

No Time for Downtime: Understanding Post-Attack Behaviors by Customers of Managed DNS Providers

Muhammad Yasir Muzayan Haq
University of Twente
Enschede, The Netherlands
m.y.m.haq@utwente.nl

Mattijs Jonker
University of Twente
Enschede, The Netherlands
m.jonker@utwente.nl

Roland van Rijswijk-Deij
University of Twente
Enschede, The Netherlands
r.m.vanrijswijk@utwente.nl

kc claffy
CAIDA / UC San Diego
La Jolla, CA, USA
kc@caida.org

Lambert J.M. Nieuwenhuis
University of Twente
Enschede, The Netherlands
l.j.m.nieuwenhuis@utwente.nl

Abhishta Abhishta
University of Twente
Enschede, The Netherlands
s.abhishta@utwente.nl

Abstract—We leverage large-scale DNS measurement data on authoritative name servers to study the reactions of domain owners affected by the 2016 DDoS attack on Dyn. We use industry sources of information about domain names to study the influence of factors such as industry sector and website popularity on the willingness of domain managers to invest in high availability of online services. Specifically, we correlate business characteristics of domain owners with their resilience strategies in the wake of DoS attacks affecting their domains. Our analysis revealed correlations between two properties of domains – industry sector and popularity – and post-attack strategies. Specifically, owners of more popular domains were more likely to re-act to increase the diversity of their authoritative DNS service for their domains. Similarly, domains in certain industry sectors were more likely to seek out such diversity in their DNS service. For example, domains categorized as *General News* were nearly 6 times more likely to re-act than domains categorized as *Internet Services*. Our results can inform managed DNS and other network service providers regarding the potential impact of downtime on their customer portfolio.

Index Terms—managed DNS, availability, DDoS mitigation, network management

1. Introduction

The domain name system (DNS) is a core component of the Internet. Its primary function is to translate human-readable domain names to IP addresses. As such, e-commerce with a web site is no longer possible if the authoritative name servers for the web site’s domain name fail to function. A managed DNS service provider offers premium services to manage DNS traffic [1], improving

reliability and performance, including potentially offering DDoS protection. Failures of managed DNS services adversely affect parties that rely on these services [2], depending on their degree of reliance.

Businesses that use a managed DNS service may adopt a more resilient strategy, i.e., use multiple DNS service providers, that ensures higher availability even when that managed DNS service provider may be unavailable. Previous studies [1], [3] have shown that when services of Dyn, a large managed DNS service provider, became unavailable in 2016 [2], many but not most domains adopted this more resilient strategy to counter the effects of downtime. Some businesses value continuous online availability enough to pay the premium charged by managed DNS service providers. It depends on the expected return on investment of such high availability of their online services. Different businesses value high availability differently, but outages clearly sometimes directly induce a loss of revenue [4].

We perform an analysis of large-scale longitudinal active DNS measurement data to identify domains that responded to the Dyn DDoS-induced outage. We map these domains to their popularity rank using historical Alexa rankings. We also obtain industry-sourced data on industry sector and risk characteristics using McAfee’s URL categorization service. Overall, we make the following contributions:

- We characterize 106K domains under .com, .net, and .org affected by the Dyn outage based on four aspects: post-attack resilience strategy, popularity, industry sector, and risk characteristics.
- We observe that domains with higher popularity (based on their Alexa ranks), more commonly adopted a more resilient strategy. We find that nearly 40% of Dyn customers in the top 100 of Alexa 1M list re-acted after the attack while in the top 500K only $\approx 15\%$ re-acted.
- Some industry sectors had a more significant correlation with post-attack response of domain managers. For example, the probability of a domain owner to adopt a more resilient strategy after the attack was nearly 9 times more if the domain belonged

This work is part of the NWO: MASCOT project, which is funded by the Netherlands Organization for Scientific Research (CS.014). This material is based on research sponsored by the National Science Foundation (NSF) grant OAC-1724853 and OAC-2131987. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF.

to *General News* category (compared to categories which did not re-act). Some industry sectors (such as *Religion*) did not have a statistically significant impact on the decision of domain owners.

- We show differences in the post-attack responses across domains categorized by their risk characteristics, based on McAfee’s URL categorization service. Nearly, 9% of the domains categorized as malicious used an additional DNS service provider after the attack, while only 1% of domains categorized as low risk took this strategy.

Our results empirically show that factors such as industry sector and popularity of domain name influence the resilience strategy adopted by domain managers in the face of unplanned unavailability of managed DNS services. We believe that our results will be useful for managed DNS providers in understanding the possible aftermath of unforeseen downtime.

We first review related work (§2), and then present our data sets and analysis methods (§3), results (§4), and discuss our conclusions and implications (§5).

We use mostly open source data-sets: Alexa top 1M historical archives from TU München¹ and URL categorization portal from McAfee [5]. OpenINTEL data are available to academic researchers upon request [6]. We use `statsmodels` package [7] in python to implement logistic regression.

2. Related Work

Earlier work has shown that even large-scale DNS deployments, such as the root zone of the DNS, can be a victim of DDoS attacks. Moura et al. provided a detailed analysis of such an attack on the root in 2015, which not only affected operations of one root server family, but led to collateral damage to other DNS operators [8].

Allman [9] reflected on the structural robustness of the DNS, and provided recommendations to improve operator practices, including a call to action to implement topological and geographical diversity in terms of authoritative name server deployment. Moura et al. studied how DNS caching and retry behavior by resolvers may hide up to 90% of DDoS attacks on the DNS from the perspective of Internet users, but also noted that this resolver behavior may impose a significant additional load on servers under attack [10].

Centralization of DNS services and resolution is increasing – up to 30% of queries to two ccTLDs originated from just five cloud providers in 2020 [11]. Such centralization enables economies of scale in provisioning infrastructure, but can also present a single point of failure [3].

The technical and business implications of the attack on Dyn were discussed extensively [3], [12]. Antonakakis et al. analyzed the emergence of Mirai, the botnet used for attacking Dyn, while also presenting the possibility that the attackers were targeting not only Dyn but also Sony Playstation infrastructure [12]. Bates et al. analyzed the market share of Dyn within the top 1000 Alexa rank domains (that belonged to `.com/.net/.org` TLDs) and

found that its market share dropped from 10% before the attack to 7.5% within the month after the attack [3].

In this paper, we extend on previous work in which we analyzed the immediate impact of the attack on Dyn, and an earlier attack on another managed DNS provider, NS1 [1]. In the previous paper, we showed that in terms of risk management, using multiple DNS providers is a good strategy. In this paper, we extend and analyze factors that potentially drove the decision-making of Dyn customers and characterize the behavior of different types of customers when faced with the consequences of such a large-scale attack.

3. Data Sources and Methodology

We outline the data sources that we use in this study, and we present our methodology in detail.

3.1. Data Sources

We primarily used three data sources for this study: longitudinal DNS measurement data, website popularity rankings, and URL categorization data.

Longitudinal DNS measurement data: We obtained data from the OpenINTEL [6] project, which measures a large part of the global DNS namespace on a daily basis. OpenINTEL collects, through active querying, resource records configured for domain names. The collected data include: (1) mappings of domain names to authoritative name server records; (2) encountered CNAME records along with their expansion; and (3) mappings of domain names (including name server names) to IP addresses. Using these data points, we indicate the authoritative name servers of each domain name. We consider domain names under `.com`, `.net`, and `.org` that used Dyn at any point in time on one day prior to the Dyn attack (Oct. 20, 2016) and monitor them for 20 days following the attack (Nov. 10, 2016). We also identify the use of non-Dyn servers by the same names.

Website popularity ranking: We measure domain popularity on the basis of the Alexa Top 1 million website list [13]. We extract the daily rank of each domain name from the Alexa Top 1M list for 90 days prior to the incident and calculate the median value. We create six cumulative groups of domains based on median rank, namely: Top 100, Top 1K, Top 10K, Top 100K, Top 500K, and Top 1M. For example, the group Top 1K consists of domains in the dataset with a median rank between 1 and 1K, inclusive.

URL Categorization: For our analysis, we use a URL categorization service from McAfee [5]. This service uses human-supervised machine learning to classify domain names based on security-related features as well as hosted content (where present).² These classifiers are able to classify URLs into 104 different fine-grained industry

2. We also considered URL categorization services from other providers i.e., FortiGuard and Zvelo. However, FortiGuard only provides a single category label for each domain name while we believe that a domain name might belong to multiple categories such as `microsoft.com` (can be *Business*, *Software*, and *Hardware*). Meanwhile, Zvelo limits the request to 10K hits per month for non-commercial licenses. We use the service from McAfee since it addresses both of the issues.

1. <https://toplists.net.in.tum.de/archive/alexal>

categories. Later in the paper, we discuss the categories that were most influential in deciding a resilience strategy.

The service from McAfee also provides a categorization based on reputation score (RS) of individual domain. RS indicates the level of risk a domain poses to a user’s network, computer, or information [5]. McAfee divides URLs into 4 *Risk Levels* according to their RS from low to high: *Minimal Risk* ($RS < 15$), *Unverified* ($15 \leq RS < 30$), *Medium Risk* ($30 \leq RS < 50$), and *High Risk* ($RS \geq 50$).

Category labels from the URL categorization service indicate the functionalities of the domains based on their content. When used outside of the context, the domains with certain characteristics might lead to unintended consequences for the client’s network and in turn, have negative impact on businesses or individuals. For example, if all employees of a company with a limited bandwidth start using a streaming service, it may deteriorate the quality of network services within this company. McAfee classifies URL categories into *Risk Groups* based on the main aspect that is most at risk when this unintended situation occurs.³ There are 7 risk groups that comprise of *domains that pose a risk on*:

Productivity if they might reduce one’s productivity at work, such as social media.

Information as they allow users to access work-irrelevant information.

Liability if accessing them might be criminal, such as pirated contents.

Security of victimizing the users, such as phishing sites.

Propriety of serving non-criminal adult-only content, such as sexual material.

Bandwidth of providing bandwidth-consuming web pages, such as video streaming.

Communication as they allow direct communication with other users.

3.2. Key Concepts and Methodology

We track the resilience strategy adopted by the domains that were customers of Dyn one day prior to the attack to analyze the change in relationship between Dyn and its customers as a consequence of the attack. We define two key concepts to systematically evaluate the changes in their approach to using managed DNS services: *Domain Status* and *Domain Movement*. We define the

3. A more detailed description of these classifications is in https://trustedsource.org/download/ts_wd_reference_guide.pdf

TABLE 1. ALL POSSIBLE REACTIONS (R0 – R3) FOR DYN CUSTOMERS. *Initial* SHOWS DOMAIN STATUS PRIOR TO THE ATTACK (OCT. 20, 2016), *First Change* SHOWS THE FIRST OBSERVED STATUS CHANGE, AND *Final Change* SHOWS THE LAST.

Reaction	Domain Status		
	Initial	First Change	Final Change
R0: Do nothing	ex non-ex	ex non-ex	ex non-ex
R1: Become non-exclusive	ex	non-ex	ex or non-ex*
R2: Unsubscribe and return	ex or non-ex	unsub	ex or non-ex*
R3: Unsubscribe and not return	ex or non-ex	unsub	unsub

ex: exclusive; non-ex: non-exclusive; unsub: unsubscribed.

*We do not distinguish final domain statuses in these categories.

reactions we observe of these domains based on these two concepts.

Domain Status shows the type of relationship between Dyn and a domain. We identify the *Domain Status* for a domain on a certain day from OpenINTEL data by looking at the presence of Dyn name servers (`*.dynect.*`) in the domain NS records on that day. *Domain Status* for a customer of Dyn on a given day can either be:

Exclusive domain Domain status that indicates the use of *only* name servers from Dyn.

Non-exclusive domain Domain status that indicates the use of name servers from Dyn *and at least one other* than Dyn.

Unsubscribed This status represents that the domain used Dyn name servers one day prior to the attack but stopped using them on or before a certain day. For example, if `example.com` was a customer of Dyn one day before the attack, but stopped using its name servers 4 days after the attack then it would have a status of *Unsubscribed* on and after that day, until it started using Dyn name servers again.

Domain Movement captures the change in *Domain Status* between two consecutive days. As Dyn allows users to list multiple name servers, *Exclusive* domains can later become *Non-exclusive* and vice versa. At the same time, a domain that is previously using Dyn services may decide to stop using its services and fully switch over to another provider, i.e., no longer shows Dyn name servers in NS records. We do not track domains that moved from being non-exclusive to exclusive as we are interested in domains that adopted a more resilient strategy ([1] showed that a substantial number of domains did not change status from non-exclusive to exclusive during this period).

We attribute *Domain Movement* based on the change in *Domain Status* for a domain in the 20 days⁴ after the attack as compared to the *Domain Status* one day prior to the attack.⁵ We define the following movements, in order to capture the transitions between domain statuses:

Moving to non-exclusive Domains that changed the *Domain Status* from exclusive to non-exclusive, i.e., use redundant services.

Unsubscribe Domain response following the attack was to stop using Dyn DNS service.

Return Domains that changed the status from Unsubscribed to Exclusive or Non-exclusive domain, i.e. again became Dyn customers.

Based on the *Domain Movement* associated with each domain that was a customer of Dyn one day prior to the attack, we attribute the following reactions to the domain:

Do nothing Domains which *did nothing* in response to the attack and hence had the same *Domain Status* for the 20 days after the attack as one day before the attack (**R0**).

Become non-exclusive Exclusive Dyn customers one day before the attack whose *Domain Status* changed to non-exclusive (**R1**) at least once within 20 days after the attack.

4. We limit the analysis window to 20 days because after that news broke of Dyn being taken over by Oracle was, and hence we cannot attribute any change in *Domain Status* after 20 days solely to the attack.

5. We do not include the day of the attack because the attack may have overlapped with part of the DNS measurement.

Unsubscribe and return Domains that used Dyn name servers one day before the attack, but unsubscribed and later returned within 20 days after the attack (**R2**).

Unsubscribe and not return Domains that unsubscribed from Dyn services within 20 days after the attack but did not return as Dyn customers (**R3**).

In this paper, we refer to domains that took reaction **R0** as *non-reactive domains* and domains that took reactions **R1/R2/R3** as *re-active domains*. Table 1 shows all possible reactions associated to the observed change in *Domain Status*.

There is a possibility of noise when we interpret *Domain Movement* based on OpenINTEL data. Specifically, inconsistent configuration between multiple authoritative for the same domain could lead to mis-identification. To account for this possibility, we consider only status changes that persist for three days or longer. Moreover, we define a domain’s reaction as its first movement we observe within 20 days after the attack.

As indicated in Table 1, we do not further distinguish status of exclusive domains after they became non-exclusive, as well as of those that returned after unsubscribing. Recall that becoming an exclusive customer involves having no (back-up) servers at another provider. Meanwhile, unsubscribing from Dyn service may avoid the risk of having a second or third service interruption from the provider. To be non-exclusive is relatively the most resilient option by having DNS service redundancy, even though it costs more or may add privacy risks.⁶

We categorize domains based on four characteristics, i.e., *popularity*, *industry sector*, *risk level*, and *risk group*, that may influence the decision of Dyn customers to use a more resilient strategy. To analyze this, we follow this step-wise approach:

- Step 1** We identify domains using Dyn one day prior to the attack based on OpenINTEL data.
- Step 2** We map domain characteristics of the domains identified in **Step 1** based on McAfee URL categorization data. We add historical Alexa ranks (median ranks of 90 days prior to the attack) of domains to measure domain popularity.
- Step 3** We evaluate the impact of website popularity on resilience strategy irrespective of the industry sector.
- Step 4** We use a multivariate logistic regression model [14] to distinguish among industry sectors that are statistically significant factor in determining *Re-active Domains* and *Non-reactive Domains*. We evaluate in more detail the resilience strategy of the 20 most statistically significant industry sectors.
- Step 5** We analyze the differences in resilience strategies of Dyn customers categorized by risk level and risk group.

4. Result and Discussion

We identify nearly 168K labeled domains with 94 unique industry categories from the categorization process that were Dyn customers on the day prior to the DDoS attack. Here we present the results of our analysis.

6. By using open public DNS.

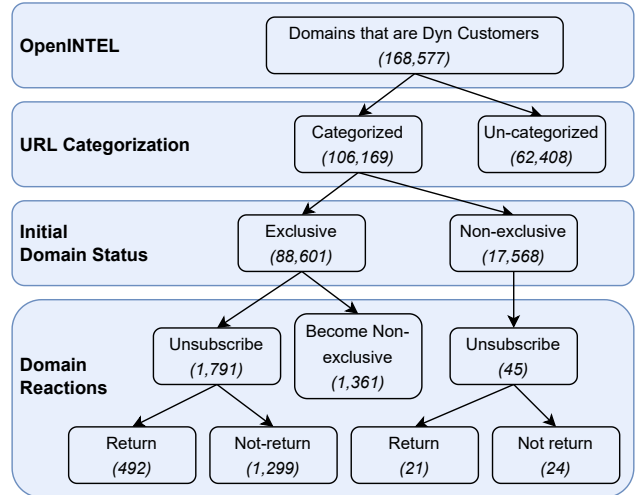


Figure 1. Dyn customer subdivision in category, domain status and reaction. We identify 169K customers and categorize them per our methodology. We distinguish exclusive and non-exclusive domains and subdivide them into reaction. ~17K non-exclusive domains, 45 unsubscribed following the attack, and 24 out of 45 did not return to Dyn later.

4.1. Categorization of Dyn Customers

We query category labels and reputation scores of the initial set of domains using a publicly available URL categorization service provided by McAfee “Real-Time Database” [15]. ~62% of the domains were categorized by McAfee.

Some domain names could not be categorized using McAfee; they are either no longer in operation or do not have enough features available for McAfee’s classifier to make a deterministic classification. We refer to these domains as *Un-categorized domains*. In this paper, we limit our analysis only to the domain names that McAfee successfully categorized, i.e., around 106K unique domains or nearly 62% of all unique domains that used Dyn services one day prior to the attack (Fig. 1).

Some domain names are mapped to multiple industry sectors and risk groups. We consider all retrieved category labels for our analysis. Hence, the total number of category labels retrieved for the domains exceeds the total number of unique domains.

We retrieve daily *Domain Status* for all categorized domains from the measurement data for 20 days after the attack to observe status changes which we use to capture *Domain Movements* as discussed in §3.

The number of Non-reactive Domains is higher than Re-active Domains. The number of *Non-reactive Domains* is significantly higher than those who became non-exclusive or unsubscribed (Fig. 1). We observed that more *Exclusive Domains* either became non-exclusive or unsubscribed from Dyn services, as they had a greater probability of experiencing downtime due to the lack of redundancy in authoritative name servers. Among non-exclusive domains, only 45 out of nearly 17K domains reacted. More than 85K out of ~89K of *Exclusive Domains* did not show any reaction (Fig. 1). This may be explained by the presence of cached NS records on DNS resolvers keeping several domains accessible during the outage [10]. Also, for some organizations the short term




Figure 2. Reactions of Dyn customers – re-active or non-reactive – among different Alexa popularity groups. For example, among *top100* Alexa customers, 42.9% re-acted to the attack. (Note that larger groups are supersets, i.e., the *top100* are included in the *top1k*)

cost of moving to another provider or adding redundancy may not justify benefits.

The number of domains that became non-exclusive is not significantly different from the number of domains that unsubscribed. *Re-active* domains showed two types of movement in terms of their decision to become non-exclusive or to unsubscribe from Dyn (Fig. 1). Among 3,197 re-active domains, 1,361 became non-exclusive and 1,836 unsubscribed. There is no clear preference in the type of movement among domains. This observation motivates us to zoom into domains with similar characteristics, i.e., popularity or industry sectors, to discover specific preferences in the resilience strategy among similar Dyn customers.

4.2. Popularity Influence on Decision to Re-act

We then incorporate Alexa rank data into our data set to indicate how popular each domain is. Out of $\sim 106K$ domains, we manage to retrieve popularity ranks of $\sim 5K$ domains that belong to the Top 1M Alexa list and ignore the ones which do not. We observe that domains with higher Alexa rank have a higher proportion of re-active domains (e.g., 42.9% in *top100*) compared to those with lower rank (e.g., 14.3% in *top1m*) as shown in Fig. 2: the percentage of re-active domains decreases as more domains with lower rank are included. Popularity is thus one factor that appears to influence the decision to re-act. A higher rank means higher traffic to these domains which may result in a higher revenue [16]. Hence, owners of domains with higher rank will prefer minimizing downtime of their domains to avoid revenue loss. Meanwhile, owners of domains with lower rank may not be sensitive to short-term downtime. Some reactions involve an extra cost, e.g., using multiple DNS service providers, which may not be worthwhile for businesses with smaller revenue.

Key takeaway: *We observe that customers running domains for more popular websites (based on Alexa rank) reacted to the attack with higher probability.*

4.3. Industry Sector Influence on Re-action

We discuss in more detail the Industry Sectors that had a significant impact on the resilience strategy adopted by the domain managers.

TABLE 2. GLM COEFFICIENTS FOR DOMAINS SHOWING MOST PROMINENT INDUSTRY SECTORS THAT HAD A STATISTICALLY SIGNIFICANT INFLUENCE ON DOMAIN REACTIONS (ORDERED BY THE NUMBER OF DOMAINS IN EACH INDUSTRY SECTOR).

Industry Sector	Coefficient*	z-value
Internet Services	0.47	6.80
Travel	1.25	14.61
Content Server	1.26	11.00
Motor Vehicles	0.78	6.26
General News	2.17	24.05
Job Search	1.30	9.40
Streaming Media	2.15	6.57
Gambling	1.60	8.97
Portal Sites	1.74	6.51
PUPs (potentially unwanted programs)	2.83	12.28
Malicious Sites	2.61	8.91
Social Networking	2.47	7.75
Personal Pages	1.87	5.37
Provocative Attire	2.71	4.86
For Kids	3.07	4.93
Interactive Web Applications	2.61	7.97
Weapons	1.98	5.32
Professional Networking	3.58	9.34
Visual Search Engine	4.12	7.42
Instant Messaging	4.03	5.32

* $p < 0.001$

To determine the impact of individual industry sectors on domain reactions, we model the resilience strategy adopted by the domains as a function of their industry sector using a multivariate logistic regression analysis. We model the binary decision of domains to re-act (re-act: **R1/R2/R3** or not react: **R0**) as a dependent variable (1).

$$\text{logit}(\pi_b) = \log\left[\frac{\pi_b}{1 - \pi_b}\right], \quad (1)$$

where $\pi_b \in [0, 1]$ represents the probability of a domain reacting and can be estimated using (2).

$$\pi_b = \frac{\beta_0 + \sum_i \beta_i x_i}{1 + (\beta_0 + \sum_i \beta_i x_i)}, \quad (2)$$

where $x_i (i = 1, 2, 3, \dots, 94)$ represents each industry sector as a separate independent variable. x_i has a value of 1 if a domain belongs to this industry sector, otherwise it has a value of 0. β_i is the odds ratio and signifies the strength of correlation between this independent variable and the dependent variable. To implement logistic regression, we use `statsmodels` package in python [7].

The regression model outputs three main metrics to indicate the influence of each category (shown in Table 2). *Coefficient* (β_i) is used to calculate the factor by which the probability of a domain re-acting to the attack increases, if this domain belongs to Industry Sector i , i.e., if the regression coefficient is estimated as β_i then the probability of post-attack domain re-action increases by e^{β_i} times as compared to the base case.

For example, domains in the *Professional Networking* with a coefficient value 3.58 are 16 times (i.e. $\frac{e^{3.58}}{e^{0.78}}$) more likely to re-act than domains in the *Motor Vehicles*, such as web sites of car manufacturers and sales, with a coefficient value 0.78. Hence, we expect the proportion of re-active domains is higher in the former than the latter. *p-value* indicates if the coefficient of a domain is statistically significant or not. *z-value* is calculated as the




Figure 3. Distribution of z-values (signifies reliability) for industry sector coefficients in the linear regression model. We discuss in more detail the reactions of domains that belong to 20 industry sectors with the largest z-values, indicated by the shaded area.

number of standard deviations the average decision of a certain industry sector is away from the average decision of the whole data set [17]. It indicates the reliability of the coefficient computed by the logistic regression model.

We discuss the type of reactions observed for domains that belong to any of the top 20 industry sectors whose coefficient was both highly reliable (greater than the median z-value as shown in Fig. 3) and statistically significant ($p\text{-value} < 0.05$).

The results indicate that some categories have more significant correlation to the resilience strategy than others. Domains in industry sectors with higher coefficient have a higher probability to re-act following outages, such as *General News* and *Streaming Media*. DNS providers should anticipate the reaction from customers in these industry sectors more than other sectors.

Key takeaway: *By using a multivariate logistic regression model, we were able to calculate and reveal that for 34 out of 94 industry sectors, there is statistically significant influence on the post-attack reaction.*

We also observe a number of industry sectors which **do not** have statistically significant influence on the decision to re-act such as *Religion*, *Ideologies*, and *Politics*. Religion and Ideologies sectors include web pages that discuss spirituality and beliefs while Politics domains host discussions of political parties, views, and opinions. We suspect that domains in these sectors do not require high resilience.

Key takeaway: *Some industry sectors do not seem to have a statistically significant influence on the decision to re-act after the attack.*

We zoom into a number of significantly influential industry sectors (Table 2). Fig. 4 shows how domains in these industry sectors reacted. The *top plot* shows proportions of re-active domains in each industry sector, characterized by their decision to become non-exclusive or to unsubscribe after the attack. The diminishing line in the top plot indicates the size, i.e., number of domains, of each industry sector on a \log_{10} scale. We observe that sector

size does not directly correlate with the proportion of re-active domains in the sector. However, we demonstrate that reactions vary among industry sectors in two ways: the proportion of *re-active* domains and the dominant type of reactions.

From the *top plot* in Fig. 4, we show that most domains in some industry sectors *became non-exclusive* following the incident, i.e., used additional DNS providers, such as domains in *General News*, *Potentially Unwanted Programs (PUPs)*, *Malicious Sites*, and *Visual Search Engine*. Domains in *PUPs* and *Malicious Sites* categories are made with an intention to infect user’s computer with software or other malicious activity, e.g., Trojan horses and viruses. Domains in *Visual Search Engine* include web pages that provide searching functionality with image-specific results. The reaction to use multiple DNS services might indicate that these domains have a high – most probably financial – value for the owners, hence, it is important to keep them up at all times. After the attack, these domains used back-up services to be more resilient against future outages. Therefore, we assume that the revenue from the domains is large enough to compensate for the cost of using an additional DNS service, if present.

Meanwhile, domains in some industry sectors showed a significant preference to **unsubscribe**, i.e., they stopped using Dyn service, after the attack (Fig. 4). Among them, we observe domains intended for adult audiences – *Gambling* sites, web pages about *Weapons*, and sites showing *Provocative Attires* – and the domains intended for more general audiences such as *Motor Vehicles* sites, *Streaming Media*, and *Social Networking* sites. The reaction to unsubscribe after the attack might also indicate that domain owners value their domains and find it necessary to re-act after the attack. However, we suspect that the cost of using a secondary DNS service is not worth the revenue from the domain, or the owner considers Dyn no longer reliable.

Key takeaway: *We see a difference in the type of reaction based on the industry sector. In some sectors, domains for the most part became non-exclusive following the attack (e.g., General News), while for others, customers largely unsubscribed (e.g., Gambling).*

The *middle plot* shows the proportion of return among domains that unsubscribed in each sector. This metric is important to measure the business impact of an outage incident as it indicates how many customers the provider loses for a longer term. In some industry sectors such as *Travel* (e.g., ticketing sites), *Motor Vehicles*, and *Personal Pages* (e.g., blogs), more domains **did not return** than those did return. On the contrary, the behavior was quite the opposite in *Gambling*, *PUPs*, and *Malicious Sites* sectors, i.e., most unsubscribed domains **returned** later. In *Visual Search Engine* sector, there is no domain that unsubscribed, hence, no data is shown in the plot for that sector.

Key takeaway: *In some industry sectors, domains that unsubscribed mostly did not return (e.g., Travel), while in others, a higher proportion of domains unsubscribed and returned (e.g., Gambling). This suggests that the industry sector affects decision-making.*




Figure 4. Behaviors of domains in selected industry sectors. **Top Plot** shows the proportions of domains with different reactions in each sector along with the sector sizes. **Middle Plot** shows the proportions of unsubscribed domains that did not return. **Bottom Plot** shows the distribution of Alexa ranks among domains that belong to a certain industry sector.

The *bottom plot* in Fig. 4 shows the distribution of Alexa ranks of the domain names – where present – within each industry sector categorized by their responses after the attack i.e., re-active or non-reactive. This plot reveals insights into the influence of popularity on the choice of resilience strategy adopted by domains in each industry sector. For most industry sectors shown in Fig. 4, the average popularity of *re-active* domains was higher than the average popularity of non-reactive domains. This is consistent with our previous findings in §4.2 that higher popularity influences the domain decision to react. However, we see that popularity was a more significant factor for some industry sectors than the others. **Non-overlapping inter quartile ranges (IQR)** for the box plots in *Job Search* and *Gambling* sectors indicate a strong correlation between popularity and the decision to react. Meanwhile, in some industry sectors such as *Internet Service*, domain popularity does not seem to have a strong correlation with the decisions to react, considering the substantial **overlap in the IQR** for the box plots of reactive and non-reactive domains.

Key takeaway: In some industry sectors (e.g., *Job Search*), the influence of popularity on the decision to re-act is stronger than in others (e.g., *Internet Service*).

4.4. Risk Level Influence on Decision to React

We analyze how risk level of the domains influences their reaction following the attack. Fig. 5 shows the behavior of domains with different risk categories, characterized by risk level and risk group. Recall that McAfee categorizes domains based on the reputation score into four risk levels as shown in §3. Most domains using Dyn services are categorized as ‘Minimal Risk’. Out of 106K domains of Dyn customers, 104K are in Minimal Risk while domains in Unverified, Medium Risk, and High Risk categories only account for 2K domains (Fig. 5).

Some domains with higher risk levels also need high availability, hence, using the managed DNS service from Dyn. As illustrated by the *top plot* in Fig. 5, some Dyn customers are domains in *medium* and *high* risk levels and they show a similar preference to become non-exclusive after the attack, i.e., use redundant DNS services. Moreover, most domains with medium and high risks returned if they unsubscribed, (the *middle plot*). We also observed this behavior among domains intended to victimize the users, e.g., PUPs and Malicious Sites (§4.3).

We see a significant **overlap in the IQR** in the box plots of Alexa ranks for domains categorized as minimal risk (see: *bottom plot* of Fig. 5). This shows that popularity (for this aggregate set of domains) does not play a major role in deciding the post-attack reactions. The number of

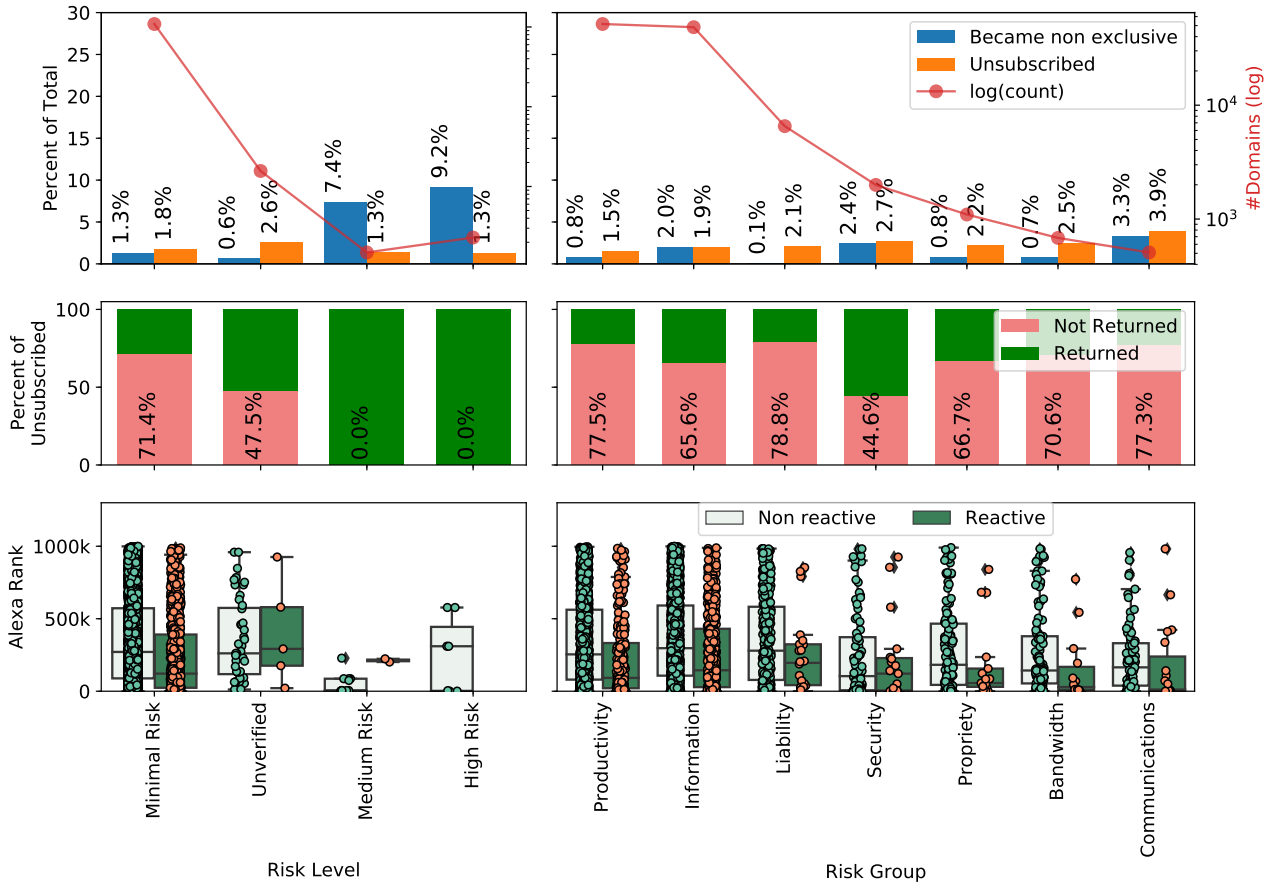


Figure 5. Behaviors of domains with different *Risk Levels* and *Risk Groups*. **Top Plot** shows the proportions of domains with different reactions in each group along with the group sizes. **Middle Plot** shows the proportions of unsubscribed domains that did not return. **Bottom Plot** shows the comparison in the distribution of Alexa ranks of re-active and non-reactive domains within each category.

re-active domains categorized as unverified or medium risk are too low to derive any conclusions about the influence of popularity on the observed reaction. Meanwhile, no domain categorized as a high risk belonged to Alexa Top 1M list. This observation is again consistent with our result in §4.3 with domains categorized as *malicious domains*.

Key takeaway: *Customers with higher risk levels mostly became non-exclusive after the attack and the majority of those who unsubscribed instead also later returned.*

In addition, considering the *top* and *middle plots* in Fig. 5, we observe that majority of the domains with minimal risk stopped using Dyn services permanently following the attack, i.e. nearly 1K domains or 71.4% of those that unsubscribed did not return. This shows that once customers of DNS service providers unsubscribe there is less chance for them to re-subscribe to the service. Other DNS providers may use similar heuristics to estimate customer loss following an outage.

Key takeaway: *Our results show that if domains in the minimal risk category unsubscribe, the probability that they stay away for a longer time is higher.*

4.5. Risk Group Influence on Decision to React

Fig. 5 shows the re-actions among domains that belong to different *Risk Groups*. According to the *top plot*, domains with *Liability*, *Bandwidth*, and *Propriety* risks show a slight preference to unsubscribe after the attack. More than 60% of the unsubscribed domains in these categories did not return.

Meanwhile, as shown in the *middle plot*, we observe similar return behaviors of domains in different risk groups, i.e., more than half of them did not return after they unsubscribed, except for domains in *Security* risk group.

As shown in the *bottom plot* of Fig. 5, the popularity of re-active domains in all risk groups (except *Security*) is higher than the popularity of the domains that do not re-act. For 5 of the 7 risk groups, we observe that the re-active domains are concentrated in relatively higher Alexa ranks. This indicates that for these domains (i.e. *Liability*, *Security*, *Propriety*, *Bandwidth*, and *Communications*), popularity played an important role in the decision of domain owners to re-act. Interestingly, these domains also include websites providing streaming and messaging services.

This result is consistent with our observation in §4.2 that higher popularity shows a correlation with the decision to re-act after the attack. In *Productivity* and *Infor-*




Figure 6. Distribution of variety in re-actions among Domains mapped to an AS.

tion risk groups (eg. social media websites), we see a substantial **overlap in the IQR** of the re-active and non-reactive box plots, which indicates a weaker correlation of the domain popularity with the decision to react.

Key takeaway: Domains in some risk groups (e.g., Liability) mostly unsubscribed after the attack.

4.6. Analysis on Domain Autonomous Systems

A single party can operate multiple domains that rely on Dyn for DNS. For example, Alphabet Inc. operates, among others, google.com and gmail.com. So far we have not considered the potential effects of a single party making decisions (i.e., driving domain reactions) for multiple domains. As a first step towards exploring if multiple re-actions can be correlated with decisions made by a single party, we measure the variety of reactions of Dyn customers within each Autonomous System (AS).

We associate domains to their hosting network at the time of the Dyn attack by mapping IP addresses (i.e., A records) to AS numbers, using IP-prefix-to-AS data from CAIDA [18]. We then calculate the percentage of re-active domains that re-acted in each AS. In other words, if two domains (both Dyn customers) mapped to a single AS and one re-acted, the percentage of re-active domains for that AS is 50%. Fig. 6 shows the resulting distribution. We only include ASes that were mapped to at-least 1 re-active domain in this plot. Nearly 1.6K different ASes, each in which on average 14 domains were hosted, showed no re-active domains at all. We observe that in the case of about 200 smaller ASes (one to only a few domains hosted), the domains all showed the same behavior (i.e., 100% re-acted). This suggests that parties that operated several domains drove the reaction for multiple domains, but on the order of a few hundred domains in total based on our initial view. This approach of course does not account for cases where a single party hosts their domains in multiple networks, which we leave for future work.

Key takeaway: Domains mapped to the same AS showed different re-actions.

5. Conclusion and Implications

Unplanned outages of managed DNS services can lead to financial losses for customers. Depending on how

customers value online availability, outages could prompt customers to switch to a different DNS service provider or to add a (secondary) provider to add resilience. We leveraged large-scale DNS measurement data to infer 168K Dyn customers at the time of the infamous 2016 DDoS attack and tracked the choices made by customers following the attack. We also used a publicly available URL categorization service and website popularity metrics to characterize customers. We demonstrated the impact of factors such as domain popularity, industry sector, risk level, and risk group on the adopted resilience strategy.

Our findings are as follows:

- We confirmed the intuition that domains with a higher popularity (i.e., larger user base) are more likely to re-act (40% of customers on the Alexa Top 100 in contrast to only 15% in the Top 500K);
- We inferred 94 industry sectors for customers and revealed that 34 sectors statistically significantly influenced post-attack customer behavior;
- We unveiled that customers in different sectors adopted different resilience strategies. For example, in some sectors customers typically stopped using Dyn altogether whereas in others they added an additional DNS provider to reduce the risk of unavailability;
- We showed that the popularity of domain names is more tightly related to post-attack behaviors in some industry sectors over others;
- A large proportion of domains categorized as ‘minimal risk’ that unsubscribed from Dyn, did not return (more than 71% did not return). We observe that domains categorized as medium or high risk prefer to become non-exclusive after the attack (9% became non-exclusive while 1% unsubscribed). We also see that domains in some risk groups such as Liability (e.g., pirated movie portals) showed a slight preference to unsubscribe after the attack, with 0.1% becoming non-exclusive and 2% unsubscribed.

The results of our study could help managed DNS providers to estimate the impact of an unplanned outage based on their customer portfolio. Businesses such as cloud service providers could make use of a methodology similar to ours to understand the impact of unplanned downtime on their customer portfolio [19].

We show that *domain popularity*, *industry sector*, and – to some extent – *risk groups & risk level*, are factors that can be used to understand the choice of resilience strategy.

5.1. Limitations and future research

The nature of our data sources and assumptions that we need to make lead to a few limitations. *First*, the URL categorization service we use in this study might exclude domains without hosted content. Recall that the domain content is a feature in our categorization process. Moreover, one should be aware that all existing URL categorization services, including McAfee’s, have accuracy concerns [15]. *Second*, Alexa does not rank non-website domains (i.e., other networked services), despite their high traffic. Moreover, Alexa ranks the Top 1 million most popular websites only. Consequentially, most Dyn

customers that we infer ($\sim 163\text{K}$ out of $\sim 168\text{K}$) do not have an associated rank. In addition, Alexa Top 1M is prone to bias and a rigorous alternative exists, i.e., Tranco [20]. However, Tranco did not yet exist when the attack happened in 2016. Last, the insights we draw are based on the analysis of data connected to a single attack event. Hence, we understand that similar behavior may not be observed for another event. However, researchers could apply our methodology to other similar attacks and outages to improve the generalizability of our findings.

As one area of future work, we plan to analyze the response following outages (attack related or otherwise) in other online services such as CDN or web hosting using a similar approach. Even though these outages would also lead to downtime, the behavior observed may be different given that the operator/manager might need to employ greater effort to opt for a more resilient strategy. In addition, considering the development of the DNS ecosystem since 2016, we plan to analyze more recent incidents to gain updated insight into the topic and study what changed in the last six years.

References

- [1] A. Abhishta, R. van Rijswijk-Deij, and L. J. M. Nieuwenhuis, "Measuring the impact of a successful DDoS attack on the customer behaviour of managed DNS service providers," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 5, pp. 70–76, Jan. 2019. [Online]. Available: <http://doi.org/10.1145/3310165.3310175>
- [2] R. Bolstridge, "Dyn ddos attack: Wide-spread impact across the financial services industry (part 1). blog," 2016.
- [3] S. Bates, J. Bowers, S. Greenstein, J. Weinstock, Y. Xu, and J. Zittrain, "Evidence of decreasing internet entropy: the lack of redundancy in dns resolution by major websites and services," National Bureau of Economic Research, Tech. Rep., 2018.
- [4] "Ponemon institute and akamai reveal the trends in the cost of web application and denial of service attacks: Asia-pacific," 2018. [Online]. Available: <https://www.akamai.com/newsroom/press-release/ponemon-institute-and-akamai-reveal-the-trends-in-the-cost-of-web-application-and-denial-of-service-attacks-asia-pacific>
- [5] "Check Single URL." [Online]. Available: <https://trustedsource.org/>
- [6] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras, "A high-performance, scalable infrastructure for large-scale active dns measurements," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 6, pp. 1877–1888, 2016.
- [7] S. Seabold and J. Perktold, "statsmodels: Econometric and statistical modeling with python," in *9th Python in Science Conference*, 2010.
- [8] G. C. Moura, R. d. O. Schmidt, J. Heidemann, W. B. de Vries, M. Muller, L. Wei, and C. Hesselman, "Anycast vs. ddos: Evaluating the november 2015 root dns event," in *Proceedings of the 2016 Internet Measurement Conference*, ser. IMC '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 255–270. [Online]. Available: <https://doi.org/10.1145/2987443.2987446>
- [9] M. Allman, "Comments on DNS Robustness," in *Proceedings of the Internet Measurement Conference 2018*, ser. IMC '18. New York, NY, USA: Association for Computing Machinery, Oct. 2018, pp. 84–90. [Online]. Available: <http://doi.org/10.1145/3278532.3278541>
- [10] G. C. M. Moura, J. Heidemann, M. Müller, R. de O. Schmidt, and M. Davids, "When the dike breaks: Dissecting dns defenses during ddos," in *Proceedings of the Internet Measurement Conference 2018*, ser. IMC '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 8–21. [Online]. Available: <https://doi.org/10.1145/3278532.3278534>
- [11] G. C. M. Moura, S. Castro, W. Hardaker, M. Wullink, and C. Hesselman, "Clouding up the internet: How centralized is dns traffic becoming?" in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 42–49. [Online]. Available: <https://doi.org/10.1145/3419394.3423625>
- [12] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis *et al.*, "Understanding the mirai botnet," in *26th USENIX security symposium (USENIX Security 17)*, 2017, pp. 1093–1110.
- [13] "Alexa - Top sites." [Online]. Available: <https://www.alexa.com/topsites>
- [14] J. Gill, J. M. Gill, M. Torres, and S. M. T. Pacheco, *Generalized linear models: a unified approach*. Sage Publications, 2019, vol. 134.
- [15] P. Vallina, V. Le Pochat, Feal, M. Paraschiv, J. Gamba, T. Burke, O. Hohlfeld, J. Tapiador, and N. Vallina-Rodriguez, "Mis-shapes, Mistakes, Misfits: An Analysis of Domain Classification Services," in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC '20. New York, NY, USA: Association for Computing Machinery, Oct. 2020, pp. 598–618. [Online]. Available: <http://doi.org/10.1145/3419394.3423660>
- [16] H. Truong, "Relationships between Web Traffic Ranks and Online Sales Revenue of E-Retailers in Australia," Ph.D. dissertation, Feb. 2018.
- [17] S. McLeod, "[Z-Score: Definition, Calculation & Interpretation]," 2019. [Online]. Available: <https://www.simplypsychology.org/z-score.html>
- [18] . The CAIDA UCSD, "Routeviews Prefix to AS mappings Dataset (pfx2as) for IPv4 and IPv6," Jul. 2008. [Online]. Available: <https://www.caida.org/catalog/datasets/routeviews-prefix2as/>
- [19] M. Y. M. Haq, A. Abhishta, and L. J. M. Nieuwenhuis, *The Effect of Consumer Portfolio on the Risk Profile of Cloud Provider*. New York, NY, USA: Association for Computing Machinery, 2021, p. 18–20. [Online]. Available: <https://doi.org/10.1145/3472716.3472851>
- [20] V. Le Pochat, T. Van Goethem, S. Tajalizadehkhoob, M. Korczynski, and W. Joosen, "Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation," in *Proceedings 2019 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2019. [Online]. Available: https://www.ndss-symposium.org/wp-content/uploads/2019/02/ndss2019_01B-3_LePochat_paper.pdf