A Cloud Provider's View of EDNS Client-Subnet Adoption

Matt Calder Xun Fan Liang Zhu
Microsoft

***Abstract*—Directing users to a nearby, high-performing front-end is core to the business of content delivery networks (CDNs). CDNs which use DNS to direct users to servers face the challenge of making decisions based at the LDNS-level, not based on the client's IP address, and, in many cases, an LDNS is not representative of the clients it serves. The EDNS Client Subnet specification provides a solution by embedding a portion of the client's IP address in the DNS query to help CDNs make better redirection decisions, but both the LDNS and authoritative resolver (CDN side) must support the standard. While there has been well-publicized adoption of Client Subnet by authoritative CDN resolvers, adoption rates across LDNSes are unknown.**

In this work, we examine Client Subnet adoption in LDNSes. We analyze DNS queries captured over one month from Microsoft's Azure Cloud platform. We find that adoption on the Internet is very low across ISPs but query volume is relatively high due to the popularity of public DNS services. We discover high network adoption rates in China and reveal that Chinese public DNS services deploy LDNSes deep into end-user networks.

I. Introduction

Content delivery networks (CDN) are global scale networks designed to improve user web experience by delivering low-latency and highly available content, primarily by serving content close to users. For large content providers, the performance gains offered by CDNs are critical because slow performance directly impacts revenue. For example, both Amazon [20] and Yahoo [27] have disclosed that hundreds of milliseconds of additional latency have a major impact on revenue.

The means by which a CDN directs users to a nearby locations is called *redirection*. A common approach to use for latency-sensitive workloads such as Search is *DNS* (§II-1).[1] While most of the time DNS-redirection performs well, the "client-LDNS mismatch" problem (§II-2) hurts performance by directing users to distant servers [26].

EDNS Client Subnet (ECS) [14] is a specification to address the client-LDNS mismatch problem by allowing an to LDNS forward a portion of the client's IP address to an authoritative DNS server. The authoritative DNS can use that client-specific information to make accurate per-client decisions rather than by LDNS. For ECS to function, both the client's LDNS and the authoritative resolver must implement the ECS specification.

ECS adoption on the authoritative end by CDNs and cloud providers has been quick, wide-spread, and well advertised. Networks such as Google, Amazon, CloudFlare, CacheFly, EdgeCast [1], Akamai [5], [13], Microsoft [3], and others have publicly advertised their adoption of ECS. This makes sense given CDNs, cloud providers, and ad-revenue driven content providers have strong financial incentives to sell and operate performant services. But without widespread adoption in LDNSes, the benefit of ECS is limited.

Outside of the two largest public DNS services, Google Public DNS and OpenDNS [1], the state of ECS adoption in LDNSes is unknown, even though the majority of end-user ISPs continue to operate their own LDNS services(§IV-A).

In this work, we quantify the adoption of ECS in LDNSes on the Internet. We are motivated by importance of accurate DNS-redirection performance to CDNs, and therefore the entire Internet ecosystem. As ECS is currently the most promising solution to the client-LDNS mismatch problem, measuring ECS prevalence is critical to understanding the factors impacting performance of CDNs that control the majority of content on the Internet.

Our work contributes the following results:

- We demonstrate that the client-LDNS mismatch problem is not only limited to public DNS services. The average distance from 15% of client/non-public LDNS pairs is 900km or more.
- We find that adoption of ECS in LDNSes across the Internet is minuscule; observing only 76 ASes sending any ECS queries.
- ECS query volume seen by authoritative resolvers is large, around 40% of total volume, due to the popularity of Google Public DNS and OpenDNS.
- Our results reveal that 39% of ASes originating ECS queries are from China. We discover this is due to Chinese public DNS services deploying their resolvers into end-user networks; a practice more commonly associated with large CDNs [10], [31].

II. Background

In this section, we discuss the background on DNS-redirection, the motivation behind ECS, and how ECS works.

1) DNS-based Redirection: A CDN has geographically distributed server clusters, called *front-ends* (a.k.a. edges or proxies), serving nearby users to shorten paths and improve performance. CDNs direct clients to a front-end by *redirection.* DNS-based redirection has been a historically popular choice with several large CDNs such as Akamai [23] and Google [10], [15].

[1] While there are several other popular approaches to redirection, such as anycast [11] or HTTP, ECS is only relevant for DNS redirection so we do not discuss others in this work.

978-3-903176-17-1 / © 2019 IFIP

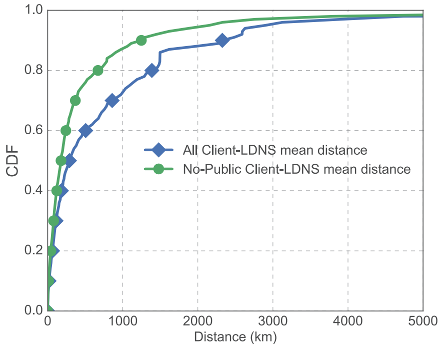

Fig. 1: Distribution of average distance between Microsoft users and their LDNS.

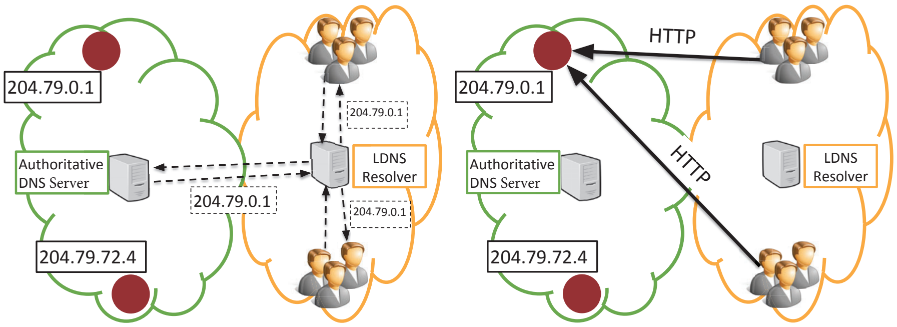

(a) **(b)**

Fig. 2: The DNS resolution process simplified. CDN authoritative DNS server can only make per-LDNS decisions **(a)**, so some clients may be directed to a distant front-end **(b)**.

Figure 2(a) shows the DNS-redirection process[2]. When a client wants to access a web resource (e.g. an image) hosted by the CDN, they send a DNS request to their local DNS (LDNS) resolver, which, if lacking a cached response, forwards the request to the CDN's *authoritative* server. In general, the authoritative server can only see that the request originated from the LDNS IP address, not a specific client. The CDN's authoritative server then returns its estimate of the best performing front-end for the clients served by the LDNS. This works well if those clients are geographically and topologically proximate to each other, but this is often not the case [13], [16], [24].

2) Client-LDNS Mismatch: Figure 2(b) illustrates the client-LDNS mismatch problem. By distance, the figure shows the best performing front-ends for the lower and upper client groups are `204.79.72.4` and `204.79.0.1` respectively. Since an authoritative server can only make a per-LDNS decision, there is no single response that will provide good performance to all clients and the lower clients are directed to a distant front-end. Next we discuss the ECS solution.

3) EDNS client-subnet-prefix: ECS exposes a portion of the client's IP address to an authoritative server, allowing CDNs to make a client-specific redirection choices [14]. For ECS to work, both the LDNS and authoritative server must have support. When an LDNS receives a client DNS request, it embeds the client prefix in the request and forwards it to the authoritative server. The specification recommends that the client prefix length be at most /24 for privacy purposes and previous work observed this is largely the case in practice [10], [28]. If the authoritative server supports ECS, it will send a response with the *scope* that the authoritative server's response covers. For example, the client prefix length may have been sent as /24 but the authoritative server decision may have been scoped to /22. The response scope plays a critical role in LDNS response caching. If the scope is large, the cached response is valid for more clients, but the response may be imprecise for some clients. Smaller scopes generate more cache entries, potentially increasing eviction and decreasing hit rate, but provide precise client responses.

This process is typically transparent to the end-user. Public DNS resolvers such as Google Public DNS and OpenDNS do not forward client-specified ECS queries to authoritative servers. Some ECS-enabled authoritative servers will response to client-specified ECS request sent directly to them.[3] Other authoritative servers employ whitelists to prevent enumeration [10], and so only incorporate ECS information in resolution when the resolvers are in their whitelists, such as Google Public DNS or OpenDNS.[4]

III. Motivation

ECS solves the client-LDNS mismatch problem. Previous work from Akamai showed that for clients using public DNS resolvers, such as Google Public DNS and OpenDNS, the mean distance between clients and their servers was over 3,200km before deploying ECS. After deploying ECS, the mean distance dropped to around 400km [13]. For countries that have very high public DNS usage, they observed a 50% reduction in RTT, from 200 to 100 milliseconds. This demonstrated that poor performance due to client-LDNS mismatch was fixed by ECS, but left open the question of how much opportunity there is for ECS to improve CDN redirection for non-public resolvers.

Figure 1 shows the CDF of mean client distances between Microsoft users and their LDNS from the Odin [12] dataset (§IV-A). Using the methodology described in Section 3.3 of Chen et al. [13], the mean client distance for an LDNS is computed by first computing the centroid of an LDNS - the average location of all clients that it serves, and then finding the mean of the distances between all clients and the centroid. The "All client-LDNS mean distance" line shows similar results to what Akamai reports for all client-LDNS pairs. The "No-Public" line is the client-LDNS mean distances

[2] We simplified DNS resolution process here to focus on ECS related DNS components.

[3] e.g.$ dig +subnet=1.2.3.4/22 @ns1.google.com www.google.com

[4] Akamai falls into this category [18].

excluding Public DNS resolvers.[5] This line shows that around 15% of client-LDNS pairs are around 900km or more from each other, meaning that large proximity between clients and their LDNS is not only a public resolver issue. Given that ECS can improve performance for a significant fraction of non-public LDNSes, we seek to understand Internet-wide ECS adoption in LDNSes.

IV. Methodology and Datasets

This section describes the methodology and datasets we use to evaluate ECS adoption in LDNSes on the Internet.

A. Datasets

ADNS Logs. Our primary dataset is a one month snapshot of production Microsoft Azure authoritative DNS (ADNS) logs collected during January 2019. ADNS are globally distributed across Azure's cloud infrastructure and host DNS records of Azure customers. For each DNS request, we use the timestamp, LDNS IP address from which the query was received, and the client-subnet prefix if the query was ECS-enabled. Our trace contains over 155 billion queries over 5 million LDNS IP addresses from 45,366 ASes, covering nearly 70% of ASes advertised on the Internet.

Odin. The second dataset is a two week collection of client-LDNS association logs from Microsoft's client-side Internet measurement platform called Odin [12]. Odin runs in popular Microsoft end-user applications and measures performance, such as latency and availability, of Microsoft's network. Part of Odin's measurement suite includes a technique borrowed from Mao et al. [22] to uncover client-LDNS associations by generating a unique hostname (e.g. **$(Rand)**`.contoso.com`) that gets resolved by Odin's authoritative DNS and also gets uploaded by the measurement client over HTTP(S). A join between HTTP and DNS logs on **$(Rand)** provides the client-LDNS association.

Our datasets have a few limitations. First, both our ADNS and Odin datasets contain only IPv4 LDNSes. Second, although unlikely, it is possible we may misclassify LDNSes as non-ECS-enabled if they only forward ECS information to authoritative DNS servers in a whitelist that does not include Microsoft.[6] Third, we do not claim to have complete coverage of LDNSes or ECS-enabled LDNSes on the Internet. As a large cloud provider, Microsoft's view of the Internet is broad, but limited to Microsoft's customers and customers of applications built on top of Azure. We further discuss coverage in the next section.

B. LDNS Coverage

Our data contains LDNSes from over 45,000 ASes and 5 million LDNS IP addresses which we believe is a large enough sample to provide insight into the state of ECS adoption. For reference, Akamai's evaluation of client-LDNS proximity [13] uses 584,000 distinct LDNSes which account for 84.6% of Akamai's global client demand.

Network Type	Count	Percent	Examples
CDN	1	1.3	Akamai
Cloud / Hosting	26	34.2	Amazon, Microsoft, Oracle, OVH, Rackspace, Packet Host
Edu	2	2.6	Oklahoma Network for Education Enrichment, China Education and Research Network Center
Misc	6	7.8	AppRiver, Alibaba, Tencent
Public DNS	3	3.9	Google Public DNS, OpenDNS
Telco	38	50	China Telecom, Korea Telecom, Frontier, Vodafone Kabel, Telefonica Germany
Total	76	100%	

TABLE I: Network type of ASes with ECS-enabled LDNSes. Due to the small number of ASes, we investigated and labeled each by hand.

RIR	Country	Count	Percent
AFRINIC	South Africa	1	1.3
APNIC	China	30	39.4
	Hong Kong	2	2.6
	India	1	1.3
	South Korea	1	1.3
ARIN	United States	3	3.9
LACNIC	Brazil	2	2.6
	Chile	1	1.3
	Ecuador	1	1.3
RIPE	Georgia	1	1.3
	Germany	5	6.5
	Moldova	1	1.3
	Netherlands	4	5.2
	Russia	2	2.6
	United Kingdom	1	1.3
Global		20	26.3
Total		76	100%

TABLE II: The geographic footprint of ASes that host ECS-enabled LDNSes. The "Global" field indicates ECS-enabled LDNS services spanning multiple continents such as Google Public DNS.

We examine our LDNS coverage by comparing against the ground-truth of two large public DNS services, OpenDNS and Google Public DNS, and coverage of LDNSes used by a popular measurement platform, RIPE Atlas [9]. From snapshots taken on 2019-02-12, our datasets have 100% coverage of all 107 prefixes listed on the Google Public DNS FAQ page [4] and all 31 prefixes listed in the OpenDNS data center location page [7]. For RIPE Atlas, we discover 6,552 distinct public LDNSes across all 10K active probes using the methodology described in §V-D. Our dataset overlaps with 97.9% of these LDNSes.

Given that LDNSes are shared Internet resources across hundreds or thousands of individual users, we believe our adoption results in §V are largely generalizable because there is likely high overlap of LDNSes seen across other large cloud networks. The exception is §V-B because these results are heavily dependent on the user population mix, which is likely very different across providers.

V. Results

Having described the need for ECS, and its importance for CDN performance, we now focus on the state of ECS adoption in LDNSes.

[5] Based on usage of Google, CloudFlare, OpenDNS, Quad9, DynDNS, UltraDNS, and Yandex public DNS services.

[6] OpenDNS is a service that uses an ECS whitelist [8].

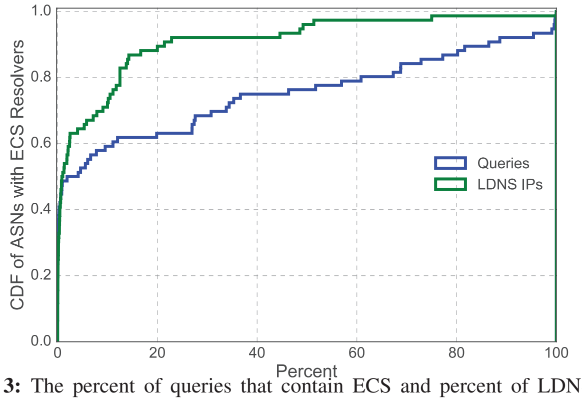

Fig. 3: The percent of queries that contain ECS and percent of LDNS IPs that send ECS queries across all ASes with ECS-enabled LDNSes.

A. *Have Networks Widely Enabled ECS for LDNS?*

Using the ADNS dataset, we first investigate how many ASes have deployed ECS-enabled LDNSes. We find that only **76** ASes host LDNSes that sent any ECS queries to Azure hosted domains or cloud resources over the one month period, indicating that ECS adoption across ASes is very small (0.1%). We exclude the full list of ISPs for space, but is available at https://aka.ms/ecslist.

Table I shows the ASes broken down by network type. The two dominant network types are Telco (50%) and Cloud/Hosting (34.4%). Since Telco networks provide connectivity to most end-users, and end-users are sensitive to CDN performance, it is intuitive that they would host the largest number of ECS-enabled LDNSes. Cloud environments are also a logical choice for ECS-enabled LDNSes, but for services instead of end-users. One reason is that many cloud providers offer services that are not available in all regions. ECS enables multi-region deployments to access services in the nearest available region. Another reason is businesses can have inter-cloud architectures. ECS facilitates a service running in one cloud provider region to get directed to a nearby region in another cloud provider. We observe many cloud providers present in our dataset including Amazon, Microsoft, Google, Oracle, and OVH.

Intra-AS ECS adoption is variable. Figure 3 shows the percent ECS queries and ECS-enabled LDNS IPs per AS. For nearly 50% of ECS-enabled ASes, ECS queries and ECS-enabled LDNS IPs make up less than 1% of the total DNS queries and LDNS IPs observed. Low adoption by queries include Amazon and Microsoft from the Cloud side, Telcos such as Vodafone Libertel(AS33915), Korea Telecom(AS4766), and several regional China Telecom networks(AS58563,AS58466). There are a number of possible reasons for this behavior. For cloud providers, it may be that the ECS-enabled LDNSes are hosted VMs for another network. For low adoption Telco networks, ECS may be part of a small-scale trial period or a small number of individual users hosting their own ECS-enabled LDNSes. If Telcos provide business connectivity, it may be that enterprises host their own ECS-enabled LDNSes but on the AS's IP space.

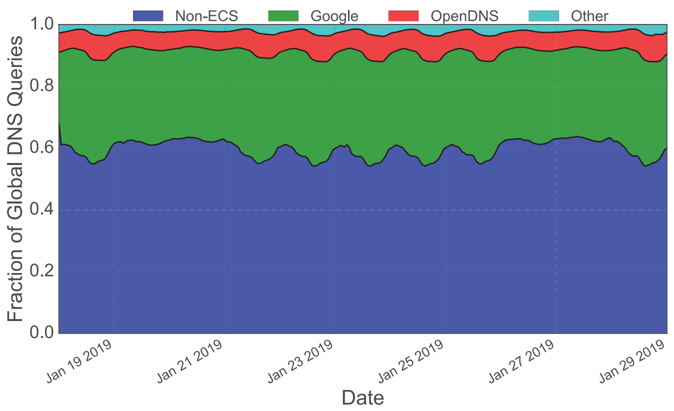

Fig. 4: ECS and Non-ECS global DNS request traffic observed from Azure authoritative DNS servers during a 10 day period in January 2019.

Adoption by query volume is not indicative of adoption by LDNS IP. In the case of Google, we see 99% of queries as ECS-enabled but only 13.7% of LDNS IPs observed send any ECS queries. For OpenDNS's primary AS we observe nearly 50% of LDNS IPs as ECS-enabled but by query volume it is 99.5%.

Next, we examine the geographic footprint of ASes with ECS-enabled LDNSes. Table II shows the breakdown by Regional Internet Registry (RIR) and country. Around 26% of ASes are marked as "Global", meaning that they offer services in multiple continents. This includes most cloud providers and public DNS services. In the RIPE region, the Netherlands and Germany have relatively strong adoption compared to other countries with a combined 11.2% of ASes. In LACNIC, Brazil(AS28299,AS61813), Chile(AS28099), and Ecuador(AS27947) all have regional Telco networks with ECS resolvers. In North America, there are only three domestic ASes, including Frontier, one of the largest broadband providers in the United States. The most striking result is that almost 40% of all ASes hosting ECS-enabled LDNSes are in China. We present a detailed investigation of ECS in China in Section V-D.

Our results show that across ISPs on the Internet, ECS adoption is very low. However, as we observe in the next section, some networks contribute a disproportionate volume of DNS traffic, revealing a different view of adoption from the authoritative server perspective.

B. *How Prominent are ECS Queries for a Cloud Provider?*

Next, we explore ECS query volume observed in Azure as compared to non-ECS queries.

We find that despite low per-network adoption on the Internet, ECS queries make up a large percentage of all authoritative DNS requests in Azure. Figure 4 shows the breakdown of ECS and non-ECS traffic globally over 10 days in January 2019. The shaded blue portion shows the percentage of non-ECS traffic is only slightly higher than ECS traffic. We observe a strong diurnal pattern of higher non-ECS traffic during North American weekday business hours that fluctuates between 55% and 65% percent of total DNS traffic. We also

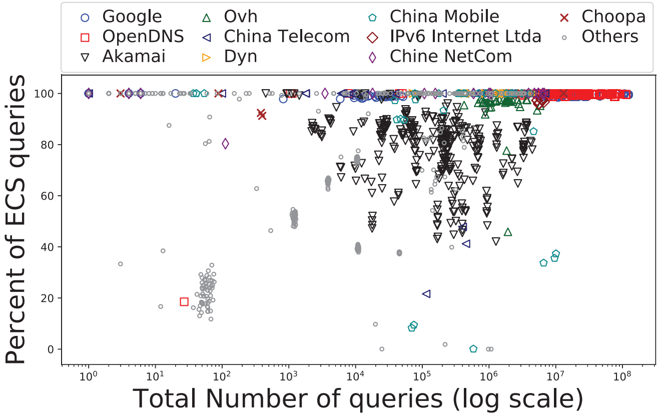

Fig. 5: Total query volume and the percent of ECS-enabled queries for the top 10 ASes by query volume. Each point represents a single LDNS IP observed in the dataset.

observe a plateau in ratio changes due to weekend traffic patterns.

Out of ECS queries in Figure 4, Google's public DNS traffic dominates, making up 30-35% of all DNS queries and around 75-80% of global ECS DNS traffic. OpenDNS is the second largest band with 5-8% and all other 73 ECS-enabled ASes combined make up 2-3% of DNS queries.

While we initially assumed that LDNSes would either fully enable ECS or not at all, we found that most ECS-enabled LDNSes show mixed support: 85% of the 3221 ECS-enabled LDNS IP addresses also send non-ECS queries, including large DNS and CDN providers. Figure 5 shows all LDNS IP addresses that send mixed ECS queries by AS. We observe that over 98% of queries from Google and OpenDNS ECS-enabled LDNS IPs addresses contain ECS. We verify that 97% of Google ECS-enabled LDNS IP addresses observed in our dataset belong to Google Public DNS prefixes [4]. We also observe that Akamai has a wide range of mixed ECS support (45% to 98% ECS-enable queries) for their LDNS addresses (black triangle in Figure 5).

We were unable to find a clear explanation for mixed ECS behavior. One possibility is that some LDNSes in our dataset are DNS forwarders with ECS and non-ECS LDNSes behind it. Another is that ECS enabled by default but disabled for certain hostnames or clients for privacy reasons. ECS data might be also dropped by middlebox such as a DNS forwarder, when a DNS packet is large or the middlebox does correctly support EDNS0.

Our results in this section showed that the volume of ECS queries served by our authoritative cloud DNS servers is quite large and originate mostly from Google and OpenDNS public DNS services. Next, we look at how this high ECS query volume translates to adoption across individual end-users on the Internet.

C. How Prevalent is ECS in End-user Networks?

In this section, we examine LDNS ECS adoption across end-users. For this analysis we rely on the client-LDNS association data from Odin (§IV) because we have client information for all LDNSes, not just those that are ECS-enabled.

Non-ECS	ECS	Google	OpenDNS	Other
78.7 %	21.3%	19.1%	2%	0.2%

TABLE III: Global client IP address adoption of ECS-enabled LDNSes from our Odin dataset.

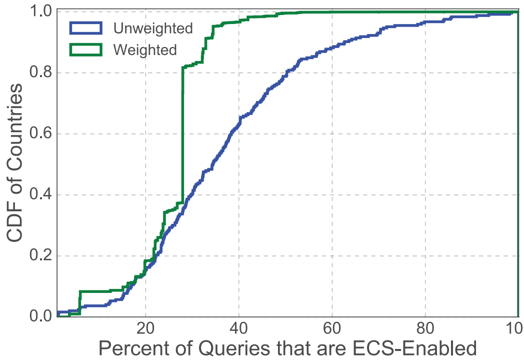

Fig. 6: The percent of ECS traffic across client countries. The green line is weighted by per-country end-user traffic volume (HTTP requests).

The large volume of ECS query volume seen in Figure 4 shows a skewed version of ECS adoption. Table III shows the percent of global client IP addresses served by ECS or non-ECS LDNSes. Only around 22% of client IPs observed in our data use ECS enabled resolvers. Google public DNS serves 19% of all client IP addresses and 90% of those using ECS-enabled resolvers. Two percent of clients use OpenDNS and only a 10th of a percent of clients use ECS-enabled LDNSes in other ASes.

Figure 6 shows the distribution of end-user ECS usage across countries. In the weighted distribution, each country is weighted by its daily average client HTTP request volume. Around 95% of weighted countries have 35% or less of their DNS traffic ECS-enabled. In contrast, around 95% of unweighted countries have 80% or less ECS-enabled traffic. This difference is due to many small nations with high dependence on public DNS services and low traffic volume.

Figure 7 gives us a global view of end-user ECS usage to contrast the difference in adoption across regions. In several

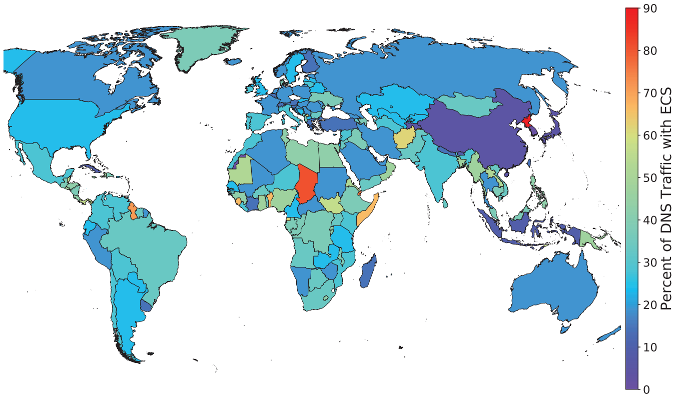

Fig. 7: The percent of ECS traffic by client country.

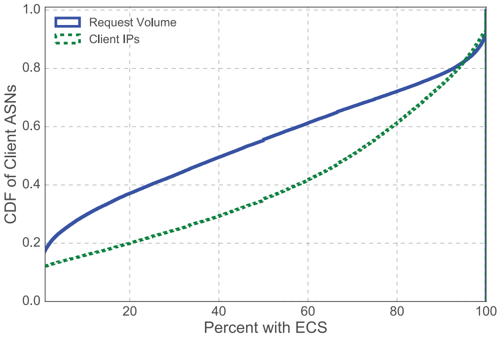

Fig. 8: The distribution of ECS usage across end-user ASes. The "Client IPs" line shows distinct client IPs within an AS that use ECS enabled LDNSes, and the "Request Volume" line weights each client IP by its HTTP(S) traffic volume in Azure.

populous Asian countries, such as China, Japan, and Korea, ECS usage is 6% or less. In Southeast Asia, Indonesia stands out having low adoption compared to its neighbors. Much of North and South America are in the 30-35% range, with exceptions of lower usage in Uruguay and Peru, and very high usage in Guyana. Much of Europe and Russia are in the 15-20% range. Table IV summarizes the lowest and highest ECS adopters by per-country query volume.

Bottom 10		**Top 10**	
Country	**ECS Query Volume**	**Country**	**ECS Query Volume**
Western Sahara	0%	Saint Pierre and Miquelon	98.49%
Saint Helena	0.75%	North Korea	98.03%
Wallis and Futuna	1.3%	Cape Verde	93.58%
Norfolk Island	1.38%	Niue	90.95%
South Korea	5%	Tokelau	85.59%
Aland Islands	5.57%	Cook Islands	85.36%
Japan	5.86%	Chad	84.01%
China	5.99%	Liechtenstein	83.83%
Tajikistan	7.08%	Djibouti	79.8%
Greenland	11.47%	Marshall Islands	79.13%

TABLE IV: The highest and lowest ECS-enabled countries.

Figure 8 shows the percent of clients using ECS-enabled LDNSes across all ASes. We see that 88% of client ASes have some clients using ECS-enabled LDNSes. Slightly less than 10% of ASes are exclusively using ECS-enabled LDNSes and 12-18% have none.

Considering the small number of ISPs hosting ECS-enabled resolvers, a significant number of clients around the world are impacted by ECS. Interestingly, from Table IV, we observe that even though China has the largest number of ASes with ECS-enabled resolvers, their ECS query volume is 7th lowest in the world. In the next section we examine this discrepancy.

D. ECS Adoption in China

Table II shows that out of the 76 ASes which contain ECS-enabled LDNSes, more than one third (30, 39.4%) are from China, while Figure 7 shows the percentage of ECS-enabled DNS queries in China is just around 5%. Table V lists the ISP names of these ASes. China Telecom(CT), China Unicom(CU), China Mobile(CM), and CERNET are major Chinese ISPs while the rest are among major Chinese IT companies. The presence of ECS-enabled LDNSes in major Chinese ISPs is in sharp contrast to adoption observed in the rest of the world. We next describe our investigation to understand this difference.

We first attempted to use reverse DNS lookup and CHAOS class DNS queries [32] to gather information about the ECS-enabled LDNS IPs. However, these methods were unhelpful due to low response rate.

Since public DNS services are the well-known adopters of ECS, we next examined Chinese public DNS services. Table VI shows there are many Chinese public DNS services (details available at https://aka.ms/cnpubdns), with 4 out of 7 supporting ECS. Public DNS+ documents that they support ECS, while we identified three others (114 DNS, DNS Pai and OneDNS) by sending ECS queries from an Azure VM to their anycast serving IP and confirmed ECS options present in the reply.

In order to understand if the ECS-enabled public DNS services contribute to the ECS-enabled LDNSes in our ADNS dataset, we need to know the *backend IPs* of these public DNS services. A backend IP is a unicast IP address that forwards the request to the authoritative server. LDNS requests do not originate from the anycast IP address because the authoritative server's response may be directed to a different LDNS site.

However, unlike Google and OpenDNS who publish their backend unicast IP ranges, the Chinese public DNS services do not. As a result, we use Chinese RIPE Atlas [9] nodes to collect the backend IPs of these LDNSes. We selected 14 active nodes in 4 major ISPs of China: 5 from China Telecom (CT), 4 from China Unicom (CU), 4 from China Mobile (CM), and 1 from CERNET. We query "whoami.akamai.net" [6] and specify each of the Chinese public DNS's serving IP as the target nameserver to collect their backends. The "whoami.akamai.net" hostname is a service run by Akamai that returns the IP address of the querying LDNS as seen by Akamai's authoritative server. Our measurement ran every 2 hours from Jan 9th to Jan 13th, 2019 to cover peak/normal hours and weekday/weekends.

Table VII shows the backend IPs collected from the Chinese RIPE Atlas nodes. In total we found 237 unique backend IPs from the seven Chinese public DNS services. We then compare this list to the list of Chinese ECS-enabled LDNSes in our ADNS dataset and found 56 overlapping IPs from the four public DNS services: 114 DNS, Public DNS+, OneDNS and DNS Pai. That means 47% (56/118) of the Chinese ECS-enabled LDNS in our ADNS dataset belong to public DNS services. Given our limited vantage points of 14 RIPE Atlas nodes, we were only able to verify that a subset of the Chinese ECS-enabled LDNSes in the ADNS dataset belong to public DNS services, but believe there are more.

We also examine the ASNs of all 237 backend IPs we found in RIPE Atlas data and list the ISP names in the last column of Table VII. Interestingly, most of these backend IPs of Chinese public DNS services are from the major Chinese ISPs' ASNs.

Name	# ASNs
China Telecom (CT)	13
China Unicom (CU)	6
China Mobile (CM)	5
CERNET	1
Alibaba	2
Tencent	2
Huawei	1
Total	30

TABLE V: Chinese ISPs that send ECS. Some ISPs regional ASNs we count towards primary ISP. For example, AS51466 is CT's network for Guangdong province.

Name	ECS?
Baidu DNS	No
sDNS	No
114 DNS	Yes
Ali DNS	No
Public DNS+	Yes
OneDNS	Yes
DNS Pai	Yes

TABLE VI: Chinese Public DNS services

	Backend IPs	**Overlap w/ Chinese ECS-enabled LDNS**	**Backend ISP**
Total	**237**	**56**	
ECS enabled	**166**	**56**	
114 DNS	35	3	CT, CU, CM
Public DNS+	66	49	CT, CU, CM
OneDNS	47	22	CT, CU, CM
DNS Pai	52	1	CT, CU, CM, CERNET
No ECS	**71**	**0**	
Baidu DNS	29	0	CT, CU, CM, Baidu
sDNS	8	0	CU, Alibaba, CNNIC
Ali DNS	34	0	CT, CU, CM, CERNET

TABLE VII: Backend IPs of Chinese Public DNS services. ISP abbreviations are in Table V.

This explains why many Chinese ASNs show up in Table II. Even though there are only four Chinese public DNS services that support ECS, their backends are actually hosted in major ISPs. This makes attributing LDNSes to their actual services challenging.

To our knowledge, this is the first work to measure public DNS services deployed in other ISPs' networks. This deployment strategy was previously associated with large CDNs such as Google and Akamai [10], [31]. Previous work has studied China's network topology and reports a highly hierarchical structure [30] and inter-ISP links are known bottlenecks that cause high latency [33]. In addition, there is limited peering infrastructure in China [21] and the first IXP in China was built only in 2016 [2]. Under such topology, public DNS services must be present in major ISPs to get competitive latency.

Another interesting result is that different Chinese public DNS services can share the same backend IPs. For example, we observed 61.151.186.143 shared by DNSPod and OneDNS. While we don't exactly know why, one hypothesis is that public DNS services share infrastructure to achieve better customer performance when they have different footprints in ISP networks.

Lastly, since we also collected the backend IPs of default LDNS resolver of all Chinese RIPE Atlas nodes, which should be the ISPs' LDNS, we can check if any ISPs' LDNS show up as ECS enabled. Of the 59 default LDNSes from RIPE Atlas nodes, we observe 30 in the ADNS dataset as non-ECS only, suggesting that none of the ISP operated LDNSes support ECS.

To summarize, we verified that at least half of the Chinese ECS enabled LDNSes belong to Chinese public DNS services. These public DNS services deploy their service inside other ISPs' network so many Chinese ASNs show up as ECS enabled. We found no evidence that Chinese ISPs support ECS on their own LDNSes.

VI. Discussion

A. Lack of Adoption

The lack of ECS adoption in LDNSes means that many end-users receive sub-optimal performance from CDNs. In this section, we discuss possible reasons for limited adoption.

Akamai's previous work described ECS scaling challenges from the CDN's perspective [13]. They point out that with an order of magnitude more clients than LDNSes, a CDN mapping system must perform DNS resolution at much finer granularity. This requires more memory in DNS servers to support records at a client prefix level and more precise measurements to capture per client prefix performance. The second challenge is supporting additional DNS query volume, with an increase roughly on the order of the number of /24s served per LDNS.

LDNSes share a different set of challenges than authoritative servers. It is well understood that DNS resolution contributes to user-perceived latency. For an LDNS provider to offer the same resolution performance with ECS, they must support a much larger cache for per client-prefix DNS responses. If additional cache is not available, the cache hit rate will drop and user-perceived latency may increase. Whether the latency benefit of the CDN's improved redirection is enough to overcome the penalty of a full DNS resolution remains an open problem.

A second issue is cost. To avoid degrading existing user performance, an LDNS provider may face with the cost of upgrading a fleet of LDNS machines to support larger caches. In addition, if their existing DNS software does not support ECS, there is an engineering cost to add support or a cost to switch DNS software. Additional concerns may include end-user privacy and exposing proximity or topological relationships between clients and LDNSes to third parties. Given these factors, there is no obvious incentive for end-user ISPs to invest in supporting ECS. In contrast, popular public DNS services are motivated to improve end-user performance with ECS support because their usage provides valuable business intelligence such as what popular domains individual IP addresses query for.

B. Future Work

Our results have strong implications for research and operations in the CDN space. ECS is a solution that is not being widely adopted but performance problems from client-LDNS mismatch remain. Additional work is needed to either solve limitations of ECS (such as cache explosion) or to explore new solutions. One solution put forth in previous work [11] showed the benefits of a hybrid anycast-DNS redirection system. Results showed that anycast could improve performance in many client-LDNS mismatch cases but the evaluation was limited and to our knowledge, has never been evaluated in production.

In our investigation of ECS in China (§V-D), we identified a challenging issue: When services, such as public DNS, deploy infrastructure in external ISPs, how can the operator of that service be identified from an IP address?

VII. Related Work

Previous work also explored ECS adopters. Most closely related to our work, Sudrajat uses active probes to study ECS adoptions on the Internet [29]. Using open DNS resolvers, they found 17 ASes with ECS-enabled LDNSes. However, this work is limited by querying only open DNS resolvers. Most broadband, mobile, and enterprise networks do not serve queries outside their own network because open DNS resolvers are often used for DDoS attacks. This means the LDNSes used by most users will not be covered. Our passive measurements at production authoritative DNS servers at Microsoft have high coverage of high traffic volume LDNSes.

Streibelt et.al [28] examined cache and user clustering strategies of several CDN ECS adopters. Otto et.al [24] presented the first evaluation of ECS's effectiveness and the adoption level by CDNs. Chen et.al [13] studied the performance impact of ECS adoption at Akamai. Sánchez et.al [25] also evaluate the effectiveness of ECS with the Dasu measurement platform.

Previous work explored DNS based server selection [15], [26] and performance of other client redirection mechanisms [11]. Other work examined the client-LDNS mismatch problem through measurements of Web and DNS traffic [17], [22] and proposed solutions of injecting client information in hostnames [16].

ECS enables new way of understanding the Internet and online services, while the capabilities also raise security and privacy concerns. Calder et.al [10] leverage ECS to study the expansion of Google's serving infrastructure. Kintis et.al [19] claim that ECS negatively impacts DNS privacy and enables targeted DNS poisoning attacks. Our work compliments these studies by better understanding the scope of ECS usage on the Internet.

VIII. Conclusion

In this paper we examined the adoption of ECS in LDNSes from the perspective of a large cloud provider. We find that even though a large percent of non-public LDNSes would be benefit from ECS, we observe only only 76 ASes with any ECS queries over a one month period. In contrast, ECS query volume is quite high due to the popularity of public DNS services. We speculate that adoption is limited by cache explosion and a lack of incentives. Our findings motivate the need for continued research to improve client redirection.

References

[1] A Faster Internet: The Global Internet Speedup. http://afasterinternet.com/participants.htm.

[2] Building CHN-IX: the first IXP in mainland China. https://blog.apnic.net/2016/04/22/building-chn-ix-first-ixp-mainland-china/ Accessed 2018-02-22.

[3] EDNS Client Subnet Support in Azure Traffic Manager. https://azure.microsoft.com/en-us/blog/edns-client-subnet-support-in-azure-traffic-manager/.

[4] Google Public DNS. https://developers.google.com/speed/public-dns/faq Accessed 2019-02-12.

[5] Google Public DNS and Location-Sensitive DNS Responses. https://webmasters.googleblog.com/2014/12/google-public-dns-and-location.html?language=en_US.

[6] Introducing a New whoami Tool for DNS Resolver Information. https://developer.akamai.com/blog/2018/05/10/introducing-new-whoami-tool-dns-resolver-information.

[7] OpenDNS Data Center Locations. https://www.opendns.com/data-center-locations/ Accessed 2019-02-12.

[8] OpenDNS EDNS Client Subnet FAQ. https://support.opendns.com/hc/en-us/articles/227987647-EDNS-Client-Subnet-FAQ Accessed 2019-02-23.

[9] RIPE Atlas. https://atlas.ripe.net/.

[10] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the Expansion of Google's Serving Infrastructure. In *IMC 2013*.

[11] M. Calder, A. Flavel, E. Katz-Bassett, R. Mahajan, and J. Padhye. Analyzing the Performance of an Anycast CDN. In *IMC*, 2015.

[12] M. Calder, M. Schröder, R. Gao, R. Stewart, J. Padhye, R. Mahajan, G. Ananthanarayanan, and E. Katz-Bassett. Odin: Microsoft's Scalable Fault-Tolerant CDN Measurement System. In *NSDI*, 2018.

[13] F. Chen, R. K. Sitaraman, and M. Torres. End-user mapping: Next Generation Request Routing for Content Delivery. In *SIGCOMM 2015*.

[14] C. Contavalli, W. van der Gaast, D. Lawrence, and W. Kumari. Client Subnet in DNS Queries (RFC7871), July 2015.

[15] X. Fan, E. Katz-Bassett, and J. Heidemann. Assessing Affinity Between Users and CDN Sites. In *TMA*, 2015.

[16] C. Huang, I. Batanov, and J. Li. A Practical Solution to the Client-LDNS Mismatch Problem. *CCR*, 42(2), 2012.

[17] C. Huang, D. A. Maltz, A. Greenberg, and J. Li. Public dns system and global traffic management. In *INFOCOM 2011*.

[18] M. Jansen. EDNS0 Client-Subnet for DNS based CDNs. SANOG 24 2004.

[19] P. Kintis, Y. Nadji, D. Dagon, M. Farrell, and M. Antonakakis. Understanding the privacy implications of ECS. In *DIMVA 2016*.

[20] G. Linden. Make Data Useful. http://sites.google.com/site/glinden/Home/StanfordDataMining.2006-11-28.ppt, 2006.

[21] Y. Z. Liu. Internet Landscape in China. https://www.globalpeeringforum.org/pastEvents/gpf12.0/wed_liu_china.pdf Accessed 2018-02-22, 2017.

[22] Z. M. Mao, C. D. Cranor, F. Douglis, M. Rabinovich, O. Spatscheck, and J. Wang. A Precise and Efficient Evaluation of the Proximity Between Web Clients and Their Local DNS Servers. In *ATC*, 2002.

[23] E. Nygren, R. K. Sitaraman, and J. Sun. The Akamai Network: A Platform for High-performance Internet Applications. *SIGOPS 2010*.

[24] 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 *IMC 2012*.

[25] M. A. Sánchez, J. S. Otto, Z. S. Bischof, D. R. Choffnes, F. E. Bustamante, B. Krishnamurthy, and W. Willinger. Dasu: Pushing experiments to the internet's edge. In *NSDI 2013*.

[26] A. Shaikh, R. Tewari, and M. Agrawal. On the effectiveness of DNS-based server selection. In *INFOCOM 2001*.

[27] S. Stefanov. Yslow 2.0. In *CSDN SD2C*, 2008.

[28] F. Streibelt, J. Böttger, N. Chatzis, G. Smaragdakis, and A. Feldmann. Exploring EDNS-client-subnet Adopters in Your Free Time. In *IMC*, 2013.

[29] F. U. Sudrajat. The State of Adoption of DNS ECS Extension on the Internet, 2017.

[30] Y. Tian, R. Dey, Y. Liu, and K. W. Ross. Topology mapping and geolocating for china's internet. *IEEE Transactions on Parallel and Distributed Systems*, 2013.

[31] F. Wohlfart, N. Chatzis, C. Dabanoglu, G. Carle, and W. Willinger. Leveraging Interconnections for Performance: The Serving Infrastructure of a Large CDN. In *SIGCOMM 2018*.

[32] S. Woolf and D. Conrad. Requirements for a Mechanism Identifying a Name Server Instance (RFC4892), June 2007.

[33] J. Xue, D. Choffnes, and J. Wang. CDNs meet CN - An Empirical Study of CDN Deployments in China. *IEEE Access*, 2017.