



Cache Me Outside: A New Look at DNS Cache Probing

Arian Akhavan Niaki¹(✉), William Marczak^{2,3}, Sahand Farhoodi⁴,
Andrew McGregor¹, Phillipa Gill¹, and Nicholas Weaver^{3,5}

¹ University of Massachusetts Amherst, Amherst, MA, USA
{arian,mcgregor,phillipa}@cs.umass.edu

² Citizen Lab, Toronto, Ontario, Canada

³ University of California Berkeley, Berkeley, CA, USA
wrm@cs.berkeley.edu

⁴ Boston University, Boston, MA, USA
sahand@bu.edu

⁵ International Computer Science Institute, Berkeley, CA, USA
nweaver@icsi.berkeley.edu

Abstract. DNS cache probing infers whether users of a DNS resolver have recently issued a query for a domain name, by determining whether the corresponding resource record (RR) is present in the resolver’s cache. The most common method involves performing DNS queries with the “recursion desired” (RD) flag set to zero, which resolvers typically answer from their caches alone. The answer’s TTL value is then used to infer when the resolver cached the RR, and thus when the domain was last queried. Previous work in this space assumes that DNS resolvers will respond to researchers’ queries. However, an increasingly common policy for resolvers is to ignore queries from outside their networks. In this paper, we demonstrate that many of these DNS resolvers can still be queried indirectly through open DNS forwarders in their network. We apply our technique to localize website filtering appliances sold by Netsweeper, Inc and, tracking the global proliferation of stalkerware. We are able to discover Netsweeper devices in ASNs where OONI and Censys fail to detect them and we observe a regionality effect in the usage of stalkerware apps across the world.

Keywords: DNS · Internet measurement · Censorship.

1 Introduction

Many connections on the Internet rely on the DNS protocol to resolve a domain name into a set of IP addresses. For performance reasons, DNS resolvers typically have a *cache* of recently resolved domain names that is shared amongst all of the resolver’s users [21,22]. Unsurprisingly, this shared state exposes a side-channel by which a user of a resolver can figure out if some other user has issued a query for a specific domain name. This process is called DNS cache snooping

© Springer Nature Switzerland AG 2021

O. Hohlfeld et al. (Eds.): PAM 2021, LNCS 12671, pp. 427–443, 2021.

https://doi.org/10.1007/978-3-030-72582-2_25

(or probing) [18]. Prior work has presented various applications of this technique, including measuring the size of botnets and proliferation of malware [29], inferring web usage patterns [37] and providing a lower-bound estimate of the popularity of rare applications [30]. These prior studies assume that researchers can elicit answers by directly issuing queries to resolvers. However, most DNS resolvers nowadays do not respond to queries from outside their network. This is partly as a countermeasure to DNS amplification attacks [28], where an attacker can trick a resolver into sending a large response to a target of the attacker’s choosing by spoofing the query source address.

In this paper, we instead probe the caches of ISP DNS resolvers through *DNS forwarders* on ISP networks, devices that may be misconfigured customer-premises equipment. Prior to our work, accessing these ISP DNS resolvers has been a challenge for Internet measurement researchers. We develop and validate a tool, *dmap*, that can probe resolvers through these forwarders. We demonstrate the applicability of our technique via two case studies, (1) Netsweeper device localization, and (2) tracking the global proliferation of stalkerware.

Case Study: Netsweeper Appliance Localization. Netsweeper, Inc., is a company that provides Internet filtering devices that has received considerable recent attention, because their appliances appear to be used for Internet censorship of political and LGBTQ content in a number of repressive countries [12,13]. Measuring the proliferation and use of these tools can help hold companies to account for uses of their technology that may violate the right to free expression, and can sometimes expose cases where technology is resold or transferred to third parties [20]. Previous work on localizing Netsweeper devices [12] typically focuses on fingerprinting *block pages* that the appliances inject, and globally accessible *admin pages* used to configure the system. We show that DNS cache probing can be a complementary measurement strategy, because it indicates Netsweeper activity that these other techniques miss.

Case Study: Tracking the Global Proliferation of Stalkerware. Stalkerware are a type of spyware that have powerful surveillance capabilities and are marketed as monitoring software used for stalking [27]. Previous work has investigated the technical aspects of stalkerware and the protections that anti-virus and app stores can offer [27]. However, there has been little quantification of their prevalence across the world. Only a recent work has studied the popularity of stalkerware apps in the United States by cache probing public DNS services [30]. Since these spyware can be fingerprinted by the unique domain names they resolve [34], we measure the global proliferation of stalkerware by leveraging DNS cache probing.

Through our case studies, we make the following key observations:

1. Expanding our view of Netsweeper globally. We identify Netsweeper devices in 18 ASNs which were not identified by related efforts [1,14]. We are also able to confirm Netsweeper activity in 42% of the ASNs identified by related efforts [1,14].

2. Shedding light on stalkerware. We perform one of the first global characterizations of stalkerware using DNS cache probing. Through this analysis, we find 22 stalkerware apps active in 79 countries. The top countries are the United States, Brazil, and Germany. We observe a regionality effect in the prevalence of stalkerware apps, where apps in the Russian language are more prevalent in Russia and Ukraine.

In ongoing work, we are examining how our method can be applied to other devices and applications that perform DNS queries.

2 Background

In this section, we provide background on the operation of DNS caches, as well as prior investigation of DNS cache probing.

2.1 DNS Caching and Recursion

The mapping between a domain name and some other information (such as an IP address) is called a Resource Record (RR). The TTL field of a DNS RR is set by its authoritative nameserver and indicates how long resolvers should cache the RR [22]. If an RR for a DNS query is not cached, the resolver will try to “recursively” resolve the domain name. Once the resolver obtains the RR(s), it will send the answer to the user, and add the RR(s) to its cache for the number of seconds specified by the TTL.

A DNS query may set the Recursion Desired (RD) bit to indicate that the DNS server should attempt recursive resolution [21]. If the bit is unset ($RD = 0$), the DNS server answers the query using local information alone. In practice, some resolvers ignore this flag and always perform recursive resolution on every query. We discuss these “ill-behaved” resolvers in Sect. 3.1. If the resolver answers an $RD=0$ query from its cache, we can use the TTL value it returns to infer the arrival time of the query that caused the answer to be cached. This process is known as DNS cache probing (or snooping) [18].

2.2 DNS Cache Probing

The original treatment of DNS cache probing [18] discusses various alternate ways to infer DNS caching beyond the $RD = 0$ technique, including measuring the DNS resolver’s response time. Further, they propose a set of recommendations to mitigate DNS cache snooping, such as restricting cache access to local users, an approach that is popular today.

A popular application of DNS cache probing has been understanding the usage and popularity of networked services. Rajab et al. [29] apply DNS cache probing to estimate the density of clients accessing a network service. They measure the relative popularity of websites using this approach, but do not mathematically validate their approach. Similarly, Wills et al. [37] characterize the relative popularity of Internet applications using cache probing of 20 Local DNS

servers. Akcan et al. [4] takes a similar approach, but leverages geographically distributed open DNS resolvers to extract web usage patterns.

A second popular application of DNS cache probing has been to understand the prevalence of bots and malware that often query distinct domains via DNS. Rajab et al. [3] perform DNS cache probing on 800K DNS resolvers to infer the footprint of a botnet. Their study is based on the fact that the botnet in question issued DNS queries to resolve the Internet Relay Chat (IRC) servers used for command-and-control. They considered their result to be a lower bound on the botnet population. Randall et al. [30] perform DNS cache probing on four large public DNS resolvers (Google, Cloudflare Quad1, OpenDNS, and Quad9) and infer their caching architecture. Finally, they use their tool to estimate the number of filled caches for each resolver with a relative error of 10%–50% and present a lower-bound estimate of 22 stalkerware apps in the U.S.

These prior approaches assume direct access to the DNS server or use open DNS resolvers. Nowadays, these resolvers appear to be overwhelmingly configured to respond to queries from only clients on their network (Sect. 3.1). Thus, the techniques from these prior approaches are becoming increasingly less applicable to today’s Internet. Our leveraging of *DNS forwarders* for probing resolver (DNS backend) caches unlocks a vast trove of data missed by directly probing resolvers alone.

Furthermore, to the extent that previous work have performed DNS cache probing, there is no indication that they have distinguished between DNS forwarders and DNS backends. DNS forwarders are included in consumer NAT/gateway devices in order to respond to DNS queries within the LAN, while DNS backends are recursive DNS resolvers. This distinction is necessary for having a reliable measurement and preventing double counting.

3 Revisiting DNS Cache Probing

In this section, we describe how we leverage DNS forwarders to enable DNS cache probing. DNS forwarders are necessary to probe DNS resolvers that only respond to local clients. We first quantify the prevalence of resolvers that only respond to local clients where we find that 75% of resolvers likely respond to only their local clients (Sect. 3.1). Using DNS forwarders, local to the resolver of interest, we are able to get around this limitation. DNS forwarders are hosts that forward a DNS query to their ISP’s recursive DNS resolver [31]. This is usually the consequence of poorly engineered or misconfigured consumer NAT/gateway devices. We describe how we identify these DNS forwarders (Sect. 3.1) and how we use them to perform measurements (Sect. 3.2). We validate the set of forwarders in Sect. 3.3 and discuss the potential ethical implications of our method in Sect. 3.4.

3.1 Locating DNS Forwarders

Consumer NAT/gateway devices include a DNS forwarder so they can provide a DHCP lease (which requires specifying the DNS resolver’s IP address) to

clients before the gateway itself obtains a DHCP lease. These DNS forwarders are intended to only respond to DNS queries from within the LAN but many are improperly firewalled, and will also forward external DNS requests to the ISP’s recursive resolver.

The steps we take to identify DNS forwarders are as follows:

Step 1: Scanning the Internet’s IPv4 Space for DNS Resolvers. We begin by extracting the results of the October 5 to 11, 2020 Censys [14] DNS scans from the Censys dataset on Google’s BigQuery platform. Censys’ DNS scans send an RD=1 DNS query to the entire IPv4 address space. The nameserver of the *scan domain* name included in the Censys scan’s DNS question will always return two answers: a fixed IP address (the *control answer*) used to establish that the host correctly resolves DNS queries, and the source IP address from which the nameserver received the DNS query packet (we call this the *resolver address* or the *backend address*).

We process the Censys results as follows. First, we filter out any IPs from the Censys results that did not respond correctly. An IP responded correctly if it answered Censys’ DNS question with exactly two answers, where one answer is the control answer. Second, we attempt to exclude *shared DNS services*, such as Google’s 8.8.8.8 or OpenDNS, by including only those IPs that are in an AS categorized as “Access/Transit” by CAIDA’s *AS Classification* dataset [6], and who responded with a *resolver address* that is also in an “Access/Transit” AS. We exclude shared DNS services because their users may be globally distributed, making location inference challenging. In other words, since users from different geographical locations can send queries to shared DNS services, the fidelity of the information we get from shared DNS services will not indicate specific countries or ISPs. Furthermore, previous work have shown that the majority of end-user ISPs continue to operate their own LDNS services [7].

Step 2: Determining which Resolvers are Suited to Cache Probing. We are only interested in DNS forwarders that forward to DNS resolvers that respect the RD=0 flag, i.e., they will not perform resolution on a DNS query containing an RD=0 flag. We are also only interested in *caching* DNS resolvers that are likely to have interesting things in their caches. To find the set of DNS forwarders that exclusively forward to *caching* well-behaved DNS resolvers, we perform our own scanning to filter the list of IPs from Step 1. We run some experiments from a single vantage point in the United States using our own scan domain, whose nameserver is configured identically to the Censys scan domain and is hosted from the West Coast of the United States. We use a timeout of 20s throughout the process of our measurement. In particular, our nameserver will return exactly two answers: a control answer, and the resolver address. We filter the list of IPs from Step 1 to include only those IPs that:

- Respond four times to RD = 0 requests to unique subdomains of our scan domain with zero answers.
- Respond four times to RD = 1 requests to unique subdomains of our scan domain with a resolver address in a single “Access/transit” AS, and the resolver address returned with *approximately full TTL*.

- Respond to at least one of ten $RD = 0$ requests for [google.com](https://www.google.com) with an IP in Google’s AS (AS15169)¹.

We consider DNS forwarders that meet the criteria set out in Step 1 and Step 2 to be “well-behaved”. Table 1 shows how many forwarders passed each phase of our filtering process on seven consecutive days in October 2020.

Table 1. Number of DNS forwarders passing each stage of our filtering process during the week of October 5–11, 2020.

Forwarders filtered	10/5	10/6	10/7	10/8	10/9	10/10	10/11
Filtered Censys Scan	811,914	814,863	817,935	823,345	790,313	793,807	811,783
$RD = 0$ check	468,882	450,421	434,773	426,936	461,981	444,785	426,350
Forward check	311,140	295,560	282,458	277,183	307,889	293,075	276,150
Google check	246,710	233,441	223,014	218,417	244,032	230,042	216,049

Since [google.com](https://www.google.com) is regularly the number one domain name on the Alexa Top Sites list [5], and the Cisco Umbrella 1 Million list [11], we would expect a correct answer for this domain to typically be present in most caching DNS servers with a significant number of users (with the notable exception of countries that inject fake answers for [google.com](https://www.google.com), such as Iran and China [24,36]). We also would not expect our scan subdomains to be present in any caches, since we freshly generate a unique subdomain for each measurement, thus we expect them to be returned with *approximately full TTL* when queried with $RD = 1$. As we have configured our scan domain’s DNS server to return answers with $TTL=60$, we define approximately full TTL as either $TTL=59$ or $TTL=60$. We believe our results are not impacted by EDNS client subnet prefix per-prefix caching, since previous work have shown [7] that there is little adoption outside of Google’s Public DNS and OpenDNS, which we have excluded.

We repeat the measurements to get a sense of the behavior of the *universe* of resolvers that a forwarder may use for DNS resolution. During our DNS cache probing, we continually validate the behaviors of respecting the $RD=0$ flag, and forwarding to only a single “Access/transit” AS, as forwarder behavior may change over time. We also take privacy precautions about these DNS forwarders, as some of them might be pointing to caches of home routers. In this case, when querying our scan domain from the DNS forwarders, the answer returned by our nameserver will include the control answer and the DNS forwarder’s address instead of the resolver’s address. This indicates that the request is not being forwarded. Thus, we remove these DNS forwarders from the set of “well-behaved” forwarders. The output of our process is a set of (forwarder, resolver) pairs.

Population of Forwarders/Resolvers. Table 2 presents the breakdown of DNS forwarders that responded correctly to a query for our control domain and

¹ We analyze Farsight Security’s Passive DNS Project data [2], and the responses they observed for [google.com](https://www.google.com) since March 2, 2018, all belong to AS15169.

Table 2. Number of DNS forwarders and the number of countries and ASes on each continent where we have access to DNS resolvers (aggregated over a week). AF = Africa, AS = Asia, EU = Europe, NA = North America, OC = Oceania/Australia, SA = South America.

	AF	AS	EU	NA	SA	OC
All Forwarders	66,626	531,867	392,148	263,730	120,505	14,988
After filtering	7,890	63,411	87,826	137,341	17,337	4,883
Resolvers	419	2,609	7,545	5,671	2,238	475
Resolver countries	42	40	48	32	12	14
Resolver ASes	152	550	2,347	1,095	624	137

the set of DNS forwarders after our filtering process, across continents. After obtaining the set of resolvers each forwarder talks to, we present the population of DNS resolvers we have access to in Table 2. As shown, we have more coverage in Europe and North America in comparison to other continents. Our dataset allows us to access at least 3 DNS backends in 84% of the countries (188) and at least 2 ASNs in 74% of the countries (188).

Availability of DNS Resolvers. A relevant question is to what extent the use of DNS forwarders provides any appreciable benefit over just directly querying DNS resolvers that answer external queries. We measured this by taking our list of (forwarder, resolver) pairs, and measuring what proportion of resolvers answered queries from our measurement machine located in the US. In the measurement we ran, there were 25,665 distinct resolvers. 75% of the resolvers were not responsive to a query for our scan domain when asked directly, but did respond when we queried them via a forwarder.

3.2 Probing DNS Forwarders

We use Google’s *gopacket* library [17] to develop our DNS cache probing tool, *dmap*. As input, *dmap* takes a list of (forwarder, resolver) pairs from our filtering process, a list of domains to probe, an *exclude-list* of IP ranges of owners who have complained and chosen to be excluded from such probing, and an *interval* (which should be chosen less than the smallest authoritative TTL for any of the domain names being probed). Alternatively, a *dmap* user can specify a different TTL for each domain name, such as in the case where domain names have vastly different TTLs.

From the (forwarder, resolver) pairs, *dmap* maintains a subset of *active forwarders* that may change over time. At any given time, *dmap* tries to have two active forwarders for each resolver. If a forwarder goes offline, or is detected misbehaving (resolving an RD = 0 query for a random subdomain of our control domain, returning resolver addresses in two different ASes, or returning a resolver address with a TTL that is not approximately full), then *dmap* removes it from the active forwarders list. For each resolver associated with this

forwarder, *dmap* activates an additional forwarder in its list that talks to the same resolver.

dmap sends a DNS query packet for each domain name to each *active forwarder* every *interval* seconds. *dmap* probes at a constant rate and iterates over the space of (forwarder, domain) pairs in a random order using the method of *zmap* [15]: generating a cyclic multiplicative group ($\mathbb{Z}/p$ for a prime p larger than twice the product of the number of forwarders and domains). It is unlikely, but should the size of the set of active forwarders increase beyond this prime, a new random order will be chosen. At the same interval, *dmap* continues probing our scan domain on each forwarder (as in **Step 2** in Sect. 3.1) to determine whether each forwarder continues to respect the RD = 0 flag, and continues to forward only to resolvers in a single (“Access/transit”) AS. *dmap* remembers any new resolvers discovered for a forwarder, and uses this information when maintaining the set of active forwarders.

At the same time as *dmap* is sending DNS query packets, it is listening for DNS responses. DNS responses are filtered to ensure their relevance. DNS responses containing no answers, or no answers for the exact domain name in the question are discarded. DNS responses containing error codes are discarded. All other DNS responses are recorded in a JSON format consistent with that generated by *zmap*’s DNS module, except responses to the RD=0 and forwarding behavior validation queries. The TTL values in the DNS responses allow us to infer the date and time when the domain name was added to the DNS cache, by subtracting the response TTL from the record’s authoritative TTL (measured by a direct query to the domain’s authoritative nameserver).

To ensure that new forwarders to probe are discovered in a timely fashion, we re-process the latest Censys scan results and re-load these into *dmap* every 24 h. When these new (forwarder, resolver) pairs are loaded, *dmap* may begin probing new forwarders as necessary to ensure that at least two forwarders are being probed for each resolver.

3.3 Ground Truth Validation

To validate that our methodology can detect and infer timestamps for a nontrivial amount of DNS lookups, we performed a two-part ground truth experiment on March 23, 2019. We used $\sim 1,000$ RIPE Atlas nodes across 106 countries to send recursive queries to a single subdomain of our scan domain once per hour with random start times (using RIPE Atlas’ random function) for 72 h. At the same time, we used *dmap* to probe for the same subdomain (across approximately 16,000 DNS forwarders in 187 countries) for a period of 26 h.

In our experiment, only 1,473 unique forwarders ever returned an answer (*i.e.*, they contacted a resolver that had received a query for our scan subdomain). These forwarders used 1,247 unique resolvers in 64 countries.

Multiple Caches are Common. If a forwarder only ever uses a single resolver, we would expect to observe our domain cached for a total of TTL seconds per hour in a given resolver. However, we observed this in only 60 of the resolvers we

study, with DNS forwarders using four DNS resolvers on average, with a median value of 2 resolvers.

Because our scan domain is specially configured to return the resolver address, we can see which resolver’s cache answered a given probe. However, when probing a domain with a standard nameserver, we cannot see which resolver’s cache answered a given probe. Since many forwarders use multiple caching resolvers (all in a single AS), we must be careful when making cache inferences.

Timestamp Validation. We cross-checked the timestamps inferred from our DNS cache probing results with *ground truth* timestamps from our DNS nameserver logs that show when a resolver actually contacted our nameserver, and timestamps from our RIPE Atlas measurement logs.

The forwarders showing cache hits in our experiment queried 1,247 unique resolvers, including resolvers that never handled our queries. We found 1,198 of these resolver IP addresses in our nameserver logs, and compared their log timestamps with the timestamps inferred from DNS cache probing. We found that our timestamp inference was accurate to 5 s for 97% of the resolvers we probed (1,166).

RIPE Atlas Requests. The data from our RIPE Atlas measurement includes: (1) The RIPE node’s IP address and (2) The base64 encoded DNS question and answers; recall that a correct answer for our scan subdomain always includes the resolver address. Thus, the RIPE Atlas data effectively tells us which resolver contacted our nameserver at which time.

The RIPE Atlas data reflects that queries sent by the 1,000 RIPE nodes over our 72 h experiment caused 5,451 distinct resolvers to query our DNS server. Of these DNS resolvers, the *dmap* output reflects that we received responses from forwarders talking to 1,142 of these resolvers. Again, our inferred timestamps are accurate (per the RIPE Atlas data) to 5 s for 97% (1,100) of the resolvers.

3.4 Ethics

Since our study uses hosts on the Internet that accept queries from arbitrary sources, care must be taken to avoid overloading (or otherwise causing trouble for) the hosts. This is especially true because many DNS forwarders are in residential networks [31].

Sending Queries at a Low Rate. In our experiments, we probe each forwarder once per DNS TTL period for the set of domains we measure, which results in a maximum probing of 28 times per minute.

This is less than one query per second. We estimate that this results in less than 1 KB/sec of bandwidth usage on the forwarder, including the forwarder receiving our query, the forwarder sending our query to the resolver, the forwarder receiving the resolver’s response, and the forwarder sending the resolver’s response to us.

Because of the low rate, we do not expect our queries to cause a notable loss in performance for the host we are probing, or use a significant portion of

the host’s bandwidth allotment, or trigger any unwanted attention from ISPs or network administrators. We note that normal user activity, such as visiting a website, can sometimes result in multiple DNS queries in a short period, far in excess of our one query per second scanning rate.

Avoiding Illegal or Controversial Domains. Since we could be using residential networks to forward our queries, there is a concern over the types of domains we query. Querying a domain name containing controversial or illegal content may invite unwanted attention from authorities who erroneously interpret our query as evidence of the forwarder’s intent to participate in illegal activity. Thus, we are careful to exclude any domains that may include objectionable or censored content, or any domains associated with products or software that might be illegal in a given jurisdiction.

Privacy Issues. Although we are leveraging end-user systems, our probes are typically answered from the caches of ISP resolvers. We are thus unable to determine whether a particular end-user has looked up a particular domain.

4 Case Studies

4.1 Case Study: Netsweeper Device Localization

We applied our DNS cache probing to identify the location of Internet filtering devices sold by Netsweeper. While these devices are marketed for use in schools, libraries, and enterprise settings, previous technical work has established that these devices are also used to block political and human rights content on major consumer-facing ISPs in several repressive countries, including Bahrain, the UAE, Somalia, and Sudan [12]. The previous work used strategies including scanning the Internet for administrator login pages associated with Netsweeper deployments, and looking for Netsweeper blockpages in data collected by OONI [1] to localize these devices.

While these techniques produce useful results, they may fail to detect devices configured to *drop* Internet traffic rather than inject a blockpage attributable to Netsweeper, and may not detect installations configured without a globally accessible administrator login page. This may be especially true going forward, given increasing security concerns about exposing these login pages: an April 2020 unauthenticated remote code execution vulnerability in Netsweeper’s administrator login page would have allowed an attacker to hijack a Netsweeper installation and redirect users to malicious websites [26].

Measuring the proliferation of commercial censorship tools like Netsweeper’s product can help hold companies to account for selling these tools to abusive customers, and can sometimes expose cases where technology is resold or transferred to third parties [20]. Finding additional strategies to localize these devices is thus highly desirable.

In addition to blocking websites specified by operators, Netsweeper devices can communicate with Netsweeper’s servers to download and block lists of “objectionable” content, such as pornography and gambling sites. Netsweeper’s

system documentation [23] mentions that Netsweeper installations run a daemon called *freshnsd* that attempts to download updated versions of these URL categorization lists from `update.netsweeper.com` (the *Netsweeper update domain*). We performed a one-week measurement looking for cache hits on the Netsweeper update domain. We considered a backend to have a *Netsweeper activity* if there were cache hits for the Netsweeper update domain on at least six of the seven days of our scan.

Results. We compared our cache probing results to results from Censys [14] during our scan period. We queried Censys using Netsweeper fingerprints from a previous Citizen Lab study [12]. The Netsweeper activity was matched by IP addresses in 70 ASNs. Of these 70 ASNs, our DNS cache probing was able to probe at least one backend in 24 of the ASNs. We found Netsweeper activity in 10 of these ASNs (roughly 42%). Our cache probing also found Netsweeper activity on backends in 18 ASNs that did not show up in the Censys results. We show our results in Table 3, locations are inferred (where possible) from PTR records of DNS resolver addresses.

Table 3. DNS resolvers with Netsweeper activity.

Coun try	ASN	Censys?	Organization	Location(s)	Coun try	ASN	Censys?	Organization	Location(s)
Australia	1221		Telstra		USA	209		CenturyLink	
	4739		iiNet	Adelaide, SA Hobart, TAS		2572	×	MOREnet	MO
Austria	8447		A1 Telekom			2914		NTT America	
Bahrain	5416		Batelco			7018		AT&T	TX
Canada	35457	×	Etisalatcom			7022		Comcast	Beaverton, OR
	852		TELUS	Edmonton, AL			Boston, MA		
Colombia	19429		ETB				Denver, CO		
India	17426	×	Primenet				Wilmington, DE		
	17753		Data Ingenious Global				Newark, NJ		
	54410	×	Vodafone Idea				Lancaster, PA		
Ireland	25441	×	Imagine		UK	702		UUNET	
Kuwait	9155		QualityNet			2856	×	British Telecom	
New Zealand	23655	×	2degrees			5089		Virgin Media	
Sudan	15706	×	Sudatel			44611	×	Wavenet	Manchester
UAE	15802		Du			206747		NCSC	
Vietnam	45543		Saigontourist Cable		Yemen	30873	×	Yemennet	

One of the puzzling ASNs in which we found Netsweeper activity was ASN 206747, listed as “UK Ministry of Defence,” where we found 64 backends with Netsweeper activity. On closer inspection, the IP addresses were in a range belonging to the UK’s National Cyber Security Center (NCSC), which offers a “Protective DNS” service [8] for national and local government agencies in the UK. An NCSC blog post explains that the service is designed to detect and block malware, and that as of 2017, 44 organizations were using the service [9]. Some UK government agencies use Netsweeper, according to Censys scanning, including the Lancashire and Essex local councils. We suspect that these councils (or other government Netsweeper users in the UK) are using the NSCS’s Protective DNS service.

Discussion. Of the ASNs showing Netsweeper activity in our DNS cache probing, but not matching any Netsweeper fingerprints on Censys, some are known to

be using Netsweeper based on data from OONI. For example, OONI data shows evidence of Netsweeper use on Bahraini ISP Batelco and UAE ISP Du, though no Batelco or Du IPs were seen matching Netsweeper fingerprints in a previous Citizen Lab study [12]. In Kuwait, two ISPs (FASTtelco and Zain) are known to use Netsweeper per OONI data and previous Censys scans [12], though there appears to be relatively little OONI testing on QualityNet, per OONI’s explorer tool [25]. There appears to be limited (or no) OONI data for some other ISPs, such as Saigontourist, and Data Ingenious Global Limited [25].

Of course, similar to OONI and Censys data, it is hard to conclude based on DNS data alone whether Netsweeper installations are deployed ISP-wide, or within an institutional or enterprise setting using the ISP’s DNS servers. Nevertheless, the fact that DNS cache probing can detect Netsweeper activity that is not connected to known Censys or OONI results shows that it can be useful as an additional measurement tool for studying Internet filtering and censorship. In future work, we plan to examine the update infrastructure associated with several additional censorship and DPI products.

4.2 Case Study: Tracking the Global Proliferation of Stalkerware

We also applied our DNS cache probing to track the global proliferation of *stalkerware*, a type of generally available spyware that allows an operator to covertly monitor a target’s devices [19]. While stalkerware applications are often marketed as “employee monitoring” or “child safety monitoring” tools, they also enable Intimate Partner Surveillance or Violence (IPS or IPV) [10, 19, 27]. In the case of IPS, an abuser first installs the stalkerware on the victim’s mobile phone. The installation of the app may cause data from the phone to be sent to the stalkerware company’s servers, where the abuser can log in to access it. Stalkerware applications are generally able to gather data including text messages, location, and logs of phone calls.

Statista reported that around 3.5 billion people have smartphones in 2020 [33]. Mobile devices are generally acknowledged to contain a vast treasure trove of information about their owners. Identifying widely used stalkerware tools can help focus advocacy efforts on specific companies and specific geographic areas, and highlight the scale of the stalkerware problem.

Previous work have studied how stalkerware is used in IPS [16, 32, 35] and highlighted that while these apps often are marketed for ostensibly legal purposes, they can be easily employed for abusive ones [10]. One recent work estimated the prevalence of stalkerware apps on four shared DNS resolver services in the US, though none have done so globally.

We first obtained a set of domain names associated with the network activity of 46 stalkerware apps [34]. We filtered out domain names that appear to host the stalkerware company’s website, as these are likely to experience DNS lookups unrelated to stalkerware activity, and obtained a final list of domains representing 22 apps. We used *dmap* to perform a one-week measurement looking for cache hits on these stalkerware domain names. Similar to the Netsweeper

study in Sect. 4.1, if there were cache hits at a backend resolver for a stalkerware domain during at least six out of the seven days of our measurement, we hypothesize that a user behind that backend resolver has the stalkerware app installed.

Global Proliferation. Our cache probing found stalkerware activity on backends in 432 ASNs and in 79 countries. The top-five stalkerware apps are shown in Table 4 based on their activity in the most number of countries. The complete results can be found in Appendix A. The Cocospy app is the most prevalent stalkerware app found in 71 countries, 239 ASNs, and 889 backend resolvers. We have also listed the top-three countries for each of the apps that had the most number of backends showing stalkerware activity. For instance, the Cocospy app is observed from 200 backend resolvers in the US. We can observe that the United States and Brazil are always among the top-three countries. Figure 1 presents the number of stalkerware applications we see across the world. The United States (21), Brazil (19) and Germany (18), Great Britain (17), and Russia (16) are the top-five countries with the most number of active stalkerware apps.

Table 4. The top-five stalkerware apps prevalent in the most number of countries.

Apps	Countries	ASNs	Backend resolvers	Top countries
Cocospy	71	239	889	BR (207), US (20), GB (84)
XNSpy	60	207	981	BR (255), US (176), GB (85)
Hoverwatch	59	187	789	US (154), BR (136), GB (69)
Spyzie	57	222	757	BR (250), US (109), DE (70)
Snoopza	53	174	673	BR (106), US (88), GB (73)




Fig. 1. Heatmap of the number of stalkerware applications observed over the world.

Regionality of the Apps. We observe a direct relationship between the language of these stalkerware apps and the region where we see the app being most active according to our measurements. For instance, two stalkerware apps, “Repliculus” and “Talklog” are both Russian products and are mostly seen in backends in Russia and Ukraine. Further, “Espiao Android” and “Meuspy” which are mostly active in Brazil, are primarily available in the Portuguese language. Although a public ground truth dataset about stalkerware prevalence does not exist, this finding validates our measurements to some extent.

5 Conclusion

In this paper, we revisit DNS cache probing and show that DNS forwarders can enable DNS cache probing, even in light of resolvers only responding to local clients. We leverage these DNS forwarders to probe DNS resolver caches that were otherwise not feasible. We then develop a formulation that allows us to infer the number of network devices behind a given DNS serve and validate this technique via controlled experiments. Further, we present two case studies, (1) we localize Netsweeper devices based on a daemon available on these devices that attempts to download updated versions of URL categorization list from Netsweeper’s update domain, and (2) we study the global proliferation of stalkerware using their known indicators. In ongoing work, we are examining how our method can be applied to other applications that perform DNS queries.

Acknowledgments. We would like to thank our shepherd, Matt Calder, and all of the anonymous reviewers for their feedback on this paper. We also thank Amin Nejabatkhsh, Armin Niaki, Ilia Shumailov, Milad Nasr, Mohammad Motiei, and Negar Ghorbani for helpful comments and suggestions.

This research was financially supported by the National Science Foundation, United States, under awards CNS-1740895 and CNS-1719386. The opinions in this paper are those of the authors and do not necessarily reflect the opinions of the sponsors.

A Global Tracking of Stalkerware Apps

The 22 stalkerware apps are shown in Table 5 based on their activity in the most number of countries.

Table 5. The 22 stalkerware apps prevalent in the most number of countries.

Apps	Countries	ASNs	Backend resolvers	Top countries
Cocospy	71	239	889	BR (207), US (20), GB (84)
XNSpy	60	207	981	BR (255), US (176), GB (85)
Hoverwatch	59	187	789	US (154), BR (136), GB (69)
Spyzie	57	222	757	BR (250), US (109), DE (70)
Snoopza	53	174	673	BR (106), US (88), GB (73)
Free Android Spy	47	136	493	US (129), GB (59), BR (57)
HighsterMobile	34	95	417	US (161), GB (57), DE (44)
GuestSpy	28	50	144	GB (25), US (21), IT (18)
Easy Logger	29	87	290	US (118), GB (64), BR (12)
AndroidMonitor	27	96	328	US (87), DE (40), RU (31)
FoneTracker	19	44	127	BR (33), GB (30), US (16)
Catwatchful	18	32	88	MX (22), US (16), GB (14)
mobispy	17	19	35	DE (8), RU (4), US (3)
TheTruthSpy	16	25	49	IT (9), DE (8), US (7)
Reptculus	16	62	132	RU (62), UA (32), BY (13)
TalkLog	15	83	209	RU (92), DE (35), UA (32)
Copy9	15	20	43	UA (10), NL (10), CA(7)
iSpyoo	7	12	19	US (5), CH (4), BR (3)
Espiao Android	6	54	355	BR (336), US (11), DE (3)
mxspy	5	6	13	DE (6), GB (3), RU (2)
HelloSpy	4	4	6	PL (2), US (2), BR (1)
Meuspy	2	27	256	BR (249), US (3)

References

1. Open observatory of network interference. <https://ooni.torproject.org/>
2. Farsight security (2020). <https://www.farsightsecurity.com/solutions/dnsdb/>
3. Abu Rajab, M., Zarfoss, J., Monroe, F., Terzis, A.: A multifaceted approach to understanding the botnet phenomenon. In: Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement, New York, NY, USA, pp. 41–52. IMC 2006, ACM (2006). <https://doi.org/10.1145/1177080.1177086>. <http://doi.acm.org/10.1145/1177080.1177086>
4. Akcan, H., Suel, T., Brönnimann, H.: Geographic web usage estimation by monitoring dns caches. In: Proceedings of the First International Workshop on Location and the Web, LOCWEB 2008, New York, NY, USA, pp. 85–92. ACM (2008). <https://doi.org/10.1145/1367798.1367813>. <http://doi.acm.org/10.1145/1367798.1367813>
5. Alexa: The top 500 sites on the web. <https://www.alexa.com/topsites>
6. CAIDA: As classification (2017). <http://www.caida.org/data/as-classification/>. Accessed April 2019
7. Calder, M., Fan, X., Zhu, L.: A cloud provider’s view of EDNs client-subnet adoption. In: 2019 Network Traffic Measurement and Analysis Conference (TMA), pp. 129–136. IEEE (2019)
8. Centre UNCS. Protective DNS (PDNS). <https://www.ncsc.gov.uk/information/pdns>

9. Centre UNCS. Protective DNS service for the public sector is now live. <https://www.ncsc.gov.uk/blog-post/protective-dns-service-public-sector-now-live>
10. Chatterjee, R., et al.: The spyware used in intimate partner violence. In: 2018 IEEE Symposium on Security and Privacy (SP), pp. 441–458. IEEE (2018)
11. Cisco: Cisco umbrella 1 million. <https://umbrella.cisco.com/blog/cisco-umbrella-1-million>
12. Dalek, J., et al.: Planet netsweeper (2018). <https://citizenlab.ca/2018/04/planet-netsweeper/>
13. Dalek, J., et al.: A method for identifying and confirming the use of URL filtering products for censorship. In: ACM Internet Measurement Conference (2013)
14. Durumeric, Z., Adrian, D., Mirian, A., Bailey, M., Halderman, J.A.: A search engine backed by internet-wide scanning. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS 2015, New York, NY, USA, pp. 542–553. ACM (2015). <https://doi.org/10.1145/2810103.2813703>. <http://doi.acm.org/10.1145/2810103.2813703>
15. Durumeric, Z., Wustrow, E., Halderman, J.: Zmap: fast internet-wide scanning and its security applications, pp. 605–620 (2013)
16. Freed, D., Palmer, J., Minchala, D.E., Levy, K., Ristenpart, T., Dell, N.: Digital technologies and intimate partner violence: a qualitative analysis with multiple stakeholders. In: Proceedings of the ACM on Human-Computer Interaction 1(CSCW), pp. 1–22 (2017)
17. Google: gopacket: Provides packet processing capabilities for Go. <https://github.com/google/gopacket>
18. Grangeia, L.: DNS cache snooping. Technical report, Securi Team-Beyond Security (2004)
19. Heasley, C.: Watching The Watchers: The Stalkerware Surveillance Ecosystem (2020). <https://github.com/diskurse/android-stalkerware>. Accessed Oct 2020
20. Marquis-Boire, M., et al.: Planet blue coat: Mapping global censorship and surveillance tools (2013). <https://citizenlab.ca/2013/01/planet-blue-coat-mapping-global-censorship-and-surveillance-tools/>
21. Mockapetris, P.: Domain names - concepts and facilities. RFC 1034, RFC Editor, November 1987. <http://www.rfc-editor.org/rfc/rfc1034.txt>
22. Mockapetris, P.: Domain names - implementation and specification. RFC 1035, RFC Editor, November 1987. <http://www.rfc-editor.org/rfc/rfc1035.txt>
23. Netsweeper: Netsweeper 6.3 Documentation: List Management - Freshnsd. https://helpdesk.netsweeper.com/docs/6.3/#t=List_Management_Docs%2FFreshnsd%2FFreshnsd.htm
24. Niaki, A.A., Hoang, N.P., Gill, P., Houmansadr, A., et al.: Triplet censors: demystifying great firewall's DNS censorship behavior. In: 10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 2020) (2020)
25. OONI: OONI Explorer. <https://explorer.ooni.org/>
26. Osborne, C.: Severe Netsweeper zero-day leaves gaping hole in users networks. <https://portswigger.net/daily-swig/severe-netsweeper-zero-day-leaves-gaping-hole-in-users-networks>
27. Parsons, C., et al.: The predator in your pocket: A multidisciplinary assessment of the stalkerware application industry
28. Paxson, V.: An analysis of using reflectors for distributed denial-of-service attacks. SIGCOMM Comput. Commun. Rev. **31**(3), 38–47 (2001). <https://doi.org/10.1145/505659.505664>. <http://doi.acm.org/10.1145/505659.505664>

29. Rajab, M.A., Monrose, F., Provos, N.: Peeking through the cloud: client density estimation via DNS cache probing. *ACM Trans. Internet Technol.* **10**(3), 9:1–9:21 (2010). <https://doi.org/10.1145/1852096.1852097>. <http://doi.acm.org/10.1145/1852096.1852097>
30. Randall, A., et al.: Trufflehunter: Cache snooping rare domains at large public DNS resolvers. In: *Proceedings of the ACM Internet Measurement Conference*, pp. 50–64 (2020)
31. Schomp, K., Callahan, T., Rabinovich, M., Allman, M.: On measuring the client-side DNS infrastructure. In: *Proceedings of the 2013 Conference on Internet Measurement Conference, IMC 2013, New York, NY, USA*, pp. 77–90. ACM (2013). <https://doi.org/10.1145/2504730.2504734>. <http://doi.acm.org/10.1145/2504730.2504734>
32. Southworth, C., Finn, J., Dawson, S., Fraser, C., Tucker, S.: Intimate partner violence, technology, and stalking. *Violence Against Women* **13**(8), 842–856 (2007)
33. Statista: Number of smartphone users worldwide from 2016 to 2021. <https://www.statista.com/statistics/330695/number-of-smartphone-users-worldwide/>
34. Te-k: Indicators on Stalkerware (2019). <https://github.com/Te-k/stalkerware-indicators>. Accessed Oct 2020
35. Tseng, E., et al.: The tools and tactics used in intimate partner surveillance: an analysis of online infidelity forums. In: *29th USENIX Security Symposium (USENIX Security 2020)*, pp. 1893–1909. USENIX Association (2020). <https://www.usenix.org/conference/usenixsecurity20/presentation/tseng>
36. Wander, M., Boelmann, C., Schwittmann, L., Weis, T.: Measurement of globally visible DNS injection. *IEEE Access* **2**, 526–536 (2014)
37. Wills, C.E., Mikhailov, M., Shang, H.: Inferring relative popularity of internet applications by actively querying DNS caches. In: *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement, IMC 2003, New York, NY, USA*, pp. 78–90. ACM (2003). <https://doi.org/10.1145/948205.948216>. <http://doi.acm.org/10.1145/948205.948216>