ECSeptional DNS Data: Evaluating Nameserver ECS Deployments with Response-Aware Scanning

PATRICK SATTLER, Technical University of Munich, Germany
JOHANNES ZIRNGIBL, Max Planck Institute for Informatics, Germany
FAHAD HILAL, Max Planck Institute for Informatics, Germany
OLIVER GASSER, IPinfo, United States
KEVIN VERMEULEN, LIX, CNRS, Ecole Polytechnique, France
GEORG CARLE, Technical University of Munich, Germany
MATTIJS JONKER, University of Twente, The Netherlands

DNS is one of the cornerstones of the Internet. Nowadays, a substantial fraction of DNS queries are handled by public resolvers (e.g., Google Public DNS and Cisco's OpenDNS) rather than ISP nameservers. This behavior makes it difficult for authoritative nameservers to provide answers based on the requesting resolver. The impact is especially important for entities that make client origin inferences to perform DNS-based load balancing (*e.g.*, CDNS). The EDNS0 Client Subnet (ECS) option adds the client's IP prefix to DNS queries, which allows authoritative nameservers to provide prefix-based responses. Previous work showed the potential of data collected during ECS scans. Infrastructure can be uncovered, and operators' subnet-specific behavior can be observed.

In this study, we introduce a new method for conducting ECS scans. Our method significantly reduces the required number of queries by up to 97 % compared to state-of-the-art techniques and allows us to provide new insights into ECS behavior. Our approach is also the first to facilitate ECS scans for IPv6. Due to its vast address space, we have developed and analyzed different IPv6 scanning approaches. We conduct a comprehensive evaluation of the ECS landscape, examining the usage and implementation of ECS across various services. Overall, 53 % of all nameservers support prefix-based responses. Furthermore, we find that Google nameservers do not comply with the Google Public DNS guidelines. Additionally, we observe that certain operators (*e.g.*, AWS Route53) exclusively employ a single specific scope prefix length without aggregation, potentially affecting resolver cache efficiency. Lastly, we make our tool and data publicly available to foster further research in the area.

CCS Concepts: • **Networks → Naming and addressing**; **Location based services**.

Additional Key Words and Phrases: EDNS0 Client Subnet extension (ECS); DNS load balancing

ACM Reference Format:
Patrick Sattler, Johannes Zirngibl, Fahad Hilal, Oliver Gasser, Kevin Vermeulen, Georg Carle, and Mattijs Jonker. 2025. ECSeptional DNS Data: Evaluating Nameserver ECS Deployments with Response-Aware Scanning. *Proc. ACM Netw.* 3, CoNEXT2, Article 11 (June 2025), 25 pages. https://doi.org/10.1145/3730977

Authors' Contact Information: Patrick Sattler, sattler@net.in.tum.de, Technical University of Munich, Munich, Germany; Johannes Zirngibl, jzirngib@mpi-inf.mpg.de, Max Planck Institute for Informatics, Saarbrücken, Germany; Fahad Hilal, fhilal@mpi-inf.mpg.de, Max Planck Institute for Informatics, Saarbrücken, Germany; Oliver Gasser, oliver@ipinfo.io, IPinfo, Seattle, United States; Kevin Vermeulen, kevin.vermeulen@laas.fr, LIX, CNRS, Ecole Polytechnique, Palaiseau, France; Georg Carle, carle@net.in.tum.de, Technical University of Munich, Munich, Germany; Mattijs Jonker, m.jonker@utwente.nl, University of Twente, Enschede, The Netherlands.

ACM 2834-5509/2025/6-ART11
https://doi.org/10.1145/3730977

1 Introduction

Popular services on the Internet are commonly served by multiple CDN edge servers. CDNs offer load distribution and low latency services to ensure user satisfaction and improve conversion rates [6, 11, 64]. Therefore, service operators aim to optimize the load distribution.

Two well-known techniques for operators to direct the initial connection of a user to a nearby vantage point are IP anycast and DNS load balancing. This work evaluates a particular aspect of DNS-based load balancing: the EDNS0 Client Subnet (ECS) option [21]. The ECS option is a solution to the hidden client problem for authoritative nameservers. Without it, the authoritative nameserver only knows the recursive resolver IP address, which originated the DNS query. Traditionally, it could only use that recursive resolver IP address to perform topology- or geolocation-based load balancing. With ECS, the resolver includes the client's subnet in the query. Therefore, authoritative nameservers can use this information to provide tailored responses to the client.

Major cloud providers such as Google [32], Amazon [8], Akamai [2], and Cloudflare [20] all provide ECS query responses and offer it as a product to their customers. On the software side, ECS is also widely adopted, *e.g.*, BIND [41], Knot DNS [49], and PowerDNS [63] support scoped responses using ECS. Hence, ECS is available to a large customer base. Additionally, prominent public recursive resolvers such as Google Public DNS and Cisco's OpenDNS also support ECS. The importance of recursive resolvers supporting ECS is well-documented [16]. Existing work has confirmed ECS's effectiveness in reducing client latency [61, 66, 75]. Others [15, 61, 73] have looked at ECS support on the authoritative DNS infrastructure of top-list domains.

Using ECS requires major effort by providers to effectively map clients to Point of Presences (PoPs). Determining how they use ECS, their mapping strategy, and whether they conform to the standard is important. This knowledge can help resolvers, network operators, and clients to better understand the network and perceived quality of experience. The information on how providers map clients to PoPs can be extracted using ECS scans. The scan results can be used to better understand the provider's infrastructure and serve as a data source for various use cases.

This paper presents a novel approach to perform IPv4 and IPv6 ECS scans. Unlike existing approaches, we keep the state from responses already received. By doing so, we are able to significantly reduce the number of required queries, alleviating the load on third-party infrastructure. We reveal and investigate previously unseen load balancing behaviors. Google uses daily patterns to update its client-server mapping. Moreover, we leverage our efficiency-increasing approach to revisit ECS support among popular domain names a decade after prior work [15, 61, 73] and for a more diverse set of top lists. Our results show a significant increase in ECS adoption among popular domains: 79 % signal support for ECS, and 40 % also provide subnet-specific responses.

The main contributions of this work are:

- *(i)* We present a novel ECS scanning approach, the first that supports ECS scans for IPv6 (see Section 4). Due to its response awareness, our approach significantly reduces the number of queries to cover the address space. Additionally, it provides a fine-grained configuration of query limits per prefix length.
- *(ii)* We use our scanning approach to evaluate the ECS behavior of nine popular domains served by four providers (see Section 5). Our approach reduces the number of queries needed for the IPv4 address space by up to 97 % over the state of the art. Based on our analysis, the best set to uncover IPv6 infrastructure is a combination of BGP and IP geofeed prefixes.
- *(iii)* We revisit and explore the prevalence of ECS among top-listed domains and provide insights into different nameserver operator behaviors (see Section 6). We find that Cloudflare provides ECS load balancing as a service for their customers but also responds with scoped answers for their own anycast addresses.

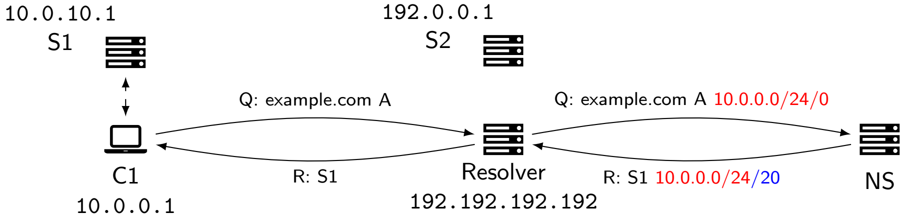

Fig. 1. Example ECS flow: The client C1 is topologically close to the server S1 but uses the resolver topologically close to S2. With the client subnet information in the query (IP address 10.0.0.0 with source prefix length 24), the nameserver can select the correct response and send it back with the original source prefix length 24 and the scope prefix length of 20. The resolver can thus cache the response for all clients within 10.0.0.0/20.

(iv) We perform a detailed analysis of ECS properties for Google, Meta, and AWS Route53 authoritative nameservers (see Section 7). We uncover RFC-violating behavior by Google, which its own public DNS resolver does not accept.

(v) We publish the scanner's source code [69] and make our scanning data publicly available [70].

2 EDNS0 Client Subnet Option

The separation of DNS into clients, recursive resolvers, and authoritative nameservers results in a fundamental problem if DNS is used for (geo-)load balancing: The authoritative nameserver is aware of the resolver and its IP address (thus location). However, the resolver does not have to be geographically or topologically close to the client (*e.g.*, using public resolvers).

To overcome this problem, RFC7871 [21] describes ECS as an Extension Mechanisms for DNS (EDNS0) [22] option. This option allows a resolver to add a client's subnet to a query. Using the client's subnet instead of its IP address preserves privacy, reduces the load on nameservers, and allows for appropriate caching. Authoritative nameservers can use this subnet information to tailor responses to the client. Figure 1 shows the general functionality and information flow of ECS.

The ECS option consists of four properties: The IANA address family [38]; an address; and a source and scope prefix length. While the address and its family are self-explanatory, we delve deeper into the properties of the prefix length. The resolver selects the source prefix length based on its ability to cache responses without resource exhaustion. RFC7871 suggests a maximum value of 24 for IPv4 and 56 for IPv6 to preserve user privacy.

The authoritative nameserver uses the client information to compose its response. It sets the so-called scope length to the prefix length for which the answer is valid. The resolver must cache the response accordingly. A scope prefix length less specific or equal to the source prefix length indicates that the response is valid for this (less specific) prefix. A scope prefix length of 0 covers the complete address space. If the scope prefix length is more specific than the source prefix length, the authoritative nameserver indicates that the provided prefix length is not specific enough. The resolver can cache based on its capabilities (the initially sent source prefix length) or rerun the query with a more specific source prefix length. In Figure 1, the resolver chooses a source prefix length of 24, while the nameserver responds with a less specific scope prefix length of 20.

Terminology: In the remainder of this paper, we use the term *scoped response* to indicate a response that contains an ECS scope larger than zero. This response can only be cached for the prefix within the ECS option. We identify nameservers and domains exhibiting such behavior as ECS *enabled*/*supporting*. They are *using* ECS iff they have ECS *enabled and* if different Resource Record Sets (RRsets) are returned for distinct ECS prefixes.

3 Related Work

We identify three directions relevant to this research: The latency reduction achieved through ECS, previous active ECS scanning approaches, and the research use cases.

ECS Latency Reduction: Several research groups [15, 61, 66, 73] showed the potential of ECS to reduce end-to-end latency. In 2012, Otto et al. [61] performed an evaluation of ECS and found that 9 % of Alexa top 1k sites used the ECS option. They demonstrated that using the ECS option substantially benefited the client's end-to-end latency. The median improvement for services with Akamai was 40 %. In 2013, Sánchez et al. [66] confirmed these results with 90 k users on their measurement platform. They showed that users from Oceania benefited the most from using ECS. More recent research by Calder et al. [16] in 2019 evaluated the adoption of ECS at local resolvers by analyzing authoritative nameserver logs from the Azure cloud. 40 % of observed queries included the ECS option. These results show that services benefit from deploying ECS and that resolvers support it. For large providers, this benefit outweighs the client-server mapping complexity introduced by ECS. As a result of this previous research, we expect to find ECS deployments within top lists.

ECS Scanning: Streibelt et al. [73] and Calder et al. [15] performed active IPv4 ECS scans to uncover service infrastructure in 2013. Streibelt et al. found that 13 % of Alexa domains seemed to enable the ECS option, but only 3 % actually used it. Nevertheless, they showed the importance of ECS by analyzing an Internet Service Provider (ISP) network with more than 10 k end-users. In total, 30 % of traffic volume involves services using ECS. They used different prefix lists, including BGP dumps from RIPE RIS [65] and Routeviews [1] as target prefixes for their ECS scans.

Calder et al. [15] collected ECS scan data over ten months to map Google's serving infrastructure. Their scans to map Google's serving infrastructure used all routable /24 prefixes from Routeviews. According to their results, ECS scans found around 20 % more addresses than those collected from open resolvers. Moreover, their long-running scan shows that a single scan does not reveal all addresses, and only an accumulated dataset can provide a more complete view. Calder et al. use Google's own resolver as an intermediary to forward the ECS queries.

Kountouras et al. [52] evaluated the use of ECS at authoritative nameservers. They see the continuous growth of ECS usage in general and also for sites that do not profit from ECS. Kountouras et al. conduct a small-scale active measurement by issuing queries to recursive resolvers, which add ECS information to the iterative queries towards authoritative nameservers. This scan aimed to determine the support of ECS by domains using a single query. They determined that the support across the Alexa Top 1M domains increased from 161 k in 2015 to 418 k in 2019. They did not provide data on nameserver adoption.

Related to these efforts, we revisit and substantially extend results for a diverse set of top lists in this paper. We devise and publish a significantly better-performing scanning approach for IPv4 (see Section 5). To our knowledge, our ECS scanner is the first public tool specifically dedicated to perform ECS scans. Moreover, we introduce support for structural IPv6 ECS scans for the first time. We perform scans for popular domains and compare our results to those of previous work. Additionally, we perform infrastructure coverage validation scans using RIPE Atlas.

Research Use Cases: ECS also enables researchers to learn more about service deployments. Jiang et al. [47] use ECS queries to resolvers to determine if subnets contain end users. They validate their results using Microsoft CDN ground truth data. Streibelt et al. [73] and Calder et al. [15] map the responses to users and cluster them to better understand the service's load balancing behavior. ECS scans can also be used to uncover IPv6 offnets, as presented by Hilal et al. [33]. Sattler et al. [67] use full address space ECS scans to obtain the ingress nodes of iCloud Private Relay [40].

Calder et al. developed the *client-centric front-end geolocation (CCG)* approach to localize the front-end infrastructure via ECS scanning. It uses geolocation databases to localize the client prefixes that trigger the responses. In a follow-up work, Fan et al. [26] measured the affinity of users to front-end clusters by evaluating ECS scan data. By applying the *CCG* approach, they find that users are remapped to different clusters, often more than 1000 km apart. Warrior et al. [75] presented their client-side approach to find the service's front-end with the lowest latency based on ECS. They and others [3] bring up the issue of operators (*e.g.*, Akamai) restricting ECS access to pre-approved recursive resolvers. Their evaluation and interpretation show that unrestricted ECS adoption is substantially more beneficial compared to its drawbacks.

This list of use cases shows the importance of ECS scanning to research. Our dedicated scanner can help to promote new research due to its ease of use. Our approach offers a more efficient and ethical foundation for future research that relies on ECS scans. In this work, we explore the prevalence of restricted ECS in top lists and perform scans towards Google's open resolvers and authoritative nameservers in Section 7.1. While it is not the primary focus of this work, we shortly highlight potential use cases in Appendix C.

4 Response-Aware ECS Scanner

Previous work indicates significant support for ECS at nameservers. To better understand the ECS ecosystem, an efficient scanning technique is needed. Existing tools that can be used to send ECS queries, including dig [42] and ZDNS [79], are either not scalable (dig), or only allow static prefixes to be sent (dig and ZDNS). Therefore, we created a dedicated ECS scanner called *ECSplorer* [69]. Our scanner is a highly configurable stateful tool that also supports IPv6. We aimed to solve two concrete goals with our dedicated ECS scanner:

(i) Uncovering the service's infrastructure by collecting all available IP addresses.
(ii) Understanding ECS load balancing behavior.

While the first goal aims to use a small set of queries to find all possible addresses, the second seeks detailed information on all distinct ECS scopes the authoritative nameserver uses. Appendix C includes preliminary evaluations showing future potential for some ECS research use cases. We additionally aim for an efficient and ethical scanning approach to reach these goals. Our second goal, in particular, requires sending many queries. Therefore, we take extra caution to not perform unnecessary queries to lower the load on nameservers.

IPv4 Approach: Similar to related work [15, 73], we iterate over the BGP-announced address space and perform queries with all relevant subnets. We obtain the BGP-announced address space by extracting all announced prefixes from BGP dumps (*e.g.*, from RouteViews [1]). Our scanner supports arbitrary source prefix lengths. In this work, we choose a source prefix length of 24 bits. In contrast to previous work, our approach does not need to scan all announced /24 IPv4 prefixes. Instead, it respects the scope prefix length by the authoritative nameserver and skips the covered address space (*e.g.*, in Figure 1 we would stop querying the full /20).

Moreover, a per-prefix length threshold can be provided to limit the number of queries for each prefix of the configured length (see Appendix B). Our scanner does not increase the specificity of the source prefix length over the configurable value. Thus, more specific scoped responses (*e.g.*, a scope of 28 in Figure 1) would not trigger more specific queries when the source prefix length is set to 24. The scanner can also be configured to perform queries for a limited number of prefixes in unrouted and RFC1918 [60] address space.

In order to efficiently implement this approach, we build a Patricia trie [50] mapping out the BGP-announced address space. Each node contains all relevant information for the scanner (*e.g.*, number of scans). This way, the scanner can walk through the address space to compute recursively if any prefix threshold has been reached (see Appendix B). We use Go and its parallelization features to implement our approach efficiently.

To overcome nameserver misbehavior, such as eventual /0 scope prefix lengths, we limit the accepted scope prefix length to a minimum of /8, *i.e.*, the maximum assigned prefix length from RIRs to Autonomous Systems (ASes). Therefore, we ignore answers with a scope prefix length less specific than /8 when deciding which subnet to scan next. This property also enables us to parallelize domains per /8, further increasing the scan performance. In reality, we never had to parallelize a single domain with more than four /8 scans to stay within our query limits (see Appendix B). Our implementation runs on low-end hardware, and we also ran it from a virtual machine with 2 cores and 2GB of RAM without any issues.

IPv6 Approach: To account for the vast IPv6 address space, we extend our scanning approach with IPv6-tailored features. We aim to cover as much of the BGP-announced address space as possible and in as much detail as possible. While IPv4 has approximately 12 M routable /24 prefixes, IPv6 has 15 billion /48 prefixes. RFC7871 [21] even suggests using a maximum of 56 for the scope prefix length, resulting in 3.8 trillion possible prefixes. Sending over a thousand times more queries compared to IPv4 is neither reasonable nor ethical. Hence, we choose a source prefix length of 48 and seed our scan with a predefined list of prefixes to limit the address space. This seed list limits the address space to scan for, and the scanner will select random subnets inside this space. The scanner is still response-aware and does not issue queries for subnets that are inside a returned scope. The scan completes when there is no more seeded address space to be scanned within the defined limits.

We also added an option to scan every prefix from the seed list at least once, *e.g.*, to scan all BGP-announced prefixes. This option ensures that at least one query is sent for each announced prefix, even if prefix limits have already been reached. Other possible seed list sources are the Internet Routing Registries (IRRs) or the IPv6 Hitlist [27, 82]. We evaluate and compare the lists in Section 5.2. Like regular port scanning in IPv6, selecting the target list is important for IPv6 ECS scans (see Section 5).

Our scanner is also capable of only issuing queries for a predefined list of ECS subnet values. This feature can be especially useful to scan a limited set of IPv6 subnets that are deemed to be representative of the address space that will be analyzed. A static input list can be used to obtain comparable results over time.

Ethics: At the core of our approach are ethical considerations to reduce the number of queries while obtaining the same information. We do not process any user data, and we apply strict rate limiting to our active scans. A detailed description of our measures can be found in Appendix A.

5 Evaluation of Response-Aware Scanning Technique

In this section, we analyze the impact of our approach on the number of issued queries and observed responses. We select nine top-listed domains to perform full IPv4 address space scans (no limits applied) using our stateful scanning approach. Additionally, we evaluate different IPv6 address space scanning approaches and compare our implementation with existing techniques. Finally, we evaluate the returned scope prefixes to gain insights into different rates of query savings.

5.1 IPv4 Results

For our evaluation, we analyze a scan from March 20, 2024 towards the authoritative nameservers of nine popular domains supporting ECS (see Table 1). The nameservers of these domains are provided by Google, Meta, Amazon Web Services (AWS), and Wikimedia. While the domains associated with Google, Meta, and Wikimedia are served by nameservers of the sites own organization, AWS's Route53 also provides DNS service for tiktok.com. Not only TikTok but any AWS customer can use Route53 to provide subnet-scoped ECS responses [8]. Cloudflare's customers can use a similar feature when using Cloudflare load balancers [20].

Table 1. Overview of IPv4 response-aware address space scans. An RRset represents the combination of resource records within each answer. The ECS scopes column contains the observed number of distinct ECS response scope lengths. tiktok.com is not hosted on AWS, but uses Route53 for DNS services.

	Queries	RRsets	Addresses	ECS Scopes	NS Operator
en.wikipedia.org	0.4 M	5	5	26	Wikimedia
m.facebook.com	1.2 M	139	139	17	Meta
web.whatsapp.com	1.2 M	139	139	17	
www.instagram.com	1.2 M	139	139	17	
www.google.com	1.7 M	32 531	2195	26	Google
www.youtube.com	1.7 M	322 534	1839	26	
tiktok.com	11.8 M	3687	636	1	AWS Route53
www.amazon.com	11.8 M	323	323	1	
www.primevideo.com	11.8 M	161	161	1	

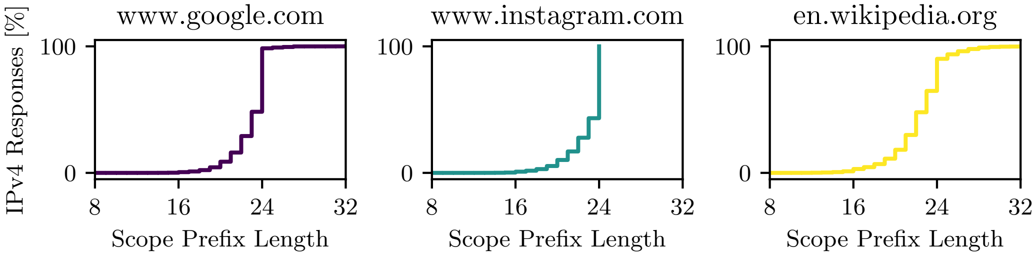

Fig. 2. Scope prefix lengths from our IPv4 response-aware scans. AWS always uses /24. See Figure 5 for IPv6.

Scan Volume Analysis: Table 1 shows the number of issued queries per domain. We can clearly see the uniform behavior of nameservers for each operator. Our approach needs nearly the same number of queries for all domains hosted by the same operator, even though RRsets and uncovered addresses differ. Our scanner issues the fewest queries (0.4 M) towards Wikimedia's nameservers. According to Wikimedia's documentation [77], the MaxMind geolocation database [57] and the gDNS geoip plugin [29] are used to deploy ECS. This relatively low number of queries is confirmed by Wikimedia only using six different sites [76] at the time of writing, for which less fine-grained results are needed. While we see all six Wikimedia data centers with other measurements, the Carrollton, Texas one was out of service [78] during this scan.

With Meta and Google, our scanner sends 1.2 M and 1.7 M queries to cover the full address space. While Meta responds with a single resource record per response and provides exactly 139 addresses for each domain, Google does respond with multiple resource records and has different address mappings for Google and YouTube. Conversely, AWS nameservers always respond with a scope prefix length of 24. Therefore, our approach is forced to scan each announced /24 prefix while ignoring IANA special-purpose [37] and non-announced prefixes.

TikTok always returns four records per RRset, and we observe 3687 RRsets. Google and YouTube exhibit a substantially larger number of RRsets than the number of distinct addresses found (*cf.*, Table 1). While Google uses one or six address records in its responses, we observe answers with everything between one and 16 records for YouTube. Therefore, the number of possible combinations leads to this large number of observed RRsets. Our second goal is to understand the load balancing strategy of nameservers. The number of combinations hints at the complexity of the mapping used by the nameserver. While all other operators stick to a 1:1 mapping between vantage points and client subnets, Google, YouTube, and TikTok use a 1:n mapping.

Query Savings: As AWS is always responding with a scope prefix length of 24, we can use AWS's query volume to determine the query savings our response-aware approach eliminated in comparison to scanning all /24 prefixes as done by Calder et al. [15] or Streibelt et al. [73]. Our approach eliminates 97 % of queries for Wikimedia, 90 % for Meta, and 86 % for Google. These numbers demonstrate the importance of using a response-aware scanner to eliminate a large part of superfluous DNS queries.

Next, we analyze the different scope prefix lengths returned for Google, Instagram (Meta), and Wikipedia (Wikimedia). As shown in Figure 2, Wikimedia returns more answers with lower scope prefix lengths than Google and Meta. 64 % of Wikimedia responses are less specific than a /24. Google responds in 48 % of cases with a prefix length less specific than 24, and Meta does so in 43 % of cases. It follows from these results that Wikipedia needs the fewest, and Meta has the largest number of queries.

While Meta does not use any scopes more specific than /24 prefixes, both, Google and Wikimedia do so. RFC7871 [21] does state that recursive resolvers should use a maximum prefix length of 24 to preserve user privacy. These more specific responses do not affect resolvers adhering to RFC suggestions (*e.g.*, Google Public DNS and the OpenDNS) as they cache responses only for the full /24. Google's own resolver also refuses to forward recursive queries with ECS scopes larger than 24, responding with an extended DNS error [53] pointing to their guidelines [30].

5.2 IPv6 Results

As described in Section 4, we cannot simply perform a full address space scan in IPv6 similar to IPv4. Instead, we use a prefix-seeded randomization approach with a scope prefix length of 48. Similar to IPv4, we use the prefixes announced in BGP as seed list. Additionally, we compare the BGP seed list to lists from IRR and /56 prefixes from the IPv6 Hitlist [27, 82]. To evaluate the usefulness of such an approach, we compare it to scanning the different static lists of prefixes at the same time: BGP prefixes obtained from a local border router on April 26, 2024 (189 k), prefixes in `route6` and `inet6num` objects from all RIRs' IRR databases on April 26, 2024 (1.3 M), and all /56 prefixes from the responsive IPv6 Hitlist address set from April 1, 2024 (6.0 M). We conduct our scans with a single list of prefixes combined from all sources but evaluate the results for each prefix list source individually in the following. The complete list of prefixes contains 7.0 M distinct IPv6 prefixes used as ECS subnet parameters in our scan.

Scan Approach Comparison: Table 2 compares the response-aware IPv6 scanning approach with the static prefix list approach. Due to ethical considerations, we apply query thresholds (see Appendix B) to limit the number of queries in the worst case to the same number as with the full IPv4 scan. In contrast to IPv4, our scanner executes more queries for Meta domains than Google domains. While the number of queries for Wikipedia also increases, our limits prevent the scan for Amazon from having as large a number of queries as for IPv4. With the help of these limits, we keep the number of queries well below the ones performed by the prefix list scan.

Meta returns 137 addresses, and both approaches uncover the complete set. We find 1.8 k and 1.6 k addresses for Google and YouTube, respectively, with fewer requests compared to Meta domains with our response-aware approach. However, for Google domains, the approach with the static prefix list reveals 12.9 % more IPv6 addresses but also requires more than five times as many queries.

As already seen in the IPv4 analysis, Google and YouTube also use multiple addresses inside a single response for `AAAA` queries. Only Amazon changes its behavior from a 1:1 mapping in IPv4 to a 1:n mapping in IPv6. We compare the efficiency of the two main scan approaches by evaluating

Table 2. Comparison of the response-aware IPv6 approach with the prefix list (7.0 M) comprised of BGP (189 k), IRR (1.3 M), and the IPv6 Hitlist (6.0 M) prefixes. We shorten prefixes to /56 where necessary. `www.primevideo.com` and `tiktok.com` do not return AAAA records.

	Resp.-aware		BGP		IRR		IPv6 Hitlist		Comb. Lists	
	#Qs	#Addrs	#Qs	#Addrs	#Qs	#Addrs	#Qs	#Addrs	#Qs	#Addrs
en.wikipedia.org	1.3 M	7	189 k	6	1.3 M	6	6.0 M	6	7.0 M	6
m.facebook.com	1.6 M	137	—"—	137	—"—	137	—"—	137	—"—	137
web.whatsapp.com	1.6 M	137	—"—	137	—"—	137	—"—	137	—"—	137
www.instagram.com	1.6 M	137	—"—	137	—"—	137	—"—	137	—"—	137
www.google.com	1.3 M	1823	—"—	1803	—"—	1739	—"—	2057	—"—	2058
www.youtube.com	1.3 M	1587	—"—	1556	—"—	1505	—"—	1690	—"—	1691
www.amazon.com	1.8 M	23 266	—"—	23 610	—"—	25 917	—"—	28 756	—"—	28 973

the number of queries needed to uncover a new RRset (*i.e.*, RRsets per query). The response-aware approach needs only 25 % of queries per RRset compared to the prefix list scan for Google and 33 % for YouTube, indicating a query reduction of 3× and 2×, respectively. In absolute terms, the response-aware approach uncovers more RRsets than the BGP and the IRR list. Only the IPv6 Hitlist prefixes uncover more RRsets, at the cost of using 262 % (4.7 M) more queries.

In contrast to the RRset results, which are an indicator of the usability of our scan for our load balancing goal, the infrastructure uncovering goal focuses on the number of distinct addresses obtained. Therefore, we first analyze the different prefix sources in Table 2 and find that the IPv6 Hitlist alone finds all but one address, but it is also the largest source of prefixes (6 M prefixes). The IRR prefix list has more ECS target prefixes than the BGP list but reveals fewer server addresses for both Google domains. The BGP prefix list (189 k) still identifies 97 % of addresses collected by the response-aware approach and 88 % of all addresses with only 3 % of queries. Therefore, we analyze the addresses unique to the IPv6 Hitlist prefix queries to understand the impact and whether we can create a limited set of ECS prefixes to uncover the domain infrastructure efficiently. Moreover, such an improved prefix set can be used as a seed set for the response-aware approach.

Google: 100 % of the `www.google.com` server addresses unique to the IPv6 Hitlist prefixes are caused by prefixes announced by Cloudflare. With `www.youtube.com`, these Cloudflare prefixes cover 72 % of the unique addresses in responses. The other 28 % of prefixes are inside Google's own AS396982. According to WHOIS information, the addresses in this AS are used by Google Cloud customers. A closer look at the Cloudflare prefixes reveals that the IPv6 Hitlist contains 30 k /56 prefixes inside `2a09:bac2::/32` and `2a09:bac3::/32`. Further investigation reveals that the hitlist addresses inside these prefixes are iCloud Private Relay [40] egress nodes operated by Cloudflare. iCloud Private Relay is a two-hop relay service offered by Apple to protect the user's privacy in the network. The egress nodes are operated by third parties such as Cloudflare. iCloud Private Relay uses these addresses for traffic on the path between the service's egress and the destination servers. Apple publishes a list of all egress ranges used by iCloud Private Relay as a geolocation feed [5]. Therefore, Google seems to use Apple's geolocation feed to optimize traffic for iCloud Private Relay users by sending dedicated responses to them.

Consequently, we employ a new list incorporating other IP geolocation feeds (geofeeds) [48, 54]. We gather feeds using the *geofeed-finder* [56], which leverages RFC9092 [12] to identify geofeeds. The initial test with the combined set (189 k BGP, 26 k iCloud Private Relay, 48 k geofeed prefixes trimmed to /56) yields the same set of Google server addresses as the combined prefix list scan.

Therefore, combining BGP, iCloud Private Relay, and geofeed prefixes (245 k) proves highly effective for ECS scans. Geofeeds, a relatively recent concept, could become a more valuable resource for future ECS scans.

Amazon: The prefix list scan reveals 24.5 % more addresses than our response-aware approach for Amazon. 28.2 k of the returned IPv6 addresses are within the fully responsive prefix (FRP) dataset provided by the IPv6 Hitlist service [28]. The remaining addresses are in seven /48 prefixes and are not part of the IPv6 Hitlist. AWS's encompassing /32 prefix contains 1.2 M FRPs of varying length. We conduct an FRP detection for these additional prefixes based on the approach by Gasser et al. [27]. Five out of the seven /48 are identified as such. We could likely find more specific prefixes for the remaining two, which would also be fully responsive based on the existing data.

Scope Prefix Lengths: While Amazon uses a static scope prefix length of 48, 51 % of Meta's responses use a maximum scope prefix length of 32 bits (see Appendix Figure 5). Google and Wikipedia have 62 % and 92 % of responses scoped at 32 bits length or more specific. This result shows that Meta provides much more fine-grained responses in IPv6 than the other two and, therefore, needs the largest number of queries from the three (*cf.*, Table 2). Our scanner can be configured to perform fine-grained queries as necessary to the respective use case. In our evaluation, we want to provide an overview of different deployments and, therefore, decide not to increase the prefix limits. Future work can adapt these limits to obtain even more relevant data. Similar to the observations for IPv4, we find that Google and Wikipedia return scope prefix lengths larger than 56, while Meta does not. Google even returns scope prefix lengths of 128, and—again similar to IPv4—its own resolver refuses query scopes that are more specific than /56.

5.3 Key Takeaways

Our scanning approach reduces the number of queries needed for a full IPv4 address space scan by 86 % to 97 %, depending on the domain. Reducing queries lessens the burden on the nameservers and enables timelier follow-up actions (*e.g.*, application layer scans). In IPv6, we find an efficient prefix list consisting of BGP prefixes and prefixes extracted from geofeeds to uncover all observed addresses. The analysis of the response scope prefix lengths shows that we do not have a snapshot of all relevant client-to-vantage point pairs, even if we obtain all observed distinct addresses. Our scanner can use the prefix list as a seed to obtain in-detail insights into the nameserver's mappings. Moreover, its efficiency in analyzing load balancing behavior is better than that of the prefix list.

6 Evaluation of the ECS Landscape

To confirm previous work and expand the scope of their analysis, we analyze the usage of ECS in the wild. This analysis helps us to understand the prevalence of nameservers and domains supporting ECS and what researchers can expect when using our tool. We perform and evaluate two types of scans: First, we send four queries per IP protocol version to all authoritative nameservers of domains in the Cloudflare Radar [18], CrUX [31], Majestic Million [55], and Umbrella [17] domain lists (Section 6.1). Second, we conduct a more in-depth scan with client prefixes from all countries and a diverse set of ASes to get a more detailed view of ECS-enabled domains.

6.1 ECS Usage on the Internet

During this first scan, we send DNS queries for each domain to their authoritative nameservers. We send individual queries with ECS information for four IPv4 prefixes and four IPv6 prefixes. We select a diverse set of prefixes based on location and topology (see Appendix B). Response-aware address space scans, as evaluated in Section 5, show that this combination provides the best coverage of locations and corresponding distinct responses.

Table 3. Top 10 nameserver ASes for `A`/`AAAA` queries. For CNAMEs, we follow the chain and query the resulting domain. Since domains can share the same canonical name, the set of queried domains, *cDomains* is smaller. Shares are relative to Total. Domains can have nameservers in multiple ASes (shares do not add up to 100 %)

		IPv4 ECS Family				IPv6 ECS Family			
Name	ASN	#NS	cDomain	Domain	ECS used	#NS	cDomain	Domain	ECS used
Cloudflare	13335	8024	73.2 %	76.9 %	2.2 %	7400	86.4 %	89.9 %	3.9 %
Amazon	16509	4188	10.5 %	9.7 %	64.9 %	4092	4.4 %	3.9 %	53.2 %
China Mobile	9808	109	3.3 %	3.5 %	0.4 %	67	0.1 %	0.0 %	0.6 %
Google	15169	48	5.0 %	3.5 %	9.2 %	32	6.7 %	4.7 %	21.6 %
Kurun.com	8796	28	3.2 %	3.4 %	0.0 %	–	–	–	–
Chinanet	134763	25	3.2 %	3.4 %	0.0 %	–	–	–	–
Alibaba	37963	238	2.3 %	2.3 %	2.5 %	158	0.2 %	0.2 %	2.0 %
Alibaba	45102	86	1.3 %	1.1 %	1.9 %	42	0.2 %	0.1 %	0.7 %
Alibaba	134963	24	0.9 %	0.8 %	0.5 %	18	0.1 %	0.1 %	0.1 %
Incapsula	19551	124	0.8 %	0.7 %	5.1 %	–	–	–	–
Total		15.8 k	1.1 M	1.2 M	184 k	13.0 k	794 k	853 k	64 k

Scanning Setup: ECS only provides value to domains hosted on distributed deployments, so we focus on domains in top lists. These popular domains usually require some sort of load balancing to handle the received load. With this measurement, we want to understand how many domains and nameservers enable and use ECS. The selected lists use different collection approaches (DNS-based for Cloudflare Radar and Umbrella; web statistics/user-based for Majestic and CrUX). Therefore, they cover a large set of relevant domains. The four domain lists have 3 M combined unique domain names with resolvable authoritative nameservers.

We limit our scan to `A` queries with IPv4 client subnets and `AAAA` queries with IPv6 clients. If the domain name resolves to a `CNAME`, we follow the chain and perform our ECS queries towards the authoritative nameserver of the resulting domain name. Multiple domains of our input set can resolve to the same canonical name, *e.g.*, domains hosted by Cloudfront. For those domains, load balancing based on ECS is done for the A or AAAA resolution of the canonical name. Thus, we test the behavior for those domains, reducing our set of targeted domains to 2.7 M domain names. In the following evaluation, we refer to the original input set based on the top lists as *Domains* while referring to the scanned set of domains with resolved canonical names as *cDomains*.

ECS Support in the Internet: The authoritative nameservers of the analyzed domain names are distributed across 196 k IPv4 and 39 k IPv6 addresses. These nameserver IP addresses are located in 21 k ASes (20.8 k IPv4 and 5.1 k IPv6). These numbers show that our target set is diversified over many ASes.

1.1 M *cDomains* (41 %, *cf.*, Table 3) have ECS enabled on their authoritative nameserver. These *cDomains* are used by 1.2 M domains from our initial input set. This result is similar to what Kountouras et al. [52] found in 2019. While we receive relatively fewer successful responses to our IPv6 ECS queries, our scan still reveals 0.8 M *cDomains* (30 %) supporting ECS. In total, only 8.9 k IPv4 (449 ASes) and 6.9 k IPv6 (140 ASes) nameserver addresses (15.8 k in total) are authoritative for these ECS-enabled *cDomains*. These are only 5 % and 18 % of queried IPv4 and IPv6 nameservers, respectively.

Only 6 % of domains (2 % for ECS with IPv6) respond with at least two different RRsets (see Table 3). This is 15 % and 11 % of ECS-enabled domains for IPv4 and IPv6, respectively. We use this as an indicator of whether ECS is actually used. Deploying ECS has preconditions (*e.g.*, a distributed setup) and requires substantial effort.

Nameserver Operators: Table 3 lists the 10 ASes containing the nameservers authoritative for the largest number of domains that provide ECS support. Cloudflare serves 76.9 % of all domains supporting ECS over IPv4 and 89.9 % over IPv6. However, only a small share of these domains result in two different answers for the four queries per IP version. To better understand how nameservers use ECS, we perform a detailed analysis of more specific scans in Section 6.2. Kurun.com, Chinanet, and Incapsula are the only ASes in the top ten that do not support ECS with AAAA queries.

By a substantial margin, AWS has the largest share (65 %) of domains using ECS with more than one answer. AWS provides a simple configuration mechanism that allows its users to configure the exact answers per client subnet [9]. This simple and fully customizable configuration, together with the primary cloud products by AWS, might be an important factor for the large number of ECS-using domains with AWS. The AS responsible for the second-largest share of such domains is Google, with only 9 % of domains, followed by Incapsula with 5 %.

43 % of IPv4 nameservers and 53 % of IPv6 nameservers enable and use ECS only for select domains. 55 % of the responsible nameservers are within the prefixes `205.251.192.0/21` and `2600:9000:5300::/45`. Both prefixes are owned by AWS and used for Route53 [10]. They are identified as being fully responsive [71, 82], *i.e.*, all usable addresses inside these prefixes are responsive. The second largest AS is Cloudflare, covering an additional 27 %. It is reasonable for these providers to enable and use ECS only for select domains as they charge for the service, and not every domain owner uses this feature.

6.2 Scanning with More Client Subnets

To better understand how nameservers make use of ECS and to extend our previous scan to identify domains and nameservers generally supporting ECS, we perform a scan with more prefixes. Due to the large number of domains, we cannot perform a full address space scan for each. If available, we select three prefixes per ISO 3166 country code [45] to obtain a geographically and topologically diverse set of prefixes. We use data provided by IPinfo [44] to geolocate all routed IPv4 /24s by their network address. For IPv6 we used /48s from our seed lists (BGP, IRR, and IPv6 Hitlist [27, 82]). Our algorithm also ensures never to select prefixes originated by the same AS. These two decisions provide us with a geographically and topologically diverse sample set to efficiently obtain an overview of how nameservers use ECS without putting too much load on them. This selection algorithm results in 698 IPv4 and 609 IPv6 prefixes.

Table 3 shows that only 184 k out of 1.2 M domains (for IPv6 64 k out of 853 k), thus less than 10 %, use the ECS option (*i.e.*, we observe more than a single RRset). In contrast to the scans in Section 6.1, in this scan, we use 698 prefixes to determine if nameservers actually return more than one RRset. We perform this additional scan for all domains that appeared to use ECS during our initial scan. Additionally, we scanned a sample of 40 k ECS enabled but not using domains. We focus our evaluation on IPv4 results as IPv6 results are similar.

We argue that if we only observe a single RRset for this larger ECS prefix set for a given domain, the domain is not using the ECS option in a way to the benefit the user. Moreover, the resolver must also provide needless caching space for these domains. Therefore, the user ends up losing its privacy to the authoritative nameserver without any gains while the resolver's caches suffer.

ECS-enabled Analysis: In Figure 3, we show the cumulative distribution of domains and their corresponding number of observed RRsets. For 37 % of the sampled ECS-enabled domains, we still observe only a single set of returned records. 66 % of these domains have a Cloudflare authoritative nameserver and 10 % a Google nameserver. This seems counterintuitive at first glance, as ECS load balancing is a paid service according to the documentation [20], but the domains seem to not make use of it. We use a free tier Cloudflare account with a registered domain name and find that Cloudflare's authoritative nameservers return scoped answers even though we did not enable it.

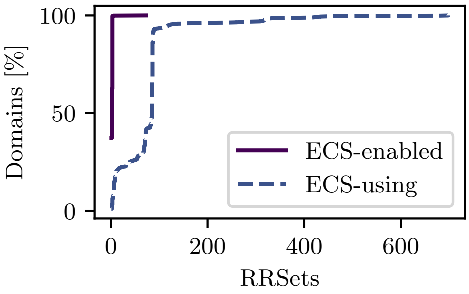

Fig. 3. ECDF for the number of observed RRsets in the country prefix list scan.

Since our domain cannot actually make *use* of this feature without paying for it, our domain does not benefit from this behavior. When considering this Cloudflare specialty, we find that 99.7 % of sampled ECS-enabled domains are not *using* ECS.

ECS-using Analysis: More than 99 % of domains that appeared to use ECS in our initial scan also did so in this follow-up scan (see Figure 3). Cloudflare's behavior of few distinct responses is present but only marginally visible (less than 1 %). 20 % of domains have ten or fewer distinct record sets. We observe two significant jumps: The first and smaller one contains 12 % of domains with 67 to 73 RRsets, and the second one between 81 and 88 sets covers 50 % of all scanned domains. AWS authoritative nameservers cause both. The records also point to addresses inside the AWS address space. This leads us to the conclusion that these are two different default AWS configurations to perform global DNS-based load balancing using the Amazon CloudFront CDN.

Besides the set of ECS-enabled domains with only a single response, Cloudflare is also responsible for 99 % of domains with two or three distinct RRsets. Interestingly, most domains with a single RRset do not resolve to Cloudflare IP addresses, while Cloudflare announces more than 99 % of the addresses from domains with multiple RRsets. This could indicate that Cloudflare actively uses ECS for domains hosted within their network. However, according to data by IPinfo [43], except for one, all IP addresses returned by Cloudflare are inside IP anycast prefixes. Furthermore, our scan reveals only three RRsets, while Cloudflare has more than 300 global PoPs [19]. Therefore, we assume that domains with two and three responses are also not actually based on ECS information. Instead, we assume that Cloudflare is balancing our requests using a strategy that does not rely on ECS. The observed ECS behavior would then be similar to the ECS-enabled domains. However, this behavior can fill up the cache of ECS-using recursive resolvers, and it allows Cloudflare to collect data on all request-originating clients' subnets of such a resolver.

6.3 Key Takeaways

Our evaluation finds that 40 % of the analyzed top-list domains have ECS enabled for IPv4 subnets on their authoritative nameservers, but these domains are served by only 5 % of queried nameserver IP addresses. The ECS deployment is highly skewed towards large DNS providers. Scanning with four prefixes revealed that only 15 % of these domains actually use ECS. A scan with 698 ECS prefixes validated these results and identified Cloudflare as the major player responsible for the delta between domains with ECS enabled and those using it. Cloudflare enables ECS for domains even if they do not pay for the service, but it will not tailor the response records for these. This result shows that a single query is not representative of determining if ECS is actually used. According to our results, four prefixes are sufficient to precisely categorize ECS-using domains.

7 Validating ECS Properties of Nameservers

Our results show that ECS is used by a considerable number (184 k) of domains. Our scanning approach helps collect addresses and understand the domains load balancing behavior. As part of this section, we examine important properties of the domains authoritative nameservers. These findings should be considered when performing ECS scans and analyzing its results.

7.1 Public Resolvers vs Authoritative Nameservers

Related work [3, 75] mentioned that some authoritative nameservers provide ECS responses only to a restricted set of resolvers. We re-evaluate this statement to validate if we can directly query authoritative nameservers instead of resolvers. This would have the ethical positive side effect of not unnecessarily filling up the public resolvers' caches. As we aim to find the most ethical scanning approach, we evaluate the benefits and drawbacks of using public resolvers. Therefore, we issue the same queries to Google's public resolver and authoritative nameservers. We chose Google's resolver as it is one of the most well-known ECS-supporting public resolvers, and it was listed as one of the few with access to restricted deployments [75]. Thus, if a service deploys a restricted ECS setup to its authoritative nameservers, it will most likely enable it for Google's resolver.

We limit this scan to a single IPv4 ECS prefix and to domain names in Cloudflare Radar and CrUX (2.1 M domains in total) to reduce the load on the resolver. A single query is sufficient to observe if ECS responses of certain domains are restricted to an allowed set of public resolvers. 824 k domains (40 %) return a scoped answer either on the authoritative nameserver scan or the recusive resolver scan or both. 2.2 k (0.3 %) of these ECS-enabled domains provide ECS information only to Google's public resolver and not to our direct queries to their authoritative nameservers. The Soprado GmbH (AS20546) is the most important provider, with 230 domains (10 %) returning ECS scopes only to queries from Google Public DNS. On the other hand, for 5.2 k (0.6 %) domains, our scanner reports scoped answers while the Google resolver does not. Cloudflare's authoritative nameservers are responsible for 3.9 k of these domains with missing ECS information in the Google resolver's responses. The information is either not added by the nameservers or removed by the resolver. Google's guidelines list several rules a nameserver has to follow for Google to send ECS queries (*e.g.*, overlapping ECS responses, as analyzed in Section 7.4). This evaluation shows that the prevalence of restricted ECS deployments is no longer a widespread phenomenon.

However, we wish to emphasize a scenario where recursive resolvers prove more advantageous as targets: `CNAME` chain resolution. Unlike our scanner, which issues specific queries but does not perform resolution tasks required for `CNAME` records[1], a recursive resolver can resolve `CNAME` chains. While this ensures receipt of the desired record type for the domain name, it also relinquishes control over the resolution path, leaving us unaware of the authoritative nameservers employed. For `CNAME` resolution needs, one could opt for a local recursive resolver on the scan machine as an alternative to public resolvers.

In summary, the prevalence of restricted ECS deployments is negligible (0.3 %). Therefore, the necessity of using public resolvers is generally not given anymore. Directly querying the authoritative nameserver reduces the impact on and by resolver caches. Nevertheless, the use of recursive resolvers should be considered when `CNAME` resolutions are relevant.

---

[1] In addition to our ECS scanner *ECSplorer* [69], we release a tool to perform `CNAME` resolution and obtain the authoritative nameservers of the domain to query. The tool produces an input list for the ECS scanner to use. The drawback compared to a recursive resolver is that this resolution is not done as part of each ECS query, and we observe rare cases where nameservers change behavior during scanning periods.

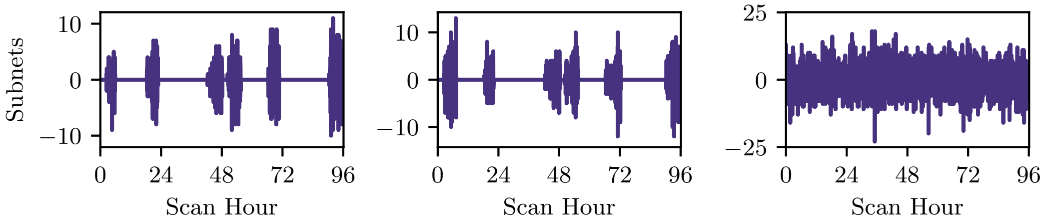

(a) www.google.com (b) www.youtube.com (c) www.instagram.com (Meta)

Fig. 4. Scope prefix changes observed in 96 h. Positive values are changes to more specific responses and negative ones to less specific. The scan started on Friday, May 10, 2024, at 15:45 UTC, at our vantage point in Germany. See Figure 6 for IPv6 plots.

7.2 Impact of Scanning Location

To validate whether a single vantage point is representative to uncover a service's infrastructure (the first goal of our approach), we perform parallel scans from four additional vantage points. While our main scanning location is in an academic network in Western Europe (Germany), we added another academic one in Europe (France) and three vantage points from a cloud provider in Frankfurt - Germany, Singapore, and Newark, New Jersey - United States. We start the scans at the same time, ensuring the same query order.

Top List Scan Impact: Due to the larger number of concurrent queries, we limit the scan to Cloudflare Radar domains and four IPv4 ECS prefixes. 380 k Cloudflare Radar domains have ECS enabled. For 31 % of these, we observe at least two different RRsets at different vantage points. 98 % of these domains with multiple RRsets are served by Cloudflare (AS13335), AWS Route53 (AS16509), or Meta (AS32934) authoritative nameservers. The responses for these names contain different nameserver identifier (NSID) [7] option values at different vantage points[2], indicating that we hit different instances of an IP anycast deployment. Nevertheless, this affects only 7 % of all domains using an AWS Route53 nameserver and 34 % of domains served by Cloudflare nameservers. Cloudflare exhibits a unique behavior: while it provides scoped answers, the response varies based on the requester's IP address and not the client subnet. This finding is analogous to Section 6.2, where we scan from a single vantage point, and most domains managed by Cloudflare do not effectively use ECS, always responding with the same RRset.

To evaluate this hypothesis, we use CAIDA's Archipelago Globally Distributed Active Measurement Platform (Ark) [13], which offers hundreds of vantage points around the globe and has the ability to send ECS queries [14]. We reimplemented our scanning approach [68] in Python to make use of Ark. The validation scan uses 130 Ark vantage points to issue a predefined DNS ECS query from all vantage points towards the same authoritative nameserver IP address. Across a set of 30 different subnets in the ECS extension, Cloudflare selects the same IP address for a specific vantage point disregarding the actual ECS information, while claiming the response is only valid for the specific ECS subnet. This behavior is only visible for domains which are actually load balanced via DNS. These results confirm our hypothesis on Cloudflare's behavior of ignoring ECS information.

Address Space Scan Impact: We use a distributed full IPv4 address space scan (same methodology as in Section 5) to evaluate the prevalence of mismatches within such a scan. The answers from

---

[2] For Meta, we decoded the Base64-encoded value to extract the cluster value. This is necessary as Meta nameservers encode a timestamp and the ECS value inside the NSID option value.

Google's nameservers match for 97 % of all queries at all locations. For Meta domains, 81 % of queries, and for Amazon, only 68 % of queries return an equivalent response. Disregarding the concrete responses to queries, we uncover the same aggregated set of server IP addresses at all locations for Google and Meta. Amazon provides some IP addresses only to certain locations. Nonetheless, we find that the reverse DNS names of the locations point to the same set of PoPs (see also Appendix C.1). As Amazon is known to have a distinctive use of its address space [28] to balance the load, we attribute our observation to this behavior.

Infrastructure Coverage Validation: We issue RIPE Atlas A record DNS measurements using all available probes to obtain results from up to 11.6 k vantage points. We schedule the measurements to run simultaneously as our response-aware scan. Before evaluating the results, we remove any bogus answers. The measurements from probes located in China return Meta-owned addresses. The behavior results from the Great Firewall of China (GFW) injecting forged answers for queries carrying censored domains [34, 72]. Therefore, we do not count these as real answers by the authoritative nameserver and disregard them in our comparison.

Amazon answers have already shown a distinctive load balancing behavior depending on the location. We validate Wikipedia answers using their published dataset [76]. Thus, we limit this validation to Google and Meta. For Meta, we find 135 IPv4 and IPv6 addresses, all of which are seen in our response-aware scan. The measurements for Google and YouTube collect 1.2 k addresses for each domain for both IPv4 and IPv6. Therefore, we covered 33 % fewer addresses than our response-aware scan while not contributing any previously unknown addresses.

In summary, our results show that the two most important ASes do not provide stable responses when evaluating a distributed scan. However, this only affects 31 % of ECS-using domain names. Our full address space analysis also shows differences in the responses depending on the scanning location. Nevertheless, from an infrastructure coverage view it finds as many addresses as actual distributed scans do.

7.3 Stability of Responses

As previously established, the purpose of ECS is to balance the load based on client prefixes. Load balancing is typically a dynamic process involving several aspects, making it a complex system to understand. To validate how often responses change, we evaluate the load balancing update behavior by zooming in on changes in the returned scope prefix length. We use the country prefix list to better understand when and how the scopes change and run ECS measurements every three minutes for 96 hours. We limit this scan to Google, YouTube, and Instagram (Meta). Meta behaves consistently the same way for all three observed domains, while Google and YouTube do not.

In Figure 4, we visualize the scope changes over time. While Meta changes at least some scope prefix length at every measurement, Google and YouTube have a very distinctive behavior. Both exhibit changes only in specific time windows. We find a daily pattern that is approximately four hours before the end of the measurement day and lasts until its end (*e.g.*, just before the 24th scan hour). Two additional time frames start right after a new measurement day (after hour 0 and 48). Neither the ECS prefix's location nor topological properties seem to be relevant for a prefix to change the returned scope. While we leave a detailed analysis of the load balancing behavior to future work, this evaluation shows that our data reveals new load balancing aspects.

7.4 Overlapping Subnets

RFC7871 [21] states that authoritative nameservers must not overlap prefixes. Google's resolver rules [30] also include this requirement and its resolvers will stop using ECS with nameservers not adhering to it. Overlapping responses are also counterintuitive as the less specific answer can overshadow the more specific one, resulting in inconsistent states between resolvers.

Table 4. Overlapping ECS: The response for query 2 can be cached for the /23 covering the first response.

	Domain	ECS Address	Source	Scope	Response
1	`example.com`	`10.0.0.0`	24	24	`192.0.2.1`
2	`example.com`	`10.0.1.0`	24	23	`198.51.100.1`

Table 5. Overlapping response scopes in a full non-response-aware /24 subnets scan. SubASNs is the number of covering prefixes where at least one covered prefix is in a different AS. We only scan a single Meta domain as all behave the same.

	Cov. Pfxs	SubASNs	Answ	/24s	BGP Pfxs
www.instagram.com	464	352	352	333	52
www.google.com	1026	996	958	944	69
www.youtube.com	5987	5760	5760	5533	779

Table 4 shows an example of the problem. A resolver sends two queries with neighboring /24 prefixes (within the same /23). The nameserver responds to the first query with a scope of 24. However, it responds to the second one with a scope of 23, covering the first query's prefix scope. This example leads to a caching issue at the resolver: Should it cache both answers or overwrite the first one? As the RFC forbids such behavior, it does not contain a solution. If the resolver first issues the second query, it will not issue the first query as long as the overlapping response is cached.

Overlapping responses are not allowed, yet we detected overlapping scopes in our initial scans for Google, YouTube, and all domains by Meta. We can rule out any time-related effects in our results as our response-aware approach scans these prefixes sequentially within milliseconds. AWS Route53 nameservers exclusively provide /24 or /48 scopes, thereby precluding any overlapping response scopes. Conversely, we do not observe any overlapping prefix scopes for Wikipedia.

In order to quantify the problem in a structural manner we perform an additional ECS scan for all announced /24 prefixes. Our response-aware approach only detects overlapping responses if the first query results in the more specific answer and only a later answer covers the previous one. It skips the remaining ECS targets if the first query returns the less specific answer. In Table 5 we list the number of less specific answers covering more specific ones from other queries. We perform this more detailed scan only for `www.instagram.com` on Meta's side as Meta responds consistently with the same scopes for all three observed domains. In comparison, Google does behave differently for `www.google.com` and `www.youtube.com`.

We find that Meta has the smallest number of answers overlapping with more specifics. YouTube has nearly 6 k prefixes with at least one more specific response that would be overlapped. The responses provided are within different /24 subnets, and 13 % of YouTube, 7 % of Google, and 11 % of Instagram prefixes respond with addresses in different BGP prefixes. This result is especially outstanding as Google's guidelines lay out in detail that such behavior does not conform to RFC7871 [21]. Affected prefixes are even more interesting, considering that the two clients from our example could be in different ASes. We find that 76 % of affected prefixes with Meta and more than 97 % of affected prefixes with Google and YouTube contain more specific client prefixes assigned to a different AS than the covering prefix. While the total number of affected prefixes is limited, it is important to consider these cases when performing full address space scan evaluations.

7.5 Key Takeaways

In this section, we used our scanning approach to analyze four aspects of nameserver ECS behavior: *(i)* In contrast to previous work [3, 75], we find no significant indication of nameservers restricting their ECS support to specific resolver addresses. Therefore, researchers can directly query authoritative nameservers instead of using ECS-enabled resolvers, thereby reducing the footprint of ECS scanning campaigns on third-party infrastructure. *(ii)* Distributed measurements show that different IP anycast instances of nameservers provide distinct answers. While our scan significantly improves the visibility of single vantage point scans, it does not replace real distributed scans. *(iii)* Higher frequency measurements reveal a scope update behavior by Google, which seems to follow daily patterns. This finding shows potential for future research to uncover the load balancing behavior of ECS-enabled services. *(iv)* We find nameserver operators—one of which is Google—that provide overlapping responses. Google's public resolver has rules for ECS responses by authoritative nameservers that prohibit such overlapping responses. Therefore, Google's authoritative nameservers do not adhere to the rules defined by its own public resolver.

8 Conclusion

We present our novel response-aware ECS measurement approach to perform full address space scans for IPv4. It also includes a seeded response-aware variant to support IPv6 scanning. Our approach can be used to uncover the infrastructure of domain names from a single vantage point and collect information on load balancing behavior. It reduces the number of queries by up to 97 % compared to previous work. We find that 40 % of top-listed domain names enable ECS, but we only observe that 6 % use it by providing tailored responses. However, these domains contain important services managed by different operators. Our presented approach allows us to evaluate their load balancing behavior in more detail and use results for future work. We find that Cloudflare, the largest nameserver provider in our dataset, responds with scoped answers but rarely provides subnet-tailored answers. For Google, we find two of their own resolver's rules to be violated by their own authoritative nameservers. We make our response-aware scanner (https://github.com/tumi8/ecsplorer [69]) and data (https://doi.org/10.14459/2025mp1779517 [70]) available to the public.

Acknowledgments

We thank the anonymous reviewers and our shepherd Siva Kesava Reddy Kakarla for their valuable feedback. We would like to thank Matthew Luckie for implementing the ECS feature into the CAIDA Ark platform and providing support for our supplementary scanning campaign (NSF OAC-2131987). This work was partially funded by the German Federal Ministry of Education and Research under project PRIMEnet (16KIS1370).

References

[1] 2024. *University of Oregon Route Views Project.* http://www.routeviews.org/routeviews/

[2] Akamai. 2024. *What is ECS (EDNS0 Client Subnet)?* https://community.akamai.com/customers/s/article/What-is-ECS-EDNS0-Client-Subnet?language=en_US

[3] Rami Al-Dalky, Michael Rabinovich, and Kyle Schomp. 2019. A Look at the ECS Behavior of DNS Resolvers. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/3355369.3355586

[4] Johanna Amann, Oliver Gasser, Quirin Scheitle, Lexi Brent, Georg Carle, and Ralph Holz. 2017. Mission Accomplished? HTTPS Security after DigiNotar. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/3131365.3131401

[5] Apple Inc. 2024. *Access IP geolocation feeds.* Retrieved 2024-12-02 from https://mask-api.icloud.com/egress-ip-ranges.csv

[6] Ioannis Arapakis, Xiao Bai, and B. Barla Cambazoglu. 2014. Impact of Response Latency on User Behavior in Web Search. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval.* https://doi.org/10.1145/2600428.2609627

[7] Rob Austein. 2007. DNS Name Server Identifier (NSID) Option. RFC 5001. https://doi.org/10.17487/RFC5001

[8] AWS. 2024. *Geolocation routing.* https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-geo.html

[9] AWS. 2024. *IP-based routing.* https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-ipbased.html

[10] AWS. 2024. *IP Ranges.* https://ip-ranges.amazonaws.com/ip-ranges.json

[11] Lucas Bernardi, Themistoklis Mavridis, and Pablo Estevez. 2019. 150 Successful Machine Learning Models: 6 Lessons Learned at Booking.com. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining.* https://doi.org/10.1145/3292500.3330744

[12] Randy Bush, Massimo Candela, Warren "Ace" Kumari, and Russ Housley. 2021. Finding and Using Geofeed Data. RFC 9092. https://doi.org/10.17487/RFC9092

[13] CAIDA. 2007. *Archipelago Globally Distributed Active Measurement Platform (Ark).* https://www.caida.org/projects/ark/

[14] CAIDA. 2010. *Scamper: a Scalable and Extensible Packet Prober for Active Measurement of the Internet.* https://www.caida.org/catalog/software/scamper/

[15] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John Heidemann, and Ramesh Govindan. 2013. Mapping the Expansion of Google's Serving Infrastructure. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/2504730.2504754

[16] Matt Calder, Xun Fan, and Liang Zhu. 2019. A Cloud Provider's View of EDNS Client-Subnet Adoption. In *Proc. Network Traffic Measurement and Analysis Conference (TMA).* https://doi.org/10.23919/TMA.2019.8784530

[17] Cisco. 2024. *Umbrella Top 1M List.* https://umbrella.cisco.com/blog/cisco-umbrella-1-million

[18] Cloudflare. 2024. *Cloudflare Radar.* https://radar.cloudflare.com/

[19] Cloudflare. 2024. *The Cloudflare global network.* https://www.cloudflare.com/network/

[20] Cloudflare Docs. 2024. *EDNS Client Subnet (ECS) support.* https://developers.cloudflare.com/load-balancing/understand-basics/traffic-steering/steering-policies/#edns-client-subnet-ecs-support

[21] Carlo Contavalli, Wilmer van der Gaast, David C Lawrence, and Warren "Ace" Kumari. 2016. Client Subnet in DNS Queries. RFC 7871. https://doi.org/10.17487/RFC7871

[22] Joao da Silva Damas, Michael Graff, and Paul A. Vixie. 2013. Extension Mechanisms for DNS (EDNS(0)). RFC 6891. https://doi.org/10.17487/RFC6891

[23] David Dittrich, Erin Kenneally, et al. 2012. The Menlo Report: Ethical principles guiding information and communication technology research. *US Department of Homeland Security* (2012).

[24] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, and J. Alex Halderman. 2014. The Matter of Heartbleed. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/2663716.2663755

[25] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *Proc. USENIX Security Symposium.*

[26] Xun Fan, Ethan Katz-Bassett, and John Heidemann. 2015. Assessing Affinity Between Users and CDN Sites. In *Traffic Monitoring and Analysis.* https://doi.org/10.1007/978-3-319-17172-2_7

[27] Oliver Gasser, Quirin Scheitle, Pawel Foremski, Qasim Lone, Maciej Korczynski, Stephen D. Strowes, Luuk Hendriks, and Georg Carle. 2018. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/3278532.3278564

[28] Gasser, Oliver and Zirngibl, Johannes and Steger, Lion. 2024. *IPv6 Hitlist Service.* Retrieved 2024-12-02 from https://ipv6hitlist.github.io

[29] gdnsd. 2024. *gdnsd.* Retrieved 2024-12-02 from https://gdnsd.org

[30] Google. 2022. *EDNS Client Subnet (ECS) Guidelines.* Retrieved 2024-12-02 from https://developers.google.com/speed/public-dns/docs/ecs

[31] Google. 2024. *Chrome User Experience Report.* https://developer.chrome.com/docs/crux

[32] Google Cloud. 2024. *Global load-balancing architectures using DNS routing policies.* https://cloud.google.com/architecture/global-load-balancing-architectures-for-dns-routing-policies

[33] Fahad Hilal, Patrick Sattler, Kevin Vermeulen, and Oliver Gasser. 2024. A First Look At IPv6 Hypergiant Infrastructure. *Proc. ACM Netw.* 2, CoNEXT2, Article 11 (June 2024), 25 pages. https://doi.org/10.1145/3656300

[34] Nguyen Phong Hoang, Arian Akhavan Niaki, Jakub Dalek, Jeffrey Knockel, Pellaeon Lin, Bill Marczak, Masashi Crete-Nishihata, Phillipa Gill, and Michalis Polychronakis. 2021. How Great is the Great Firewall? Measuring China's DNS Censorship. In *30th USENIX Security Symposium (USENIX Security 21).* 3381–3398.

[35] Ralph Holz, Lothar Braun, Nils Kammenhuber, and Georg Carle. 2011. The SSL landscape: a thorough analysis of the x.509 PKI using active and passive measurements. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/2068816.2068856

[36] Ralph Holz, Jens Hiller, Johanna Amann, Abbas Razaghpanah, Thomas Jost, Narseo Vallina-Rodriguez, and Oliver Hohlfeld. 2020. Tracking the deployment of TLS 1.3 on the web: a story of experimentation and centralization. *SIGCOMM Comput. Commun. Rev.* (2020). https://doi.org/10.1145/3411740.3411742

[37] IANA. 2021. *IPv4 Special-Purpose Address Registry.* Retrieved 2024-12-02 from https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml

[38] IANA. 2023. *Address Family Numbers.* Retrieved 2024-12-02 from https://www.iana.org/assignments/address-family-numbers/address-family-numbers.xhtml

[39] ICANN. 2021. *Centralized Zone Data Service.* https://czds.icann.org/home

[40] Apple Inc. 2021. iCloud Private Relay Overview. (2021). https://www.apple.com/privacy/docs/iCloud_Private_Relay_Overview_Dec2021.PDF

[41] Internet Systems Consortium (ISC). 2020. *BIND: Using the GeoIP Features.* Retrieved 2024-12-02 from https://kb.isc.org/docs/aa-01149

[42] Internet Systems Consortium (ISC). 2024. *BIND 9.* Retrieved 2024-12-02 from https://www.isc.org/bind/

[43] IPinfo. 2024. *Anycast Dataset.* https://ipinfo.io/tags/anycast

[44] IPinfo. 2024. *IP geolocation database.* https://ipinfo.io/products/ip-geolocation-database

[45] ISO 3166-1:2020 2020. *Codes for the representation of names of countries and their subdivisions.* Standard. International Organization for Standardization, Geneva, CH.

[46] Jana Iyengar and Martin Thomson. 2021. QUIC: A UDP-Based Multiplexed and Secure Transport. RFC 9000. https://doi.org/10.17487/RFC9000

[47] Weifan Jiang, Tao Luo, Thomas Koch, Yunfan Zhang, Ethan Katz-Bassett, and Matt Calder. 2021. Towards Identifying Networks with Internet Clients Using Public Data. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/3487552.3487844

[48] Erik Kline, Krzysztof Duleba, Zoltan Szamonek, Stefan Moser, and Warren "Ace" Kumari. 2020. A Format for Self-Published IP Geolocation Feeds. RFC 8805. https://doi.org/10.17487/RFC8805

[49] KNOT DNS. 2024. *geoip Extension.* Retrieved 2024-12-02 from https://www.knot-dns.cz/docs/3.0/singlehtml/index.html#geoip-geography-based-responses

[50] Donald Ervin Knuth. 1973. *The Art of Computer Programming*. Vol. 3. Addison-Wesley Reading, MA. 498–500 pages.

[51] Platon Kotzias, Abbas Razaghpanah, Johanna Amann, Kenneth G. Paterson, Narseo Vallina-Rodriguez, and Juan Caballero. 2018. Coming of Age: A Longitudinal Study of TLS Deployment. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/3278532.3278568

[52] Athanasios Kountouras, Panagiotis Kintis, Athanasios Avgetidis, Thomas Papastergiou, Charles Lever, Michalis Polychronakis, and Manos Antonakakis. 2021. Understanding the Growth and Security Considerations of ECS.. In *Proc. Network and Distributed System Security Symposium (NDSS).* https://doi.org/10.14722/ndss.2021.24343

[53] Warren "Ace" Kumari, Evan Hunt, Roy Arends, Wes Hardaker, and David C Lawrence. 2020. Extended DNS Errors. RFC 8914. https://doi.org/10.17487/RFC8914

[54] Ioana Livadariu, Kevin Vermeulen, Maxime Mouchet, and Vasilis Giotsas. 2024. Geofeeds: Revolutionizing IP Geolocation or Illusionary Promises? *Proc. ACM Netw.* 2, CoNEXT3, Article 15 (Aug. 2024). https://doi.org/10.1145/3676869

[55] Majestic. 2024. *The Majestic Million.* https://majestic.com/reports/majestic-million/

[56] Massimo Candela. 2024. *geofeed-finder.* Retrieved 2024-12-02 from https://github.com/massimocandela/geofeed-finder

[57] MaxMind. 2024. *GeoIP Databases.* Retrieved 2024-12-02 from https://www.maxmind.com/en/geoip-databases

[58] MaxMind. 2024. *GeoLite2 Free Geolocation Data.* https://dev.maxmind.com/geoip/geolite2-free-geolocation-data

[59] MaxMind. 2024. *MaxMind GeoIP® Databases.* https://www.maxmind.com/en/geoip-databases

[60] Robert Moskowitz, Daniel Karrenberg, Yakov Rekhter, Eliot Lear, and Geert Jan de Groot. 1996. Address Allocation for Private Internets. RFC 1918. https://doi.org/10.17487/RFC1918

[61] John S. Otto, Mario A. Sánchez, John P. Rula, and Fabián E. Bustamante. 2012. Content delivery and the natural evolution of DNS: remote dns trends, performance issues and alternative solutions. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/2398776.2398831

[62] Craig Partridge and Mark Allman. 2016. Addressing Ethical Considerations in Network Measurement Papers. *Commun. ACM* 59, 10 (Oct. 2016).

[63] PowerDNS. 2024. *Authoritative Server Settings - edns-subnet-processing.* Retrieved 2024-12-02 from https://doc.powerdns.com/authoritative/settings.html#edns-subnet-processing

[64] Haoran Qiu, Subho S. Banerjee, Saurabh Jha, Zbigniew T. Kalbarczyk, and Ravishankar K. Iyer. 2020. FIRM: An Intelligent Fine-grained Resource Management Framework for SLO-Oriented Microservices. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20).* https://www.usenix.org/conference/osdi20/presentation/qiu

[65] RIPE NCC. [n. d.]. *Routing Information Service (RIS).* https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris

[66] Mario A. Sánchez, John S. Otto, Zachary S. Bischof, David R. Choffnes, Fabián E. Bustamante, Balachander Krishnamurthy, and Walter Willinger. 2013. Dasu: Pushing Experiments to the Internet's Edge. In *10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13).* https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/sanchez

[67] Patrick Sattler, Juliane Aulbach, Johannes Zirngibl, and Georg Carle. 2022. Towards a Tectonic Traffic Shift? Investigating Apple's New Relay Network. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/3517745.3561426

[68] Patrick Sattler and Mattijs Jonker. 2025. *Ark-interoperable Adaptation of ECSplorer.* https://github.com/sattler/ark-ecs

[69] Patrick Sattler, Roland Reif, and Patrick Grossmann. 2025. *ECSplorer - A Response-aware ECS Scanner.* https://github.com/tumi8/ecsplorer

[70] Patrick Sattler, Johannes Zirngibl, Fahad Hilal, Oliver Gasser, Georg Carle, and Mattijs Jonker. 2025. *ECS Scan Results at TUM University Library.* https://doi.org/10.14459/2025mp1779517 doi:10.14459/2025mp1779517.

[71] Patrick Sattler, Johannes Zirngibl, Mattijs Jonker, Oliver Gasser, Georg Carle, and Ralph Holz. 2023. Packed to the Brim: Investigating the Impact of Highly Responsive Prefixes on Internet-wide Measurement Campaigns. *Proc. ACM Netw.* (2023). https://doi.org/10.1145/3629146

[72] Lion Steger, Liming Kuang, Johannes Zirngibl, Georg Carle, and Oliver Gasser. 2023. Target Acquired? Evaluating Target Generation Algorithms for IPv6. In *2023 7th Network Traffic Measurement and Analysis Conference (TMA).* 1–10. https://doi.org/10.23919/TMA58422.2023.10199073

[73] Florian Streibelt, Jan Böttger, Nikolaos Chatzis, Georgios Smaragdakis, and Anja Feldmann. 2013. Exploring EDNS-Client-Subnet Adopters in Your Free Time. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/2504730.2504767

[74] TUM I8. 2024. *Goscanner.* https://github.com/tumi8/goscanner

[75] Marc Anthony Warrior, Uri Klarman, Marcel Flores, and Aleksandar Kuzmanovic. 2017. Drongo: Speeding Up CDNs with Subnet Assimilation from the Client. In *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies.* https://doi.org/10.1145/3143361.3143365

[76] Wikimedia. 2024. *Data centers.* Retrieved 2024-12-02 from https://wikitech.wikimedia.org/wiki/Data_centers

[77] Wikimedia. 2024. *Geographic DNS.* Retrieved 2024-12-02 from https://wikitech.wikimedia.org/wiki/DNS#Geographic_DNS

[78] Wikimedia. 2024. *Northward Datacentre Switchover (March 2024).* Retrieved 2024-12-02 from https://phabricator.wikimedia.org/T357547

[79] ZDNS. 2024. *ZDNS - Fast CLI DNS Lookup Tool.* https://github.com/zmap/zdns

[80] Johannes Zirngibl, Philippe Buschmann, Patrick Sattler, Benedikt Jaeger, Juliane Aulbach, and Georg Carle. 2021. It's Over 9000: Analyzing Early QUIC Deployments with the Standardization on the Horizon. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/3487552.3487826

[81] Johannes Zirngibl, Florian Gebauer, Patrick Sattler, Markus Sosnowski, and Georg Carle. 2024. QUIC Hunter: Finding QUIC Deployments and Identifying Server Libraries Across the Internet. In *Proc. Passive and Active Measurement (PAM).* https://doi.org/10.1007/978-3-031-56252-5_13

[82] Johannes Zirngibl, Lion Steger, Patrick Sattler, Oliver Gasser, and Georg Carle. 2022. Rusty Clusters? Dusting an IPv6 Research Foundation. In *Proc. ACM Internet Measurement Conference (IMC).* https://doi.org/10.1145/3517745.3561440

A Ethics

All our scans are set up based on a set of ethical measures we follow strictly. These are mainly based on informed consent [23] and well known best practices [62]. This study does not involve users, their information or sensitive data but focuses on publicly reachable and available services. To not cause harm to any infrastructure, we apply measures described by Durumeric et al. [25]. We limit the rate of our scans and use a blocklist based on requests to be excluded from our scans. We are directly registered as abuse contact for our scan infrastructure and react quickly to all requests. Furthermore, we host websites on all IP addresses used for scanning to inform about our research and provide contact information for further details or scan exclusion. We did not receive any exclusion requests during this scanning campaign.

Table 6. IPv6 scan limits for our response-aware scanner.

	Routed					Unrouted		
Prefix Size	/48	/40	/32	29	/16	/32	/16	Total
Queries	1	4	64	512	32 768	1	200	16 384

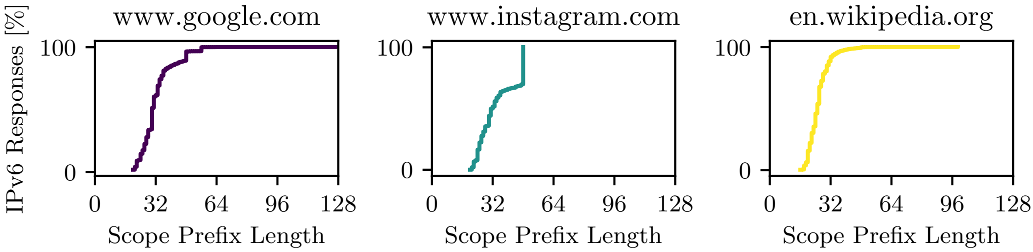

Fig. 5. Scope prefix lengths observed at our IPv6 response-aware scans. AWS Route53 always uses /48. More information can be found in Section 5.2. For comparison, Figure 2 shows results for IPv4.

Our response-aware scans are designed to reduce the query load as much as possible. Respecting the scope indicated by the nameserver allows to drastically reduce the query load compared to a scan of all /24 subnets as shown in Section 5.

B Artifacts, Configurations, and Complementary IPv6 Results

Artifacts: We will publish the source code, instructions, and used configurations for the ECS scanner and the accompanying tool to follow CNAMEs to a public Git repository. Additionally, we will push scanning result data to a public data archive.

Configurations: We use the following four IPv4 and IPv6 prefixes the ECS exploration scan in Section 6.1.

IPv4: `108.238.84.0/24`, `2.59.158.0/24`, `5.200.28.0/24`, `1.23.92.0/24`
IPv6: `2600:1700::/48`, `2a00:1630::/48`, `2a01:6f0:100::/48`, `4000:2002::/48`

The prefix limits in Table 6 are applied to each trie node, which represents a specific prefix, with a specific length. That means each /48 is allowed to have the configured amount of scans. We use limits in our IPv6 address space scans, for our response-aware scan introduced in Section 4 and used in Section 5, The first category applies to all prefixes and subprefixes in the given seed prefix set. The prefix lengths where selected after an evaluation of BGP announced prefix lengths and their prevalence. Additionally, we also allow to configure a combined (routed and not routed) limit per prefix. In our case we set these values to the sum of routed and not routed limits. While we also support limits for IPv4, full address space scans do not use any limits.

Complementary IPv6 Figures: Figure 5 shows the scopes returned by authoritative nameservers for the respective domains in our response-aware scans as discussed in Section 5.2. In comparison, results for our IPv4 scans can be found in Figure 2. Figure 6 shows the changes of response granularity for IPv6 client addresses as discussed in Section 7.3. IPv6 and IPv4 results (see Figure 4) show changes in the same time frames for Google and Youtube.

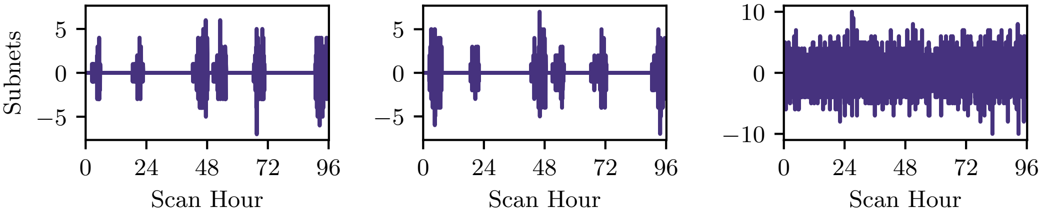

(a) www.google.com (b) www.youtube.com (c) www.instagram.com (Meta)

Fig. 6. IPv6 Scope prefix changes observed in 96 h. Positive values are changes to more specific responses and negative ones to less specific. The scan started on Friday, May 10, 2024, at 15:45 UTC, at our vantage point in Germany.

C Analyses of ECS Research Use Cases

We present use cases, including a basic analysis to show the usefulness of ECS for research, besides the general use case to load balance based on client locations. While these sections are not a central part to our contribution, they show potentially interesting future applications of our work.

Our response-aware scan can collect information on how clients are mapped to hyper-giant infrastructure. Large providers put significant effort into this mapping, resulting in precise geolocation, sometimes even better than geolocation databases. Furthermore, ECS provides possibilities to conduct *virtually distributed* scans. Therefore, it helps to improve domain-dependent scans such as QUIC, allowing for more successful handshakes. We further evaluate the usefulness of ECS scanning for IPv6 hitlists and server deployment locations.

C.1 Matching Server and Client Locations

Calder et al. [15] presented the *client-centric front-end geolocation (CCG)* approach. CCG uses majority voting to determine the location of *front-ends* (*i.e.*, the PoP's IP addresses). Their approach relies on the assumption that a nameserver will assign clients to destination addresses nearby. In contrast to their approach, to identify the server locations, we suggest to use ECS data to estimate client location based on the data collected by large providers. We rely on location hints in the reverse DNS pointer record of the PoP IP addresses to estimate the location of PoPs and the clients assigned to them. We show the potential of our suggestion by comparing our estimates to IP to location databases, and showcase an example where ECS provides a better estimate.

We assume that location hints in the rDNS name of servers can serve as a reliable indicator for mapping the server's location. Even if the exact physical location might not be fully accurate, the rDNS location hint expresses a data point controlled and actively set by the provider. We use our full address space scan results from Section 5 and manually determine the location hint patterns for the rDNS names of these results (*e.g.*, `lis` in `server.lis50.r.cloudfront.net`). We find locations for all destination addresses except for Google and YouTube where only 33 % of names contain a usable location hint.

We assign this location to clients based on ECS results and compare the location to data from IPinfo [44], and MaxMind's GeoLite City [58] and GeoIP City databases [59]. Figure 7 shows the distance between our estimate and the results from databases. The boxplots show the median as a horizontal line and the quartiles Q1 and Q3 as box. The whiskers are based on the 1.5 interquartile range value. Distances to all databases are similar. The median is between 500 km and 800 km for the three large operators and 1200 km for Wikipedia. A larger distance for Wikipedia is not

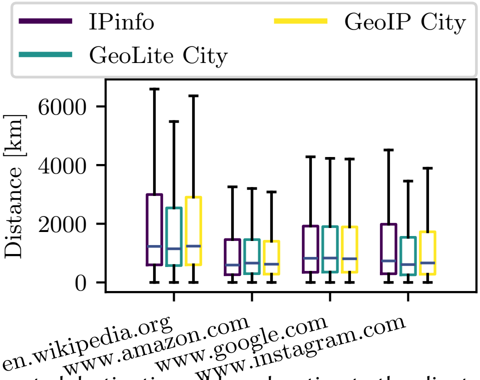

Fig. 7. Distance from the computed destination address location to the client subnet location obtained with the corresponding database.

surprising as it only had five active servers during the scanning time. Nevertheless, the data shows that 75 % of clients are within a radius of 2000 km of their assigned PoP (3000 km for Wikipedia).

Large distances between the mapping of providers and the databases indicate that one information is inaccurate. We select 25 client addresses where the distance between the database and the provider information for at least six domains is more than 10 000 km. To estimate the correct geolocation we rely on distributed round trip time (RTT) measurements with RIPE Atlas. We select probes within the country indicated by the databases and providers. Our results show that for 17 addresses, the RTT from probes close to the database location is below 40 ms while probes close to the ECS provider mapping result in RTTs larger 100 ms. Thus, due to the physical limitations of transmission distance within a certain time, results indicate a better mapping by the databases. For four addresses, results are inverted, indicating a better mapping from the ECS provider. Interestingly, probes close to both mappings for four targets result in low RTTs, indicating IP anycast deployments.

This evaluation shows, that for most client addresses, the ECS provider mapping is similar to geolocation databases and the information can be used as an alternative. In cases of large differences, both sources can be correct. Further investigation is needed to improve this mapping or improve databases based on extracted ECS provider information. Our scanner reduces the number of queries to obtain the data relevant for this use case by up to 97 %.

C.2 TLS Scanning

Scanning the Transport Layer Security (TLS) landscape is an important tool for researcher to learn about the deployment of new versions and extensions [4, 35, 36, 51, 80] and to measure how many hosts are impacted by vulnerabilities [24]. The Server Name Indication (SNI) TLS extension allows a client to indicate the domain name of the requested service to the server during the TLS handshake. For instances serving multiple domain names, this value is decisive to select the correct certificate. Zirngibl et al. [81] report that scanning with SNI is especially important with QUIC [46] (which uses TLS 1.3) as some servers even drop the connection without an error when the Client Hello does not contain an SNI value.

Therefore, we evaluate the usefulness of our *virtually distributed* ECS results. Table 1 and 2 show that our ECS scans find significantly more addresses per domain than a simple DNS resolution from one vantage point. For Google, we obtain 2.2 k IPv4 and 2.1 k IPv6 addresses while a single DNS response receives at maximum six different addresses, as it would be seen by a traditional single vantage point scan. We compare the (domain, IP address)-pairs of our ECS scans to our local DNS resolution of 638 M domain names from CZDS zone files [39], top lists, and domains in certificates from Certificate Transparency logs. Out of the combined list of 5.6 k IP addresses (see Table 1) 5.3 k do not have any associated domain within our local DNS resolution. A QUIC scan using the QScanner [80] reveals that we can successfully connect to 4.6 k of these targets when the server name is included in the SNI. However, if we remove the SNI the scan is not successful for 300 IP addresses anymore but results in a timeout. For the (domain, IP address)-pairs from our country prefix scan (see Section 6.2) we find 60 k additional IP addresses paired with 200 k domains where we can complete a QUIC connection successfully. For 5.7 k of these addresses, a handshake without SNI fails with an error and for 35.6 k addresses, the handshake even fails with a timeout. Considering the SNI problems described by Zirngibl et al. [81], we present this as a contribution to complement the existing measurement approach.

Even though TLS connections are not commonly dropped when no SNI extension is given, the server might not be able to return the appropriate certificate without it. Our QUIC scan and additional TLS over TCP scans using the Goscanner [74], find two or more valid certificates per domain for all domains except the ones served by AWS Route53 (Amazon, TikTok, and PrimeVideo). This evaluation already shows that ECS scanning is beneficial for single vantage point TLS and QUIC scans to evaluate specific providers and certificates compared to a single vantage point.

C.3 IPv6 Hitlist

The IPv6 Hitlist [27, 82] is an important measurement campaign that collects responsive IPv6 addresses. While it is collecting data since 2018 from different sources, including different vantage points, most scans are conducted from a single vantage point. Our ECS scans allow us to evaluate the coverage of the hitlist based on our *virtually distributed* scans and whether our ECS approach provides further addresses. In total, our scans reveal 73.2 M IPv6 addresses. Interestingly, 72.6 M (99.2 %) of the addresses are not part of the cumulative input of the IPv6 Hitlist. However, 99.8 % of these addresses are within prefixes announced by Amazon. According to the maintainers of the IPv6 Hitlist, Amazon is one of the main operators announcing so-called fully responsive prefixes [82]. For those prefixes, each address is responsive to ICMP and port scans. Thus, they are filtered from the hitlist service to reduce biases. We compare newly found addresses to the list of fully responsive prefixes from the service and can verify that 94.7 % are indeed within the already identified prefix and would be filtered. We assume that Amazon encodes information (*e.g.*, related to the domain or resolution time) into the returned IPv6 address. Zirngibl et al. [82] suggest that higher-layer scans should pay attention to fully responsive prefixes, but *interesting* targets should be selected. Our tool provides means to identify those targets based on domains and the load balancing strategy by the operators in addition to single query DNS resolutions.

Considering domains and results from the scans evaluated in Table 2 besides `www.amazon.com`, at least 85 % of identified addresses for each domain are already known to the IPv6 Hitlist service. Thus, their cumulative effort and combination of different sources covers most of the infrastructure of those well-known services. Nevertheless, this work shows that ECS and our tool can be used to identify the infrastructure of specific services in more detail.

Received December 2024; revised April 2025; accepted April 2025