domains. To test this hypothesis, we aggregated the ECS-enabled domains in the Alexa top million by their Alexa rank. Figures 11 and 12 present the distribution of ECS-enabled domains and authorities, respectively, according to their rank for our measurements in 2019 and 2015. The authority rank is defined as the average Alexa rank of all domains for which the given DNS server is authoritative.

By closely examining Figure 11, we see that there is indeed the trend that a larger percentage of the highly ranked domains tend to support ECS compared to the lower ranks, with some notable exceptions. Especially for 2015, this is very clear, even though the ECS adoption between ranks do not vary as greatly as 2019. The top ranks of the Alexa dataset traditionally do not change drastically over time, compared to the lower ranks. Additionally, these domains are associated with sites that are visited by millions of clients daily. Consequently, these sites often have multiple servers located all around the globe and may use CDNs to help improve network performance. Such sites represent the intended beneficiaries of ECS. Compared to 2015, it is apparent that many more domains support ECS today, even in the lower ranks. Websites are increasingly relying on CDNs to enhance their customer . This, combined with the introduction of new domains in the dataset that were not there in the past and serve different content, justifies the difference over the years.

The majority of authorities (Figure 12) in the 2019 dataset have ranks falling around the middle of the top million, while in the 2015 dataset the landscape is more balanced with ECS-enabled authorities having almost a linear distribution over the possible ranks. This is reasonable since the authority ranks are averages, suggesting that many authorities today are shared by domains spanning multiple ranks in the Alexa top million. The prevalence of shared hosting and DNS services likely explains much of this behavior.

This result highlights that there is a clear trend in ECS adoption over the years. Even though ECS is an optional standard designed to solve a particular issue, we have observed a steady increase in the adoption of this extension over time by a variety of websites without an apparent consideration as to whether the adoption of ECS would offer a performance improvement to the website.

C. Client IP Subnet Information

Using DNS logs generated from the authorities we discussed above, we were able to observe the geographic and network locality of clients. Since ECS was enabled on the authorities, recursives submitted ECS-enabled domain resolution requests, leaking the first three octets of the clients' IP addresses. Using this leaked client prefix, we were able to identify in greater detail the potential geographic location of the DNS requests than solely relying on the recursive IP. Furthermore, we could identify specific organizations and networks, many of them research institutes and government networks that were interested in the domain names in our sinkhole authority. This information was collected by solely operating an authority and would have not been available to us if it was not for ECS.

1) Client Geolocation and Network prefixes: In the absence of ECS, the visibility of the authorities would have been limited to the recursive IPs. However, when ECS is taken into consideration, we are able to observe a significantly better picture of the geolocation of the clients "behind" the recursives. Figure 13 shows the geographic distribution of the ECS enabled recursives and the clients that resolved domain names in each of the authorities. More specifically, we were able to identify 180, 231, and 204 more countries in the sinkhole, in TLD and the DNS Zones authorities respectively when we considered the geolocation of the client prefixes in the ECS enabled requests compared to only taking into account the location of the recursive IPs. The source of the DNS requests can be traced back in greater detail when ECS is enabled.

TABLE III: *Number of unique "/24" prefixes for the clients of ECS enabled requests and the recursives of legacy DNS requests for a random day in each dataset. We can see that in the TLD and the DNS Zones the ECS enabled traffic comes from more "/24s" than the traffic of the legacy DNS requests, even though the legacy DNS requests constitute the majority of the daily DNS requests.*

	TLD Authority	DNS Zones	Sinkhole Authority
ECS client subnets	660,073	771,052	1,319
Recursive client subnets	218,944	166,374	4,151

When looking into the origin network prefixes of the requests, we identified some noteworthy cases of networks that queried our domain names in the sinkhole Authority. Among the requests received by our sinkholed domains, the prefixes 180.94.82.0/24 and 180.94.94.0/24 (two networks in Afghanistan and delegated to "AFGHANTELECOM Government Communication Network.") appear to resolve two domain names related with APT activity in the past. Additionally, we also see frequent DNS queries to our APT domains from Academic networks, with 128.237.28.0/24 delegated to Carnegie Mellon University, 147.46.121.0/24 delegated to Seoul National University, and 171.67.70.0/24 delegated to Stanford University being some prominent examples. These DNS queries could be research related (e.g., by dynamically running malware that communicated with our sinkholed domains names) rather than infections. Finally, we also observe requests coming from Security vendors with 155.64.38.0/24 delegated to Symantec Corporation and 103.245.47.0/24 delegated to McAfee making requests to both our APT related domain names and typosquatting domains.

Prior to ECS, we would have only seen that "someone is using GoogleDNS" to resolve these domains. With ECS, we could identify specific networks engaged in research, security vendors, and governmental activities. We remind the reader that all this information comes from DNS alone and is off-path of any TCP analysis. However, it could be argued that even with legacy DNS requests authorities can have the same level of subnet visibility as clients that would not use ECS enabled open recursives such as Google's public DNS, would resort to a local recursive solution thus still revealing their subnet to the authorities. In order to test this argument, we randomly chose one day of DNS requests from each authority and looked at the number of ECS-client "/24" prefixes for ECS enabled requests and number of "/24" from legacy DNS requests after removing bogon IP prefixes from both datasets. Table III shows that we can see more client "/24s" in the ECS enabled traffic for the TLD and the DNS Zones authorities than in the legacy DNS traffic, even though as we have shown above that legacy DNS traffic makes up the majority of the daily DNS traffic for these authorities. While ECS only reveals up to 24 bits of the IP address of the clients, authorities can see a wider range of client subnets than legacy DNS. We can see that ECS can be used for extracting more granular information about the nature of the clients than in legacy DNS queries.

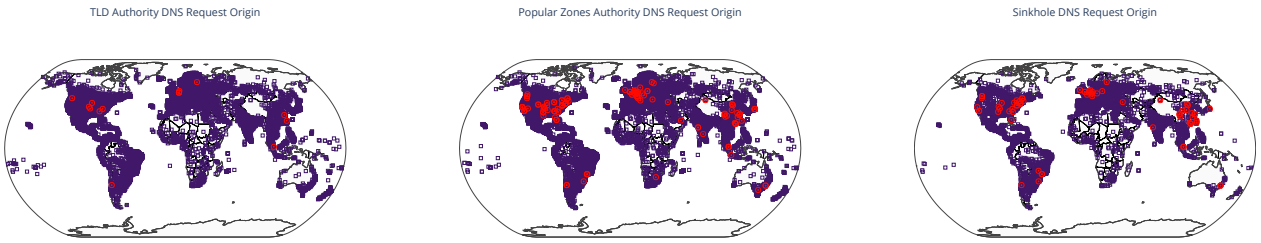

Fig. 13: *The distribution and density of the geographic location of the recursives and clients making ECS-enabled DNS requests to the authorities. The red dots show the location of the ECS recursives while the location of the clients behind the requests are in purple. We can see that by considering the geolocation of the client prefixes, which is only available in the ECS-enabled requests, an authority is getting a much more granular view of the source of the DNS requests. For the TLD and DNS Zones we calculate the distribution for a random day in June 2015 and June 2019, respectively, while in the sinkhole authority we use the full dataset.*

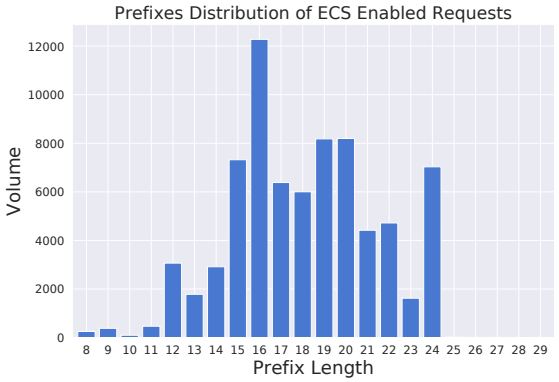

Fig. 14: *The distribution of the DNS resolution requests compared to the CIDR prefix length from where they originated. ECS could have provided the same level of service with the respective announced CIDR we see in the plot. Thus, the client's /24 was submitted with no value for the client.*

2) ECS Scope Size: Another interesting observation with respect to ECS-enabled resolution requests has to do with the size of the network the recursive is reporting to our authorities. We correlated the prefix reported by the recursives with the prefix containing the host addresses. From the eight billion ECS-enabled requests in our sinkhole authority, we could identify 75,138 unique prefixes for 99.8% of the requests. We were not able to identify 3,186 networks for 0.2% of the ECS-enabled DNS requests that correspond to IPv6 networks that are out of the scope of the current experiment and subnets that were not available on the Route Views [45] project database. From the available networks, 7,030 were "/24" delegations while 68,074 had a smaller prefix. We also note 34 cases of networks with a bigger prefix than a "/24". Figure 14, shows the distribution of the prefixes where IP addresses reported by ECS are delegated to.

Regarding the distribution of the client prefixes that recursives forward to our sinkhole authority, we see that the significant majority of the ECS client prefixes (99.8%) are "/24s" and we do not observe any prefix less specific than a "/24" exchanged. Although only the "/24" portion was forwarded for most of the clients, by considering figure 14, the origin of a request can be attributed to a *single* organization or ISP. On the other hand, we observe 130,261 queries in which the recursives respond with the full "/32" IPv4 address of the clients. By looking more closely at these queries, we see that the announced prefix for the corresponding clients is smaller or equal to a "/24". Thus there is no point in forwarding the full "/32" IPv4 address, and that the last octet from all the "/32s" forwarded was "1". The vast majority of the recursives (118 out of the 122) exhibiting this behavior were attributed to different organizations in China and did not forward any prefix smaller than a full IPv4 address. We observe queries for all of the domain names corresponding to our authority from 432 different clients. Considering that these recursives do not forward any prefix smaller than a "/32" for all the IPv4 client addresses that they serve and the fact that the last octet is "1" for all the clients served there is a high probability that these recursives are misconfigured.

To sum up, using passive DNS datasets, we showed that ECS enabled traffic makes up a considerable portion of the daily DNS traffic the past years. Contrary to legacy DNS requests, ECS enabled queries provide more granular client information to authorities. This can be a valuable tool for researchers using DNS lookup data (e.g., running a sinkhole as we have illustrated) in order to better understand the nature of the clients that are querying a domain name.

D. ECS Speakers and CDNs

We have seen that ECS has a high adoption rate among domains in the Alexa top million. Given that the goal of ECS is to improve CDN performance by enabling more accurate identification of a user's location [15], [39], this raises the question of whether these domains actually use ECS to facilitate content delivery. As CDNs rely on servers in multiple locations around the world, we expect resolutions of ECS-enabled domains from different vantage points to result in different IP addresses, exhibiting this way a consistent CDN behavior. In this subsection, we will show that only a few ECS-enabled domains appear to resolve to more than a single IP address. Consequently, there is no real performance benefit for the vast majority of domains that currently support ECS in the Alexa top million.

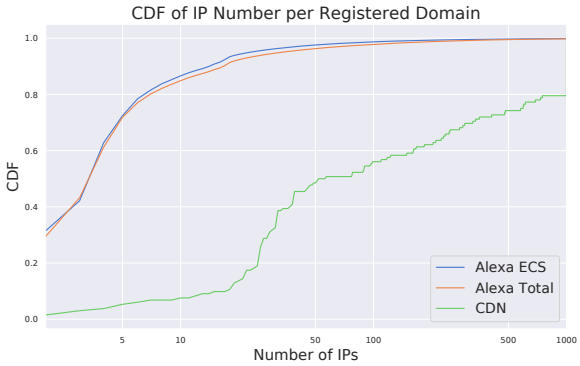

Fig. 15: *CDF of the number of IP addresses per domain name (log-scale) in the three datasets. The majority of CDNs have a much higher number of IP addresses, in contrast to ECS-enabled domains and the average Alexa domain.*

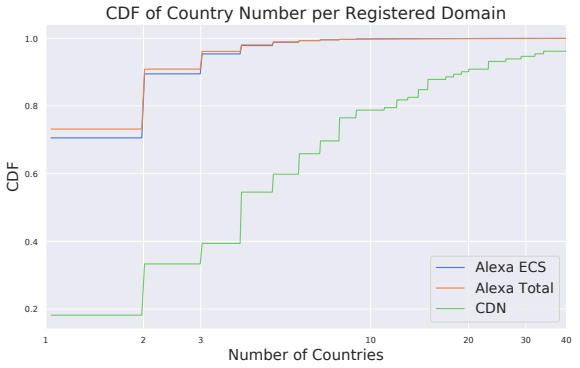

Fig. 16: *CDF of the number of distinct countries for IP addresses per domain name (log-scale) in the three datasets. CDN domains are distributed in multiple countries around the globe to better deliver their content, whereas ECS-enabled domains are mostly contained in the same country.*

To perform the experiments, we created a list of 133 verified CDNs by starting with a list of known, popular CDN domains and supplementing it with additional domains discovered in real-world network data. We used this set of CDN domains to make observations about the operation of the respective networks. Using the ISP DNS dataset collected by a large ISP in North America over the first five days of April 2019, we counted the number of IP addresses each domain (and CDN subdomains) resolves to, both in the CDN and the Alexa list. We observed that 50% of the verified CDN domains resolved to more than 50 distinct IP addresses in our passive DNS dataset. In sharp contrast, 80% of all ECS-enabled domain names from the Alexa top million resolved to less than seven distinct IP address, and less than 5% resolved to more than 50 IP addresses. Figure 15 shows the cumulative distribution function (CDF) of the number of IP addresses that a domain name resolves to for each of these data sets. Since the number of IP addresses for some domain names exceeded 400,000, we set an upper bound of 1000 IPs for each domain name for visualization purposes. Every domain with at least

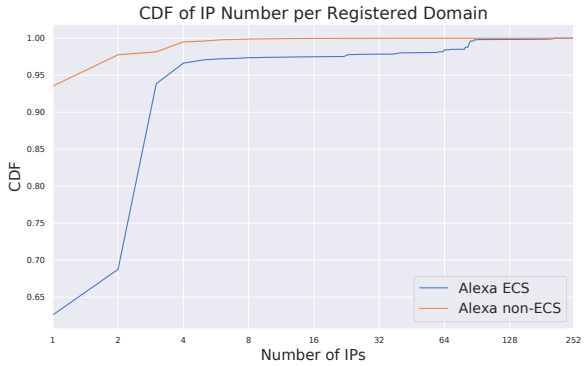

Fig. 17: *CDF of the number of IP addresses per domain name (log-scale) in our active querying experiment, notice y-axis starts at 0.625. The majority of Alexa domains have a very small number of IPs that they resolve to even when using ECS, in fact the majority over 62% only resolves in one IP, observing no benefit from the use of ECS.*

1000 associated IPs was aggregated into a single group so that the resulting plots are easier to read and understand.

We further correlate each IP address with its associated country of origin in order to provide a better understanding of the geographical diversity of the IP infrastructure that hosts each domain. For this purpose, we used the MAXMIND GeoIP2 [34] country database. For each domain, we counted the number of different countries it resolved to and presented the results in Figure 16. Taking geographic location into consideration, less than 20% of the known CDN domains and the vast majority (70%) of both the Alexa and ECS-enabled domains resolve to only a single country. Again for visualization purposes, we set the maximum number of distinct countries to be 40.

It is clear from Figures 16 and 15 that most of the Alexa domains do not share many of the CDN characteristics (the networks that ECS was originally proposed for), even though they eagerly support ECS. Also, behavior-wise, the generic Alexa domains and the ECS-enabled Alexa domains display very similar attributes (although not exactly the same).

In order to further examine the behavior of popular domains from Alexa that support ECS and to examine the utility of ECS for these domains, we conducted a further active experiment with the purpose of demonstrating the variety of RDATA when ECS is used. We take the entire Alexa 1M list of domains and submit queries using a modified resolver that allows us to specify the client prefix we will send to the authority. For a list of geographically diverse IP addresses to use as ECS prefixes, we used the publicly available AWS IP-ranges, which provides us with actual IPs that are geographically diverse as is the AWS infrastructure VIII-A. We repeat the experiment for 26 different IP ranges and present them in figures 17 and 18.

It is very clear from Figure 18 that the overwhelming majority (over 98%) of Alexa domains are hosted in only one country. This means that there are small geographical benefits from the use of ECS, even when we query the domains from client subnets that correspond to locations all around

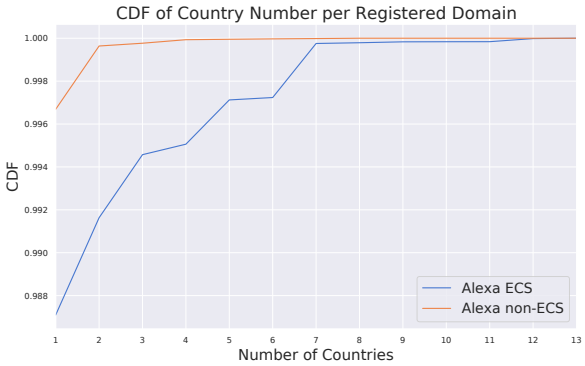

Fig. 18: *CDF of the number of distinct countries for IP addresses per domain name (log-scale) in our active querying experiment, notice y-axis starts at 0.98. In terms of variability on the country that's hosting the domain, Alexa domains exhibit even less variability and are in line with our passive measurements.*

Fig. 19: *Scatterplot of the Autonomous System Number (ASN) where the authority's IP address is being announced from and the ASN where the RDATA for a domain name resides into. The diagonal corresponds to authority-domain pairs that reside in the same Autonomous System.*

the world. Similarly Figure 17 verifies that our passive DNS measurements are consistent with active global measurements and shows that the hosting infrastructure of popular Alexa domains is not particularly diverse, especially compared with the CDN diversity we observe at Figure 15. The vast majority of Alexa domains, even those that support ECS, only utilize one IP address, and 95% of them use less than four IP addresses.

Based on the observed behavior, it appears that the benefit from the use of ECS is not significant (or apparent) for a large number of these domains. This reinforces our intuition that ECS is sometimes misused. It is also apparent that even for the case of the limited number of Alexa domains that point to many IP addresses, these IPs are not necessarily located in different places around the world. On the contrary, most of the them can be found in the same region. In these cases, the users' anonymity could be waived without any benefit for them. Given that, in the following subsection, we will measure the diversity of the infrastructure of these domains to understand how users' information travel during a DNS resolution request and in which cases other entities can obtain this information.

E. Infrastructure Diversity

To present how different entities are involved in the resolution of a domain, we analyze the infrastructure that hosts ECS-enabled domain names. Since routing on the Internet is based on Autonomous Systems (AS), BGP announcements, and peering agreements between ASes, we focus on the distance of the ASes that host ECS-enabled authorities from the ASes that host the respective services for those domains (i.e., RDATA). Ideally, we would want to know the different hops a packet will make before reaching the authority and the respective service. However, network packets are expected to take several different paths, depending on factors like peering agreements, congestion, load balancing, etc. which make it particularly hard to predict [17], [18], [33], [37].

In order to demonstrate that the DNS packets traveling to an authority are likely to take a different routing path from consecutive communication with the actual service the

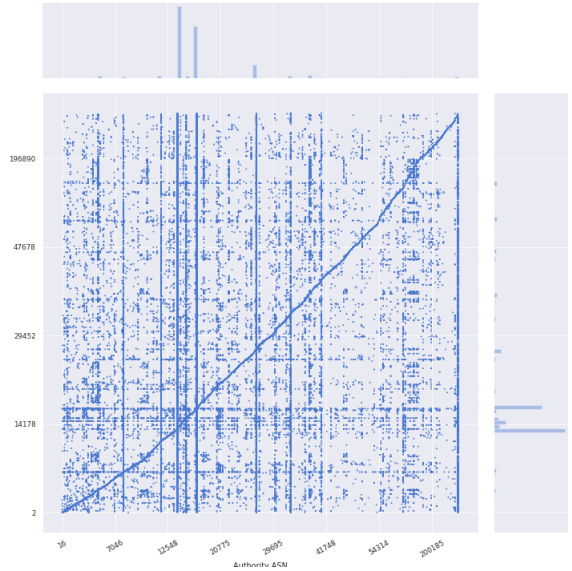

Fig. 20: *A different visualization of Figure 19 showing the joint distribution and collapsing the empty space. This distorts the diagonal because different ASNs are present in each axis. The diagonal is now a crooked line.*

domain offers (e.g., web server that serves HTML context), and therefore reveal information about the client to multiple other entities, we base our analysis on the ASes that host the authority and the returned IP address for an ECS-enabled domain. We also show that there are entities positioned on the path between a global recursive and the authority of multiple domain names, which are in a position to collect all clients' information just from DNS resolution requests.

For a given ECS-enabled domain name, if both the authority and the respective RDATA [35], [36] — referred to hereon as the *service* — are hosted within the same AS, then there is a probability that DNS leaks will be limited to the same path as the TCP connection that will follow. Currently,

however, it is often the case that the DNS packets will take a completely different routing path than the subsequent service connection (as shown in Figure 2). In that case, the ECS subnet information is leaked to all ASes between the resolver and the authority.

When a given ECS-enabled domain is served by an authority in a different AS than the service, then inevitably the DNS packet will take a different route and pass through at least one different AS (that of the authority). Thus, in the best case, one more AS will have information about the client (in reality, more ASes are likely to exist on the path). Figure 19 shows the relation between the AS of the authority (x-axis) and the AS of the service (y-axis). The diagonal on the plot depicts cases where the authority and the service are located within the same AS. To generate Figure 19 and Figure 20, we used the 2019 Alexa data from subsection IV-B. We associated each authority and domain IP with the ASN according to the Route Views dataset [45].

One thing that stands out from Figure 19 is that this kind of visualization is problematic because there is a lot of empty space (i.e., ASN numbers that are not used either in the authority or the RDATA axis). For that reason, we also created Figure 20, which is practically Figure 19 without the empty space (collapsing the plot to only include valid data points). Since each of the axes has different ASes (because some of the authority ASN may not have complete overlap with the RDATA ASN and vice versa), the actual diagonal is comprised of different ASNs per axis. However, if both the ASN for RDATA and authority are the same, it will be a dot near the diagonal. The ideal diagonal now looks like a crooked line but still stands out.

For reference, we examine some of the top 10 authority ASNs (that are related to the most RDATA ASN). Namely, AS 13335 belongs to CloudFlare, AS 26496 to GoDaddy, AS 16509 to Amazon, and AS 396576 to VeriSign. Obviously, these organizations are affiliated with CDNs, cloud services and domain name registrations (and thus parking) and that behavior is expected. In any case, every point in the plot apart from the diagonal in Figure 19 and the "crooked" diagonal in Figure 20 corresponds to cases where the RDATA (web server) and the authority (DNS Server) ASN are separate. Therefore there is potential information leakage to a different entity. There is no arguing that in the above cases, this occurs predominantly. The outsourcing phenomenon is a characteristic of the modern web.

Finally, to estimate the potential of a leak when a DNS resolution request arrives at any of the authorities, we measured the number of ASes that the authority's AS peers with. We use the Shadowserver [42] API to identify peers for the CIDRs that announce the IP addresses of the authorities for ECS-enabled domains. Figure 21 shows the distribution of peers for each domain name. The majority of the domains are served by authorities that are located in ASes with three, four, or eight peers. Any of those peers, along with other ASes until a packet reaches them, is a potential collector of activity from ECS-enabled DNS packets.

Essentially, we find that a large number of the domain names that utilize ECS use third-party DNS providers. This means that the DNS infrastructure of these domains resides in a separate network with a different AS and administrator. Thus, the IP information included in the new ECS-enabled DNS packets is shared with third parties unknown to the client for no immediately discernible reason. Considering the lack of a diverse hosting infrastructure for these domains, there is no benefit from enabling ECS. Similarly, ECS enabled domains provide IP information to third parties on-path during the resolution process. This partial information (e.g., a /24) would otherwise be unavailable to anyone other than the recursive itself.

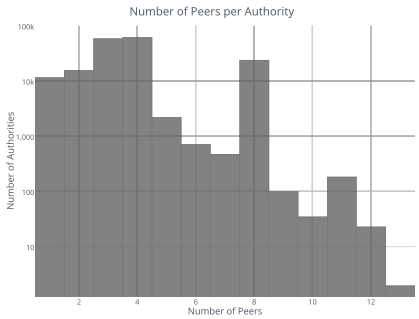

Fig. 21: *The distribution of the number of peers per Autonomous System that hosts an ECS-enabled authority. The vast majority of the authorities reside in ASs that have three, four, or eight peers, which can be potential alternative paths for a DNS resolution request and one more collection point for entities involved.*

V. Related Work

The interaction of DNS and anonymity networks has been well studied. Krishnan et al. [29] have shown how DNS prefetching can leak information regarding users' activity online to the degree that information regarding web searches can be inferred by simply logging a browser's resolution requests. Zhao et al. [46] perform a deep analysis on each step of a domain name resolution process, showing information that can be inferred from users' private data by only looking at public data. They also propose a simple range query scheme that can be used to protect the user. In the same context, Guha and Francis [22] describe an attack against the DNS, by passively monitoring DNS related traffic, that can provide a variety of information about a user that includes location, habits, and commute patterns. Moreover, Bortzmeyer, in RFC 7626 [11] attempts to enumerate the attacks and privacy implications, aggregated into six different categories, made possible only using DNS; they concluded their work with several security considerations on the matter. Lastly, Bortzmeyer also describes potential privacy issues and attacks via monitoring DNS traffic and examining the domain names included in packets, which can be solved by implementing RFC 7816 [12].

On the other hand, ECS is a relatively new technology and is motivated by the performance challenges related to the growing use of public recursives [16], as discussed by Huang et al. [26]. The guidelines in the corresponding RFC [15] provide a general outline on the how ECS should be deployed and how ECS-enabled servers should be operated. Streibelt et al. [43] demonstrate how one could utilize ECS-enabled authorities to uncover details about the infrastructure of an ECS-enabled

zone and how it is being used by the owner. Recently, Al-Dalky et al. [7] study a more specific aspect of ECS that has to do with the caching behavior of DNS resolvers when it comes to ECS enabled answers and the variety of different caching behaviors that can be examined. Our work focuses on a long-term study of the behavior and adoption of ECS. Lastly, Otto et al. [40] have measured how the adoption of ECS can increase the accuracy, with which authorities can identify a client's geographic location and provide better content delivery.

VI. Discussion

Considering that currently, ECS is enabled by default depending on the recursive used, the user has limited ability to control the amount of information shared using ECS, and so we would like to discuss the options available to the users. The RFC mentions that the user can signal the maximum resolution of the scope netmask that can be used by setting it in the initial request to the recursive, and the recursive should follow the resolution that the user's stub set. By setting a scope netmask of /0 the user can effectively opt out of using ECS while also not taking advantage of the benefits that ECS provides. Another option that the user has is to set a netmask more coarse than the default used /24 resolution. That will balance privacy and allow for more content delivery optimization by services that benefit from ECS. The issue with this approach is that no user-facing stub resolver currently allows for this setting level. Support for ECS scope netmask setting needs to be added to stub resolvers. Another potential issue is that currently, not all the recursives implement the RFC correctly but default to a different netmask, disobeying the netmask set by the user, similar to cases mentioned in Section IV-C2.

From the side of website operators, we can only comment that they should only enable ECS responses when they perform some form of traffic optimization. Considering the number of domains that seem to support ECS but do not benefit from the protocol, we believe that a large number of managed domain hosting enables ECS by default. Another interesting approach is the discussion around more privacy-minded ECS implementation that was presented in a publication by NextDNS's Olivier Poitrey [41]. This approach relies on the geographical awareness of the Autonomous Systems that the recursive resolver serves and depends on providing a geographically relevant IP portion to the authority instead of the user's IP address portion. As for the more privacy-conscious user, the standard privacy-preserving methods of browsing the web such as VPNs and the Tor network will still provide the user the ability to hide their IP from an ECS enabled authority. A more straightforward solution would be to manually set the DNS servers that the user prefers and thus choose a set of recursive resolvers that do not send ECS information. On the other hand, with any of these solutions, the user will not be able to take advantage of ECS's benefits.

VII. Conclusion

In this paper, we presented a longitudinal study measuring the adoption of a DNS extension called ECS. Given the widespread usage of DNS in IP based networks, the goal of our work was to identify how changes introduced by this extension affect network communications that rely on DNS. This analysis serves as a case study that explores the unintended consequences, both good and bad, of introducing small changes to fundamental network protocols.

The primary goal of ECS was to optimize CDN selection through the use of DNS, but our analysis found that most sites in the Alexa top million do not receive any benefit from ECS (Section IV-D). This result demonstrates how new functionality may not always get used as intended, and therefore, it is essential to consider potential unintended consequences. For example, we identified that most authoritative DNS servers using ECS adhere to the proposed defaults and set an IP subnet mask of /24 (Section IV-A). The use of small subnet masks results in the sharing of fine-grained client information with DNS nameservers above the recursive DNS server. We found that the majority of ECS-enabled domain names outsource their DNS infrastructure (Section IV-E). As a result, more networks now have fine-grained client information for DNS on-path DNS communication. Thus, we find ECS potentially exacerbates the effects of existing threats such as DNS leaks.

These consequences raise questions about the scope of impact. Our analysis finds that, despite being optional, ECS has seen steady adoption over time (Section IV-B). Thus, the unintended consequences of ECS are not limited to a small subset of Internet communications. As a result, authoritative DNS servers—and all DNS nameservers above the recursive for that matter—now have visibility about the client networks querying them. This client information enables DNS operators to track client networks and user behaviors in ways that were not possible before ECS (Section IV-C). At the same time, this same information can also help security practitioners track new threats or aid remediation efforts when local network visibility is limited.

Ultimately, we find that ECS has impacted a large volume of DNS traffic on the Internet. It is widely deployed and used by domains all across the Alexa top million. As a result, security practitioners should be aware of its pitfalls and potential uses for good.

References

[1] "IP to ASN Mapping - Team Cymru," http://www.team-cymru.org/IP-ASN-mapping.html, 2016.

[2] "A free, global DNS resolution service that you can use as an alternative to your current DNS provider." https://developers.google.com/speed/public-dns, 2020.

[3] "IMPROVE YOUR INTERNET," https://www.opendns.com/, 2020.

[4] "Internet Security & Privacy In a Few Easy Steps," https://www.quad9.net/, 2020.

[5] "NextDNS," https://nextdns.io/, 2020.

[6] P. Agten, W. Joosen, F. Piessens, and N. Nikiforakis, "Seven months' worth of mistakes: A longitudinal study of typosquatting abuse." in *NDSS*, 2015, pp. 156–168.

[7] R. Al-Dalky, M. Rabinovich, and K. Schomp, "A look at the ecs behavior of dns resolvers," in *Proceedings of the Internet Measurement Conference*, ser. IMC '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 116–129. [Online]. Available: https://doi.org/10.1145/3355369.3355586

[8] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, "Understanding the mirai botnet," in *26th USENIX Security Symposium (USENIX Security 17)*. Vancouver, BC:

USENIX Association, aug 2017, pp. 1093–1110. [Online]. Available: https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/antonakakis

[9] M. Antonakakis, R. Perdisci, D. Dagon, W. Lee, and N. Feamster, "Building a dynamic reputation system for dns." in *USENIX Security Symposium (SECURITY)*, 2011.

[10] M. Antonakakis, R. Perdisci, Y. Nadji, N. V. II, S. Abu-Nimeh, W. Lee, and D. Dagon, "From throw-away traffic to bots - detecting the rise of dga-based malware," in *USENIX Security Symposium (SECURITY)*, 2012.

[11] S. Bortzmeyer, "DNS Privacy Considerations," RFC 7626, Internet Engineering Task Force, 2015. [Online]. Available: https://tools.ietf.org/html/rfc7626

[12] ——, "DNS Query Name Minimisation to Improve Privacy," RFC 7816, Internet Engineering Task Force, 2016. [Online]. Available: https://tools.ietf.org/html/rfc7816

[13] C. Contavalli, W. V. D. Gaast, D. Lawrence, and W. Kumari, "Client Subnet in DNS Requests," RFC 7871, Internet Engineering Task Force, May 2015. [Online]. Available: https://tools.ietf.org/html/draft-ietf-dnsop-edns-client-subnet-00

[14] ——, "Client Subnet in DNS Requests (draft-ietf-dnsop-edns-client-subnet-00)," 2015. [Online]. Available: https://datatracker.ietf.org/doc/draft-ietf-dnsop-edns-client-subnet/00/

[15] ——, "Client Subnet in DNS Queries," RFC 7871, Internet Engineering Task Force, May 2016. [Online]. Available: https://datatracker.ietf.org/doc/rfc7871/

[16] D. Dagon, N. Provos, C. P. Lee, and W. Lee, "Corrupted DNS Resolution Paths: The Rise of a Malicious Resolution Authority." in *NDSS*, 2008.

[17] Feamster, Nick and Rexford, Jennifer, "Network-wide prediction of BGP routes," *IEEE/ACM Transactions on Networking (TON)*, vol. 15, no. 2, pp. 253–266, 2007.

[18] Feamster, Nick and Winick, Jared and Rexford, Jennifer, "A model of BGP routing for network engineering," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 32, no. 1. ACM, 2004, pp. 331–342.

[19] Google, "Introduction to Google Public DNS," https://developers.google.com/speed/public-dns/docs/intro, accessed: 2015-04-07. [Online]. Available: https://developers.google.com/speed/public-dns/docs/intro

[20] B. Greschbach, T. Pulls, L. M. Roberts, P. Winter, and N. Feamster, "The effect of dns on tor's anonymity," *arXiv preprint arXiv:1609.08187*, 2016.

[21] G. Gu, R. Perdisci, J. Zhang, and W. Lee, "Botminer: Clustering analysis of network traffic for protocol and structure-independent botnet detection," in *USENIX Security Symposium (SECURITY)*, 2008.

[22] Guha, Saikat and Francis, Paul, "Identity trail: Covert surveillance using DNS," in *Privacy Enhancing Technologies*. Springer, 2007, pp. 153–166.

[23] S. Hao, M. Thomas, V. Paxson, N. Feamster, C. Kreibich, C. Grier, and S. Hollenbeck, "Understanding the domain registration behavior of spammers," in *Proceedings of the 2013 Conference on Internet Measurement Conference*, ser. IMC '13. New York, NY, USA: Association for Computing Machinery, 2013, pp. 63–76. [Online]. Available: https://doi.org/10.1145/2504730.2504753

[24] T. Holz, C. Gorecki, K. Rieck, and F. C. Freiling, "Measuring and detecting fast-flux service networks." in *NDSS*, 2008.

[25] R. Housley, J. Curran, G. Huston, and D. Conrad, "The Internet Numbers Registry System," RFC 7020 (Informational), Internet Engineering Task Force, Aug. 2013. [Online]. Available: http://www.ietf.org/rfc/rfc7020.txt

[26] C. Huang, D. A. Maltz, J. Li, and A. Greenberg, "Public DNS System and Global Traffic Management," in *INFOCOM, 2011 Proceedings IEEE*. IEEE, 2011, pp. 2615–2623.

[27] P. Kintis, Y. Nadji, D. Dagon, and M. Antonakakis, "Understanding the privacy implications of ecs," in *Detection of Intrusions and Malware, and Vulnerability Assessment: 13th International Conference, DIMVA 2016, San Sebastián, Spain, July 7-8, 2016, Proceedings*, vol. 9721. Springer, 2016, p. 343.

[28] Kintis, Panagiotis and Miramirkhani, Najmeh and Lever, Charles and Chen, Yizheng and Romero-Gómez, Rosa and Pitropakis, Nikolaos and Nikiforakis, Nick and Antonakakis, Manos, "Hiding in plain sight: A longitudinal study of combosquatting abuse," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 569–586.

[29] S. Krishnan and F. Monrose, "DNS prefetching and its privacy implications: When good things go bad," in *Proceedings of the 3rd USENIX conference on Large-scale exploits and emergent threats: botnets, spyware, worms, and more*. USENIX Association, 2010, pp. 10–10.

[30] M. Leech, M. Ganis, Y. Lee, R. Kuris, D. Koblas, and L. Jones, "SOCKS Protocol Version 5," RFC 1928 (Proposed Standard), Internet Engineering Task Force, Mar. 1996. [Online]. Available: http://www.ietf.org/rfc/rfc1928.txt

[31] C. Lever, R. Walls, Y. Nadji, D. Dagon, P. McDaniel, and M. Antonakakis, "Domain-z: 28 registrations later," in *IEEE Symposium on Security and Privacy (Oakland)*, 2016.

[32] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkitasubramaniam, "l-diversity: Privacy beyond k-anonymity," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 1, no. 1, p. 3, 2007.

[33] Madhyastha, Harsha V and Anderson, Thomas and Krishnamurthy, Arvind and Spring, Neil and Venkataramani, Arun, "A structural approach to latency prediction," in *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*. ACM, 2006, pp. 99–104.

[34] MAXMIND, "GeoIP2: Industry Leading IP Intelligence," 2015. [Online]. Available: https://www.maxmind.com/en/geoip2-services-and-databases

[35] P. Mockapetris, "Domain names - concepts and facilities," RFC 1034 (INTERNET STANDARD), Internet Engineering Task Force, Nov. 1987, updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592, 5936. [Online]. Available: http://www.ietf.org/rfc/rfc1034.txt

[36] ——, "Domain names - implementation and specification," RFC 1035 (INTERNET STANDARD), Internet Engineering Task Force, Nov. 1987. [Online]. Available: http://www.ietf.org/rfc/rfc1035.txt

[37] Mühlbauer, Wolfgang and Feldmann, Anja and Maennel, Olaf and Roughan, Matthew and Uhlig, Steve, "Building an AS-topology model that captures route diversity," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4, pp. 195–206, 2006.

[38] OpenDNS, "The OpenDNS Global Network Delivers a Secure Connection Every Time. Everywhere." http://info.opendns.com/rs/opendns/images/TD-Umbrella-Delivery-Platform.pdf, 2010. [Online]. Available: http://info.opendns.com/rs/opendns/images/TD-Umbrella-Delivery-Platform.pdf

[39] ——, "A Faster Internet: http://www.afasterinternet.com," 2011. [Online]. Available: http://www.afasterinternet.com

[40] J. S. Otto, M. A. Sánchez, J. P. Rula, and F. 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*. ACM, 2012, pp. 523–536.

[41] O. Poitrey, "How we made DNS both fast and private with ECS," https://medium.com/nextdns/how-we-made-dns-both-fast-and-private-with-ecs-4970d70401e5, 2019. [Online]. Available: https://medium.com/nextdns/how-we-made-dns-both-fast-and-private-with-ecs-4970d70401e5

[42] Shadowserver, "Shadowserver Foundation," https://www.shadowserver.org/, 2016. [Online]. Available: https://www.shadowserver.org/

[43] F. Streibelt, J. Böttger, N. Chatzis, G. Smaragdakis, and A. Feldmann, "Exploring EDNS-Client-Subnet Adopters in your Free Time," in *Proceedings of the 2013 conference on Internet Measurement Conference*. ACM, 2013, pp. 305–312.

[44] J. Szurdi, B. Kocso, G. Cseh, J. Spring, M. Felegyhazi, and C. Kanich, "The long "taile" of typosquatting domain names," in *23rd {USENIX} Security Symposium ({USENIX} Security 14)*, 2014, pp. 191–206.

[45] R. Views, "University of oregon route views project," 2000.

[46] F. Zhao, Y. Hori, and K. Sakurai, "Analysis of Privacy Disclosure

in DNS Query," in *Multimedia and Ubiquitous Engineering, 2007. MUE'07. International Conference on*. IEEE, 2007, pp. 952–957.

VIII. Appendix

A. Active probing subnets

CIDR	Country	Region
150.222.81.0/24	Irland	West Europe
64.252.84.0/24	United Kingdom	West Europe
52.95.224.0/24	Spain	South Europe
52.94.18.0/24	Spain	South Europe
52.219.168.0/24	Germany	Central Europe
64.252.88.0/24	Germany	Central Europe
13.248.100.0/24	Sweden	North Europe
15.177.72.0/24	Sweden	North Europe
150.222.78.0/24	Singapore	Southeast Asia
64.252.104.0/24	Singapore	Southeast Asia
13.248.117.0/24	India	South Asia
150.222.235.0/24	India	South Asia
52.95.226.0/24	Hong Kong	East Asia
54.240.241.0/24	Hong Kong	East Asia
15.221.34.0/24	Japan	Northeast Asia
150.222.116.0/24	South Korea	Northeast Asia
15.230.137.0/24	United States	North America East
13.248.103.0/24	United States	North America East
99.77.132.0/24	United States	North America West
52.95.247.0/24	United States	North America West
15.230.138.0/24	South Africa	South Africa
52.95.180.0/24	South Africa	South Africa
99.77.147.0/24	Bahrain	Middle East
13.248.106.0/24	Bahrain	Middle East
64.252.78.0/24	Brazil	South America
150.222.12.0/24	Brazil	South America

TABLE IV: *CIDRs and their respective countries and regions selected for the active probing of the Alexa 1M domains for ECS optimized responses. The CIDRS are networks belonging to Amazon AWS based on publicly available data. The countries are geolocation of the CIDRs based on Amazon's published network information, available at: https://docs.aws.amazon.com/general/latest/gr/aws-ip-ranges.html*