Measuring the Impact of a Successful DDoS Attack on the Customer Behaviour of Managed DNS Service Providers*

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

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

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

ABSTRACT

Distributed Denial-of-Service (DDoS) attacks continue to pose a serious threat to the availability of Internet services. The Domain Name System (DNS) is part of the core of the Internet and a crucial factor in the successful delivery of Internet services. Because of the importance of DNS, specialist service providers have sprung up in the market, that provide managed DNS services. One of their key selling points is that they protect DNS for a domain against DDoS attacks. But what if such a service becomes the target of a DDoS attack, and that attack succeeds?

In this paper we analyse two such events, an attack on NS1 in May 2016, and an attack on Dyn in October 2016. We do this by analysing the change in the behaviour of the service's customers. For our analysis we leverage data from the OpenINTEL active DNS measurement system, which covers large parts of the global DNS over time. Our results show an almost immediate and statistically significant change in the behaviour of domains that use NS1 or Dyn as a DNS service provider. We observe a decline in the number of domains that exclusively use NS1 or Dyn as a managed DNS service provider, and see a shift toward risk spreading by using multiple providers. While a large managed DNS provider may be better equipped to protect against attacks, these two case studies show they are not impervious to them. This calls into question the wisdom of using a single provider for managed DNS. Our results show that spreading risk by using multiple providers is an effective countermeasure, albeit probably at a higher cost.

CCS CONCEPTS

• **General and reference** → **Empirical studies**; **Measurement**; *Validation*; • **Networks** → **Denial-of-service attacks**;

KEYWORDS

DDoS Attacks, Dyn, NS1, Domain Name System, Customer Behaviour, Economic Impact

1 INTRODUCTION

Distributed Denial-of-Service (DDoS) attacks continue to pose a serious threat to the availability of Internet-based services. In the last decade we have seen a constant increase in the intensity of these attacks [1–3]. An immediate impact of a successful DDoS attack is the unavailability of services provided by the victim to its customers. For instance, for an e-commerce firm this unavailability might result in decrease of sales during the attack and can also cause damage to the reputation of the victim [8].

These attacks also threaten the availability of services that support the Internet usage for an everyday user. One of the core services on which the Internet is built is the Domain Name System (DNS). DNS is responsible for translating easy to remember domain names into machine readable IP addresses. Thus, unavailability of the DNS leads to unavailability of web services for most users. On several occasions, attackers have targeted the DNS with a DDoS attack to bring down web services. Hence, it is important for firms that prioritise availability to choose a DNS provider that is resilient in the face of DDoS attacks. There are several managed DNS providers that provide DDoS resilient services. NS1 and Dyn are two such managed DNS (MDNS) service providers. On May 16th, 2016 and October 21st, 2016, DDoS attacks targeted NS1 [6] and Dyn [12] respectively. The attacks were successful in hindering the services provided by NS1 and Dyn for the better part of a day.

While much has been said about the impact of especially the Dyn attack, one aspect of these attacks has received far less attention, namely: *What is the impact of such an attack on the behaviour of customers of affected MDNS providers?* In this paper, we study this impact by looking at the DNS configuration of domains in a large DNS dataset. This allows us to answer questions such as: do customers continue to use the services of the attacked MDNS after the attack or not? If they remain a customer, do they change their behaviour?

Our contributions are as follows:

- We provide a framework for measuring the behaviour of domains using an MDNS service provider.
- We use this framework to analyse the impact of successful DDoS attacks on NS1 and Dyn on the behaviour of domains that use their services.
- We show statistically significant changes in customer behaviour after the attacks, such as, e.g., adding a second DNS provider for a domain.
- We show that most customers that start using a second provider use another MDNS service provider as a secondary DNS to further reduce the risk of downtime.

2 DNS AS A RESOURCE

In order to understand the behaviour of customers after a DDoS attack, it is important to first understand the additional benefits of the service provided by NS1 & Dyn to its customers. In this section we look at the Domain Name System (DNS) as a resource [5] and explain its benefits [28] with the help of a so-called value network. DNS is one of the core services that supports the Internet. It translates human readable *domain names* (e.g. `www.example.com`) into

*This is a slightly revised version of the paper "Measuring the Impact of a Successful DDoS Attack on the Customer Behaviour of Managed DNS Service Providers" that was initially presented at the SIGCOMM'18 workshop on traffic measurements for cybersecurity.

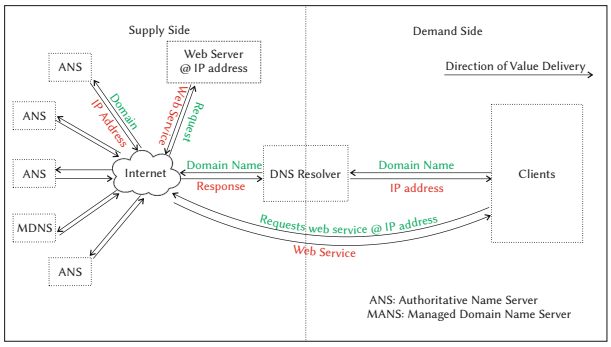

Figure 1: Value network of web service delivery showing the role of various components of the DNS

machine readable *IP addresses* (e.g. `93.184.216.34`) [30]. Hence, it is safe to categorise DNS as a resource that facilitates the delivery of other web based services (e.g. e-banking) to the customers of a bank.

The DNS itself is hierarchically organised: *root level, top level domains, public suffixes, second level domains, third level domain and so on.* The data (a tuple of domain name and IP address) on a domain name server is distributed according to zones and is stored in zone files. The authority for the records related to a domain name is delegated to a so-called authoritative name server (ANS) with the help of so-called `NS` records.

Value of managed authoritative name servers for firms: A *value network* – with related concepts such as actors, roles and value adding activities – can be used to describe and analyse a specific product or service offering in a detailed way [11]. A value network shows the value adding actors involved in the service delivery process and their relationships. Such a network helps in understanding the benefits and roles of each of the actors in the process. A value network is defined as *"a spontaneous sensing and responding spatial and temporal structure of largely coupled value proposing social and economic actors interacting through institutions and technology, to: (1) co-produce service offerings, (2) exchange service offerings, and (3) co-create value"* [17] . Figure 1 shows the value network of a web service delivery.

We can understand the value network shown in Figure 1 by considering an example of a customer who wishes to transfer money to another account without physically visiting a bank. In this case the customer first needs to log on to the e-banking website of their bank using a web browser. Once the customer requests the e-banking website, the web browser then queries the DNS resolver of its network for the *IP address* associated with this *domain name*. In case the response to this query is not present in the cache of the DNS resolver, it retrieves the *IP address* from the authoritative name server (*ANS*) associated with this domain name and forwards the response to the web browser. The web browser then connects to the server located at the *IP address* and in-turn provides the web service to the customer. With the help of this example, it is evident that unavailability of *ANS* can lead to potential unavailability of the e-banking service.

On multiple occasions [16, 23, 24], criminals have targeted the availability of various components of the value network as described above using DDoS attacks. The need for availability of *ANS* has created a market for managed domain name service providers (*MDNS*). These *MDNS* provide the following benefits in addition to the features of an *ANS* [32]:

(1) faster response times;
(2) load balancing;
(3) and DDoS Protection.

Hence, for a *domain* that forms a source of revenue for a company, a *MDNS* promises greater availability and helps the company in efficiently catering to the needs of its consumer.

It is common practice for domain owners to specify multiple ANS for their domain. DNS resolvers may then query each of these authoritative name servers, although they will have a preference on the basis of metrics (e.g. round trip time) [22]. In the context of the use of MDNS, this practice has additional consequences. A domain owner can choose to *exclusively* use multiple ANS from a single provider. If this provider then somehow goes down, the domain owner will suffer unavailability as a consequence. Another option is for the domain owner to procure services from multiple MDNS providers and thus to *non-exclusively* use MDNS services. While this makes DNS management a bit more complex for the domain owner, and potentially comes at a higher cost, it has one significant benefit: if one MDNS provider goes down, the domain will still be available under the assumption that the other MDNS provider(s) are still operational.

MDNS	Dataset	Start Date	End Date
NS1	OpenINTEL	29th October 2015	5th June 2016
Dyn	OpenINTEL	4th of April 2016	11th November 2016

Table 1: Details of Dataset

3 IMPACT OF A DDOS ATTACK

A successful DDoS attack hampers the availability of an MDNS provider. As the added value of using MDNS is DDoS protection, a successful attack can lead to loss of customers in a market where availability is of great importance [9, 36]. In this section we introduce a framework that can capture the behaviour of domains using an MDNS provider. We use this framework to study two DDoS attack events: (1) on NS1 on 16th May 2016 and (2) on Dyn on 21st October 2016. For our analysis, we make use of a large longitudinal dataset that is introduced in the following section.

3.1 Dataset

We use the OpenINTEL dataset as source data to analyse the impact of DDoS attacks on the behaviour of domains using an MDNS provider. The OpenINTEL project collects unique long-term datasets with daily DNS measurements for all domains under the main top-level domains on the Internet (including .com, .net and .org). Currently, OpenINTEL covers over 60% of the global DNS name space every 24 hours. Van Rijswijk-Deij *et al.* [31] explain the data collection method in detail.

We use data for the domains in three generic top level domains (gTLDs) .com, .net and .org. In order to get a list of domains that use Dyn/NS1 on a given day we query the dataset for all domains that use Dyn/NS1 name server addresses in their NS records on that day. We use the measurements in the OpenINTEL dataset for time intervals as shown in Table 1.

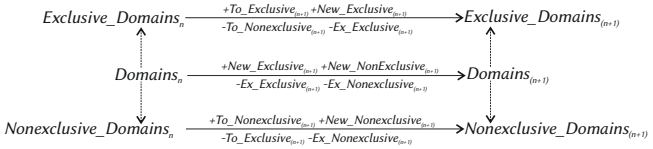

Figure 2: Relationship between the behaviour variables showing the changes in variable from day n to day $n+1$

3.2 Type of domains

On the basis of the number of different service providers found in NS records for a domain, we categorise domains into two types:

Exclusive: A domain is categorised as exclusive if it uses only Dyn/NS1 name server addresses in its NS records.

Non-exclusive: A domain is categorised as non-exclusive if it uses name server addresses from multiple providers including Dyn/NS1.

This categorisation is of great importance for this study as a non-exclusive domain will not experience an inferior service quality during an attack . Now, in order to measure the change in the behaviour of domains we need to first define *behaviour*.

3.3 Measuring the impact

We define a step by step procedure that we use to perform our analysis. In order to measure the impact of the DDoS attack on the domains using NS1/Dyn as an MDNS provider, we use an approach similar to event studies [18]. We use a five step approach to measure the impact as described below:

Step 1: Define variables representing the behaviour of domains.
Step 2: Define a trend period and an event period.
Step 3: Measure behaviour in the trend period.
Step 4: Measure behaviour in the event period.
Step 5: Analyse any changes in behaviour.

Step 1 is discussed in Section 3.3.1, Step 2 in Section 3.3.2, Steps 3 and 4 in Section 3.3.3, and Step 5, the analysis, is discussed in Section 4.

3.3.1 Behaviour of domains. In this study we define the behaviour of domains that use NS1/Dyn's MDNS infrastructure on the basis of the following variables:

$Domains_n$ Total number of domains using NS1/Dyn on day n.
$Exclusive\_Domains_n$ Total number of domains exclusively using NS1/Dyn on day n.
$Nonexclusive\_Domains_n$ Total number of domains that are non-exclusively using NS1/Dyn on day n.
$To\_Exclusive_n$ Total number of domains that move from being non-exclusive to exclusive users of NS1/Dyn on day n.
$To\_Nonexclusive_n$ Total number of domains that move from being exclusive to non-exclusive users of NS1/Dyn on day n.
$New\_Exclusive_n$ Total number of *new* domains that became a new exclusive users of NS1/Dyn on day n (did not use NS1/Dyn on day $n-1$).
$New\_Nonexclusive_n$ Total number of *new* domains that became a new non-exclusive users of NS1/Dyn on day n (did not use NS1/Dyn on day $n-1$).

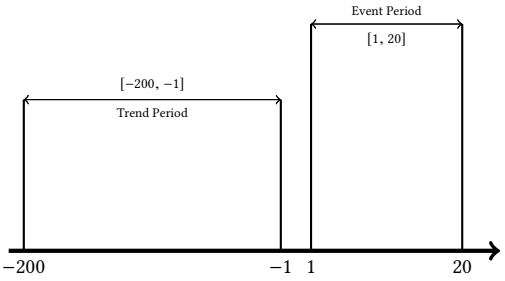

Figure 3: Trend and Event Periods.

$Ex\_Exclusive_n$ Total number of exclusive domains that stopped using NS1/Dyn on day n.
$Ex\_Nonexclusive_n$ Total number of non-exclusive domains that stopped using NS1/Dyn on day n.
$Ex\_Domains_n$ Total number of domains that stopped using NS1/Dyn on day n.

Figure 2 shows the relationship between the behaviour variables. Daily measurements of each of the behavioural variables provide us with a time series. In order to analyse this time series we calculate the daily change and 10-day cumulative average of the behavioural variables. For example, the change in variable $Domains_n$ represented by variable $\Delta Domains_n$ can be calculated with the help of Equation 1.

$$\Delta Domains_n = Domains_n - Domains_{(n-1)} \quad (1)$$

Calculating a 10-day cumulative average for the behavioural variables helps us to measure the net behaviour over a 10 day period [27]. It also filters any short term effects of random events from the time series. A 10-day cumulative average variable will not show changes due to an event whose effects disappear in less than 10 days. The use of cumulative averaging of time series is common practice in statistics to filter out noise. We can calculate the net cumulative average of a behaviour variable for day i as shown in Equation 2.

$$Cumulative\_Variable_i = \frac{1}{10}\sum_{n=9}^{0} Behaviour\_Variable_{i-n} \quad (2)$$

3.3.2 Trend and event period. The trend period is the interval before the attack date that we analyse to study the usual tendency of behaviour variables. This gives us a measure of behavioural variables without the influence of a large DDoS attack event. In this paper we study the usual behaviour of the behaviour variables for 200 days before the DDoS attacks. The trend period considered by us is consistent with the studies that analyse the impact of events on stock prices of the event stakeholders [4]. Similarly, the event period is the interval after the attack date that we analyse to study the deviations from the usual tendency (measure of behavioural variables under the influence of a large DDoS attack event). For this study we chose an event period of 20 days. Relatively soon after the second attack event we are analysing (i.e. the Dyn attack on 21st October 2016) a news article regarding the sale of Dyn to Oracle was published (on 11th November 2016). As this is another major event that may influence customers of Dyn, we run into the risk that any analysis of the behaviour variables beyond the 20

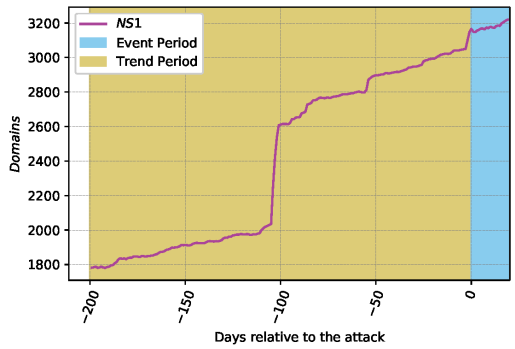

(a) Total Domains NS1

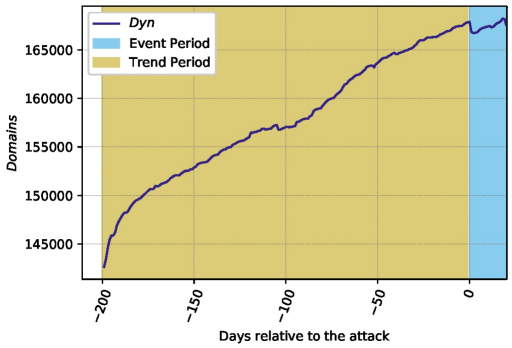

(b) Total Domains Dyn

Figure 4: Total Domains using NS1 and Dyn

day window will be biased as it will also show effects that are a consequence of the takeover by Oracle. In order to keep this event window consistent for both the measurements we consider a 20 day event period for NS1 as well.

3.3.3 Measurement of behaviour variables. We measure the behaviour of domains by calculating the daily values for the behaviour variables that are described before. We do this with the help of the OpenINTEL dataset. $Domains_n$ is computed for each day of both the trend and the event period on the basis of the number of domains having Dyn or NS1 name server addresses in their NS records. If a domain had only Dyn/NS1 NS addresses then it was counted in variable $Exclusive\_Domain_n$, else it was counted in variable $Nonexclusive\_Domain_n$. We also calculated the daily changes in these variables as explained previously. We plot a time series of each of these variables in Figure 5. We discuss the interpretation of these plots in Section 4.1.

Next, we measure the activity of these domains on the basis of the difference in the domains using Dyn/NS1 on two consecutive days. If a Domain was a user of Dyn or NS1 on day $n-1$ but not a user on day n we count it in variable $Ex\_Exclusive_n$ or $Ex\_Nonexclusive_n$ depending on the state of the domain on day $n-1$. For example, if a domain `www.example.com` is an exclusive user of Dyn on day $n-1$ but does not use the services of Dyn on day n, then it will be counted in variable $Ex\_Exclusive_n$. In another case, if a domain was exclusive on day $n-1$ and non-exclusive on day n, we count it in variable $To\_Nonexclusive_n$. If a domain moved from being non-exclusive to exclusive on the next day we count it in variable $To\_Exclusive_n$. Some new domains also start using services of Dyn or NS1 each day, we count them in variable $New\_Exclusive_n$ or $New\_Nonexclusive_n$ depending on their joining status.

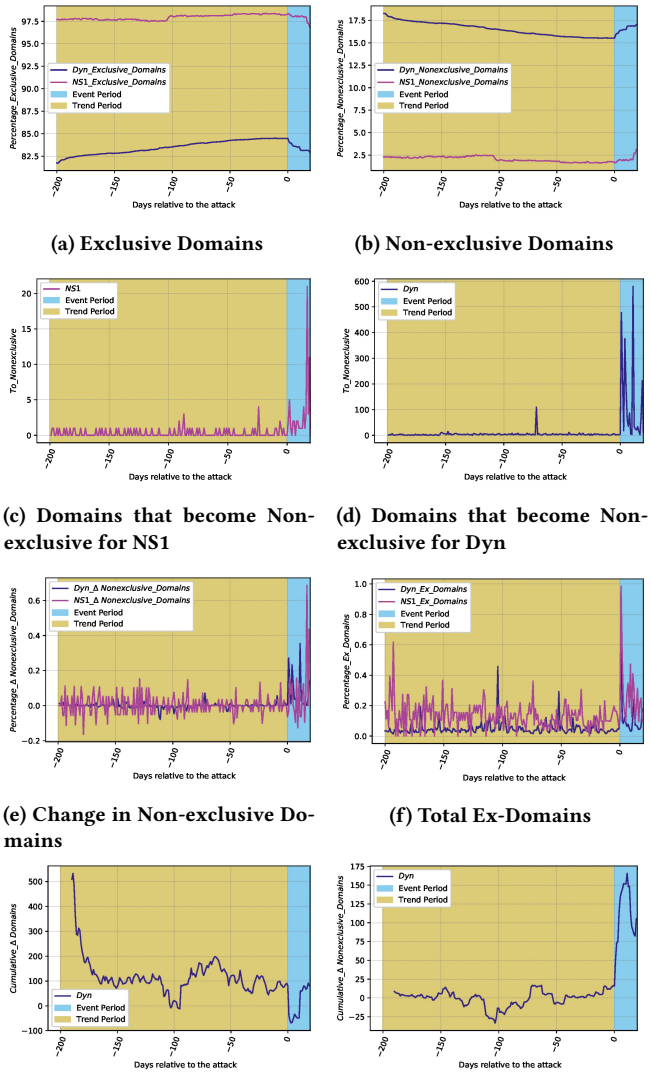

(a) Exclusive Domains

(b) Non-exclusive Domains

(c) Domains that become Non-exclusive for NS1

(d) Domains that become Non-exclusive for Dyn

(e) Change in Non-exclusive Domains

(f) Total Ex-Domains

(g) Cumulative change in Domains for Dyn

(h) Cumulative change in Non-exclusive Domains for Dyn

Figure 5: Time Series of Behaviour Variables

4 ANALYSIS AND RESULTS

We study the change in behaviour of the domains in three stages. First, we present the time series analysis of behaviour variables in Section 4.1. Then, we examine the statistical significance of the changes observed, in Section 4.2. Finally, we study the choice of secondary DNS service provider for the domains that become non-exclusive in Section 4.3.

4.1 Observations

Due to the DDoS attacks on May 16th and October 21st, 2016 the service provided by NS1 and Dyn respectively was interrupted and

the availability of the domains that used NS1/Dyn was threatened. In this section we discuss the interpretation of the changes in time series of the behaviour variables observed during the event period for both NS1 and Dyn.

Figures 4a and 4b show a sudden drop in number of domains using NS1 and Dyn just after the DDoS attack. These figures also show that NS1 was a much smaller MDNS (in terms of number of domains) than Dyn. The drop in the case of Dyn is much more observable than in the case of NS1. The drop shows that some domains stopped using the services of NS1 and Dyn after the attack and moved to a different MDNS. However, we observe a recovery pattern after the attack as the total number of domains starts increasing again a day after the attack. This indicates that some of the domains that stopped using NS1/Dyn, return when the services provided by NS1/Dyn are no longer affected due to a DDoS attack.

On the other hand, we do not see a similar recovery pattern for *Exclusive_Domains* (Figure 5a). The lack of recovery pattern for domains that use NS1 and Dyn exclusively can be attributed to the sudden and continuous rise in the number of domains using NS1 and Dyn non-exclusively. This sudden rise in the number of non-exclusive domains can be seen in Figure 5b. This shows that exclusive customers start using services from additional providers (become non-exclusive) in order to diversify the risk posed by DDoS attacks on their MDNS provider. The increase in the number of domains using NS1 and Dyn non-exclusively can be more clearly observed with the help of Figures 5c and 5d respectively. The notable change in preference of domains from using NS1/Dyn exclusively to non-exclusive use of their services after the attack can be clearly observed in Figure 5e. The percentage of total domains that choose to be non-exclusive in a single day in the event period is considerably higher than the trend period for both attacked MDNS providers.

Figure 5f shows a large number of domains leaving NS1/Dyn after the attack. During the event period, in case of NS1, 63.5% of the total domains that left using its services were exclusive users. In case of Dyn, 96.7% of the total users that stopped using its services during the event period were exclusive.

Zooming in on the larger of the two attack events, on Dyn, we can see the severity of the impact of the DDoS attack on Dyn with the help of the time series of cumulative variables. In Figure 5g we observe a strong negative cumulative impact on the total number of domains using Dyn in the event period (relative to the trend period). The only negative dip in the trend period can be attributed to a large number of non-exclusive domains leaving Dyn in the period 80 to 120 days before the attack (July-August 2016) as seen in Figure 5h. Contrastingly, in the event period we observe a sharp increase in the number of non-exclusive domains in Dyn. This behaviour is consistent and helps us re-emphasise the fact that domains tend to become non-exclusive users of an MDNS provider after the attack.

4.2 Statistical significance of the change in behaviour variables

With the help of the time-series plots we can observe the changes in the behavioural variables. In this section we test for statistical significance of the changes observed in the time series for both MDNS providers. The null hypothesis considered to examine the change in behaviour of the domains is as follows:

Variable	Trend Period Mean		Event Period Mean		t-statistic	
	Dyn	NS1	Dyn	NS1	Dyn	NS1
Δ *Domains*	127.05	6.87	-9.545	3.42	2.229*	1.45
Δ *Exclusive_Domains*	126.985	6.80	-127.82	1.42	3.16*	2.18*
Δ *Nonexclusive_Domains*	0.065	0.07	118.27	2	−3.341*	−1.42
Ex_Exclusive	66.63	2.85	212.59	5.47	−2.595*	−2.02*
Ex_Nonexclusive	10.68	0.24	7.682	3.19	1.93	−7.32*
New_Exclusive	194.29	9.68	195.4	8.90	-0.057	0.40
New_Nonexclusive	10.07	0.29	15.32	3.19	−2.49*	−8.1*
To_Nonexclusive	3.8	0.3	114	3	−3.12*	−2.57*
To_Exclusive	3.1	0.27	3.36	1	-0.44	-5.1

*p -value ≤ 0.05

Table 2: Results of T-test on behavioural variables

H_{a1}: There is no change in the behaviour of domains that use an MDNS provider after a DDoS attack.

In context with the measurement variables considered in this study we can reformulate the null hypothesis as follows:

H_{a2}: There is no change in the mean of behaviour variables in the trend and the event periods.

We evaluate the null hypothesis by comparing the mean values of behavioural variables for both MDNS providers in the trend and the event period with the help of a t-test [25]. We consider the change in variables with a p -value ≤ 0.05 to be statistically significant. Table 2 shows the test statistics for each variable.

We find that the mean values for the change in total domains and change in exclusive domains during the trend period were significantly (statistically) higher than during the event period. The negative mean values for daily change in domains and daily change in exclusive domains shows that domains leave Dyn after the attack in the event period. On the other hand, the number of domains using Dyn non-exclusively witness a significant growth in the event period. We notice a similar statistically significant increase in the non-exclusive users on NS1.

We also find the change in variable *Ex_Exclusive* to be statistically significant for both Dyn and NS1. This demonstrates that an abnormally large number of exclusive domains stopped using their services in the event period.

We do not observe any change in the average number of new exclusive domains joining the attacked MDNS in the event period. However, we notice an abnormally large number of new non-exclusive domains joining the MDNS. This can be an indication that a number of exclusive domains that leave Dyn after the attack returned as non-exclusive. Looking at the results of the t-test for variables *To_Nonexclusive* and *To_Exclusive* we can say that a large number of exclusive domains became non-exclusive in the event period but the trend of non-exclusive domains becoming exclusive did not really change.

4.3 Choice of Secondary DNS

Given that a significant number of Dyn and NS1 customers become non-exclusive users, we also analysed which secondary providers they choose. In order to understand the choices made by non-exclusive domains using NS1 and Dyn before the attack we evaluate the secondary `NS` addresses in the `NS` records of these domains one day before the attack. Figures 6a and 6b show the top secondary DNS choices for non-exclusive domains using NS1/Dyn one day before each attack. Most of the very few non-exclusive domains of NS1 used another MDNS provider for secondary DNS. However, in

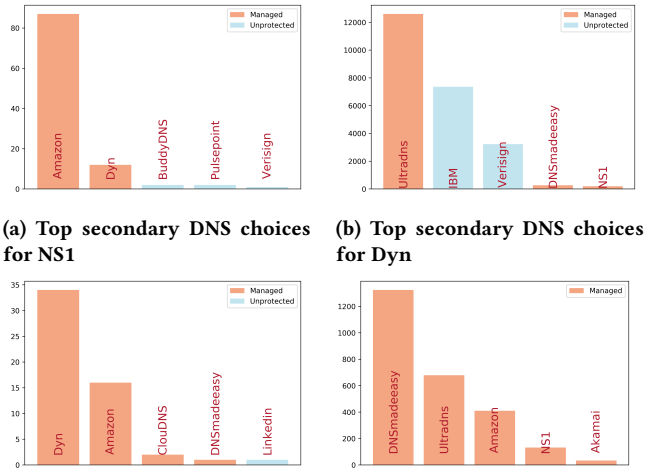

(a) Top secondary DNS choices for NS1

(b) Top secondary DNS choices for Dyn

(c) Top changes after attack NS1

(d) Top changes after attack Dyn

Figure 6: Secondary DNS choices for attacked MDNS

the case of Dyn we see that a remarkable number of domains used non-managed DNS service providers as a secondary choice.

After the DDoS attacks, we can observe with the help of Figures 6c and 6d that most of the users of NS1 and Dyn that became no-exclusive over a period of 20 days after the attacks added another MDNS service provider. Since, it is highly unlikely that two MDNS service providers fail due to a DDoS attack at the same time it underlines the fact that in terms of risk management, using multiple providers is a good strategy.

5 RELATED WORK

Distributed Denial-of-Service (DDoS) attacks have been the subject of intense study. Studies of the technical aspects of DDoS attacks have shown that there are myriad strategies for conducting an attack. The booter phenomenon has made DDoS attacks accessible to every one [26]. Studies have also shown reflection and botnet based attacks to be extremely effective [34]. Characterisation of DDoS attacks has been done by studies on the basis of intensity, source and event ports [13, 15, 20, 29]. At the same time, various DDoS mitigation techniques have been suggested by multiple researchers [19, 35]. Studies have also been conducted in order to evaluate the effectiveness of mitigation techniques [10].

Focusing specifically on the DNS, Moura *et al.* [21] evaluate the Nov. 30 and Dec. 1, 2015 events on the Root of the DNS. They show that large attacks can overwhelm some sites of some root letters. In addition, they also provide evidence that high traffic on one service can result in collateral damage to other services, possibly in the same data centre. In the event analysed in that study the overall DNS service was resilient to the DDoS attack. In case of the events evaluated by us in this paper (Dyn and NS1 attack), the overall DNS service provided by Dyn and NS1 was not able to absorb the attack.

Jonker *et al.* [14] study the adoption of DDoS protection services in general, using active DNS measurements. They observe that there are generally three strategies for mitigation in the face of an attack, two that use redirection via the DNS and one that redirects traffic using the Border Gateway Protocol (BGP). Furthermore, Jonker *et al.* observe that there are two general types of customer behaviour: one group of customers uses on-demand DDoS protection, only switching it on in case of an actual attack. The other group chooses to enable DDoS protection permanently, always routing traffic via the DDoS protection service. In this paper, we study a particular case of the latter, in which the DNS for a customer is supposed to always be protected against DDoS attack by making use of Dyn or NS1's managed DNS service.

Finally, industry reports [7, 33] from DDoS protection firms have studied the impact of DDoS attacks on the customers of Dyn. But in contrast to the framework used by us in this paper, they did not consider the domain segmentation (exclusive and non-exclusive) or the return behaviour of domains. These measurements form an integral part of such an analysis.

To the best of our knowledge, our paper is the first to empirically measure the direct impact of a successful DDoS attack on the behaviour of the victim's customers.

6 CONCLUSIONS

In this paper, we set out to study the effects of a successful DDoS attack on a managed DNS provider. Using data from the OpenINTEL platform, we test if the fallout of a successful attack results in changes in customer behaviour. We introduce a novel framework that measures the decisions a customer of an MDNS provider can take. We then use this model to analyse the change in customer behaviour after a successful DDoS attack.

According to the observations from our datasets, we can identify two types of customer behaviour. Most Dyn and NS1 customers use the MDNS *exclusively*, that is: they only configure authoritative name servers provided by Dyn or NS1 for their domains. A small, but non-trivial fraction of customers use the MDNS services *non-exclusively*. This means that they configure some of the authoritative name servers for a domain to be from Dyn/NS1, and some from other providers, or managed by themselves. In the period leading up to the attack, we observe a gradual growth in the use of services provided by both Dyn and NS1. Furthermore, we observe no significant changes in customer behaviour from using Dyn/NS1 exclusively to non-exclusively for both existing and new customers. If we then focus on the aftermath of the attack, we observe a number of statistically significant changes:

- A significant number of MDNS customers that were using Dyn's or NS1's service exclusively switch to non-exclusive use in the aftermath of the attack. Furthermore, our analysis shows that in most cases this change is lasting, that is: in the period analysed the majority of domains that switch from exclusive to non-exclusive remain in that configuration.
- We observe no significant changes in the behaviour of Dyn customers that were already non-exclusive users. While this result was to be expected – since they were likely not affected by the attack – it underlines the fact that in terms of risk management, using multiple providers is a good strategy.
- Lastly, we observe that most of the newly non-exclusive customers after the attack on Dyn and NS1 use an MDNS service provider as a secondary DNS to further reduce the risk of downtime.

Summarising, our study shows that our model captures significant changes in customer behaviour in the wake of a large, successful DDoS attack on a provider whose business model includes protecting customers against such attacks. Furthermore, these changes in behaviour are not just temporary, but we observe lasting changes in customer behaviour and permanent loss of customers.

7 FUTURE WORK

In this paper we showed *that* there is a change in customer behaviour, and especially that customers choose to hedge their bets by starting to use multiple managed DNS service providers. The next step is to understand *why* customers change their behaviour, and especially why they make specific choices, such as starting to use multiple providers. Intuitively, one might assume that using more than one provider leads to a cost increase, so it would be valuable to understand if this is the case, and if so, what rationale customers have to make this choice, and whether they have an upper bound on an increase in cost. To study this, we believe it is necessary to conduct a qualitative study, where decision makers at organisations affected by an attack are consulted about their decision-making process. The outcome of such a study may also be valuable for future decision making when organisations plan to outsource DNS to a managed DNS provider, and may even have wider applicability in cloud outsourcing strategies. It is clear from the examples of NS1 and Dyn that when taking into account that even large providers may be taken down by DDoS attacks that there are serious risks when outsourcing to a single provider.

ACKNOWLEDGMENTS

This work is part of the NWO: D3 project, which is funded by the Netherlands Organization for Scientific Research (628.001.018).The research leading to the results presented in this paper was made possible by OpenINTEL, a joint project of SURFnet, the University of Twente and SIDN. We would also like to thank Dr. Reinoud Joosten for his comments on the statistical analysis.

REFERENCES

[1] 2015. Worldwide Infrastructure Security Report, Arbor Networks.
[2] 2016. Worldwide Infrastructure Security Report, Arbor Networks.
[3] 2017. Worldwide Infrastructure Security Report, Arbor Networks.
[4] Abhishta, Reinoud Joosten, and Lambert J.M. Nieuwenhuis. 2017. Analysing the Impact of a DDoS Attack Announcement on Victim Stock Prices. In *2017 25th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)*. IEEE Press. https://doi.org/10.1109/PDP.2017.82
[5] Jay Barney, Mike Wright, and David J Ketchen Jr. 2001. The resource-based view of the firm: Ten years after 1991. *Journal of management* 27, 6 (2001), 625–641.
[6] Kris Beevers. 2016. A note from NS1's CEO: How we responded to last weeks's major, multi-faceted DDoS Attacks. Blog. http://ns1.com/blog/how-we-responded-to-last-weeks-major-multi-faceted-ddos-attacks
[7] Rich Bolstridge. 2016. Dyn DDoS Attack: Wide-Spread Impact Across the Financial Services Industry (Part 1). Blog. https://blogs.akamai.com/2016/10/dyn-ddos-attack-wide-spread-impact-across-the-financial-services-industry-part-1.html
[8] Huseyin Cavusoglu, Birendra Mishra, and Srinivasan Raghunathan. 2004. The Effect of Internet Security Breach Announcements on Market Value: Capital Market Reactions for Breached Firms and Internet Security Developers. *International Journal of Electronic Commerce* (2004).
[9] J Joseph Cronin, Michael K Brady, and G Tomas M Hult. 2000. Assessing the effects of quality, value, and customer satisfaction on consumer behavioral intentions in service environments. *Journal of retailing* 76, 2 (2000), 193–218.
[10] Christoph Dietzel, Anja Feldmann, and Thomas King. 2016. Blackholing at ixps: On the effectiveness of ddos mitigation in the wild. In *International Conference on Passive and Active Network Measurement*. Springer, 319–332.
[11] Michel Ehrenhard, Bjorn Kijl, and Lambert Nieuwenhuis. 2014. Market adoption barriers of multi-stakeholder technology: Smart homes for the aging population. *Technological Forecasting and Social Change* 89, Supplement C (2014), 306 – 315. https://doi.org/10.1016/j.techfore.2014.08.002
[12] Scott Hilton. 2016. Dyn Analysis Summary Of Friday October 21 Attack. Blog. https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/
[13] M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto, and A. Dainotti. 2017. Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem. In *Internet Measurement Conference (IMC)*. ACM.
[14] Mattijs Jonker, Anna Sperotto, Roland van Rijswijk-Deij, Ramin Sadre, and Aiko Pras. 2016. Measuring the Adoption of DDoS Protection Services. In *Proceedings of ACM SIGCOMM Internet Measurement Conference 2016*. ACM Press, Santa Monica, CA, USA.
[15] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. 2015. Amppot: Monitoring and defending against amplification ddos attacks. In *International Workshop on Recent Advances in Intrusion Detection*. Springer, 615–636.
[16] Brian Krebs. 2016. KrebsOnSecurity Hit With Record DDoS. https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/
[17] Robert F. Lusch, Stephen L. Vargo, and Mohan Tanniru. 2010. Service, value networks and learning. *Journal of the Academy of Marketing Science* 38, 1 (01 Feb 2010), 19–31. https://doi.org/10.1007/s11747-008-0131-z
[18] A. Craig MacKinlay. 1997. Event Studies in Economics and Finance. *Journal of Economic Literature* 35, 1 (1997), 13–39.
[19] Jelena Mirkovic and Peter Reiher. 2004. A taxonomy of DDoS attack and DDoS defense mechanisms. *ACM SIGCOMM Computer Communication Review* (2004).
[20] David Moore, Colleen Shannon, Douglas J Brown, Geoffrey M Voelker, and Stefan Savage. 2006. Inferring internet denial-of-service activity. *ACM Transactions on Computer Systems (TOCS)* 24, 2 (2006), 115–139.
[21] Giovane C.M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Muller, Lan Wei, and Cristian Hesselman. 2016. Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In *Proceedings of the 2016 Internet Measurement Conference (IMC '16)*. ACM, 255–270. https://doi.org/10.1145/2987443.2987446
[22] Moritz Müller, Giovane C. M. Moura, Ricardo de O. Schmidt, and John Heidemann. 2017. *Recursives in the Wild: Engineering Authoritative DNS Servers*. Technical Report ISI-TR-720. Available: https://www.isi.edu/~johnh/PAPERS/Mueller17a.pdf. USC/Information Sciences Institute.
[23] Pierluigi Paganini. [n. d.]. ProtonMail paid a $6000 Ransom to stop DDoS AttacksSecurity Affairs. http://securityaffairs.co/wordpress/41775/cyber-crime/protonmail-paid-ransom-ddos.html
[24] Janene Pieters. 2015. Ziggo: More Cyber Attacks Expected. Blog. https://nltimes.nl/2015/08/20/ziggo-cyber-attacks-expected
[25] Thomas J Rothenberg. 1984. Approximating the distributions of econometric estimators and test statistics. *Handbook of econometrics* 2 (1984), 881–935.
[26] J. J. Santanna, R. van Rijswijk-Deij, R. Hofstede, A. Sperotto, M. Wierbosch, L. Z. Granville, and A. Pras. 2015. Booters: An analysis of DDoS-as-a-service attacks. In *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*. 243–251. https://doi.org/10.1109/INM.2015.7140298
[27] Abraham Savitzky and Marcel JE Golay. 1964. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry* 36, 8 (1964), 1627–1639.
[28] Paul P Tallon, Kenneth L Kraemer, and Vijay Gurbaxani. 2000. Executives' perceptions of the business value of information technology: a process-oriented approach. *Journal of Management Information Systems* 16, 4 (2000), 145–173.
[29] Daniel R Thomas, Richard Clayton, and Alastair R Beresford. 2017. 1000 days of UDP amplification DDoS attacks. (2017).
[30] R. van Rijswijk-Deij. 2017. *Improving DNS Security: A Measurement-Based Approach*. Ph.D. Dissertation. University of Twente.
[31] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras. 2016. A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements. *IEEE Journal on Selected Areas in Communications* 34, 6 (June 2016), 1877–1888. https://doi.org/10.1109/JSAC.2016.2558918
[32] Verisign. 2017. VERISIGN: MANAGED DNS SERVICES. https://www.verisign.com/assets/pdf/resource-center/datasheet-mdns-overview.pdf
[33] Stephanie Weagle. 2017. Financial Impact of Mirai DDoS Attack on Dyn Revealed in New Data. Blog. https://www.corero.com/blog/797-financial-impact-of-mirai-ddos-attack-on-dyn-revealed-in-new-data.html
[34] Arne Welzel, Christian Rossow, and Herbert Bos. 2014. On measuring the impact of DDoS botnets. In *Proceedings of the Seventh European Workshop on System Security*. ACM, 3.
[35] Saman Taghavi Zargar, James Joshi, and David Tipper. 2013. A survey of defense mechanisms against distributed denial of service (DDOS) flooding attacks. *IEEE Communications Surveys and Tutorials* (2013).
[36] Valarie A Zeithaml, Leonard L Berry, and Ananthanarayanan Parasuraman. 1996. The behavioral consequences of service quality. *the Journal of Marketing* (1996).