Unsupervised Learning Techniques for Malware Characterization: Understanding Certain DNS-based DDoS Attacks

RENÉE BURTON, Infoblox

This article details data science research in the area of Cyber Threat Intelligence applied to a specific type of Distributed Denial of Service (DDoS) attack. We study a DDoS technique prevalent in the Domain Name System (DNS) for which little malware have been recovered. Using data from a globally distributed set of a passive collectors (pDNS), we create a statistical classifier to identify these attacks and then use unsupervised learning to investigate the attack events and the malware that generates them. The first known major study of this technique, this work demonstrates that current attacks have little resemblance to earlier published descriptions and identifies several features of the attacks. Through a combination of text and time-series features, we are able to characterize the dominant malware and demonstrate that the number of global-scale attack systems is relatively small.

CCS Concepts: • **Security and privacy** → **Malware and its mitigation; Denial-of-service attacks**; • **Computing methodologies** → **Unsupervised learning**;

Additional Key Words and Phrases: Domain name service (DNS), data science, malware, ddos attacks, botnet, clustering, threat intelligence

ACM Reference format:

Renée Burton. 2020. Unsupervised Learning Techniques for Malware Characterization: Understanding Certain DNS-based DDoS Attacks. *Digit. Threat.: Res. Pract.* 1, 3, Article 14 (August 2020), 26 pages.
<https://doi.org/10.1145/3377869>

1 INTRODUCTION

In the field of Cyber Security, there is a rich history of characterizing malicious actors, their mechanisms, and victims, through the analysis of recovered *malware*¹ and related event data. Reverse engineers and threat analysts take advantage of the fact that malware developers often unwittingly leave fingerprints in their code through the choice of libraries, variables, infection mechanisms, and other observables. In some cases, the Cyber Security Industry is able to correlate seemingly disparate pieces of information together and attribute malware to specific malicious actors. Identifying a specific technique or actor involved in an attack, allows the Industry to better understand the magnitude of the threat and protect against it. The study of cyber actors, their motives, and

¹We use the term malware here to mean the software used for malicious means, not limited to such software on compromised devices.

Authors' address: R. Burton, Infoblox, 2106 Pacific Ave 6th floor, Tacoma, WA 98402; email: rburton@infoblox.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2576-5337/2020/08-ART14 \$15.00

<https://doi.org/10.1145/3377869>

mechanisms, falls within Cyber Threat Intelligence, an area traditionally dominated by security analysts utilizing deep domain expertise, but increasingly open to data science methodologies. The malware itself is critical to the vast majority of work of this type.

Here, we address a class of Distributed Denial of Service (DDoS) attacks for which malware is rarely obtained, but which remains a persistent presence on the Internet. Motivated by the need to help threat analysts more easily identify and characterize threats, the results of this research demonstrate a dramatic change in these attacks in recent years, made visible through statistical analysis and unsupervised machine learning. Leveraging over six months of real data, we identify a number of features that allow us to cluster the majority of attacks into a handful of families and in some cases completely reverse the algorithm for generating the attack. Data analysis shows that actors have changed their modus operandi, making detection and correction more difficult for Internet providers. We developed a highly accurate statistical classifier to identify the attacks and engineered features not previously mentioned in the literature.

The attack, known as a Slow Drip, Random Subdomain, and Water Torture Attack, among other names, targets the Domain Name System (DNS), a globally distributed database that provides critical functionality for the Internet. First observed in 2009 [17], the attacks emerged with tremendous strength in early 2014. While their scale has varied tremendously, they remain a daily fixture of the cyber landscape today. The DNS community has focused primarily on mitigating ongoing attacks, largely through scaling the ability of systems to handle additional load. Research that attempted to address detection based on attack characteristics was hindered by the limited reporting available or focused on single systems. This large-scale study, utilizing current, real data, is the first of its kind.

By early 2018, threat analysts noted that it was increasingly difficult to determine whether high-volume traffic was part of a Slow Drip attack or due to some other errant Internet phenomena.² One analyst provided examples of spikes in AirBnB-related domains that appeared suspicious in terms of their volume and variety, but in which many of the subdomains could be interpreted as locations in China. In other examples, the subdomains contained a hodgepodge of unusual, but not random, words. Some victim domains, like AirBnB, were quite popular, while others were obscure. Were these attacks all related somehow?, they asked, and were they Slow Drip attacks at all? The extraordinarily large volume and variety of data made human analysis time-consuming and faulty. Our research sought to support these analytic questions and to understand the attack technique and related malware, as it appears on the Internet today. Specifically, we looked to address common threat intelligence questions:

- How many such malware systems exist?
- How can we attribute two attacks to the same malware?
- Has the threat landscape evolved over time, and how?
- Can we characterize the attack generation algorithms?

We begin by introducing DNS, the Slow Drip attack technique, and discussing prior work in Section 2. The data and methodologies used in this work are then detailed in Section 3. After describing the underlying datasets for this work, in Section 4, we review our method for detecting attacks in passive data. That same section provides a cursory analysis of the attack landscape over a six-month period in 2018. The article then transitions from identifying attacks to describing signatures within attacks. Section 6 discusses feature engineering and includes details of a number of novel features found in the data. Those features are used to cluster attacks in Section 7. Those results demonstrate that the features are meaningful, stand over time, and address some of the intelligence questions. Finally, we discuss the applicability of our approach to evolving threats, changes observed in 2019, and unresolved questions regarding these attacks in Section 8.

²Personal conversations between this author and professional cyber threat analysts.

2 BACKGROUND

In this section, we describe the background necessary to understand the Slow Drip attack mechanism and its impact on the Domain Name System. We introduce vocabulary that is used throughout the rest of the article and review the previous literature. Terms that are written in **bold face** are key concepts that will be used repeatedly through the article; other terms that are helpful but not used throughout are written in *italic* font.

2.1 The Domain Name System (DNS)

The Domain Name System (DNS) is a global, hierarchical, distributed database that serves, among other things, to map domain names to Internet Protocol (IP) addresses. While relatively straightforward in concept, in practice the global DNS is complex to the point of being arcane [10]. For the purpose of this article, we introduce a limited scope and vocabulary; the interested reader can find more depth in the operation and security of DNS in References [15] and [13].

The domain name system operates as a query-response protocol, in which a query for a **Fully Qualified Domain Name (FQDN)** is made by a client, or *endpoint*, and is answered via an iterative process known as *resolution*. An FQDN is made up of a series of text **labels** separated by periods. For example, the FQDN `www.google.com` has three labels [`www`,” `google`,” `com`].

From a hierarchical perspective, the right-most label in an FQDN is the *Top-Level-Domain (TLD)* and each subsequent label represents a **subdomain** of the FQDN created from all the prior labels. For instance `www.google.com` is a subdomain of `google.com`, which in turn is a subdomain of `com`. The term **domain** is often used to refer to both an FQDN and the scope of its possible subdomains. Thus, `www.google.com`, `mail.google.com`, `inbox.google.com`, and `photo.google.com` are all subdomains of the domain `google.com`. The TLD is a single label and is considered *public* in the sense that its subdomains are available for registration and not controlled by the TLD. The TLDs are limited in number and controlled by the Internet Corporation for Assigned Names and Numbers (ICANN). In some cases, subdomains of a TLD are also managed as public domains, most notably domains like `co.uk`, and are considered *public suffixes* and create extended TLDs (eTLD).³ A **Second-Level-Domain (SLD)** is privately owned and the direct subdomain of a public suffix, or eTLD. Examples of second-level-domains include `google.com` and `google.co.uk`. The SLD is sometimes referred to as the base domain.

Within the Domain Name System, *authoritative name servers* are servers that hold authoritative, or definitive, answers for a certain portion of the database, generally a specific domain. If these authoritative servers are not functioning properly, Internet traffic to the domains for which they are authoritative may be interrupted or completely disrupted. For this reason, companies often have multiple authoritative name servers for their domains.

While it is possible for an endpoint to resolve DNS queries themselves, in practice, most devices rely on large **recursive resolvers** to perform resolution on their behalf. Internet Service Providers provide recursive resolvers for their customers, for example. Many recursive resolvers are configured to answer queries only for devices in their network, thereby limiting the resource demands on those network appliances. There are *public resolvers*, such as those operated by Google, openDNS, and Cloudflare, designed to handle recursion for any endpoint selecting their service.⁴ However, there are also a large number of devices on the Internet that, through misconfiguration or otherwise, will act as recursive resolvers for any client but are not announced as public resolvers. These are known as **open resolvers** and are frequently leveraged by cyber actors to anonymize and amplify DDoS traffic.

In another effort to ensure efficiency, the Domain Name System employs *caching* of records at the recursive resolvers. Caching of records prevents unnecessary Internet traffic by caching the records for popular domain

³A list of public suffixes is maintained by Mozilla and found at publicsuffix.org.

⁴Endpoints are typically configured with at least one recursive resolver to which they forward their queries.

Fig. 1. Slow Drip attack mechanisms.

names, such as `www.google.com`. For security purposes, many recursive resolvers also employ negative caching, in which they remember, for some short period of time, that a given FQDN has no answer in the DNS. Queries for non-existent domains typically return what is referred to as an NXDOMAIN response, which triggers this form of caching.

2.2 Slow Drip DDoS Attacks

Given the fundamental role the DNS plays in the functioning of the Internet, attackers are known to abuse it. In the technique considered in this article, malicious actors generate a massive number of queries for non-existent domains. Through the resolution process, these queries are forwarded to the authoritative name servers, which may become overwhelmed with the unexpected volume. The attack generates non-existent subdomains of a common SLD, with little to no repetition in the queries. This serves to counter caching at recursive resolvers and force as much traffic as possible to the authoritative name servers. All of the queries within an attack will be subdomains of a common SLD, which we refer to as the **attack domain**. Authoritative servers may attempt to mitigate an attack in a number of ways, including dropping requests from recursive resolvers that transmit too many requests.⁵ In a counter play, attackers may utilize open resolvers to diffuse their traffic over a large IP space, maximizing transmission paths and reducing the likelihood of blocked traffic. Some attackers utilize spoofed⁶ IP addresses in the request packets, which help both anonymize and diffuse traffic, hindering mitigation. An overview of the attack is shown in Figure 1.

This attack technique was first recorded in 2009 in China [17], and again in a large attack against the French Network Information Center, AFNIC, in October 2013 [2]. Beyond these two attacks, there is no indication the technique was actively used until February 2014 when it became a daily phenomenon that disrupted the global DNS and had particularly damaging collateral damage on Internet Service Provider (ISP) infrastructure [22]. During the early years, the attack used queries for pseudo-random subdomains of primarily Chinese-owned domains. The terms Random Subdomain attack and Slow Drip, or Water Torture, attack were derived from these characteristics. While targeted at authoritative name servers, these attacks are often more damaging to the

⁵This is known as Response Rate Limiting (RRL).

⁶IP Spoofing is the creation of Internet traffic with fake IP addresses.

Internet's middle infrastructure [20], including recursive resolvers within ISPs, as they were not provisioned for unusually high traffic volumes. Solutions like response-rate-limiting were designed for use at the authoritative server and are less effective at recursive resolvers where the distribution of client IPs may not trigger such protections even if they were implemented. For practical reasons, most previous analyses of these attacks focused on mitigation.

While the common domain for an attack is considered the attack domain, the nature of the attack creates a resource exhaustion of the authoritative name servers, not the attack domain or its Internet resources, such as web pages. Beyond potentially disabling the authoritative name server from answering legitimate queries, e.g., for their website, the attacked domain suffers no direct denial of service traffic from these particular attacks. From this perspective, the attack domain itself does not play a significant role in the attack, only the authoritative name server. It is also possible for the attack domain to be operated by the attacker as a mechanism to target a particular name server.

If a name server's resources become exhausted, it may be unable to serve any of its customers, regardless of the requested domain. In reality, many different types of cyber attacks may occur simultaneously, including ones that target the web servers of the attack domain, for example. In the famous Mirai attacks of 2016, a wide range of attacks were deployed simultaneously, and in series, against Dyn authoritative name servers, effectively crippling a large number of popular websites [3]. One of the types of attacks used was Slow Drip. In other cases, malware has been discovered that contains different DDoS mechanisms for domains known to be attacked via Slow Drip; it seems plausible that these attacks were used in conjunction with one another [14].

More precisely, an attack consists of queries for a very large number of non-existent FQDNs within a common **attack domain**. Queries that are part of the attack will have one or more subdomain labels. If we consider the attack domain as an SLD, *attack_domain*, then queries that are part of the attack traffic will have the form:

$$label_n.label_{n-1}...label_1.attack_domain. \quad (1)$$

The string $label_n$ is referred to as the **prefix**. The attack domain is an SLD, most often containing two labels. The **suffix** of an FQDN is the parent domain of that FQDN, determined by removing the prefix, and thus of the form $label_{n-1}...label_1.attack_domain$. Slow Drip attacks are described in the literature as a set of pseudo-random prefixes of a single suffix [4, 6]. Previous published examples of confirmed attack domains include 111f.com [4] and www.dafa888.com [2], each of which contained a single pseudo-random label.

2.3 Prior Work

The vast majority of literature on this specific type of DDoS attack has focused on detection and mitigation. This includes methods of determining an ongoing attack from the rate of queries [1] and the randomness of the FQDNs [21]. In Reference [1] the authors propose an algorithm for identifying ongoing attacks within a recursive resolver or name server, which leverages a fixed size cache to identify domains that are seen frequently and, for the case of Slow Drip attacks, suddenly frequently. Their solution for inline detection promises to require fewer resources and afford more timely identification than traditional approaches. In Reference [21], they propose a means to detect attacks based on the queried FQDN alone and introduce a set of text-based features for this purpose. Unfortunately, this approach made assumptions about the randomness in the traffic not found in real data and therefore is of limited use operationally.

Only a small percentage of the literature contains analysis of the attack mechanisms and their victims, and none has undertaken a global study over a long period of time. An early expository [12] studied several months of traffic from a personal network in 2014. The author found distinct signatures in the IP packets of an attack and demonstrated large-scale IP spoofing.⁷ They augmented statistics about the attacks with some investigation

⁷The primary source for this analysis was ICMP error messages, which allowed the author to make inferences about the attack over time.

of the attack domains, showing that the majority of these were Chinese-owned. Patterns within the characters of attacks were described in a blogpost [6] in 2014.⁸

In work that considered several DDoS techniques, Qihoo 360 [16] used unsupervised learning to study traffic observed at honeypot⁹ sensors to understand “who is being attacked by what botnet families under which c2 controllers with what set of attack parameters.” They began tracking DDoS botnet families in 2014 and covered different forms of DDoS attacks, including two that are DNS-based. Their approach used clustering on packet-level features. One of these clusters, which they call *dns_cls1*, is a random subdomain attack. The examples they provide are consistent with those found by Reference [12] and confirmed by Reference [4] as part of a single attack system.¹⁰

In Reference [4], a single DDoS attack system is studied in depth. First documented by Reference [12], this attack system, dubbed *ExploderBot* by the authors, was initially seen in February 2014. These attacks were observed almost continuously through mid-2016, after which their activity was interrupted by multiple long periods of silence. The analysis of Reference [4] concluded that:

- the attacks contained a single pseudo-random alphabetic label of length 1–16, which followed a precise pattern;
- the traffic was not generated by a consumer-device botnet;
- the packets utilized very-large scale IP spoofing;
- the attacks did not appear to favor open resolvers to diffuse traffic; and
- the attacks were the work of a single actor.

The vast majority of domains either had Chinese-operated authoritative names servers or were registered in China [4, 12]. We independently confirmed the pseudo-random attack pattern and the attack tempo seen in References [4], [12], and [6] using historic data from 2016–2017.

In the vast majority of reporting on Slow Drip DDoS attacks, the examples cited to demonstrate the attacks can be tied back to the ExploderBot system. The Mirai malware is also known to contain this technique, as documented in Reference [3] and uses a 12-character alphanumeric label. However, the Mirai malware does not transit the DNS normally, instead sending traffic directly to IP addresses configured in the malware, and we have not observed this traffic in our data. Beyond these two systems, little is known about the malware and actors that generate these attacks.

3 DATA AND METHODOLOGIES

This section describes the data used in our research, as well as limitations of that data and key methodologies. Our work focused on threat intelligence questions, which requires the ability to large-scale retrospective analysis. We use passive DNS for this purpose and attempt to counter the inherent issues of bias by incorporating independent sources collected at different points in the DNS resolution path. These data were then used to detect attacks using the technique described in Section 4.

3.1 Data Sources

This research relied on two independent accesses to *passive DNS* (pDNS) records, obtained by recording DNS query-response events at recursive resolvers and including over 12B passively observed records per day. The majority of the data was collected between large recursive resolvers and authoritative name servers. A smaller source included data from several open resolvers.¹¹ The data sources included recursive resolvers present on four

⁸We are able to verify that the attacks discussed in the blog are all related to the actor discussed in Reference [4].

⁹A honeypot is a server set up as a decoy to lure cyber actors and to detect, deflect, or study cyber attacks.

¹⁰In Reference [16], they also claim that these prefixes are consistent with the Elknot/BillGates malware.

¹¹These are recursive resolvers that will resolve queries for any client but are not publicly advertised.

continents. The smaller open resolver sources collected approximately 10M records daily. The data used were obtained over seven months, from June through November 2018, and again in the month of January 2019.

To build a dataset for analysis, we developed a statistical classifier, described in Section 4, to identify attacks within normal traffic. Attacks are detected using *events* composed of query-response pairs aggregated by the SLD. For each detected attack, we extract the following fields from all events related to the domain for that day, regardless of their resolution:

- UTC timestamp of observation,
- fully qualified domain name (FQDN, qname),
- query type (qtype),
- response code (rcode).

While References [4] and [12] used other elements of the IP and UDP header, such as the time-to-live and port fields, these were not available to us.¹²

3.2 Collection Bias and Noise

Collection bias, also called observation bias, is a potential source of problems for this type of research. Without access to the actual malware, we infer information based on observable data, which is inherently limited. While our data collectors are globally distributed and cover a range of network types, it is possible that large-scale attacks still do not pass our sensors or occur at volumes too low for our detector threshold. As we attempt to study attacks over time, our collection is also susceptible to changes in the underlying traffic due to variability in Internet routing, among other reasons.

We also have to contend with a wide range of *noise* in the data. DNS traffic is naturally prone to noisy data, as queries are passed through the resolution process largely unchecked until they reach the authoritative server. The multitude of applications running on endpoints can create illegitimate queries.¹³ Moreover, users can accidentally drive traffic through the DNS.¹⁴ In addition, while we expect a single attack query to create a single event in the data, recursive resolvers and other devices in the path of the query may replicate the query or even initiate new queries. In some portions of this analysis, successful resolution of a domain is considered noise in the sense that it further skews the data and makes analysis more difficult. The impact of noise on our analysis is largely dependent on the popularity of the domain, as more popular domains are more likely to produce noisy data. The problem of noise was addressed through normalization and considering only non-resolved queries.

3.3 Methodologies

Attack data were obtained through the creation of a statistical classifier, as described in Section 4, that identified anomalies in traffic within and across days. The classifier was tested on previously published attacks that occurred between January 1st, 2017, and May 31st, 2018, to ensure its reliability, and the results of newly identified attacks were manually analyzed. From the resulting detection data, events were summarized into a small set of fields for the purpose of analysis and feature engineering.

As detailed in Section 6, exploratory analysis of the resulting attack records led to a number of strong features that segregated attacks into several distinct groups. While a wide array of potential features were analyzed for their ability to distinguish attack algorithms from one another, the most promising were derived from lexical and time-series analysis, along with the variety of query types used during an attack. Anomalies in the attack traffic were discovered by measuring the attack events against a baseline of global DNS traffic over time. Character distributions proved to be particularly useful in separating attacks. To take advantage of this, archetypal attacks

¹²UDP is the User Datagram Protocol used most often to encode DNS queries for transmission.

¹³Google Chrome is known to create unresolvable pseudo-random queries, for example.

¹⁴For example, a user might mistype a website name in their browser.

were identified through unsupervised machine learning and then used to create a distance measure used in more comprehensive clustering later.

Features were derived using three months of data and clustered as described in Section 7. Finally, to address the potential of transient features derived from the Summer 2018 data, the same features and cluster analysis was applied to January 2019 data and the results compared for consistency.

4 ATTACK DETECTION

Slow Drip attacks are traditionally easily detected at the authoritative name server and can sometimes be detected in live traffic at recursive resolvers with methods like those described in Reference [1]. However, our research is based on passive DNS analysis. Detecting attacks in passive DNS collection is made difficult by the scale and variety of global DNS traffic, as well as the nature of the attack. Traffic created by anti-virus products and content delivery networks (CDNs) has patterns that resemble a number of cyber attacks, including Slow Drip DDoS attacks. Attempts to identify attacks using simple measures on the FQDN, as in Reference [21], will fail in this environment, creating a large number of false positives, as will approaches based on measurements of traffic volumes within a day. To overcome these issues, we use both current and historical data in the detection algorithm.

The detector is a statistical classifier that identifies potential attacks retrospectively within daily traffic. It contains two stages that identify anomalous domain activity within a day of pDNS records, which is also anomalous in comparison to the domain's prior history. To trigger the detector a domain (SLD) must have an unusually high number of unresolved queries within a day and, additionally, the number of unique subdomains must be abnormal in comparison to the domain's recent past.¹⁵ To test performance, the classifier was run on historical data covering 10 days in 2017 and 2018 for which Slow Drip attacks were reported in Reference [4] and which the openDNS Twitter feed DNSStream contained alerts for Resource Exhaustion Attacks, generally synonymous for Slow Drip attacks in DNS. The results were manually reviewed for accuracy.

The detector first groups all of the events for a given data source by domain, or SLD, and counts the number of distinct unresolved queries.¹⁶ In the larger data source, the initial 12B records are reduced to approximately 225M domains. Given the distribution of unresolved query counts across this entire set of domains, we determine the quartiles, $Q1$ and $Q3$, that represent the number of events below which the bottom 25% of domains, and above which the top 25%, lie, respectively. The Inter-Quartile Range (IQR) is defined as $Q3 - Q1$, and *outliers* are those domains, d , for which the number of unresolved queries, $N(d)$, satisfies

$$N(d) > 1.5 * IQR + Q3. \quad (2)$$

This formulation identifies a set of domains that have an unusually high number of unresolved queries, with respect to all other domains on a given day. As concrete examples, on August 1st, 2018, the detector found unusually high query volumes for `uberinternal.com` and `icelandairlabs.com`. When compared to previous day's volumes, `uberinternal.com` had a 5,600-fold increase in activity and `icelandairlabs.com` showed no queries on the prior day.

Unfortunately, determining statistical outliers in this way is not sufficient for identifying Slow Drip attacks in passive DNS collection. Used alone, this will lead to a number of false positives, particularly resulting from anti-virus product and content delivery networks. To address this issue, we also compare the traffic with the recent past. Attacks often cross a date boundary, and detection will be missed if the comparison is only made with the previous day. However, it is important to compare recent activity to ensure that the large number of queries represents a true change. For this reason, we use a two-day separation for detection. The second stage of

¹⁵We have also observed these attacks against so-called wildcard domains, in which every subdomain is resolved. However, we do not consider that variation in this study.

¹⁶We consider all unresolved queries, not just NXDOMAIN responses.

the detector looks for a dramatic increase in the number of subdomains observed for domains raised during the first stage. This is accomplished by computing the change in the number of unique subdomains between the two dates for all second-level-domains, calculating the quartiles for the resulting distribution, and determining the outlier threshold analogous to that in Equation (2). Domains raised in the first stage, which are also anomalous in the second stage, are considered likely Slow Drip attacks. As a concrete example, `uberinternal.com` showed a 214k-fold increase in the number of subdomains queried between days.

These two stages create thresholds above which unresolved queries for a domain are considered a Slow Drip attack. The thresholds for an attack vary for each data source and are dependent on the normal volume observed in that source. The combination of anomaly detection with historical perspective serves to weed out false positives. In evaluating nearly 2K attacks over a six-month period, across multiple data sources, no false positives were found. However, this approach can have false negatives; that is, it might miss attacks for which we do not observe enough traffic.

5 ATTACK LANDSCAPE

To understand the overall attack landscape, meaning the tempo and general characteristics of Slow Drip attacks, we analyzed attacks discovered with the detector described in Section 4 over a six-month period (June–November 2018), where an **attack** was identified by a domain and a date. During this time, there were 1,949 separate attacks detected, covering 1,418 domains, with a median of seven attacks daily and eight days with over 25 targeted domains. On days with a large number of attacks, the attacks often contained groups of related domains. For example, we found 10 attacks on August 12, 2018, targeting domains related to Western Union. On other days, the relationship between domains was not readily apparent. The volume of attacks in the open resolver set is, as expected, much smaller, and with more variance. It includes 221 attacks covering 143 unique domains over the same time period.

We were particularly interested in finding attacks observable in independent data sources, including data from *open resolvers*. The use of open resolvers by attackers makes both detection and source identification of attacks more difficult by distributing the attack.¹⁷ Discovering the same attack across sources lends support to the accuracy of our detector and indicates the likely breadth of an attack. Nearly 100 attacks and a total of 48 domains were observed in multiple data sources that included at least one open resolver. This overlap signifies an actor tactic to leverage open resolvers, representing nearly 25% of the attacks in the sample. The domains in the intersection are both unsurprising, such as `airbnb.com`, and curious, such as a series of domains beginning with `sc-` seen in June 2018 that appear to have no legitimate purpose or common relationship.

In spite of a very strong signature for ExploderBot Slow Drip traffic, we found no attacks of this kind during the evaluation period. We detect this actor consistently through historical data, and so while it remains possible that our data sources are simply not observing the attacks, it appears more plausible that the actor has remained quiet through the latter half of 2018.

This research focused on the malware, specifically the generation of subdomains used in the attacks, and we did not do an in-depth analysis of the victims. We did find a prevalence of high-profile domains, such as those owned by AirBnB, Sony, DotDash, and Toyota present in the attacks. In the final clusters, described in Section 7, we see groups of well-known domains such as these, as well as clusters containing Pharmaceutical and Banking Industry domains. At the same time, almost 38% of the domains were not in the top 1M most commonly observed domains within our datasets, and many showed little sign of normal traffic. Some examples of these domains include `nspk.ru` and `verimi.de`. In these cases, the motivation for a DDoS attack is particularly unclear. In contrast to observations of Slow Drip attacks in 2014–2015, we did not find a preponderance of Chinese domains and name servers.

¹⁷These resolvers are not publicly announced as recursive resolvers, but are easily detected by daily scanning of the Internet by organizations such as Censys.

Table 1. Features Used for Clustering Attacks, Including the Number of Features of Each Type and a Description

Count	Name	Description
1	<i>label_ratio</i>	percentage of unique labels in attack
1	<i>overlap_ratio</i>	percentage of label overlap with DNS enumeration dictionary
6	<i>qtype = n</i>	percentage of query type <i>n</i> records in attack, <i>n</i> in References [1, 5, 12, 15, 16, 28, 33, 43]
1	<i>active_min</i>	number of minutes where NXDOMAIN events greater than median of all events per minute
1	<i>lex_ratio</i>	estimated percentage of lexicographically ascending subdomains
9	<i>unigram_distn</i>	J-S Distance from Unigram Char0 Distribution for these attacks as (date, <i>attack_domain</i>): (2018-08-13, shave.io), (2018-08-19, informatica.com), (2018-07-24, toyota.com), (2018-07-31, ecooltra.com), (2018-08-07, centralnic.com), (2018-07-12, verimi.de), (2018-08-09, norwichpharma.com), (2018-08-02, ruckuslbs.com)

6 ATTACK FEATURES

Without a copy of the malware, it might be impossible to fully determine how specific attack traffic is generated. Given that attack generators are software written by individuals, the hope is that the traffic carries the fingerprints of the designer and is distinct enough that it can be correlated across attacks.¹⁸ We use **unsupervised learning** techniques to group the attacks into subsets that have similar characteristics. We do not know in advance the true nature of the groups, or clusters, and the algorithm learns these from the data itself.

In **clustering**, we identify characteristics of the data, called **features**, that help group subsets of data and isolate groups from one another. A feature in the data is considered a *strong feature* if it can reliably separate groups of data; and a *weak feature* if it demonstrates a measurable characteristic for separating data but cannot be used alone to do so. A number of weak features might be combined to reliably separate data. These terms are subjective, but useful in understanding how features help in the machine learning process. **Feature Engineering** is the iterative process of identifying, evaluating, and selecting features. Successful use of machine learning to solve problems is highly reliant on proper feature engineering, much more so than the choice of the machine learning algorithm itself. Feature engineering is also a laborious process, and researchers often rely on previous work to obtain features for their experiments.

In the case of Slow Drip attacks, little insight into the attack features exists in the literature. As a result, we undertook the feature-engineering process from scratch. Using a set of 435 attacks identified between June and August 2018 and over 1B records, we identified a set of 20 strong features. In the remainder of this section, we first review the results of our exploratory data analysis as a high-level description of the data and then cover several of the more complex features in depth, independently. The final feature set is summarized in Table 1.

6.1 Exploratory Data Analysis

During the initial review of the data, it was readily apparent that the structure of the more recent attacks was quite different than that previously reported. This section details our initial exploration of the data and describes the observations that led to our ultimate feature set. We expected the data to contain random subdomains of an attack domain, that is, queries for non-existent FQDNs of the form

$$random.attack_domain, \quad (3)$$

¹⁸Two different pieces of malware may share the same generator, particularly if it is simplistic, and two different cyber actors may utilize the same malware for different attacks.

Fig. 2. The distribution of first characters seen in an attack on hfax.com shows that prefixes most often started with digits.

Fig. 3. The distribution of first characters seen in an attack on dollarshaveclub.com shows that prefixes started with letters.

where *random* is a single pseudo-random label and *attack_domain* is an SLD. It is also commonly described, as in Reference [6], as containing a second label, leaving queries of the form

$$random.fixed_label.attack_domain. \quad (4)$$

The data we studied, however, contained what appeared to be dictionary terms, rather than, or in addition to, pseudo-random values in the subdomain labels. Additionally, many of the attacks contained FQDNs with several labels, in stark contrast to the single pseudo-random label noted in the literature. Here are a few examples:

- ent254.sharepoint.hp.com,
- passyourdrugtest.airbnb.com,
- ltrefgt.byairbnb.com,
- rtx.bjbgp.bjbgp.bjbgp.91y.com,
- godoid-028.prod.ap-northeast-1.int.vidible.tv,
- www.diskgas.api.csd.bitwala.com.

Even this limited set of samples reveals the problems experienced by threat analysts in evaluating suspected Slow Drip traffic.

When considered in aggregate, the prefixes in the sample data had strong characteristics that distinguished one attack from another and were unlike pseudo-randomly generated data. This feature was quantified by creating **unigram character distributions** for each attack. All of the prefixes observed in an attack were treated as a set of individual characters, called a unigram, and the number of times each character was present in the set was counted, creating a distribution. The resulting count distributions can be compared, both visually and statistically, for similarity. As an example, two sample character distributions are shown in Figures 2 and 3. These distributions, derived from millions of observations, are very clearly distinct. The process of generating unigram character features for use in clustering attacks was a multi-stage process and is detailed in Section 6.2.

Also notable was the large number of labels commonly seen in normal DNS traffic, such as prod or cloud. A single attack often included both common terms and a strange assortment of uncommon labels. This exemplified another challenge for threat analysts in reviewing attack queries. In an effort to determine whether these attacks utilized the same generators, we extracted very long labels that were present across multiple attacks. Doing so led to a handful of surprising labels, e.g., 007dapotianmuweijiaomenduchang, common to several seemingly independent attacks, and the hypothesis that an unusual dictionary was used in the construction of some attacks. The resulting feature is described further in Section 6.3.

Time-series analysis, that is, the analysis of attack events considered in time order, proved one of the most fruitful approaches to explore features within the sample data. Visualization of various summary statistics of the events illuminated distinguishing characteristics between attacks. As an example, some attacks contained FQDNs with a wide range of prefix lengths, and the proportion of each length remained relatively fixed over time, while other attacks contained only prefixes of a fixed length at any given time, as seen in Figure 8. Most prevalent among these characteristics was the lexicographic ordering of the FQDNs, when considered as a series of strings ordered in time, in many attacks. Deriving a quantifiable measurement of that feature proved challenging, however, as seen in Section 6.4.

In the remainder of this section, we detail how these observations were converted to clustering features.

6.2 Unigram Character Features

To take advantage of the similarity, and differences, observed in character distributions between attacks, we designed a process that allowed us to measure the difference between **unigram character distributions** for the prefixes in each attack against a set of fixed distributions, creating a feature that could be computed reliably over time. Given two distributions, the **Jensen-Shannon (J-S) distance** [7] is a distance measure with a real number value in the interval $[0, 1]$, where smaller values indicate that the two distributions are more similar.¹⁹ To create a feature to be used in the final clustering of attacks, we first used the Jensen-Shannon distance alone to identify a set of **archetype attacks**. The distance between two attacks is defined as the J-S distance between their respective distributions. For each of the 435 attacks, unigram character distributions were computed and the pairwise Jensen-Shannon (J-S) distance between these distributions calculated. We hypothesized that attacks created by the same malware, or attack generator, would have similar character distributions and therefore small J-S distances. Computing the pairwise distance between all attacks over time, however, is unwieldy, at best, making *pairwise* distance an unreasonable feature. Instead, we clustered the initial attacks using only this distance measure and identified a small set of attacks that represented large clusters. The specific process for this feature engineering is described in detail below. This approach does not attempt to represent all attack generators, but instead provides a statistical feature that measures the distance from a given attack to a handful of other fixed attacks. We have not seen this approach of converting distribution similarities into a distance metric that is then utilized as a clustering feature described elsewhere in the literature. This section elaborates on the process used to select archetype attacks.

We considered two different unigram distributions for each attack. The first was the overall unigram character distribution of the set of unique *prefixes*, or *label_n*, over all FQDNs observed in an attack. We refer to this as the **overall character distribution**. We also considered the distributions of only the first character of each prefix within an attack. We refer to these distributions as **char0 distributions**. We then consider the underlying dictionary of characters observed across all the attacks. A total of 66 distinct characters were present in the labels across all attacks.²⁰

DBSCAN [8] was used to cluster the attacks via pairwise distances. This algorithm was chosen for clustering attacks by character distribution, because it does not require the number of clusters to be selected *a priori* and it identifies *outliers*. Each cluster is composed of *core points* and *boundary points*. The algorithm has two parameters, *eps*, which specifies the distance the algorithm searches for additional cluster members from core points, and *min_points*, the number of points that must exist within epsilon range for a point in the cluster to be considered a core point. Any point that is further than *eps* away from all others is an *outlier*. In the application of DBSCAN to the unigram character distributions, the **distance between two attacks** is the J-S distance between the associated character distributions.

¹⁹The Jensen-Shannon distance between two probability distributions is the square root of the Jensen-Shannon divergence of the distributions.

²⁰Many of these characters are not valid in DNS, however, the queries will still transit the DNS for resolution.

Fig. 4. Clusters of attacks based on the J-S Distance between the char0 distributions.

The goal was to identify tight clusters that might represent attacks created by the same attack generator. For this reason, we did not attempt to account for all of our data and were unconcerned about outliers.²¹ The most central distribution within a cluster became the representative, or **archetype**. The feature for use in more comprehensive clustering was the distance of each attack from a small set of archetype attacks.

This approach proved very successful. While the clustering was based on attacks, uniquely identified by a date and domain pairing, the clusters *effectively* represented domains, meaning that a given domain, even if attacked on several days, was usually found in only one cluster. This is consistent with attacks being created by the same generator, or set of generators, for a given domain over time. In other words, if we observe attacks against a domain such as `uberinternal.com` or `airbnb.com` over several distinct days, the prefixes generated have similar character distributions.

The overall unigram distributions led to four clusters, and using the char0 distributions, we found eight clusters containing more than 10 attacks. A visualization of the attacks, clusters, and the points that were not labeled, is found in Figure 4. Attacks that could not be clustered are labeled in the figure as cluster -1. For each cluster the archetype attack domain is noted.

These clusters also picked up sets of *related* domains, which again is consistent with a single attack generator. As an example, seen in Figure 6, one cluster includes numerous domains owned by the company AirBnB, e.g., `atairbnb.com` and `withairbnb.com`, but also includes the AirBnB content-delivery-network

²¹Outliers in clustering algorithms are subjective and a result of thresholds applied to define cluster boundaries. An outlier is any point that lies outside of all cluster boundaries and can be thought of as points that are dissimilar, according to the features used, from all others.

Fig. 5. Sample cluster derived from the char0 distributions. This cluster contains a distinct set of domains owned by AirBnB. These attacks have low overlap with the DNS enumeration list.

Fig. 6. Sample cluster derived from the char0 distributions. This cluster contains a distinct set of domains owned by AirBnB. These attacks have very high overlap with the DNS enumeration list.

(CDN), muscache.com. This same cluster contains numerous domains owned by the media company DotDash.²² Another cluster is dominated by a combination of Pharmaceutical Industry domains and domains registered in Iceland. Each of these clusters represents attacks over numerous days within the approximately 60 days represented by the data sample. Two clusters contained domains owned by AirBnB, but that did not overlap; see Figures 5 and 6.

There are a few obvious downsides to this approach of feature engineering. First, we made an assumption that while the volume of data we observe may vary widely between attacks, our collection perspective relative to the attack generators is unchanged, and therefore the character distributions from the same attack generator will be relatively consistent over time. This is a useful but flawed assumption that could lead us to separate attacks that are created by the same generator. Second, we treat each set of distributions as derived from a single generator. If multiple generators are used during an attack, each generator will contribute to a mixed character distribution. The expected result is that, in that case, we observe more clusters than attack generators that truly exist.

Considering character distributions at-scale confirmed our suspicion that most attacks in our dataset do not use a pseudo-random prefix generator. In nearly 450 attacks, only 8 of the attacks had distributions similar to a uniform random distribution. Even in those cases, it was not apparent that the prefixes were pseudo-randomly generated. Attacks on the domain `nspk.ru` proved to be constructed from incrementally adding characters sorted in lexicographical order. These attacks included queries for domains such as `a.nspk.ru`, `b.nspk.ru`, `...aa.nspk.ru`, `ab.nspk.ru`. While the character distribution was the same as a pseudo-random generator, the attack was clearly not created that way. An attack on the domain `verimi.de` observed on July 6, 2018, exhibited a nearly uniform random distribution and was used for an archetype attack for clustering.

Having identified promising clusters, the most central attack within each of these clusters is determined by the shortest distance to every other attack within the cluster. The associated distribution for this central attack is the representative for the cluster, and the attack is the archetype for that cluster. In the final clustering, one feature is generated for each of these archetype representatives: the Jensen-Shannon distance between an attack and the archetype. We also included the distance to the `verimi.de` attack as a uniform random representative. The distance between cluster centers is shown in Appendix Figure 14.

6.3 DNS Enumeration

During exploratory data analysis, there was a surprising combination of relatively common labels and seemingly unrelated and uncommon labels within the same attack. We isolated long labels, greater than 20 characters, which

²²Formerly about.com, DotDash domains include thoughtco.com and lifewire.com, among many others.

were present in multiple attacks. These included a strange mix of dictionary-type labels, composite terms, and multiple languages, such as

- mobilebusinessapplicationdevelopment,
- 007dapotianmuweijiaomenduchang, and
- caracepatmembuatwebsitegratis.

Surprisingly, a Google search for one of these terms led to a GitHub repository [19] containing a list of over 420K terms culled from DNS, a so-called **DNS enumeration list**. This kind of list is generated by cyber actors to assist in reconnaissance and formulating attacks, but the use has not been previously reported in Slow Drip attacks.

We hypothesized one of the attack generators utilized a DNS enumeration list as a dictionary, creating FQDNs during an attack by picking words from the dictionary and continually appending them as labels. For each attack the intersection of its labels and the GitHub list was computed. Over 15% of them drew 80% or more of their labels from this dictionary, all but confirming that some kind of DNS enumeration list was used as a dictionary for an attack generator, given the volume of unique labels found in attacks. A significant number of attacks contained 30%–60% overlap, as seen in Figure 9, which could be consistent with either multiple attack generators in use for those attacks or a larger underlying dictionary.

Manual inspection of the domains found in each grouping showed correlation with the clustering based on character distributions. In particular, the two clusters generated by the Jensen-Shannon distance between character distributions containing AirBnB domains correlated to a high and low inclusion of the DNS enumeration terms in their labels, as noted in Figures 5 and 6.

The overlap with the DNS enumeration dictionary is not notable in any of the attacks found in data from open resolvers. It is possible that actors using a DNS enumeration dictionary are not leveraging open resolvers in their attacks.

6.4 Time-series Features

We found during the exploratory data analysis phase that when considered as strings in time order, the FQDNs used in an attack were often in **lexicographical order**, that is, a query for aa-dev.airbnb.com would be transmitted prior to one for ab-dev.airbnb.com. However, capturing this observation into a numerically accurate feature is difficult. For example, an attack that is generated by drawing prefixes from a sorted dictionary and transmitting traffic from multiple locations will be observed as an interleaved set when aggregated in time. The more complex the generator and the more sources of traffic, the more obscured the original ordering becomes. A number of factors could influence ordering, in particular:

- Observations are not directly from the attack generators and may contain the interleaving of the attack traffic of several devices.
- Internet routing may affect the timing of events at sensors and result in observations that are out of order from their original creation.
- The sensors themselves are often part of load-balanced systems, which further diffuse arriving traffic.
- Additional queries may be created by Internet appliances along the routing path, adding noise to the data.
- A large number of events occur every second, the smallest granularity of our observation period, and we are unable to determine the order in which those are generated.
- It is extremely unlikely all events are observed.

To address these challenges, we **approximate the lexicographic ordering** by drawing a single event for each **suffix**²³ per second, creating a time series for each suffix. The suffix is fixed, because exploratory analysis led to

²³Recall, the suffix of an FQDN is found by removing the prefix *label_n*.

Fig. 7. An approximation of the lexicographic ordering of events within attacks estimated by considering a single FQDN per suffix per second.

Fig. 8. Prefix lengths observed over time for an attack against hfax.com. For a fixed suffix, the prefix lengths are initially quite variable, but transition to length three, and then length four, for extensive time periods.

the hypothesis that subdomains may be created incrementally by increasing the number of labels in the FQDN. Given a fixed suffix, and a sample for each second, we calculate the percentage of these that are in order. For each attack, the average of this percentage over all suffixes is an approximate measure of the lexicographic ordering. We represent this as a ratio of events in the interval $[0, 1]$, as shown in Figure 7. Somewhat surprisingly, almost no attacks contain randomly ordered prefixes. A majority of attacks have a strong tendency toward lexicographic ordering, while 27 of the attacks demonstrate reverse ordering.

Another feature drawn from the exploratory data analysis was the length of each attack. While most attacks were relatively short, lasting under 15 minutes, several attacks lasted hours and some nearly a full day. Obtaining an accurate measure of the attack length was hindered by noise in the traffic and skew in the distribution of events per minute for each domain. Popular domains are likely to have a constant stream of unresolvable queries, due to factors such as user typos and the use of content delivery services. Measuring an attack by the number of minutes for which there were more unresolvable queries than the mean, a common benchmark statistic, led to inaccurate results given these factors. We found the number of minutes in which attack events exceeded the median number of non-existent subdomains (NXDOMAIN events) measured over the full day more accurate. This value was used as the **attack length** feature in clustering.

Time-series analysis indicated there was a reasonable likelihood that multiple attack generators were used in some attacks. The number of labels and variety of prefix lengths observed per second were illuminating. In a small subset of attacks there were distinct periods of variable behavior. An example of this is shown in Figure 8, in which a period of widely varying prefix lengths is followed by long periods, with far greater activity, containing a single prefix length and a smooth transition between lengths. Quantifying these behaviors over all attacks, however, would be extremely difficult, and these observations were not used in the final clustering algorithms.

6.5 Qtype Features

Another striking observation from exploratory data analysis was the use of multiple DNS query types (*qtype*) during an attack. This has not been reported previously. As mentioned in Section 2.1, while DNS is most recognized for its translation of domain names to IP addresses, it can deliver a multitude of other types of information. The type of data sought is included in the query. For example, a query for type “MX” would request the mail server domain, rather than the IP address of a domain. While IPv4 address requests, called “A” records, are by far the most common request seen in DNS traffic, the protocol allows for 65,536 different data types. Other commonly seen query types include the IPv6 address, type “AAAA,” and the authoritative name server, type “NS.” Slow Drip attacks are known to request IPv4 (A) records [4].

Many attacks included queries for less common query types, and these occurred in distinct combinations. Six record types were often simultaneously utilized, though some variations existed. These six record types were

Fig. 9. Percentage of labels per attack that overlap with GitHub DNS enumeration list.

Fig. 10. Correlation of 20 features used for clustering.

query types: 1 (IPv4 address), 5 (CNAME), 12 (PTR), 16 (TXT), 28 (IPv6 address), and 33 (SRV). Of these, PTR, CNAME, and SRV are the most unusual. We generated **qtype features** for each of these six by computing the percentage of events in an attack of each type.

We also observed large queries for a handful of other unexpected qtypes. A comparison of attacks using multiple qtypes to those with only “A” record queries is shown in Appendix Figure 15. Both the combination of query types and the subdomains found within each type were rare within normal DNS traffic. From an attacker’s perspective, this variety of query types serves to further diffuse the attack and ensure that the majority of requests are forwarded to the authoritative name server.

6.6 Feature Correlation

The final feature set, provided in Table 1, considered the uniqueness of the domain labels, their overlap with the DNS enumeration dictionary,²⁴ the distance between unigram character distributions and nine established distributions (including one uniform random distribution), the length of the attack, the variety of query types (qtype), and an estimate of the lexicographic ordering of the queries. The correlation of the feature set is shown in Figure 10.

7 FINAL CLUSTER ANALYSIS

An unsupervised learning technique, *hierarchical density-based spatial clustering of applications with noise (HDBSCAN)* [5], an extension of DBSCAN [8], was used to cluster attacks based on the features in Table 1.²⁵ None of the features consider either the date of the attack or the attacked domain itself, which are the two components used to uniquely label an attack. Thus, the clusters describe attacks that have generative features in common and represent attacks that are likely created by the same attack generator. This clustering was performed on the set of 435 attacks found from June–August 2018 from the larger data source and independently applied to the attacks found in the open resolver set for source comparison. Finally, the same features were used to cluster data from January 2019 to ensure continued relevance of the features and clusters.

7.1 Results

The clustering revealed nine large groups of attacks, accounting for 279 of the 435 attacks. These clusters and their dominant features are described in Table 2. Using uniform manifold approximation and projection (UMAP) [9, 18], we can reduce the features into two-dimensional space for the purpose of visualization. This projection allows us to see how similar attacks are with one another, as seen in Figure 11. Attacks that could not

²⁴Here, we mean the list found on GitHub discussed earlier.

²⁵A minimum cluster size of 10 was used.

Table 2. Clusters of at Least 10 Attacks within a Dataset of 435 Attacks Selected over 60 Days Using 20 Features in Table 1

No.	Size	Days	Sample Attack Domains	Dominant Features
0	20	6	AirBnB-related, Ruckus Wireless	high DNS overlap, five qtypes, lex-ascending, moderate attack lengths
1	11	5	DotDash-related	redundant labels, one qtype, high DNS enumeration overlap, short attacks
2	28	5	Pharmaceutical Industry and Icelandic	unique labels, lex-descending, five qtypes, moderate length attacks
3	46	28	Indonesian domains, universities, payment sites, Tinder	very lex-ascending, five qtypes
4	47	19	Toyota, Starbucks, Paypal	high label reuse, high DNS overlap, one qtype
5	10	9	dollarshaveclub.com, shave.io	no DNS enumeration overlap, one qtype, strong unigram pattern, shorter attacks
6	27	8	starting sc-, AirBnB-related, PlayCanvas	moderately duplicated labels, one qtype, no DNS enumeration overlap, lex-ascending, strong unigram pattern
7	55	12	Western Union, Harry's	~50% label reuse, one qtype, short attacks, high lex-ascending
8	11	6	crowdshield.com, coupa.com	one qtype, high lex-ascending, far from unigram models

be clustered automatically are in cluster -1, indicated as pale blue points. From the plot, we can see that although the clustering algorithm was unable to include these points into a specific cluster, in most cases they appear likely to belong to one of the named clusters. Were the attack generators highly variable, we would observe isolated, or outlier, points scattered across Figure 11.

These clusters are similar, but not identical, to those found using only Jensen-Shannon distance between unigram character distributions, as evidenced by two of the unigram attacks falling into the same cluster, and two others as outliers. Each cluster crosses a range of dates, but often contains obviously related attack domains. Several smaller clusters were also apparent, and if we allow clusters as small as five attacks, the clusters represent 293 of the attacks. We suspect that a large number of attacks are created by a single system, but also that two generators are in use at the same time, hindering easy identification.

The clustering of attacks found within open resolver data is also illuminating. If we consider attacks that are observed at *both* multiple data sources, we find that they almost entirely fall within a single cluster, cluster 6 in Table 2, represented by playcanvas.com. Broadening the scope to include common attack domains, and not specific attacks, the open resolver attacks primarily fall into two clusters: cluster 6 and cluster 4.

Clustering results indicate that attack generators are limited to a relatively small number of algorithms. Multiple clusters shown in Table 2 indicate the use of several query types; moreover, they contain a combination of uncommon query types. We suspect that there are commonalities in these generators, even where their sub-domain features are distinct, based on this query type pattern. Combining the clustering results with subject matter expertise, there are likely four to eight attack generators in use that are able to create global-scale attacks on repeated days. The projection of the 20 features into two-dimensional space, Figure 11 is consistent with this conclusion.

Some attacks appear to be a composite of multiple generators. Figures 8 and 17 show two examples where we see abrupt changes in statistical features of the attack. Some of the attacks appear to be a composite of a DNS enumeration dictionary generator and another dictionary source. In addition, rapid changes in the characteristics

Fig. 11. Clusters in a sample of 435 attacks using features from Table 1 and labeled with sample domains. Visualization via UMAP [9] projection.

of the FQDNs, such as the prefix changes seen in Figure 22, were fairly common and often overlapping. Further, the projection of clusters in Figure 11 shows that many of the attacks that were not clusters lie in between those that were, indicating they might be a mix of the two cluster generators.

7.2 Resilience to Model Drift

A significant potential Achilles heel for machine learning is known as *model drift*, in which the underlying data landscape changes and a model that previously performed well begins to fail. In the context of Slow Drip attacks, this would be the case for any detection based on the assumption of pseudo-random subdomains, for example, as the threat landscape has evolved. For this reason, models need to be regularly checked for continued relevance to the environment. The original feature engineering and clustering was computed from data for attacks that occurred in June–August 2018. We extracted these same features from nine days of attacks in January 2019 to determine whether they remained valid. During this time there were 106 attacks detected on 83 domains. Far more attacks in this new dataset overlapped with the DNS enumeration list [19], with over 25% of the attacks drawing more than 80% of their labels from the list.

To evaluate the relevance of the features in Table 1 to the new attacks, HDBSCAN was used to cluster the new attacks, as well as the union of the two datasets. When considered alone, 93 of the new 106 attacks fall into two clusters, indicating that the features remained relevant for clustering the new data. When combined with the earlier dataset, 64 fell into previously defined clusters, indicating stability in the clusters over six months. Most of the attacks in the January 2019 data fall into cluster 0 in Table 2, while the others fall into cluster 6 and cluster 4.

8 CONCLUSION

The results of this research and development support both threat analysts and data scientists. An anomaly-based statistical classifier provides a means to detect these attacks with high reliability in both very large-scale and moderate-scale data environments. Feature engineering led to the discovery of notable characteristics of several generators, including the use of DNS enumeration dictionaries and unusual query type combinations, that can help analysts recognize patterns in putative attacks. It further confirmed their concerns that the threat landscape had changed, making the attacks harder to recognize. Unsupervised machine learning identified nine major clusters and the associated important features for each cluster. Moreover, the features remained relevant after six months.

There are two major components of this work as an application to cyber intelligence research and industry: the detector, which allows us to locate the activity in a sea of noise, and the clustering, which allows us to understand the activity at some higher level. As actors adapt naturally, and in reaction to public information about their techniques, will these approaches remain relevant? Assuming that the attacker still creates a reasonably large number of subdomains, beyond the normal activity for the domain and without expected resolution, the statistical detector will identify them. While described in this article as a daily detection, the observation window for determining anomalous volumes can be shortened for a more rapid discovery. Detection is ultimately dependent on the ability to observe a statistical change in volume, which is dependent both on the observation point and the diffusion of traffic by the attacker.

In contrast, the attack generation algorithms could be easily changed and no longer be similar to the clusters identified in Table 2. The attack could, for example, change dictionaries. When comparing new traffic with the established baselines, it will drift away and appear as outliers. This is similar to the shift of domain names generated by the original ExploderBot system to those more recently observed. The unsupervised learning of pseudo-labels, or groups, described here is intended largely as a retrospective analysis and a mechanism for identifying changes in the attack landscape. With regular retraining, changes in dictionaries, query types, or other characteristics should be identified, although they may ultimately require additional feature engineering. We are fortunate here that within this type of attack there are only limited dimensions in which the attackers can make changes: ultimately, they must make a large number of unresolvable valid DNS queries.

With 18 months of observed attacks as of late 2019, we have found the volumes to ebb and flow, but not disappear. During 2019, we observed a change in the attacks to include known malicious domains. A variety of exposed MageCart command-and-control domains were first observed in January 2019, and other domains, from different actors and exploits, appeared occasionally throughout the year. These trends highlight the uncertain motivation of the attackers and raise further doubt that the intention is a denial of service. Aside from the unusual victim domains, the low volumes, far lower than the original ExploderBot attacks, make successful denial of service unlikely. At the same time, the structured query names have little room to contain an information channel. Without access to the malware that is generating the traffic, the question of motivation may be left unanswered.

A APPENDIX: SUPPORTING VISUALIZATIONS AND ADDITIONAL FEATURES

We have included visualizations and further details in the Appendix that might appeal to certain readers interested in the feature engineering performed. In addition, this Appendix includes a number of features we evaluated in the course of our research but did not use in the final clustering.

A.1 Unigram Distributions as Features

Visualizing the similarity, or J-S distance, between distributions with a heat map confirms that this distance measure can be used to separate attacks. The distance between attacks is the J-S distance between the respective character distributions for each attack and is in the range $[0, 1]$. Figure 12 shows a heat map of pairwise distances,

with smaller distances represented by darker pixels. In this random sample of 150 attacks, subsets of similar attacks are seen as dark groupings along the diagonal.²⁶ Attacks that have dissimilar char0 distributions will have light-colored pixels. Considering this, in Figure 12, we see a few outlier attacks in the top left, some smaller groups of similar attacks, and a larger set that shows core similarity with broader distribution.

Fig. 12. Pairwise Jensen-Shannon distances between char0 distributions of prefixes in 150 randomly selected attacks, sorted to clustering of similar distributions.

Fig. 13. Relative pairwise Jensen-Shannon distance between overall unigram distributions of prefixes in 150 randomly selected attacks. While clusters exist, they are far less distinct than those found using the char0 distributions.

If we compare the char0 distribution similarities with the overall distribution similarities, we see a dramatic difference. The distance matrix for overall unigram distributions of the same set of 150 attacks is illustrated in Figure 13. Here, we observe much less variance in the distributions.

Recall from Section 6.2 that the attacks were clustered via their char0 distributions and eight clusters identified, the centers of which were considered *archetype* distributions. The distances between these archetype distributions is shown Figure 14. This allows the reader to see that the clusters are reasonably separated, though a few are closer and are possibly related distributions. As these distributions are later used in clustering all of the attacks, we want to observe some separation in the data. Our distributions include a nearly uniform random one, generated by an attack on `verimini.de`, that is furthest from all other distributions.

A.2 Query Types as Features

Most attacks contained either a single query type, qtype 1, or a specific, surprising combination of five query types. In Figure 15 a scatter plot shows the proportion of qtype 1 queries in attacks. Notice that there are two large groupings, but also a small set of attacks with a wide degree of variance.

Another observation of the exploratory data analysis was the transition in some attacks between the use of several query types to a single query type over time. This phenomena, shown in Figure 17, serves to hinder using the query type alone as a strong feature. In this case, six query types are used simultaneously for part of the attack, five in even proportions, and then a transition occurs to a single query type. This type of phenomena occurs in a small, but notable, portion of the attacks.

Nearly half of the attacks contained a substantial number of qtype 12 (PTR) records requests. This is particularly surprising, as PTR records are primarily used to locate a hostname from an IP address in a specially formatted FQDN. An example of the domains attacked using PTR records are shown in Figure 16. Notably, none of the AirBnB-related attacks contain PTR records, whereas attacks against universities, e.g., `rit.edu`, all contain PTR records.

²⁶This visualization also uses agglomerative clustering of the results to sort the attacks.

Fig. 14. The Jensen-Shannon distances between char0 cluster centers indicates good separation between the clusters.

Fig. 15. The relative type distribution between type 1 (“A”) record queries and other types over 444 attacks, including small jitter to the x- and y-coordinates to visually separate attack points.

Fig. 16. Domains attacked using PTR queries sized by the number of PTR records observed.

A.3 DNS Enumeration as a Feature

Those attacks containing a very large portion of GitHub DNS Enumeration dictionary [19] contained domains that were notably different than seen in other features. While `byairbnb.com` is present, most of the domains are less recognizable, and quite a few of the larger attacks are against Russian or Chinese domains. Figure 18 shows a word cloud containing the attacked domains sized by the percentage of overlap between attack prefixes and the DNS enumeration dictionary.

While the overlap with this specific list was present in a substantial number of domains, the true dictionary is unknown. We attempted to “bootstrap” from the attack labels to the underlying dictionary by taking the union of all labels from attacks with greater than an 80% overlap to the GitHub list. This did increase the dictionary size, but not substantially. More importantly, it did not impact the overall distribution of the overlap statistic.

Fig. 17. An example of the distribution of query types transitioning over time during an attack from six types to a single type.

Overlap with DNS Enumeration Dictionary

Fig. 18. Attacked domains, sized by their overlap with the percentage of prefixes found in the DNS enumeration dictionary.

Fig. 19. Distribution of attack lengths as measured by the number of minutes in which nxdo-main events exceeded the median for the attack domain.

Fig. 20. Maximum depth of labels observed during attacks.

A.4 Attack Length as a Feature

As seen in Figure 19, most attacks are quite short. This is consistent with other studies of DDoS attacks writ large [11]. Accurately estimating the length of an attack using a single computation across domains is difficult due to the variance in baseline traffic for each domain, which creates noise. The common approach of using a mean will be inaccurate due to the skew within the distribution of events per minute for various types of domains. We evaluated multiple measures and ultimately used the number of minutes for which NXDOMAIN responses for a domain were above the median level for that domain during the day. This measure tends to underestimate attack lengths for rare domains, e.g., hfax.com, and overestimate lengths of popularly domains, e.g., airbnb.com, but creates a reasonable relative measure. The vast majority of long-lived attacks are unpopular domains, as measured by global DNS requests over time.

A.5 Other Potential Features

We evaluated a wide range of features, many of which had promising characteristics, before settling on those used in this article. In this section, we provide an overview of some of the other features we considered and that might be relevant for understanding actors and attack generators.

A.5.1 Open Resolvers Usage. While we did not utilize the presence of an attack as an explicit feature in our clustering, the use of open resolvers by an attacker to diffuse their packets is a tactic that may serve to fingerprint the attack generator. Our results found that the attacks observed at an open resolver were all within one cluster in Table 2.

Fig. 21. The number of labels over time for one attack on airbnb.com.

Fig. 22. Prefix lengths over time during one attack on airbnb.com. The proportion of lengths remains relatively constant although the volume changes.

For attacks created through consumer-based botnets, there is little advantage to using open resolvers and there are evident risks.²⁷ In many commercial environments, DNS requests to external IP addresses will be blocked, thereby thwarting an attack emanating from these networks. Additionally, the space of open resolvers changes and the bot members need to be regularly updated. A simpler solution is to have bot members transmit requests through their normal recursive resolver. However, attacks generated from an actor-controlled infrastructure, such as cloud services provider, benefit from scattering the initial DNS requests across open resolvers without much risk.

Our focus in this research was on global-scale attacks, from which it is relatively difficult to discern how widely open resolvers are utilized given the natural transit of traffic through the Internet; most of our observations occur at a point where the connection with the initial request to an open resolver would be lost. A more in-depth look at all attacks found at various open resolvers may shed further light on generators.

A.5.2 Label Depth. While the Slow Drip is typically characterized as a single random label on a fixed attack domain, as shown in Equation (1), we found that current attack systems frequently used more labels. The *depth* of an FQDN is the number of labels it contains. If we remove the attack domain from each FQDN within an attack, we can consider how many labels, in terms of depth, were used in the attack. Instead of a single random label, we find the most attacks contain more depth, and that eight labels is particularly common; see Figure 20. The variance seen here indicates that this feature can be used to separate some of the attacks; however, overall, it is a weak feature. We also considered averaging the unique levels per minute to incorporate time, and alternate descriptive statistics such as the variance over time. As seen in Figure 21, however, this simple descriptive statistic can obscure strong time-series features.

A.5.3 Prefix Lengths. Another natural set of features to investigate surround the label or prefix lengths. The distribution of these lengths appears to vary with some consistency across attacks. We examined various descriptive statistics for the prefix lengths, including the standard deviation and co-variance per attack, but ultimately did not use these in our clustering.

A.5.4 Time-series Features. During our data exploration and feature engineering, time-series analysis again and again revealed characteristics in the data not otherwise visible. Some of these features are seen in the main article, including the lexicographic ordering in attacks. There were a number of time-related features that we did not capture in our clustering, but found particularly compelling. Some generators appear to use a fixed number of labels at any time, which changes over the attack, as seen in Figure 21, presumably from building attack queries by continually appending or removing, labels. In this particular example, the number of labels descends in time,

²⁷By consumer-based botnets, we mean botnets of compromised devices found in a home consumer environment, such as laptops, phones, and home routers, as well as Internet of Things devices such as alarm and lighting systems.

but in many other examples it is seen to ascend. This would be consistent with the use of a dictionary, which is used to select a label, and creating queries by continually appending labels.

In the same attack shown in Figure 21 on `airbnb.com`, we find that the relative distribution of prefix lengths remains relatively constant over time, even as the number of labels changes. In Figure 22, we see a condensed view of this effect; for space, we have limited the prefix length to 20. In reality, the longest prefix is 55 characters, the maximum length permissible. This is consistent with the use of an underlying dictionary.

Features considered to incorporate time-series elements of the attacks included averaging descriptive statistics such as minimum or maximum of different components, e.g., labels or prefix lengths, per minute over an attack, as well as variance and co-variance measures. While promising, these features proved more difficult to assess and scale given the volume of data and were not used in the final clustering. Given the strong time elements in DDoS attacks, improved approaches to leveraging time-series features at scale would likely prove valuable to understanding the malware and actors.

ACKNOWLEDGMENTS

The author is thankful for the insights of Infoblox Threat Intelligence Analysts that helped motivate and refine this work. Additionally, thoughtful reviews of the article by Mike Last, Emma Lanier, Chris Heckman, and Laura da Rocha, as well as anonymous reviewers greatly helped focus the content.

REFERENCES

- [1] Yehuda Afek, Anat Bremner-Barr, Edith Cohen, Shir Landau Feibish, and Michal Shagam. 2016. Efficient distinct heavy hitters for DNS DDoS attack detection. *arXiv e-prints*, Article arXiv:1612.02636 (Dec. 2016).
- [2] AFNIC. 2014. Random Qnames—Dafa888 DoS Attack, Presentation on October 2013 Attack on AFNIC. Retrieved from [ps://indico.dns-oarc.net/event/20/contributions/278/attachments/242/452/dafa888-DosAttack.pdf](https://indico.dns-oarc.net/event/20/contributions/278/attachments/242/452/dafa888-DosAttack.pdf).
- [3] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J. Alex Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. 2017. Understanding the Mirai botnet. In *Proceedings of the USENIX Security Symposium*. USENIX Association, 1093–1110.
- [4] Renée Burton and Cameron Switzer. 2018. *ExploderBot: A Slow Drip Attack System*. National Security Agency Technical Report. National Security Agency, Fort Meade, MD. DOI: <https://doi.org/10.13140/RG.2.2.32214.29760>
- [5] Ricardo J. G. B. Campello, Davoud Moulavi, and Joerg Sander. 2013. Density-based clustering based on hierarchical density estimates. In *Advances in Knowledge Discovery and Data Mining*, Jian Pei, Vincent S. Tseng, Longbing Cao, Hiroshi Motoda, and Guandong Xu (Eds.). Springer Berlin, 160–172.
- [6] Emmet Cassidy. 2014. A DNS Cache-Busting Technique for DDoS-style Attacks Against Authoritative Name Servers. Retrieved from <https://blog.cloudmark.com/2014/10/07/character-analysis-of-slow-drip-attacks>.
- [7] D. M. Endres and J. E. Schindelin. 2003. A new metric for probability distributions. *IEEE Trans. Inf. Theor* 49, 7 (2003), 1858–1860. DOI: <https://doi.org/10.1109/TIT.2003.813506>
- [8] Martin Ester, Hans-Peter Kriegel, Jiirg Sander, and Xiaowei Xu. 1996. A density-based algorithm for discovering clusters in large spatial databases with Noise. In *Proceedings of the SIG-KDD*.
- [9] L. McInnes and J. Healy. 2018. UMAP: Uniform manifold approximation and projection for dimension reduction. *ArXiv e-prints* 1802.03426, 2018.
- [10] Bert Hubert. 2018. Herding the DNS Camel. Retrieved from <https://www.ietf.org/blog/herding-dns-camel/>.
- [11] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. 2017. Millions of targets under attack: A macroscopic characterization of the DoS ecosystem. In *Proceedings of the Internet Measurement Conference (IMC'17)*. ACM, New York, NY, 100–113. DOI: <https://doi.org/10.1145/3131365.3131383>
- [12] Michael Joost. 2014. About DNS Attacks and Destination Unreachable Reports. Retrieved from <https://michael-joost.de/dnsterror.html>.
- [13] Anestis Karasaridis. 2012. *DNS Security: In-depth Vulnerability Analysis and Mitigation Solutions*. Amazon.com Services LLC, ASIN B007ZW50WE. <https://www.amazon.com/DNS-Security-depth-Vulnerability-Mitigation-ebook/dp/B007ZW50WE>.
- [14] Antiy Labs. 2017. A Description of Magic Ferret, a.k.a. Weasel, DDOS malware. Retrieved from www.antiy.com/response/weasel.html.
- [15] Cricket Liu and Paul Albitz. 2006. *DNS and BIND*. O'Reilly.
- [16] Ya Liu. 2016. Improve DDoS Botnet Tracking with Honeypots. Retrieved from <https://www.botconf.eu/botconf-2016/>.
- [17] Ziqian Liu. 2009. Lessons Learned from May 19 China's DNS Collapse. Retrieved from https://www.dns-oarc.net/files/workshop-200911/Ziqian_Liu.pdf.

- [18] Leland McInnes, John Healy, and James Melville. 2018. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv e-prints*, Article arXiv:1802.03426v2. (Dec. 2018).
- [19] n00py. 2018. All Wordlists for Every DNS Enumeration Tool... Ever. Retrieved from <https://gist.github.com/n00py/2cba6990e4dacc52c5536346338f6f1e>.
- [20] Kei Nishida. 2015. Water Torture: A Slow Drip DNS DDoS Attack on QTNt. Retrieved from <https://www.slideshare.net/apnic/dnswatertortureonqtnet-1425130417-1425507043>.
- [21] Yuya Takeuchi, Takuro Yoshida, Ryotaro Kobayashi, Masahiko Kato, and Hiroyuki Kishimoto. 2016. Detection of the DNS water torture attack by analyzing features of the subdomain name. *J. Inf. Proc.* 24, 5 (2016), 793–801. DOI : <https://doi.org/10.2197/ipsjjip.24.793>
- [22] Ralf Weber. 2014. Latest Internet Plague: Random Subdomain Attacks. Retrieved from <https://www.youtube.com/watch?v=BDa2akVgbLg>.

Received April 2019; revised December 2019; accepted December 2019