

Enterprise DNS Asset Mapping and Cyber-Health Tracking via Passive Traffic Analysis

Minzhao Lyu, Hassan Habibi Gharakheili, Craig Russell, and Vijay Sivaraman

Abstract—The Domain Name System (DNS) is a critical service that enables domain names to be converted to IP addresses (or vice versa); consequently, it is generally permitted through enterprise security systems (e.g., firewalls) with little restriction. This has exposed organizational networks to DDoS, exfiltration, and reflection attacks, inflicting significant financial and reputational damage. Large organizations with loosely federated IT departments (e.g., Universities and Research Institutes) often are not fully aware of all their DNS assets and vulnerabilities, let alone the attack surface they expose to the outside world.

In this paper, we address the “DNS blind spot” by developing methods to passively analyze live DNS traffic, identify organizational DNS assets, and monitor their health on a continuous basis. Our contributions are threefold. First, we perform a comprehensive analysis of all DNS traffic in two large organizations (a University Campus and a Government Research Institute) for over a month, and identify key behavioral profiles for various asset types such as recursive resolvers, authoritative name servers, and mixed DNS servers. Second, we develop an unsupervised clustering method that classifies enterprise DNS assets using the behavioral attributes identified, and demonstrate that our method successfully classifies over 100 DNS assets across the two organizations. Third, our method continuously tracks various health metrics across the organizational DNS assets and identifies several instances of improper configuration, data exfiltration, DDoS, and reflection attacks. We believe the passive analysis methods in this paper can help enterprises monitor organizational DNS health in an automated and risk-free manner.

Index Terms—DNS analysis, host monitoring, network security

I. INTRODUCTION

ENTERPRISE networks are large in size with many thousands of connected devices and dynamic in nature as hosts come and go, and servers get commissioned and decommissioned to adapt to the organization’s changing needs. Enterprise IT departments often track such assets manually today, with records maintained in spreadsheets and configuration files. This is not only cumbersome, but also error-prone and almost impossible to keep up-to-date. It is, therefore, not surprising that many enterprise network administrators are not fully aware of their internal assets [40], and consequently do not know the attack surface they expose to the outside world. The problem is even more acute in University and Research

M. Lyu, H. Habibi Gharakheili, C. Russell and V. Sivaraman are with the School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, NSW 2052, Australia (e-mails: minzhao.lyu@unsw.edu.au, h.habibi@unsw.edu.au, craig.russell@unsw.edu.au, vijay@unsw.edu.au).

This article is an extended and improved version of our paper presented at the Passive and Active Measurement (PAM) 2019 conference [35].

Institute campus networks [20] that host a wide variety of sensitive/lucrative data and provide an attractive high-speed network infrastructure for hackers to use as potential attack launchpads. More importantly, their networks are operated relatively loosely and under decentralized control to cater to the fast-changing and diverse needs of various departments and research groups. While it is common practice, distributed (but not well-orchestrated) operation of IT resources can lead to inefficient utilization of resources and/or security loopholes.

In this paper, we focus on DNS, a protocol of choice exploited by cyber-crimes and botnets as it can readily bypass firewalls and security middleboxes. Due to the open nature of DNS, it is common for organizations to apply few (if any) restrictions (e.g., firewall rules) to DNS traffic. Thus, it is unsurprising to see the increasing frequency and quantity of malware compromised devices and attacks that leverage DNS protocol [24], [51], [25], [15], [6], such as DDoS attack, DNS tunneling and sensitive data exfiltration.

Enterprises typically host various kinds of DNS assets. They often host a few recursive resolvers that proxy DNS requests from internal hosts to the respective DNS servers (internal or external) and cache the returned results to reduce the number of repeated queries. Individual hosts may choose to over-ride the enterprise recursive resolvers, such as by manually changing their preferred resolver to a public one (such as Google’s and CloudFlare’s public resolvers), but in general, a majority of hosts will use the default recursive resolver provided by their organization. In addition, enterprises typically host a number of authoritative name servers to serve the various domains belonging to the organization. For example, organization-wide services (like email, VPN, etc.) may be managed by central IT. At the same time, each department may operate its own authoritative name server to resolve department-specific web pages. It is not uncommon for the various IT entities to operate in silos, often unaware of the assets being managed by the other. To make matters worse, on-campus retail stores (bookshops, food outlets, etc.) that lease connectivity from the campus may also be housing their own DNS assets, which are often poorly secured as they lack the skills.

According to the IT department of the two networks studied in this paper, existing commercial products and tools (border firewalls, SIEM¹ platforms, or network management middleboxes) often fall short of expectations, particularly in providing fine-grained visibility into DNS asset maps and DNS asset behaviors, and hence creating significant blind spots during

¹Security Information and Event Management.

operation [16]. There exists a significant body of academic research on DNS traffic analysis and security. Existing works either focus on forensic analysis of logs collected from DNS servers, such as recursive resolvers on the Internet [7], [29], [38], [12] and domain registrars [31], or characterizing (malicious) query names in DNS packets [34], [2], [4], [48]. Our prior work in [35] is among the few to analyze DNS traffic in enterprise networks with a view to identifying and monitoring DNS assets to facilitate better security management.

In this paper, we develop a data-driven method to automatically map and track cyber-health of DNS assets in an enterprise network. Commercial security appliances can consume insights obtained from our methods. Our contributions are three-fold:

- In §III, we perform a comprehensive forensic analysis of DNS traffic from two large organizations collected over 32 days, comprising nearly a billion queries/responses. We examine their network properties (IPv4/v6, UDP/TCP, etc.), functional properties (unpaired queries/responses, errors, etc.), and service properties (lookup types, record types, etc.). These enable us to build behavioral profiles of how various DNS assets (recursive resolvers, authoritative name servers, and mixed servers) behave.
- In §IV, we use the behavioral characteristics learned in §III to identify key attributes of DNS assets, extract such attributes from network/transport layer header fields without payload inspection, and develop an unsupervised machine learning model using clustering algorithms to dynamically and continuously classify asset types, including recursive resolvers, authoritative name-servers, mixed DNS servers, and regular end-host clients that may or may not be subject to enterprise network address translation (NAT). We apply our method to identify over 100 different DNS assets across the two organizations, and validate our results by cross-checking with IT staff. Our method further identifies assets that were commissioned/decommissioned or changed use during the monitoring period, further validating its utility in dynamically evolving environments.
- In §V, we develop data-driven metrics that commercial SIEM platforms can consume to track the cyber health of DNS assets or identify their anomalous behaviors – our metrics are inspired by the insights obtained from §III. Our methods reveal a prevalence of poor server configurations in both organizations, allowing attackers to exploit them for reflection attacks. Further, our methods are able to identify the organizational DNS assets that are complicit in scans, DDoS, and data exfiltration. We also give proposals on how these DNS threats can be mitigated. Our prototype runs on a commodity server (with a four-core 2.10GHz CPU, 48GB RAM, and 10Gbps network interfaces) and is ready to be deployed at the border of an enterprise network.

This paper is an extension of our previous work presented in [35]. Extensions and new contributions can be summarized as follows. First, we have improved our dataset by collecting more inclusive traffic (*i.e.*, IPv4 and IPv6) with a 32-day

duration (*i.e.*, from 3rd June 2019 to 4th July 2019), while in [35] the dataset only contained IPv4 traffic with a 7-day duration. Second, we have enriched our baseline analysis to highlight fine-grained characteristics, namely, “network”, “functional” and “service” properties of DNS traffic measured at the border of enterprise networks (§III-B); added discussions on the composition of normal and abnormal DNS lookups (*i.e.*, contents and query names) across DNS assets (§III-C), while in [35], the analysis was only limited to volumetric DNS activities (*i.e.*, Fig. 4). Third, while the methodology for DNS asset classification remains unchanged, we have updated our results using the 32-day dataset (§IV). Finally, motivated by insights pertinent to network security and operations obtained from the analysis phase, we have developed new metrics (not presented in [35].) to monitor the “general” cyber health of individual DNS assets (§V). We demonstrate how they can trigger further investigations for specialized inference on DNS assets.

Taken together, our contributions help enterprises address their current blind spot in monitoring their DNS assets and the threats they are exposed to. The passive analysis we propose is automated, risk-free, and particularly beneficial to large organizations with numerous assets managed by diverse personnel.

II. RELATED WORK

We now discuss related works on DNS traffic analysis (§II-A) and DNS attacks (§II-B) with highlights on the novelty of our work and details of extensions from our previous conference paper [35].

A. Analysis of DNS Traffic

DNS traffic has been analyzed for various purposes, ranging from measuring performance (effect of Time-to-Live of DNS records) [29], [44], [4] to identifying malicious domains [30], [31], [2] and the security of DNS [14], [39], [50]. In this paper we have profiled the pattern of DNS traffic for individual hosts of two enterprise networks to map DNS assets to their function and thereby identify their relative importance and health for efficient monitoring and security.

DNS data can be collected from different locations (such as from log files of recursive resolvers [29], [13] or authoritative name servers) or with different granularity (such as query/response logs or aggregated records). Datasets used in [39], [14] contain DNS traffic for top level domains such as `.com`, and `.net`. The work in [52] studies the root cause of query failures by analyzing DNS logs collected from recursive resolvers operated by three Internet service providers. We collect our data at the edge of an enterprise network, specifically outside the firewall at the point of interconnect with the external Internet. We note that while using data from resolver logs can provide detailed information about end hosts and their query types/patterns, this approach limits visibility and may not be comprehensive enough to accurately establish patterns related to the assets of the entire network. Specifically, recursive resolvers configured by the central IT department do not handle inbound DNS lookups toward authoritative

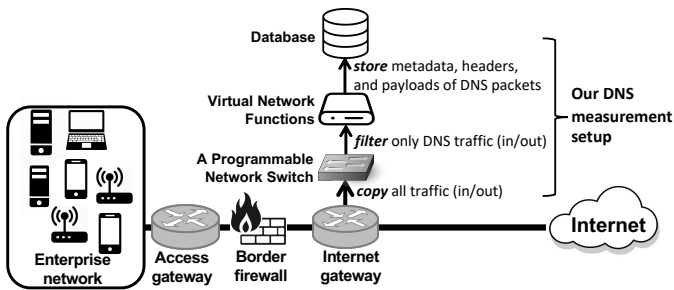


Fig. 1: Our DNS measurement setup.

name servers. Additionally, they may miss outbound lookups made by those internal hosts that are configured to use public resolvers (e.g., Google, Cloudflare) on the Internet – some of these examples will be discussed in §III-C.

There are also studies characterize malicious domain names. [30] inspects DNS traffic from top-level domain servers to detect abnormal activity and identifies key characteristics of malicious domains in terms of their resource records and lookup patterns, PREDATOR [31] derives domain reputation using registration features to enable early detection of potentially malicious DNS domains without capturing traffic, and [47] gives practical recommendations for using public domain ranking lists in security research, based on their temporary changes. As for detection of such suspicious domain names, [38] investigates into the coexistence of names in distributed DNS recursive resolvers, [27] explores the value of game theory in detecting malicious domain names generated by hidden Markov models and probabilistic context-free grammars, which can bypass legacy detection methods.

Prior works analyzed DNS traffic collected from different vantage points with various objectives. This paper measures traffic at the edge of an enterprise network where border firewalls are typically deployed. We profile the behavior and traffic health of enterprise DNS assets by identifying patterns in DNS communications and distributions of DNS packet types. We also highlight some observations, such as benign query names like “google.com” are misused in cyber-attacks (scans and query-floods) targeting enterprise networks and services.

B. Studies on DNS Security

There are protocols developed by the community to address security issues of DNS. For example, DNSSEC [8] was proposed more than two decades ago to protect data integrity, preventing attacks such as cache poisoning [49]. Authors of [14] conducted a large-scale measurement study on the adoption of DNSSEC in 2017 and found that only 1% of domains implement this secure protocol due to difficulties in the registration process as well as operational challenges involved. We will, in §II-A, highlight our observations (of course, in the context of network traffic for the two enterprises analyzed in this paper) on how a tiny fraction of DNS traffic is mapped to DNSSEC.

In addition to efforts to embed security measures into the protocol, the frequent use of DNS in volumetric attacks has

raised concerns. Authors of [50] reported the amplification factor (i.e., the ratio between the size of a DNS response and that of its corresponding query) of DNS service. Works in [39] focus on authoritative name servers used as reflectors in DNS amplification attacks – this indicates certain vulnerabilities of enterprise DNS servers that may be misused in DDoS attacks. The work in [36] proposed a hierarchical graph structure with anomaly detection models to identify distributed DNS attackers outside an enterprise network at various levels of aggregation (e.g., host, subnet, and AS²).

Existing works focus on highlighting certain vulnerabilities of the DNS protocol or developing methods for detecting DNS attacks. Our work, instead, systematically profiles and tracks the DNS behavior of active hosts in an enterprise by analyzing DNS traffic from the network border. The system we develop and the insights we draw will help IT departments better map their assets, discover potential DNS vulnerabilities, and identify misbehaved (potentially infected) hosts within their network.

III. ANALYSIS OF DNS TRAFFIC FROM TWO ENTERPRISES

In this section, we analyze the characteristics of DNS traffic collected from the border of two enterprise networks, a large university campus and a national research institute. We start by introducing our measurement setup in §III-A. In §III-B, we discuss the “network”, “functional”, and “service” properties of one-week DNS packets collected from both organizations to highlight their normal and abnormal profiles. We then (in §III-C) focus on the distribution of DNS packets among each enterprise host to reveal their DNS behavioral patterns and unhealthy traffic compositions.

A. Measurement Setup

In both organizations, the corresponding IT department provisioned a full mirror (both inbound and outbound) of their Internet traffic (each on a 10 Gbps interface) to our data collection system, shown in Fig. 1, from their border routers (**outside** of the firewall). Therefore, we focus on DNS communications that cross the border between enterprise hosts and outside servers on the Internet. It is important to note that our measurement setup at the border would not see requests of internal hosts for internal DNS servers, but they certainly exist and are handled internally processing internal communications is beyond the scope of this work. Appropriate ethics clearances for this study are granted³. We extracted DNS packets in real-time by packet mirroring rules for IPv4 and IPv6 TCP/UDP packets with port 53 on an OpenFlow-enabled network switch (i.e., NoviFlow 2122 [45]). In this paper, we focus on unencrypted DNS traffic via its typically assigned port 53, while it is worth noting that a tiny fraction of (encrypted) DNS lookups between resolvers and clients might be carried by TLS [22] and HTTPS [32] that are beyond the scope of this paper. It is important to note that our clustering model in §IV and six of eight health metrics (QSRI, QSRO,

²Autonomous System.

³UNSW Human Research Ethics Advisory Panel approval number HC17499, and CSIRO Data61 Ethics approval number 115/17.

TABLE I: Network properties of DNS packets in our dataset.

			Incoming			Outgoing		
			query	response	malformed	query	response	malformed
University	IPv4	TCP	258,315	217,210	298,979 (38%)	244,633	553,097	4,824(0.3%)
		UDP	166,492,688	181,610,373	56,665,050 (14%)	190,974,279	38,321,129	2,158,933(0.9%)
	IPv6	TCP	1,223	23,080	5,261(17%)	25,525	1,203	38(0.1%)
		UDP	10,989,944	53,592,304	200,323(0.3%)	54,673,191	7,182,025	207(0.0006%)
Rsrch. Ins.	IPv4	TCP	25,829	175,786	18,542(8%)	200,269	28,421	3,375(1.4%)
		UDP	48,629,262	53,423,998	1,034,531(1.0%)	59,638,578	22,344,154	2,493(0.003%)
	IPv6	TCP	425	11,445	14,394(55%)	19,708	338	274(1.3%)
		UDP	5,889,648	14,068,272	82,050(0.4%)	16,455,764	6,502,566	224(0.0009%)

QRI, RRO, QRO, and RRI) in §V employ volumetric DNS traffic attributes and are agnostic to packet content. Hence, they can still be applied to encrypted DNS traffic. That said, two of the health metrics, namely LEF and NELF, require DNS payload and are hence inapplicable to encrypted traffic [37]. The mirrored DNS traffic is processed by a virtual network function running on a generic server with DPDK libraries [23] which reassemble packet streams, parse headers (network, transport, application) extract the payload of each DNS packet, and stores them into our database.

Time span of our dataset: This paper considers the data collected from both organizations for 32 days, from 3 June to 4 July 2019 (i.e., the first few weeks of an academic term in the university). In §III, we use the first-week data (from 3 June to 9 June) for a comprehensive analysis. In §IV and §V, we apply our asset classification and health tracking methods to the entire (32 days) dataset. It is important to note that the insights we draw from our measurement in 2019 pertain to specific behaviors DNS assets display on the network and how accurately modeling these behaviors can highlight certain issues related to configurations, performance, or cybersecurity of these assets. Compared to what it was in 2019, the volume of network traffic in universities and research institutes has slightly dropped (observed from our ongoing traffic measurement now in 2022) due to increased working/learning from home. Our work (methods, metrics, and insights) is sufficiently generic for analyzing DNS traffic (regardless of traffic volume and frequency of anomalies) that assists network operators in systematically managing their DNS assets and infrastructure.

B. Understanding DNS Traffic at Enterprise Network Border

We begin by examining “network”, “functional”, and “service” properties of DNS packets, which provide answers for the following three questions related to DNS traffic profiles of an organization. How does each DNS packet get carried at network-level (§III-B1)? How are DNS queries responded to with or without errors (§III-B2)? What is the service type of each DNS packet (§III-B3)?

1) *Network Property:* DNS packets can be carried by either TCP or UDP at the transport layer (TCP/53 [21] or UDP/53) via IPv4 or IPv6 protocols. Table I summarizes DNS packets are distributed by transport-layer and network-layer properties.

IPv4 versus IPv6: Unsurprisingly, the majority of DNS packets are carried by IPv4 protocol, and it is clear that

the adoption of IPv6 in DNS communications has become non-negligible in both organizations. We found that 21.03% and 21.84% of outgoing⁴ DNS packets in the university and research networks, respectively, are IPv6, while this measure for incoming DNS traffic of the two organizations is 13.78% and 16.26%.

TCP versus UDP: Considering the transport layer, DNS over UDP seems to be default for enterprise hosts, accounting for more than 99% of outgoing and incoming packets in both organizations, while DNS over TCP is still staying minority (less than 0.3%). We also observe that DNS responses over TCP often have larger sizes (i.e., consist of many resource records) than their UDP counterparts. Such a distinction between TCP-based and UDP-based DNS responses is more pronounced in DNS responses carried by IPv6 in our dataset.

Queries versus Responses: Focusing on the correlation between DNS queries and responses, we highlight four pairs of query/response in Table I, as examples – each pair is color-coded for identification. It can be seen that the number of outgoing queries is slightly higher than the number of incoming responses, suggesting unanswered DNS lookups made by enterprise hosts (green and purple pairs in Table I). We also observe that count of incoming queries over IPv4 UDP is more than double the count of outgoing responses in both organizations (e.g., red and yellow pairs in Table I), highlighting the prevalence of DNS scans and floods on enterprise networks. However, this is not substantiated in IPv6 packets.

Lastly, let us look at two specific categories of DNS packets measured from the border of the university network, as highlighted by a pair of gray cells in Table I. We note that the count (i.e., more than half a million) of outgoing TCP-based response packets, those sourced from TCP port 53, over IPv4 is more than double the number (i.e., about a quarter a million) of incoming TCP-based query packets, those destined to TCP port 53. Manual investigations revealed that 53% of those outgoing TCP responses are single ACK packets on their TCP flow, without having any corresponding incoming TCP packet. We also found that those single TCP ACK packets are generated by 640 enterprise IP addresses (that are end-hosts) – none of them are classified as authoritative name server or recursive resolver later in §IV. Note that such behaviors are

⁴This paper uses the terms “outgoing” and “incoming” to denote the direction of packets, respectively, “exiting” and “entering” the border of the enterprise network.

TABLE II: Functional properties of DNS packets.

		Incoming		Outgoing	
		IPv4	IPv6	IPv4	IPv6
University	Unanswered qry.	130,683,135	3,813,677	11,431,123	1,105,818
	Unsolicited resp.	2,039,794	22,486	2,806,358	5,738
	NXDOMAIN pairs	7,493,599	1,885,742	5,164,713	1,532,410
	SERVFAIL pairs	3,897,549	34,643	1,363,391	112,618
	REFUSED pairs	24,820,580	16,409,541	2,102,724	26,130
	OtherError pairs	113,291	90	794	0
	Non-enterprise pairs	9,153,748	252,860	178,234,417	53,096,280
	Enterprise pairs	26,914,120	6,924,630	1,553,372	496,618
	Enterprise pairs	29,159,158	182,886	9,912,604	2,843,892
Research Institute	Unanswered qry.	29,159,158	182,886	9,912,604	2,843,892
	Unsolicited resp.	3,673,541	448,137	2,876,642	795,717
	NXDOMAIN pairs	2,730,158	974,480	3,775,508	1,011,591
	SERVFAIL pairs	248,275	19,715	2,389,070	5,601
	REFUSED pairs	1,390,138	245,390	781,259	133,599
	OtherError pairs	17,061	50	621	230
	Non-enterprise pairs	2,071,310	558,205	48,754,035	13,275,249
	Enterprise pairs	17,424,623	5,148,982	1,172,208	356,331

often seen in malicious TCP activities such as ACK-based host scans [11] or ACK flooding attacks [19] on DNS services. Therefore, we suspect those single outgoing ACK packets in our dataset are crafted (perhaps by malware) to look like DNS responses (sourced from port 53), bypassing the enterprise firewalls.

Malformed Packets: We found that 9.8% and 0.5% of total DNS packets in the university and research institute datasets, respectively, are malformed. Malformed packets cannot be correctly parsed since their header information do not match their payload content (*e.g.*, the number of resource records indicated in the header is inconsistent with the actual content in the payload). There are various reasons [9], [10], [42] for malformed packets such as broken software, packet truncation/distortion during transmission, or malicious traffic crafted by attackers. It can be seen that there are more malformed incoming packets compared to outgoing packets, as highlighted by percentage values (computed per each row per direction) under malformed columns in Table I. We note that all malformed packets result in no response (*i.e.*, probably they get filtered by the border firewall or dropped by the destination host).

Another observation is that malformed DNS packets are more likely carried over TCP. For example, an inbound packet over IPv4 TCP in the university network is malformed with a probability of 38%, while that is 14% over IPv4 UDP (bold text in Table I). Besides, when comparing the two organizations, we observed that the university network sends more malformed packets in total fraction than the research institution, particularly for outbound IPv4 UDP packets (0.9% versus 0.003% for the university and the research institute, respectively). It possibly indicates malicious activities originated from university hosts, as the university network is open and less restricted, while the research institute does not allow BYOT (bring-you-own-technology) devices and has strict enforcement for network security.

2) *Functional Property:* In terms of functional property, we categorize DNS packets into three clusters: (a) unpaired packets (*i.e.*, queries with no reply or responses without a corresponding query), (b) DNS lookups with a reply containing response code other than NOERROR, and (c) successful DNS lookups.

TABLE III: Service properties of DNS packets in our dataset.

		Incoming		Outgoing		
		IPv4	IPv6	IPv4	IPv6	
University	A pairs	19,986,211	3,692,671	111,896,351	29,538,537	
	AAAA pairs	7,782,897	2,014,541	32,223,426	6,615,980	
	PTR pairs	2,927,101	594,635	24,749,068	15,775,549	
	MX pairs	1,413,019	210,452	831,571	192,365	
	SPF pairs	43,943	8,600	109	28	
	TXT pairs	723,796	64,690	4,415,435	659,723	
	CNAME pairs	79,235	23,408	11,708	1,693	
	SRV pairs	599,022	197,023	2,678,513	103,494	
	SOA pairs	220,711	88,752	714,524	299,316	
	NS pairs	1,057,700	223,808	727,438	358,427	
	ANY pairs	1,205,822	46,315	114,584	9,592	
	Other pairs	21,990	10,305	1,209,553	3,754	
	Research Institute	A pairs	7,664,442	1,585,811	21,571,867	6,174,823
		AAAA pairs	2,287,039	755,134	23,774,650	6,107,818
PTR pairs		7,677,620	2,998,552	2,040,424	599,030	
MX pairs		441,075	117,015	301,651	84,904	
SPF pairs		3,782	662	15,974	3,984	
TXT pairs		120,308	13,198	1,099,786	342,250	
CNAME pairs		43,399	5,556	19,933	116	
SRV pairs		230,046	33,483	364,039	79,366	
SOA pairs		222,796	40,515	0	0	
NS pairs		683,641	132,410	532,757	179,340	
ANY pairs		101,867	23,012	830	447	
Other pairs		18,054	988	1,512	165	

Unpaired Packets: This category is captured by two rows labeled as “*unanswered qry.*” and “*unsolicited resp.*” in Table II. Unanswered queries (highlighted by red cells in Table II), carried over both IPv4 and IPv6, contribute to a large fraction of total incoming DNS packets – 40.4% and 30.8% in the university and research institute, respectively – this is mostly due to frequent DNS scans and query floods targeting enterprise DNS infrastructure. They are identified by their behavioral patterns, such as periodic, focused, slow-rate, and distributed that are typically expected from external malicious sources toward enterprise hosts. Temporal characteristics and behavioral patterns of network-based DNS attacks have been extensively analyzed by our previous work [36] that specifically develops methods for detecting distributed DNS attacks. On the other hand, unanswered outgoing queries only account for a relatively smaller fraction in each organization (*i.e.*, 2.4% and 7.9%). We note that unanswered queries are fairly normal in modern networks for a number of reasons, such as service outages, mis-routes, or dropped packets. However, persistently observing such issues for a certain host within an enterprise network is worth further investigations from security and/or operational viewpoints. Moving to unsolicited responses, their fraction in both inbound and outbound traffic are quite similar. This is mainly because of packet drop during transmission, misconfiguration of external DNS servers or DNS-based reflection attacks [28],[5] from/to the enterprises.

DNS Lookups without NOERROR Response Code: Now we focus on those DNS lookups whose response code (in the header of their corresponding response packet) is not NOERROR (response codes like NXDOMAIN, REFUSED). It is important to note that response codes other than NOERROR do not necessarily indicate malicious activities. Top three popular response codes in both enterprises are listed as NXDOMAIN, SERVFAIL and REFUSED in Table II – all other codes are grouped under OtherError. NXDOMAIN (not-existent domain) is

returned if the requested domain name is incorrect (does not exist). They are often the result of a typo in the web address, or they might be an attempt to access a website that no longer exists. This response code indeed highlights a negative response (not necessarily an error) saying the name a client asked for does not exist. That said, it is important to note that persistent `NXDOMAIN` messages are early indicators of security issues like malicious queries (*e.g.*, command-and-control and data exfiltration) sent by malware-infected hosts [7], [48]. `SERVFAIL` and `REFUSED` indicate that target DNS servers are unable to provide a resolved answer for various reasons such as zone restrictions or incorrect query formats. A frequent occurrence of those errors could indicate improper configurations on hosts/servers, or DNS attacks. As highlighted by yellow cells in Table II, for the university network, `REFUSED` and `NXDOMAIN` are the most popular error types of incoming and outgoing lookups, respectively; while `NXDOMAIN` dominates the error types of both incoming and outgoing lookups in the research institute.

Successful DNS Lookups: Given a successful (*i.e.*, with `NOERROR` flag) pair of DNS query and response packets, their requested domain name is either “relevant” (*i.e.*, belonging to services provided by the organization) or “irrelevant” to the enterprise. Large enterprises often operate authoritative name servers for their own domain names. In addition, some like universities and research institutes may support other namespaces (*e.g.*, corresponding to various groups and projects). We create allowlists of domain names whose authoritative name servers are managed within these two enterprises – we refer to corresponding inbound queries as “enterprise lookups”. Therefore, inbound DNS lookups are expected to be relevant to all domains managed by the corresponding enterprise network.

We found a non-negligible portion (21.7% and 10.4%, respectively) of inbound non-error DNS lookups that ask for irrelevant domain names in both enterprises. These inbound queries (asking for irrelevant domains) could be attributed to misconfigurations in the origin networks or even malicious traffic (*e.g.*, scans or floods) sourced from external entities. Surprisingly, some of the enterprise DNS servers (probably misconfigured) resolved those irrelevant questions. We will (in §V) take into account inbound DNS queries asking for irrelevant domain names, which enterprise authoritative name servers should not resolve, thus, are flagged as unwanted.

Also, we found that a tiny portion (0.8% and 2.4%) of outbound DNS lookups in both networks contain questions for their corresponding enterprise services. Further investigations revealed that the top destinations of those queries are public recursive resolvers such as `8.8.8.8` and `8.8.4.4` operated by Google. Contacting public DNS resolvers could have been configured manually by the user of those internal hosts or automatically by a departmental DHCP server (not necessarily managed centrally). Although best practice guidelines do not recommend such configurations, they are not necessarily malicious.

3) *Service Property:* Successful DNS lookups are asking for various types of services, such as IPv4 address (A type), IPv6 address (AAAA type), and reverse lookup for domain names (PTR type). We now focus on the question type spec-

ified in the query header of each successful lookup. Statistics for success inbound/outbound lookup pairs are shown in Table III.

Successful Lookups: We start with successful lookup types in both networks. As highlighted by blue cells in Table III, as expected, we observe that requests for IPv4/IPv6 addresses and reverse lookups for domain names are most common types of both inbound and outbound DNS traffic in both organizations. Besides, non-negligible amounts of email-related (*i.e.*, `MX` and `SPF`), text exchange (`TXT`), authoritative name service-related (*i.e.*, `CNAME`, `NS` and `SOA`), and service location (`SRV`) lookups are observed in both organizations, indicating the popularity of their corresponding services. Note that `CNAME`, `NS` and `SOA` are generated by authoritative name servers to indicate canonical names, name servers, and authoritative domains, respectively. On the University network, those outgoing responses are generated by 25 internal hosts that all will be classified as authoritative name servers later in §IV – two main servers contributed to more than 95% of those packets. Apart from these top contributing types consistently seen in both networks, we observe some different service profiles across the two organizations. For example, as highlighted by yellow cells in Table III, no outbound DNS lookup is found for `SOA` (that asking for authoritative information of a zone) in the research institute, while a few outbound lookups for `SPF` (requesting authorized email servers of a domain) are seen in the university network.

Uncommon Services: Some types of DNS lookups are found to be relatively infrequent (*e.g.*, less than 1% of total DNS lookups) in the two enterprise networks. For example, we observe inbound DNS lookups with the query type “`ANY`” in both organizations, particularly for the university network (marked as bold red text in Table III), that get answered by the enterprise servers. Though `ANY` queries can be legitimately used for debugging and checking the state of a DNS server for a particular name, DNS operators are recommended by the Internet community [1], [18] to take some prudent measures on how to handle these specific queries due to possible exploits or vulnerabilities (*e.g.*, amplification attacks). Focusing on the outbound traffic, university hosts sent out many `ANY` type requests. In contrast, hosts in the research institute rarely had such activities (relevant cells are marked as red). Besides, a small number of `A6` (deprecated version of lookups for IPv6 address) and `NAPTR` (mapping domain names to host URLs) are found in outbound requests in the university network, respectively contributing to 0.15% and 0.31% of the total count of outbound queries.

Adoption of DNSSEC: We now look at statistics of successful DNS lookups related to DNSSEC in both organizations. DNSSEC [8] was proposed almost two decades ago to provide origin authentication and integrity assurance services for DNS data. Authoritative servers may or may not implement DNSSEC and hence indicate it in their responses to resolvers or clients. Prior measurement studies [14] on domain registrars resulted that the adoption of such an extension is still in the early stage. Fairly similar observations were made in our dataset. In our dataset from the university campus, we found tiny fractions of both inbound (0.005%) and outbound (0.1%)




Fig. 2: University campus: outgoing queries and incoming responses, measured during 3 June to 9 June 2019.

lookups associated with DNSSEC services, including `DNSKEY`, `DS`, `RRSIG`, `NSEC`, `NSEC3`, and `DLV`. These measures indicate that the adoption of DNSSEC (within these organizations as well as across the Internet) is relatively slow. We made similar observations in the dataset from the research institute, where 0.005% of inbound and 0.2% of outbound lookups pertain to DNSSEC services.

C. Profiling DNS Behaviors of Enterprise Hosts

Enterprises typically operate two types of DNS servers: (a) **recursive resolvers** are those that act on behalf of end-hosts to resolve the network address of a domain name and return the answer to the requesting end-host (recursive resolvers commonly keep a copy of positive and negative responses in a local cache to reduce frequent recursions and prevent certain types of DNS-related DDoS attacks [43]), and (b) **authoritative servers** of a domain/zone are those that receive queries from anywhere on the Internet for the network address of a sub-domain within the zone for which they are authoritative (*e.g.*, `organizationXYZ.net`).

In order to better understand the DNS behavior of various hosts (and their role) inside an enterprise network, we divide the DNS dataset into two categories: (a) DNS queries from enterprise hosts that leave the network towards a server on the Internet along with DNS responses that enter the network (§III-C1), (b) DNS queries from external hosts that enter the network towards an enterprise host along with DNS responses that leave the network (§III-C2).

This analysis helps us identify important attributes related to host DNS behavior, characterizing its type/function, including authoritative name server, recursive resolver, or end-host inside the enterprise that may not always be fully visible to the network operators. This also enables us to capture the normal pattern of DNS activity for various hosts and identify the abnormal traffic status of DNS infrastructures.

1) *Outgoing Queries & Incoming Responses*: Fig. 2 shows a time trace of DNS outgoing queries and incoming responses for the university campus⁵, with granularity over 10-minute intervals on a typical semester week.

The university network handles on average 417 outgoing queries and 408 incoming responses per second. As discussed

⁵We omit results for the research institute in this section, as fairly similar observations were made.





Fig. 3: Outbound lookups for the university network: (a) CCDF of # unwanted (outgoing queries and incoming responses) DNS packets, and (b) composition of ↑lookups per enterprise host.

in Table I, 4.9% of outgoing queries are “unanswered” (*i.e.*, 12.5M out of 256.2M) during the week. And 2.06% of incoming responses to the university campus network (*i.e.*, 2.1M out of 99.9M) are “unsolicited” on the same day.

Query Per Host: We now consider individual hosts in each enterprise. Unsurprisingly, the majority of outgoing DNS queries are generated by only two hosts, A and B, in the network, *i.e.*, 66.8% of the total in the university campus (shown by blue and yellow shades in Figures 2(a)). These hosts are also the primary recipients of incoming DNS responses from the Internet. We have verified with the IT department of the enterprise that both hosts are primary recursive resolvers of this organization. In addition to these recursive resolvers, we observe a number of hosts shown by red shades in Fig. 2(a) that generate DNS queries outside of the enterprise network. The 6,089 other University hosts in Fig. 2(a) are either: end-hosts configured by public DNS resolvers that make direct queries out of the enterprise network, or secondary recursive servers operating in smaller sub-networks at the department level. We found that 301 of these 6,089 University hosts actively send queries (at least once every hour) over the day and contact more than 10 Internet-based DNS servers (resolvers or name-servers). These 301 hosts display the behavior of recursive resolvers but with fairly low throughput; thus, we deem them secondary resolvers. The remaining 5,788 hosts are only active for a limited interval (*i.e.*, between 5 min to 10 hours) and contact a small number of public resolvers over the day.

Response Per Host: Considering incoming responses in Fig. 2(b) for the university network, a larger number of “other” hosts in the organization are observed – approximately 196K IP addresses corresponding to the three subnets of size /16 owned by the university. Most of these “other” hosts (*i.e.*, 97%) are the destinations of unsolicited responses, which indicates that either misconfiguration of external DNS servers, or the university network is suffering from DNS reflections.

Unwanted DNS Packets Per Host: To better understand these potentially abnormal unanswered outgoing queries, unsolicited incoming responses, and error outgoing DNS lookups, we analyze their distribution among hosts in the two enterprises.



(a) Incoming DNS queries.


(b) Outgoing DNS responses.

Fig. 4: University campus: incoming queries and outgoing responses, measured during 3 June to 9 June 2019.

Fig. 3(a) shows the CCDF plot of the distributions per host for the university campus. All enterprise IP addresses in our dataset received unsolicited responses, and it is clear from the blue line that 99.9% of them are associated with 10 to 100 such packets – they did not have any outbound queries over the week. We observe that the hosts that have sent outbound queries to the public Internet received more unsolicited responses than those hosts that have never sent any DNS lookup. Outbound unanswered queries and lookups without `NOERROR` response code are more concentrated on a small fraction of hosts, as shown in the tail of black and red lines. 2,140 and 1,812 (out of 6,091) hosts sent unanswered queries or lookups without `NOERROR` response code – possibly due to packet drops during forwarding, typos in domain names, or malicious activities such as generating scans and DoS attacks.

Unsurprisingly, the primary recursive resolvers in both organizations are top sources and targets. In the University campus, hosts A and B respectively are the sources of 4M (33%) and 3M (25%) unanswered queries, 12M (22%) and 10M (18%) lookups without `NOERROR` response code, and are the destinations of 66K (3%) and 42K (2%) unsolicited responses.

Outbound Lookups Composition of Each Host: Let us have a closer look at the composition of outbound lookups along with inbound responses for selected hosts on the university network, as shown in Fig. 3(b). These hosts are among the top 100 in terms of outbound lookups (more than 35K over a week) with no error replies. Each bar represents an individual host. Note that each normal (green bars) or error (red bars) lookup refers to a pair of an outbound query and its corresponding inbound response – two-way communication. On the other hand, unanswered outbound queries (black bars) miss their corresponding responses, and unsolicited inbound responses (blue bars) miss their corresponding queries – one-way communication. Seventy-three of these hosts have more than 80% normal DNS packets in their outbound queries and inbound responses. The major unwanted DNS packet type is lookups without `NOERROR` response code (red shades), such as `NxDomain`, `ServerFailure` and `QueryRefused`. It might be because of typo error in domain names or malicious DNS activities such as DoS attack or contacting remote attackers using random domain strings [48]. Unanswered queries (black



(a) CCDF: # unwanted DNS pkts per host. (b) ↓ lookups composition per host.

Fig. 5: Inbound lookups for the university network: (a) CCDF of # unwanted (incoming queries and outgoing responses) DNS packets, and (b) composition of ↓lookups per enterprise host.

shades) sent to external IP addresses that do not get a reply back are the second popular reason. We focused on hosts 13, 14, 15, and 19, which are found to have respectively 23.5%, 23.2%, 23.3%, and 50.1% of their outgoing queries unwanted, and investigated their packet traces. These specific hosts seem to be likely infected servers or hosts that generate DNS scans or DoS attacks based on their traffic patterns. They consistently sent repetitive queries to many different external IP addresses or a surge of queries to an external DNS server. As we will explain later, repetitive queries and responses are patterns that are commonly found in malicious activities like scans and DoS attacks. More examples and additional insights will be provided in §V. Besides, three university hosts (index 56, 60, and 62) are also suffering from many unsolicited responses, occupying 9.29%, 3.96%, and 6.92% of their total number of packets for outbound queries and inbound responses. After manually investigating packet traces, they are found to be the target of activities resembling small-scale DNS reflection attacks. They received surges of repetitive unsolicited responses from external IP addresses during short periods; for example, 99.04% of unsolicited responses destined to host 56 were sourced from a recursive resolver belonging to a private company in China.

2) *Incoming Queries & Outgoing Responses:* Enterprises commonly receive DNS queries from the Internet that are addressed to their authoritative name servers.

It can be seen that two hosts of the University campus (*i.e.*, hosts C and D in Fig. 4(b)) are the dominant contributors to outgoing DNS responses – we have verified (by reverse lookup) that these hosts are indeed the name servers of the organization. Interestingly, for both organizations, we observe that a large number of hosts (*i.e.*, 197K IP addresses (shown by red shades in Fig. 4(a) for the university network) receive queries from the Internet. Still, a significant majority of them are unanswered (*i.e.*, 75.6%). These hosts are supposed to neither receive nor respond to incoming DNS queries, highlighting the amount of unwanted DNS traffic that targets enterprise hosts for scanning or DoS purposes.

Unwanted DNS Packets Per Host: To better understand hosts involved in incoming queries and outgoing responses, we show the distribution of inbound unanswered queries, lookups

TABLE IV: Samples of host attributes.

	QryFracOut	fracExtSrv	fracExtClient	actvQryOutTime
Univ name serv. (host C)	0	0	0.26	0
Rsch main name server	0	0	0.42	0
Univ rec. resolv. (host A)	1	0.23	0	1
Rsch main recurs. resolv.	1	0.43	0	1
Univ mixed DNS Server	0.31	0.02	0.03	1
Rsch mixed DNS Server	0.23	0.0003	0.0013	1
Univ end-host	1	0.00001	0	0.041
Rsch end-host	1	0.00001	0	0.25

replied without `NOERROR` responses and unsolicited outgoing responses from hosts inside the two enterprises.

Fig. 5(a) shows the CCDF plot of the distributions per host for the university campus. More than 99% enterprise IPs (including unassigned IP addresses) received unanswered queries from the Internet. As shown as the black line, almost all IPs are targeted by a small number (*i.e.*, less than 100) of such queries over a week – it indicates active and frequent DNS scans toward the organization. Some internal hosts received a massive amount of inbound queries at a high packet rate, located at the tail of the black line in Fig. 5(a), are likely to be victims of query flooding attacks. For example, a mixed DNS server (*i.e.*, performs as both authoritative name server and local recursive resolver) operated by a school in engineering faculty received 102M (75.5% of all unanswered incoming queries) lookups asking for non-enterprise services such as “google.com” and “163.com”.

Moreover, 59 hosts sent unsolicited outbound responses (due to server misconfiguration, used as a reflector by internal attackers or packet drop); 47 hosts sent responses without `NOERROR` (due to typos in domain names by outside users or being as victims in query-based attacks). In Fig. 5(a), the hosts that send unsolicited outbound responses are shown as blue dots, and the hosts that send responses without `NOERROR` are shown as red dots. The top 3 hosts that sent most of the unsolicited responses (86.1%) are all servers operated by sub-department (verified by reverse lookups), and the organizational IT department does not have knowledge and control over them, highlighting the security blind spots for a large enterprise network.

Inbound Lookups Composition of Each Host: Similar to what we saw earlier in Fig. 3(b), we now illustrate the composition of inbound lookups along with outbound responses for selected hosts on the university network in Fig. 5(b). These 47 hosts are among those that sent at least one outbound response over the week. Only six hosts are associated with more than 80% normal inbound lookups, and 45 hosts have error inbound lookups with response code other than `NOERROR`. Interestingly, 2,083 out of 2,085 outbound responses from the 45th host are labeled as lookups without `NOERROR` response code. This host could possibly be an authoritative name server dedicated for internal use, which received irrelevant questions such as “researchscan541.eecs.umich.edu”, “www.qq.com” and “www.wikipedia.org” and respond with `REFUSED`. Again, we acknowledge that our measurement setup at the border would not see requests of internal hosts for internal DNS servers, but they certainly exist and are handled internally. Three hosts (ranked 26, 27, and 44 in terms of the number of outgoing responses) are occupied by more than 90% unso-

TABLE V: University campus: host clusters (3 June 2019).

	Count	QryFracOut	fracExtSrv	fracExtClient	actvQryOutTime
name server	24	0.0004	1e-5	0.03	0.04
recursive resolver	21	0.99	0.04	6e-5	0.77
mixed DNS srv.	22	0.57	0.008	0.01	0.64
end-host	2,518	1.00	3e-5	0.00	0.24

TABLE VI: Research institute: host clusters (3 June 2019).

	Count	QryFracOut	fracExtSrv	fracExtClient	actvQryOutTime
name server	13	0.00	0.00	0.07	0.00
recurs. resolv.	25	1.00	0.03	0.00	0.86
mixed DNS srv.	2	0.81	0.05	0.04	0.54
end-host	245	1.00	5e-4	0.00	0.17

lited responses. They are all operated by sub-departments and are potential error-configured (such as unsynchronized timing) or reflecting DNS responses for internal attackers, as we observed a significant amount of unsolicited responses for question name `miiep` under the deprecated service type `ANY` and other irrelevant to the enterprise zone. Finally, three internal hosts suffered from a large fraction (more than 50%) of unanswered queries, especially for the 7th host – it is the mixed DNS server in engineering faculty as mentioned above, which was consistently under DoS attacks by irrelevant queries. The exhaustion of server resources led to it becoming unresponsive to most incoming queries (and only about 1% of queries got answered, including relevant and irrelevant questions).

IV. CLUSTERING ENTERPRISE DNS ASSETS

In this section, we firstly articulate key attributes that can effectively differentiate types of DNS-related enterprise hosts (§IV-A). We then develop a unsupervised clustering technique to determine if an enterprise host with a given DNS activity is a “name server”, “recursive resolver”, “mixed DNS server”, or a “regular end-host” (§IV-B). We then rank the enterprise DNS servers into “name server” and “recursive resolver” by their importance, whereas mixed DNS servers are ranked in both types (§IV-C). Finally, the regular end-hosts can be further clustered as “NATed” or “not-NATed” based on their DNS activities as described by the proposed attributes (§IV-D).

Our proposed system automatically generates lists of active servers into three categories located inside enterprise networks and rankings in terms of their name server and resolver functionalities, with the real-time DNS data mirrored from the border switch of enterprise networks. The system first performs “*Data cleansing*” that aggregates DNS data into one-day granularity and removes unsolicited responses and unanswered queries (*i.e.*, step 1); then “*Attribute extraction*” in step 2 computes attributes required by the following algorithms; “*Server mapping*” in step 3 classify DNS assets of various types; and finally “*Server ranking*” in step 4 ranks their criticality. The output is a classification and a ranked order of criticality, which an IT manager can then use to accordingly adjust management and security policies.

A. Attributes

Following the insights obtained from the DNS behavior of various hosts, we now identify attributes that help automati-

cally (a) map a given host to its function including authoritative name server, recursive resolver, mixed DNS server (*i.e.*, both name server and recursive resolver), or a regular client; and (b) rank the importance of DNS servers. All attributes are computed from DNS packets’ metadata (*i.e.*, headers) without inspecting their payload, resulting in a cost-effective inference method.

1) *Dataset Cleansing*: We first clean our dataset by removing unwanted (or malicious) records including unsolicited responses and unanswered queries – it removes the large fraction of unassigned or inactive IP addresses that are only associated with incoming DNS traffic. This is done by correlating the transaction ID of responses with the ID of their corresponding queries. In the cleaned dataset, incoming responses are equal in number to outgoing queries, and similarly for the number of incoming queries and outgoing responses.

2) *Functionality Mapping*: As discussed in §III-C1, recursive resolvers are often fairly active in terms of queries-out and responses-in, while name servers, on the other hand, are typically found with a high volume of queries-in and responses-out. Hence, a host attribute defined by the *query fraction of all outgoing DNS packets (QryFracOut)* should distinguish recursive resolvers from name servers. As shown in Table IV, this attribute has a value close to 1 for recursive resolvers and a value close to 0 for name servers.

Also, some end-hosts, configured to use public resolvers (*e.g.*, 8.8.8.8 of Google), contribute to parts of DNS queries out of the enterprise network. We note that these end-hosts ask a limited number of Internet servers during their activity period whereas the recursive resolvers typically communicate with a larger number of external servers. Thus, we define a second attribute as the *fraction of total number of external servers queried (fracExtSrv) per individual enterprise host*. As shown in Table IV, the value of this attribute for end-hosts is much smaller than for recursive resolvers. Similarly for incoming queries, we consider a third attribute as the *fraction of total number of external hosts that initiate query in (fracExtClient) per individual enterprise host*. Indeed, this attribute has a larger value for name servers compared with other hosts, as shown in Table IV.

Lastly, to better distinguish between end-hosts and recursive resolvers (high and low profile servers), we define a fourth attribute as the *fraction of active hours for outgoing queries (actvQryOutTime)*. For each host, this attribute indicate the fraction of time it sends outgoing queries. Regular clients have a smaller value of this attribute compared with recursive resolvers and mixed DNS servers, as shown in Table IV.

3) *Importance Ranking*: Two different attributes are used to rank the importance of name servers and recursive resolvers respectively. Note that we rank mixed DNS servers within both name servers and recursive resolvers for their mixed DNS behaviour. For recursive resolvers, we use *QryFracHost* defined as the *fraction of outgoing queries* sent by each host over the cleaned dataset. And for name servers, we use *RespFracHost* as the *fraction of outgoing responses* sent by each host.

B. Host Clustering

We note that the task of grouping instances (network hosts in our case) can be done using multi-class classification or clustering algorithms. Multi-class classifiers often need to be trained by a sufficient amount of labeled data to yield a decent accuracy. Given the limited number of hosts with ground-truth labels in both networks studied in this paper, we employ clustering techniques to identify groups of hosts that display distinct patterns in their DNS traffic.

1) *Selecting Algorithms*: We considered three common clustering algorithms, namely Hierarchical Clustering (HC), K-means and Expectation-maximization (EM). HC is more suitable for datasets with a large set of attributes and instances that have logical hierarchy (*e.g.*, genomic data). In our case however, hosts of enterprise networks do not have a logical hierarchy and the number of attributes are relatively small, therefore HC is not appropriate. K-means clustering algorithms are distance-based unsupervised machine learning techniques. By measuring the distance of attributes from each instance and their centroids, it groups data-points into a given number of clusters by iterations of moving centroids. In our case there is a significant distance variation of attributes for hosts within each cluster (*e.g.*, highly active name servers or recursive resolvers versus low active ones) which may lead to mis-clustering.

The EM algorithm is a suitable fit in our case since it uses the probability of an instance belonging to a cluster regardless of its absolute distance. It establishes initial centroids using a K-means algorithm, starts with an initial probability distribution following a Gaussian model and iterates to achieve convergence. This mechanism, without using absolute distance during iteration, decreases the chance of biased results due to extreme outliers. Hence, we choose an EM clustering algorithm for “*DNS Host Clustering Machine*”.

2) *Number of Clusters*: Choosing the appropriate number of clusters is the key step in clustering algorithms. As discussed earlier, we have chosen four clusters based on our observation of various types of servers. One way to validate the number of clusters is with the “elbow” method. The idea of the elbow method is to run k-means clustering on the dataset for a range of k values that calculates the sum of squared errors (SSE) for each value of k. The error decreases as k increases; this is because as the number of clusters increases, the SSE becomes smaller so the distortion also gets smaller. The goal of the elbow method is to choose an optimal k around which the SSE decreases abruptly (*i.e.*, ranging from 3 to 5 in our results, hence, $k = 4$ clusters seems a reasonable value for both the university and the research institute).

3) *Clustering Results*: We tuned the number of iterations and type of covariance for our clustering machine to maximize the performance in both enterprises. Tables V and VI show the number of hosts identified in each cluster based on data from 3 June 2019. We also see the average value of various attributes within each cluster. For the cluster of name servers, *QryFracOut* approaches 0 in both organizations (some name servers performed outbound DNS lookups for its own operational purposes), highlighting the fact that almost all outgoing DNS packets from these hosts are responses rather than queries, which matches with the expected behavior. Having a high




Fig. 6: Hosts clustering results across 32 days.

number of external clients served also indicates the activity of these hosts – in the University campus and research institute respectively 24 and 13 name servers collectively serve 81.6% and 91% (*i.e.*, $24 \times 3.4\%$ and $13 \times 7\%$) of external hosts.

Considering recursive resolvers in Tables V and VI, the average $QryFracOut$ is close to 1 for both organizations as expected. It is seen that some of these hosts also answer incoming queries (from external hosts) possibly due to their mis-configuration. However, the number of external clients served by these hosts is very small (*i.e.*, less than 5 per recursive resolver) leading to an average fraction near 0. Also, looking at the number of external servers queried (*i.e.*, $fracExtSrv$), the average value of this attribute for recursive resolvers is reasonably high, *i.e.*, 21 and 25 hosts in the University and the research network respectively contribute to 83% and 89% of total $fracExtSrv$ – this is also expected since they commonly communicate with public resolvers or authoritative name servers on the Internet.

Hosts clustered as mixed DNS servers in both organizations have a moderate value of the $QryFracOut$ attribute (*i.e.*, 0.57 and 0.81 for the University and the research network respectively) depending on their varying level of inbound/outbound DNS activity. Also, in terms of external clients and servers communicated with, the mixed servers lie between name servers and recursive resolvers. Lastly, regular end-hosts generate only outbound DNS queries (*i.e.*, $QryFracOut$ equals to 1), contact a small number of external resolvers, and are active for shorter duration of time over a day (*i.e.*, $actvQryOutTime$ less than 0.5).

4) *Interpreting the Confidence of Clustering*: Our clustering algorithm also generates a confidence level as an output. This can be used as a measure of reliability for our classifier. If adequate information is not provided by attributes of an instance then the algorithm will decide its cluster with a low confidence level – this can be interpreted as an “unknown” cluster. The average confidence level of the result clustering is 98.13% for both organizations, with more than 99% of instances classified with a confidence-level of more than 85%. This indicates the strength of our host-level attributes, enabling the algorithm to cluster them with a very high confidence-level.

5) *Server Clusters Across 32 Days*: We now check the performance of our clustering algorithm over 32 days. Fig. 6 shows a heat map for clusters of servers. Columns list server hosts that were identified in Tables V and VI (*i.e.*, 66 hosts in the University network and 40 hosts in the research network).

Rows display the cluster into which each server is classified. The color of each cell depicts the number of days (over 32 days) that each host is identified as the corresponding cluster – dark cells depict a high number of occurrences (approaching 32), while bright cells represent a low occurrence closer to 0.

In the University network we identified 25 name servers, shown by H1 to H25 in Fig. 6(a); the majority of which are repeatedly classified as a name server over 32 days, thus represented by dark cells at their intersections with the bottom row, highlighting the strong signature of their profile as a name server. Exceptions is H25, which was only active for 7 days as name server and 1 day as end-host. It is an IP address belong to school of physics under department of science, as verified by reverse lookups.

Among 21 recursive resolvers of the university campus, shown by H25 to H46 in Fig. 6(a); 7 of them (including hosts A and B in Fig. 2) are consistently classified as recursive resolver, and the rest are re-classified as end hosts (due to their varying activity). Lastly, 20 mixed servers, shown by H46 to H66 in Fig. 6(a), are classified consistently though their behavior sometimes is closer to a end-host or a name server.

Our results from the Research Institute network are fairly similar – Fig. 6(b) shows that hosts H1-H13 are consistently classified as name servers, while hosts H14-H38 are recursive resolvers and H39-H40 are mixed servers. Unlike the University Campus, 9 recursive resolvers are classified as mixed-server from 1 to 6 days. They are owned by business units in the organization, revealing the dynamicity of their DNS infrastructures.

6) *IT Verification of Clustering Results*: The IT department in both organizations verified the top-ranked DNS resolvers and name-servers found across the 32 days, meaning 100% accuracy for ground-truth DNS assets, as they are directly configured and controlled by the IT departments. For the university campus, three authoritative name servers and three recursive resolvers directly operated at the university level are consistently labeled as their true types. One mixed DNS server configured by our research group is also being classified correctly. As for the research institute, same results are obtained for the two authoritative name servers and one recursive resolver by the organization, and one mixed DNS server by our lab. Additional to the known DNS assets, we revealed unknown servers configured by departments of the two enterprises (we verified their functionality by reverse DNS lookup and their IP range allocated by IT departments).




Fig. 7: CCDF: fraction of active hour per day for NATed and not-NATed end-host IP addresses.

Interestingly, three of the department-level name servers our method identified were involved as reflectors in a DNS amplification attack, and IT was able to confirm that these were managed by affiliated entities (such as retail stores that lease space and Internet connectivity from the University) - this clearly points to the use of our system in identifying and classifying assets whose security posture the network operators themselves may not have direct control over.

C. Server Ranking

Our system discovered 46 authoritative name servers and 43 recursive resolvers in the University (a mixed DNS server are treated as both name server and recursive resolver), and 15 authoritative name server and 27 recursive resolvers at the Research Institute. However, only 6 top ranked DNS servers, in each organization, contribute to more than 90% of outgoing queries and responses. Servers ranking provides network operators with the popularity of their DNS assets.

D. Clustering of End-hosts: NATed or Not?

We note that NAT gateways (their IP addresses) appear in our dataset as enterprise assets because of the associated DNS traffic. We, therefore, believe that NAT gateways (though they are not directly indicative of end-hosts or servers) should be considered and classified for comprehensive asset monitoring. Determining whether an asset is NATed or not-NATed would help the network operator (IT department) better choose inference metrics (will be discussed in §V), indicating the end-host is performing healthily. For example, the operator may choose relatively less tight thresholds (a wider range of acceptable values) for NATed hosts than not-NATed ones as they represent a collection of end-hosts. We, therefore, applied our clustering algorithm (using the same attributes introduced in §IV-A) to those IP addresses identified as endhosts, determining whether they are behind a NAT gateway or not (*i.e.*, two clusters: NATed and not-NATed).

We, therefore, applied our clustering algorithm (using the same attributes introduced in §IV-A) to those IP addresses identified as end-hosts, determining whether they are behind a NAT gateway or not (*i.e.*, two clusters: NATed and not-NATed). In both networks, all WiFi clients are

behind NAT gateways. Additionally, some specific departments of the two enterprises use NAT for their wired clients too. We verified our end-host clustering by reverse lookup for each enterprise network. Each NATed IP address has a corresponding domain name in specific forms configured by IT departments. For example the University campus wireless NAT gateways are associated with domain-names as “SSID-pat-pool-a-b-c-d.gw.unsw.edu.au”, where “a.b.c.d” is the public IP address of the NAT gateway, and “SSID” is the WiFi SSID for the University campus network. Similarly, in the Research institute NAT gateways use names in form of “c-d.pool.rsch-primary-domain” where “c.d” is the last two octets of the public IP address of the NAT gateway in the Research institute.

1) *Clustering Results*: On 3rd June 2019, our end-host clustering shows that 337 and 42 of end-hosts IP addresses are NATed in the University campus and the Research institute respectively. We note that the two clusters of end-hosts are distinguished primarily by two attributes, namely *actvTimeFrac* - a NATed IP address (representing a group of end-hosts) is expected to have a longer duration of DNS activity compared to a not-NATed IP address (representing a single end-host), as illustrated in Fig. 7, and *numExtSrv* - a NATed IP address is expected to have more than one queried public DNS resolvers, as it represents many individual hosts each connect with their selected resolvers on the Internet. All classified not-NATed hosts contacted less than 10 external DNS servers in both organizations during 3rd June, while 54% and 26% NATed IPs in the university and research institute were queried more than 10 public servers.

2) *IT Verification by Domain Names*: We verified their corresponding domain names configured by their IT departments. Some IPs with domain-names of NAT gateways are incorrectly classified as not-NATed end-hosts - this is because their daily DNS activity was fairly low, *i.e.*, less than an hour with only one external resolver contacted. On the other hand, not-NATed end-hosts, expected to be less active but with long duration and high volume of DNS activity (*i.e.*, almost the whole day), were misclassified. While we have limited ground-truth data on DNS assets, parts of our classifications are verified by their DNS names assigned by IT departments. For assets classified as NATed end-hosts, we managed to verify our classification for 77.2% of them on the university campus and 75.0% in the research institute that have domain-name patterns dedicated for NAT gateways by IT departments. For assets classified as not-NATed end-hosts, the majority of them (more than 90%) in both organizations do not map to any organizational domain-names, hence implying typical (not-NATed) end-hosts. It is important to note that an IP address classified as NAT or not-NATed but could not be verified by the respective IT department does not mean false classification. That address could have been allocated by sub-departments or groups. For example, in our research laboratory, we have configured several WiFi routers to connect experimental and commercial devices (*e.g.*, IoT) to the Internet via the campus network. Those public IP addresses do not necessarily map to any domain names managed by our university IT department.




Fig. 8: CCDF: Consistency of end-hosts clustering.

3) *Clustering Results Across 32 Days*: Looking into the consistency of end-hosts clustering across 32 days, we note that more than 90% end-hosts in the University campus are consistently labeled as NATed over 7 days (as show in Fig. 8(a)). 52% end-hosts are classified as NATed from 7 days to 15 days. Those IP addresses are owned by sub-departments in the university, and re-shuffled within their subnets by the organizational DHCP servers periodically. As for the University IP addresses get classified as not-NATed (e.g., desktops with public IP addresses through wired connection), majority (63%) of them only appear once during 32 days. It is because of their low-profile activities and daily IP re-shuffling.

Similar observations were obtained from the research institute (shown in Fig. 8(b)), except there are 5 IP addresses appeared as NATed across the 32 days – they belongs to IT infrastructures controlled by critical scientific basements such as Australia Telescope National Facilities, which are separated controlled with more freedom thus not affected by periodically DHCP reallocation.

V. MONITORING DNS ASSET HEALTH

Having shown how DNS assets in an enterprise network can be identified and classified based on their network behavior, we now extend the study to monitor their health continuously. The objective is to detect *anomalous* behavior, indicating that the asset is being misused or attacked, and identify the root cause of such deviations in behavior. We begin in §V-A by providing two examples of observable anomalies from our dataset – one attributable to poor configuration and the other subject to a DDoS attack. Inspired by these examples, in §V-B we develop a set of *health metrics* that can track the behavior of each asset along various dimensions, and in §V-C develop a method to label and warn anomalous behaviors based on these health tracking metrics. Finally, in §V-D we apply our methods to the 32-day dataset from the two organizations and present results into misuse and attack patterns detected by our methods.

A. Examples Illustrating Anomalous DNS Asset Behavior

By manually inspecting our dataset, we could identify several behavioral patterns that seemed unusual. In total, we identify six types of anomalous DNS behaviors, which will be discussed later in §V-B and Table §VIII. We now begin by providing a couple of illustrative examples of anomalous

behavior and subsequently develop methods to automatically detect misbehaviors by tracking various health metrics.

Example 1 – DNS Misuse: We found that one of the authoritative DNS servers at the Research Institute was responding with an unexpected high number of “NXDOMAIN” messages, indicating that corresponding queried domain names do not exist. Manual investigation revealed that those queried names were irrelevant (e.g., “www.taobao.com”) to the organization. In fact, almost a third of DNS queries to this server were irrelevant. Note that the enterprise network does not manage the authoritative name server of those queries (as discussed in §III-B2). However, all of them got responded. We also found that about 15% of incoming queries were asking for “com”, which is irrelevant to the enterprise. The enterprise server responded with an “NXDOMAIN” message. Note that network administrators may employ certain policies to better manage their service infrastructure. DNS policies allow administrators to configure their DNS server to respond in a custom manner based on DNS queries and DNS clients that send queries [41]. An authoritative name server may, by default, give NXDOMAIN responses to irrelevant queries or can be configured in a way to drop certain queries. In the context of this example, those NXDOMAIN responses which are unnecessary and relatively large in volume could have been prevented by appropriate policy configurations. This example demonstrates how a poorly configured server can behave outside its intended function. Such a vulnerability exposes the server to attackers who aim to launch a denial-of-service attack or use it as a reflector for attacking others.

Example 2 – DNS Flood Attack: We found one of the authoritative DNS servers in the University dataset to show a sustained 142% increase in inbound query rates over a 10-day period (7-Jun 0:17am till 17-Jun 4:43pm). By investigating packet traces for this interval, we found 3.3M queries with the same query name “aids.gov” sourced from 974 external IP addresses with a certain pattern of activity – each external source (on average) launched about 300 queries within a 20-second period, and then went idle. We also verified that this domain name “aids.gov” is irrelevant to our enterprise network by way of performing a name server lookup. We note that repetitive queries and/or responses are known patterns in volumetric attacks like DoS and scan. As a result, the DNS server struggled to keep up with high rates of requests and managed to process only about 70% of incoming queries. Lastly, we found 40.9% of the responses during this period correspond to queries irrelevant to the organization where almost half of these (irrelevant) responses contained an “NXDOMAIN” response code.

B. DNS Traffic Health Metrics

Having seen some examples of poor behavior from DNS servers, we now propose monitoring metrics that can be used to track the health of each DNS asset in the organization. Insights drawn from a comprehensive analysis of DNS asset behaviors (particularly §III-B) inspired the design of four categories of metrics encompassing four aspects of behavior, namely, service, functional, network, and volumetric. In other




Fig. 9: Examples of observed DNS anomalies and their corresponding health alerts.

words, our metrics aim to measure the “baseline” behavior of DNS assets and track their “general” health on the network. They are able to raise flags indicative of possible performance, configuration, or cybersecurity issues for those assets, triggering “specialized” investigation and/or inspection (*e.g.*, involving in DNS exfiltration [2] or DGA-based [3] malware activities).

Service Behavior: From a border perspective, authoritative name servers are expected to only serve DNS queries seeking to resolve domains relevant to the enterprise. Conversely, recursive resolvers should only send outbound queries for domains outside of the enterprise – queries for internal domains are internally sent to the enterprise authoritative name servers without crossing the network border. We therefore define **Non-Enterprise Lookup Fraction (NELF)** as the fraction of query names that are irrelevant to the enterprise services. A properly configured authoritative name server should have NELF of 0, while for a recursive resolver this metric should be 1. Note that “Non-Enterprise Lookups” are for those domains that are truly beyond the operational scope of the enterprise network (*e.g.*, taobao.com, google.com, and umich.edu, as discussed in §III-C2 and §V-A). In practice, enterprise IT departments could construct and maintain an allowlist to dynamically add/remove domains managed in their networks, maximizing the accuracy of NELF. A practical method would be (reactively) performing name server lookups (*i.e.*, NS type) for unseen domains to verify whether their address spaces belong to the enterprise or not.

Functional Behavior: Under ideal conditions, responses of a properly functioning DNS server are expected to carry “NOERROR” as response code. However, a DNS query can fail due to various reasons (discussed in §III-B2), such as the domain name queried may not exist, an answer cannot be given, or the server refuses to answer due to some policies configured. Therefore, we define **Lookup Error Fraction (LEF)** for a DNS server as the fraction of its responses that carry a response code other than NOERROR. A large value (configurable, say, >30%) for this metric indicates potential misbehavior that requires further diagnosis.

Network Behavior: Under normal circumstances a query is associated with a response. However, the network trace often reveals inbound responses with no outbound queries (*e.g.*, a reflective attack to a victim whose IP address was spoofed), as well as outbound queries with no inbound response (*e.g.*, a malicious internal host launching a DoS attack via the DNS

TABLE VII: Alerts and occurrence frequency (in the fraction of epochs) for our two example DNS assets.

Direction	Profile	Alert	Exp. 1	Exp. 2
In ↓	Service	high NELF	83.7%	85.6%
↓	Functional	high LEF	6.6%	0.1%
↓	Network	low QSRI	94.5%	15.6%
↓	Network	high QSRO	0.1%	0.3%
↓	Volmetric	high QRI	0.0%	29.7%
↓	Volmetric	high RRI	5.1%	9.8%
Out ↑	Service	low NELF	0.0%	0.0%
↑	Functional	high LEF	100.0%	20.2%
↑	Network	high QSRI	0.8%	0.3%
↑	Network	low QSRO	0.0%	84.4%
↑	Volmetric	high RRO	5.6%	7.7%
↑	Volmetric	high QRO	1.0%	31.1%

cache/proxy). To track such anomalous network behavior, we define the Query Service Ratio Inbound (QSRI), *i.e.*, ratio of outbound responses to inbound queries, and Query Service Ratio Outbound (QSRO), *i.e.*, ratio of inbound responses and outbound queries. All DNS assets should ideally have these two metrics as 1, showing the balanced profile of queries and responses.

Volmetric Behavior: A sudden increase in DNS packet rates certainly highlights an abnormal incident that may impact enterprise assets. It is important to note that our primary aim is to ensure classified assets display healthy behavior (in terms of the volume of their activities) at relatively slower time scales instead of looking for anomalies in (close to) real-time. We, therefore, track hourly counts (rates) of inbound and outbound queries and responses of individual DNS assets by four metrics, namely, **QryRateIn (QRI)**, **RespRateOut (RRO)**, **QryRateOut (QRO)**, and **RespRateIn (RRI)**. That way, the number of alerts becomes more manageable and meaningful for a high-level analysis of asset cyber-health (whether they are involved in volumetric attacks or data exfiltration). Note that a network operator may configure a different time scale (*e.g.*, 10 minutes) for these metrics, obtaining insights into activities relatively faster. These metrics will flag those epochs in which rates increase more than a configurable threshold value (a threshold of 30% is adopted in our implementation as will be discussed in §V-C), suggesting possible volumetric misbehaviors. Note that the increase in each time epoch is computed with respect to the previous epoch, representing the time derivative of rates.

C. Using Health Metrics to Detect Anomalies

Using the health metrics identified above, we build a simple mechanism to detect and alert various anomalous behaviors of DNS assets. The set of anomalies we consider in this paper, illustrated pictorially in Fig. 9, include:

- **Misconfiguration:** Consider an authoritative DNS server that has been poorly configured and resolves queries for domains that it has no authority over (*i.e.*, do not belong to the enterprise). The exploitation of this by attackers (*e.g.*, as a reflector) will manifest in an alert when the NELF metric becomes high, while LEF could also be

high (in case the queries are malformed or non-existent). Conversely, a misconfiguration alert is triggered when the NELF metric falls below a threshold value for a poorly configured recursive DNS resolver. We acknowledge that misconfigurations may not necessarily indicate security events. However, they can highlight an unhealthy state of operation for the respective DNS asset, and hence worthwhile to get flagged for further investigations and/or remedial actions if possible.

- **DDoS Attack (query/response/reflector):** A distributed denial-of-service attack on an enterprise DNS server will manifest in the form of a volumetric rise in QRI, potentially accompanied by a high value in NELF and/or LEF. Most queries in DDoS attacks tend to be either fixed or random domains instead of customizing query names specific to the victim enterprise. In DDoS, an infected host could directly generate volumetric queries, responses, or act as an attack reflector.
- **DNS Exfiltration:** An infected enterprise host, attempting to exfiltrate data via DNS, will cause QRO to rise, potentially accompanied by unanswered queries (rise in QSRO) and/or lookup errors (rise in LEF). A combination of these metrics can be used as triggers to conduct a deeper investigation into exfiltration, *e.g.*, using the method developed by [2]. One may argue that QRO is expected to be relatively high for legitimate recursive resolvers. Therefore, we infer from a combination of metrics, each with specific thresholds (value ranges) to cater for some reasonable deviations (discussed in §V-D).
- **Scans:** The presence of malware in the enterprise that performs outbound scans can be detected by monitoring for a rise in outbound queries (QRO), potentially accompanied by unanswered queries (rise in QSRO) and/or lookup errors (rise in LEF).

In what follows we continuously track the health metrics of the various DNS assets identified in the two enterprise networks by our earlier clustering algorithm, and evaluate our ability to identify anomalous behaviors indicative of misconfigurations and/or attacks. Note that our proposed metrics and alerts from DNS behavioral monitoring could be consumed by SIEM platforms and/or combined with security appliances to verify whether an enterprise host is indeed involved in malicious communications or not. Such combined inferences are beyond the scope of this paper.

D. Insights from Two Enterprise Networks

We applied the proposed traffic health metrics to our 32-day DNS traces captured from both organizations, comprising the assets as identified earlier in Tables V and VI for the University (67 DNS assets) and Research Institute (40 DNS assets) respectively. The metrics are computed each epoch (of one hour), and our first step is to identify epochs wherein the health metrics deviate significantly from their expected values. In general, DNS assets in the University raise more alerts than the research institute. In order to limit the number of alerts, we choose a margin value that is at the elbow points in a curve, which is at around the 30% mark. This is also

consistent with the threshold values used by many state-of-the-arts security appliances, *e.g.*, from Palo Alto [46], Fortinet [26] and Cisco [17]. While organizations are free to tune the threshold alerting values for each health metric to suit their environments, in this work for simplicity we will maintain it at 30%. In what follows we first examine two DNS assets that exhibited high rates of alerts (as shown in Table VII), followed by a general overview of alerts across the two organizations. We then design an inference engine that combines the health metric alerts and deduces the nature of the underlying anomaly causing these alerts using the relationships identified earlier in §V-C.

Example 1: A DNS server in the University Law Department serves as both authoritative name server and recursive resolver. It exhibited unhealthy elevated NELF metric for 83.7% of epochs, and unhealthy depressed QSRI for 94.5% of epochs, indicating its **misconfiguration** was being exploited by attackers for a potential **DDoS attack**. Queries for “`d.c.b.a.in-addr.arpa`” were coming from many external IP addresses, and the server was responding to a vast majority (over 90%) of them, thereby wasting its resources. The asset also exhibited an unhealthy LEF metrics for a non-negligible fraction (6.6%) of epochs, indicating possible proxying **scans**. On 29-Jun, this server sent queries to 131 external IP addresses, of which 18 responded – this asset is likely being utilized as a proxy to perform slow reconnaissance scans to discover the availability of DNS servers on the Internet, as analyzed and explained in detail in our other work [36].

Example 2: A DNS server in the University Engineering Department also exhibited many inbound health alerts, such as high NELF for 85.6% of epochs, low QSRI for 15.6% of epochs, and high QRI for 29.7% of epochs. Investigation confirmed that it was **misconfigured** and exploited by attackers using it to launch reflection attacks with queries for domain names such as `dnsscan.shadowserver.org`, `researchscan541.eecs.umich.edu`⁶, and `nil`. The server was also giving outbound alerts for high LEF, QSRO, and QRO, resembling **DNS exfiltration** behaviors. Indeed, our post-hoc analysis showed that on 30-Jun it sent out 709K DNS queries with pattern `SARICA[10digits].com` toward an IP address in Turkey, and on the next day, another 964K DNS queries to the same server with pattern `akbank[9digits].com.tr` – those random 9/10 digits are very likely encoded version of the exfiltrated data, as highlighted in [2].

Alerts across the two organizations: Certain DNS assets – 35% in the University and 13% in the research institute – were consistently flagged by alerts in each epoch. These turn out to be largely Authoritative DNS servers that are publicly facing, and hence exposed to inbound DNS attacks (interesting, most of these were managed by sub-departments or third-parties, rather than central IT in the organization). Recursive resolvers in both organizations raised relatively fewer alerts, typically in QRO and RRI during some epochs.

⁶Although some of the domain names seem to be designed for research-based scans, they are indeed misused by malicious actors to launch the reflection attacks on 30-Jun as discussed in Example 2.

TABLE VIII: DNS anomalies considered in this paper, their indicative alerts, and required post-hoc analysis.

DNS Anomaly Type	Indicative Alerts	Post Analysis
A1: Misconfiguration	↑ NELF & LEF	None
A2: Query DDoS	↑ QSRI & QRI	Flow profile
A3: Response DDoS	↓ QSRO & ↑ RRI	Flow profile
A4: Attack reflector	↑ QRI & RRO	Flow profile
A5: Generating scan	↑ LEF & ↓ QSRO	Flow profile
A6: Data exfiltration	↑ LEF & QRO & ↓ QSRO	Query content
A4': Reflector (after fix)	↑ QRI & RRO	Flow profile

Inferring anomalies from alerts: Tracking the health metrics (aka “symptoms”) allows us to make inferences about the underlying anomalies (aka “diseases”). We built a simple inference engine using the Codebook Correlation technique used extensively in Network Management for event correlation [33]. A causality graph (as shown in Fig. 9) was built, a codebook correlation model was derived, and then “alerts” from the 32-day dataset were looked up in the codebook to infer the underlying “anomaly” – some examples are listed in Table VIII. The outcomes, in terms of the health of the DNS assets across the two organizations, are shown in Fig. 10. By employing these health metrics, one could isolate certain assets that exhibit anomalous DNS behaviors. It is important to note that determining the actual nature of such anomalies (whether misconfiguration or security events) would certainly require further analysis of corresponding flows and/or packet contents (like works in [36] and [2]), which is beyond the scope of this paper.

Our first observation is that misconfiguration is a significant problem across both organizations – 56% and 33% of DNS assets in the University and research institute, respectively, serve DNS queries not relevant to the enterprise. This is a serious concern – Authoritative DNS servers are resolving non-enterprise queries and thereby being exposed to random queries, which can lead to denial-of-service; while recursive resolvers are resolving queries for non-enterprise hosts, thus being made available to attackers as reflectors for DDoS attacks on spoofed victims. Indeed, our analysis shows that if these DNS configurations were to be rectified, the number of DNS assets being used as reflectors falls from 25% to 3% in the University, and from 20% to 0% in the Research Institute (shown as the rightmost bars of Fig. 10(a) and 10(b)).

The second most significant concern is that there is evidence of scans emanating from both organizations, as indicated by epochs of high lookup failures (LEF) and low success of responses (QSRO). These could indicate the possibility of malware lurking within organizations and using DNS to perform scans on other Internet hosts. Further investigation of root causes and confirming whether flagged hosts are possibly infected or misconfigured are beyond the scope of this paper.

Finally, we note that there are epochs in which some data exfiltration-like symptoms [2] are seen in the DNS behavior of university assets. Such concerns have been brought to the attention of our central IT department, which is keen to obtain any sign of DNS malware across the enterprise network. Again, knowing the hosts complicit in this may




Fig. 10: Severity of DNS anomalies of each enterprise asset in both organizations.

require analysis of traffic within the organization (our traffic feed at the border does not tell us which internal host made the DNS request to the organizational cache/proxy), which is beyond the scope of this paper. Similarly, a few assets in the Research institute occasionally display anomalous volumetric patterns resembling DDoS attacks on external victims.

While we do not intend to diagnose and confirm every DNS problem, which may require a comprehensive specialized post-analysis, our method continuously assesses the general health of each DNS asset in the organization. It flags potential issues that can be investigated further by the network operator. It provides them with actionable intelligence to rectify misconfigurations, amend firewall policy rules, rate-limit query rates, etc., to better protect their DNS assets and infrastructure. As an example, our system revealed volumetric and reflection-like misbehaviors from as well as misconfigurations in poorly managed DNS servers owned and operated by affiliated entities (e.g., retail stores) and sub-departments. The IT department thereafter communicated those issues with the respective teams instead of directly taking remedial actions. Recent measurements indicate some of those behaviors have been corrected.

VI. CONCLUSION

Enterprise networks are often vulnerable to DNS-based cyber attacks due to insufficient monitoring of DNS traffic. In this paper, we have developed methods to classify enterprise assets and continuously track their cyber-health by passively analyzing DNS traffic crossing the network border of an organization. We performed a comprehensive analysis of DNS packets from two large organizations to identify asset profiles by network, functional, and service characteristics. We highlighted the behavior of enterprise hosts, either benign and anomalous. We then trained unsupervised machine learning models by DNS traffic attributes that classify the DNS assets, including authoritative name server, recursive resolver, mixed DNS server, and end-hosts behind or not behind the NAT. Lastly, we developed metrics to track the cyber health of enterprise DNS assets continuously. We identified several instances of improper configurations, data exfiltration, DDoS, and reflection attacks. Results of our real-time application have been verified with IT departments of the two organizations while revealing unknown knowledge that helps them enhance

their security management without incurring risks and excessive labor costs.

REFERENCES

- [1] J. Abley, O. Gudmundsson, M. Majkowski, and E. Hunt, "Providing Minimal-Sized Responses to DNS Queries That Have QTYPE=ANY," RFC 8482, Jan 2019, doi: 10.17487/RFC8482.
- [2] J. Ahmed, H. Habibi Gharakheili, Q. Raza, C. Russell, and V. Sivaraman, "Monitoring Enterprise DNS Queries for Detecting Data Exfiltration From Internal Hosts," *IEEE Transactions on Network and Service Management*, Sep 2020.
- [3] J. Ahmed, H. Habibi Gharakheili, C. Russell, and V. Sivaraman, "Automatic Detection of DGA-Enabled Malware Using SDN and Traffic Behavioral Modeling," *IEEE Transactions on Network Science and Engineering*, May 2022.
- [4] M. Almeida, A. Finamore, D. Perino, N. Vallina-Rodriguez, and M. Varvello, "Dissecting DNS Stakeholders in Mobile Networks," in *Proc. ACM CoNEXT*, Incheon, Republic of Korea, Dec 2017.
- [5] M. Anagnostopoulos, G. Kambourakis, S. Gritzalis, and D. K. Y. Yau, "Never say never: Authoritative TLD nameserver-powered DNS amplification," in *Proc. IEEE/IFIP NOMS*, Taipei, Taiwan, Apr 2018.
- [6] M. Anagnostopoulos, G. Kambourakis, P. Kopanos, G. Louloudakis, and S. Gritzalis, "DNS Amplification Attack Revisited," *Comput. Secur.*, Nov 2013.
- [7] M. Antonakakis, R. Perdisci, Y. Nadji, N. Vasiloglou, S. Abu-Nimeh, W. Lee, and D. Dagon, "From Throw-away Traffic to Bots: Detecting the Rise of DGA-based Malware," in *Proc. USENIX Security*, Bellevue, WA, USA, Aug 2012.
- [8] R. Arends, R. Austein, D. M. M. Larson, and R. Rose, "DNS Security Introduction and Requirements," RFC 4033, Mar 2005, doi:10.17487/RFC4033.
- [9] M. Bykova, S. Ostermann, and B. Tjaden, "Detecting Network Intrusions via a Statistical Analysis of Network Packet Characteristics," in *Proc. IEEE SST*, Athens, OH, USA, Mar 2001.
- [10] M. Bykova and S. Ostermann, "Statistical Analysis of Malformed Packets and Their Origins in the Modern Internet," in *Proc. ACM IMC*, Marseille, France, Nov 2002.
- [11] CAPEC, "CAPEC-297: TCP ACK Ping," <https://capec.mitre.org/data/definitions/297.html>, 2021, accessed: 2021-12-18.
- [12] Y. Chen, M. Antonakakis, R. Perdisci, Y. Nadji, D. Dagon, and W. Lee, "DNS Noise: Measuring the Pervasiveness of Disposable Domains in Modern DNS Traffic," in *Proc. IEEE/IFIP DSN*, Atlanta, Georgia, USA, Jun 2014.
- [13] H. Choi and H. Lee, "Identifying Botnets by Capturing Group Activities in DNS Traffic," *Computer Networks*, vol. 56, no. 1, pp. 20–33, Feb 2012.
- [14] T. Chung, R. van Rijswijk-Deij, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, and C. Wilson, "Understanding the Role of Registrars in DNSSEC Deployment," in *Proc. ACM IMC*, London, UK, Nov 2017.
- [15] CISA, "Alert (TA13-088A) DNS Amplification Attacks," <https://www.us-cert.gov/ncas/alerts/TA13-088A>, 2018, accessed: 2018-05-01.
- [16] Cisco Blog, "Overcoming the DNS Blind Spot," <https://blogs.cisco.com/security/overcoming-the-dns-blind-spot>, 2016, accessed: 2019-05-15.
- [17] Cisco Systems, "Protection Against Distributed Denial of Service Attacks," <https://bit.ly/2WUbvVk>, 2018, accessed: 2018-11-2.
- [18] Cloudflare, "What Happened Next: The Deprecation of ANY," <https://blog.cloudflare.com/what-happened-next-the-deprecation-of-any/>, 2019, accessed: 2019-6-17.
- [19] CloudFlare, "What is an ACK Flood DDoS Attack?" <https://www.cloudflare.com/en-au/learning/ddos/what-is-an-ack-flood/>, 2021, accessed: 2021-12-18.
- [20] Deloitte. (2018) Elevating Cybersecurity on the Higher Education Leadership Agenda. <https://bit.ly/36w2pLx>.
- [21] J. Dickinson, S. Dickinson, R. Bellis, A. Mankin, and D. Wessels, "DNS Transport over TCP - Implementation Requirements," RFC 7766, Mar 2016, doi:10.17487/RFC7766.
- [22] S. Dickinson, D. Gillmor, and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS," RFC 8310, doi:10.17487/RFC8310.
- [23] DPDK Project, "Developer Quick Start Guide Learn How To Get Involved With DPDK," <https://www.dpdk.org>, 2020, accessed: 2020-01-24.
- [24] EfficientIP, "A New Era Of Network Attacks," Global DNS Threat Report, 2018.
- [25] C. Fachkha, E. Bou-Harb, and M. Debbabi, "Fingerprinting Internet DNS Amplification DDoS Activities," in *Proc. NTMS*, Dubai, United Arab Emirates, Mar 2014.
- [26] Fortinet, "FortiDDoS and Verisign DDoS Protection Service," <https://bit.ly/2DsDObH>, 2018, accessed: 2018-11-2.
- [27] Y. Fu, L. Yu, O. Hambolu, I. Ozcelik, B. Husain, J. Sun, K. Sapra, D. Du, C. T. Beasley, and R. R. Brooks, "Stealthy Domain Generation Algorithms," *IEEE Transactions on Information Forensics and Security*, Jun 2017.
- [28] H. Gao, V. Yegneswaran, Y. Chen, P. Porras, S. Ghosh, J. Jiang, and H. Duan, "An Empirical Reexamination of Global DNS Behavior," in *Proc. ACM SIGCOMM*, Hong Kong, China, Aug 2013.
- [29] H. Gao, V. Yegneswaran, J. Jiang, Y. Chen, P. Porras, S. Ghosh, and H. Duan, "Reexamining DNS From a Global Recursive Resolver Perspective," *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 43–57, Feb 2016.
- [30] S. Hao, N. Feamster, and R. Pandrangi, "Monitoring the Initial DNS Behavior of Malicious Domains," in *Proc. ACM IMC*, Berlin, Germany, Nov 2011.
- [31] S. Hao, A. Kantchelian, B. Miller, V. Paxson, and N. Feamster, "PREDATOR: Proactive Recognition and Elimination of Domain Abuse at Time-Of-Registration," in *Proc. ACM CCS*, Vienna, Austria, Oct 2016.
- [32] P. Hoffman and P. McManus, "DNS Queries over HTTPS (DoH)," RFC 8484, Oct 2018, doi:10.17487/RFC8484.
- [33] S. Kliger, S. Yemini, Y. Yemini, D. Ohsie, and S. Stolfo, "A Coding Approach to Event Correlation," in *Proc. IEEE/IFIP IM*, Santa Barbara, CA, USA, May 1995.
- [34] Y. Lee and N. Spring, "Identifying and Analyzing Broadband Internet Reverse DNS Names," in *Proc. ACM CoNEXT*, Incheon, Republic of Korea, 2017.
- [35] M. Lyu, H. Habibi Gharakheili, C. Russell, and V. Sivaraman, "Mapping an Enterprise Network by Analyzing DNS Traffic," in *Proc. PAM*, Puerto Varas, Chile, Mar 2019.
- [36] —, "Hierarchical Anomaly-Based Detection of Distributed DNS Attacks on Enterprise Networks," *IEEE Transactions on Network and Service Management*, Mar 2021.
- [37] M. Lyu, H. Habibi Gharakheili, and V. Sivaraman, "A Survey on DNS Encryption: Current Development, Malware Misuse, and Inference Techniques," *ACM Comput. Surv.*, Jul 2022.
- [38] X. Ma, J. Zhang, J. Tao, J. Li, J. Tian, and X. Guan, "Dnsradar: Outsourcing malicious domain detection based on distributed cache-footprints," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 11, pp. 1906–1921, Nov 2014.
- [39] D. C. MacFarland, C. A. Shue, and A. J. Kalafut, "The best bang for the byte: Characterizing the potential of DNS amplification attacks," *Computer Networks*, Apr 2017.
- [40] S. Marshall, "CANDID: Classifying Assets in Networks by Determining Importance and Dependencies," University of California at Berkeley, Electrical Engineering and Computer Sciences, Tech. Rep., May 2013.
- [41] Microsoft Docs, "Use DNS Policy for Applying Filters on DNS Queries," <https://docs.microsoft.com/en-us/windows-server/networking/dns/deploy/apply-filters-on-dns-queries>, 2021, accessed: 2022-4-2.
- [42] J. Mirkovic and P. Reiher, "A Taxonomy of DDoS Attack and DDoS Defense Mechanisms," *ACM SIGCOMM Computer Communication Review*, Apr 2004.
- [43] G. C. M. Moura, S. Castro, J. Heidemann, and W. Hardaker, "TsuNAME: Exploiting Misconfiguration and Vulnerability to DDoS DNS," in *Proc. ACM IMC*, Virtual Event, Nov 2021.
- [44] M. Müller, G. C. M. Moura, R. de O. Schmidt, and J. Heidemann, "Recursives in the Wild: Engineering Authoritative DNS Servers," in *Proc. ACM IMC*, London, United Kingdom, Nov 2017.
- [45] NoviFlow, "NoviSwitch 2122 High Performance Open-Flow Switch," <https://noviflow.com/wp-content/uploads/NoviSwitch-2122-Datasheet-1.pdf>, 2018, accessed: 2018-28-1.
- [46] Palo Alto Networks, "DoS and Zone Protection Best Practices," <https://bit.ly/2HQOMwU>, 2018, accessed: 2018-28-1.
- [47] W. Wweyemamu, T. Lauinger, C. Wilson, W. K. Robertson, and E. Kirida, "Clustering and the Weekend Effect: Recommendations for the Use of Top Domain Lists in Security Research," in *Proc. PAM*, Puerto Varas, Chile, Mar 2019.
- [48] S. Schüppen, D. Teubert, P. Herrmann, and U. Meyer, "FANCI: Feature-based Automated NXDomain Classification and Intelligence," in *Proc. USENIX Security*, Baltimore, MD, USA, Aug 2018.
- [49] S. Son and V. Shmatikov, "The Hitchhiker's Guide to DNS Cache Poisoning," in *Proc. SecureComm*, Singapore, Sep 2010.

- [50] R. van Rijswijk-Deij, A. Sperotto, and A. Pras, "DNSSEC and Its Potential for DDoS Attacks: A Comprehensive Measurement Study," in *Proc. ACM IMC*, Vancouver, BC, Canada, Nov 2014.
- [51] J. Vijayan, "Frequency & Costs of DNS-Based Attacks Soar," <https://ubm.io/2Nxx5Cr>, 2018, accessed: 2018-05-16.
- [52] D. Yang, Z. Li, and G. Tyson, "A Deep Dive into DNS Query Failures," in *Proc. USENIX ATC*, Virtual Event, Jul 2020.



Minzhao Lyu received his B.Eng. (First Class Hons.) and Ph.D. degree from the University of New South Wales, Sydney, Australia in 2017 and 2022 respectively. He has worked at CSIRO's Data61, Sydney, Australia as a student fellow and at National Telemedicine Center of China as a research intern. He is currently a Postdoctoral Research Associate at the University of New South Wales, Sydney, Australia. His research interests include network data analytics, network security, programmable networks, and applied machine learning.



Hassan Habibi Gharakheili received his B.Sc. and M.Sc. degrees of Electrical Engineering from the Sharif University of Technology in Tehran, Iran in 2001 and 2004 respectively, and his Ph.D. in Electrical Engineering and Telecommunications from the University of New South Wales (UNSW) in Sydney, Australia in 2015. He is currently a Senior Lecturer at UNSW Sydney. His research interests include programmable networks, learning-based networked systems, and data analytics in computer systems.



Craig Russell received his Ph.D. in Applied Mathematics from Macquarie University, Sydney in 1997. He is currently Director of Engineering at Canopus Networks and Adjunct Senior Lecturer at UNSW. He was a principal research engineer at CSIRO's Data61, Sydney, Australia, and has previously held commercial roles in the telecommunications and software industries. He has design, implementation and operational experience in a wide range of advanced telecommunications equipment and protocols as well as experience in developing software applications.

His research interests are in software-defined networking and the application of machine learning techniques to solve problems in network security.



Vijay Sivaraman received his B. Tech. from the Indian Institute of Technology in Delhi, India, in 1994, his M.S. from North Carolina State University in 1996, and his Ph.D. from the University of California at Los Angeles in 2000. He has worked at Bell-Labs as a student Fellow, in a silicon valley start-up manufacturing optical switch-routers, and as a Senior Research Engineer at the CSIRO in Australia. He is now a Professor at the University of New South Wales in Sydney, Australia. His research interests include Software Defined Networking, network architectures, and cyber-security particularly for IoT networks.

network architectures, and cyber-security particularly for IoT networks.