

**TOWARDS UNDERSTANDING THE LIFECYCLE OF MALICIOUS NETWORK
INFRASTRUCTURE**

A Dissertation
Presented to
The Academic Faculty

By

Athanasios Avgetidis

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
Computer Science
College of Computing

Georgia Institute of Technology

August 2025

Copyright © Athanasios Avgetidis 2025

TOWARDS UNDERSTANDING THE LIFECYCLE OF MALICIOUS NETWORK INFRASTRUCTURE

Thesis committee:

Dr. Manos Antonakakis, Advisor
School of Electrical and Computer Engineering & School of Computer Science
Adjunct
Georgia Institute of Technology

Dr. Angelos Keromytis, Co-Advisor
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Fabian Monroe
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Roberto Perdisci
School of Computing
University of Georgia

Dr. Alberto Dainotti
School of Computer Science
Georgia Institute of Technology

Date approved: June 3rd 2025

”Time is the wisest counselor of all”

Pericles

For my family

ACKNOWLEDGMENTS

Throughout my PhD journey, I was fortunate enough to meet and work with many outstanding people to whom I owe deep gratitude. First of all, I would like to thank my advisor, Manos Antonakakis, who has been there all along, showed me how to do research, battle through adversity, and most importantly, be a better person. I would also like to deeply thank my co-advisor, Angelos Keromytis, who assisted me along the process, has always been helpful and supportive, and has surprised me with his out-of-the-box way of thinking and innovative ideas.

Aside from my advisors, I was lucky enough to work with leading faculty in my area, from whom I have learned a lot. I want to thank Fabian Monrose, who made me a better researcher by paying attention to detail and communicating my thoughts and research more effectively. I am also deeply grateful to Roberto Perdisci, who taught me to keep high standards and be more thorough in my work. Additionally, I want to thank Zane Ma for assisting me in various parts of my research and helping me see my ideas in a new light.

My research would not be possible without my collaborators, who have been a big support in my research, day-to-day task execution, "brainstorming", and more importantly, "venting". I want to especially thank my friend and collaborator Panagiotis Kintis, who has been there since my first days in research and has helped me in every part of my PhD journey. I am also deeply grateful to Omar Alrawi and Chaz Lever for advising, assisting me, and enabling much of my research. My special thanks to all collaborators, co-workers, and friends in the lab, Thomas Papastergiou, Aaron Faulkenberry, Kevin Valakuzhy, Thanasis Kountouras, Logan O'Hara, Tillson Galloway, Vinny Adjibi, Kevin Sam Tharayil, Athanasios Moschos, Georgios Kokolakis, Miuyin Yong Wong, Konstantinos Karakatsanis, Kleanthis Karakolios, Vaggelis Froudakis, William Garrison, Alex Neal, Eligio Makransky, Michael Mitchel, Joseph Reilly, and Nikkia Green. I will remember your help, and you made this journey more meaningful.

I would like to thank and dedicate this dissertation to my family and my dedicated partner, Glacy, who have helped me at various challenging steps of PhD life and have assisted me all along. Last but not least, I want to thank the committee members who, with their suggestions, helped to make this dissertation better and more cohesive.

TABLE OF CONTENTS

Acknowledgments	v
List of Tables	xi
List of Figures	xiii
Summary	xvi
Chapter 1: Introduction	1
1.1 Introduction	1
1.2 Thesis Statement	5
1.3 Contributions	6
1.4 Dissertation Overview	8
Chapter 2: Background and Previous Work	9
2.1 Background	9
2.1.1 Malicious Network Infrastructure	9
2.1.2 The Domain Name System (DNS) Hierarchy	11
2.2 Previous Work	13
2.2.1 Domain Name Lifecycle and Characterization	13

2.2.2	Malware and Upper DNS Hierarchy Infrastructure Characterization Studies	14
2.2.3	Cybercrime and Threat Actor Characterization Studies	15
2.2.4	APT and Sophisticated Attack Detection and Characterization	15
Chapter 3: Understanding Malware Infrastructure from the Upper DNS Hierarchy		17
3.1	Motivation	17
3.2	Datasets and Methodology	19
3.2.1	Datasets	19
3.2.2	Methodology and Validation	21
3.3	Hosting Infrastructure	24
3.4	Malware Clients	27
3.4.1	Industry Sectors	29
3.5	Malware Lifecycle	32
3.6	Vantage Point Comparison	35
3.6.1	Measurement Planes	35
3.6.2	Comparison	38
3.7	Summary	40
Chapter 4: Understanding the Interaction of Malicious Actors With Their Infrastructure Through an Empirical Analysis of Password Stealers		42
4.1	Motivation	42
4.2	Data and Methodology	44
4.2.1	Data Sources	44

4.2.2	Data Validation	48
4.3	Ethical Considerations	55
4.4	Analysis Results	56
4.4.1	Stealers on the Internet	56
4.4.2	Characterization of Operators	63
4.5	Discussion	67
4.6	Summary	69
Chapter 5: Understanding the Lifecycle and Infrastructure of APT Domain		
	Names	71
5.1	Motivation	71
5.2	Challenges in Domain Lifecycle Analysis	74
5.2.1	Challenges	75
5.2.2	Placing The Challenges In Context: The SolarWinds Attack	77
5.2.3	Observations and Takeaways	78
5.3	Datasets and Methodology	79
5.3.1	OSINT Datasets	79
5.3.2	DNS Datasets and Threat Reports IP Visibility	82
5.3.3	Measurement Methodology	84
5.4	Evaluation	90
5.4.1	Training and Evaluation Datasets	90
5.4.2	Experimental Results	93
5.4.3	Infrastructure Expansion and Lifetime Characterization	96
5.5	Infrastructure Analysis	97

5.5.1	Infrastructure Utilization	98
5.5.2	Infrastructure Lifecycle	102
5.6	Discussion	107
5.7	Summary	107
Chapter 6:	Conclusion	109
6.1	Summary	109
6.2	Limitations	111
6.2.1	Malware Infrastructure Study Limitations	111
6.2.2	Password Stealers Study Limitations	112
6.2.3	APT Domain Study Limitations	113
6.3	Future Work	114
6.4	Closing Remarks	117
Appendices		118
Appendix A:	Atropos Generalization For Threat Hunting	119
References		120

LIST OF TABLES

2.1	Methodologies and systems characterizing the lifecycle of malicious domain names. Atropos, presented in chapter 5, is the first methodology to historically identify the likely attack-utilized IP addresses of known malicious domain names.	13
3.1	Top 15 malware families based on the number of malicious domains in our dataset.	23
3.2	Unique clients querying malicious domains and the number of malware families in each industry (ISIC section).	29
3.3	Client distribution across sectors for seven-day samples starting 2017-03-01, 2018-03-01, 2019-03-01, and 2020-03-01. <i>All</i> represents the complete <i>AuthDNS</i> dataset while <i>Mal</i> represents only malicious domains.	30
3.4	Request characterization for three temporal windows. While most clients and sectors are observed between malicious domain detection and expiration/takedown, many client ASNs and countries (ASCCs) first connect after expiration/takedown.	33
3.5	Most popular ASNs first observed in each temporal window. Scanners and AV vendors appear mostly during and after the detection of a malicious domain, while hosting networks are most prevalent during the setup of the domain.	33
4.1	A list of data sources used in this study.	44
4.2	A list of top password stealers found in our dataset.	45
4.3	<i>Stealer</i> dataset fields summary.	47
4.4	Top 10 user agents and related statistics.	50

4.5	Top 10 TLDs and hosting networks for panel hosting server domains.	58
4.6	Networks resolving stealer domains by country for residential, business, and government networks.	60
4.7	Top 10 hosting networks querying stealer domains.	61
4.8	Top panel operator device types and operating systems.	64
4.9	Top 10 countries of operator IP addresses and their proxy and tor networks.	66
5.1	Coverage of IOCs for the top 10 publishers in terms of reports. Overall, we utilize a total of 2,188 APT reports.	81
5.2	Datasets utilized in the study.	81
5.3	A and AAAA resource record visibility after enriching the known APT domain names with our DNS data sources. APT IPs appearing on threat reports can only characterize 23.52% of APT FQDNs in popular DNS datasets.	82
5.4	The features of Atropos. Atropos utilizes 22 features from four distinct classes.	86
5.5	Average 10-fold cross-validation performance of Atropos on the PR dataset. Atropos achieves at best a 99.86 ROC AUC score when utilizing Virus Total DNS data and training on the PR dataset utilizing a Random Forest Model.	93
5.6	Out-of-distribution test set evaluation of Atropos. Atropos achieves an over 91.00% accuracy across the two evaluation datasets, demonstrating generalization.	94
5.7	Atropos MDI Feature Importance when trained on PR dataset and utilizing Active DNS data with an 80-20% split.	96
5.8	Number of network IoCs associated with the actors from the OSINT threat reports and identified by Atropos for the top 10 actors, and overall. Atropos provides three times the IP visibility of threat reports and contextualizes three times more domain names than threat reports.	96
A.1	Evaluation of Atropos trained with and altered PR dataset.	119

LIST OF FIGURES

2.1	The hierarchy of the Domain Name System (DNS). The different vantage points of DNS have often been utilized as measurement planes to study the Internet as well as its security.	11
3.1	Distribution of the number of malware samples and servers associated with each domain in <i>AuthDNS</i>	24
3.2	Daily Volume <i>AuthDNS</i> for malware domains. ECS-enabled requests (orange) average 17.9% of daily requests.	24
3.3	Geographic distribution of the IP addresses of hosting and clients for malware-related domains. Darker colors indicate more malware families are associated with the country through the hosting server (left) or querying client (right).	25
3.4	Distribution of the number of malware families per IP, Prefix, ASN, and country that have resolved malicious domains in <i>AuthDNS</i> . Most network infrastructure and targeted networks are not strongly correlated with a single malware family.	25
3.5	Correlation between different measures of hosting infrastructure. Higher numbers of samples per malware family correlate with more domains. Higher domain utilization correlates with more server IPs and more hosting countries.	26
3.6	Spearman correlation between different measures of potential victims. Higher numbers of clients querying for malware family-related domains correlate strongly with a more diverse set of impacted countries and economic sectors.	28

3.7	Malware measurement can occur 1) in the global DNS plane, 2) at or near the local client, or 3) at the malware infrastructure. Each location has specific sub-components that interact via request (filled arrow) and response (empty arrow) protocols (e.g., DNS, C2 protocols). Existing research has studied many of the depicted components with host-based (yellow) or network-based (blue) techniques.	36
4.1	An overview of <i>Stealer</i> data collection.	46
4.2	Distribution of <i>AME</i> per top largest group.	53
4.3	Panel signature generation and identification.	54
4.4	Distribution of panels and associated malware.	57
4.5	Distribution of domain events and detection.	58
4.6	Distribution of the time delta for events and detection.	59
4.7	The diurnal analysis for the top 20 countries of operator device activity (dark more active and light less active).	65
5.1	Lifecycle of an actor-controlled domain name. Multiple owners and infrastructure types complicate forensics.	73
5.2	Domain and IP lifecycle of <i>deftsecurity[.]com</i> and <i>incomeupdate[.]com</i> sunburst domain names initially reported in [2]. In this work, we seek to automatically identify the actor-utilized IPs (colored in green). The numbers inside the parentheses reflect the number of unique IPs of each category. . .	75
5.3	An overview of <i>Atropos</i> . <i>Atropos</i> utilizes OSINT datasets and historical DNS data to label and filter APT infrastructure in a 3-step process.	85
5.4	Daily active APT actors for the top 15 most utilized provider ASes in the last decade. Cloudflare utilization for domain name hosting has increased drastically over the years, making forensic analysis and attribution of IP infrastructure harder.	97
5.5	AS re-utilization among APT actors.	99
5.6	Number of actor-utilized IPs per country for the top affiliated countries of the APT actors.	100

5.7	Number of IPs per category of infrastructure for the top affiliated countries of the APT actors.	101
5.8	The unique lifecycle of the infrastructure associated with APT domains compared to the first public report date.	103
5.9	Number of days that actor-utilized IPs were first and last observed before their domain name public disclosure.	105

SUMMARY

Network infrastructure is an important component of malicious cyber operations. From novice attacks conducted by script kiddies to highly sophisticated threats backed by nation-states, network infrastructure is being utilized for command and control, data exfiltration, malware hosting, and social engineering, among others. Over the years, while there have been several studies that have focused on detecting, blocking, and characterizing malicious infrastructure, the temporal dynamics of how this infrastructure changes over time and the characteristics of the stakeholders interacting with it have often been overlooked. This thesis shows that the temporal analysis of malicious infrastructure reveals network attributes that can characterize the stakeholders that interact with it. The systematic analysis of such network attributes can aid the accurate discovery of previously unreported malicious infrastructure and increase our awareness of the behaviors of the stakeholders that interact with it.

Through longitudinal empirical studies and novel methodologies, this thesis demonstrates the importance of accounting for the temporal dynamics of malicious network infrastructure. Specifically, it introduces a novel methodology that accurately identifies historically utilized IP infrastructure from domain names of sophisticated threats, which expands the publicly reported IP knowledge by 3.06 times. It also showcases how the temporal analysis of malicious network infrastructure can help threat analysts and security practitioners better understand the quantitative distributions of the network interactions of the stakeholders (i.e., scanners, security vendors, victims, and threat actors). More precisely, this thesis pinpoints the minimum network log retention window for uncovering at least 90% of the infrastructure of sophisticated attacks down to 25 months and characterizes for the first time the lifecycle of network requests into malware-related domain names from the upper DNS hierarchy. These insights have applicable takeaways for log retention policies for network data and victim and infrastructure analysis studies using DNS datasets.

CHAPTER 1

INTRODUCTION

1.1 Introduction

Cyberattacks rely heavily on network infrastructure in order to be successful. From the early days of Back Orifice and Netbus, [1] to the SolarWinds supply chain attack [2], network infrastructure is pivotal to facilitate key objectives of many cyberattacks, such as remote access, command and control (C2), hosting of malicious and phishing content, and data exfiltration. Phishing pages [3], malware C2s [4], proxy nodes [5, 6], and reconnaissance scanners [7] are common ways that malicious threat actors choose to use network infrastructure for the conduct of cyberattacks. This network infrastructure is usually comprised of servers and computers that are purchased, leased, or even hacked by malicious threat actors and can span multiple geographies, even in the course of a single attack.

During or after the identification or exposure of cyberattacks, researchers and forensic analysts conduct investigations in order to understand the breadth and depth of such attacks, to block and sanitize infrastructure, and often in cases of large or targeted attacks, to identify their source (attribution). Such forensic investigations often uncover new attack infrastructure [8] or unveil new findings regarding the velocity of such attacks [9] and the timeliness of the security response against certain threats and types of attacks [10, 11]. The results of such analyses enable more comprehensive cyberattack responses, offer a better understanding of the lag and the gaps of the security community, and deduce where more resources need to be invested.

Despite their usefulness, conducting such analyses and investigations is significantly challenging. Researchers and analysts often need to have access to historical datasets that provide sufficient coverage of the threats they are investigating, while at the same time,

they need to understand the limitations and biases that these datasets introduce in order to come to accurate conclusions. Additionally, the introduction of new network protocols and changes to existing policies can reduce the value of previously important methodologies and datasets (e.g., WHOIS with the adoption of GDPR [12] or below the recursive DNS visibility with the adoption of DNS over HTTPS (DoH [13])). Furthermore, network analysts face the additional challenge that the association of network infrastructure with malicious threat actors is often transitory, and further analysis is usually needed to pinpoint the time window that the threat actors likely used it. Acknowledging these challenges, previous measurement studies of malicious infrastructure have provided significant insights for the security community, however, they have mainly focused on studying particular malware threats [9, 14, 11], certain types of attacks and targets [10, 15, 16], or had visibility into only certain parts of the infection lifecycle [17]. More importantly, despite the plethora of prior works, not enough emphasis has been given to understanding the temporal dynamics of malicious network infrastructure with the key stakeholders that interact with it (i.e., malicious threat actors, cyberattack victims, researchers, security professionals, and scanners). This dissertation characterizes the lifecycle of malicious network infrastructure through three longitudinal measurement studies, with an emphasis on understanding the interactions of the malicious threat actors with it and discovering the time window it was first provisioned and utilized in cyberattacks.

In order to narrow this gap between prior work and the challenges of measuring important aspects of malicious infrastructure, in our first work, we conduct the first large-scale analysis of the malware ecosystem through the lens of an understudied vantage point, that of Authoritative DNS (AuthDNS) of a popular domain name registrar. This vantage point exhibits unique advantages that are lacking from other commonly used datasets or previous studies. AuthDNS features a wide and geographically diverse malware hosting and infection visibility, a full temporal view of malicious domains from registration to detection and post-takedown, while retaining DoH and DoT request visibility as it is placed above the

recursive. Utilizing this vantage point, we observe malware heterogeneity (202 families), global infrastructure (399,830 IPs in 151 countries), and infection (40,937 querying Autonomous Systems (ASes)) visibility, as well as breadth of temporal coverage (2017–2021). Our analysis reveals that malware communications are temporally sensitive: over 90% of ASes first query a malicious domain after public detection, and a median of 38.6% ASes only query after domain expiration or takedown, highlighting that client infection estimation studies need to account for non-victim traffic on every phase of the malware lifecycle. This highlights that forensic analysts conducting infection estimation tasks need to be aware of scanners and security vendors, mainly after public detection. To fit AuthDNS into the broader context of malware research and forensic analysis, we compare AuthDNS with other vantage points on four qualitative aspects and discuss their advantages and limitations. Ultimately, we establish AuthDNS as a unique and valuable measurement and analysis perspective that enables a wide-reaching view of historic and emerging threats of the malware ecosystem.

While the AuthDNS study has provided us with a wide and high-level view of the global malware infection lifecycle and the interactions of various stakeholders with it, it did not provide insights into how malicious threat actors interact with their infrastructure throughout their operations. A more in-depth look at such interactions is critical for both forensic analysts and researchers because they can lead to insights that could enable more timely responses and a better understanding of the attacker’s incentives and operations. To answer such questions, we collaborated with Malbeacon[18], a threat intelligence company, and analyzed the activities of Password Stealers (*Stealers*). *Stealers* is a leading source for gathering stolen credentials, which are heavily being used in modern cyberattacks [19]. Malbeacon tracks the interactions of cybercriminals with their C2 panels, offering visibility on the operators’ HTTP requests with their infrastructure. In this work, we conduct the first longitudinal study of *Stealers* and their operators spanning over 20 months and 4,500 operator devices. We find that operators heavily use proxies, including traditional VPNs,

residential proxies, mobile proxies, and the Tor network, when managing their botnet. We find that on average, public blocklists detect Stealer domains 74 days after the initial registration, with a median of 11 days, while almost 70% of the operators stop accessing their panels within a month of the first public detection. This detection gap gives *Stealers* ample time to infect and harvest credentials from various networks highlighting that the security community needs to invest more resources in more promptly addressing such threats as it has a significant impact on the continuation of their operations.

While the aforementioned works have provided us with unique insights into the temporal behavior of malware communications as well as a deeper understanding of the network modus-operandi of malicious threat actors, they have only partly characterized the infrastructure lifecycle, as they have mainly focused on communications towards the infrastructure servers –targeting– and the interactions of the actors with their servers while their operations were active on certain threats. As we have previously discussed, one of the major problems analysts face is that they often have to manually investigate when malicious infrastructure has been active, which has been regarded as ”a manual and time-consuming process” in prior work [20]. One of the leading reasons behind this is that Indicators of Compromise (IoC), among which are malicious infrastructure such as domain names and IPs, are usually shared in bulk without much or any temporal context and thus require manual inspection. However, this temporal context can enable several capabilities for analysis such as: identifying how long different actors take to provision infrastructure, identifying unknown historical infrastructure that was not mentioned publicly, and measuring the security response post-identification of a threat. To this end, we propose Atropos, a novel supervised learning system that automatically and accurately discovers the actor-controlled infrastructure and active window of historical domain names based on publicly available datasets. We apply Atropos, to over 31.000 domains of Advanced Persistent Threat (APT) attacks, which are particularly hard to detect [21]. Atropos increases our knowledge of the APT IP infrastructure by 3.06 times, compared to that of public threat reports. We evaluate

Atropos on two expert-labeled datasets and find it to be highly accurate, with a mean 10-fold cross-validation accuracy of 98.90%, and robust against practical adversarial attacks. Utilizing the discovered infrastructure that Atropos provides, we conduct the largest APT infrastructure characterization study to date, spanning over a decade. We find that APT actors reuse the same networks across multiple years, they frequently utilize infrastructure that is close to their targets, and APT attacks take on average 317 days to be reported from the first day of attack infrastructure provisioning. Finally, we recommend network operators retain network traffic logs for at least 19 months, to improve to probability of detection for the majority of an APT infection infrastructure.

Through our longitudinal empirical studies and our proposed methodologies, we demonstrate the importance of taking into account the temporal dynamics of malicious network infrastructure. We are the first to present a methodology that accurately identifies historically utilized IP infrastructure from domain names of sophisticated threats and expands the publicly reported IP knowledge by 3.06 times. We also showcase how taking the temporal dynamics of malicious infrastructure into consideration, can help us better understand the quantitative distributions of the network interactions of the stakeholders (i.e., DNS scanners, security vendors, victims, and threat actors), which give us important takeaways regarding how long threat actors utilize their infrastructure, how long before the public reporting of an attack they provision their IP addresses and how victim analysis utilizing DNS datasets could be impacted.

1.2 Thesis Statement

The systematic characterization of malicious network infrastructure in historical network datasets by taking into account its temporal dynamics increases our understanding of the network communications of the stakeholders that interact with it and enables a more comprehensive identification of sophisticated attack infrastructure. Specifically, this thesis demonstrates (1) a novel methodology that enables the accurate discovery of over 3 times

more IP addresses that are associated with the activities of sophisticated threat actors compared to those that are publicly reported and (2) pinpoints the minimum network log retention window for uncovering at least 90% of the infrastructure of sophisticated attacks down to 25 months.

1.3 Contributions

The first large-scale characterization of the lifecycle and infrastructure of malware-related domain names from the upper DNS hierarchy. Utilizing a longitudinal dataset from a popular DNS registrar spanning four years and 12,212 registered domains hosting 202 unique malware families, we measured the lifecycle of the networks that have historically queried these domain names. While malware hosting domain names are mainly distributed among countries with large hosting providers, the distribution of the clients querying them is much wider, with 71.3% of the malware families hosted in such domains being queried by over 100 countries. Our lifecycle analysis unveils that on the median malware hosting domain name, 54.5% and 62.5% of the unique Autonomous Systems (ASes) and countries, respectively, will first be observed after its detection. Even more alarmingly, 38.6% of the unique ASes will be first observed after take-down or expiration highlighting the importance of the discovery of DNS scanners and security vendors – interacting with such domains after detection – towards the accurate estimation of potentially infected populations querying such domain names.

The first quantification of the lifecycle of the interactions of cybercriminals with their password-stealing management infrastructure. By collaborating with a security vendor, we were able to study the lifecycle of the interactions of cybercriminals with their infrastructure that manages infected victims of password-stealing malware. We find that cybercriminals provision their infrastructure fast, within two weeks after registration, but it takes on average 64 days for the security community to detect them. Our lifecycle analysis indicates that 69.03% of the cybercriminals operators stop accessing their management

infrastructure within 30 days of detection, highlighting how important such an event is to curb the operations of cybercriminals. We suggest that the community should frequently scan known domain names and IPs for the appearance of such infrastructure panels in order to reduce the 64-day lag of detection.

The first system that accurately identifies and contextualizes new IP infrastructure associated with sophisticated cyber actors that is not publicly reported. We developed a new system, which we call Atropos, that combines historical DNS logs and threat report information to identify new, unreported IP infrastructure of known domain names of Advanced Persistent Threats (APTs) and contextualizes their lifecycle. We deploy Atropos on a dataset of 31,398 APT domain names and over 120 million DNS records spanning over a decade from April 2013 to January 2025. Atropos is able to accurately uncover 3.06 more IP addresses related to known APT domain names, compared to those publicly shared on APT reports, and achieves an accuracy of over 91.00% in two evaluation datasets. Additionally, Atropos can characterize and contextualize the lifecycle of three times more domain names than those that can be characterized only with the infrastructure publicly provided in threat reports.

The first large-scale characterization of the lifecycle and infrastructure of sophisticated cyber actors. Utilizing the network infrastructure provided in public threat reports and that which our novel system, Atropos, provides, we conduct the largest and most comprehensive APT infrastructure analysis to date, spanning over a decade and 405 APT actors. We quantify the time window during which organizations need to keep network logs to identify the vast majority of the infrastructure of an APT attack. Our results show that the network logs should be preserved for at least 19 to 25 months. Furthermore, we find that while APT actors utilize a plethora of different hosting providers, they only re-use a small portion of them, and that over the years, the use of cloud-fronting has increased significantly. These findings verify expert knowledge [22] and make network forensics and attribution harder.

1.4 Dissertation Overview

The remainder of this dissertation is organized as follows. Chapter 2 presents the background and foundational previous work that contextualizes the concepts and ideas described in the next sections. Chapter 3 presents the characterization of the infrastructure and lifecycle of malware-hosting domain names. It describes the methodology we utilized to identify domain names related to malware, exposes the difference between the infrastructure that hosts these domain names and the infrastructure querying them, and lastly characterizes the distribution of the querying networks across the lifecycle of each malware domain name. In Chapter 4, we describe our empirical analysis and characterization of the interactions of cybercriminals with the management infrastructure of password-stealing operations. We detail the HTTP dataset as well as the merging methodology we utilize to identify operator devices, and we characterize both the infrastructure and the lifecycle of these operators across a period of 20 months. Chapter 5 details our novel tool, Atropos, which accurately expands the IP infrastructure of known APT actors by a factor of 3.06. Additionally, we present the largest infrastructure characterization study of APT actors to date, spanning over a decade and 405 APT actors. Lastly, Chapter 6 concludes the thesis, summarizing the contributions, detailing the limitations of the presented work, and presenting to the reader the closing remarks.

CHAPTER 2

BACKGROUND AND PREVIOUS WORK

2.1 Background

2.1.1 Malicious Network Infrastructure

Since the early days of Back Orifice and Netbus, [1], the network has played a pivotal role in the successful conduct of cyberattacks. Network and Internet access is what enables remote access tools and malware to be controlled by their operators[23], worms to spread and infect other computers[24], and distributed denial of service (DDoS) attacks to be effective in taking down critical infrastructure [9].

The hourglass model design of the Internet emphasizes the role of Internet Protocol (IP) addresses in network communications, and as such, their importance towards network security. IP addresses characterize the network infrastructure that malicious threat actors utilize for a plethora of important tasks, from reconnaissance (e.g., scanning for potential victims) to command and control (C2) of their infected bots, exfiltration of sensitive information, hosting social engineering, and luring content, down to simple access to the Internet. As such, IP addresses have become an important staple of network defense that is often used to blacklist but also characterize all of the aforementioned kinds of malicious network infrastructure.

Despite their importance, IP addresses lack dynamicity and human interpretability. The Domain Name System (DNS), often regarded as the phonebook of the Internet, maps IP addresses to human-understandable domain names (e.g., google.com), but also offers the ability to dynamically change the IP addresses that map to a domain name. This capability has been widely utilized by malicious threat actors who choose to utilize domain names to dynamically change their pool of IP addresses and make their operations more resistant to

IP blocklists. Additionally, it has opened new doors of abuse and evasion, with threat actors depending on the similarity of their utilized domain names with those of popular legitimate domain names in order to trick unsuspected victims into visiting their malicious infrastructure and threat analysts to ignore it during investigations. Due to their popularity among cyber attacks, domain names are also an important and common indicator in network defense that is being used to blocklist and characterize malicious network infrastructure.

The rise of web browsers and web content made even more specific indicators important for the Internet, but also for cybercriminals. The Uniform Resource Locator (URL) is a key mechanism that is used to retrieve specific content from IP addresses and domain names. A URL is a unique indicator that is used to refer to specific content of a network address (e.g., an image or a script of a website). In cybersecurity, URLs are also major indicators of malicious network infrastructure that is being used to refer to specific malicious or social engineering content "under" a website, with phishing detection and blocklisting usually heavily relying on URL blocklists. Cybercriminals also usually surgically compromise and insert malicious content in specific URLs of otherwise benign domain names and web pages [25, 26], abusing their residual trust.

All three aforementioned indicator types of malicious network infrastructure: IP addresses, domain names, and URLs, are highly important for daily cybersecurity tasks, such as blocking malicious traffic in firewalls, identifying and taking down the source of infections [27], and detecting new sources of abuse. Their importance to network security is also demonstrated by the fact that they are the three most popular shared network indicators in malicious threat reports [28], which are frequently being used by security practitioners to bolster the network defenses and expand our understanding of security threats. This dissertation explores and characterizes the malicious network infrastructure utilizing these three key indicator types.

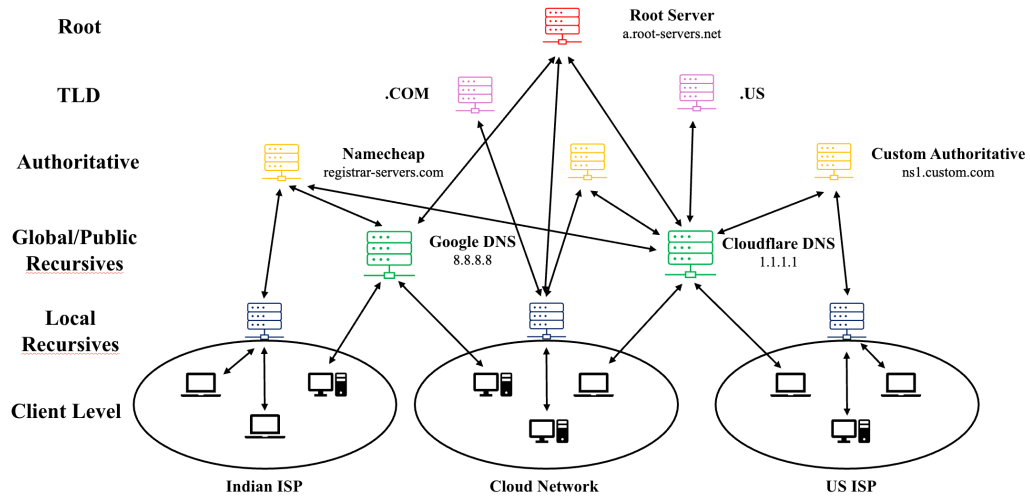


Figure 2.1: The hierarchy of the Domain Name System (DNS). The different vantage points of DNS have often been utilized as measurement planes to study the Internet as well as its security.

2.1.2 The Domain Name System (DNS) Hierarchy

The domain name system (DNS) [29] is one of the core components of the Internet. DNS translates semantic domain names that offer more context and are more memorable to Internet users into IP addresses, making it a useful data source for observing Internet communication. For more details and a broader view of DNS, we refer the reader to [30, 31], and in this section, we will focus on the DNS hierarchy to provide a background of the DNS vantage points that have been utilized in this dissertation.

The structure of the domain name system is hierarchical and can be summarized in the illustration of Figure Figure 2.1. At the bottom of the hierarchy of DNS(Client Level), we find endpoints often referred to as clients, which are individual computers and servers that issue DNS requests in order to get the IP address of a server they want to visit. These clients can belong to different networks and Internet Service Providers (ISPs) that are distributed around individual countries and the world. The DNS requests of these clients typically reach the next level of the hierarchy, which is the DNS recursives. The DNS recursives are tasked to find and ask the relevant DNS nameservers in order to get an answer for the

domain names that the clients have asked them. DNS recursives can be grouped into two categories, as showcased in Figure Figure 2.1. Local recursives are typically DNS recursives that exist either on the same or a nearby network of the clients that issue the DNS requests (e.g., the DNS recursive of the ISP of the clients). Global or public recursives typically exist outside of the network of the querying clients and can be set up to be the preferred recursives of the users or be utilized by specific software (e.g., Google DNS in Chrome browsers). The visibility of local and global recursives in querying clients is obviously different, as the local recursives will have a narrower breadth that is limited mainly to the clients that exist on the same network or ISP, while global and public recursives can have visibility into any client that chooses to query them. Assuming that DNS recursives do not know the answer to the DNS question that the querying clients have asked them (due to DNS caching), they will then ask the root DNS servers that exist at the top of the hierarchy. The root DNS servers typically will not have the answer for such granular information and will therefore redirect the DNS recursives to the TLD nameservers. The TLD nameservers have a narrow scope of knowing the authoritative DNS nameservers of the zone (i.e., TLD) they serve. Next, the TLD nameservers will again redirect the DNS recursive to the authoritative DNS server that has the answer for the individual domain names that the querying clients have asked for. The authoritative DNS nameservers can be either big nameservers of DNS registrars or DNS management services (e.g., Namecheap or CloudFlare), or custom nameservers that the owners of the domain names can choose to provision. Big authoritative nameservers like DNS registrars can have access to all the domain names of the users that utilize them (e.g., all domain names registered in a DNS registrar and not moved to another authoritative DNS server), while smaller custom authorities will only have limited visibility to the domain names that exist in their "zone". During this dissertation, we utilize DNS datasets from both the authoritative DNS servers of a popular DNS registrar (i.e., in chapter 3) and that of a local recursive DNS scanner (i.e., ActiveDNS [32] in chapter 5) that scans millions of domain names of over 1,100 gTLDs

every day. In chapter 3 and chapter 5 we discuss how these different DNS vantage points offer us adequate and even more comprehensive DNS visibility compared to prior work in order to study the lifecycle of malicious infrastructure.

2.2 Previous Work

2.2.1 Domain Name Lifecycle and Characterization

Table 2.1: Methodologies and systems characterizing the lifecycle of malicious domain names. Atropos, presented in chapter 5, is the first methodology to historically identify the likely attack-utilized IP addresses of known malicious domain names.

Prior Work	Is Historically Applicable	Characterizes Individual IPs
Lever [33]	✓	-
Affinito [34]	✓	-
LLoyd [35]	-	✓
Sebastian [36]	-	-
Atropos	✓	✓

Understanding the lifecycle of domain names has been the subject of prior works in the security and measurement community. Lever et al.[33] offered an alternative to WHOIS and tried to identify domain ownership changes using Alembic, a lightweight algorithm that utilized passive DNS data. However, their methodology was aimed at identifying changes of ownership and not the utilized IP infrastructure of malicious domains. Affinito et al. [34] studied the lifecycle of domains and malicious domains in blocklists utilizing zone file data and, similarly to Lever, developed a methodology to bound the lifecycles of domain names, but not to label their infrastructure. Lloyd et. al. [35] developed a methodology to classify domain names as "active", "no-IP", or "inactive", with an aim to find domain names serving content under the registrant's control. However, this methodology is not applicable to historically malicious domain names and mainly relies on parking infrastructure lists that are not sufficient to incorporate unknown parking IPs and sinkhole IPs that some malicious domain names are usually pointed at after detection. Sebastian et al. developed

an automated approach to attribute domain names to their most likely owner [36], however, this work did not classify the IP infrastructure associated with a domain name. In this thesis, in chapter 5 we present the first methodology that identifies the IP infrastructure and the time window in which historically reported malicious domain names of APT attacks were most likely provisioned and utilized.

2.2.2 Malware and Upper DNS Hierarchy Infrastructure Characterization Studies

Prior works using passive authoritative DNS data have focused on detection and measurement in the upper DNS hierarchy. Antonakakis et. al., [37] proposed Kopsis, a supervised learning system that passively monitors domain names in the upper DNS hierarchy and detects malware-related domain names. Thomas et al., [38] analyzed the DNS traffic of several TLD nameservers and identified strongly connected components related to malware domain names. Hao et. al., [39] studied the initial DNS requests of malicious domain names and found that most domain names are involved in attacks shortly after registration. However, none of these studies focused on characterizing the malware infrastructure of diverse malware across the years.

Other studies have focused on characterizing malware-related domain names and their network infrastructure, utilizing recursive, sinkhole, and endpoint agent vantage points. Lever et. al., [40] studied the infrastructure of 26.8 million malware samples and found malware to re-use IP infrastructure. Kotzias et. al., [41] studied the impact of different malware on over 28 thousand enterprises for nearly three years. Rezaeirad et. al., [17] profiled the stakeholders visiting malware infrastructure after sinkholing, and identified that 99% of the IP addresses are not victim-related. However, none of these works explored the communications of malware infrastructure in an end-to-end lifecycle perspective from registration to post-takedown. In chapter 3, we present such an analysis that is not limited by these shortcomings, and we complement the contributions of all the prior works by utilizing the perspective of a popular DNS registrar.

2.2.3 Cybercrime and Threat Actor Characterization Studies

Several studies have analyzed different cybercrime operations and their actors in order to understand their incentives. These cybercrimes include pharmaceutical spam [42, 43], spam botnets [44], spam life-cycle [10], targeted attacks [15], click-fraud bots [45], ransomware [46], and RATs [14]. Moreover, prior work [47] has explored cybercrime business relationships and their collaboration. Franklin et al [48] investigated the financial aspect of cybercrime by analyzing transactions on IRC servers. Studying cybercrime operators requires various techniques that include honeypots [49], internet-wide scanning [17, 16], seizing malware infrastructure [50, 44, 51], tracking underground activities [52, 53], analyzing recovered credentials [54], and a combination of diverse data sources [47, 40]. Other works relied on honey tokens to study URL shortening services [55], email typosquatting [56], social media manipulation [57], detect intrusions [58, 59], and vet malicious browser extensions [60]. While these works provide a valuable perspective into cybercrime tactics, However, they have not thoroughly investigated how cybercriminals manage and operate password-stealing campaigns. In chapter 4, we present the first large-scale characterization of password-stealing operations, emphasizing the temporal dynamics of the operator's interactions with their botnet management panels.

2.2.4 APT and Sophisticated Attack Detection and Characterization

In network-based detection systems, network traffic data and domain lifecycle analysis have also been used as the means of APT detection. Alageel et. al., [20] proposed Hawk-Eye, an APT command and control domain detection system that utilizes PCAP data. Oprea et al [61] propose a framework for early-stage APT detection, by modeling the network communications of the internal hosts of an enterprise with outside hosts and utilizing belief propagation. Lamprakis et al [62] suggest a system capable of detecting APT commands and controlling traffic in an unsupervised fashion utilizing host weblogs. Chiba et al proposed a detection system that is based on domain name lifecycle analysis [63]. Other

studies suggested techniques for the detection of lateral movement that are applicable to APTs [64, 65, 66, 67], while a large amount of work has focused on provenance detection and investigation systems [68, 69, 70, 71, 72, 73, 74]. Such studies are orthogonal to our scope as they are aimed at the detection of APT domains rather than the investigation of their network infrastructure over the years.

Several measurement studies have analyzed APT actors and sophisticated attacks over the years. Marczak et al. [16] were among the first to empirically measure and characterize the modus operandi of nation-state actors. Le Blond et al [15] characterized targeted APT attacks against NGO members, finding the actors to utilize recently disclosed vulnerabilities in their malware. Urban et. al [75] analyzed 93 APT reports and found that 80% of the APT actors start their attacks by sending phishing emails. Saha et al., have conducted a user study utilizing 15 APT expert practitioners and have identified that current tools and practices in APT analysis feature significant challenges for threat hunting and attribution [22]. In our study in chapter 5, we reinforce the findings of these prior works regarding the difficulty of the analysis and the sophistication of APT threats by analyzing for the first time comprehensively the infrastructure of 405 actors.

CHAPTER 3

UNDERSTANDING MALWARE INFRASTRUCTURE FROM THE UPPER DNS HIERARCHY

3.1 Motivation

Malware is a pervasive and growing problem [76, 77]. To counter this rising tide, the security community has performed extensive research into understanding malware and has devised techniques for detection, mitigation, and prevention. Unfortunately, malware is extremely diverse—it spans potentially unwanted programs (PUPs), ransomware, and rootkits—making it difficult to generalize results and defenses based on individual malware families.

Ecosystem-wide analysis of malware is necessary to understand broad malware characteristics and to enact appropriate high-level protections and policies. For example, Lever et al. [40] noted heavy malicious usage of popular cloud hosting services which introduced the need for stricter vetting and policing by providers. As another example, Kotzias et al. [41] found that different industries have highly variable infection rates (76% versus 16% for Electrical Equipment compared to Banking), which either suggests targeted attacks by malware operators or indicates that security policies for some industries are more effective than others. Macro-level analysis of malware at large can lead to solutions with far-reaching impact.

Although prior work has explored many aspects of the malware ecosystem, existing research perspectives only have partial visibility into when and where malware infections occur. With the exception of peer-to-peer networks, malware sandboxes cannot observe infected hosts in the wild. The visibility of passive recursive DNS [40] is limited to a handful of collaborating networks. Host-based measurement [41] is often biased or dependent on

pre-installed software and challenging to scale globally. Sinkholes [78, 17] miss infection phases prior to infrastructure takedown. Studies focused on individual malware families (e.g., Mirai [79], ransomware [80]) may have nearly complete visibility, but the lack of malware heterogeneity precludes broader malware ecosystem insight.

This work explores passively collected authoritative DNS (*AuthDNS*) server logs as a new vantage point for characterizing the broader malware ecosystem. The ubiquity of DNS for network communications and its hierarchical nature creates an opportunity to examine malware across four dimensions: malware family diversity, full lifecycle time span, and global visibility into both malware infections and infrastructure. Leveraging data from one of the twenty largest top-level DNS authority zones¹, we study the extent to which *AuthDNS* can replicate previous research findings and also further expand our understanding of the malware ecosystem.

We perform three case studies from the *AuthDNS* perspective. First, looking at malware infrastructure, we find substantial overlap in the networks utilized by different malware families. In the most extreme case, we observe an AS hosting 715 domains associated with 94 distinct malware families. This observation supports prior work [40, 81, 82], which show malware hosting is often interlaced with legitimate infrastructure. We perform a detailed comparison to understand the nuances of each measurement perspective.

Second, we examine the breadth of global malware infections. Previous works studying a wide set of malware have detailed visibility into a specific subset of affected clients (e.g., enterprise networks protected by a specific AV vendor [41]). The *AuthDNS* vantage point provides a slice of global visibility. After looking at all querying clients, we find that targeted malware infections are not apparent for most malware families. Instead, we find that infection rates per country or sector correlate (≥ 0.95 Spearman's ρ) with overall network activity. The vast majority of clients fall under the *Information & Communication* or *Wholesale & Retail Trade* (due to how Amazon's space is classified) industry sectors.

¹Undisclosed due to data sharing agreements.

Third, we examine an under-measured aspect of the malware ecosystem: the full lifecycle of malware communications, from domain registration to blocklisting, and ultimately, expiration. We find that most malicious domains are set up and detected quickly, within five days for 50% of new registrations. Furthermore, we observe a multitude of scanners that emerge after a domain’s detection, as well as a median 38.6% of new client networks first querying malicious domains *after* their expiration. Two explanations for this phenomenon are persistent infections on mobile clients that migrate ASes, or scanners and security professionals querying expired domains [17]. Estimating malware infections from a network perspective after a domain’s expiration should be done with caution.

This study comprises the central pillars of malware epidemiology: the infrastructure that spreads and controls malware, and the location and timing of client infections. To understand how *AuthDNS* supplements existing research, we discuss the advantages and limitations of each vantage point. We then categorize the general types of ecosystem properties (e.g., malware variants, victim targeting, etc.) and provide guidelines for which perspectives will yield meaningful measurements. Ultimately, this work establishes *AuthDNS* as a unique outlook on the malware ecosystem, replicates prior results on malware infrastructure, expands our understanding of malware epidemiology, and introduces a framework to contextualize existing and future research.

3.2 Datasets and Methodology

This section details *AuthDNS* and supporting datasets and describes our methodology.

3.2.1 Datasets

Passive Authoritative DNS (*AuthDNS*). We collaborate with a domain registrar that collects DNS data at the authoritative DNS nameservers used by the top-level zones that it serves. Our DNS data spans 2017-02-09 to 2021-06-30 and includes all DNS packets sent or received by the authority. We extract the IP address of the recursive resolver, the domain

name resolved, the response from the authority, and the client IP subnet for ECS-enabled queries.

Malware DNS (*MAL*). We collect malware domains from a data partner [83] that executes suspicious Windows binaries in an isolated malware sandbox. The malware executions span from January 2018 to April 2021 and amount to 30,302,106 executions. We obtain the communications in PCAP form and extract the DNS traffic.

VirusTotal (*VT*). We query VirusTotal [84] to collect malware family classification labels for malware samples in our *MAL* dataset, and we use AVClass 2 [85] to identify the most relevant label. While *VT* offers results from a plethora of antivirus engines, we only use AV detection results from 17 popular antivirus (AV) vendors that we have found provide stable labels. Additionally, we utilize *VT* to extract historical data for malware samples, malicious domains, and the dates that they were first labeled as malicious.

IP Whois (*IPWHOIS*). We use the *Prefix-to-AS* dataset available from CAIDA [86] to annotate the networks initiating DNS requests. We joined this data with the ASN-to-AS organization delegations provided by the Regional Internet Registries (RIRs). When discussing the *IPWHOIS* dataset, we are referring to the union of these datasets. We utilize this dataset to map IPs to the organizations (and countries) that announce their prefixes.

Industries (*IND*). In order to link an IP address to its industry, we use a commercial IP intelligence dataset. While the dataset is imperfect—a portion of Amazon’s IP space is labeled as *Wholesale and Retail Trade*, which is partially accurate since Amazon’s retail business utilizes its own cloud infrastructure—it represents one of the best labelings available. Open-source solutions such as ASdb [87] provide AS-level granularity that is too coarse for our purposes. *IND* includes organizational property information based on the “International Standard Industrial Classification of All Economic Activities” (ISIC). The Statistics Division of the United Nations (UNSD) [88] provides the mappings of ISIC codes and business categories. We intersect the two datasets to attribute an IP address to a specific business based on the UN standard. We refer to different industries as *sectors* through the

rest of the paper.

3.2.2 Methodology and Validation

This section presents the methodology used to generate and validate our malware domain dataset. Algorithm 1 summarizes our methodology:

Algorithm 1: Labeling and Validation Methodology of Malware Samples and Malicious Domains

Input: MAL (malware DNS), AuthDNS (authoritative DNS), VT (VirusTotal)

Output: Labeled malware samples and validated malicious domains

- 1 **Step 1: Generating Malware Domain Dataset;**
 - 2 Extract overlapping DNS queries between MAL and AuthDNS;
 - 3 Filter out top-ranked domains using the Tranco list;
 - 4 Extract effective 2LDs (e2LDs) $\rightarrow \mathcal{D}_{mal}$;
 - 5 Submit domains to VT to obtain the number of vendors labeling them as malicious;
 - 6 **Step 2: Expand Sample Set;**
 - 7 **foreach** $d \in \mathcal{D}_{mal}$ **do**
 - 8 | Query VT for additional malware samples;
 - 9 Merge with original MAL samples $\rightarrow \mathcal{S}_{total}$;
 - 10 **Step 3: Malware Sample Labeling;**
 - 11 **foreach** $s \in \mathcal{S}_{total}$ **do**
 - 12 | Query VT for AV labels;
 - 13 | Apply AVClass2 to normalize labels;
 - 14 | **if** *label is generic or AVs disagree* **then**
 - 15 | | Label s as SINGLETON;
 - 16 | **else**
 - 17 | | Assign dominant family label to s ;
 - 18 **Step 4: Malicious Domain Validation;**
 - 19 **foreach** $d \in \mathcal{D}_{mal}$ **do**
 - 20 | Check historical URL reputation in VT;
 - 21 | Confirm if the majority of queries are from MAL samples only;
 - 22 **Step 5: IP Labeling and Enrichment;**
 - 23 **foreach** *client/server IP* **do**
 - 24 | Use IPWHOIS to map to ASN and organization;
 - 25 | Use IND to map IP to the industry sector;
-

Generating Malware Dataset. To obtain a set of malware-related domains, we first find the overlap between our *MAL* and *AuthDNS* datasets. Malware samples may query

benign domains to check for network connectivity. Similar to Lever et al. [40], we filter out top-ranked domains in the Tranco list [89]. This filtering yielded 12,212 effective second-level domains (e2LDs), which capture the registrable portion of a domain name. For example, in the fully qualified domain name *www[.]example[.]co[.]uk* the e2LD is *example[.]co[.]uk*, while the second level domain name is *co[.]uk*.

The 12,212 malicious e2LDs are associated with 174,112 malware samples from *MAL*. We submit them to *VT* for scanning and find that 98.96% of the samples are known to *VT*, and 99.97% of known samples are marked as malicious by five or more AV vendors. Finally, we expand our dataset by querying *VT* for all malicious samples communicating with the 12,212 malicious domains. This reflection yields an additional 70,898 samples, for a total of 245,010 samples.

Malware Sample Labeling. Different AV vendors offer divergent labels for a malware sample [90]. We use AVClass2 [85] and a malware encyclopedia [91] to resolve these aliases (e.g., *bladabindi* to *njrat*) when possible. We keep the top malware family label by AV vendor agreement and disregard generic labels or labels where the AV vendors cannot agree (SINGLETON). Following this methodology, we discard 81,750 samples (33.63%) assigned the label SINGLETON. The 161,322 (66.37%) successfully labeled samples represent 202 distinct malware families. No malware families appear to have an outsized representation in our datasets, and we summarize the top 15 malware families by the number of domains in Table 3.1.

Figure 3.1 shows the cumulative distribution of the number of malware samples and hosting servers per domain in our dataset. Most domains are associated with only a handful of malware samples, with 57% of the domains related to less than three samples. A similar trend holds for the number of servers resolved by a given domain. These distributions are consistent with those in prior large-scale malware measurement studies [40].

Malicious Domain Validation. To validate the maliciousness of the 12,212 e2LDs,

Table 3.1: Top 15 malware families based on the number of malicious domains in our dataset.

Family	Malware		Server		Client		
	Domains	Samples	IPs	CC	Count	CC	Sectors
darkkomet	3,578	16,441	175K	140	2,187K	232	20
njrat	1,924	10,596	195K	129	1,970K	229	21
cybergate	1,181	2,546	38K	100	931K	219	19
xtrat	946	2,801	62K	89	1,108K	222	19
bifrose	700	1,432	11K	62	497K	211	18
razy	667	1,139	107K	110	1,508K	225	18
remcos	563	39,279	61K	103	1,028K	221	18
nanocore	501	2,112	72K	116	1,446K	227	19
ponystealer	450	4,891	49K	93	106K	222	17
gamarue	410	761	53K	97	1,523K	225	19
poison	355	1,018	18K	75	692K	212	18
vobfus	282	3,843	36K	89	936K	219	19
nymeria	279	966	39K	101	838K	215	18
zbot	229	24,736	9K	61	945K	220	20
netwire	228	634	34K	82	859K	223	18

we query *VT* and find that 76.7% of the malicious e2LDs have at least one historical URL labeled as malicious. 87.5% of the filtered malware samples only queried domains in our *AuthDNS* and no additional domains. This combination of factors gives us high confidence in our malicious domain dataset.

Figure 3.2 shows the aggregate daily query volume of malicious domains, as seen in *AuthDNS*. Our vantage point provides a stable view throughout the four years of our study, except for three dips related to collection issues. On average 17.9% of daily requests are ECS-enabled, allowing us to learn the clients’ subnets in addition to the IP address of the recursive. Similar to Kountouras et al. [92], we define a client as the client subnet when ECS is enabled and the recursive’s IP address when ECS is not enabled. We use this client definition for our experiments in section 3.4 and section 3.5. Finally, we apply *IPWHOIS* and *IND* to servers and clients in order to identify relevant ASNs, organizations, countries, and industry sectors.

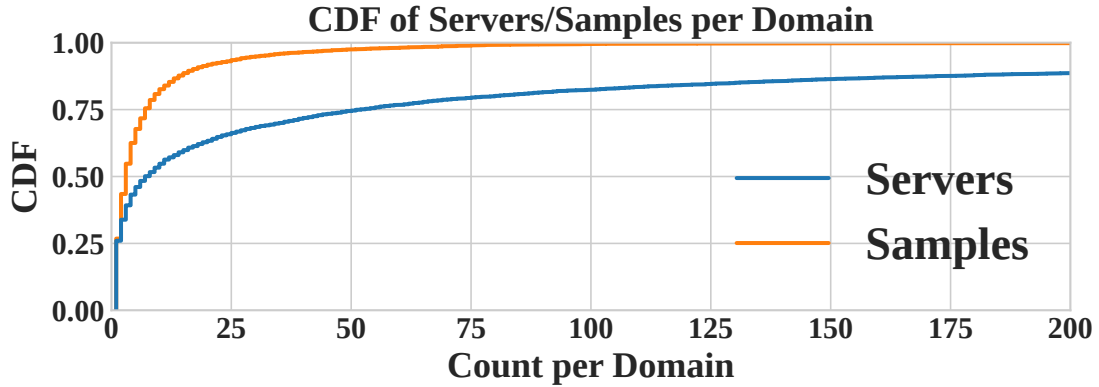


Figure 3.1: Distribution of the number of malware samples and servers associated with each domain in *AuthDNS*.




Figure 3.2: Daily Volume *AuthDNS* for malware domains. ECS-enabled requests (orange) average 17.9% of daily requests.

3.3 Hosting Infrastructure

The hosting infrastructure used by cybercriminals is an essential aspect of malware communication. Understanding how malicious actors distribute and coordinate malware enables the security community to take more effective remediation steps and can focus resources on areas of frequent abuse. To study this infrastructure, we consider a set of 6,400 domains representing the intersection of domains with malware family labels and *IPWHOIS* labels for the IP addresses resolved by those domains. In aggregate, this set of domains points to 399,830 different IP addresses in 151 countries.

First, we consider where malware is hosted. Figure 3.3a shows a map of all the countries we can associate with infrastructure resolved by malicious domains, with lighter col-

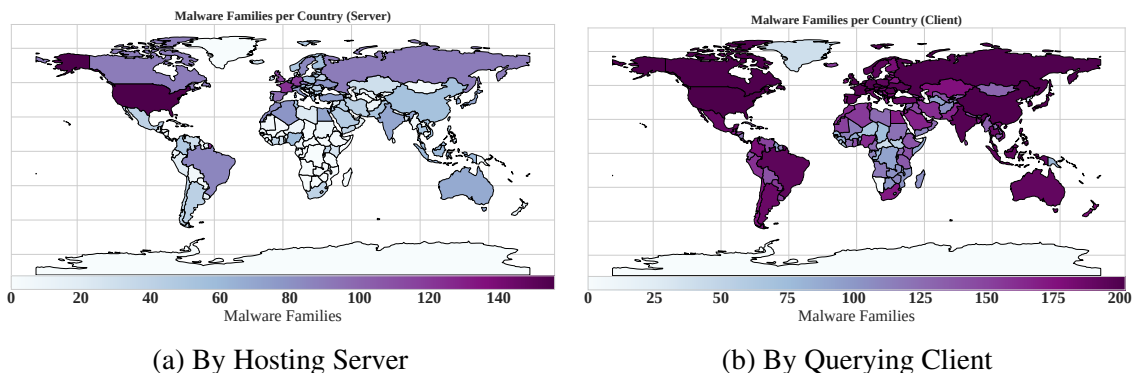


Figure 3.3: Geographic distribution of the IP addresses of hosting and clients for malware-related domains. Darker colors indicate more malware families are associated with the country through the hosting server (left) or querying client (right).



Figure 3.4: Distribution of the number of malware families per IP, Prefix, ASN, and country that have resolved malicious domains in *AuthDNS*. Most network infrastructure and targeted networks are not strongly correlated with a single malware family.

ors indicating fewer malware families. Countries home to large hosting providers—like the United States, France, and Germany—also host large numbers of malware families. Tajalizadehkhoob [81] and Mezzour [82] both found that the distribution of C2 infrastructure on legitimate hosting platforms was strongly correlated with the size of the hosting platform and weakly correlated with their security policies. Our work reiterates that hosting infrastructure may enable malware communication to hide in plain sight. For example, Lever [40] showed that PUP software is often long-lived on legitimate, commercial hosting platforms and found a growing trend of malware samples taking advantage of such hosting.

Zooming in, we examine how infrastructure is reused across different malware families. Figure 3.4a shows the distribution in the number of malware families hosted per country (corresponding to Figure 3.3a), ASN, network (BGP Prefix), and IP address. We find

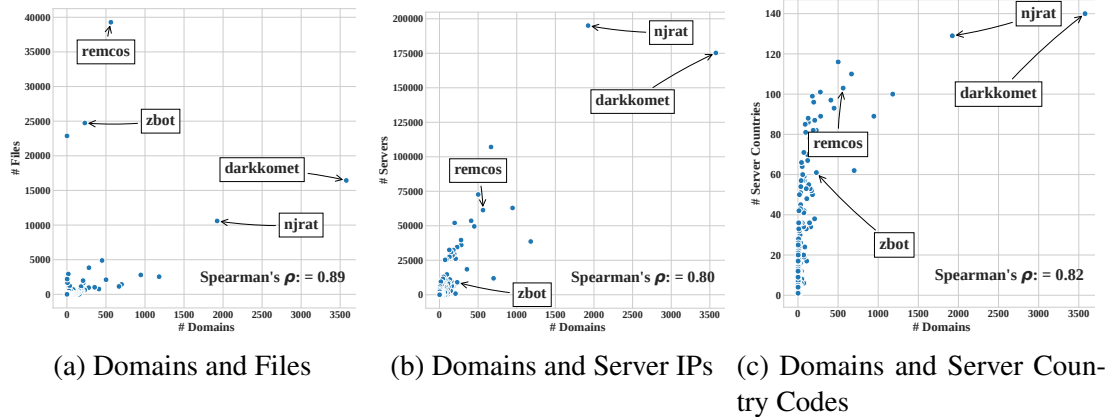


Figure 3.5: Correlation between different measures of hosting infrastructure. Higher numbers of samples per malware family correlate with more domains. Higher domain utilization correlates with more server IPs and more hosting countries.

that only 102,728 (25.7%) of malware-hosting IP addresses were associated with a single malware family. Conversely, 26,226 (6.6%) of IP addresses resolved by malware domains could be tied to ten or more families. In one case, we found that IPs belonging to AS29075 (IELO IELO-LIAZO SERVICES SAS) were pointed to by 715 domains corresponding to 94 malware families. We believe this to be the result of many malicious actors taking advantage of a proxy operated within this French ISP, demonstrating widespread reuse.

Finally, malware families often spread their hosting across multiple countries. We found that only 24 malware families have their hosting contained to a single country. To help explain the intra-family diversity of hosting, we looked at the correlation between the number of domains in *AuthDNS* contacted by each malware family and the number of samples, hosting server IPs, and hosting server countries. Figure 3.5a shows a strong correlation between the number of domains used by a given malware family and the number of unique samples in our dataset. Several outliers, such as *zbot*, for which we observed very few domains but a large number of files, reduced the Pearson correlation. However, the Spearman correlation, which is more tolerant of outliers, still showed a strong correlation. Figure 3.5b goes on to show a strong correlation between the number of domains contacted by each malware family and the number of hosting servers observed. Further,


Figure 3.5c shows that as the number of domains and hosting IPs increases, so does host country diversity. As the malware family progresses, using more domains or samples, it naturally expands. This breadth of hosting infrastructure contributes significantly to the security community’s challenge of attribution and takedowns.

Takeaway-1: *The view of malware-related domain hosting provided by AuthDNS largely agrees with prior work, which relies upon network data collected at different points in the DNS hierarchy. Infrastructure is reused across different malware families and is often intertwined with legitimate hosting services. Within a malware family, it is common to see many host networks and IPs deployed, often crossing geopolitical boundaries. This agreement between datasets suggests interchangeability; however, a key factor makes AuthDNS data superior when available. Non-global vantage points such as RecursiveDNS will only yield snapshots of the hosting infrastructure once customers using that recursive begin querying for a given domain. This may limit visibility during the early stages of a domain’s life, particularly before widespread infection by the corresponding malware occurs. AuthDNS does not suffer from this limitation.*

3.4 Malware Clients

AuthDNS provides a unique perspective on potential victims who query for malicious domains. In section 3.3, many of our findings concerning malicious domain hosting agreed with prior work and could be drawn from other vantage points. The same interchangeability is not valid for studying those infected by malware families. The limitations of previous techniques become apparent when we observe victims through the global perspective of *AuthDNS*.

As discussed in Figure 25, authoritative nameservers receive DNS requests from recursives rather than individual hosts. This makes the tradeoff of gaining visibility into all querying recursives, but gives up visibility into individual endpoints, which recursive DNS provides for a subset of the population. Thus, for non-ECS queries, we consider the IP



(a) Client Count and Impacted Countries (b) Client Count and Impacted Sectors (c) Impacted Countries and Impacted Sectors

Figure 3.6: Spearman correlation between different measures of potential victims. Higher numbers of clients querying for malware family-related domains correlate strongly with a more diverse set of impacted countries and economic sectors.

address of the recursive to be the client, while for ECS-enabled requests, we use the ECS netmask.

Our aim with *AuthDNS* is to study large infected populations as epidemiologists rather than infected individuals as doctors. Figure 3.3b shows the number of malware families affecting each country, with darker colors representing more malware families. A significant portion of malware families plagues nearly every country. These globally expansive infections contrast with Figure 3.3a, which showed higher concentration levels of malware family hosting in particular countries. Figure 3.4b zooms in to indicate the number of malware families that are contacted by each network or client.

Viewed from the opposite direction, Figure 3.6a shows the number of countries each malware family affected with respect to the total number of clients observed. We found that only one family, `fosniw`, had queries to related domains originating from fewer than ten countries, while 144 (71.3%) of malware families were found to be queried from 100 or more countries. This agrees with Mezzour et al. [82], which also witnessed near-universal affliction by malware in developed countries. Additionally, they found that infection rates correlated strongly with the IT resources of that country. As with the location of hosting infrastructure for malware-related domains, the spread of malware across geopolitical

boundaries complicates the security communities’ task of identifying the targeting of victims for most malware families. We see that malware families do not generally tend to target specific networks, but rather, many networks appear to be infected by multiple different malware families. Furthermore, malware family infections are not commonly confined by geography as seen from the perspective of *AuthDNS*.

3.4.1 Industry Sectors

Table 3.2: Unique clients querying malicious domains and the number of malware families in each industry (ISIC section).

ISIC Section	Clients	Malware Families
Information & Communication	3,108,546	202
Wholesale & Retail Trade	567,729	202
Education	29,741	201
Professional, Scientific & Technical Activities	11,576	196
Manufacturing	4,837	192
Government, Defence	4,697	178
Financial & Insurance Activities	3,670	183
Human Health & Social Work Activities	3,785	172
Accommodation & Food Service Activities	2,785	148
Transportation and Storage	624	155
Arts, Entertainment & Recreation	421	140
Electricity, Gas, Steam & A/C Supply	333	127
Administrative and Support Service Activities	199	141
Extraterritorial Organizations and Bodies	164	120
Other Service Activities	149	149
Real Estate Activities	96	86
Construction	74	38
Mining and Quarrying	17	23
Agriculture, Forestry and Fishing	5	18
Water Supply, Sewerage.	5	8

Another way of grouping clients is by the type of network they query from. Kotzias et al. found evidence that different industries are affected by different amounts of malware samples [41]. In that study, the authors relied on file reputation logs collected from endpoint protection software to study malware affecting customers of a large cybersecurity company.

Table 3.3: Client distribution across sectors for seven-day samples starting 2017-03-01, 2018-03-01, 2019-03-01, and 2020-03-01. *All* represents the complete *AuthDNS* dataset while *Mal* represents only malicious domains.

ISIC Section	2017		2018		2019		2020	
	All	Mal	All	Mal	All	Mal	All	Mal
Information & Communication	0.85M	68K	1.4M	121K	1.3M	136K	1.3M	117K
Wholesale & Retail Trade	5.2K	1.0K	14K	3.4K	16K	6.8K	29K	10K
Education	19K	1.5K	27K	2K	25K	2.6K	29K	1.9K
Professional, Scientific & Technical	4.6K	402	8.8K	1.1K	6.5K	1.1K	7.7K	1.1K
Manufacturing	2,109	181	3.2K	387	2.9K	383	2.7K	246
Government, Defence	3.1K	215	5.2K	347	4.3K	384	4.5K	310
Human Health & Social Work	2.1K	129	3.5K	213	3.3K	228	3.5K	163
Financial & Insurance	2.7K	157	4.1K	267	3.6K	271	3.7K	214
Accommodation & Food Service	1.5K	69	2.6K	103	2.2k	131	2.3K	75
Transportation & storage	250	34	379	58	317	59	309	25
Arts, Entertainment & Recreation	342	17	554	32	510	29	493	16
Electricity, Gas, Steam & A/C Supply	224	19	316	37	269	39	299	26
8 remaining sections	494	29	771	52	757	56	773	52
Correlation (Spearman):	0.98	0.99	0.98	0.99	0.98	0.99	0.98	0.98

They found that specific industries were affected by a disproportionate number of malware, suggesting targeting by malware families and disparity in security posture across industries. However, their vantage point was limited to customers of the cybersecurity company, which they acknowledge introduces bias. We seek to augment this work by studying how malware families infect industries from a global vantage point.

For this analysis, we use the industries (IND) discussed in subsection 3.2.1. This mapping enables us to group clients by ISIC code. ISIC provides a hierarchical classification of industries that breaks down into 21 sections, 88 divisions, 238 groups, and 419 classes. We use the ISIC section synonymously with industry sector in the remainder of this work.

We observed around 39.7B requests for malicious domains in our dataset, and were able to assign an industry label for approximately 28.08B (70.7%) requests. We only consider requests from clients to whom we can assign an industry label.

Table 3.2 shows the number of clients we observed in each sector as well as the number of offending malware families. For clarity, we rename some of the ISIC code classification labels. We can immediately see that each industry contains clients querying do-

mains associated with numerous different malware families. In fact, 15 of the ISIC sections appear to be impacted by over half of the malware families in our dataset. We note that the `Information & Communication` ISIC section, as well as `Wholesale & Retail Trade` section, contains requests from all the malware families and represents more clients than any of the remaining industries by several orders of magnitude. From the more granular ISIC divisions, we see that most queries from the `Information & Communications` industry can be attributed to wired and wireless communications due to the classification `Internet Service Providers (ISP)` and the `Residential & Business Hosting Infrastructure`. A portion of the IP address space controlled by Amazon is labeled as `Wholesale & Retail Trade`, contributing to an overestimation of the effects on this population. This also explains why these two sections contain several orders of magnitude more requests compared to other sections.

To capture the representation of clients querying for malicious domains compared to all clients in *AuthDNS*, we sampled a seven-day window for each year in our study. Table 3.3 shows the number of clients from top ISIC sections querying for any (malicious or benign) domain during this window and the subset querying for malicious domains. The final row shows the correlation (Spearman) between the sampled datasets and the dataset of malicious domains that spans the complete four-year study.

The next largest ISIC section by number of unique clients is `Education`, with roughly half of the requests in this section coming from institutes of higher education such as colleges and universities. These institutional networks typically have a wide variety of users, including students, staff, faculty, and visitors. Many such networks may not have direct control over the devices on their network. A heterogeneous base of infected devices and research-related activities provides a sensible explanation for why `Education`, and higher education in particular, accounts for so many malware-related queries. While some of the remaining ISIC sections seem like prime candidates for targeted behavior, we note that even sections associated with the government, defense, finance, and infrastructure

appear to be impacted by many different malware families.

We find that malware families generally impact multiple ISIC sections, with 72.7% of the malware families found in more than 10 sectors. Our aggregate analysis with global visibility cannot draw the same conclusions as Kotzias et al. [41], which found that 1,911 (37%) of the malware families in their study were only seen in one enterprise. Instead, Figure 3.6b and Figure 3.6c respectively show that the number of industries a malware family impacts is correlated with the overall number of clients impacted by that malware and the geographic diversity of those clients. As malware families grow, so does the diversity of their victims.

Takeaway-2: *In aggregate, we do not see malware families solely affecting individual industries. 72.7% of the malware families in our data are found to affect more than 10 distinct industry sections. While datasets derived from RecursiveDNS or host-based security products offer a view into the networking behavior of individual end-users, they can introduce biases by considering customers in a particular geographic region or those already taking steps to mitigate their online risk. Our study of clients affected by a range of malware families highlights how AuthDNS’s global vantage point can reduce these biases and lead us to draw divergent conclusions when studying malware infections from an epidemiological standpoint. Still, AuthDNS leaves ample room for studies such as Kotzias et al. [41] that provide greater visibility into individual infected hosts for a subset of the population once these potential biases are placed in the context of a global view. While AuthDNS does not provide visibility into end-users, it does offer a complete view of recursive querying domains under that authority. ECS-enabled requests further narrow this gap when looking at affected clients through the lens of AuthDNS.*

3.5 Malware Lifecycle

Utilizing the unique vantage point of *AuthDNS*, we perform a temporal analysis in order to understand the lifecycle of malicious domains. We complement the client visibility of

Table 3.4: Request characterization for three temporal windows. While most clients and sectors are observed between malicious domain detection and expiration/takedown, many client ASNs and countries (ASCCs) first connect after expiration/takedown.

Domains	Registration to Detection						Detection to Expiration/Takedown						Post Expiration/Takedown								
	ASNs	(%)	ASCCs	(%)	Sectors	(%)	Days	ASNs	(%)	ASCCs	(%)	Sectors	(%)	Days	ASNs	(%)	ASCCs	(%)	Sectors	(%)	Days
10%	0	(0.00)	0	(0.00)	0	(0.00)	0	19	(11.6)	5	(11.5)	0	(00.0)	1	10	(6.80)	1	(1.02)	0	(0.00)	238
25%	0	(0.00)	0	(0.00)	0	(0.00)	1	59	(36.2)	15	(42.9)	2	(37.5)	23	36	(21.9)	3	(9.09)	0	(0.00)	526
50%	6	(2.63)	3	(7.69)	1	(20.0)	4	101	(54.0)	26	(62.5)	4	(62.5)	30	76	(38.6)	10	(22.7)	1	(16.7)	963
75%	22	(10.0)	9	(21.8)	2	(37.5)	19	164	(70.4)	37	(79.0)	5	(80.0)	100	132	(55.3)	18	(40.0)	2	(30.0)	1,180
90%	54	(23.4)	18	(40.0)	3	(57.1)	79	369	(86.5)	54	(90.4)	7	(100)	419	229	(71.3)	28	(58.6)	3	(50.0)	1,256
max	2,243	(96.6)	136	(100)	14	(100)	1,154	11,650	(100)	187	(100)	15	(100)	1,661	4,644	(100)	95	(100)	10	(100)	1,558

previous studies that observed malicious domains after expiration [17] by considering all clients querying for a malicious domain name during three phases: registration to detection, detection to expiration, and post-expiration. We determine the date of detection as the earliest of the following dates: a malicious URL of the domain is detected by more than one vendor in VT, a malicious hash communicating with that domain is detected in VT, or a malicious hash communicating with that domain is seen in our malware DNS dataset. In order to fully observe the domain lifecycle, we only consider domains that were registered after the first day of visibility we have in *AuthDNS*. Further, we restrict our analysis to domains that have been registered only once in our *AuthDNS* dataset so that we do not observe noise from previous or subsequent registrations as domain names get repurposed. This filtering leaves us with 2, 308 domain names, 18.9% of the total domains in our dataset.

Table 3.5: Most popular ASNs first observed in each temporal window. Scanners and AV vendors appear mostly during and after the detection of a malicious domain, while hosting networks are most prevalent during the setup of the domain.

Registration to Detection		Detection to Expiration/Takedown		Post Expiration/Takedown	
ASNAME	Domains	ASNAME	Domains	ASNAME	Domains
AMAZON-AES	772	WINTEK-CORP	1,745	CNNIC-ALIBABA-US-NET-AP	Alibaba (US) Technology Co., Ltd. 1,387
CORBINA-AS PJSC "Vimpelcom"	756	GEORGIA-TECH	1,738	CNIX-AP	China Networks Inter-Exchange 1,363
GOOGLE	645	OVH OVH SAS	1,662	CHINATELECOM-TIANJIN	Tianjij.300000 1,226
LEVEL3	611	MFENET	1,649	InterConnect ML	Consultancy 1,114
AMAZON-02	602	PAN0001	1,641	FSOL-AS F-Solutions Oy	1,084

Table 3.4 summarizes the networks for the lowest 10%, lower quartile, median, upper quartile, 90%, and max of querying clients during all phases of the domain lifecycle. The registration to detection window is relatively short, lasting 19 days or less for 75% of domains. Additionally, at the median, only six networks (ASNs) and three countries queried for domains while they were in this initial phase. By comparison, the second temporal

window, detection to domain expiration/takedown is significantly longer, with at least 23 days representing the lower quartile. At the median, 54% of ASNs that will ultimately query for a domain do so for the first time during this window. The same observation holds for the countries and industry sectors of these clients. Finally, queries continue during the post-expiration/takedown period, which continues until the end of our four-year *AuthDNS* dataset for domains that are not re-registered. Interestingly, in this period, the median domain observes more than 76 unique ASNs and ten countries querying it for the first time. This represents a long tail of unique clients first seen only after a domain has expired or been taken down.

In order to understand the most popular networks in each lifecycle phase, we look at the top querying ASNs across domains. Table 3.5 shows the top five unique ASNs as seen by the number of first occurrences in each temporal window. During the registration to detection window, we first observe large hosting networks (Amazon), large recursives (Google), and large telecommunication companies (Vimpelcom and Level3). This window is related to the setup and testing of the domains by the actors and the first potential victim connections, resulting in queries from large recursives and telcos. After the domain's detection, the most common ASNs to be first observed are large scanners (GEORGIA-TECH [32]), AV companies (MFENET - McAfee and PAN0001 - PALO ALTO NETWORKS), and other large hosting networks (WINTEK-CORP and OVH), which can contain other scanners. In this window, AVs, sandboxes, and scanners query malicious domains and map their IP address space. Lastly, in the final window, post-expiration/takedown, we observe large Chinese telcos and business networks from other countries. These post-expiration queries could be due to network mobility of infected clients, new infections, or scanning.

Takeaway-3: *The view provided by AuthDNS shows that researchers need to consider a domain's lifecycle to measure infected populations accurately. Most domains in our dataset were detected as malicious soon after registration, with the median time being four days. After detection, domains will receive increased interest from scanners and AV*

vendors, which can artificially inflate infected population counts if proper filtering is not applied. Notably, there commonly exists a long tail of new client queries after a domains' expiration or takedown. Existing infections on mobile clients generate queries from new networks and may persist into this final phase. However, prior studies further suggest that scanning activity late in the lifecycle of a domain may constitute a significant portion of queries [17]. Researchers and practitioners using network data, AuthDNS or otherwise, to estimate client infections, risk obscuring malware behavior when they do not distinguish between phases of the domain lifecycle. As a community, there is room for further improvement in identifying scanners and distinguishing the lifecycle phases for domains with multiple registrations. Addressing these challenges will allow researchers to better understand and help infected populations.

3.6 Vantage Point Comparison

Thus far, we have shown *AuthDNS*'s ability to reproduce previous observations of malware infrastructure (section 3.3), add a novel perspective on the distribution of malware infections (section 3.4), and to introduce a full temporal view of the malware domain lifecycle (section 3.5). In this section, we synthesize these findings and contextualize them in the broader landscape of malware ecosystem and epidemiology research. We first enumerate related work and map the relationships between different perspectives. We then compare the perspectives along four different qualitative characteristics and highlight the appropriate role of each perspective, gaps in existing malware visibility, and avenues for future research.

3.6.1 Measurement Planes

Broadly speaking, malware utilizes three distinct network planes² (Figure 3.7), which we define as a grouping of network components based on their location and functionality

²Unrelated to control/data planes from software-defined networking.

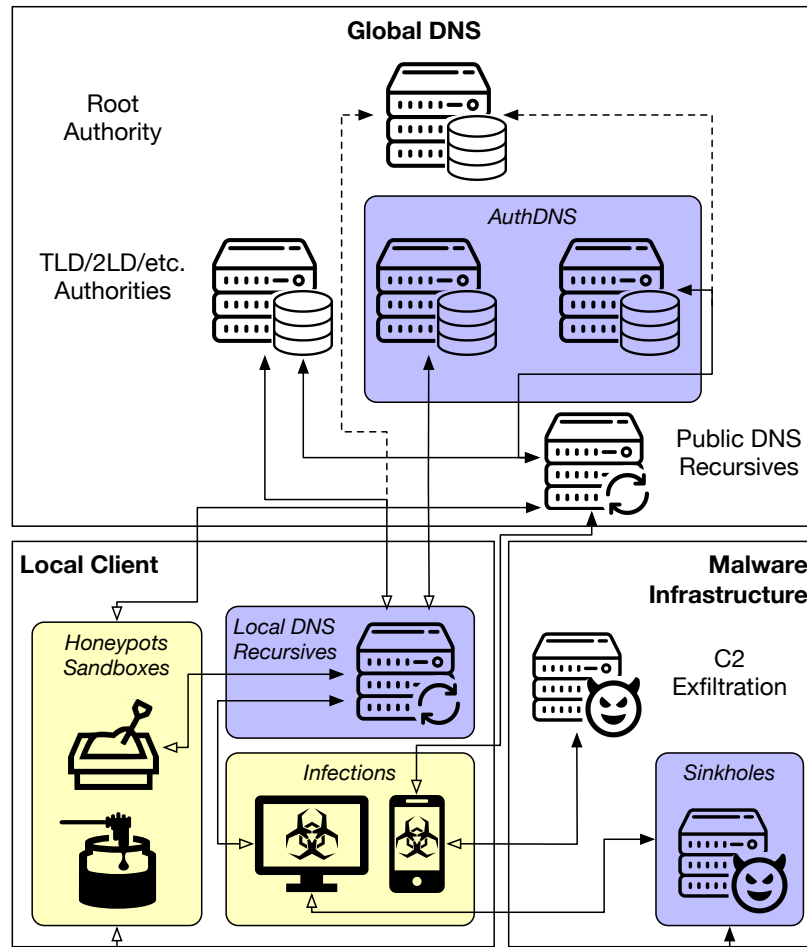


Figure 3.7: Malware measurement can occur 1) in the global DNS plane, 2) at or near the local client, or 3) at the malware infrastructure. Each location has specific sub-components that interact via request (filled arrow) and response (empty arrow) protocols (e.g., DNS, C2 protocols). Existing research has studied many of the depicted components with host-based (yellow) or network-based (blue) techniques.

within the network.

Global DNS DNS is the bootstrapping protocol used by most network communication to map domain names to IP locations; malware uses DNS extensively, as evidenced by countless domain blocklists and DGA malware. Global DNS refers to upper-hierarchy authoritative DNS servers (e.g., root, TLD, 2LD) and large public recursives (e.g., Google, CloudFlare), which share a global perspective on domain lookups. Global DNS servers can receive DNS requests from a universal set of clients and have worldwide visibility into domain usage. Prior work [37, 39, 38] utilizing global DNS authorities within the realm of malware has focused on detection, spam measurements, and one work [93] performed a case study on stalkerware based on probing of large public recursives.

Local Client The local client plane consists of malware-infected clients and the local networks in which they reside. In contrast to global DNS, which only has network-based techniques, the local client plane contains both host-based and network-based approaches. Host-based approaches include any measurements that directly observe partial or full execution of malware: interactive honeypots, malware sandboxes, and in-the-wild infections. Existing malware research has skewed heavily toward host-based analysis of the client plane. As a brief example amongst a profusion of works, sandboxes and honeypots have been used to study general Windows malware [94, 95], malware downloaders [96], Android applications [97], malware protocol reverse engineering [98, 99], exploit sites [100, 101], C2 hosting [81, 102, 103], and DDoS [104]. Two works have utilized host-based local client techniques to shed light on the broader malware ecosystem. Kotzias et al. [41] applied host-based infection measurement across 28K enterprises in 67 industries to determine enterprise malware trends. Messour et al. [82] conducted an empirical analysis of Symantec’s telemetry data to observe the distribution of different malware types across countries.

The primary network-based approach within the local client plane is the collection of

DNS data from local recursives that handle a network’s DNS traffic and are often set by default via Dynamic Host Configuration Protocol (DHCP). Alrawi et al. [11] used recursive passive DNS data to estimate infections by IoT malware, while Lever et al. [40] focused on more general malware, including PUPs.

Malware Infrastructure Malware relies on infrastructure most commonly for command and control (C2), but can also use separate infrastructure for hosting, data exfiltration, or other functions. Measurement of malware infrastructure IP addresses can occur from global DNS and local client planes, but to collect communication data between infections and malware infrastructure, researchers have developed sinkholes. Sinkholes allow a researcher to operate or imitate malware infrastructure and collect richer data about connecting clients. Several works have utilized sinkholes to study specific phenomena (e.g., remote-access trojans (RAT) [17], botnets [50]); one prior work by Alowaisheq et al. [78] studied sinkhole domain behavior across all types of malware, but did not operate any sinkholes, since they require malware-specific configuration.

3.6.2 Comparison

Infection Visibility Infection visibility is the capability of a vantage point to assess all infections of a threat globally and temporally. This work shows *AuthDNS* datasets yield high global infection visibility as they provide access to all DNS requests made to a malicious domain, across all locations and time. Datasets that are not based on network infrastructure are limited in infection visibility as they can only observe a subset of clients based on data source (e.g., AV vendors, ISP clients, recursive clients, email clients). Domain sinkholes provide global location visibility, but partial temporal visibility, as they are limited to the post-expiration period of a domain. Infrastructure takeover can provide global and temporal infection visibility guarantees; however, this is difficult to execute and scale.

Infection Precision Infection precision is the capability of the dataset to accurately estimate the validity and type of infections. Passive DNS datasets contain noisy infection data that is muddled with traffic from scanners, malware sandboxes, or security professionals. Thus, users of passive DNS datasets should filter clients based on their behavior and network origin when estimating infections. Additionally, many different malware samples and families can be hosted on the same domain name, and the type of infection per client cannot be guaranteed. Client-side antivirus datasets provide higher infection precision for the type of client and the existence of a specific malware sample; infrastructure takeover datasets can provide the highest precision by looking at the collected infected system data. Lastly, domain sinkholing initially provides partial visibility, since a domain will continue to receive queries after its detection period, but sinkholing data can be enhanced for a better infection estimation as shown by Rezaeirand et al. [17]. Email datasets provide the lowest precision as they observe the targeting aspect of an attack rather than the infection.

Client Granularity Client granularity is the capability of the dataset to trace the infections down to single clients or users. Authoritative pDNS datasets are limited in this regard since clients are obscured by recursive DNS servers and caching. However, as shown in this study, researchers can use the ECS field of an ECS-enabled request to obtain higher precision client granularity. Recursive pDNS datasets yield even higher client granularity as they can observe all the clients under the recursive making requests for a malicious domain. Client-side AV datasets, infrastructure takedown datasets, ISP network logs, and email datasets provide high guarantees of client granularity as they can observe a unique client or user.

Malicious Infrastructure Visibility Malicious infrastructure visibility is the capability of the dataset to observe what infrastructure the malware actors have used to perform their campaign. *AuthDNS* is, by definition, the authoritative source of the mapping of domains to IPs for a malicious domain. DNS rewriting, for the profit of the recursive operator [105],

or for the protection of customers, may limit the hosting infrastructure visibility provided by a RecursiveDNS dataset. Infrastructure takedowns provide the highest guarantees as they provide direct access to infrastructure; however, it is not scalable. Infrastructure visibility via recursive DNS and ISP network logs depends on the volume and consistency of communications by infected clients within the measured networks. The remaining datasets cannot provide any insights regarding the infrastructure used by the malicious actors.

Takeaway-4: *AuthDNS has several advantages and disadvantages when compared to the vantage points used in previously published research. We find global, temporal, and infrastructure visibility to be the biggest advantages of our dataset. Thus, we position our measurements along these advantages and we study each aspect in depth in section 3.3, section 3.4, and section 3.5 and report our most insightful results. AuthDNS has limited client granularity and limited infection precision. Future work can be aimed at addressing this issue.*

3.7 Summary

Understanding malware lifecycles is vital in the fight against Internet threats. This work presents a longitudinal study analyzing the network communication of 202 different malware families from the perspective of a popular authoritative DNS server. We observed billions of resolutions over four years at our authoritative collection point, enabling temporally complete and global visibility into malicious domain usage. *AuthDNS* simultaneously solidifies prior findings while also shedding new light on the epidemiology of malware. First, different malware families often reuse the same network infrastructure, so threat intelligence needs to label malicious infrastructure cautiously. Second, malware families, when analyzed in aggregate from an *AuthDNS* vantage point, do not appear to target specific networks or industries. Instead, they spread to many different industries with high regularity over time. Third, our temporal analysis shows that newly registered malicious

domains are set up and detected quickly. Due to network noise from scanners and AV vendors, both the temporal and organizational properties of network clients should be considered when estimating malware infections from a network perspective. Finally, we introduce a brief taxonomy of malware measurement perspectives and discuss the advantages and disadvantages across four primary measurement goals. By broadening our understanding of global malware infections, this work serves as a stepping stone to making malware characterization more accurate and, ultimately, to making mitigation more effective.

CHAPTER 4

UNDERSTANDING THE INTERACTION OF MALICIOUS ACTORS WITH THEIR INFRASTRUCTURE THROUGH AN EMPIRICAL ANALYSIS OF PASSWORD STEALERS

4.1 Motivation

In chapter 3, by utilizing the *AuthDNS* dataset and its end-to-end temporal visibility in the lifecycle of malware-hosting domain names, we have explored the network interactions (i.e., DNS requests) all the stakeholders that communicate with these domain names, from victims to security vendors and DNS scanners and identified key takeaways towards the quantitative aspects of the network interactions throughout the domain name lifecycles and the qualitative aspects of the visibility that the *AuthDNS* dataset provides. Despite that, a major perspective that was not explored but has a significant impact on understanding the lifecycle of malicious network infrastructure is that of the interactions of the malicious threat actors with their own infrastructure.

Throughout the lifecycle of a malicious domain name, many different stakeholders interact with it in the form of DNS requests, from DNS scanners like the ActiveDNS project [32] to security vendors and, more importantly, victims. While the distribution of new and unique networks first querying the domain names is skewed towards the end of the lifecycle, there are many requests distributed through the lifecycle that are hard to characterize for the stakeholder responsible for them. This is particularly difficult when seeking to characterize the interactions of the malicious threat actors that utilize these malicious domain names in order to understand how long they operate, which networks they operate from, and what impact detection has on their operations. The utilization of public and large DNS recursives, VPNs, and proxies makes the interactions of the threat actors blend

in with DNS scanners, security vendors, and victim traffic, making it almost impossible to characterize them utilizing this vantage point alone.

In this chapter, we characterize the interactions of malicious threat actors with the infrastructure they utilize for managing and performing their cyber attacks, with the goal of understanding: 1) the lifecycle of their interactions, 2) the network infrastructure they utilize, and 3) the impact that malicious detection has on their operations. To do so, we partnered with *MalBeacon*, a threat intelligence company, and we studied the activities of 4,586 *Stealer* operators through their devices, over a period spanning 20 months (Apr 2019 - Dec 2020). *Stealers* are specialized commodity malware that harvest credentials from infected hosts. *Stealers* utilize many attack vectors, including drive-by download, application repackaging, remote exploitation, social engineering, and phishing. Stolen credentials – the main goal of *Stealers* operations – are the primary way for cybercriminals to gain initial access to a network and their utilization has risen five times since 2021 [19, 106]. Given the importance of stolen credentials to cybercriminals and the breadth of unique *Stealers* devices in our studied dataset, we believe that our study can offer valuable insights into how malicious cyber threat actors utilize network infrastructure to manage their operations.

The rest of this chapter is organized as follows: section 4.2 describes the datasets and methodology we utilized in this work in order to label and identify unique operator devices, section 4.3 describes the ethical considerations of this work and section 4.4 presents the results of this work answering the research questions we have posed in this section. In section 4.5, we discuss the practical takeaways of this work. Finally, section 4.6 summarizes the main findings of this chapter. Our work accompanies six months of the *Stealers* dataset and the implementation code to foster reproducibility and transparency¹.

¹<https://github.com/Astrolavos/stealer-sec23>

Table 4.1: A list of data sources used in this study.

Dataset	Description	Source
Stealers	Stealer tracker	<i>MalBeacon</i>
Active DNS	Domain reg./resolution	ActiveDNS Project [32]
Passive DNS	Recursive and authority domain lookups	US ISP, Global Recursives, Nameserver Authority, TLD Authority
Threat Intelligence	Malware and domain intel.	URLScan [107], VirusTotal [84] IP Reg. [108], bot tracker [109, 110, 111] residential and mobile proxies [6, 112]

4.2 Data and Methodology

In collaboration with *MalBeacon*, we had initially set out to answer our research questions and gain insights that can help researchers develop better defenses (detection and prevention) and aid law enforcement in pursuing cybercriminals more effectively (deterrence). Unfortunately, the *Stealer* dataset alone does not allow us to explore these questions thoroughly; therefore, we must augment the dataset with external data sources. We rely on DNS and threat intelligence. The DNS dataset characterizes DNS records, volumetrics, and client resolutions. The threat intelligence datasets enrich, validate, and identify additional artifacts of malicious infrastructure. Table 4.1 summarizes our data sources.

Scope. Our work investigates the *harvesting phase* of the credential theft lifecycle. The resale and distribution of the credentials throughout the underground forums or other illicit markets are out of scope. Specifically, this work studies one harvesting channel, namely *Stealer* malware, their *Stealer* operators, and the service providers. Readers can refer to prior works [113, 10, 53, 44, 48, 52] on credential theft profits.

4.2.1 Data Sources

Stealers Dataset. *MalBeacon*, a threat intelligence company, provided us access to their commercially available *Stealers* dataset. *MalBeacon* tracks many *Stealer* families,

Table 4.2: A list of top password stealers found in our dataset.

Family	First Sold	Price	Leaked	Panels ($N = 5,295$)	Hosts ($N = 2,602$)
LokiBot [114]	2015	\$80-\$300	✓	3,613 (68.23%)	1,952 (75.01%)
Formbook [115]	2016	\$29-\$299		1,195 (16.62%)	285 (5.32%)
Amadey [116]	2018	\$600	✓	56 (1.05%)	44 (1.70%)
Baldr [117]	2019	\$100-\$150		32 (0.6%)	32 (1.22%)
Blacknet [118]	2019	Open Source		12 (0.22%)	12 (0.46%)
AZORult [119]	2016	\$100	✓	8 (0.15%)	8 (0.31%)
Neutrino [120]	2013	\$200-\$500	✓	9 (0.17%)	8 (0.31%)
Agent Tesla [121]	2014	\$12-\$69		5 (0.09%)	5 (0.19%)
Nexus [122]	2020	\$100		5 (0.09%)	5 (0.19%)
KPOT [123]	2018	\$85		2 (0.03%)	2 (0.08%)

which are listed in Table 4.2. In our initial analysis, we noticed a skewness in the dataset that can potentially be attributed to the malware’s (Lokibot, Formbook, AZORult) popularity in the wild [124, 115, 119], limitation of the data collection process, or a combination of both. *MalBeacon* uses a *proprietary* pixel-tracking technique, similar to email marketing, embedded into artificial credentials, documents, and other sensitive information that *Stealers* target. When the operator views the stolen information, the browser will request the embedded pixel from *MalBeacon*’s server and reveal information about their device (IP address and user-agent).

Figure 4.1 is an overview of how *MalBeacon* collects the *Stealer* dataset. Step ❶ the *Stealer* infects a system and sends stolen artificial data with the embedded pixel (❷) to the *C&C* server, which is committed to the backend storage (❸). Next, when the operators use their device (❹) to connect to the *C&C* server (❺), they log in to the management panel (❻), where the embedded pixel gets rendered (❼). Before the pixel can render, the operator’s browser will connect to the pixel server (❽) to retrieve the pixel. The pixel server logs the HTTP request from the operator’s browser into an activity log database and generates a unique random long-lived cookie ID that is sent back in the response header. Any subsequent requests by the operator would include the cookie ID, which enables tracking

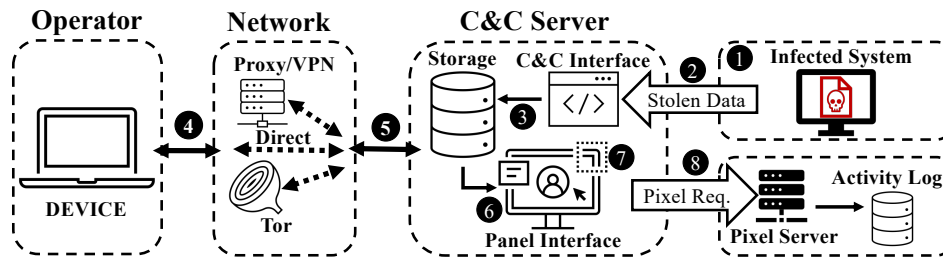


Figure 4.1: An overview of *Stealer* data collection.

operators across different panels. Table 5.2 summarizes the dataset fields and their counts.


MalBeacon did not disclose the proprietary implementation details for their system, but we demonstrate how to collect a similar dataset using the approach found in Nachum et al. [125]. In brief, Nachum et al. modify stolen system artifacts by inserting an HTML image tag alongside the original in the following format: Original Value + Image Tag, i.e., “DESKTOP-UU1KCDG.” When the stolen artifacts are rendered in the HTTP panel interface (*C&C*), the operator’s browser will callback to the image hosting server, and the hosting server will log the IP address, user-agent, and HTTP headers. To test this hypothesis, we implemented the system found in Nachum et al. and tested five *Stealer* malware families (Amadey, Azorult, BlackNet, LokiBot, and Neutrino) for the following browsers: Chrome 96.0.4664.45, Firefox 94.0.2, and Edge 95.0.1020.44.

Table 4.3: *Stealer* dataset fields summary.

Field Name	Description	Unique
Timestamp	The time a tracking event was observed.	202,538
IP Address	IP address used by the operator to access the panel server.	21,812
User-Agent	User-agent string associated with the operator’s device.	1,484
Cookie ID	A session identifier set by the tracker for the operator’s browser.	5,552
Panel Web Address	The referrer field sent to the tracker.	27,823

We collected the same fields (IP address, user-agent, HTTP header) by using a Windows 10 virtual machine and hooked system calls to modify values such as the IP address (Amadey, Neutrino), Computer Name (Azorult, BlackNet), Global Unique Identifier (Lokibot), and Bot Name (Neutrino).

We can utilize additional fields to insert the pixel code, but we leave that for future work. We induced a pixel callback and cookie ID persistence for all families across all three browsers. When testing with private browsing, we observe that the cookie IDs are cleared after each session. Our testing found that privacy features on modern browsers trim the entire referrer field. Specifically, we observed that starting with Firefox 87 and Chrome 89, the path and the query string information of the referrer field are missing [126]. The privacy feature impacts the future collection of similar datasets and limits our cookie merging and malware labeling methodology. However, this work collected the *Stealers* dataset before the browser privacy change (March 2021).

DNS Datasets. We use the aDNS from the ActiveDNSProject [32]. The project resolves over 1,100 zones and includes resolutions for Alexa’s Top 1M and public blocklists. Each domain is resolved twice during 24 hours. We use aDNS to investigate the *Stealer* infras-

structure by enumerating relationships between observed IPs and domains. Furthermore, we use three passive DNS (pDNS) datasets from a US-based internet service provider (ISP), geographically distributed local and global DNS resolvers, and an authoritative nameserver responsible for several zones and a top-level domain (TLD) authority. The pDNS datasets are anonymized to exclude any customer-related information. We use pDNS to amplify the coverage of the stealer domain resolutions and estimate potentially infected networks resolving the stealer domains. Combining these datasets, we get global visibility from over 80 million internet-connected devices.

Threat Intelligence Datasets. We use eight threat intelligence sources, namely URLScan [107], VirusTotal [84], IP Registry [108], residential and mobile proxy dataset [6, 112], and botnet trackers [109, 110, 111]. URLScan implements a website scanning engine to analyze JavaScript, HTML, and embedded content to detect malicious code. VirusTotal (VT) is a threat-sharing platform used by hundreds of commercial companies and thousands of security researchers to share malicious indicators. IP Registry is an IP intelligence service that collects and correlates data from partner networks and public sources like BGP tables, regional internet registry databases, internet service provider data, geofeeds, and latency measurements. The data covers 99.9% of the IPv4 space but excludes loopback, link-local, multicast, private, site-local, and wildcard IPs. The botnet trackers use open-source threat intelligence to track *C&C* servers. The residential and mobile proxy datasets are sourced from an academic study [6, 112] that includes 6.42M residential IPs collected between May 2017 and February 2018 and 8M mobile proxy IPs collected between April and August 2019.

4.2.2 Data Validation

The raw pixel server logs contain HTTP request records where each record has a timestamp, the source IP address, and the HTTP header. *MalBeacon* processes the HTTP headers into three fields: user-agent (UA), cookie ID, and referer field. The final dataset format is

a JSON file that contains the fields in Table 5.2. Our initial analysis of the *Stealer* dataset aims at validating the dataset by inspecting the consistency of user agents, the persistence of cookie IDs, the identification of *C&C* instances, and the labeling by malware families.

User-Agent Validation. We analyze the number of unique browsers and operating systems per cookie ID to investigate if the UA strings are potentially spoofed. If UA spoofing were present, the browser and operating system of the UA per cookie ID would change. We found six (0.01%) cookie IDs with more than one unique browser and 25 (0.45%) with more than one operating system. Manual inspection of those records reveals six cookie IDs with multiple versions of the Windows OS, four cookie IDs with multiple versions of macOS, and 12 cookie IDs with other operating systems (Linux and Android), which suggests potential UA spoofing.

On the other hand, 99.55% of the cookie IDs have only one operating system and browser, with 73.23% having only one browser version. The rest change their browser version, but they are consistent with the release of browser updates. For example, 50% of the devices update their browser within 21 days or sooner, and 75% update their browser version within 41 days or sooner. However, a set of records from Firefox has versions before the update release, which can indicate spoofing or beta/early testing. Those UAs were associated with 145 cookie IDs and 6,068 records. In total, the potentially spoofed UAs account for 6,243 (3.0%) records associated with 170 (3.0%) cookie IDs. We discard those records when we perform operator device measurements.

Lastly, we analyze the top 10 UAs in the dataset and present the results in Table 4.4. We group by OS and browser and count the associated cookie IDs, *C&C*, and the average days between a browser update release and a UA change. The most popular OS is Windows, and the most popular browsers are Chrome and Firefox. The most popular UAs appear to follow a uniform distribution with respect to the associated cookie ID count, which implies that those UAs are not spoofed. We found, on average, 1.25 cookie IDs are associated per *C&C*, while 75% of the *C&C* instances are associated with a single cookie ID. Although these

Table 4.4: Top 10 user agents and related statistics.

OS	Browser	Cookie IDs	C&C	Update (Days)
Windows 7	Chrome 75.0.3770.100	116	119	22.50
Windows 10	Chrome 79.0.3945.130	112	110	24.15
Windows 10	Firefox 68.0	112	140	36.32
Windows 10	Firefox 69.0	111	113	47.23
Windows 10	Chrome 75.0.3770.142	109	120	53.54
Windows 10	Chrome 75.0.3770.100	108	122	21.74
Windows 10	Chrome 73.0.3683.103	95	88	28.57
Windows 10	Chrome 74.0.3729.169	88	109	22.31
Windows 10	Firefox 70.0	82	96	24.95
Windows 7	Chrome 75.0.3770.142	80	72	17.46

statistics imply that the overwhelming majority of the UAs are not spoofed, an operator can still spoof the most popular UAs to masquerade their actual device fingerprint. This is an artifact limitation that we can not verify from the dataset. Realistically, an operator must know the most popular UAs in use with a particular *C&C* panel to spoof a popular UA.

Cookie ID Persistence. The cookie IDs associated with each request may not be persistent if operators clear their browsing history or use private browsing sessions. We refer to the ephemeral cookie IDs as *cookie churn*, where a device is assigned multiple cookie IDs over time because they are not persistent. We find the ratio of cookie IDs per *C&C* panel to be, on average, 1.59 with a median of one and a maximum of 67, which implies that cookie churn is present in a subset of the dataset. We address the *cookie churn* problem by applying a similar technique to the work of Dasgupta et al. [127]. Briefly, Dasgupta et al. address cookie churn for *user-modeling* and *reach-frequency* in the context of online advertisement. User-modeling refers to estimating how many users visit a particular site (users per *C&C* panel), whereas reach-frequency refers to how often an individual user visits a particular site. Our study focuses on user modeling to address the cookie churn problem.

We use the OS, browser, and panel URL as device profiles. In addition, we use two

cannot-link constraints, namely cookie lifespan overlap and browser version. Cannot-link constraints are logical constraints that can disambiguate distinct but similar device profiles. For example, the cookie ID’s lifespan interval (last seen - first seen) cannot overlap. If two device profiles use Windows 10 and the Chrome browser, but the lifespan of their cookie IDs overlaps, then we assume that they are distinct since they access the same *C&C* from similar devices but use different cookies. The browser version constraint merges cookie IDs if and only if the browser version in later records is greater than or equal to the browser versions in earlier records per *C&C* panel.

We design and implement algorithm 2 to analyze and reconcile multiple cookie IDs belonging to the same device. The input takes a set of *C&C* panels and retrieves a set of devices that access the panels (line 2). A device is a tuple of UA string and cookie ID, where the UA is parsed for the OS, browser, and browser version. Once we have a set of devices (*D*), we group the records by the OS and browser and sort them by the first seen date (lines 3 and 4). For each group (*g*), we iterate through the cookie IDs and either allocate a new cluster (line 9) or merge on the profile features and cannot-link constraints (line 16). Since we lack the ground truth to evaluate the accuracy of Algorithm algorithm 2, we define an error metric called *ambiguous merge error* to quantify missed merges. Our merge policy coalesces cookie ID candidates with the earliest cluster (first seen), and therefore, the metric captures how many other clusters the candidate cookie ID could have merged with.

We calculate the *ambiguous merge error* (*AME*) using the following formula: $AME = \frac{|collision|}{|g_i.GetClusters()|}$. Specifically, we calculate the *AME* per group (*g_i*) since the merge error can only occur when the profile features and cannot-link constraints are met for more than one cluster per group. We found 872 groups with at least two cookie IDs. We skip groups with one cookie ID since they cannot be merged. Out of the 872, we detected merge misses in only 29 groups. Furthermore, 19 out of the 29 groups with merge misses are in the top 100 largest groups. The largest *AME* value is 1.22, which indicates that the merge is ineffective, i.e., merge error over 100%. This merge error belongs to the 89th largest

Algorithm 2: Merge device’s cookie IDs.

Input: A set of unique $C\&C$ (C^2)
Result: Merged Cookie ID Clusters

```
1  $Merged \leftarrow \{\}$ 
2  $D \leftarrow \text{GetAssociatedDevices}(C^2)$ 
3  $G \leftarrow \text{Group}(D, \text{by}=[\text{OS}, \text{Browser}])$ 
4 for  $g$  in  $G.\text{sortAsc}(\text{firstSeen})$  do
5     for  $i=0$  to  $g.\text{size}$  do
6         if  $g_i$  in  $Merged$  then
7              $\text{continue}$ 
8          $Merged.\text{addNewCluster}(g_i)$ 
9         for  $j=i+1$  to  $g.\text{size}$  do
10            for  $c$  in  $Merged.\text{GetClusters}()$  do
11                if  $g_j.\text{lifespan}$  not overlap  $c.\text{lifespans}$ 
12                    and  $|g_j.C^2 \cap c.C^2| \geq 1$ 
13                    and  $g_j.\text{browser\_ver} \geq c.\text{browser\_ver}$  then
14                         $\text{MergeWithCluster}(c, g_j)$ 
15 return  $Merged$ 
```

group, which had nine unique cookie IDs and 11 possible merge combinations (ambiguous merges).

We discard groups with large AME values (more than 0.20) for the analysis. We summarize the distribution of AME for the largest top 10, 100, and all groups in Figure 4.2. Eight out of the ten largest groups have less than 0.1 AME rate. Additionally, five out of the ten largest groups have a 0.0 AME rate, which gives us confidence in the results since these groups have many cookie ID nodes. For example, group two has 68 unique cookie IDs and a merge collision count of 0. Beyond the AME metric, we manually inspected the top 100 groups to ensure that Algorithm algorithm 2 correctly coalesced cookie IDs and accounted for merge misses.

C&C Instance Identification and Labeling. The *Stealer* dataset does not contain any malware family labels or panel instance distinction, which makes our analysis challenging. Identifying and labeling the panel instances is essential for us to discern between different malware families and hosting infrastructure. We perform three labeling tasks: identification of panel instances, panel malware families, and panel dynamic DNS domains. A single host can serve multiple panels. We define a panel instance (Π) by the domain or IP address (δ) and URL path (ρ). More formally, $\Pi = \{\delta, \rho\}$ where the δ can be a domain or an IP




Figure 4.2: Distribution of AME per top largest group.

address and ρ is the URL path starting from the domain/IP to file name and extension (γ). For example, the following illustrates the components of a panel URL address:

$$\text{http://}\overbrace{\text{domain.tld}}^{\delta}\overbrace{\text{/path/file.ext}}^{\gamma}\text{?param=1}$$

ρ

We label records that do not contain URL paths as unknown and exclude them. Next, we assign a malware family label to the panel instances. We rely on the panel’s URL components, such as the path (ρ), file name, and extension (γ), and parameters. We manually create *Stealer* family label signatures based on leaked source codes and panel tracker services [111, 109, 110]. Figure 4.3 presents our labeling process. In step one (❶), we extract URL patterns and labels from our source code and panel trackers. Next (❷), we use the strings and their order to generate a fingerprint for each *Stealer* family. In step three (❸), we store the signatures and the family labels in the database. Finally, in step four (❹), we label the panel instances based on the derived signatures. The signatures are in the form of regular expressions. From the 202, 538 records in the *Stealers* dataset, 15, 237 (7.5%) are

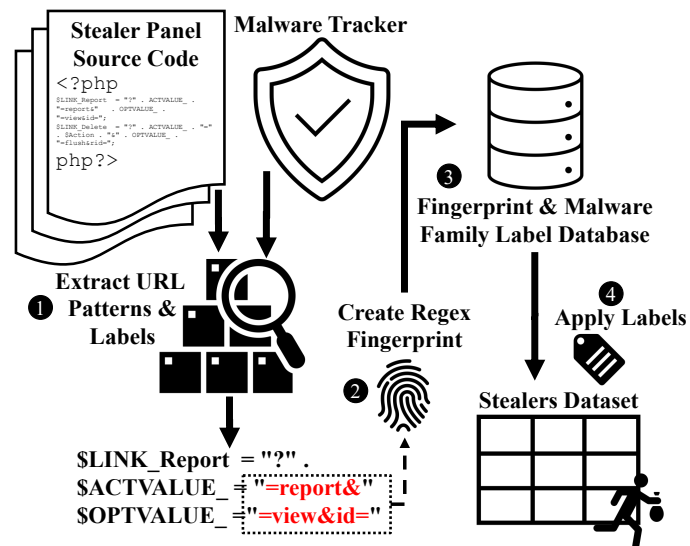


Figure 4.3: Panel signature generation and identification.

associated with 357 (6.7%) panel instances with unknown labels. We attempted to use the AV labels from the malware files associated with each panel instance; however, we found them unreliable and noisy [90]. For Effective Second-Level (E2L) Dynamic DNS domains (DDNS), we manually verify them to ensure there are no false positives, and we use pDNS to identify domains with 50 or more subdomains.

4.3 Ethical Considerations

We take our ethical and legal responsibility seriously and ensure our study does not violate widely accepted norms. Our institute reviewed our request for an IRB and concluded that we do not require an IRB review. We also presented our study to the institute’s Office of Cybersecurity for compliance, and they had no concerns. This study uses data collected by *MalBeacon*, a US-based commercial company that operates and adheres to the Computer Fraud and Abuse Act (CFAA). The collection technique *does not* actively scan, exploit, or social engineer the malware operators in any way, and an external legal review committee reviewed *MalBeacon*’s tracking method and deemed it compliant with the Computer Fraud and Abuse Act (CFAA) and the Directive on attacks against information systems. The approach relies on honey tokens that many prior works use [55, 56, 128, 58, 59, 60, 57], which date back to 2004. Moreover, our dataset analysis follows the precedence of prior works that study malware operator activities [54, 50, 43].

Research of criminal activity often involves deception or clandestine research activity [129, 130], so requests for waivers of both informed consent and post-hoc debriefing may be relatively common as compared with research studies of non-criminal activity. Support for such waivers is recommended when the research involves no more than minimal risk to the subjects, and the research could not be carried out without the waiver. Deception is necessary for the *Stealers* dataset to obtain data that characterizes the *Stealer* ecosystem. Such studies are considered permissible when (1) the research addresses important questions of public concern, (2) the research cannot be conducted if the subjects must pro-

vide consent, and (3) involving subjects in the research without their permission does not significantly compromise their autonomy. This study meets all three criteria, and the scope follows well-established Menlo guidelines. Furthermore, our study analyzes a commercial dataset (passive observations) and does not directly implicate any malware operators or cause direct harm.

Finally, the data contains no personally identifiable information (PII). The IP address can be considered PII with additional auxiliary data, but not by itself. From a law-enforcement perspective, an IP address can be subpoenaed by the ISP to get PII information about the person leasing the IP address at a given time. We do not have legal authority or access to auxiliary information to identify individuals. Despite well-established guidelines on deceptive studies and issues regarding PII, we note that computer security research is more like behavioral research because the risks typically are not physical and can be challenging to quantify. Although evidence indicates that harm resulting from deceptive experiments is minimal and transient, it is still incumbent upon us to identify and minimize potential harm. We reiterate that we take the responsibility seriously and ensure our study does not violate ethical norms.

4.4 Analysis Results

To answer our research questions, we study how *Stealers* use internet infrastructure and analyze how *Stealer* operators administer their botnets by characterizing their devices, networks, and activities.

4.4.1 Stealers on the Internet

Our analysis of the *Stealers* public code shows that *Stealers* require minimal hosting infrastructure. We further seek to characterize *Stealer* hosting on the internet. Specifically, we characterize the domains and hosting networks of *Stealers*, quantify the detection delay between infrastructure setup and blocklist detection, and assess the potential infections

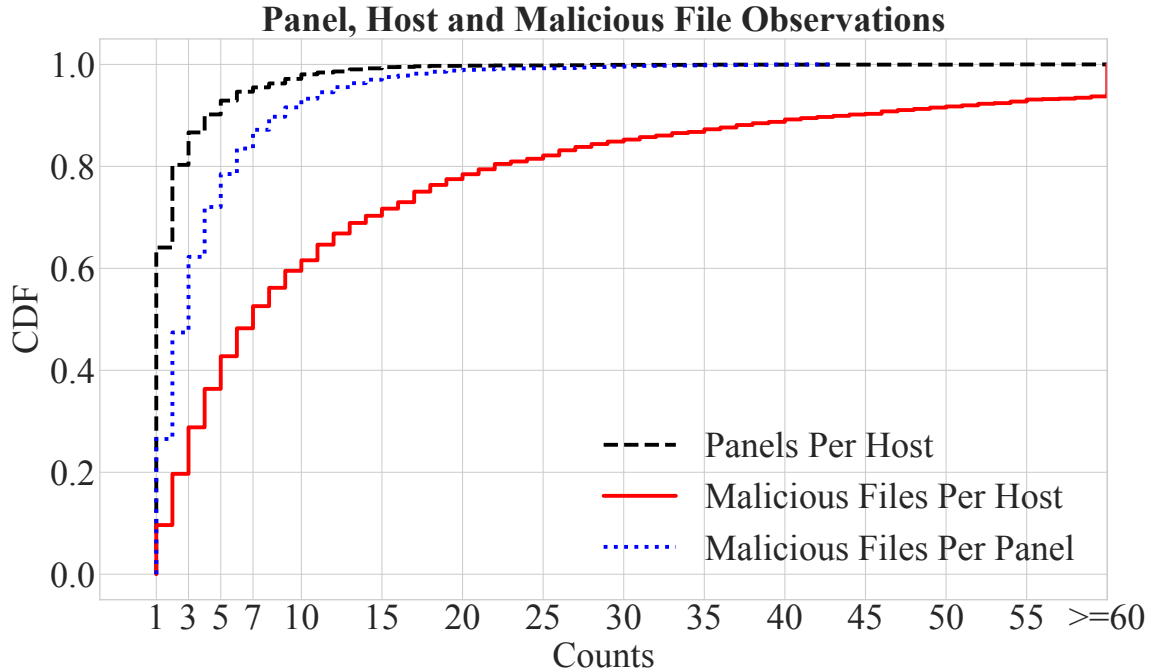


Figure 4.4: Distribution of panels and associated malware.

indirectly through the DNS dataset.

Internet Infrastructure. The *Stealers* dataset contains 2,187 registered domains, out of which 78 are DDNS and web hosting domains, and 281 panel hosting IP addresses for a total of 2,468 unique panel servers (hosts). This count excludes the two bogon panel IP addresses and three popular non-malicious domains in the Alexa top 100K [131]. Table 4.5 summarizes the top 10 top-level domains (TLD) count for effective second-level domains (E2LD)s of the *C&C* panels. For the panel domains, 41.4% use the COM TLD, followed by 19.0% that use free country code domains (ccTLDs) like TK, ML, CF, and GA. Free ccTLDs are known to be heavily abused by malware [132]. The right side of Table 4.5 summarizes the top ten network names for the *C&C* panels, which account for 70% of the hosts. About 30.9% use US-based hosting (Cloudflare, Namecheap, and Unified Layer), 15.8% use Russian-based hosting (Reg.ru, SelecTel, Mail.Ru, The First, and IHOR-AS), and 12.2% use Chinese-based hosting (Alibaba Cloud and Tencent).

In Figure 4.4, we present the distribution of panels and associated malware files per host




Figure 4.5: Distribution of domain events and detection.

and panel, respectively. Note that we differentiate between the host and the panel since a host can serve multiple panel instances. We observe that 64% of the hosts serve a single panel, 26% of the hosts serve between two and four panels, and 9.8% of the hosts serve five or more panels; the largest host has 71 panel instances. We find that 61.5% of the hosts have 10 or fewer malicious files associated with them. The number of malware files per panel and host has a maximum value of 43 and 249, respectively.

Table 4.5: Top 10 TLDs and hosting networks for panel hosting server domains.

TLD	Domain (%)	Type	Reg. Cost	Network	Domain (%)
COM	874 (41.5%)	Commercial	\$8.38	CLOUDFLARENET	308 (14.1%)
GA	107 (5.0%)	Country Code	\$0	NAMECHEAP-NET	263 (12.0%)
XYZ	105 (4.9%)	General	\$0.99	CNNIC-ALIBABA-US-NET-AP	197 (9.0%)
ML	97 (4.6%)	Country Code	\$0	UNIFIEDLAYER-AS-1	105 (4.8%)
INFO	94 (4.4%)	Information	\$2.99	SELECTEL OOO	86 (3.9%)
TK	79 (3.7%)	Country Code	\$0	AS-REGRU	79 (3.6%)
ICU	73 (3.5%)	Business	\$1.99	TENCENT-NET-AP-CN	71 (3.2%)
CF	66 (3.1%)	Country Code	\$0	Mail.Ru LLC	64 (2.9%)
TOP	61 (2.9%)	General	\$0.99	THEFIRST-AS JSC The First	61 (2.8%)
GQ	56 (2.6%)	Country Code	\$0	IHOR-AS Ihor Hosting	57 (2.6%)

Detection of Stealer Hosting. Next, we want to assess if public blocklists detect *Stealer*

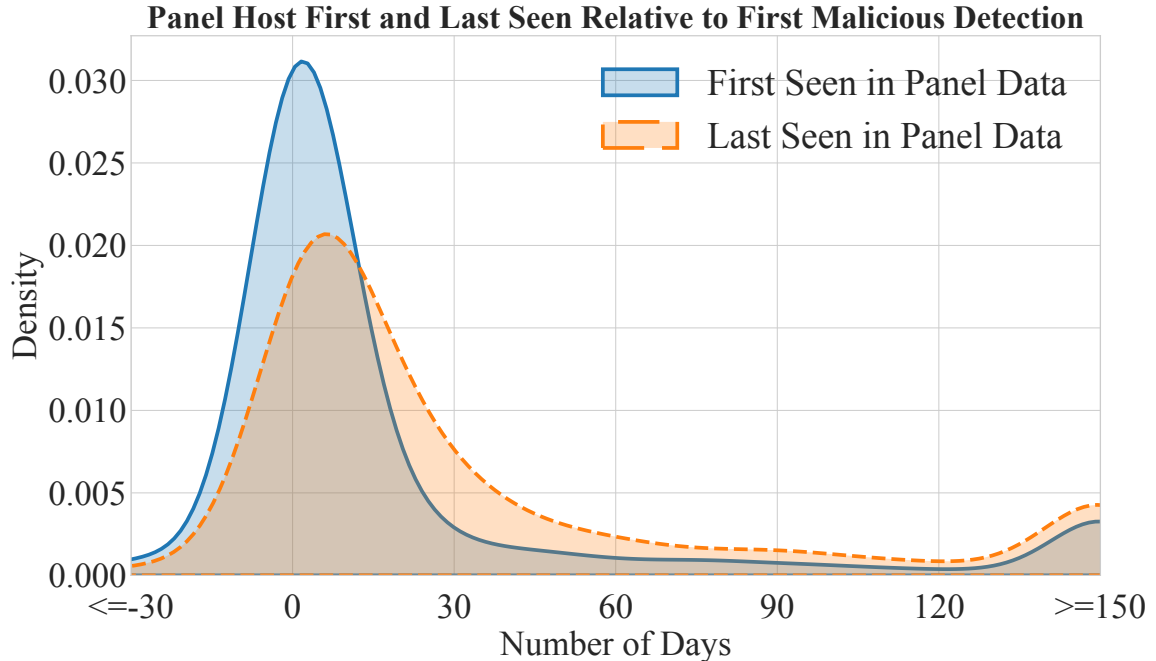


Figure 4.6: Distribution of the time delta for events and detection.

infrastructure, and if they do, what is the time delta between the domain setup and detection. The time delta can inform us of the current defense efficacy against *Stealers* and identify limitations researchers can address. We find that 95% of the *Stealer* hosts appear on the VT historical blacklist. Surprisingly, 123 hosts do not appear on public blocklists. We investigated the 123 hosts and found no notable difference from the detected domains. Figure 4.5 quantifies the detection timeline for 52.58% of the newly registered *Stealer* domains with no prior DNS history (first-time registration). The plot shows the distribution of the events for new DNS records (solid blue line), malicious detection (dotted orange line), and the first operator activities in the *Stealers* dataset (dashed green line).

The average and median time for the first observed DNS record is 15 and two days, respectively. The pDNS data shows that operators set the DNS records within the first week after registration for 77% of the domains. We find the average and median time for detection is 74 and 11 days, respectively. Notably, the operators continue to access the *Stealer* hosts even after detection for an average of 74 days. On the other hand, 53.26% and 69.03% of the *Stealer* hosts stop operating 14 and 30 days after appearing on blocklists,

Table 4.6: Networks resolving stealer domains by country for residential, business, and government networks.

Type	Client Networks		Residential Networks				Business Networks				Government Networks					
	Count (%)	Countries	Count (%)	Vol.	Days	Vol/Day	Countries	Count (%)	Vol.	Days	Vol/Day	Countries	Count (%)	Vol.	Days	Vol/Day
Hosting	67,958 (40.5)	China	4,187 (14.1)	607,282	473	1,283	United States	25,315 (92.8)	1,441,020	500	2,882	United States	113 (54.6)	40,161	328	122
ISP/Telco	37,463 (22.3)	Morocco	3,313 (11.2)	47,854	351	136	Vietnam	619 (2.2)	2,004,091	348	5,758	Canada	14 (6.7)	405	25	16
Residential	29,595 (17.6)	India	2,556 (8.6)	135,815	466	291	United Kingdom	309 (1.1)	1,652,777	420	3,935	China	8 (3.8)	604	139	4
Business	27,269 (16.1)	United States	2,293 (7.7)	195,714	481	406	S. Korea	152 (0.5)	18,798	276	68	Italy	6 (2.9)	265	60	4
Education	5,143 (3.0)	Iran	1,479 (5.0)	16,929	429	39	India	117 (0.4)	5,399	287	19	Indonesia	5 (2.4)	7	6	1
Government	207 (0.1)	Mexico	1,410 (4.7)	75,469	403	187	Nigeria	108 (0.4)	7,615	212	36	Israel	4 (1.9)	235	57	4
Health	188 (0.1)	Indonesia	1,360 (4.6)	48,557	352	137	China	69 (0.2)	182,895	361	506	India	4 (1.9)	4,264	80	53

respectively. For 43% and 28% of the newly registered panel domains, we find that they are detected within one week and after two months, respectively. The remaining *Stealer* domains go undetected for an average of 64 days and a median of six days after their first DNS resolution. Within the undetected domains, 33% remains undetected for over a month.

We observe, on average, 87 days between registration and first appearance in the *Stealers* dataset, with a median of 20 days. *MalBeacon* integrates with VT to share samples, which may correlate with the median time to detection (20 days). Additionally, Figure 4.6 shows the time window distribution for the first and last seen activity from the *Stealers* dataset centered around the first malicious detection of a panel host observed in VT. Almost 70% of the panel hosts appear in the *Stealers* dataset within seven days or less after their first detection. In summary, operators provision *Stealer* hosts within two weeks. They appear on blocklists within 74 days on average. Operators continue to access the *Stealer* hosts for an average of 74 days after their detection.

Assessing Victim Targeting.

We estimate the number of targeted victims to understand the impact of *Stealers*. To get an accurate estimate, we would require direct access to the *C&C* server, which we do not have. Instead, we use the pDNS dataset to estimate the number of potential infections by analyzing the DNS resolutions. We quantify the number of DNS resolutions by network types and countries during the active time frame of each domain in the *Stealers* dataset. We define a network by the EDNS Client Subnet (ECS) [92, 133] found in the DNS resource records for clients resolving domains above the recursive, where the DNS recursive query the upper DNS hierarchy (i.e., TLDs and authoritative name servers). It is important

to note that the results are associated with subnets, not IP addresses, which can underestimate the number of targeted victims. Moreover, we base the analysis on potential, not confirmed, infections.

We observe a total of 255,925 unique networks, but we can only label 167,989 (65.6%) of them. Table 4.6 presents the results. The table has four parts, namely the *Client Networks*, *Residential Networks*, *Business Networks*, and *Government Networks*. The *Client Networks* is a breakdown of all 167,989 labeled networks. The *Residential Networks* category is a breakdown of the networks that belong to residential subnets grouped by country. The *Business Networks* is a breakdown of the labeled business subnets grouped by country. The *Government Networks* category is a breakdown of the labeled government subnets grouped by country. For each network label, we show the network count (Count), lookup volume (Vol), days queried (Days), and lookup volume rate (Vol/Day).

Table 4.7: Top 10 hosting networks querying stealer domains.

Hosting AS	Networks
AMAZON-AES	30,705
AMAZON-02	12,515
CLOUDFLARENET	5,890
MICROSOFT-CORP-MSN-AS-BLOCK	4,708
OVH OVH SAS	1,623
DIGITALOCEAN-ASN	728
MAXIHOST	543
M247 M247 Ltd	536
SOFTLAYER	461
UK2NET-AS UK-2 Limited	332

We find that 40.5% of the resolutions originate from *Hosting* networks. These networks appear to be associated with virtual private server (VPS) providers, cloud providers (i.e., AWS, OVH, Azure), and content delivery networks (CDNs), as shown in Table 4.7. The rDNS records show that VPS and cloud networks account for virtual private network (VPN)

services. Moreover, a portion of cloud networks and most of the CDNs appear to be internet scanners or security tools. These observations align with prior works on malicious domain sinkhole analysis [17]. However, many hosting networks are unlikely to be infected clients.

We observe ISP/Telco as the second and Residential as the third most popular networks. The residential networks are more likely to be victims since ISPs designate the space for home users. For the *Residential Networks*, we observe that Chinese clients make up 14.1% of the potential infections, followed by Morocco (11.2%), India (8.6%), and the United States (7.7%). Notably, we find 207 government networks resolving *Stealer* domains. We took a closer look at the 113 U.S. government networks and found a mix of federal (24), state (32), and local (58) government networks. At the federal level, we found high-profile government networks like the U.S. Social Security Administration (4), the U.S. House of Representatives (2), and the U.S. Senate (2).

Investigating further, we found a total of 107 DNS responses for 27 different *Stealer* domains from August 2019 to November 2020. More specifically, for the U.S. Senate network, we observe a total of 12 distinct resolutions for nine domains from January 2020 to July 2020. These DNS resolutions originate from what appear to be the DNS recursive servers for the U.S. Senate network. These resolutions suggest that there may be more infections because the DNS resolutions are typically cached. Nevertheless, the sensitivity of these government networks, including the U.S. Social Security Administration, demonstrates the wide reach and impact of *Stealers*. Finally, the infection period for all 28 domains extends over a year, giving operators ample time to execute other capabilities (e.g., keylogging, dropping malware, and reverse shell).

Takeaway: We find *Stealer* infrastructure requires minimal hosting resources and abuse services such as free ccTLDs and cloud-fronting. Moreover, on average, public blocklists detect *Stealer* domains 74 days after the initial registration with a median of 11 days. This detection gap gives *Stealers* ample time to infect and harvest credentials from various networks. Their long-lived activities may be problematic, as they allow operators

time to exercise other malware capabilities (i.e., install ransomware [134]).

4.4.2 Characterization of Operators

The *Stealers* dataset provides a unique vantage point to characterize how *Stealer* operators manage their botnet using the *C&C* panels. We take a closer look at how operators interact with the *C&C* panels through their devices and shed light on their tactics.

Device and Network Characteristics. Characterization of the device and network association can inform researchers about common patterns cybercriminals use. These characteristics can help build heuristic-based defenses that profile device and network properties to flag suspicious and unauthorized access. On average, operator devices access panels using 6.66 IP addresses that belong to 1.95 autonomous systems (ASNs). The largest number of IP addresses associated with an operator device is 230, belonging to nine ASNs. Moreover, the standard deviation for operator device IP addresses is almost double the average (12.7). When looking at how operators access their *C&C* panels, we find, on average, operator devices access 1.62 unique panel instances, 1.51 unique domains, and manage 1.04 malware families. The operator device with the most panel instances accesses 57 unique panels hosted on 42 distinct domains. We take a closer look at this particular example and find that the 42 distinct domains use algorithmically generated domains (DGA).

After applying the cookie merging algorithm (Algorithm algorithm 2), we find operator devices associated with 1.17 cookie IDs on average. The operator device with the most cookies has 55 unique cookie IDs. For over a month, this device used the same operating system, browser, and browser version to access the same panel with 55 unique non-overlapping cookies, suggesting cookie churn.

In the entire *Stealers* dataset, 465 (10.14%) operator devices have more than one cookie ID. We find, on average, 5.7 more IP address associations for these devices. Cookie merging (Algorithm algorithm 2) helped us build a complete profile for these operator devices and uncover related IP addresses and ASN associations that we would have missed

otherwise. Cookie churn fragments access patterns, and we must address the churn to build a more accurate device profile. Additionally, operator device profiles are diverse and can help distinguish between operators.

Operator Devices Types. Next, we take a look at the type of devices used by operators to understand the type of devices they utilize. Table 4.8 summarizes the overall statistics for the operators’ devices. There are 4, 586 unique operator devices associated with two types (desktop and mobile), 19 different device vendors, 70 different device models, eight different operating systems, and 8 different browsers. We find 4, 467 (97.40%) of the operator devices are personal computers (PCs), while the rest 2.6% are mobile devices. Among the PC operator devices, 4, 282 (95.85%) use a version of Microsoft Windows, 135 (3.02%) use a version of Apple’s macOS, and 50 (1.12%) use a Linux/Unix system. Among the mobile operator devices, 95 (91.34%) are Android while the rest 9 are iOS. In terms of browsers, 3, 232 (70.47%) operator devices use a version of Google Chrome, 1, 293 (28.19%) use a version of Mozilla Firefox, while the remaining 40 use other browsers. These results suggest that most *Stealer* operators conduct their work from personal computers (desktops or laptops) using popular browsers, but they are utilizing 3 times more the Mozilla Firefox browser compared to the average user[135].

Table 4.8: Top panel operator device types and operating systems.

OS	Desktop		OS	Mobile	
	Ver.	Count (%)		Ver.	Count (%)
Windows	10	2,148 (46.84)	Android	9	22 (0.48)
	7	1,112 (24.25)		8.1	18 (0.39)
	8.1	893 (19.47)		7	17 (0.37)
MacOS	10.14	47 (1.02)		8	9 (0.19)
	10.15	29 (0.63)		10	9 (0.19)
	10.13	26 (0.56)		6	6 (0.13)
Linux	All	50 (1.10)	iOS	12	8 (0.17)

Networks Access Patterns. We analyze the network types, the use of proxies, and the




Figure 4.7: The diurnal analysis for the top 20 countries of operator device activity (dark more active and light less active).

localized diurnal access times to investigate the access patterns. In total, operator networks originate from 135 different countries with different network classifications. The network classifications include ISP (11.55%), ISP-Mobile (55.14%), and hosting networks (31.71%). Interestingly, over half of the operator networks are classified as ISP-Mobile. The bar graph in Figure 4.7 presents the top 20 countries for operator device networks. Most ISP (80.32%) and ISP-Mobile (84.72%) networks are located in Nigeria. Revealingly, 99% of the internet broadband in Nigeria relies on mobile wireless connections [136]. Using the residential and mobile proxy dataset [6, 112], we intersect the timestamp and IP address of each operator device and find 882 (4.04%) mobile proxy records matching against the operator IP addresses. However, omitting the timestamp field to only match against the IP, we find 1,785 (8.43%) and 5,667 (26.76%) matches for residential and mobile proxies, respectively.

Furthermore, we analyze the number of Tor exit nodes associated with the operator networks and present the overlap per country in Table 4.9. Nigerian IP addresses make up

Table 4.9: Top 10 countries of operator IP addresses and their proxy and tor networks.

Country	IPs	Mobile Proxy (%)	Residential Proxy (%)	Tor Exit Node (%)
Nigeria	11,375	4,326 (38.03%)	1,181 (10.38%)	0 (0%)
United States	1,936	161 (8.32%)	36 (1.86%)	15 (0.77%)
Great Britain	908	153 (16.85%)	65 (7.16%)	7 (0.77%)
South Korea	812	170 (20.93%)	14 (1.72%)	0 (0%)
Germany	496	40 (8.06%)	47 (9.47%)	10 (2.01%)
Netherlands	418	33 (7.90%)	31 (7.41%)	5 (1.20%)
Turkey	291	19 (6.52%)	16 (5.50%)	0 (0%)
Canada	279	23 (8.24%)	24 (8.60%)	3 (1.07%)
France	231	28 (12.12%)	21 (9.09%)	2 (0.86%)
Norway	222	4 (1.80%)	4 (1.80%)	0 (0%)

about 53.73% of the operator networks, and 42.55% were observed as proxy networks. Additionally, the top networks classified as hosting are also strongly associated with VPN services. For example, we find most hosting networks to be located in the US, Great Britain, Germany, and the Netherlands, and the top 3 ASNs are: AS9009 M247 Ltd, AS198605 AVAST Software s.r.o. and AS205016 HERN Labs belong to VPN services [137, 138, 139]. We cross-checked the hosting networks with IP intelligence feeds and found that IPRegistry [108] labels them as VPN networks. These findings suggest that *Stealer* operators use proxy networks like residential, mobile, Tor, and traditional VPN services when accessing the management panel. These findings demonstrate that operator profiling can be involved and naively using the operator networks to attribute cybercriminals can be inaccurate.

Operator Device Diurnality. Diurnal analysis can provide another perspective on the nature of operator device access. We can use the analysis as an additional confluence signal for the geographical location. We quantify the access frequency for only *ISP-based* IP addresses not found on the proxy lists. The time zones for the diurnal analysis are *localized* to the geographical region associated with the operator’s IP address. Figure 4.7 presents the diurnal access patterns for ISP-based (Mobile and Landline) operators. We present the top 20 countries, which account for 95.60% of the ISP-based operator device IP addresses in the dataset, and make up 63.70% of the IP addresses of the entire dataset.

The time zone localization shows higher activity on weekdays than the weekends for most countries. For example, the Nigerian diurnal profiles have double the weekday activity compared to the weekends.

The results suggest that most operator devices are more active on weekdays, regardless of the potential victim connections. Those diurnal activities can imply that operators manage *Stealer* as a full-time job as they are mostly connecting during weekdays. The higher activities observed on the weekend for some regions (Russia, Spain, Namibia) can suggest that these operators use proxy networks and do not necessarily reside there. More importantly, when combined with other signals (device fingerprint, network, and access profiles), these observations can provide higher confidence in the operator device profiles.

Takeaway: Operators use proxy services ranging from traditional VPNs to mobile and residential proxies to Tor networks. In particular, the mobile and residential proxies can cause misdirection when characterizing operator profiles. The cookie IDs are reasonably persistent with the majority of the devices in the dataset, but for some operators, private browsing results in ephemeral cookie IDs. The diurnal analysis suggests that operators administer their botnet as a full-time job.

4.5 Discussion

During this chapter, we investigated how malicious threat actors conducting *Stealers* operations utilized network infrastructure. Our findings have several practical applications that could be utilized by network analysts and law enforcement.

Actionable Insights. How can researchers and law enforcement act on these insights? For researchers, we empirically document that *Stealers* have defensive tactics to prevent active scanning and identification of *C&C* panels. Researchers can incorporate this information to build a tailored internet-wide scanning system to find *C&C* panels. For example, a scanner can scan a target host twice, once to trigger a block and a second time to check if the connection is blocked. This approach turns the *Stealer* defense system against itself

and allows researchers to detect possible *C&C* panel hosts. Additional insights, such as geographical distribution of infrastructure, ASN association, and infrastructure characteristics, can inform researchers to design and evaluate adequate active *Stealer* infrastructure detectors.

Law enforcement can apply our operator device profiling techniques to characterize cybercriminals accurately. We show operators use private browsing and diverse proxy services to masquerade their fingerprints. However, law enforcement can build a more accurate device timeline and *C&C* panel access as forensic evidence using our cookie churn merging algorithm and diurnal analysis. Moreover, the affiliation analysis can identify cybercriminal groups and pinpoint their top active participants, which can help law enforcement efficiently go after influential operators. Similarly, our findings can help researchers to identify active *Stealer* infrastructure and prioritize their cleanup. For example, researchers and law enforcement can collaborate to takedown domains with large clusters of operator activities. Lastly, our infection analysis can lead law enforcement to investigate sensitive networks with potential *Stealer* infections.

Operator Attribution Attribution can be of two types, namely, virtual or physical. Physical attribution requires jurisdiction and legal access to private information. Additionally, an ethical aspect of physical attribution must adhere to some acceptable policies and norms. This work focuses on virtual attribution to identify operator affiliation, albeit these techniques are meant to complement and enhance existing methods instead of being used independently. Virtual attribution deals with identifying and tracking different threat groups based on indicators of compromise (IoC). However, we caution the reader that attributing to a specific group is complex, and we avoid making speculative judgments. For instance, our dataset shows that a significant number of activities come from Nigeria, but this can be misleading for a forensic analyst because it does not represent the whole picture. Although this observation is suggestive, we observe that many Nigerian operator networks are mobile or residential proxies. Enigmatically, these proxies appear to be part of anonymity networks

(similar to Tor), where participants may be willingly or unknowingly tunneling traffic [6, 112]. Nevertheless, law enforcement could incorporate our techniques to improve virtual and physical attribution.

4.6 Summary

Our empirical analysis of *Stealers* sheds light on the infrastructure and the lifecycle of the interactions of cybercriminals with it. We found that operators quickly provision their C2 infrastructure within 14 days after registration to their domain names and much of the *Stealers* infrastructure to be long undetected, with public blocklists detecting *Stealer* domains on average 74 days after initial domain registration, which gives operators plenty of time to infect more victims. *Stealers* operators conduct their campaigns utilizing minimal hosting resources and abuse services such as free ccTLDs and cloud-fronting. Operators use proxy services ranging from traditional VPNs to mobile and residential proxies to Tor networks, and the mobile and residential proxies they utilize can cause misdirection when characterizing their profile; thus, law enforcement and security analysts have to be careful in their attributions. The diurnal analysis of the operators suggests that they administer their botnet as a full-time job. Last, we find that 69.03% of the operators stop utilizing their panels within 30 days of a detection event, suggesting that while they do not immediately abandon their operations after detection, the detection event is critical to curb their operations. Future works need to invest more resources in more prompt detection of such management panels and counter the anti-scanning defenses that operators place in order to avoid being easily identified.

While so far, chapter 3 and chapter 4 have provided us with unique insights into the temporal behavior of malware communications as well as a deeper understanding of the network modus-operandi of cybercriminals, they have only partly characterized the infrastructure lifecycle, as they have mainly focused on communications towards the infrastructure servers –targeting– and the interactions of the cybercriminals with their servers while

their operations were active. In the next chapter, we focus on studying the temporal dynamics of malicious domain names by identifying when they have been historically active without the need for proprietary techniques as showcased in this chapter, and demonstrate how that can enable network and security analysts to get a more comprehensive view of the IP infrastructure of historical cyber attacks.

CHAPTER 5

UNDERSTANDING THE LIFECYCLE AND INFRASTRUCTURE OF APT DOMAIN NAMES

5.1 Motivation

In this chapter, we focus on characterizing the temporal dynamics of malicious domain names and propose a novel methodology for the accurate discovery of historically utilized IP infrastructure. Since prior work [10] and our takeaways from chapter 3 and chapter 4, point out that less sophisticated and commodity threats register and utilize domain names rapidly, we focus our scope for this work on sophisticated threats. Such threats have been demonstrated to strategically age their domain names [140, 141] in order to bypass common detection systems [142] that focus on features such as the age of the domain name [143, 144]. Although sophisticated threats do not exclusively utilize strategically aged domain names, given their sophistication, they are more likely to do so, thus their infrastructure is a more applicable target for historical and temporal characterization relative to common threats that rapidly utilize domain names after detection.

Advanced Persistent Threats (APTs) are attacks conducted by well-organized, well-funded, and technically sophisticated actors [145]. The term APT, likely coined in 2006 by analysts of the United States Air Force [146], is used to differentiate commodity and low-sophistication operations (e.g., script kiddies) from those that are more complex and often backed by nation-states and sophisticated crime syndicates. The sophisticated and unique *modus operandi* of these actors—as captured by MITRE’s cyber kill chain [147]—has led to specialized mechanisms for APT detection and investigation [68, 69, 70, 71, 72]. Despite active APT research, recent attacks have continued to cause widespread damage, such as the SolarWinds supply chain attack that forced more than 18,000 customers (including the

US government) to install malicious code [148] or the 2025 *Bybit hack* [149] that stole \$1.5 billion worth of digital tokens.

Prior work on APTs has been mainly focused on detection and investigation systems [68, 69, 70, 71, 72, 73, 74] either aiming at identifying APT attacks in real-time or aiding with forensic investigations. Measurement studies have focused on understanding the attack surfaces of organizations targeted by APT actors [75], the vulnerabilities they use [150], or the tactics, techniques, and procedures (TTPs) they employ [151], or sophisticated attacks against specific targets [15] and regions [16]. Despite the prior work to understand and combat APT attacks, APT investigations still remain a highly manual effort done by experts [22]. Among the top challenges expert APT analysts currently face is that the *“lack of automation and validation in data ingestion impacts the use of historical threat data [22].”* While these challenges are evident across different signals of APT investigations, such as TTPs and malware, they also pose a major problem in the utilization of Indicators of Compromise (IoCs), such as domain names and IPs, which remain the primary signals for APT attribution [22]. Aside from aiding expert APT analysts in investigation and attribution efforts, characterizing and contextualizing the network infrastructure (i.e., domains and IPs) of APTs, which has been demonstrated to be lacking from public reports and threat intelligence [152, 153], can help us answer and quantify research questions that are still largely unanswered. For instance, the network infrastructure comprehensiveness of public threat reports, the longevity of APT infrastructure before disclosure, and the infrastructure utilization trends and similarities of APT groups over the years are still open research questions. Answering these questions can help the community devise more comprehensive defensive strategies, develop more effective attribution systems by utilizing network attributes, and understand how long organizations need to keep network logs in order to detect whether they have been a victim of an APT attack, considering that APTs are particularly persistent compared to commodity threats, thus requiring higher log retention windows.

One of the main challenges in trying to answer the aforementioned research questions

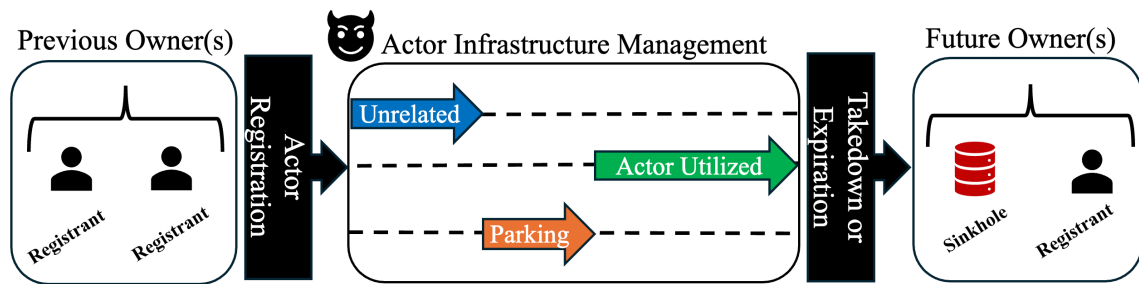


Figure 5.1: Lifecycle of an actor-controlled domain name. Multiple owners and infrastructure types complicate forensics.

is the fact that the relationship between the network infrastructure used to orchestrate an attack and the APT actors is transitory, as is evident in their domain names [33, 34] and as shown in Figure 5.1. For instance, an APT actor can register a previously expired domain name, park it at parking infrastructure, point it to their attack infrastructure for a few days, and then let it expire or be taken down. Another challenge is the fact that APT attacks can persist for years, and the actors can dynamically change the IP addresses utilized by their domain names. Thus, to comprehensively and accurately identify the network infrastructure associated with an APT domain and its lifetime, a forensic analyst needs access to a dataset that is capable of witnessing the historical IP changes, has to filter out unrelated and noisy infrastructure (e.g., parking and sinkhole, etc.), and finally pinpoint the infrastructure and period of time in which each domain name was "active". These challenges diminish the usefulness of networks IoCs just as they are extracted from threat reports, and require analysts to invest manual effort to enrich, contextualize, and validate them, which is time-consuming [20], and is typically conducted on a per-incident basis [154].

In this work, we reduce the knowledge gap in the network infrastructure of APT attacks by performing the first longitudinal study of APT infrastructure used by 405 APT actors over a period spanning a decade. We focus on measuring and expanding the comprehensiveness of the publicly known IP infrastructure of APT attacks, by enriching known, high-confidence APT domain names appearing across 2,188 APT reports with historical DNS data. To this end, and considering the measurement challenges we discussed, we utilize

two historic DNS datasets [155, 32] that witness changes to over 1,100 generic top-level domains (gTLDs) daily, and a novel measurement methodology that automatically and accurately characterizes historic APT infrastructure. Our novel measurement methodology, called Atropos, filters and labels domain-to-IP mappings – Resource Records – related to known domains of APT actors, while discarding IP addresses that are unrelated to APT attacks (i.e. parking, sinkholes, etc.), providing needed automation that expands, validates and contextualizes historical threat data, which has been recently characterized as a major challenge by APT experts [22]. Our contributions are as follows:

- A novel measurement methodology that expands and contextualizes the network infrastructure of known APT domain names offering three times the IP visibility and domain contextualization than that of public threat reports. We will make the code of Atropos available.
- The largest and most comprehensive APT infrastructure analysis to date, spanning over a decade and 405 APT actors.
- We quantify the time window during which organizations need to keep network logs to identify the vast majority of the infrastructure of an APT attack. Our results show that the network logs should be preserved for at least 19 to 25 months.
- We find that while APT actors utilize a plethora of different hosting providers, they only re-use a small portion of them, and that over the years, the use of cloud-fronting has increased significantly. These findings verify expert knowledge [22] and make network forensics and attribution harder.

5.2 Challenges in Domain Lifecycle Analysis

APT network forensic investigations are often conducted as more of an art than a science. Among the many challenges that network forensic investigators must address, identifying the period of time in which the APT attack was *active* has traditionally been a highly

manual process. In the following, we outline the main challenges investigators face in temporally bounding the active period of the APT attack (subsection 5.2.1), then put these challenges into perspective using the SolarWinds attack as a case study (subsection 5.2.2) and finally outline the scope and requirements we need to measure the lifecycle of APT domains (subsection 5.2.3).

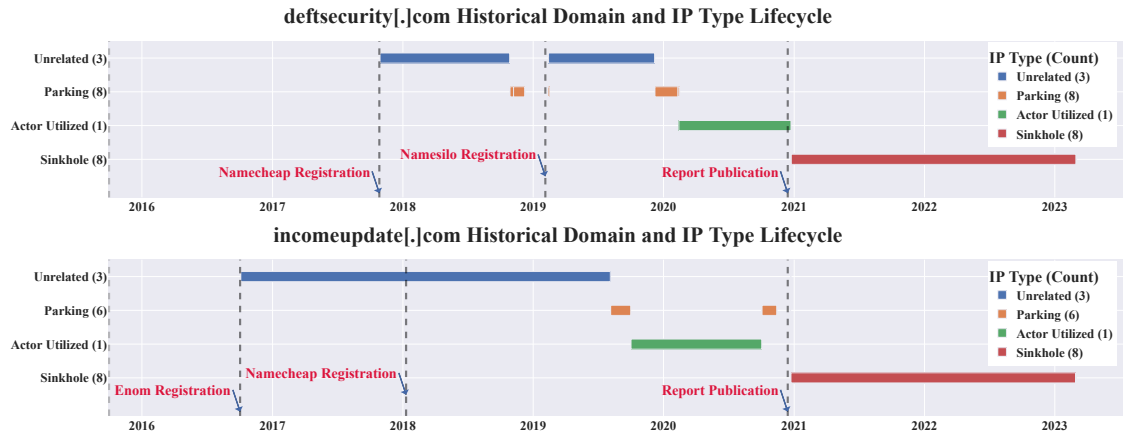


Figure 5.2: Domain and IP lifecycle of deftsecurity[.]com and incomeupdate[.]com sunburst domain names initially reported in [2]. In this work, we seek to automatically identify the actor-utilized IPs (colored in green). The numbers inside the parentheses reflect the number of unique IPs of each category.

5.2.1 Challenges

We focus on the challenges related to identifying the network infrastructure provisioned for APT domains. There are two major classes of such domains: newly registered domains and re-registered domains (previously owned). It is fairly typical that after registration, a domain could point to its registrar’s default parking infrastructure for a period of hours, days, or months [156]. The actors also may choose to *park* the domain at a benign IP outside their control (e.g., an IP with a positive Internet reputation) for “aging” reasons and to establish network reputation, since newly registered domains with no network history are often more suspicious than long-lived ones [142]. Other actors may choose to *point and periodically move* the domain to arbitrary infrastructure in order to inject deliberate noise in passive and active DNS repositories. In either case, however, when a domain name

is effectively used in an attack (i.e., to deliver exploits, as a social engineering domain, command and control, or exfiltration point), it must point to the actor-utilized infrastructure for the attack to be effective.

After detection or disclosure, an APT domain could be taken down, sinkholed, or left to expire until it is re-registered by some other legitimate or malicious entity [157]. While some or all of these events may occur, there is no definitive lifecycle of an APT domain. As a forensic investigator tries to piece together the timeline of the APT attack, we can assume that they will uncover a combination of the actor-utilized hosting infrastructure used in the attack, infrastructure belonging to previous owners of the domain, parking infrastructure, sinkholes, and even deliberate noise added by the actors.

Considering the aforementioned challenges, it is clear that extracting the actor-utilized hosting infrastructure of a domain name is no easy task. To make matters worse, it is also particularly hard to collect a clean and complete picture of the infrastructure to conduct further analyses, as APT domain names are usually utilized by nation-states and high-profile adversaries that do not — obviously — share any information about their operations. This work aims to address the problem of discovering **unknown actor-utilized** APT attack infrastructure in a rigorous and automated fashion by temporally bounding the active period of the APT attack.

Next, we summarize four main challenges that network forensic investigators face as they analyze APT attacks. Across all these challenges, we use the term “domain lifetime” to reflect all observable time periods in which a domain name existed.

Non-Actor Ownership. This is the period in the domain lifetime in which an actor does not own a domain due to it belonging to a different owner. As illustrated in Figure 5.1, this can occur either before the actor registers the domain, or after the domain expires or is taken down.

Sinkholing. This is the period in which the domain points to sinkhole infrastructure. Sinkholing occurs after domain detection and results in a malicious domain pointing to

infrastructure controlled by security companies and professionals, registrars, and law enforcement agencies [157].

Various Forms of Parking. This is the period in which the domain points to various forms of parking infrastructure. This infrastructure may be placeholder registrar-controlled parking infrastructure shortly after the domain name is registered, or parking infrastructure where the actors can point their domain names to age them until they use them for their operation.

Deliberate Noise Injection. Actors can point their domain to infrastructure that is not under their control to gain a positive (benign) reputation before they use the domain in an attack. Such an action could easily inject noise in passive and active DNS repositories, effectively making the network forensic investigation of the APT attack significantly harder.

5.2.2 Placing The Challenges In Context: The SolarWinds Attack

Next, we put the four major classes of challenges into perspective using the attack against SolarWinds as an example. Figure 5.2 showcases the historical life-cycles of two domains used in the *SolarWinds* attack, `deftsecurity[.]com` and `incomeupdate[.]com`. It also depicts the corresponding IP infrastructure that can be discovered from publicly available Active DNS datasets [32]. For this example, these IP addresses were labeled taking into account the public reporting of the attack [2, 158] and manual threat analysis from multiple analysts. Although the two domains were used in the same attack, their lifecycles differ in registrar, IP infrastructure, and pre-registration activity. These differences alone make the analysis of the two APT domains used in the same attack and controlled by the same actor non-trivial.

Starting with the domain registration patterns, APT actors re-registered domains from two different registrars that had completely different hosting history (possibly because of their previous domain owners). By just utilizing WHOIS data and manually trying to pinpoint the most likely window of actor registration, a forensic analyst would not be able to

identify which IPs were the ones utilized by the adversaries and used in the attack. That is because there are multiple parking and other unknown IPs in the historical DNS data — even in the period where the actors likely re-registered these two domains. Thus, to identify the actor-utilized IPs, an analyst would need more than just a temporal window of interest.

One solution to filter out the non-actor-utilized IPs would be to use publicly available lists of various parking IPs and DNS nameserver infrastructure [35]. By doing so, utilizing the IP and DNS nameserver data from [156], manually inspecting the DNS nameservers, and identifying various parking infrastructure, we could only additionally remove a subset of the publicly known parking infrastructure (colored in orange).

While this methodology has reduced the amount of infrastructure to inspect, it is still not sufficient, as the domain names have been pointing to cloud infrastructure (Amazon and Unified Layer, colored in blue after their latest registration in Figure 5.2) which has not been attributed to the SolarWinds attack due to its big temporal distance (many months before the attack took place). An analyst, knowing the timeline of the attack and manually inspecting the properties of this unknown cloud infrastructure, would filter out these IPs as likely parking and inactive infrastructure and yield only the actor-utilized IPs as they have been publicly reported [159, 2]. The practice of registering domain names years before their utilization and strategically aging them to infrastructure other than the attack infrastructure has been documented in prior reports [141, 140]. Evidently, filtering out all of these non-attack-related IPs is a non-trivial and labor-heavy process, often left to expert analysts. In this work, we seek to automatically identify the actor-utilized IPs of historical APT domain names in a transparent way and with a low false positive rate.

5.2.3 Observations and Takeaways

By taking into consideration the challenges and the lessons learned from the SolarWinds campaign, we arrive at the following three observations: first, APT domains feature unique lifecycles that can differ even within the same campaign, second, these lifecycles can last

multiple years, and domain registrations, and third, in time they can be associated with a diverse set of infrastructure (e.g., parking, sinkhole, etc.) that is often not associated with the actor-utilized IPs. Thus, measuring the lifecycle of APT domains requires:

- A historical dataset that observes and logs the infrastructure changes in APT domains across the years.
- A methodology that filters and labels the IP infrastructure associated with the APT domains and considers the diverse infrastructure types we discussed.
- A methodology that is applicable on a per-domain basis.

To satisfy the aforementioned requirements we take the following steps: first, we utilize two historical DNS datasets that span a decade and capture the changes in DNS resolutions of 405 APT actors and second, we develop a novel system that filters and labels these historical DNS resolutions taking into account the diverse infrastructure we encountered on our case study and operates on a per domain basis with high accuracy. Next, we discuss the datasets and measurement methodology in more detail.

5.3 Datasets and Methodology

This section introduces the OSINT datasets (subsection 5.3.1) we use to study 405 APT groups as outlined by our visibility in Table 5.3. Then we proceed by diving deep into Atropos (subsection 5.3.3), its modules, and how these modules enable Atropos to reliably and accurately identify actor-utilized infrastructure.

5.3.1 OSINT Datasets

Threat Report Data. Threat reports have been highly utilized in prior works to gather IoC datasets related to APT threats [160]. In this work, we utilize two major threat report datasets to gather APT domain names. The first data source we employ is AlienVault

Open Threat Exchange (OTX) [161]. AlienVault OTX is a large open threat intelligence community that has released more than 19 million IoCs to date. In our study, we only consider APT IoCs that are vetted by AlienVault’s internal threat research team. These APT IoCs are extracted from threat reports that leading security vendors release and disclosures from reputable threat researchers. The second APT data source we employ is CyberMonitor [162]. CyberMonitor is an aggregation of popular APT threat reports and datasets such as APTnotes and others, that have been heavily used in former works [163, 164, 160]. We manually parse threat reports from this data source that were published between 2007 to June 2019 with the intent to extract four attributes: APT domains, APT IPs, publication date of the report, and name of the APT actor that is associated with the domain names and IPs. For both datasets, we filter out any report that does not involve a threat actor found in Malpedia [91]. Malpedia maintains an up-to-date collection of APT groups and sophisticated actors that reflects the APT threat actor grouping done by the MISP project [165] and is more comprehensive than that of MITRE [166]. Utilizing this classification, we map any alias of the same actor to the standard domain name used by Malpedia and remove reports containing multiple actors. The final dataset consists of 2,188 APT reports, larger than previous APT studies [20, 75].

Table 5.1 shows the top 10 publishers in terms of the IoCs we utilize in this study. It is important to note that most of our indicators come from reputable security vendors, and we do not consider IoCs that come without a published report in order to minimize potential noise in our APT datasets from unreliable sources as random users in the AlienVault community [167].

Historical Active DNS. We use historical DNS resource records for all of the APT domain names in our dataset that are provided to us by the Active DNS [32] project. The Active DNS project daily scans millions of domain names from over 1,100 gTLDs and has been utilized in many prior measurement works [9, 172, 6]. The historical DNS records span from January 2016 to January 2025 and include A, AAAA, NS, NX, MX, and SOA query

Table 5.1: Coverage of IOCs for the top 10 publishers in terms of reports. Overall, we utilize a total of 2,188 APT reports.

Publisher	No. of Reports	No. of APTs	No. of IOCs		
			Domain	e2LD	IP
Palo Alto Networks	133	81	3024	2738	706
Kaspersky Lab	126	81	2574	1803	392
Trend Micro	90	63	1386	960	441
ESET	80	51	679	607	459
FireEye	65	50	1373	1228	151
Symantec	63	50	972	903	241
Proofpoint	57	45	1091	622	105
Talos	52	38	1906	1680	211
SentinelOne	40	32	940	835	159
Tencent	37	25	308	262	41

Table 5.2: Datasets utilized in the study.

Dataset	Time Span	Number of Records
[168, 161] APT Domains	2013-04-13 to 2025-03-01	31,398
[169] Compromised Domains	2013-03-20 to 2025-04-08	132,210
[156, 157, 170] Parking and Sinkholes	2007-07-18 to 2024-03-13	85,509
[32] Historical Active DNS	2016-01-01 to 2025-01-31	119,959,784
[171] Historical Virus Total DNS	2013-04-01 to 2025-03-06	480,093
[168, 161] Threat Report IoCs and Data	2013-04-13 to 2025-03-01	144,239

responses.

Virus Total (VT). To complement the coverage of ActiveDNS, we use a premium Virus-Total API access to gather historical DNS resource records for all of our APT domain names. Additionally, we query Virus Total to gather other domain and IP-related data in order to generate features for Atropos’ feature extraction module, which we detail in subsection 5.3.3.

Parking and Sinkholes. We utilize parking IPs and DNS nameservers from both an academic publication and Maltrail [156, 170], as well as manually labeling the DNS name-

Table 5.3: A and AAAA resource record visibility after enriching the known APT domain names with our DNS data sources. APT IPs appearing on threat reports can only characterize 23.52% of APT FQDNs in popular DNS datasets.

Visibility Metrics	Threat Reports	Active DNS and VT	Report IPs Matched on Active DNS and VT
Timespan	2013-04 2025-03	2013-04 2025-03	2013-04 2025-03
FQDNs	31,398	28,524 (90.84%)	7,386 (23.52%)
E2lds	22,691	20,975 (92.43%)	5,392 (23.76%)
APT Actors	413	405 (98.06%)	278 (67.31%)
RRs	N/A	1,004,614	51,891

servers of APT records to identify parking ones. Additionally, we utilize sinkhole IPs and DNS nameservers from an academic publication and a public list [157, 170], as well as manually labeling the DNS name servers of APT records to identify sinkholes.

Compromised Domain List. To filter out compromised domain names, we remove the APT domains that were mentioned to be compromised in the reports they were published in from the CyberMonitor [162] source. Additionally, we also filter out compromised domain names based on an aggregation of compromised domain lists [169], which includes various reputable sources such as abuse.ch and SANS.

5.3.2 DNS Datasets and Threat Reports IP Visibility

So far we have demonstrated the challenges that analysts face when studying domain lifecycles and introduced the datasets we will utilize in the study. Since we are mainly interested in identifying actor-utilized IPs in order to study the infrastructure they utilize, we can just gather the high-confidence domain names and IPs that appear in our 2,188 APT reports and utilize our DNS datasets to match them. This way, we will only utilize known domains and IPs that threat analysts in reputable reports have identified. Table 5.3 presents the visibility of our DNS data sources on the reports FQDNs, E2lds(i.e., the registrable portion of a do-

main name), APT actors, and resource records (RRs) after removing NX records and bogon IPs [173] (e.g., unroutable, private, loopback networks). As we can see, both these DNS sources together can provide at least one IP for 90.84% of the APT FQDNs and 98.06% of the APT actors, showcasing that our DNS datasets have a significant IP coverage on the APT domain names.

With this DNS visibility, we can now match the APT domains and IPs that get shared on APT reports and see what percentage of domain names threat reports can characterize with an IP. When we do so, we can see that only 23.52% of the FQDNs and 67.31% of the APT actors can be characterized as demonstrated in Table 5.3. Clearly, if we just utilize the reported domain and IPs, we can only characterize less than a quarter of the historical APT domains, even with DNS data sources that have over 90% APT domain coverage. Evidently, there are legitimate reasons why the IoCs shared on APT reports are not comprehensive. For example, report authors may not have IP-level visibility of the domain names they have identified during their analyses or may choose not to share all the IPs they have identified. For example, the IPs that the APT actors utilized may belong to virtual hosting or cloudfronting IPs and serve both benign and malicious domains at the same time, and report authors want to avoid readers blacklisting such IPs and causing harm. Additionally, threat report authors may lack the historical DNS datasets to identify the active IPs of the operation. We find that the percentage of APT reports that both share domains and IPs is only 44.22% of all reports that share at least one domain. The size imbalance has also been demonstrated in prior work [28]. Despite all the legitimate reasons that inhibit threat report authors from identifying or sharing comprehensively the IP addresses related to the domain name IoCs, answering the research questions regarding the longevity of APT network infrastructure cannot be done thoroughly just by utilizing threat report information.

Takeaways: Simply matching known APT domains to known APT IPs from APT reports using popular DNS data can only characterize **23.52%** of the APT domains. We find that only **44.22%** of the APT reports sharing domains also share IPs, which further

substantiates the coverage concerns of threat intelligence that have been raised in prior works [64, 152].

5.3.3 Measurement Methodology

Considering the lack of comprehensive OSINT visibility in domain-to-IP mappings and infrastructure coverage, we need to develop a measurement methodology in order to expand the APT infrastructure coverage and conduct a representative measurement study. However, as we have described in section 5.2, identifying the actor-utilized IPs on an APT domain is challenging. Previous works have tried to address similar problems [34, 35, 33], but they suffer from specific shortcomings, which we demonstrated with the Solarwinds case study subsection 5.2.2, as they are largely not applicable to address all the challenges we have showcased in a historical timeframe. To address these shortcomings and characterize more domains than those that simply APT reports can, we develop a supervised model that we call Atropos. Atropos automatically filters historical DNS records and identifies actor-utilized IP addresses of known APT domains. More specifically, Atropos ingests domain-to-IP mappings (i.e., DNS Resource Records — RRs) from DNS data, only for domains that appear in APT reports, and identifies which RRs correspond to infrastructure likely used by the APT actors. Atropos uses a combination of different OSINT datasets and three inline analytical modules, which are described below.

Enrichment and Filtering Module

To discover new APT infrastructure, we need to use a dataset that provides both historical and wide visibility among many different actors across different geographies. The initial step of our methodology is to enrich the APT domains in OSINT reports with historical DNS data [32, 155]. As different APT domains have been utilized in different time periods and are often registered by multiple Internet users (e.g., in the event of re-registration), this module aims to gather all the related historical IP infrastructure of the APT domains in our

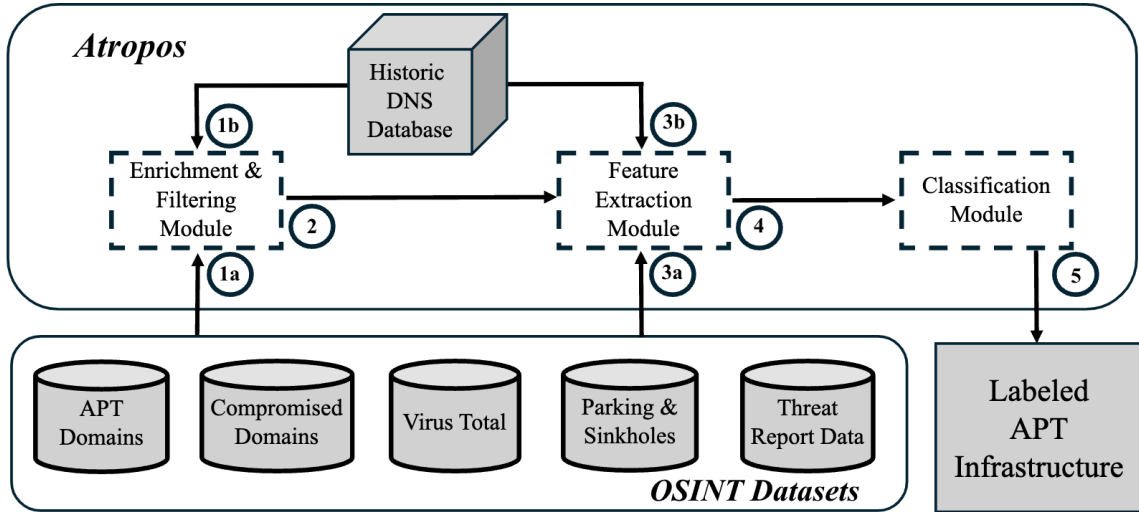


Figure 5.3: An overview of Atropos. Atropos utilizes OSINT datasets and historical DNS data to label and filter APT infrastructure in a 3-step process.

datasets.

In this module, Atropos will filter out all known compromised domains as well as bogon IPs [173] and non-existent domains (NXDOMAIN). This filtering is necessary as these records will not be related to infrastructure provisioned by the APT actors for an attack campaign. Additionally, we filter all domain names related to DNS fast-fluxing. Fast-fluxing is the process that involves the frequent change of the RRs of a domain name to many different IPs that can span hundreds or even thousands [174]. We consider such domain names out-of-scope of Atropos, as during development we found that the tactics and techniques of some fast-fluxing actors make them feature different lifecycles than those of typical APT domains, as discussed in section 5.2 and thus require dedicated models. For example, the gamaredon group has been demonstrated to keep utilizing the same detected and reported domain names, long after the reporting of its attacks, thus having malicious activity after its detection [175, 176]. We leave the development of dedicated systems for such lifecycles to future work. To remove such domains from our dataset, inspired by [174], we count the number of distinct IPs per domain and filter out the top 5% of the domain names in our dataset. This methodology filters out 1,497 domains with 544 IPs on

Table 5.4: The features of Atropos. Atropos utilizes 22 features from four distinct classes.

#f	Feature	Class	#f	Feature	Class
f_1	Detection and IP Fseen Delta	Temporal	f_{12}	IP Reputation	OSINT
f_2	Detection and IP Lseen Delta	Temporal	f_{13}	Number of Malicious Votes	OSINT
f_3	IP Lifetime	Temporal	f_{14}	Number of Harmless Votes	OSINT
f_4	Number of Historic Domains on IP	Infra.	f_{15}	Number of Malicious Analyses	OSINT
f_5	Mean Concurrent Domains on IP	Infra.	f_{16}	Number of Suspicious Analyses	OSINT
f_6	Median Concurrent Domains on IP	Infra.	f_{17}	Number of Undetected Analyses	OSINT
f_7	Number of IP Communicating Files	Infra.	f_{18}	Number of Harmless Analyses	OSINT
f_8	IP is Known Parking	OSINT	f_{19}	Num. of Domain Communicating Files	Domain
f_9	Nameserver is Known Parking	OSINT	f_{20}	Num. of Files Downloaded From Domain	Domain
f_{10}	IP is Known Sinkhole	OSINT	f_{21}	Number of Domain Subdomains	Domain
f_{11}	Nameserver is Known Sinkhole	OSINT	f_{22}	Number of Domain Certificates	Domain

average per domain, with 82.69% of the domain names belonging to the gamaredon group, which, as we have described, is known for fast fluxing activities [175, 176]. Since only 5% of the domains have been filtered, we do not consider the impact of this filtering significant for the generalization of our methodology, as we will showcase later.

Table 5.3 illustrates that the system’s visibility in the APT Fully Qualified Domain Names (FQDNs), APT effective Second Level Domains (e2LDs) [177], number of APT actors, and RRs that Atropos gathers after the DNS filtering and enrichment is significant. After all enrichment and filtering, our DNS visibility spans over a decade, with at least one resource record for 26,615 (84.76%) of the FQDNs and 402 (97.33%) of the APT actors in all the reports published between April 2013 and March 2025.

Feature Extraction Module

The next step of our methodology is to extract the features needed in order to train our models. Table 5.4 illustrates the features of Atropos. We utilize a total of 22 features from four classes, namely temporal, infrastructure, OSINT, and domain name features. We pick our features based on historical forensic experience and argue about their utility. We outline the four main classes of our features as well as the intuitions behind them.

Temporal Class (3 features):

- (f_1) **Domain Detection and IP First Seen Date Delta:** The time delta (in days)

between the first day the domain name was reported in a threat report and the first day that the domain first pointed to the IP. This feature aims to identify the IPs close to the detection of the domain that are more likely to be associated with the actor and remove older or newer IPs that are likely associated with previous or future owners of the domain name.

- **(f_2) Domain Detection and IP Last Seen Date Delta:** The time delta (in days) between the first day the domain name was reported in a threat report and the last day the domain first pointed to the IP. Since the disclosure of the APT domains to the public does not always happen right after their detection or sinkholing, this feature is meant to identify sinkhole and parking infrastructure that an APT domain has been pointed to before its detection and persisted months or even years after its public disclosure.
- **(f_3) IP Lifetime:** The number of days that the domain pointed to the IP address. This feature can help differentiate between short-lived placeholder and testing IPs and longer-lived APT-controlled IPs and parking. For example, in Figure 5.2, the domains pointed to placeholder parking IPs [156] for a median of 41 days compared to 314 and 366 days for the actor-controlled IPs.

Infrastructure Class (4 features):

- **(f_4) Number of Historical Domains Pointed to the IP:** The total number of historical domains ever pointed to the given IP according to the DNS data source. Similar to our example, this feature is meant to find parking and sinkhole IPs.
- **(f_5, f_6) Mean/Median of Concurrent Domains Pointed to the IP:** The mean and median number of other domains pointed to the IP during the period that the given domain is pointed to the IP. Since IP addresses are volatile over time, these features are meant to capture the infrastructure behavior of a given IP only at the time when the domain was pointed to it.

- **(f_7) Number of Historical Files communicating with the IP:** The total number of historical files that have been communicating with the given IP according to VirusTotal. In our example (Figure 5.2), sinkhole IPs have a median number of 83,128 communicating files on VirusTotal compared to a median of zero for the APT-controlled IP and parking infrastructure. This usually happens because malware dynamic execution will occur after a domain has been sinkholed and VirusTotal will only see the sinkhole IP.

OSINT Class (11 features):

- **Parking Features.** (f_8) Known Parking IP: Whether the IP appears on known parking lists. (f_9) Known Parking Nameserver IP Overlap: Whether the domain is served by a known parking nameserver at the same time as the domain points to the IP for at least 70% of the time (a percentage we manually pick after multiple tests).
- **Sinkhole Features.** (f_{10}) Known Sinkhole IP: Whether the IP appears on known sinkhole lists. In our example, this time period is illustrated by the red-colored infrastructure (Figure 5.2). (f_{11}) Known Sinkhole Nameserver IP Overlap: Whether a known sinkhole nameserver is serving a domain at the same time as the domain points to the IP for at least 70% of the time.
- **IP Reputation:** These features (f_{12} : IP Reputation, f_{13} : IP Votes Malicious and f_{14} : IP Votes Harmless) take into account the publicly known reputation of an IP based on the votes from the VirusTotal community [155]. Despite these scores not being perfect, they do help in some instances to identify benign IPs that malware actors can point their domain names to gain residual trust.
- **IP Analyses:** These features (f_{15} : IP Analyses Malicious, f_{16} : IP Analyses Suspicious, f_{17} : IP Analyses Undetected, and f_{18} : IP Analyses Harmless), compute the number of URL scanners in VirusTotal that have flagged an IP with the given label.

Domain Name Class (4 features).

- (f_{19}) **Number of Communicating Files:** The number of files that VirusTotal has found to have communicated with the domain.
- (f_{20}) **Number of Downloaded Files:** The number of files that were available to be downloaded by the given domain name according to Virus Total.
- (f_{21}) **Number of Subdomains:** The number of subdomains that were seen according to VirusTotal under the given domain name.
- (f_{22}) **Number of Certificates:** The number of SSL certificates that have been associated with the domain name at some point in time according to VirusTotal.

Classification Module

The final step in our methodology is to feed the feature vectors to our model. The classification module consists of a binary classifier that ingests the 22 features we have described and classifies each resource record as actor-utilized (True) or non-actor-utilized (False). During the development of Atropos, we experimented with various machine learning methods, including heuristics, Decision Trees, Support Vector Machines, Random Forests, XGBOOST [178], and Multi-Layer Perceptrons. During our experimental analysis, while other models had a great performance, we found the Random Forest classifier to offer the best ROC AUC performance across datasets while offering decision interpretability; thus, we picked this model over the rest. During its development, we trained and fine-tuned the hyperparameters only utilizing our training dataset – to prevent data snooping [179] – described in subsection 5.4.1, optimizing for ROC AUC with grid search. To showcase generalization, we tested Atropos on two out-of-distribution datasets. Finally, to demonstrate transferability across different DNS datasets, we train and test Atropos utilizing different models on each DNS dataset (ActiveDNS and VirusTotal) and show that accuracy is similar.

5.4 Evaluation

In this section, we discuss the training and performance evaluation of Atropos. Atropos is trained and fine-tuned on a training dataset based on the public knowledge of public threat reports, which we call the Public Reports Dataset (**PR**). After Atropos is trained and fine-tuned, it is tested on two different test datasets that were not considered during development, with the aim of evaluating our methodology against potential sampling bias and overfitting. Atropos achieves 10-fold cross-validation accuracy scores of 98.16% and 98.90% on Active DNS and VirusTotal DNS datasets, respectively, demonstrating transferability, and accuracy scores of 91.39% and 96.00% when evaluated on the test datasets (**EA**) and (**FR**), respectively, demonstrating generalization.

5.4.1 Training and Evaluation Datasets

Collecting ground truth regarding the infrastructure of APT actors is very challenging. Two of the main reasons that contribute to this are that APT actors will not share their attack playbooks with the public and the fact that APT attacks are, by definition, sophisticated. Thus, in order to create our training and evaluation datasets, we take two steps. First, we utilize the public knowledge of domains and IPs existing in public threat reports, and second, we utilize three analysts for manual labeling. These analysts consist of two PhD students with seven and four years of experience in APT network forensics (*JA1* and *JA2* respectively), and one senior APT network analyst with over 20 years of experience (*SA*). The instructions given to the analysts were the following:

- You are given DNS resource records (RRs) of historical APT domains.
- Your task is to label these RRs as actor-utilized (True) or non-actor-utilized (False).
- An RR is actor-utilized when the IP corresponding to the domain is the infrastructure utilized in the APT operation.

- You can utilize any tool at your disposal to do so.
- Deliver a file with every RR you can confidently label.

Aside from the RRs, the analysts are also provided with open Internet access along with all the features generated, and they are allowed to perform any tasks to validate the correctness of their decision (e.g., reverse IP lookups, searching IPs in IP intelligence and other reports, etc.). Next, we provide more details regarding each labeled dataset.

(Training) Public Reports Dataset (PR)

This set incorporates the public knowledge from APT reports. As APT actors will not share their infrastructure with the public, the next most accurate set that can be utilized is that of report authors who have manually labeled the infrastructure and openly shared it in threat reports. For this dataset, we utilize all the APT domain to IP mappings (RRs) that have been publicly mentioned in the APT reports of our APT data sources described in subsection 5.3.1, and have been matched together in Active DNS. However, these records only represent the positive class (i.e., actor-utilized) of the ground truth. To generate the negative class (i.e., not actor-utilized), and avoid class imbalance [179], we pick an equal amount of other random resource records from Active DNS, for the same domains that have a positive class record, and give all these records for manual labeling to analyst *JAI*. Analyst *JAI* confidently labels 1,915 out of 2,027 RRs and marks 1,065 RRs as actor-utilized and 851 RRs as non-actor-utilized. While the class distribution is not equal, the final dataset does not suffer from class imbalance [179], with 55.61% actor-utilized RRs and 44.43% non-actor-utilized RRs. Overall, this dataset consists of 1,915 resource records from 938 domains of 94 APT actors, from threat reports spanning from 2014-02-11 to 2023-04-13. To further evaluate *JAI* records for label inaccuracies [179], after *JAI* has completed the manual labeling, we give the same set of records and instructions to another junior analyst *JA2* from the same organization as *JAI* for labeling. After their inspection,

we quantify the level of agreement between the two analysts by computing the Cohen’s kappa [180] for the records they both successfully labeled. We find a Cohen’s kappa score of 0.9820, suggesting almost perfect agreement, thus giving us confidence that the *PR* dataset has a very high level of agreement among analysts.

(Evaluation) Senior Expert Analyst Dataset (EA)

Despite that the *PR* incorporates the public reports’ APT infrastructure labels and the fact that the two analysts reached a high confidence agreement level in manually labeling it, sampling bias could still be apparent [179]. To better understand the potential sampling bias of the *PR* dataset that will be used for training, we ask an expert analyst with over 20 years of experience, from a separate organization of *JAI* and *JA2*, to manually label a second completely disjoint ground truth from that of *PR*. This set consists of all the RRs found in Active DNS for one random domain name per APT actor, totaling 2,293 RRs. *SA* was able to confidently label 831 from the 2,293 RRs and marked 155 RRs as actor-utilized and 683 RRs as non-actor-utilized. The dataset *SA* labeled is not as balanced as *PR*, since *SA* was given all the historical RRs for each domain name and not a balanced set of RRs, in contrast to *JAI*. We do utilize this dataset — since the *PR* dataset is balanced — to evaluate Atropos in a scenario without base rate fallacy [179]. Overall, this dataset consists of 831 RRs from 191 domain names of 191 different APT actors.

(Evaluation) Future Records Dataset (FR)

The second test set is created after the system is completed with the intent to evaluate its performance against future distributions of RRs that were not seen during training. To do that, we pick a random sample of 100 RRs from reports spanning from 2023-05-03 to 2025-01-29, which were published after all of the reports from our training dataset. These 100 RRs correspond to 65 domains, 98 IPs, and 33 APT actors. Given the same instructions and data that were outlined in subsection 5.4.1, analysts *JAI* and *JA2* label these 100 RRs

and resolve their disagreements to arrive at a single dataset. The class distribution of this set is 73 non-actor-utilized and 27 actor-utilized RRs.

Table 5.5: Average 10-fold cross-validation performance of Atropos on the PR dataset. Atropos achieves at best a 99.86 ROC AUC score when utilizing Virus Total DNS data and training on the *PR* dataset utilizing a Random Forest Model.

DNS Dataset	ML Model	Average 10-fold X Validation Scores				
		ROC AUC	F1-Macro	Accuracy	Precision	Recall
Active DNS	Random Forest	99.82%	98.14%	98.16%	98.03%	98.60%
Active DNS	Decision Tree	97.67%	97.71%	97.72%	97.74%	98.11%
Active DNS	XGBOOST	99.52%	98.36%	98.37%	98.07%	99.00%
Active DNS	SVM	97.86%	88.20%	88.70%	82.88%	100.0%
Active DNS	MLP	96.83%	94.33%	94.37%	96.50%	93.05%
Virus Total	Random Forest	99.86%	98.86%	98.90%	98.37%	99.77%
Virus Total	Decision Tree	98.14%	98.39%	98.44%	97.62%	99.77%
Virus Total	XGBOOST	99.86%	98.66%	98.70%	98.35%	99.44%
Virus Total	SVM	97.40%	81.37%	83.44%	78.38%	99.33%
Virus Total	MLP	97.00%	95.54%	95.56%	96.30%	96.31%

5.4.2 Experimental Results

Classification Results

Table 5.5 shows the average 10-fold cross-validation performance of Atropos on the *PR* training dataset. Atropos achieves significant ROC AUC scores across all utilized machine learning models and the two DNS datasets. The best-performing model in terms of ROC AUC score is Random Forest with a score of 99.82% and 99.86% on Active DNS and VirusTotal datasets, respectively. This showcases that Atropos has high performance across models and can have high levels of transferability across different DNS datasets during our evaluation when training. Since Random Forest has the highest performing scores, we pick this model as best for our next test, out-of-distribution experiments.

Our second experiment evaluates Atropos against two test sets (**EA** and **FR**) that consist of records that were not considered during training with the intent to test Atropos performance against out-of-distribution(OOD) datasets and observe its generalization and robustness against sampling bias that has been identified as a major problem in the security

field [179]. Table 5.6 demonstrates Atropos’ performance against these two test sets and across the two DNS datasets. We observe that in all of the tests, Atropos remains highly accurate with accuracy equal to and higher than 91.00%. We also notice that the precision of Atropos drops especially in the **EA** dataset. We investigate these records and find out that the largest class of false positives comes from Cloudflare and Namecheap web-hosting IPs (35.71%), while the rest are distributed among different ASes. After debriefing the **EA** analyst, they mentioned that they do not consider any cloud-fronting and virtual-hosting IP addresses (e.g., Cloudflare, Namecheap virtual-hosting) as operation-related, as they do not provide any basis for pivoting or evidence that the actors owned the IPs, as they can belong to multiple users. Despite that **EA** analyst is correct and these IPs are not useful for pivoting and should not be considered for blacklisting, this comes in contrast with our instructions in which we outlined we wanted to identify the *IP corresponding to the domain is the infrastructure utilized in the APT operation*, regardless of whether they are cloud-fronting or virtual hosting. Despite that, the overall performance of Atropos across all tests remains very high, and this experiment showcased that its results are generalizable in (OOD) datasets. In the appendix, we demonstrate how Atropos can be adjusted to generalize in similar scenarios of labeling as those considered by the **EA** analyst.

Table 5.6: Out-of-distribution test set evaluation of Atropos. Atropos achieves an over 91.00% accuracy across the two evaluation datasets, demonstrating generalization.

DNS Dataset	Test Set	ROC AUC	F1-Macro	Accuracy	Precision	Recall
Active DNS	FR	95.47%	95.08%	95.38%	92.00%	95.38%
Virus Total	FR	95.47%	95.08%	95.38%	92.00%	95.38%
Active DNS	EA	87.13%	85.56%	91.00%	73.23%	91.00%
Virus Total	EA	88.47%	87.20%	91.39%	76.53%	91.39%

Feature Importance

By calculating the Mean Decrease of Impurity (MDI) score on an 80-20% split utilizing the *PR* dataset and Active DNS data, we rank the features that Atropos has used to find out their

importance, thus offering model interpretability. Table Table 5.7 presents these results. We observe that the top five features by MDI include the number of Historic Domains on IP (0.187), the IP first seen Delta (0.177), the number of Communicating Files on IP (0.158), and the Mean and Median Concurrent Domains on IP (0.125 and 0.149), thus highlighting importance across all feature types but specifically in infrastructure and temporal features. This fact aligns with our observations from the SolarWinds case study in subsection 5.2.2, where both the temporal (i.e., when an IP was pointed to the domain name compared to detection) and infrastructural (i.e., what kind of infrastructure that IP is), are necessary to distinguish the actor-utilized from the non-actor utilized infrastructure. Considering this, we are confident that Atropos makes decisions that follow the principles that a human analyst would also have used.

The strong performance of the top three features can be attributed to their capability to identify parking and sinkhole infrastructure. This reflects on the motivating example of *SolarWinds* we showcased in Section subsection 5.2.2, where the actor-controlled IPs had only the SolarWinds domain names pointed to them while parking and sinkhole IPs had more than nine million and 600 other domains pointed to them, respectively. The other benefit of these features is that Atropos does not only rely on parking and sinkhole IP and DNS name server lists, which are usually static and can take months or even years to be updated.

The second strongest set of features is the temporal features. This is not a surprise, because as we saw in *SolarWinds*, the APT-controlled IPs pointed to the domains a few months before the detection and continued to be the primary destination of the domains until very close to their detection. Atropos can pick up on this temporal aspect and penalize IPs of previous owners that were first seen on the domains very early and IPs of sinkholes that were first seen after the domain detection, similarly to Fig. Figure 5.2.

Table 5.7: Atropos MDI Feature Importance when trained on *PR* dataset and utilizing Active DNS data with an 80-20% split.

#f	Feature	MDI	#f	Feature	MDI
f_1	Detection and IP Fseen Delta	0.177	f_{12}	IP Reputation	0.012
f_2	Detection and IP Lseen Delta	0.050	f_{13}	# of Malicious Votes	0.038
f_3	IP Lifetime	0.007	f_{14}	# of Harmless Votes	0.009
f_4	# of Historic Domains on IP	0.187	f_{15}	# of Malicious Analyses	0.004
f_5	Mean Concurrent Domains on IP	0.125	f_{16}	# of Suspicious Analyses	0.005
f_6	Median Concurrent Domains on IP	0.149	f_{17}	# of Undetected Analyses	0.011
f_7	# of IP Communicating Files	0.158	f_{18}	# of Harmless Analyses	0.013
f_8	IP is Known Parking	0.018	f_{19}	# of Domain Communicating Files	0.001
f_9	Nameserver is Known Parking	0.019	f_{20}	# of Files Downloaded From Domain	0.003
f_{10}	IP is Known Sinkhole	0.009	f_{21}	# of Domain Subdomains	0.003
f_{11}	Nameserver is Known Sinkhole	0.000	f_{22}	# of Domain Certificates	0.002

Table 5.8: Number of network IoCs associated with the actors from the OSINT threat reports and identified by Atropos for the top 10 actors, and overall. Atropos provides three times the IP visibility of threat reports and contextualizes three times more domain names than threat reports.

Actor	IP addresses		BGP prefixes		ASN		Domain Coverage (%)	
	Reports	Atropos	Reports	Atropos	Reports	Atropos	Reports	Atropos
Lazarus Group	1,047	776	569	504	371	241	20.25%	76.12%
Gamaredon	361	1,873	130	623	25	253	20.90%	45.19%
Fin7	218	341	132	222	56	126	23.48%	60.75%
Unc1878	208	379	73	134	48	47	63.45%	96.49%
APT28	204	723	155	381	97	202	29.63%	71.89%
Muddywater	173	301	92	206	36	98	06.38%	43.20%
Winniti Group	158	206	85	126	47	74	23.64%	21.28%
APT29	157	144	123	130	93	91	28.35%	65.67%
Sandworm	132	53	93	20	70	13	21.42%	71.42%
CharmingKitten	128	554	62	188	62	81	46.84%	63.19%
Total	7,553	25,049	3,530	6,115	1,291	1,762	20.20%	61.07%

5.4.3 Infrastructure Expansion and Lifetime Characterization

Table 5.8 presents the number of IPs, BGP Prefixes, Autonomous System Numbers (ASNs), and domain coverage comparison between what is provided in threat reports and what is identified by Atropos. The table presents the coverage of the top 10 APT actors along with the total number of all the actors. Overall, Atropos provides 3.062 times more high-confidence IPs than OSINT APT reports. The added benefit for BGP prefixes and autonomous systems (ASes) is smaller as they represent bigger groupings of Internet infrastructure, but are still significant. Furthermore, Atropos provides actor-utilized IP mappings




Figure 5.4: Daily active APT actors for the top 15 most utilized provider ASes in the last decade. Cloudflare utilization for domain name hosting has increased drastically over the years, making forensic analysis and attribution of IP infrastructure harder.

for 61.07% of domains provided in APT reports, which is significantly larger than that of just matching the IPs that exist or reports with their domain names. To characterize the lifetime of IP infrastructure, we utilize the historical information provided by the Active DNS dataset [32], enabling us to build lifetimes with a daily granularity for all of the high-confidence IPs Atropos has identified associated with the APT domains in our dataset.

Takeaways: Our measurement methodology accurately expands the APT infrastructure identified by APT reports by **3.06** times and is able to characterize **61.07%** of the APT domains appearing in threat reports, thus enabling us to conduct a more comprehensive measurement study than just utilizing OSINT data.

5.5 Infrastructure Analysis

In this section, we conduct the largest APT and most comprehensive APT infrastructure analysis to date. To do so, we utilize all of the APT IoCs from our OSINT data sources described in subsection 5.3.1, as well as the new infrastructure Atropos has identified, utilizing both the ActiveDNS and VirusTotal DNS datasets. For more conservative estimations, we remove RRs on which ActiveDNS and VirusTotal models disagree. The number of these records is only 1.34% of the overall records and thus it does not bias our measurement results. We structure our analysis around the following research questions:

- Where do APT actors provision their infrastructure and do they re-use the same hosting providers over the years? (subsection 5.5.1)
- What is the lifecycle of the different infrastructure types associated with APT domains, and how does that affect forensic analysis? (subsection 5.5.2)
- How long before the public reporting of an attack are actor-utilized IPs provisioned to the domains, and what is the time window of their observability? (subsection 5.5.2)

5.5.1 Infrastructure Utilization

Hosting Provider Utilization

Figure 5.4 demonstrates the density of the daily active APT actors that utilize any of the top 15 hosting providers in our dataset. We observe that these hosting provider ASes that APTs utilize consist of a mix of cloud-fronting, CDN, and proxying providers (e.g., Cloudflare, Akamai, AWS, Google Cloud), virtual hosting providers (e.g., Vultr, DigitalOcean, OVH, Namecheap, UnifiedLayer), dedicated hosting (i.e., Hetzner, OVH, Hostkey, M247), and providers that are more tolerant to abuse (i.e., Colocrossing, Stark Industries, M247). Thus, APT actors utilize a plethora of different types of hosting providers for their domain name hosting and do not primarily choose a specific category of providers. Temporally, we observe that after 2023, CloudFlare has drastically increased in popularity among actors, with 74 different actors hosting at least one domain in their network. This increased popularity of Cloudflare over the years is well justified, as this provider offers very lucrative technologies that enhance the stealthiness of APT infrastructure, such as origin IP masking and blending with benign domain traffic behind the same virtual hosting IPs. This trend complicates network threat hunting and forensics as it diminishes the value of IP addresses for such domains, a fact that has been anecdotally verified by APT experts [22]. Another recent rising trend is that of the increased utilization of the bulletproof hosting provider




Figure 5.5: AS re-utilization among APT actors.

Stark Industries Solution after 2023. *Stark Industries Solution* is a new bulletproof hosting provider that was launched in February of 2022[181](although its IP space was used in previous attacks under different management). While we observe 23 different groups to have utilized this hosting provider since 2023, the two groups with the highest number of domains are *Fin7* and *MuddyWater* [182, 183]. Lastly, despite the aforementioned rising trends, several hosting providers (e.g., *Hetzner*, *VULTR*, *M247*) have featured a steady utilization by APT groups across the years.

Infrastructure Reuse

In order to measure the re-use of ASes among different APTs, we identify for each domain name and group the first time that domain was provisioned to each AS. Then we measure the proportion of ASes that get reused for more than one year per APT actor, and we focus this experiment only on the actors that we have, in the APT context, a significant number of domains, which is more than 20 domain names. This amounts to 135 APT actors. Figure 5.5 demonstrates the proportion of ASes that these actors reuse across the

years. We notice that most APT actors re-use a small portion of all the ASes they have provisioned their domain names historically, with the average re-use rate for all these actors being 26.20%. This means that most actors do not choose to host their domain names on the exact same set of ASes over the years; however, they do re-use a smaller portion of the same hosting ASes. When we look at the percentage of the APT groups that do re-use at least one AS for over one year, we see that it is 97.03%. Thus, APT actors do re-use network infrastructure in the same hosting providers; however, this re-use only accounts for a small portion of all the infrastructure they have used historically. Threat hunters and attribution experts need to be careful when identifying and attributing new campaigns to existing actors, mainly by network infrastructure signals, and will need to focus on the infrastructure that is consistently being reused when doing so.




Figure 5.6: Number of actor-utilized IPs per country for the top affiliated countries of the APT actors.

Infrastructure Geolocation

Another important insight that can help us characterize and compare the APT actors is the geolocation of their infrastructure. To that end, we map each actor-utilized IP address to the country where it is most likely located according to IPInfo [184] and then analyze the

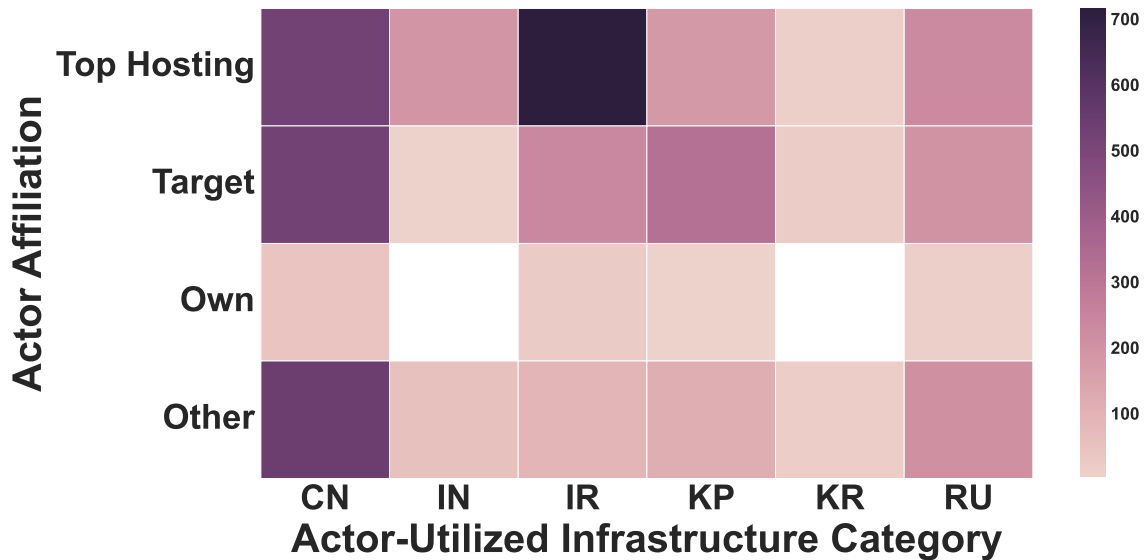


Figure 5.7: Number of IPs per category of infrastructure for the top affiliated countries of the APT actors.

correlation of the location of the infrastructure with the country affiliation of the actors.

Figure 5.6 shows a heatmap of the country an actor is affiliated with and the country where the actor-utilized infrastructure is provisioned. In the interest of space, the countries have been limited to the ones with the most publicity and references across our threat reports. We can observe that most of the actor-utilized IPs are provisioned in the USA, with other big hosting provider countries like Germany and the Netherlands to follow. Additionally, we can see that actors from different countries choose to utilize infrastructure with different patterns that, in some cases, overlap, like the Russian and Iranian APT actors. Their utilization of infrastructure among the US, the Netherlands, Germany, France, and the United Kingdom is more evident and different from that of Chinese actors, which, aside from their disproportionate use of US-based infrastructure, also utilize more infrastructure in Hong Kong, Japan, and South Korea.

These findings raise two interesting questions. First, whether the location of the actor-utilized infrastructure correlates with the location of the attack target. Second, whether the location of said infrastructure relates to countries with large hosting providers. To answer these questions, we utilize targeting data from the APT reports and match each

domain name and IP with the countries that were identified as targets in the same APT reports. We only use infrastructure for which targeting information is available in this part of our analysis. We also group together countries that are the top 10 largest hosting providers [185] to see if the infrastructure provisioning of the actors is correlated with those aspects.

In Figure 5.7, we see that APT actors from the top countries mostly provision their infrastructure either in countries that have large hosting providers or in the target countries. Chinese and North Korean actors deploy most of their infrastructure in their target countries, while Iranian and Indian groups mostly utilize countries that have large hosting providers. The Chinese actors provision infrastructure to countries labeled as "Other" which is mostly located in Hong Kong and Singapore. Finally, as expected, country-affiliated APT actors rarely provision infrastructure in their own country.

Takeaways: The infrastructure utilization analysis has demonstrated that APT actors utilize a plethora of different hosting providers, with trends changing over the years. The recent increase in cloud-fronting service utilization makes forensic analysis significantly more difficult and calls for adjustments in attribution and detection models. While the majority of APT actors re-use infrastructure, this re-utilization only occurs to a small portion of their overall infrastructure.

5.5.2 Infrastructure Lifecycle

Aside from the amplification of the known APT infrastructure that we saw in Table 5.8, Atropos also provides a plethora of new domain-to-APT IP mappings. These mappings allow us to measure the lifecycle of APT infrastructure more comprehensively using the domain names in the APT OSINT reports and the first day they were reported.

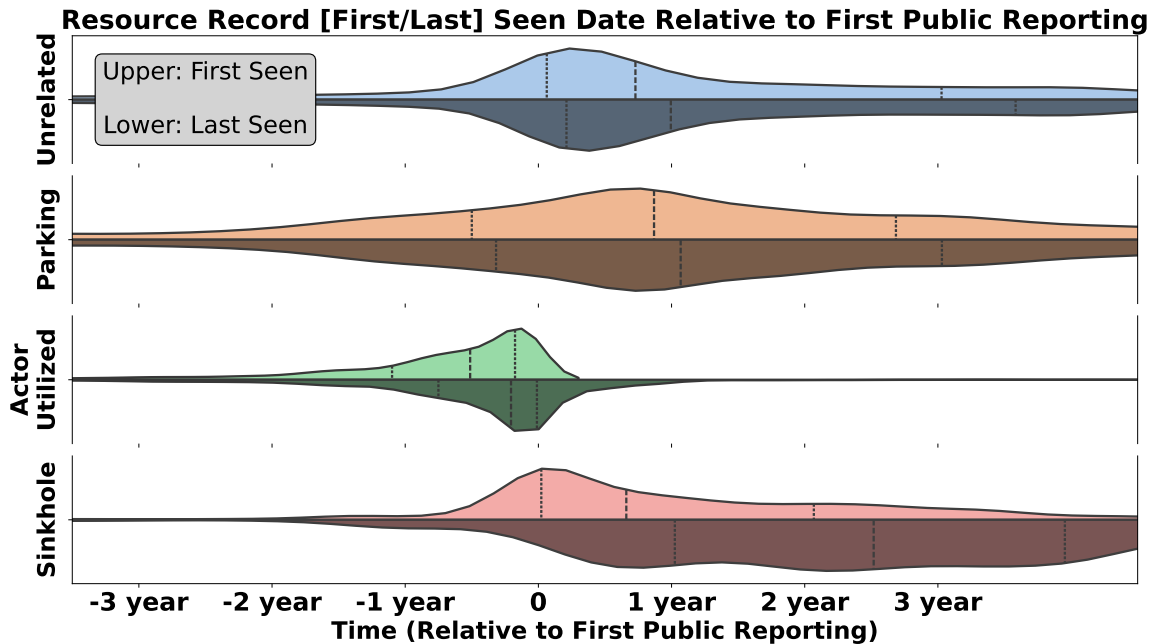


Figure 5.8: The unique lifecycle of the infrastructure associated with APT domains compared to the first public report date.

Infrastructure Type Analysis

As we have discussed in subsection 5.2.1, there are many types of infrastructure associated with APT domains that complicate forensic analysis. Figure 5.8 shows the lifecycle of all these types of infrastructure compared to the first public reporting of each of the domains they point at. We observe that the vast majority of actor-utilized infrastructure is mainly first seen on DNS records a few months before the first reporting date, with most of the IPs spanning back to within 2 years before. More interestingly, **73.6%** of the actor-utilized IPs no longer point to their domains at the detection date. This finding has practical applications for analysts and systems detecting and investigating APT infrastructure during and after the disclosure of an attack, considering the lack of comprehensive coverage that threat reports provide. *Analysts and systems that do not utilize historical DNS datasets and do not consider the lifetime of their IP infrastructure risk at best incomprehensively discovering attacker-utilized infrastructure or at worst, misclassifying parking, sinkhole, or other unrelated infrastructure that appears after detection as attacker-utilized.*

As expected, sinkhole infrastructure mainly starts being observed near the first reporting date and spans long after detection; however 17.5% of the sinkhole IPs appear before the domain first reporting, and thus, analysts investigating to find actor-utilized IPs before the domain detection have to consider them. Similarly, 35.9% of known parking IPs appear to be first pointed to the domains before publication, and as we showcased in the SolarWinds case study, they have to be considered and filtered out even before the domain reporting. Surprisingly, **31%** and **42%** of the parking and sinkhole IPs, respectively, have one or more malicious detections on Virus Total, and **18%** of the sinkhole IPs have five or more. This can be explained due to the large amount of APT and other malicious domains that end up being pointed at them, which makes some vendors flag them as malicious by association. Nevertheless, this fact highlights that researchers have to be careful and not blindly trust vendors' detections but consider the type of infrastructure when doing forensic analysis or building intrusion detection or investigation systems, especially considering that five or fewer VirusTotal malicious detections are common in malicious IP labeling [186, 187, 188, 189, 190].

Atropos filters out a lot of infrastructure that is either not actor-utilized, known sinkholes, or parking IPs. This infrastructure is primarily first pointed after the domain reporting for 75.9% of the IPs, and is mainly associated with unrelated actor-utilized infrastructure, such as future owners, parking, and sinkhole IPs unknown to the public. The top two IPs of this class are: "35[.]205[.]61[.]67", an unknown to our sinkhole list sinkhole [191], and "54[.]65[.]172[.]3", an Amazon shared hosting IP that had 995,067 domains historically pointed at it. Since these IPs are not directly utilized by the actors for the intent of their malicious operations, Atropos considers these IPs as unrelated.

Takeaways: The diverse type of IP infrastructure associated with APT domains features unique temporal lifecycles. **73.6%** of the actor-utilized IPs no longer point to their domains after their detection, highlighting the importance of historical data for comprehensive infrastructure tracking. Researchers and analysts have to be very careful not to misclassify




Figure 5.9: Number of days that actor-utilized IPs were first and last observed before their domain name public disclosure.

parking and sinkhole infrastructure as actor-utilized despite their malicious detections by vendors.

Actor-Utilized IP Activity

Figure 5.9 demonstrates the first and last seen of the APT utilized resource records among all actors as observed in ActiveDNS compared to their first public disclosure. We observe that there is a wide variation among initial provisioning delta compared to the first public threat reporting. The mean and median first IP provisioning is 317 and 187 days, respectively, before the first public disclosure, which indicates that many actor-utilized IPs remain well under the radar for months. This fact reinforces the common knowledge that APT attacks are stealthy, and it takes a significant amount of time to detect them in contrast to other cyber attacks like phishing or password stealers, which feature significantly shorter detection lifecycles of 21 hours and 11 days, respectively [10, 103]. It is important to note that this is a higher bound estimate as it includes the time for a report to be written and published; however, the difference is still significant relative to commodity threats, and the

reliance of expert APT analysts on the public reports has been recently verified [22]. The long delta between first infrastructure provisioning and public reporting of the attacks can also be explained by the fact that advanced actors have been reported to strategically age their domain names [141, 140]. Looking at the last time the actor-utilized IPs kept pointing to their domain names relative to public disclosure, we observe a mean and median time of 173 and 75 days, respectively. This fact reiterates the need for historical data in order to comprehensively track APT infrastructure, and the need for threat analysts to be careful of the type of infrastructure that gets pointed to the domains close to public release, as we have showcased that security vendors can sinkhole APT domains even before the public release of an attack.

These observations can aid network detection systems that are heavily dependent on features related to the short lifespan of malicious domain names, which have been proven not to be adversarially robust [142]. Furthermore, they have practical implications for organizations and government entities that need to forensically investigate APT attacks against them. Our results demonstrate that 90% of the actor-utilized IPs first and last get pointed to their domain names between 19 to 25 months before their public reporting. Thus, organizations that are sensitive to APT threats and network APT experts will need to keep at least 19 months' worth of historical network records to comprehensively evaluate whether they have been a target of a prior APT threat and to thoroughly investigate the network infrastructure of APT actors, respectively.

Takeaways: The lifecycle analysis indicates that APT actors first provision infrastructure on their domain names **317** days on average before the APT attack is publicly reported. This number alone provides ample time for actors to successfully conduct their operations while making detection systems that assign a positive reputation to longer-lived domains and infrastructure less effective. Organizations need to keep their network logs for at least (**19** to **25** months) to be able to identify 90% of the APT infrastructure from a DNS perspective.

5.6 Discussion

The recent increase in the utilization of cloud-fronting services like CloudFlare among APT actors makes forensics and attribution harder. For APT actors and malicious domains utilizing such infrastructure, future works could adjust DNS-based detection and attribution systems to work well beyond the infrastructure-level features of prominent works [143, 144] and emphasize lexical, registration, and temporal characteristics of domain names. Despite this trend, traditional dedicated hosting still remains prevalent, and systems can still capitalize on such features for detection and attribution. Although the focus of our work was to comprehensively measure the infrastructure utilized by actors, future works can utilize and adjust methodologies like Atropos to identify and focus on dedicated hosting and only on the infrastructure that actors choose to re-use for attribution and threat hunting, as we have demonstrated in Appendix Appendix A.

5.7 Summary

In this chapter, we analyzed the network infrastructure of 405 APT actors spanning over a decade. Utilizing our novel measurement methodology, we have expanded the IP visibility and contextualized the network infrastructure of three times more domains than that possible only with infrastructure from public threat reports. This infrastructure visibility enables us to conduct the largest APT infrastructure characterization study and pinpoint several practical findings. Our lifecycle analysis determines that organizations will need to retain network logs for at least 19 to 25 months in order to maintain comprehensive visibility in APT network infrastructure in the case of an attack. We observed that while APT actors utilize a plethora of different hosting providers, they only re-use a small portion of them over the years, and that use of cloud-fronting has increased significantly, making network forensics and attribution harder. Furthermore, we find that at the first public disclosure of APT attacks, the vast majority of their IP infrastructure no longer points to their domain

names, and considering the incomplete infrastructure sharing on public reports, this finding highlights the need for historical data retention and systems to increase the completeness of publicly known APT infrastructure. Our findings verify prior insights from experts and we hope to be the basis for increased attention from the community.

CHAPTER 6

CONCLUSION

6.1 Summary

The goal of this dissertation has been to shed light on the temporal aspects of network infrastructure utilization in cyber attacks. Each work presented in this dissertation has offered insights into the different stakeholders interacting with malicious network infrastructure from different and unique network vantage points.

Our first study presented a longitudinal study analyzing the network communication of 202 different malware families from the perspective of a popular authoritative DNS server. We observed billions of resolutions over four years at our authoritative collection point, enabling temporally complete and global visibility into malicious domain usage. *AuthDNS* simultaneously solidifies prior findings while also shedding new light on the epidemiology of malware. Our temporal analysis demonstrated that the vast majority of newly registered malicious domains are set up and detected quickly. Due to network noise from scanners and AV vendors, both the temporal and organizational properties of network clients should be considered when estimating malware infections from a network perspective. Finally, we introduced a brief taxonomy of malware measurement perspectives and discussed the advantages and disadvantages across four primary measurement goals.

The second study presented an empirical analysis of *Stealers* and shed light on the infrastructure and the lifecycle of the interactions of cybercriminals with it. We found that operators quickly provision their C2 infrastructure within 14 days after registration to their domain names and much of the *Stealers* infrastructure to be long undetected, with public blocklists detecting *Stealer* domains on average 74 days after initial domain registration, which gives operators plenty of time to infect more victims. *Stealers* operators

conduct their campaigns utilizing minimal hosting resources and abuse services such as free ccTLDs and cloud-fronting. Operators use proxy services ranging from traditional VPNs to mobile and residential proxies to Tor networks, and the mobile and residential proxies they utilize can cause misdirection when characterizing their profile; thus, law enforcement and security analysts have to be careful in their attributions. The diurnal analysis of the operators suggested that they administer their botnet as a full-time job. Last, we find that 69.03% of the operators stop utilizing their panels within 30 days of a detection event, suggesting that while they do not immediately abandon their operations after detection, the detection event is critical to curb their operations.

Our last study characterized the network infrastructure of sophisticated actors (i.e., APT and sophisticated cybercriminal groups). In this work, we proposed a novel measurement methodology that expanded the IP visibility and contextualized the network infrastructure of three times more domains than that possible only with infrastructure from public threat reports. This infrastructure visibility enabled us to conduct the largest APT infrastructure characterization study and pinpoint several practical findings. The lifecycle analysis determined that organizations will need to retain network logs for at least 19 to 25 months in order to maintain comprehensive visibility in APT network infrastructure in the case of an attack. We observed that while APT actors utilize a plethora of different hosting providers, they only re-use a small portion of them over the years, and that the use of cloud-fronting has increased significantly, making network forensics and attribution harder. Furthermore, we find that at the first public disclosure of APT attacks, the vast majority of their IP infrastructure no longer points to their domain names, and considering the incomplete infrastructure sharing on public reports, this finding highlights the need for historical data retention and systems to increase the completeness of publicly known APT infrastructure.

Next, we discuss the limitations of the studies and systems of this thesis and propose future avenues of research, considering the limitations but also the insights provided in this dissertation. Lastly, we offer the closing remarks.

6.2 Limitations

6.2.1 Malware Infrastructure Study Limitations

Recursive resolvers. The DNS protocol relies heavily on recursive resolvers, which operate between authoritative DNS servers and DNS clients. This indirection makes client estimation difficult. Although many public DNS resolvers support ECS [92, 192], lack of client support for ECS can lead to underestimation. Furthermore, if multiple infected hosts exist within the same ECS network block, *AuthDNS* cannot distinguish between them. Finally, authoritative DNS servers typically see only a portion of the DNS requests issued by individual hosts [193] due to caching by recursives. We do not utilize query volumes. Instead, we focus on the number of unique clients we observe querying for each domain in *AuthDNS* over the four years of our measurement. Due to the aforementioned limitations, our results, even with a global perspective, should be viewed as lower bounds on the overall malware ecosystem.

Noisy clients. Not all DNS lookups for malware-related domains come from the malware itself. Honeypots and network scanners may query DNS to detect malware-related infrastructure in several cases. This is a common challenge in prior malware ecosystem research that leads to overestimation [194, 17]. We do not address this limitation in section 3.3 and section 3.4, in order to perform meaningful comparisons with established alternative perspectives. However, we begin to address this challenge in section 3.5 by examining the different stages of the malware domain lifecycle and identifying likely scanners based on signals such as queries that consistently appear after new malware domains are reported/discovered on blocklists.

VPNs and proxies. Clients may utilize VPNs or proxies to hide their true network location. This can skew *AuthDNS*'s geolocation of infected populations. To approximate the

presence of proxies and anonymizing networks in our dataset, we measure the prevalence of Tor exit nodes in *AuthDNS* using historical Tor exit node lists[195], accounting for the days that each exit node is active. We find the average daily client percentage and average daily query volume percentage of Tor exit node IPs to be 0.07% and 0.001%, respectively; thus, their presence on our dataset is minimal. The low prevalence of anonymizing networks on our dataset does not guarantee the absence of other popular proxy and VPN providers. A lack of well-documented historical datasets for proxies/VPNs limits our ability to measure them more thoroughly.

Malware Visibility. The observations in our study are limited by the visibility of our datasets. More specifically, our visibility of malicious domains depends on the *MAL* dataset, which only includes Windows malware. Additionally, we intersect the malicious domains with those registered in our *AuthDNS* dataset, which removes an additional set of malicious domains. Despite these limitations, our study covers more than 200 malware families.

6.2.2 Password Stealers Study Limitations

The operational nature of the *Stealer* dataset can affect the accuracy of our results. The tracking pixel may only appear on some panel pages and therefore miss activities from operator devices. Additionally, since the data collection relies on running malware in a sandbox, the malware binary collection and analysis can create a skewed view of the malware families. However, since our dataset is large (hundreds of thousands of records), we can assume the data is statistically representative of the overall population.

The data validation analysis shows that operators may spoof their UA, use private browsing, or use multiple devices. It is difficult, if not impossible, to associate a virtual entity with a physical entity based on the current dataset. Nevertheless, we make conservative assumptions about the operators by framing the analysis as operator *devices* and extensively validating the dataset.

Another possible limitation is the effect of network address translation (NAT) traffic and aggregated pDNS data from recursive servers. These artifacts can impact our infection estimation and operator count. Additionally, operator network proxies can create ambiguities about the geographical regions of the operators.

6.2.3 APT Domain Study Limitations

Despite the increased infrastructure visibility that our measurement methodology (Atropos) provides compared to APT threat reports, it cannot identify all actor-utilized IPs for all domains, as illustrated in Table Table 5.8. Some of the APT domain names belong to ccTLDs and other TLDs that do not share their zone files, so it is difficult for DNS scanners to pick them up before their detection. APT actors may also set their name servers to respond with a valid command and control IP only to specific target networks (i.e., victims) and with invalid IPs to others, including projects like Active DNS. Additionally, some APT actors may utilize a subdomain that hasn't been observed by a DNS scanner (e.g., 3LD or 4LD) for their command and control server and park their e2LD to known parking locations, which Atropos will filter out. Despite all this, the infrastructure expansion compared to public reports for our measurement study is still significant.

As illustrated in Section subsection 5.4.2, Atropos performs very well in both evaluation datasets; however, its performance can vary by actor depending on how differently actors utilize the network infrastructure. APT actors can perform mimicry attacks or utilize fast flux [175] to induce false positives and perform *label shift* [179, 196]. Future work can build dedicated models for individual groups and their strategies to address such issues. We did not explore this avenue as the existing high-confidence ground truth for these threats is insufficient to effectively represent each APT actor in a machine-learning model without big class imbalances [179].

6.3 Future Work

While the previous subsection described the limitations of the studies and methodologies that are part of this dissertation, in this subsection, we discuss future works that could try to address these limitations, but also future works that are enabled by the innovations and insights of this thesis.

Accurate discovery of DNS scanners, security vendors, and victims in network datasets utilizing lifecycle features. The findings of chapter 3, demonstrated that while it is hard to assess the stakeholders interacting with malware-hosting domain names just by their network, due to the utilization of public recursives, VPNs and proxies, the temporal interactions of some stakeholders such as DNS scanners and security vendors start after the detection of a malicious domain name. Future works can focus on accurately identifying the different types of stakeholders interacting with malware infrastructure by taking advantage of lifecycle features, such as the first time they accessed the infrastructure, compared to the domain registration, detection, and takedown. Future works can also deploy distributed DNS and HTTP honeypots in order to identify DNS and HTTP scanners and their attributes, and filter them out when assessing victim targeting on network data.

Improving the performance and generalization of network infrastructure characterization systems. One of the limitations of Atropos is that while it provides a significant IP expansion based on what is publicly reported, it can still not characterize all known APT domain names with the IPs they likely utilized when they were active. Future works could try to apply similar methodologies to Atropos in other DNS datasets, including passive DNS in order to increase the coverage and comprehensiveness in APT domain characterization. Passive DNS datasets from various regions around the world could provide IP visibility to domain names that were not visible by the active DNS datasets we utilized in our study.

Additionally, since Atropos is a general system that works across all domain names

of APT actors, that may introduce a lower precision or comprehensiveness in the domain and IP characterizations of some actors. The network provisioning techniques can vary per actor, thus, a general model may not be precise across all actors. Future works could aim to build dedicated models that operate on particular APT actors, or groups of actors that utilize similar provisioning techniques, in order to improve the characterization performance. The same methodology could be applied to the development of dedicated models that deal with the characterization of fast-fluxing domain names, which we excluded from our study. The utilization of fast-fluxing malicious domain names even after their detection by some APT actors like the Gamaredon Group (e.g., in some ccTLDs that delay their takedown) is significantly different from most of the other actors that stop utilizing their domain names or pointing new IP infrastructure to them after detection. Future studies could aim to build models or characterize the modus operandi of such divergent groups.

Furthermore, a significant part missing from the community related to the network infrastructure of sophisticated threats is that of historical benchmark datasets of real-world utilized infrastructure. This is obviously hard, as sophisticated actors that are often backed by nation-states will not release their utilized infrastructure to the public. In our study, we had to manually extract and label historical resource records associated with known APT domains in order to create training and evaluation datasets for Atropos, which was very time-consuming and could be subject to the subjective biases of the junior and senior analysts we utilized. The community should focus on building bigger, community-evaluated ground truth datasets of the historical infrastructure associated with sophisticated threats and their domain names in order to enable consistent comparison among the systems and studies that will follow. Additionally, these datasets could also enable the creation of more dedicated systems that are specialized to particular actors or groups of actors.

Understanding and measuring the impact of lifecycle analysis in domain detection and reputation systems. In chapter 5 we have discussed how identifying when a malicious domain name was historically active is a hard process and presented as a solution to

this problem, our system, Atropos. Sophisticated actors have been utilizing the process of strategically aging their domain names in order to evade detectors that heavily rely on features such as the temporal proximity of a domain name to its registration in order to consider it as malicious [143, 144]. While recent works have demonstrated that such systems can easily be evaded [142], future works can utilize Atropos to identify the active period and infrastructure of malicious domain names and measure the impact of incorporating such features in their training datasets. As we discussed in chapter 5, a malicious domain name can also be associated with parking, sinkhole, and other IP infrastructure that is not relevant to an attack. Future works can incorporate this signal into their detectors and not indiscriminately train their models, assuming all the IP infrastructure of a malicious domain name is relevant to the attacks. For example, future works can utilize systems and methodologies like Atropos to identify the particular IPs associated with a domain name that were likely utilized during an attack and train their domain or IP reputation systems only with these IPs and not the parking, sinkhole, and other unrelated IPs (e.g., previous or future owners of a domain name) that are not associated with the infrastructure the cybercriminals used. As demonstrated in chapter 5, over 70% of the attack-utilized IP addresses in sophisticated attacks no longer point to their domains at the time of the first public reporting, so future detection systems need to take that into account when training.

Attribution of APT attacks using an expanded network infrastructure dataset and utilizing the provisioning lifecycle patterns of different threats. In chapter 4 and chapter 5, we discussed the implications of utilizing IP infrastructure for malicious threat actor attribution. In agreement with experts [22], we have found that attribution exclusively utilizing IP infrastructure exhibits many challenges due to the proxies, VPNs, and cloud-fronting that modern cybercriminals utilize. However, we have also demonstrated that cybercriminals often reuse infrastructure, and this can be the basis of future attribution work. Attribution analysts can expand known and publicly reported APT infrastructure utilizing tools like Atropos, and then focus on identifying the subset of the infrastructure that gets

reused. Additionally, future network attribution efforts can focus on metadata regarding the infrastructure provisioning of APT actors, such as the average delta of infrastructure provisioning across each actor, the parking and infrastructure that different actors utilize to strategically age their domain names [140, 141], and the lifecycle patterns of provisioning across different domain names. Such efforts could lead to more robust attribution systems than simply matching future infrastructure based on domain names and IPs alone.

6.4 Closing Remarks

This dissertation aims to provide real-world insights into how network infrastructure is being utilized by malicious threat actors throughout their operations and over the years, and how the characterization of these interactions can affect common security tasks. The findings presented in this thesis validate prior works [40, 82] with respect to the geographic distribution of malicious infrastructure and its victims, but also expert insight on attribution [22]. We have shown that victim analysis utilizing authoritative DNS datasets is temporally sensitive, and that threat analysts need to take into account DNS scanners and security vendors who are overwhelmingly scanning malicious infrastructure after detection. We conducted an empirical analysis of the interactions of cybercriminals with their botnet management panels, and we identified that the detection event is significant in curbing their operations, with 69.03% of the operators stopping accessing their panels within 30 days of their detection. Finally, we demonstrated how sophisticated malicious threat actors provision their infrastructure IPs to their domains on average 317 days before their attacks are publicly reported, and provided a measurement methodology, Atropos, capable of expanding their known IP infrastructure by over three times from that which is publicly reported in threat reports. We hope that the findings and tools presented in this dissertation will lead to the development of more robust and accurate methodologies and systems with applications to malicious infrastructure detection and attribution.

Appendices

APPENDIX A

ATROPOS GENERALIZATION FOR THREAT HUNTING

To see whether Atropos can generalize and adapt to different analyst requirements, such as threat hunting for IPs that are non-cloud-fronting and virtual hosting – similar to the labeling methodology of EA described in Section subsection 5.4.2 –, we modify the PR dataset by flipping all the labels of IP addresses with more than 200 concurrent domain names pointed to them as non-APT controlled to imitate EA labeling process, changing 63 resource records from APT-controlled to non-APT controlled. We name this dataset PR-NVH. We train our model again utilizing PR-NVH and report our results in Table Table A.1. We observe that the accuracy and precision of the new model improve compared to those presented in Table Table 5.6, meaning that Atropos can be trained on datasets with different requirements and provide accurate results for different use cases that are outside of the scope of our study.

Table A.1: Evaluation of Atropos trained with and altered PR dataset.

DNS Dataset	Test Set	ROC AUC	F1-Macro	Accuracy	Precision	Recall
Active DNS	EA	87.86%	89.03%	93.58%	85.45%	93.58%
Virus Total	EA	87.29%	88.06%	92.39%	83.22%	92.39%

REFERENCES

- [1] L. Zeltser, “The evolution of malicious agents,” *Web Report*. Available online at, 2000.
- [2] Mandiant, *Highly Evasive Attacker Leverages SolarWinds Supply Chain to Compromise Multiple Global Victims With SUNBURST Backdoor*, <https://www.mandiant.com/resources/blog/evasive-attacker-leverages-solarwinds-supply-chain-compromises-with-sunburst-backdoor/>, 2020.
- [3] R. Al Halaseh and J. Alqatawna, “Analyzing cybercrimes strategies: The case of phishing attack,” in *2016 cybersecurity and cyberforensics conference (ccc)*, IEEE, 2016, pp. 82–88.
- [4] J. Gardiner, M. Cova, and S. Nagaraja, “Command & control: Understanding, denying and detecting—a review of malware c2 techniques, detection and defences,” *arXiv preprint arXiv:1408.1136*, 2014.
- [5] A. Luotonen and K. Altis, “World-wide web proxies,” *Computer Networks and ISDN systems*, vol. 27, no. 2, pp. 147–154, 1994.
- [6] X. Mi *et al.*, “Resident evil: Understanding residential ip proxy as a dark service.”
- [7] H. Sanghvi and M. Dahiya, “Cyber reconnaissance: An alarm before cyber attack,” *International Journal of Computer Applications*, vol. 63, no. 6, 2013.
- [8] RiskIQ, *SolarWinds: Advancing the Story*, <https://community.riskiq.com/article/9a515637>, 2021.
- [9] M. Antonakakis *et al.*, “Understanding the mirai botnet,” in *26th USENIX security symposium (USENIX Security 17)*, 2017, pp. 1093–1110.
- [10] A. Oest *et al.*, “Sunrise to sunset: Analyzing the end-to-end life cycle and effectiveness of phishing attacks at scale,” in *29th {USENIX} Security Symposium ({USENIX} Security 20)*, 2020.
- [11] O. Alrawi, C. Lever, K. Valakuzhy, K. Snow, F. Monroe, M. Antonakakis, *et al.*, “The circle of life: A {large-scale} study of the {iot} malware lifecycle,” in *Proceedings of the 30th USENIX Security Symposium (USENIX Security 21)*, 2021.
- [12] C. Lu *et al.*, “From whois to whowas: A large-scale measurement study of domain registration privacy under the gdpr,” in *NDSS*, 2021.
- [13] P. Hoffman and P. McManus, *Rfc 8484: Dns queries over https (doh)*, 2018.

- [14] B. Farinholt *et al.*, “To catch a ratter: Monitoring the behavior of amateur darkcomet rat operators in the wild.”
- [15] S. Le Blond, A. Uritesc, C. Gilbert, Z. L. Chua, P. Saxena, and E. Kirda, “A look at targeted attacks through the lense of an NGO.”
- [16] W. R. Marczak, J. Scott-Railton, M. Marquis-Boire, and V. Paxson, “When governments hack opponents: A look at actors and technology.”
- [17] M. Rezaeirad, B. Farinholt, H. Dharmdasani, P. Pearce, K. Levchenko, and D. McCoy, “Schrödinger’s RAT: Profiling the stakeholders in the remote access trojan ecosystem.”
- [18] Malbeacon, *Malbeacon*, <https://www.malbeacon.com/>, 2020.
- [19] Verizon, “2023 data breach investigations report,” <https://inquest.net/wp-content/uploads/2023-data-breach-investigations-report-dbir.pdf>,
- [20] A. Alageel and S. Maffei, “Hawk-eye: Holistic detection of apt command and control domains,” in *Proceedings of the 36th Annual ACM Symposium on Applied Computing*, 2021, pp. 1664–1673.
- [21] C. Tankard, “Advanced persistent threats and how to monitor and deter them,” *Network security*, vol. 2011, no. 8, pp. 16–19, 2011.
- [22] A. Saha, J. Mattei, J. Blasco, L. Cavallaro, D. Votipka, and M. Lindorfer, “Expert insights into advanced persistent threats: Analysis, attribution, and challenges,”
- [23] H. E. Decloedt and R. P. Van Heerden, “Rootkits, trojans, backdoors and new developments,” 2010.
- [24] Z. Chen, L. Gao, and K. Kwiat, “Modeling the spread of active worms,” in *IEEE INFOCOM 2003. Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE Cat. No. 03CH37428)*, IEEE, vol. 3, 2003, pp. 1890–1900.
- [25] R. De Silva, M. Nabeel, C. Elvitigala, I. Khalil, T. Yu, and C. Keppitiyagama, “Compromised or {attacker-owned}: A large scale classification and study of hosting domains of malicious {urls},” in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 3721–3738.
- [26] S. Maroofi, M. Korczyński, C. Hesselman, B. Ampeau, and A. Duda, “Comar: Classification of compromised versus maliciously registered domains,” in *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, IEEE, 2020, pp. 607–623.

- [27] E. Alowaisheq *et al.*, “Cracking the wall of confinement: Understanding and analyzing malicious domain.”
- [28] J. Caballero, G. Gomez, S. Matic, G. Sánchez, S. Sebastián, and A. Villacañas, “Goodfatr: A platform for automated threat report collection and ioc extraction,” *arXiv preprint arXiv:2208.00042*, 2022.
- [29] P. Mockapetris, *Domain names - concepts and facilities*, RFC 1034 (INTERNET STANDARD), Updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592, 5936, Internet Engineering Task Force, Nov. 1987.
- [30] P. Mockapetris and K. J. Dunlap, “Development of the domain name system,” in *Symposium Proceedings on Communications Architectures and Protocols*, 1988.
- [31] P. V. Mockapetris, *Rfc1035: Domain names-implementation and specification*, 1987.
- [32] A. Kountouras *et al.*, “Enabling network security through active dns datasets,” in *Proceedings of the International Symposium on Research in Attacks, Intrusions, and Defenses*, Springer, 2016.
- [33] C. Lever, R. Walls, Y. Nadji, D. Dagon, P. McDaniel, and M. Antonakakis, “Domainz: 28 registrations later measuring the exploitation of residual trust in domains,” in *2016 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2016, pp. 691–706.
- [34] A. Affinito *et al.*, “Domain name lifetimes: Baseline and threats,”
- [35] S. Lloyd, C. Hernandez-Gañan, and S. Tajalizadehkhoob, “Towards more rigorous domain-based metrics: Quantifying the prevalence and implications of “active” domains,” in *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, IEEE, 2023, pp. 539–545.
- [36] S. Sebastián, R.-G. Diugan, J. Caballero, I. Sanchez-Rola, and L. Bilge, “Domain and website attribution beyond whois,” in *Proceedings of the 39th Annual Computer Security Applications Conference*, 2023, pp. 124–137.
- [37] M. Antonakakis, R. Perdisci, W. Lee, N. Vasiloglou II, and D. Dagon, “Detecting malware domains at the upper {dns} hierarchy,” in *Proceedings of the 20th USENIX Security Symposium (USENIX Security 11)*, 2011.
- [38] M. Thomas and A. Mohaisen, “Kindred domains: Detecting and clustering botnet domains using dns traffic,” in *Proceedings of the 23rd International Conference on World Wide Web*, 2014.

- [39] S. Hao, N. Feamster, and R. Pandrangi, “Monitoring the initial dns behavior of malicious domains,” in *Proceedings of the 2011 ACM Internet Measurement Conference (IMC 11)*, 2011.
- [40] C. Lever, P. Kotzias, D. Balzarotti, J. Caballero, and M. Antonakakis, “A lustrum of malware network communication: Evolution and insights,” in *Proceedings of the 2017 IEEE Symposium on Security and Privacy (S&P)*, 2017.
- [41] P. Kotzias, L. Bilge, P.-A. Vervier, and J. Caballero, “Mind your own business: A longitudinal study of threats and vulnerabilities in enterprises,” in *Proceedings of the 2019 Network and Distributed System Security Symposium (NDSS 19)*, 2019.
- [42] C. Kanich *et al.*, “Spamalytics: An empirical analysis of spam marketing conversion.”
- [43] K. Levchenko *et al.*, “Click trajectories: End-to-end analysis of the spam value chain.”
- [44] B. Stone-Gross, T. Holz, G. Stringhini, and G. Vigna, “The underground economy of spam: A botmaster’s perspective of coordinating large-scale spam campaigns.,” *4th {USENIX} Workshop on Large-Scale Exploits and Emergent Threats ({LEET} 11)*, 2011.
- [45] P. Pearce *et al.*, “Characterizing large-scale click fraud in zeroaccess.”
- [46] D. Y. Huang *et al.*, “Tracking ransomware end-to-end.”
- [47] G. Stringhini, O. Hohlfeld, C. Kruegel, and G. Vigna, “The harvester, the botmaster, and the spammer: On the relations between the different actors in the spam landscape.”
- [48] J. Franklin, A. Perrig, V. Paxson, and S. Savage, “An inquiry into the nature and causes of the wealth of internet miscreants..”
- [49] T. Barron and N. Nikiforakis, “Picky attackers: Quantifying the role of system properties on intruder behavior.”
- [50] B. Stone-Gross *et al.*, “Your botnet is my botnet: Analysis of a botnet takeover,” in *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS 09)*, 2009.
- [51] C. Y. Cho, J. Caballero, C. Grier, V. Paxson, and D. Song, “Insights from the inside: A view of botnet management from infiltration.,” *4th {USENIX} Workshop on Large-Scale Exploits and Emergent Threats ({LEET} 11)*, 2010.

- [52] K. Thomas *et al.*, “Data breaches, phishing, or malware? understanding the risks of stolen credentials.”
- [53] A. Mirian, J. DeBlasio, S. Savage, G. M. Voelker, and K. Thomas, “Hack for hire: Exploring the emerging market for account hijacking.”
- [54] T. Holz, M. Engelberth, and F. Freiling, “Learning more about the underground economy: A case-study of keyloggers and dropzones,” in *European Symposium on Research in Computer Security*, 2009.
- [55] N. Nikiforakis *et al.*, “Stranger danger: Exploring the ecosystem of ad-based url shortening services.”
- [56] J. Szurdi and N. Christin, “Email typosquatting.”
- [57] E. De Cristofaro, A. Friedman, G. Jourjon, M. A. Kaafar, and M. Z. Shafiq, “Paying for likes? understanding facebook like fraud using honeypots.”
- [58] M. B. Salem and S. J. Stolfo, “Decoy document deployment for effective masquerade attack detection.”
- [59] J. Yuill, M. Zappe, D. Denning, and F. Feer, “Honeyfiles: Deceptive files for intrusion detection,” in *Proceedings from the Fifth Annual IEEE SMC Information Assurance Workshop, 2004.*, 2004.
- [60] A. Kapravelos, C. Grier, N. Chachra, C. Kruegel, G. Vigna, and V. Paxson, “Hulk: Eliciting malicious behavior in browser extensions.”
- [61] A. Oprea, Z. Li, T.-F. Yen, S. H. Chin, and S. Alrwais, “Detection of early-stage enterprise infection by mining large-scale log data,” in *2015 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, IEEE, 2015, pp. 45–56.
- [62] P. Lamprakis, R. Dargenio, D. Gugelmann, V. Lenders, M. Happe, and L. Vanbever, “Unsupervised detection of apt c&c channels using web request graphs,” in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*, Springer, 2017, pp. 366–387.
- [63] D. Chiba, H. Nakano, and T. Koide, “Domaindynamics: Advancing lifecycle-based risk assessment of domain names,” *Computers & Security*, vol. 153, p. 104366, 2025.
- [64] B. Bowman, C. Laprade, Y. Ji, and H. H. Huang, “Detecting lateral movement in enterprise computer networks with unsupervised graph {ai},” in *23rd International*

Symposium on Research in Attacks, Intrusions and Defenses (*{RAID}* 2020), 2020, pp. 257–268.

- [65] G. Ho *et al.*, “Hopper: Modeling and detecting lateral movement,” in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 3093–3110.
- [66] I. J. King and H. H. Huang, “Euler: Detecting network lateral movement via scalable temporal link prediction,” *ACM Transactions on Privacy and Security*, vol. 26, no. 3, pp. 1–36, 2023.
- [67] J. Khoury, D. Klisura, H. Zanddizari, G. Parra, P. Najafirad, and E. Bou-Harb, “Jbeil: Temporal graph-based inductive learning to infer lateral movement in evolving enterprise networks,” in *2024 IEEE Symposium on Security and Privacy (SP)*, IEEE Computer Society, 2023, pp. 9–9.
- [68] S. M. Milajerdi, R. Gjomemo, B. Eshete, R. Sekar, and V. Venkatakrishnan, “Holmes: Real-time apt detection through correlation of suspicious information flows,” in *2019 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2019, pp. 1137–1152.
- [69] X. Han, T. Pasquier, A. Bates, J. Mickens, and M. Seltzer, “Unicorn: Runtime provenance-based detector for advanced persistent threats,” *arXiv preprint arXiv:2001.01525*, 2020.
- [70] A. Alsaheel *et al.*, “{Atlas}: A sequence-based learning approach for attack investigation,” in *30th {USENIX} Security Symposium ({USENIX} Security 21)*, 2021.
- [71] H. Irshad *et al.*, “Trace: Enterprise-wide provenance tracking for real-time apt detection,” *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 4363–4376, 2021.
- [72] M. U. Rehman, H. Ahmadi, and W. U. Hassan, “Flash: A comprehensive approach to intrusion detection via provenance graph representation learning,” in *2024 IEEE Symposium on Security and Privacy (SP)*, IEEE Computer Society, 2024, pp. 139–139.
- [73] M. N. Hossain *et al.*, “{Sleuth}: Real-time attack scenario reconstruction from {cots} audit data,” in *26th {USENIX} Security Symposium ({USENIX} Security 17)*, 2017, pp. 487–504.
- [74] F. Liu, Y. Wen, D. Zhang, X. Jiang, X. Xing, and D. Meng, “Log2vec: A heterogeneous graph embedding based approach for detecting cyber threats within enterprise,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 1777–1794.

- [75] T. Urban, M. Große-Kampmann, D. Tatang, T. Holz, and N. Pohlmann, “Plenty of phish in the sea: Analyzing potential pre-attack surfaces,” in *Computer Security—ESORICS 2020: 25th European Symposium on Research in Computer Security, ESORICS 2020, Guildford, UK, September 14–18, 2020, Proceedings, Part II 25*, Springer, 2020, pp. 272–291.
- [76] Malwarebytes, *2022 threat review*, https://www.malwarebytes.com/resources/malwarebytes-threat-review-2022/mwb_threatreview_2022_ss_v1.pdf, 2022.
- [77] Malwarebytes, *2022 global threat report*, <https://go.crowdstrike.com/rs/281-OBQ-266/images/Report2022GTR.pdf>, 2022.
- [78] E. Alowaisheq *et al.*, “Cracking the wall of confinement: Understanding and analyzing malicious domain take-downs,” in *Proceedings of the 2019 Network and Distributed System Security Symposium (NDSS 19)*, 2019.
- [79] M. Antonakakis *et al.*, “Understanding the Mirai botnet,” in *Proceedings of the 26th USENIX Security Symposium (USENIX Security 17)*, 2017.
- [80] A. Kharraz, S. Arshad, C. Mulliner, W. Robertson, and E. Kirda, “UNVEIL: A large-scale, automated approach to detecting ransomware,” in *Proceedings of the 25th USENIX Security Symposium (USENIX Security 16)*, 2016.
- [81] S. Tajalizadehkhoob, C. Gañán, A. Noroozian, and M. v. Eeten, “The role of hosting providers in fighting command and control infrastructure of financial malware,” in *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, 2017.
- [82] G. Mezzour, K. M. Carley, and L. R. Carley, “Global variation in attack encounters and hosting,” in *Proceedings of Hot Topics in Science of Security: Symposium and Bootcamp*, ACM, 2017.
- [83] G. Tech, *Gt malware passive dns data daily feed*, https://impactcybertrust.org/dataset_view?idDataset=520, 2022.
- [84] *Virustotal*, <https://www.virustotal.com>, 2022.
- [85] S. Sebastián and J. Caballero, “Avclass2: Massive malware tag extraction from av labels,” in *Proceedings of the 2020 Computer Security Applications Conference*, 2020.
- [86] CAIDA, *Routeviews prefix-to-as mappings (pfx2as) for ipv4 and ipv6*, <http://data.caida.org/datasets/routing/routeviews-prefix2as/>, 2022.

- [87] M. Ziv, L. Izhikevich, K. Ruth, K. Izhikevich, and Z. Durumeric, “ASdb: A system for classifying owners of autonomous systems,” in *Proceedings of the 2021 ACM Internet Measurement Conference (IMC 21)*, 2021.
- [88] *Unsd - statistical classifications*, <https://unstats.un.org/unsd/classifications>, 2019.
- [89] V. L. Pochat, T. Van Goethem, S. Tajalizadehkhoob, M. Korczyński, and W. Joosen, “Tranco: A research-oriented top sites ranking hardened against manipulation,” *arXiv preprint arXiv:1806.01156*, 2018.
- [90] A. Mohaisen and O. Alrawi, “Av-meter: An evaluation of antivirus scans and labels,” in *Proceedings of the International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*, Springer, 2014.
- [91] D. Plohmann, M. Clauss, S. Enders, and E. Padilla, “Malpedia: A collaborative effort to inventorize the malware landscape,” *Botconf*, 2017.
- [92] A. Kountouras *et al.*, “Understanding the growth and security considerations of ecs.,” in *Proceedings of the 2021 Network and Distributed System Security Symposium (NDSS 21)*, 2021.
- [93] A. Randall *et al.*, “Trufflehunter: Cache snooping rare domains at large public DNS resolvers,” in *Proceedings of the 2020 ACM Internet Measurement Conference (IMC 20)*, 2020.
- [94] U. Bayer, I. Habibi, D. Balzarotti, E. Kirda, and C. Kruegel, “A view on current malware behaviors,” in *Proceedings of the 2nd USENIX Conference on Large-scale Exploits and Emergent Threats (LEET 09)*, 2009.
- [95] A. Mohaisen, O. Alrawi, and M. Mohaisen, “Amal: High-fidelity, behavior-based automated malware analysis and classification,” *computers & security*, vol. 52, 2015.
- [96] C. Rossow, C. Dietrich, and H. Bos, “Large-scale analysis of malware downloaders,” in *Proceedings of the International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*, Springer, 2012.
- [97] M. Lindorfer, M. Neugschwandtner, L. Weichselbaum, Y. Fratantonio, V. Van Der Veen, and C. Platzer, “Andrubis–1,000,000 apps later: A view on current android malware behaviors,” in *Proceedings of the 3rd International Workshop on Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*, IEEE, 2014.
- [98] Z. Xu, A. Nappa, R. Baykov, G. Yang, J. Caballero, and G. Gu, “Autoprobe: Towards automatic active malicious server probing using dynamic binary analysis,”

in *Proceedings of the 21st ACM Conference on Computer and Communications Security (CCS 14)*, 2014.

- [99] A. Nappa, Z. Xu, M. Z. Rafique, J. Caballero, and G. Gu, “Cyberprobe: Towards internet-scale active detection of malicious servers,” in *Proceedings of the 2014 Network and Distributed System Security Symposium (NDSS 14)*, 2014.
- [100] Y.-M. Wang, D. Beck, X. Jiang, and R. Roussev, “Automated web patrol with strider honeymoons: Finding web sites that exploit browser vulnerabilities,” in *Proceedings of the 2006 Network and Distributed System Security Symposium (NDSS 06)*, 2006.
- [101] N. Provos, P. Mavrommatis, M. Rajab, and F. Monrose, “All your iframes point to us,” in *Proceedings of the 17th USENIX Security Symposium (USENIX Security 08)*, 2008.
- [102] A. Mohaisen and O. Alrawi, “Unveiling zeus: Automated classification of malware samples,” in *Proceedings of the 22nd International Conference on World Wide Web*, 2013.
- [103] A. Avgetidis *et al.*, “Beyond the gates: An empirical analysis of http-managed password stealers and operators,” in *Proceedings of the 32nd USENIX Security Symposium (USENIX Security 23)*, 2023.
- [104] W. Chang, A. Mohaisen, A. Wang, and S. Chen, “Measuring botnets in the wild: Some new trends,” in *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security*, 2015.
- [105] N. Weaver, C. Kreibich, and V. Paxson, “Redirecting {dns} for ads and profit,” in *Proceedings of the 2011 USENIX Workshop on Free and Open Communications on the Internet (FOCI 11)*, 2011.
- [106] P. A. Networks, *2024 Unit 42 Incident Response Report: Navigating the Shift in Cybersecurity Threat Tactics*, <https://unit42.paloaltonetworks.com/unit42-incident-response-report-2024-threat-guide/>.
- [107] SecurityTrails, “URLScan - a free service to scan and analyse websites.,” <https://urlscan.io/about/>, 2016.
- [108] IPRegistry, “IP Geolocation and Threat Detection,” <https://ipregistry.co/>,
- [109] ViriBack Tracker, “C2 Tracker,” <http://tracker.viriback.com/>,
- [110] benkow, “Panel Tracker,” <https://benkow.cc/passwords.php>,

- [111] Cybercrime Tracker, “Bot Tracker,” <https://cybercrime-tracker.net/index.php>,
- [112] X. Mi, S. Tang, Z. Li, X. Liao, F. Qian, and X. Wang, “Your phone is my proxy: Detecting and understanding mobile proxy networks.”
- [113] Blueliv, “The Credential Theft Ecosystem,” https://web.archive.org/web/20210107175227/https://www.blueliv.com/resources/reports/The_credential_theft_ecosystem.pdf, 2018.
- [114] d00rt, “Lokibot infostealer “hijacked” version,” https://web.archive.org/web/20210117055851/https://raw.githubusercontent.com/d00rt/hijacked_lokibot_version/master/doc/LokiBot_hijacked_2018.pdf, 2018.
- [115] N. Villeneuve, R. Eitzman, S. Nemes, and T. Dean, “Significant FormBook Distribution Campaigns Impacting the U.S. and South Korea,” <https://web.archive.org/web/20201211012535/https://www.fireeye.com/blog/threat-research/2017/10/formbook-malware-distribution-campaigns.html>, 2017.
- [116] KrabsOnSecurity, “Analyzing Amadey – a simple native malware,” <https://web.archive.org/web/20201107235827/https://krabsonsecurity.com/2019/02/13/analyzing-amadey-a-simple-native-malware/>, 2019.
- [117] A. Zsigovits, “Baldr vs The World,” <https://web.archive.org/web/20191217054044/https://www.sophos.com/en-us/medialibrary/PDFs/technical-papers/baldr-vs-the-world.pdf>, 2019.
- [118] Partheeban, “Dark Side Of BlackNET RAT,” <https://web.archive.org/web/20201224065951/https://labs.k7computing.com/?p=21365>, 2020.
- [119] The BlackBerry Cylance Threat Research Team, “Threat Spotlight: Analyzing AZORult Infostealer Malware,” <https://web.archive.org/web/20200920132950/https://blogs.blackberry.com/en/2019/06/threat-spotlight-analyzing-azorult-infostealer-malware>, 2019.
- [120] Malware Don’t Need Coffee, “Neutrino Bot (aka MS:Win32/Kasidet),” <https://web.archive.org/web/20201129222945/https://malware.dontneedcoffee.com/2014/06/neutrino-bot-aka-kasidet.html>, 2014.
- [121] J. WALTER, “Agent Tesla — Old RAT Uses New Tricks to Stay on Top,” <https://web.archive.org/web/20201207093051/https://labs.sentinelone.com/agent-tesla-old-rat-uses-new-tricks-to-stay-on-top/>, 2020.
- [122] fr3dhk, “Nexus - Just another stealer,” <https://web.archive.org/web/20201129072131/https://fr3d.hk/blog/nexus-just-another-stealer>, 2020.

- [123] D. SCHWARZ, “New KPOT v2.0 stealer brings zero persistence and in-memory features to silently steal credentials,” <https://web.archive.org/web/20201207040629/https://www.proofpoint.com/us/threat-insight/post/new-kpot-v20-stealer-brings-zero-persistence-and-memory-features-silently-steal>, 2019.
- [124] *Alert (AA20-266A): LokiBot Malware*, Accessed on 08/01/2021.
- [125] S. Nachum, A. Schuster, and O. Etzion, “Detection in the dark—exploiting xss vulnerability in c&c panels to detect malwares,” in *International Symposium on Cyber Security Cryptography and Machine Learning*, Springer, 2018.
- [126] Mozilla, *Firefox 87 trims http referrers by default to protect user privacy*, <https://blog.mozilla.org/security/2021/03/22/firefox-87-trims-http-referrers-by-default-to-protect-user-privacy/>, 2021.
- [127] A. Dasgupta, M. Gurevich, L. Zhang, B. Tseng, and A. O. Thomas, “Overcoming browser cookie churn with clustering,” in *Proceedings of the fifth ACM international conference on Web search and data mining*, 2012.
- [128] B. Liu, Z. Liu, J. Zhang, T. Wei, and W. Zou, “How many eyes are spying on your shared folders?” In *Proceedings of the 2012 ACM workshop on Privacy in the electronic society*, 2012.
- [129] M. Bailey, D. Dittrich, E. Kenneally, and D. Maughan, “The menlo report,” *IEEE Security & Privacy*, 2012.
- [130] L. Gelinas, A. Wertheimer, and F. G. Miller, “When and why is research without consent permissible?” *Hastings Center Report*, 2016.
- [131] Amazon Web Services, Inc., “Alexa top sites.,” <https://www.alexa.com/topsites>, 2021.
- [132] The Spamhaus Project, “The 10 Most Abused Top Level Domains,” <https://www.spamhaus.org/statistics/tlds/>,
- [133] P. Kintis, Y. Nadji, D. Dagon, M. Farrell, and M. Antonakakis, “Understanding the privacy implications of ecs.”
- [134] R. Future, *Initial access brokers are key to rise in ransomware attacks*, <https://www.recordedfuture.com/initial-access-brokers-key-to-rise-in-ransomware-attacks>, 2022.
- [135] Statcounter, *Desktop browser market share worldwide*, <https://gs.statcounter.com/browser-market-share/desktop/worldwide/monthly-201904-202012>, 2022.

- [136] O. E. Agboje, S. O. Adedoyin, and C. U. Ndujiuba, “State of fiber optic networks for internet broadband penetration in nigeria-a review,” *International Journal of Optoelectronic Engineering*, vol. 7, no. 1, pp. 1–12, 2017.
- [137] Scamalytics, *M247 ltd - fraud risk scamalytics*, <https://scamalytics.com/ip/isp/m247-ltd>, 2022.
- [138] Scamalytics, *Avast software s.r.o. - fraud risk scamalytics*, <https://scamalytics.com/ip/isp/avast-software-s-r-o>, 2022.
- [139] Scamalytics, *Hern labs ab - fraud risk scamalytics*, <https://scamalytics.com/ip/isp/hern-labs-ab>, 2022.
- [140] Infoblox, *Dog Hunt: Finding Decoy Dog Toolkit via Anomalous DNS Traffic*, <https://blogs.infoblox.com/threat-intelligence/cyber-threat-advisory/dog-hunt-finding-decoy-dog-toolkit-via-anomalous-dns-traffic/>.
- [141] P. A. Networks, *Strategically Aged Domain Detection: Capture APT Attacks With DNS Traffic Trends*, <https://unit42.paloaltonetworks.com/strategically-aged-domain-detection/>.
- [142] T. Galloway, K. Karakolios, Z. Ma, R. Perdisco, A. Keromytis, and M. Antonakakis, “Practical attacks against dns reputation systems,” in *2024 IEEE Symposium on Security and Privacy (SP)*, 2024.
- [143] M. Antonakakis, R. Perdisci, D. Dagon, W. Lee, and N. Feamster, “Building a dynamic reputation system for {dns},” in *19th USENIX Security Symposium (USENIX Security 10)*, 2010.
- [144] L. Bilge, E. Kirda, C. Kruegel, and M. Balduzzi, “Exposure: Finding malicious domains using passive dns analysis.” in *Ndss*, 2011, pp. 1–17.
- [145] A. Ahmad, J. Webb, K. C. Desouza, and J. Boorman, “Strategically-motivated advanced persistent threat: Definition, process, tactics and a disinformation model of counterattack,” *Computers & Security*, vol. 86, pp. 402–418, 2019.
- [146] B. Binde, R. McRee, and T. J. O’Connor, “Assessing outbound traffic to uncover advanced persistent threat,” *SANS Institute. Whitepaper*, vol. 16, 2011.
- [147] B. E. Strom, A. Applebaum, D. P. Miller, K. C. Nickels, A. G. Pennington, and C. B. Thomas, “Mitre att&ck: Design and philosophy,” in *Technical report*, The MITRE Corporation, 2018.
- [148] Fortinet, *Solar winds cyber attack*, <https://www.fortinet.com/resources/cyberglossary/solarwinds-cyber-attack>, 2024.

- [149] Reuters, *Crypto's biggest hacks and heists after \$1.5 billion theft from Bybit*, <https://www.reuters.com/technology/cybersecurity/fbi-says-north-korea-was-responsible-15-billion-bybit-hack-2025-02-27/>.
- [150] G. Di Tizio, M. Armellini, and F. Massacci, "Software updates strategies: A quantitative evaluation against advanced persistent threats," *IEEE Transactions on Software Engineering*, vol. 49, no. 3, pp. 1359–1373, 2022.
- [151] M. R. Rahman, S. K. Basak, R. M. Hezaveh, and L. Williams, "Attackers reveal their arsenal: An investigation of adversarial techniques in cti reports," *arXiv preprint arXiv:2401.01865*, 2024.
- [152] W. Tounsi and H. Rais, "A survey on technical threat intelligence in the age of sophisticated cyber attacks," *Computers & security*, vol. 72, pp. 212–233, 2018.
- [153] T. Geras and T. Schreck, "The" big beast to tackle": Practices in quality assurance for cyber threat intelligence," in *Proceedings of the 27th International Symposium on Research in Attacks, Intrusions and Defenses*, 2024, pp. 337–352.
- [154] Domaintools, *Unraveling Network Infrastructure Linked to the SolarWinds Hack*, <https://www.domaintools.com/resources/blog/unraveling-network-infrastructure-linked-to-the-solarwinds-hack/>, 2020.
- [155] V. Total, *VirusTotal*, <https://www.virustotal.com/en>, 2024.
- [156] J. Zirngibl, S. Deusch, P. Sattler, J. Aulbach, G. Carle, and M. Jonker, "Domain parking: Largely present, rarely considered!" In *Proc. Network Traffic Measurement and Analysis Conference (TMA) 2022*, 2022.
- [157] E. Alowaisheq *et al.*, "Cracking the wall of confinement: Understanding and analyzing malicious domain,"
- [158] Domaintools, *Continuous Eruption: Further Analysis of the SolarWinds Supply Chain Incident*, <https://www.domaintools.com/resources/blog/continuous-eruption-further-analysis-of-the-solarwinds-supply-incident/>, 2020.
- [159] Domaintools, *Domaintools whois history*, <https://research.domaintools.com/research/whois-history/>, 2023.
- [160] N. Rani, B. Saha, and S. K. Shukla, "A comprehensive survey of advanced persistent threat attribution: Taxonomy, methods, challenges and open research problems," *arXiv preprint arXiv:2409.11415*, 2024.
- [161] O. AlienVault, *The World's First Truly Open Threat Intelligence Community*, 2024.

- [162] CyberMonitor, *Apt and cybercriminal campaign collection*. https://github.com/CyberMonitor/APT_CyberCriminal_Campagin_Collections, 2023.
- [163] A. Mantovani, S. Aonzo, X. Ugarte-Pedrero, A. Merlo, and D. Balzarotti, “Prevalence and impact of low-entropy packing schemes in the malware ecosystem.,” in *NDSS*, 2020.
- [164] A. Anwar, Y. H. Chen, R. Hodgman, T. Sellers, E. Kirda, and A. Oprea, “A recent year on the internet: Measuring and understanding the threats to everyday internet devices,” in *Proceedings of the 38th Annual Computer Security Applications Conference*, 2022, pp. 251–266.
- [165] MISP, *MISP Galaxy Threat Actors*, <https://raw.githubusercontent.com/MISP/misp-galaxy/main/clusters/threat-actor.json>, 2024.
- [166] MITRE, *MITRE Groups*, <https://attack.mitre.org/groups/>.
- [167] X. Bouwman *et al.*, “Helping hands: Measuring the impact of a large threat intelligence sharing community,” in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 1149–1165.
- [168] Cybermonitor, *APT & Cybercriminals Campaign Collection*, https://github.com/CyberMonitor/APT_CyberCriminal_Campagin_Collections, 2023.
- [169] zonefiles.io, *Compromised domain list*, <https://zonefiles.io/compromised-domain-list/>, 2024.
- [170] M. Stampar and M. Kasimov, *maltrail: Malicious traffic detection system*, 2019.
- [171] V. Total, *VirusTotal Historical Whois API*, <https://docs.virustotal.com/reference/domain-resolutions>.
- [172] K. Tian, S. T. Jan, H. Hu, D. Yao, and G. Wang, “Needle in a haystack: Tracking down elite phishing domains in the wild,” in *Proceedings of the Internet Measurement Conference 2018*, 2018, pp. 429–442.
- [173] T. Cymru, *Unravelling the Mystery of Bogons: A senior stakeholder and IT professional guide*, <https://www.team-cymru.com/post/unravelling-the-mystery-of-bogons-a-senior-stakeholder-and-it-professional-guide>.
- [174] T. Holz, C. Gorecki, K. Rieck, and F. C. Freiling, “Measuring and detecting fast-flux service networks.,” in *Ndss*, 2008.

- [175] P. A. Networks, '*Russia's Trident Ursa (aka Gamaredon APT) Cyber Conflict Operations Unwavering Since Invasion of Ukraine*', <https://unit42.paloaltonetworks.com/trident-ursa/>, 2022.
- [176] P. A. Networks, '*Russia's Gamaredon aka Primitive Bear APT Group Actively Targeting Ukraine*', <https://unit42.paloaltonetworks.com/gamaredon-primitive-bear-ukraine-update-2021/>, 2022.
- [177] A. Faulkenberry *et al.*, "View from above: Exploring the malware ecosystem from the upper dns hierarchy," in *Proceedings of the 38th Annual Computer Security Applications Conference*, 2022, pp. 240–250.
- [178] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.
- [179] D. Arp *et al.*, "Dos and don'ts of machine learning in computer security," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 3971–3988.
- [180] S. Sun, "Meta-analysis of cohen's kappa," *Health Services and Outcomes Research Methodology*, vol. 11, pp. 145–163, 2011.
- [181] KrebsonSecurity, '*Stark Industries Solutions: An Iron Hammer in the Cloud*', <https://krebsonsecurity.com/2024/05/stark-industries-solutions-an-iron-hammer-in-the-cloud/>.
- [182] S. Push, '*FIN7: Silent Push unearths the largest group of FIN7 domains ever discovered. 4000+ IOFA domains and IPs found. Louvre, Meta, and Reuters targeted in massive global phishing and malware campaigns.*' <https://www.silentpush.com/blog/fin7/>.
- [183] D. Instinct, '*MuddyC2Go – Latest C2 Framework Used by Iranian APT MuddyWater Spotted in Israel*', <https://www.deepinstinct.com/blog/muddyc2go-latest-c2-framework-used-by-iranian-apt-muddywater-spotted-in-israel>.
- [184] IPinfo, '*IPinfo*'. <https://ipinfo.io/>, 2024.
- [185] Pingdom, '*The top 100 web hosting countries*', <https://www.pingdom.com/blog/web-hosting-countries-2013/>.
- [186] I. Arnaldo, A. Cuesta-Infante, A. Arun, M. Lam, C. Bassias, and K. Veeramachani, "Learning representations for log data in cybersecurity," in *Cyber Security Cryptography and Machine Learning: First International Conference, CSCML 2017, Beer-Sheva, Israel, June 29-30, 2017, Proceedings 1*, Springer, 2017, pp. 250–268.

- [187] J. Gharibshah, T. C. Li, A. Castro, K. Pelechris, E. E. Papalexakis, and M. Faloutsos, “Mining actionable information from security forums: The case of malicious ip addresses,” *From Security to Community Detection in Social Networking Platforms*, pp. 193–211, 2019.
- [188] J. Gharibshah *et al.*, “Inferip: Extracting actionable information from security discussion forums,” in *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, 2017, pp. 301–304.
- [189] L. Wang, A. Nappa, J. Caballero, T. Ristenpart, and A. Akella, “Whowas: A platform for measuring web deployments on iaas clouds,” in *Proceedings of the 2014 ACM Internet Measurement Conference (IMC 14)*, 2014.
- [190] C. Rong, G. Gou, M. Cui, G. Xiong, Z. Li, and L. Guo, “Malfinder: An ensemble learning-based framework for malicious traffic detection,” in *2020 IEEE Symposium on Computers and Communications (ISCC)*, IEEE, 2020, pp. 7–7.
- [191] Snapmaker, *Snapmaker controller reaching out to Anubis networks sink hole*, <https://forum.snapmaker.com/t/snapmaker-controller-reaching-out-to-anubis-networks-sink-hole/31250>.
- [192] B. Liu *et al.*, “Who is answering my queries: Understanding and characterizing interception of the DNS resolution path,” in *Proceedings of the 27th USENIX Security Symposium (USENIX Security 18)*, 2018.
- [193] G. C. Moura, J. Heidemann, R. d. O. Schmidt, and W. Hardaker, “Cache me if you can: Effects of DNS time-to-live,” in *Proceedings of the 2019 ACM Internet Measurement Conference (IMC 19)*, 2019.
- [194] C. Kanich, K. Levchenko, B. Enright, G. M. Voelker, and S. Savage, “The heisenbot uncertainty problem: Challenges in separating bots from chaff,” in *Proceedings of the 1st USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET 08)*, 2008.
- [195] T. T. Project, *Tordnsel’s exit lists*, <https://metrics.torproject.org/collector/archive/exit-lists/>, 2022.
- [196] F. Li, A. Lai, and D. Ddl, “Evidence of advanced persistent threat: A case study of malware for political espionage,” in *2011 6th International Conference on Malicious and Unwanted Software*, IEEE, 2011, pp. 102–109.