Booted: An Analysis of a Payment Intervention on a DDoS-for-Hire Service

Ryan Brunt
ryan.brunt@nyu.edu

Prakhar Pandey
prakhar.pandey@nyu.edu

Damon McCoy
mccoy@nyu.edu

New York University
Brooklyn, NY 11201

ABSTRACT

Distributed Denial of Service (DDoS) attacks have become a growing threat that, to a large extent, have become commoditized by DDoS-for-hire, or "booter", services. In this case study, we analyze leaked fine-grain "ground truth" data from a larger booter service, VDoS, which earned over $597,000 over two years and launched 915,000 DDoS attacks and 48 attack years (i.e., the amount of DDoS time faced by victims of VDoS). The time period of the data includes data before and after a payment intervention, providing a rare opportunity to understand how it impacted VDoS's operation. We find that VDoS's revenue and subscriber base were growing before this payment intervention and began to decline afterwards. Our analysis shows that few existing customers switch from a regulated payment method to Bitcoin. We also find that harm from VDoS in terms of attacks launched and attack time both decrease by 40% (40,000 fewer attacks and 2 fewer attack years per month) during the payment intervention. However, VDoS likely remained profitable, albeit less, until the end of our data.

1. INTRODUCTION

Cybercrime has become increasingly efficient by leveraging underground specialists selling abusive capabilities, services, and resources [27]. Understanding the economic structure of these businesses [17, 20] and impact of prior interventions [19, 25, 28, 29] could potentially assist in prioritizing more effective interventions [6]. There are a host of inference methods that can supply indications of intervention impact [9, 11, 25], but there remain many unanswered questions. Unfortunately, there is little "ground truth" data for evaluating these inference methods and providing finer-grained measurements of interventions.

This paper is an analysis of rare ground truth data from a DDoS-for-hire service that covers a time period which fortuitously encompasses before and after data from an intervention. In particular, we analyze leaked and scraped data spanning two years of operation at VDoS, which was a large-scale DDoS-for-hire service [14]. We find that VDoS earned $597,862 in revenue over two years and launched 915,287 DDoS attacks over one year with a duration of 48.2 attack years.

Launching DDoS attacks has become highly commoditized by subscription based DDoS-for-hire services, colloquially referred to as "booters" or "stressors" in underground marketplaces. For $5 a month, a technically unsophisticated attacker can launch as many DDoS attacks as they want [1].

However, this decoupling of the person launching a DDoS attack and the service that provides the infrastructure to perform the DDoS attack creates an opportunity to disrupt payments between clients and DDoS services. A prior study explored the efficacy of a payment intervention against a large number of booter services [11]. With the ground truth data available in this paper, we are able to provide comprehensive documentation on how this payment intervention impacted VDoS with respect to three key aspects:

Revenue: We find that VDoS's revenue was increasing and peaked at over $42,000/month for the month before the start of PayPal's payment intervention and then started declining to just over $20,000/month for the last full month of revenue. Once VDoS could no longer accept PayPal payments they attempted to reestablish a regulated payment channel, but these attempts were largely unsuccessful. For the last month of the data they only accepted Bitcoin payments. However, even at the end of our analysis, VDoS was likely profitable.

Subscribers: VDoS's paid subscriber base peaked the same month as their revenue at 1,781 active paid subscribers and steadily decreased to 692 during the last month of our data. Our analysis found that only 316 (11%) of their repeat customers switched from a regulated payment method that accepted credit and debit cards to Bitcoin. We found an even stronger payment channel affinity with only 70 (2%) customers that paid with Bitcoin switching to a regulated payment channel. This indicates there are likely two segments of subscribers and intuitively accepting both a regulated payment channel and Bitcoin yields the largest subscriber bases.

Attacks: We find for the most part that attack volumes track revenue with a month delay, since customers pay for their subscription upfront. Our analysis shows that as VDoS's revenue and active subscriber base dwindled, so did the amount of harmful DDoS attacks launched by VDoS. The peak attack time we found was slightly under 100,000 attacks and 5 attack years per month when VDoS's revenue was slightly over $30,000/month. This decreased to slightly under 60,000 attacks and 3 attack years during the last month for which we have attack data. Unfortunately, we have incomplete attack data and likely missed the peak of VDoS's attack volume. However, the payment intervention correlates to a 40% decrease in attack volume, which equates to 40,000 fewer attacks and 2 fewer attack years per month.

[1]These attacks are normally effective at disrupting residential and lower-end hosting internet connections without DDoS protection.

This paper represents our analysis of data from a single booter service over a limited time period. Thus, while we can perform fine-grained analysis normally not possible with inferential methods, it might not generalize to other booter services and likely does not generalize to other criminal hacking-for-hire services. However, to the best of our knowledge, ours is the first paper to analyze fine-grained data from a criminal service that includes data both before and after a payment intervention.

The rest of the paper is structured as follows. In section 2, we present background on booter services and VDoS. Section 3 presents related work and section 4 presents a description of our data set and validation. Next, we present our analysis of the fine-grained data in section 5. In section 6, we discuss the implication of our study and future work. Finally, we conclude in section 7.

2. BACKGROUND

DDoS-for-Hire services, commonly called "booters," "stressors," or "stress testers" operate on a subscription basis. A customer or subscriber [2] of a booter service can typically manually launch an unlimited or daily quota of DDoS attacks. The attack durations range from 30 seconds to several hours and are limited to one to four concurrent attacks for one to three months depending on the tier of subscription purchased. The cost of these subscriptions range from $5 to $300. Many booter services build their attack infrastructure using rented Virtual Private Servers (VPS) from hosting companies with cheap or unmetered bandwidth and lax policies against sending spoofed packets and launching DDoS attacks [11].

Booter services also operated front-end web servers that are serve as the interface between customers and the booter service. These front-end web servers enable, payment of subscriptions, notification of newly added features, status of the booter service, customer support through a standard help desk ticking system, and launching of DDoS attacks. Customers can log into the front-end web server and launch attacks from a few Mbps to a few Gbps by typing in the IP address or domain name of their victim and clicking a button. The default is normally an amplified volumetric attack, such as DNS, NTP, or SSDP amplification attack. By far the most popular method was DNS amplification, which accounted for 53% of VDoS attacks with the next most popular method, NTP, accounting for just 8%. However, these services often also support SYN flooding attacks and layer 7 HTTP attacks. Recently some of these services have begun adding what they claim are more difficult to filter attacks for sites that are using anti-DDoS protection services. As part of this study we did not investigate how these newer attack methods are implemented.

VDoS also sold "API access" to their backend attack infrastructure as a method of monetizing excess DDoS attack capacity. This allowed other booter services to focus on advertising, customer support, payment processing, and their front-end design. These booters then purchase API access from VDoS to launch attacks for their customers using VDoS's attack infrastructure through an API designed by VDoS.

[2]We use these two terms (customer/subscriber) interchangeably in this paper.

When there was no excess capacity VDoS subscriber's attacks were prioritized over third-party booter services customers' attacks.

While booters often have Terms of Service (ToS) that forbid attacking unauthorized servers [5], they are unenforced except in the rare instance as an excuse to shed a customer that is excessively launching attacks [3]. It is common knowledge on underground forums, where these service are marketed, that booters allow subscribers to launch DDoS attacks against unauthorized targets. This thin veneer of legitimacy has also not prevented booter service operators and customers from being arrested and found guilty [**?**]. These booter services effectively commoditized the ability to disrupt unprotected online services by cheaply providing powerful DDoS attacks to unsophisticated attackers.

Karami et al. [11] monitored booters who accepted PayPal for 6 weeks, tracking the merchant accounts and payment methods used by these services. At the end of 6 weeks they found that 23 booters were able to accept PayPal for at least 3 of the weeks. These booters varied in their approach to using PayPal. Some maintained a single merchant account for the entire measurement while others changed accounts approximately every 5 days. After the initial measurement period, the domains and merchant accounts were reported to PayPal. The measurements of the effectiveness of that study were largely qualitative in nature, since it lacked fine-grained data. After the study, PayPal continued to proactively identify and limit booter accounts (including VDoS's). The funds in limited PayPal accounts are frozen, which prevents the account holder from withdrawing the remaining account balance. This seizure of frozen funds coupled with the cost of establishing replacement PayPal accounts each time one becomes limited ultimately drove booters, such as VDoS, to abandon PayPal all together.

VDoS was one of the larger DDoS-for-hire services that was estimated to be earning over $17,000 per month based on scraped data from December 2014 to February 2015 [11] [4]. It operated from at least July 2014 until September 2016, when it closed down following the arrest of the two primary owners [13]. Before the service ceased operation, there were two public leaks in July 2016 of VDoS's operational backend database [14] and the HTTP logs from their attack server [1]. Based on our analysis of VDoS's database, they completely stopped accepting PayPal payments on January 1, 2016. However, using the Internet Archive [5] we found that between September 23, 2015 and October 20, 2015 PayPal was no longer a payment option listed on vdos-s.com (although options for direct Visa and Mastercard payments were added). From the ticket database we found instances of PayPal limiting VDoS's merchant account as far back as December 2014.

[3]The leaked database used in this analysis includes customer support tickets opened when a VDoS user needed assistance. Based on analysis of messages in VDoS's ticketing system, we found many instances of subscribers indicating they were attacking unauthorized servers without repercussions. Ironically, we also found one instance of a user that launched over one thousand attacks being banned on the pretenses of violating this policy.

[4]Our analysis of leaked ground truth data shows the actual revenue of VDoS over this period was $53,423, indicating this method of estimating revenue was accurate to within 4.5%.

[5]http://archive.org/web/

In August 2015, VDoS support began replying to tickets regarding PayPal with a canned response:

> We are experiencing some issues with PayPal; I do not believe it will be coming back anytime soon. However, I highly recommend you purchase using Bitcoin.
>
> You can use the following website(s) to purchase Bitcoin:
> -https://localbitcoins.com
> -https://www.deepdotweb.com/buy-bitcoins-with-paypal-credit-cards/
>
> Alternatively, we also accept credit card payments. (Note: You cannot purchase our VIP plans with this method. Only Bitcoin)

3. RELATED WORK

In this section we will present some of the closer related work focused directly on understanding DDoS attacks, booter services, and economic analysis of cybercriminal operations.

3.1 Booters

Booter services have been investigated by numerous prior studies. Initially, the leaked database of twBooter was studied in 2013 which provided a description of the basic structure, scale, and exposed the illicit nature of these booter services that attempted to maintain an air of legitimacy [10]. Follow up studies attempted to enumerate a large number of these booter services [23], analyzed additional leaked booter databases [22], and attack traffic [24]. None of these previously analyzed leaked databases spanned time periods when the booter was subject to a focused payment intervention.

Additional studies have explored supervised machine learning methods for detecting booter websites [5], demographics of booter operators [7], deploying honeypots to monitor [12] and analyze amplified attacks [18], and attribute their attack server infrastructure [15]. We do not include an analysis of victims attacked by VDoS, since our preliminary analysis indicated that the victims were largely gaming servers and residential connections similar to these prior studies. Rather the focus of our study is on performing a longitudinal analysis of the technical and financial operation of VDoS based on fine-grained leaked and scraped data.

The closest work related to our study is a study by Karami et al. [11], which focused on measuring the impact of a payment intervention launched by PayPal to disrupt revenue collected by over 40 booter services. This study included leaked databases from two other booter services. However, the leaked database in this study spans a longer timeframe and includes data before and after the intervention. This allows us to answer questions that could not have been answered in the previous study (e.g., how the subscribers responded to the intervention). It also enabled us to validate that the revenue estimation method used in the prior study was accurate to within 4.5%.Thus, our study is complementary to this prior study and provides a quantitative analysis of the impact of PayPal's payment intervention on VDoS.

3.2 Economic Analysis

Anderson [2] was one of the first to explore the economics of security. Follow up studies have broadly quantified the economics of cybercrime [3, 4, 27] based on aggregating and summarizing data from macro level economic studies. These studies have found that an understanding of the economics and operations of cybercrime can be a valuable method for crafting more effective intervention strategies, which was synthesized by Clayton et al. [6].

There have been more narrow studies focusing on inferring the conversion rates of spam into sales [8] and estimating total revenue of spamvertized-based operations [9, 16]. A pair of studies investigated the effectiveness of website takedowns for mitigating phishing attacks [21] and counterfeit goods [29] and found that, while there is some impact, it is limited. Another study by Soska and Christin [25] found that takedowns of illicit darknet market places had little lasting effectiveness at undermining this market. An analysis of the end-to-end support infrastructure of spam found that there was a potential bottleneck in the payment processing for spamvertized goods [17]. Stone-Gross et al. [26] found that fake anti-virus operators manipulate refund rates to reduce chargebacks and maintain payment processing.

Most closely related to our study are two studies that inferred economics of illicit pharmaceutical operations based on leaked back-end databases [20] and measured the outside impact of payment interventions in undermining counterfeit software and pharmaceutical operations [19]. To the best of our knowledge, our investigation constitutes the first inside look into customer behavior during the transition from regulated to unregulated payment methods for subscriptions.

4. DATA

Our analysis in this study uses both leaked and scraped data from multiple sources. Table 1 provides a brief description of these data sources. When using leaked or third-party reported scraped data there are both ethical and validity concerns. In this section we will provide an overview of the datasets, discuss how we validated some parts of the data, and finally we will talk about the ethical framework of our study.

4.1 Description

Our analysis of the leaked database found that VDoS also operated several sister booter services, including cnBooter, vStress, and uStress. These similarly ceased operation in September 2016 after the arrest of these services' co-owners. These booters were far smaller in scale with only 15 paid subscribers between the three of them. For this reason, we do not include them in our analysis. We do not include a detailed description of all 22 tables in the database. However, we performed most of our analysis using 6 tables: *users* contains registration and account status information; *payments* has transaction records for subscriptions; *sent_payments* contains records of payments for hosting service and other expenses related to the operation of VDoS; *attacklogs* contains details of user initiated attacks; *tickets* and *replys* tables have user support tickets and employee responses.

The second source of data used in our analysis is derived from Apache HTTP access logs that were leaked from https://api.vdos-s.com. The leaked database mentioned in the paragraph above is from the frontend site https://vdos-s.com. The frontend site was used to manage subscriptions as well as an interface for users to launch attacks. The API server, on the other hand, was used as the command and

Source	Fields	Explanation
Leaked Database	Users, payments, attack logs, ticketing	Backend VDoS database
Leaked HTTP Logs	timestamp, duration, target, method	Apache access logs includes attacks initiated using API
Scraped Data*	timestamp, duration, method	VDoS posted attack information

Table 1: *Summary of datasets used in our analysis. * Scraped data collected by Karami et at.*

Field	Value	Description
Timestamp	03/Sep/2015:00:34:44 +0200	
Source	78.128.92.156	Attack server belonging to VDoS
Config	index.php	Script used to launch attack
Target IP	*	Redacted victim ip
Target Port	8055	victim port
Duration	1200	Length of attack in seconds
Method	dns	Attack Method, 23 unique methods

Table 2: *Example of attack data from Apache HTTP access log.*

control center for launching attacks using the rented attack servers. The users would submit a form on the frontend website with parameters for the attack (such as target ip, attack method, duration, etc.) which would be passed to the API server and the attack would begin. As [14] points out the API server was used by more booters than just VDoS (e.g., PoodleStresser). Additionally, we find cases of the owners running custom attack scripts using the API that do not appear in the frontend database. An example of the fields and values in the HTTP log are shown in Table 2. This allowed us to reconstruct more attack logs than those contained in the leaked database.

Our final source of data was provided to us by Karami et al. [11], and consists of scraped attack details and subscriber usernames that were publicly published as a running status of current ongoing attacks on the frontend subscriber website. We were able to use this data to reconstruct additional attack data.

Table 2 shows a high-level summary of the information we were able to derive about VDoS's operation using these three data sources. It shows that VDoS earned close to $600,000 over two years and had 10,000 paying subscribers that launched over 900,000 DDoS attacks against over 270,000 unique IP addresses for a total of 48 attack years over the one year of attack data that we were able to reconstruct. This shows the large-scale harm likely caused by this single booter service.

Since there are several data sources spanning different time ranges we will quickly recap what data we have. The leaked database has payment and subscription information that ranges from July 2014 through July 2016. This includes account details and payment details. Unfortunately, the attack data in the backend database had been deleted except for three months prior to the leak (May, June, and July of 2016). In order to fill in the missing attack data we turn to other sources. First, we have data scraped from the website by [11]. This covers December 2014 and January 2015. Second, we have HTTP logs from the API server used to launch attacks. These logs cover September 2015 through September 2016. So the only overlap between these datasets is the attack data which covers May, June, and July of 2016 in the backend database and the HTTP logs.

Total subscriber revenue	$597,862
Registered Users	75,321
Paying users with at least 1 attack	10,000
Avg. Active subscribers per month	970
One time subscriber	7,250
Recurring subscriber	2,964
Number of victims	272,741
Attack time	48.2 years
Attack count	915,287

Table 3: *Summary of data used in this analysis*

This overlap is useful for validation, but we utilize the scraped data and HTTP logs to analyze attacks covering a longer time period.

4.2 Ethics

All of the data we used for our study was at one time publicly available, albeit some of it was unintentionally made public via a leak. There is likely little if any Personally Identifiable Information (PII) in this leak and we did not find any during our analysis. This was a criminal service and the usernames are pseudonyms that are intentionally difficult to link to the actual persons. Our Institutional Review Board (IRB) exempted our study, since the leaks were public and we found no PII in the data. We did not attempt to deanonymize anyone in these leaks as part of our study. We also did not include any raw information such as usernames, email addresses, or IP addresses in this paper. If we had found any PII in these leaks, we would have contacted our IRB again and submitted a revised protocol for review. With this in mind, it should be noted that residual privacy risk remains in datasets like the one analyzed in this paper. Emails and customer support tickets stored in the database are particularly risky. Although we did not come across any PII in the tickets we analyzed, private information could be shared in these correspondences especially when dealing with refunding purchases. Furthermore, the pseudonyms used as usernames in the database are not always chosen at random. Users may use the same pseudonym across many sites and could potentially be deanonymized. VDoS did

Location	# IP's in Blacklist	# of Payments
Brazil	22	12
China	20	16
Europe	26	1
Middle East	4	1
North America	15	2
Australia	1	1

Table 4: *Residential IP addresses found in black list and number of payments received from each.*

monitor IP addresses of users to prevent account sharing. Our IRB does not consider IP addresses to be PII. However, IP addresses could potentially be used to deanonymize customers; therefore we did not make use of this data in our study.

4.3 Blacklisting Analysis

VDoS maintained a table of blacklisted IP addresses and domains for which subscribers could not launch attacks. VDoS charged people $2.50 per IP address to be blacklisted. On October 10, 2015, we found tickets indicating that they increased the price to $4. In total, extortion payments only earned VDoS $373. Payments for blacklisting were only accepted in Bitcoin likely since accepting extortion payments using revisable payment methods such as PayPal would have been unwise. Given that blacklisting payments were only ever accepted in Bitcoin, this analysis is outside the primary scope of our study on the effects of PayPal's payment intervention on VDoS. We include our detailed analysis of this extortion activity, since to our knowledge no prior study has analyzed this facet of a booter server.

VDoS used blacklisting as a means of extorting victims being attacked. We found 272 IP addresses that had been blacklisted. These can be broken down into 166 hosting services, 88 residential, 15 business, 2 content delivery networks, and 1 cellular IP. VDoS only received payments from 86 hosting servers (52%), 32 residential (36%), and 10 business (67%) generating a mere $373 [6]. We also looked at country of origin for residential addresses only, since these might be more indicative of the actual location of the extorted person, and found that the majority of payments came from Brazil and China. The full results can be seen in Table 4.

Some honesty was demonstrated by VDoS. Based on our analysis of tickets, if the victim's IP address did not appear in the attack database and the victim requested to be blacklisted, VDoS would simply inform them they had no need to be blacklisted. VDoS also blacklisted all Israeli IP addresses. As reported by Brian Krebs [14], several reasons were given for this to users attempting to attack a domain hosted in Israel. Users were told that Israel was blocked for "safety" or "security" reasons. Others were told more directly that the owners were from Israel and did not want themselves or their region to be the target of DDoS attacks. Additionally, they blacklisted Cloudflare's IP addresses. This blocking of Cloudflare's IP addresses might have been motivated by them using Cloudflare's services to protect their front-end. It might also be because VDoS's operators knew that their DDoS attacks would be ineffective against Cloudflare's anti-DDoS protection.

[6]Our classification of ISP type is based on the results from the MaxMind precision insights API: https://www.maxmind.com/en/geoip2-precision-insights

4.4 Validation

As mentioned previously, VDoS was compromised and their backend database and their HTTP Logs were publicly leaked. Our other data set is scraped attack information that was publicly published by VDoS and collected by Karami et al. [11]. As with any data of unknown provenance and more especially with data from a criminal operation, it is important to investigate the accuracy and validity of the records before drawing conclusions. The main concerns are whether the database is an accurate reflection of the site's operations and that data was not fabricated. Thus, before we begin discussing our findings, we will attempt to illuminate flaws in the data where they exist and proceed with our analysis only where the data permits.

Our three primary techniques for assessing the accuracy and validity of the data were: 1) Check tables in the database for consistency; 2) Cross check overlapping data from other source(s) used in our analysis when possible; 3) Validate data against outside sources of information.

Table Agreement

The users table from the leaked database has records for 75,321 unique users. These can be broken down into three groups: 8,290 Active, 64,639 Unapproval (users who registered but never confirmed their email), and 2,392 Banned (Users who violated VDoS terms of service, many of whom were paid subscribers). Included in the users table are the staff accounts (6 admin and 3 support, one of which was banned). We find our own accounts, which we registered for a prior study, were correctly included in the user table. The payment information and type of membership for our one paid account are also correct.

A large portion of the users never launched an attack (64,901). Of these users, only 190 made a payment. They likely purchased an account but never launched a boot. On the other hand, 10,420 users have at least one recorded attack in the database, and 420 (4%) cannot be found in the payments table. Inspecting the tickets database revealed that some of these payments were handled manually. VDoS would request a screenshot of the transaction and then grant membership. In addition, we found instances of free trial accounts being occasionally provided based on an analysis of support tickets. To be conservative, we do not count any of these 420 as paid subscribers.

Of the 10,000 with a boot and recorded payment, 9,427 of these had an expiration date listed (i.e., they were not lifetime members). 2% (217) of these users do not have a payment matching the expiration date. We determined this based on date of purchase and the recorded package bought by the user. If the purchase date plus the package time offset exceeded the expiration listed in the user table, then we consider it valid.

Out of the 217 accounts without a payment matching the expiration, 180 opened tickets with VDoS customer support. These tickets are typically associated with manually adding time to users' accounts due to service downtime (a Twitter post from VDoS account offers 4 free days for service

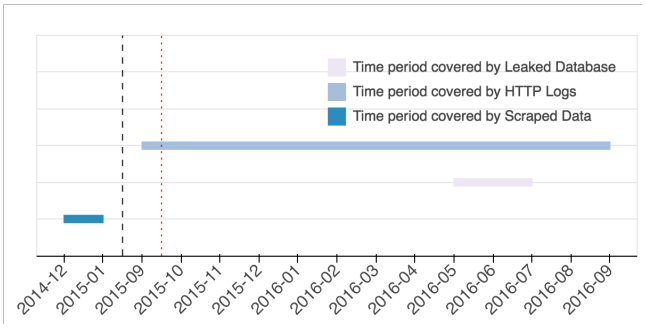

Figure 1: *Timeline of each source of attack information. Note there is a gap in the dates at the black dashed line. The red dotted line marks the culmination of the payment intervention when PayPal was removed as a payment option on the VDoS website.*

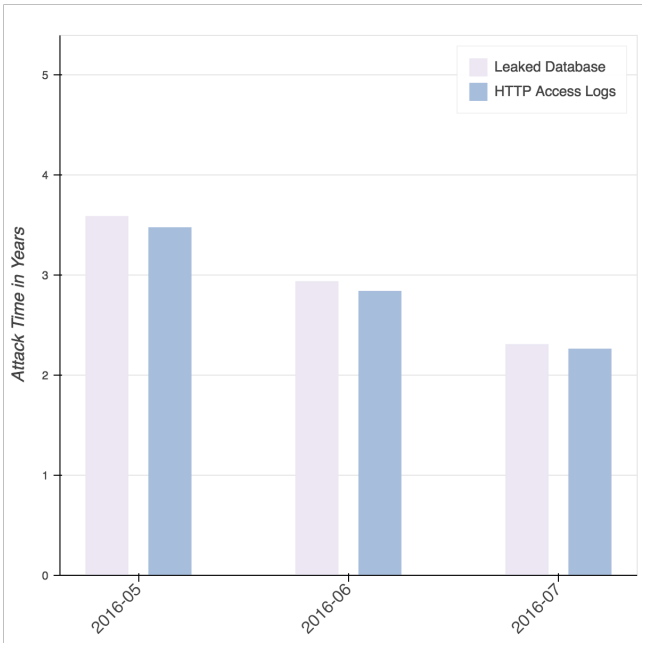

Figure 2: *This graph displays the total amount of attack time measured in years for the overlapping time period when we have both the leaked data base and HTTP logs.*

down time in February 2015 [7]), payment disputes (including refunds when PayPal was disrupted), or giving a free trial. Two account expirations were set as invalid dates, e.g., February 30th, and three had unknown expiration dates. We were unable to find a compelling explanation for the expiration dates on the remaining 32 accounts.

Additionally, 17 payments were made associated with unknown expiration times for 12 users. Several of these payments are from the same PayPal user, and all such accounts have an unapproval status.

There are 711 Lifetime accounts of which 584 are found in the payments table. However, only 573 of these have at least one attack. From this reduced set of active lifetime accounts we found 32 without record of paying for a lifetime account. Several of these accounts are administrator and support personnel who responded to tickets and were likely granted free accounts.

One interesting case is a user who purchased a lifetime account but had 11,896 Attacks. From an analysis of the ticketing data we can see that a VDoS administrator told this user that they were losing money on him and then banned him. The user then tried to make another account but was banned again since the IP address used to register was the same. We also found cases of accounts that were transferred to another person and banned. The values in the tables align with the information in the support tickets for these cases and other similar instances.

This indicates that the information in the tables is largely consistent and likely accurate within a small percentage. These errors were often due to manual updates that can be traced to communications in support tickets.

Attack Logs

We were able to extract 895,769 attack records from August 30, 2015 to September 1, 2016 from the leaked HTTP logs. The logs include more than just VDoS attack information since other booter services also used VDoS's API to launch attacks. Based on the IP address of the HTTP client sending the request, we find that 809,850 of the attacks were launched by VDoS. Poodle Stresser was another booter run by the creators of VDoS which leveraged the VDoS API to attack larger targets such as Blizzard. Unfortunately, all records from the attacklogs table in the leaked database appear to have been deleted in April 2016. Thus, the leaked database only contains 169,845 attacks that span from May 2016 to July 2016. Our final source of attack information comes from scraping data from the VDoS website where they posted attack information in a graph to flaunt their capacity. It covers December 1, 2014 until January 26, 2015. Figure 1 shows a timeline of which time periods are covered by our data sources and the overlap. Figures 2 and 3 show that each these data sources includes about the same amount and duration of attacks.

We were able to match the records in the API dataset and the attack logs recorded in the database during the overlapping time period to check for consistency. Limiting the API data to the date range of the database results in 165,992 attacks in the API data and 169,845 attacks in the database. We attempted to match the database attack records based on target, attack method, target port, and timestamp (within 10 minutes), which yielded 4,701 (3%) unmatched attacks in the API logs and 8,490 (5%) unmatched attacks in the database. The attacks missing from the database are possibly caused by the attack servers being overload at times or networking failures. The missing attacks from the HTTP logs suggests that not all attacks were launched via the API. Note that we cannot cross validate the API attack records that do not overlap with the attack records in the database. The scraped data was previously validated by matching NTP attack victims and self-attacks launched by the researchers [11].

Given the close match between the HTTP logs and the leaked database information, we believe this information is accurate enough to draw meaningful conclusions.

[7] https://twitter.com/vDosStresser/status/563733862761922560

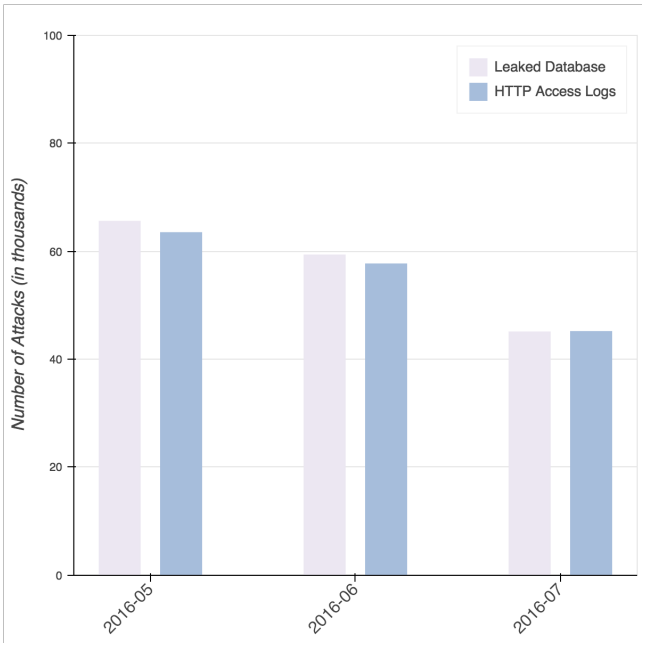

Figure 3: *This graph displays the total number of attacks for the overlapping time period when we have both the leaked data base and HTTP logs.*

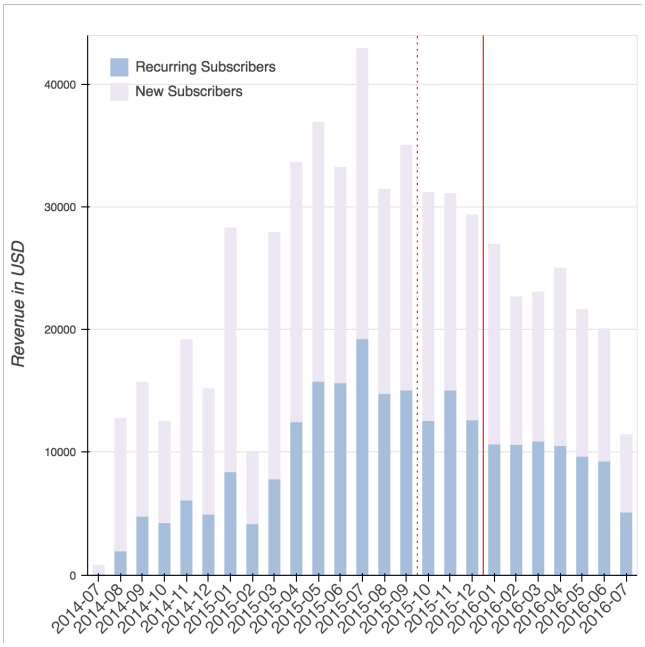

Figure 4: *Monthly revenue for VDoS broken down by recurring and new subscribers. Note that the red solid vertical line marks when PayPal was no longer accepted. The red dotted line marks the culmination of the payment intervention when PayPal was removed as a payment option on the VDoS website.*

We will perform the rest of the analysis in our paper using only the attacks in the scrape and HTTP Logs, since these provide attacks over the longest time period.

Bitcoin

Our final piece of validation focused on the users who paid with Bitcoin. For all Bitcoin payments, we checked if the transaction ID recorded in the database was confirmed on the blockchain. There are a total of 3,580 Bitcoin transactions in the database. Twenty-two transactions had invalid hashes, 19 had several transactions listed with the same hash, and 83 had a difference of more than two dollars between the amount on the blockchain and in the database. In total, 125 (3%) Bitcoin transaction records failed our validation test. We also analyzed the data to explore other incoming transactions to the primary address controlled by VDoS. We found 742 transactions not recorded in the database. Most of the discrepancies are likely an artifact of users paying through an intermediary address (hence a different transaction hash). Other Bitcoin payments were for their IP address blacklisting service, i.e., extortion.

This data provides us with an interesting population of Bitcoin users (e.g., those subscribing to a criminal DDoS service). Thus, we performed a quick analysis to understand what precautions, if any, were taken to remain anonymous. Only 29 users use the same wallet twice to make payments. Most users use a new wallet for each transaction, suggesting that a majority of the users are at least taking some steps to remain anonymous or using a third-party wallet service.

Summary

In summary, while we cannot rule out to possibility that this data was fabricated, it is mostly internally consistent. We were also able to validate a few pieces of data relating to our account with the service.

Thus, we believe with high confidence that the data is mostly valid, accurately represents the operation of VDoS, and is not a fabrication.

5. ANALYSIS

Using these data sets, we now provide a detailed assessment of the impact of PayPal's payment intervention and VDoS's transition from regulated payment methods to a largely unregulated payment method, Bitcoin. From the perspective of the data in these leaks, we consider revenue, customers, and attacks.

5.1 Revenue

VDoS was generating a median revenue stream of $25,985 each month over 24 months. The minimum revenue was $9,956 and the maximum $42,924. Figure 4 shows the breakdown of revenue from new and recurring subscribers, which made up $354,984 (59%) and $242,878 (41%) respectively. Their revenue showed steady growth for the first year followed by steady decline beginning in August 2015 after a payment intervention by PayPal [11]. VDoS's revenue declined from a high of $42,924 in July 2015 to $20,069 in June 2016, which is the last full month of information in the leaked database and 2 months before VDoS ceased operation.

Looking at the growth period from March 2015 to July 2015, revenue from subscribers paying via PayPal was relatively flat and most of the revenue growth came from customers paying with Bitcoin, as shown in Figure 5. This revenue from the Bitcoin payment channel increased from $2,674 and 10% of revenue to $12,652 and 29% of revenue over these five months.

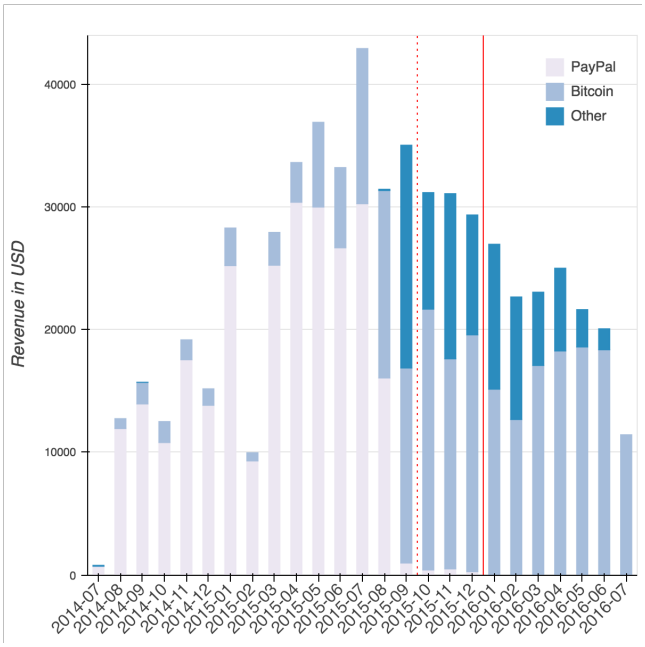

Figure 5: *Monthly revenue for VDoS by payment channel. Note that the red solid vertical line marks when PayPal was no longer accepted. The red dotted line marks the culmination of the payment intervention when PayPal was removed as a payment option on the VDoS website.*

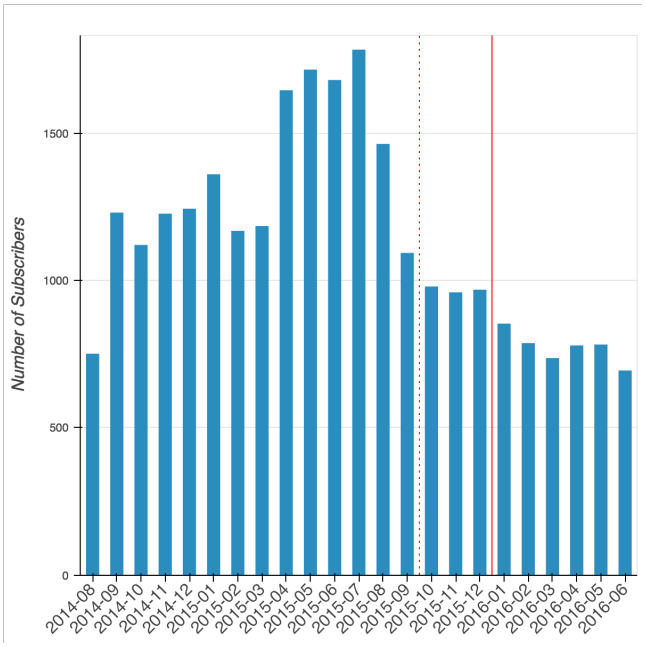

Figure 6: *Number of active subscribers each month. Active here means they were current on their payment subscription for that month. The red solid vertical line marks when PayPal was no longer accepted. The red dotted line marks the culmination of the payment intervention when PayPal was removed as a payment option on the VDoS website.*

Near the middle of August 2015[8], the payment intervention that limited VDoS's ability to accept PayPal payments [11] began to take its toll on VDoS. Disrupting VDoS's PayPal payment channel had a noticeable effect on both recurring and new revenue. By August 2015, payments from the PayPal channel decreased by $12,458 (44%) from an average of $28,523 over the previous five months. The Bitcoin payment channel increased by $6,360 (71%), but did not fully compensate for lost revenue from PayPal.

The next month, VDoS established a number of ad-hoc payment methods, such as other third-party payment processors that accept credit card payments. Most of these methods were short lived, likely due to the payment processors learning about the nature of their illicit DDoS service and terminating their accounts. The revenue from these other regulated payment channels dwindled over a ten month period from $18,167 in September 2015 to $1,700 during June 2016. The last month of the database leak in July 2016 shows no other forms payments other than Bitcoin.

Based on communications with PayPal, we know that some of VDoS's accounts were frozen. However, PayPal did not provide us with any estimates of the amounts in the frozen accounts. It is also likely that some of the payments accepted via other third-party payment processors were seized and not paid out to VDoS, which might have led to the abandonment of regulated payment channels.

[8]The exact timing of the intervention is unclear. Based on the support tickets being opened regarding payments this is when the support staff began redirecting users to Bitcoin. See the discussion at the end of section 2.

Virtually all third-party processors, including PayPal, hold back some of a merchant's revenue for weeks to months in case of chargebacks and fines.

Unfortunately, the information in the leaked database does not provide any insight into how much revenue from regulated payment channels was realized and how much was seized. Thus, we conservatively assume that none of it was seized. This causes it to be difficult to know the true lost revenue from no longer accepting regulated payment channels based on the leaked payment information. Ultimately, VDoS's operators' attempts to reestablish regulated payment channels demonstrates that they perceived this would increase their actualized profits. It is therefore safe to assume that PayPal's payment intervention and VDoS's subsequent inability to establish a new regulated payment channel decreased the scale and profitability of VDoS.

5.2 Subscriber Behavior

In order to understand the impact of PayPal's payment intervention, we also perform an analysis of the subscribers' payment behavior. There were 10,190 total paid subscribers (190 did not launch any attacks) of which 7,250 made only one payment. Of these one time customers, 5,531 paid using PayPal or some other regulated payment channel and 1,719 used Bitcoin. Figure 6 shows that the number of paid subscribers began to decrease once PayPal started their payment intervention in July 2015 and continued to fall throughout the rest of our dataset. This shows a strong preference for regulated payment channels, such as PayPal, over Bitcoin. Likely there are many reasons for this preference. For instance, one possible explanation is that regulated payment methods allow for recovery of funds if the purchased

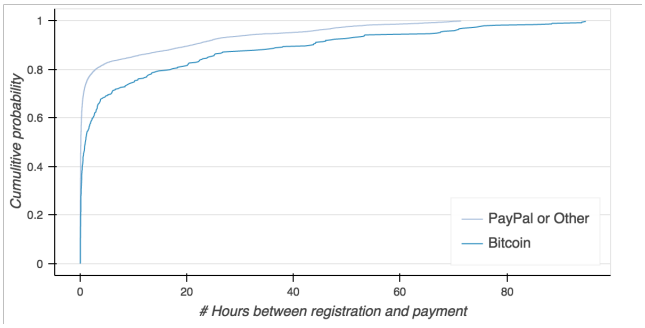

Figure 7: *This is the empirical CDF of the time between registration and initial payment. The darker blue line is users paying with Bitcoin for their first payment. This illustrates the delay caused by the difficulty of obtaining and using Bitcoin for payments. The median time is over an hour longer for Bitcoin users. Note that outliers were removed by cutting off anything above 1.5 times the Inter-Quartile Range.*

service is not delivered. Another reason might be the logistical and technical difficulties of purchasing and using Bitcoin. Figure 7 hints at some of the difficulties customers encounter by showing the increased median delay of 781% (2.7 hours) when using Bitcoin verses regulated payment channels.

2,964 subscribers made at least two payments and, of these, 2,069 customers exclusively paid via a regulated payment channel and 509 used only Bitcoin. Only 316 (11%) repeat subscribers initially paid using a regulated payment channel and then switched to Bitcoin, while only 70 (2%) repeat customers switched to regulated payment channels from Bitcoin. Again, repeat customers demonstrate a preference for regulated payment channels over Bitcoin. However, Figure 8 shows a slight increase of subscribers that switched from regulated fiat payment channels to Bitcoin after PayPal was largely unavailable.

The final subscriber payment behavior we analyzed is the registered user to paid subscriber conversion rate. Overall, registered users converted to paying users at a rate of about 13%. Before the PayPal cutoff in January 2016, the conversion rate was about 14% vs 11% after January 2016. Interestingly, the conversion rate dropped in August 2015 to 5%, when the site announced it would no longer be accepting PayPal, and then rebounded by the end of 2015. While the intervention did cause a disruption in the conversion rate, it was short lived.

5.3 Attacks

The attack data shows the scale of the attacks was significant, with a total of 915,287 attacks and 48.2 attack years over the one year of attack data that we analyzed [9]. This equates to an average of over 60,000 attacks launched and a total attack time exceeding 3 attack years each month.

As Figures 9 and 10 show, the attacks and attack time both increased from the initial scrape data that spans December 2014 to January 2015 and the leaked HTTP logs

[9]We noted that two attacks were launched via the API in September and October of 2015 with overly long duration times (31,710 years and 32 years respectively). We removed these two attacks from our analyses, since they were launched by administrators and obviously did not actually last for the entire scheduled duration.

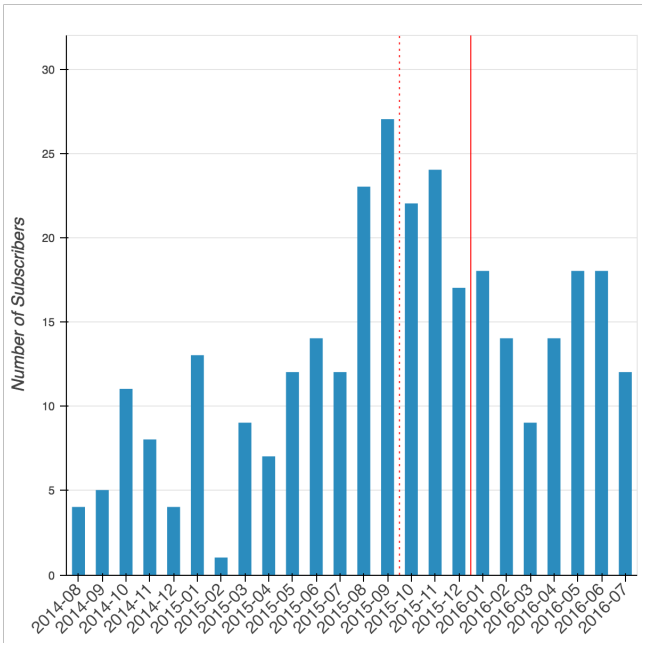

Figure 8: *Number of subscribers that switched from regulated fiat payment channels to Bitcoin. The sold red vertical line marks when PayPal was no longer accepted. The red dotted line marks the culmination of the payment intervention when PayPal was removed as a payment option on the VDoS website.*

which span from September 2015 to August 2016. It is noteworthy that attack volume (both number and time) dropped significantly in January of 2016 after the payment intervention, from an average of 79,203 attacks and 3.9 hours per month to an average of 54,782 attacks and 2.8 hours per month once PayPal was no longer accepted. This represents a 31% decrease in attacks and a 28% decrease in attack time.

There is a discrepancy between the drop in the number of active subscribers and attack time. Active subscribers began to decline almost in sync with the payment intervention, while the drop in attacks came about 3 months later. This possibly suggests that the initial users that dropped off were not launching many attacks. We begin to see a decline in attack time as the subscriptions of users that launched attacks expired and were not renewed.

Roughly speaking, revenue follows the number of attacks and attack time with approximately a one month lag, due to the fact that subscribers make payments upfront for subscriptions. Figure 11 shows the relationship between revenue and attack volume. This again demonstrates the effectiveness of the intervention which disrupted the payment infrastructure of VDoS. Attack number and attack time are a good metric to demonstrate the effectiveness of a payment intervention, since ultimately a decline in these represents a decrease in harm caused by VDoS.

5.4 Estimated Profit Margins

While there is some cost data in the database (sent_payments database table), it only contained data from July 28, 2015 to August 26, 2015. The total amount over this one month was $10,789. About $2,379 (22%) of this was for hosting services. It is difficult to determine the destina-

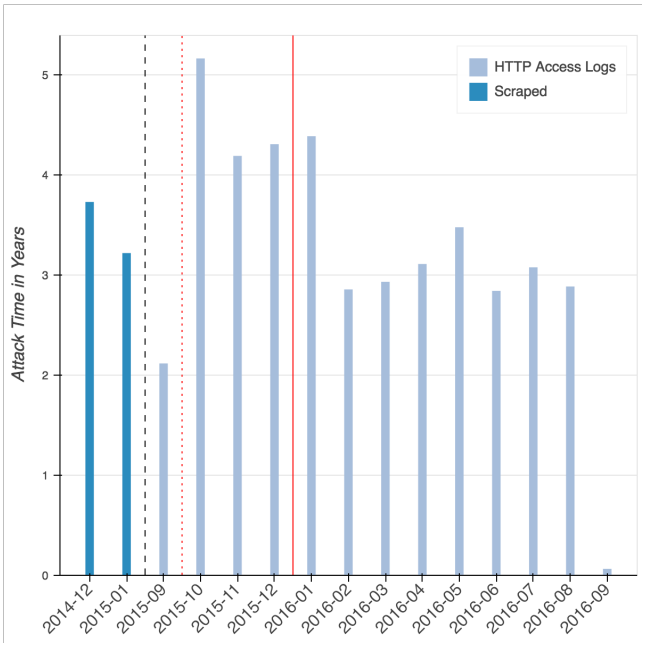

Figure 9: *This graph displays the total amount of attack time measured in years. Note there is a gap in the dates at the black dashed line. The solid red line represents the point when VDoS stopped accepting PayPal. The red dotted line marks the culmination of the payment intervention when PayPal was removed as a payment option on the VDoS website.*

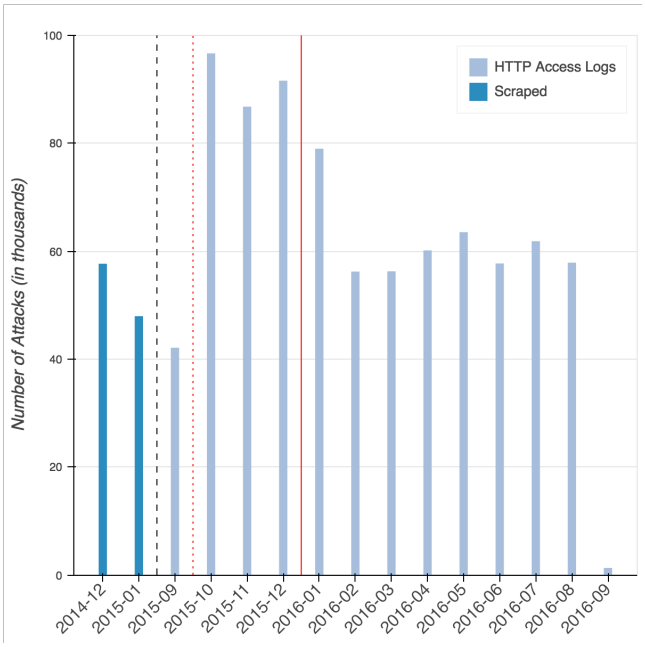

Figure 10: *This graph displays the total number of attacks. Again, note there is a gap in the dates at the black dashed line, and the solid red line represents the point when VDoS stopped accepting PayPal. The red dotted line marks the culmination of the payment intervention when PayPal was removed as a payment option on the VDoS website.*

tion of the other payments as the labels are not descriptive. For example, $3,858 (in 72 separate payments) was sent to an entity which might correspond to payment for support services. Based on what we found, it appears VDoS's major costs were hosting and customer support.

However, we are likely missing costs such as maintaining their scripts and adding new attacks. We are also missing some sources of revenue, such as fees collected to allow other booter services to use their attack API for launching attacks. VDoS also likely paid around 3% in payment processing charges or Bitcoin to fiat conversion fees. Even with this conservative cost estimate, VDoS was likely operating at a profit before and after the PayPal intervention, assuming their costs were relatively stable.

6. DISCUSSION

In this study, we have an inside view into the effects of PayPal's payment intervention on a single booter service, VDoS. While it is unclear if the effects we find generalize to other criminal services or booter services, our study constitutes an initial attempt to understand the effects of disrupting regulated payment channels on booter services.

Based on our analysis we found that there appear to be two disjointed sets of customers: one that has access to regulated payment methods such as PayPal, and another segment that has access and is willing to use unregulated payment methods such as Bitcoin. Only 386 (13%) of all repeat subscribers changed between the two payment methods.

The time period when VDoS was able to reliably accept PayPal and Bitcoin payments represents the high point of VDoS's revenue.

Our analysis of VDoS indicates that, currently, a booter service which is unable to accept PayPal or other regulated payment channels will lose revenue from the segment of customers that are unable or unwilling to use Bitcoin. However, it is unclear if these customers simply switch to another booter service that does accept a regulated payment channel and what would happen if all booter services were unable to accept regulated payments, such as debit or credit cards.

As unregulated payment channels such as Bitcoin become more mainstream, it will be important to understand user response to transitions such as the one described in this paper. Currently, we hypothesize that the difficulty in purchasing, and the perceived threat of theft keep some segment of users from adopting Bitcoin as a payment channel, even for criminal activity such as booting. It should be noted that although we found users were unlikely to switch from PayPal, this may not be the case for other services or even other booters as this was a case study of VDoS in particular. As future work we hope to opportunistically obtain and analyze additional information from booters and other criminal services that have their regulated payment channels disrupted and transition to unregulated payment channels, such as Bitcoin. This will hopefully improve our understanding of the efficacy of payment interventions as a means for disrupting booter services and other criminal services.

However, VDoS likely remained profitable, albeit less so, after they abandoned regulated payment channels. This in-

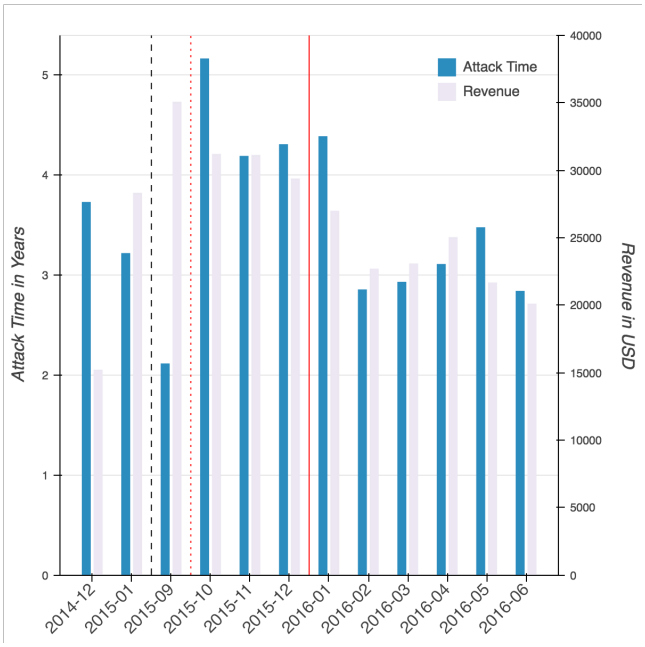

Figure 11: *Revenue tracks closely with attack time suggesting both are tightly correlated with the number of users. The final month of revenue data (July 2016) is not displayed since only revenue for half of the month is included in the leaked database. Again, note there is a gap in the dates at the black dashed line, and the solid red line represents the point when VDoS stopped accepting PayPal. The red dotted line marks the culmination of the payment intervention when PayPal was removed as a payment option on the VDoS website.*

dicates that we must explore other methods in addition to payment interventions as a means to further disrupt harmful booter services. Some of these might be improved attribution and investigation methods that assist law enforcement agencies in prosecuting operators and subscribers of booter services. This, in turn, might have a deterrence effect that reduces the number of subscribers and operators. Another avenue of exploration is the identification and disruption of a booter service's technical infrastructure, such as domain names, frontend website, and backend attack infrastructure takedowns that might drive up the costs of their operations.

7. CONCLUSION

This paper provides a fortuitous view inside the economics and operation of a DDoS-for-hire service that was impacted by a payment intervention. Among the results of this work, we have shown that VDoS was a growing business before the payment intervention, showing the increased demand for DDoS services. Additionally, we have confirmed the large-scale harm caused by booter services such as VDoS. We have also shown that this payment intervention against VDoS likely impacted their revenue and decreased the amount of harm caused by VDoS. Our analysis shows that some segment of booter customers are unwilling to adopt Bitcoin payment methods. However, VDoS continued to cause a large amount of harm and be profitable after the payment intervention.

Our assessment is that, currently, payment interventions will likely decrease the revenue and harm caused by booter services, though they are not a complete solution to mitigating the threat from these services. As such, the research community needs to continue to develop other intervention strategies and methods of measuring the effectiveness of these strategies.

Acknowledgments

The authors thank the reviewers for helpful feedback. In addition, we thank PayPal for their assistance with this project.

This work was supported in part by the National Science Foundation under contract 1619620 and a gift from Google. The opinions in this paper are those of the authors and do not necessarily reflect the opinions of any funding sponsor.

8. REFERENCES

[1] vDOS/PoodleCorp attack servers. https://ddosinvestigations.wordpress.com/2016/09/19/vdospoodlecorp-attack-servers/, 2016.

[2] R. Anderson. Why Information Security is Hard-An Economic Perspective. In *Proceedings of the 17th Annual Computer Security Applications Conference*, ACSAC '01, 2001.

[3] R. Anderson, C. Barton, R. Böhme, R. Clayton, M. J. G. van Eeten, M. Levi, T. Moore, and S. Savage. Measuring the Cost of Cybercrime. In *The Economics of Information Security and Privacy*, pages 265–300. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[4] R. Anderson and T. Moore. The Economics of Information Security. *Science*, 314(5799):610–613, 2006.

[5] J. J. Chromik, J. J. Cardoso de Santanna, A. Sperotto, and A. Pras. Booter websites characterization: Towards a list of threats. In *Proceedings of 33rd Brazilian Symposium on Computer Networks and Distributed Systems, SBRC 2015, Vitoria, Brasil*, pages 445–458. Brazilian Computer Society (SBC), May 2015.

[6] R. Clayton, T. Moore, and N. Christin. Concentrating Correctly on Cybercrime Concentration. In *14th Workshop on the Economics of Information Security*, 2015.

[7] A. Hutchings and R. Clayton. Exploring the Provision of Online Booter Services. *Deviant Behavior*, 37(10):1163–1178, 2016.

[8] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage. Spamalytics: An Empirical Analysis of Spam Marketing Conversion. In *Proceedings of the 15th ACM Conference on Computer and Communications Security*, CCS '08, pages 3–14, New York, NY, USA, 2008. ACM.

[9] C. Kanich, N. Weavery, D. McCoy, T. Halvorson, C. Kreibichy, K. Levchenko, V. Paxson, G. M. Voelker, and S. Savage. Show Me the Money: Characterizing Spam-advertised Revenue. In *Proceedings of the 20th USENIX Conference on Security*, SEC'11, pages 15–15, Berkeley, CA, USA, 2011. USENIX Association.

[10] M. Karami and D. McCoy. Understanding the Emerging Threat of DDoS-as-a-Service. In *Presented as part of the 6th USENIX Workshop on Large-Scale*

Exploits and Emergent Threats, Washington, D.C., 2013. USENIX.

[11] M. Karami, Y. Park, and D. McCoy. Stress Testing the Booters: Understanding and Undermining the Business of DDoS Services. In *World Wide Web Conference (WWW)*. ACM, 2016.

[12] L. Krämer, J. Krupp, D. Makita, T. Nishizoe, T. Koide, K. Yoshioka, and C. Rossow. AmpPot: Monitoring and Defending Amplification DDoS Attacks. In *Proceedings of the 18th International Symposium on Research in Attacks, Intrusions and Defenses*, November 2015.

[13] B. Krebs. Alleged vDOS Proprietors Arrested in Israel. https://krebsonsecurity.com/2016/09/alleged-vdos-proprietors-arrested-in-israel/, 2016.

[14] B. Krebs. Israeli Online Attack Service 'vDOS' Earned $600,000 in Two Years. https://krebsonsecurity.com/2016/09/israeli-online-attack-service-vdos-earned-600000-in-two-years/, 2016.

[15] J. Krupp, M. Backes, and C. Rossow. Identifying the Scanners and Attack Infrastructure behind Amplification DDoS attacks. In *Proceedings of the 2016 ACM Conference on Computer and Communications Security*. ACM, 2016.

[16] N. Leontiadis, T. Moore, and N. Christin. Measuring and Analyzing Search-Redirection Attacks in the Illicit Online Prescription Drug Trade. In *USENIX Security Symposium*. USENIX Association, 2011.

[17] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Félegyházi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu, D. McCoy, N. Weaver, V. Paxson, G. M. Voelker, and S. Savage. Click Trajectories: End-to-End Analysis of the Spam Value Chain. In *Proceedings of the 2011 IEEE Symposium on Security and Privacy*, SP '11, pages 431–446, Washington, DC, USA, 2011. IEEE Computer Society.

[18] D. Makita, K. Yoshioka, and M. van Eeten. Who Gets the Boot? Analyzing Victimization by DDoS-as-a-Service. In *Research in Attacks, Intrusions, and Defenses: 19th International Symposium, RAID 2016, Paris, France, September 19-21, 2016, Proceedings*. Springer, 2016.

[19] D. McCoy, H. Dharmdasani, C. Kreibich, G. M. Voelker, and S. Savage. Priceless: The Role of Payments in Abuse-advertised Goods. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12, pages 845–856, New York, NY, USA, 2012. ACM.

[20] D. McCoy, A. Pitsillidis, G. Jordan, N. Weaver, C. Kreibich, B. Krebs, G. M. Voelker, S. Savage, and K. Levchenko. PharmaLeaks: Understanding the Business of Online Pharmaceutical Affiliate Programs. In *Proceedings of the 21st USENIX Conference on Security Symposium*, Security'12, pages 1–1, Berkeley, CA, USA, 2012. USENIX Association.

[21] T. Moore and R. Clayton. Examining the impact of website take-down on phishing. In L. F. Cranor, editor, *APWG eCrime Researchers Summit*, volume 269 of *ACM International Conference Proceeding Series*, pages 1–13. ACM, 2007.

[22] J. J. Santanna, R. Durban, A. Sperotto, and A. Pras. Inside booters: An analysis on operational databases. In *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, pages 432–440, May 2015.

[23] J. J. Santanna and A. Sperotto. *Characterizing and Mitigating the DDoS-as-a-Service Phenomenon*, pages 74–78. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.

[24] J. J. Santanna, R. van Rijswijk-Deij, R. Hofstede, A. Sperotto, M. Wierbosch, L. Z. Granville, and A. Pras. Booters 2014; An Analysis of DDoS-as-a-service Attacks. In *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, pages 243–251, May 2015.

[25] K. Soska and N. Christin. Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem. In *Proceedings of the 24th USENIX Conference on Security Symposium*, SEC'15, pages 33–48, Berkeley, CA, USA, 2015. USENIX Association.

[26] B. Stone-Gross, R. Abman, R. A. Kemmerer, C. Kruegel, D. G. Steigerwald, and G. Vigna. The Underground Economy of Fake Antivirus Software. In B. Schneier, editor, *Economics of Information Security and Privacy III*, pages 55–78. Springer New York, New York, NY, 2013.

[27] K. Thomas, D. Y. Huang, D. Y. Wang, E. Bursztein, C. Grier, T. Holt, C. Kruegel, D. McCoy, S. Savage, and G. Vigna. Framing Dependencies Introduced by Underground Commoditization. In *14th Annual Workshop on the Economics of Information Security, WEIS 2015, Delft, The Netherlands, 22-23 June, 2015*, 2015.

[28] M. Vasek and T. Moore. Do malware reports expedite cleanup? An experimental study. In *Proceedings of the 5th USENIX Workshop on Cyber Security Experimentation and Test*, CSET'12, Berkeley, CA, USA, 2012. USENIX Association.

[29] D. Y. Wang, M. Der, M. Karami, L. Saul, D. McCoy, S. Savage, and G. M. Voelker. Search + Seizure: The Effectiveness of Interventions on SEO Campaigns. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, IMC '14, pages 359–372, New York, NY, USA, 2014. ACM.