

Booters - An Analysis of DDoS-as-a-Service Attacks

José Jair Santanna*, Roland van Rijswijk-Deij*[†], Rick Hofstede*, Anna Sperotto*,
Mark Wierbosch*, Lisandro Zambenedetti Granville[‡], Aiko Pras*

*University of Twente, The Netherlands

{j.j.santanna, r.m.vanrijswijk, r.j.hofstede, a.sperotto, a.pras}@utwente.nl

m.b.wierbosch@student.utwente.nl

[†]SURFnet bv, The Netherlands

roland.vanrijswijk@surfnet.nl

[‡]Federal University of Rio Grande do Sul, Brazil

granville@inf.ufrgs.br

Abstract—In 2012, the Dutch National Research and Education Network, SURFnet, observed a multitude of Distributed Denial of Service (DDoS) attacks against educational institutions. These attacks were effective enough to cause the online exams of hundreds of students to be cancelled. Surprisingly, these attacks were purchased by students from websites, known as *Booters*. These sites provide DDoS attacks as a paid service (*DDoS-as-a-Service*) at costs starting from 1 USD. Since this problem was first identified by SURFnet, Booters have been used repeatedly to perform attacks on schools in SURFnet’s constituency. Very little is known, however, about the characteristics of Booters, and particularly how their attacks are structured. This is vital information needed to mitigate these attacks. In this paper we analyse the characteristics of 14 distinct Booters based on more than 250 GB of network data from real attacks. Our findings show that Booters pose a real threat that should not be underestimated, especially since our analysis suggests that they can easily increase their firepower based on their current infrastructure.

I. INTRODUCTION

Since the Spring of 2012, the Dutch National Research and Education Network, SURFnet¹, has been confronted with a series of Distributed Denial of Service (DDoS) attacks against Dutch schools. These attacks typically last between 30 minutes and 1 hour and are mostly targeted at schools where students aged between 16 and 20 follow professional education. Interestingly, a correlation was found between the occurrence of one particular set of DDoS attacks and class schedules for certain courses. This led to an investigation that uncovered a student who was responsible for attacks against his home institution. When questioned, the student admitted to having used a *Booter* – a website that offers *DDoS-as-a-Service*, usually for a paid fee – to attack his own school.

According to SURFnet the attacks generated by Booters are not as massive as those that recently targeted organizations such as CloudFlare [1] and Spamhaus [2]. However, the attacks are sufficiently powerful to isolate schools from the Internet, despite the fact that these schools are often connected to the Internet via links of 1 Gbps or more. Additionally, reports indicate that Booters are also being used to launch attacks against personal websites, government agencies, and even other Booters [3], [4]. This makes Booters a worrying phenomenon, especially considering that very little is known

about the characteristics of the attacks that they perform, which is essential knowledge for mitigating their attacks.

The goal of this paper is to create awareness around Booter attacks. In our study, we investigate the characteristics of Booter attacks in terms of the volume of generated traffic as well as the service and networking infrastructure used by Booters. Finally, based on our measurements, we discuss possible defense mechanisms and the relationship between Booters and DDoS protection services. We performed measurements to analyze the attacks generated by Booters on our own infrastructure. We investigated more than 250 GB of traffic. We intend to make all data acquired during our experiments available to interested researchers.

Although there is a vast amount of literature on DDoS attacks and mitigation techniques [5], [6], [7], [8], [9], this paper is, to the best of our knowledge, the first to present a structured analysis of several *DDoS-as-a-Service* providers. Karami et al. [10], [3] investigate Booters as well, using a similar approach but focus on a single Booter only. Our study, instead, presents the recent Booter landscape, showing trends and common characteristics of the Booter market. Indirectly related to Booters, several authors [11], [12] refer to cyber-crime as being an organized market (*Crimeware-as-a-Service*), showing how common customers can easily access and acquire crimeware. However, their contribution is mainly focused on botnets and phishing, and does not include Booters.

Our main contribution is that we analyse the attack characteristics of 14 distinct Booters. The results of our analysis are a valuable aid to attack targets in mitigating and preventing these attacks and help other researchers in understanding how Booters work. In addition to this, our algorithm to compensate for missing network traffic is valuable for researchers that do not have enough network capacity to measure large scale attacks. Finally, the data we collected, which we pledge to share with interested researchers², is a valuable resource for the network security research community, for example by using it to validate new DDoS detection and protection approaches.

The remainder of this paper is organized as follows. In Section II, we describe the methodology we used to analyze Booter attacks and the algorithm used to compensate traffic of DDoS attacks. In Section III, we report on our experiences

¹<http://www.surf.nl/en/about-surf/subsidiaries/surfnet>

²http://www.simpleweb.org/w/index.php/Traces#Booters_-_An_Analysis_of_DDoS-as-a-Service_Attacks

with Booter attacks, highlighting the major characteristics of the attack traffic. We discuss the relationship between Booters and DDoS Protection Services (DPSes) in Section IV, and provide conclusions in Section V.

II. METHODOLOGY

In this section we discuss our approach to investigating Booters. We start by describing how we localize Booters. Next, we give a step-by-step description of how we prepare our experiment and how we launched the Booters. We end the section by discussing the compensation algorithm used to overcome limitations in our measurement setup.

A. Finding Booters

To find Booters, we use a two step process. First, a crawler, previously introduced in [13], is used to generate a weekly list of candidate URLs. Every URL can be a Booter, a web page (e.g., blogs, video services, reports) containing the name of a Booter, or a regular web page unrelated to Booters, which in this case is discarded. To generate the list of candidate URLs, our crawler searches with Google’s Custom Search for the following keywords: “Booter”, “Stresser”, “DDoS”, “DDoS-as-a-Service”, and “DDoS-for-hire”. In our experience, this is a comprehensive set of keywords for finding Booters.

Second, we perform a manual investigation of every candidate URL on our weekly list to check whether an entry is, or refers to, a Booter. The result is a second list containing the Booters we found. Our crawler and manual investigation of URLs has been performed since July 2013. Note that although it is easy to find Booters on the Internet, we decided to anonymize the names of Booters used in this paper for ethical reasons. The complete list of Booters is provided upon request. Another note is that although we are aware that our list does not retrieve some Booters advertised in sites not indexed by Google, such as hacker forums, our goal with this approach is to retrieve a comprehensive list of Booter that can be easily found by any user on the Internet.

B. Measurements

To investigate the characteristics of Booter attacks, we purchased DDoS attacks from 14 Booters that were online and operational on 14 and 15 August 2013. The goal of our experiment was to determine how much traffic Booters are able to generate and the geographical distribution of the systems misused by Booters to perform attacks. In collaboration with SURFnet and the University of Twente (UT), we launched a series of attacks on network infrastructure specifically dedicated to this experiment at the UT. Although our list of Booters at that moment was composed of 21 online Booters, 7 of them had a faulty payment system that did not allow us to purchase packages of attacks.

For each of the 14 Booters investigated we: 1) create an account, 2) purchase an attack package; and 3) launch UDP-based DDoS attacks against a null-routed IP address at the UT. Although Booters offer several types of DDoS attacks, as we present in Section III-A, for these experiments we concentrated on volumetric attacks based on UDP because no service running on the target system is required, and the only potential bottleneck is the network link capacity.

Algorithm 1: Traffic compensation for sets of packets, returning a time-series.

```

input : pkts, threshold, bin_size
1 bin_start_time  $\leftarrow$  pkts[0].time
2 bin  $\leftarrow$  0; bin_data  $\leftarrow$  0; gap_time  $\leftarrow$  0
3 for  $i \in [1, \text{pkts.length}]$  do
4   if pkts[i].time - bin_start_time  $\leq$  bin_size then
5     bin_data  $\leftarrow$  bin_data + pkts[i].size
6      $\Delta t \leftarrow$  pkts[i].time - pkts[i - 1].time
7     if  $\Delta t >$  threshold then
8       gap_time  $\leftarrow$  gap_time +  $\Delta t$ 
9   else
10    compensated[bin] =  $\frac{\text{bin\_data}}{(\text{bin\_size} - \text{gap\_time})}$ 
11    bin_start_time  $\leftarrow$  bin_start_time + bin_size
12    bin  $\leftarrow$  bin + 1
13    bin_data  $\leftarrow$  pkts[i].size
14    gap_time  $\leftarrow$  0
15 return compensated

```

C. Compensating DDoS attack traffic

During the attacks, we captured raw packet data at the UT using dedicated hardware, capable of capturing traffic at 10 Gbps. To ensure that attacks did not hinder the functioning of our or other networks, and that the attack traffic rate remains below the maximum network link rate (10 Gbps), SURFnet and the Computer Security Incident Response Teams (CSIRTs) from SURFnet and the UT were informed and actively collaborated in monitoring the attack traffic. Unexpectedly, when we compared the traffic rate at SURFnet to the rate measured at the UT, the latter had a lower traffic rate indicating that our equipment was overloaded during the measurement. By investigating artifacts in the data, we realized that both the PCIe-bus and RAID system of the monitoring system were overloaded during the attacks. Because SURFnet only retains visual statistics instead of packet captures, we could not use their measurements to analyze the attacks in detail.

To still be able to use the UT dataset, we developed an algorithm to compensate for lost traffic. Analysis shows that the algorithm is very accurate in relation to what was measured at SURFnet. Since we concentrate only on volumetric attacks, our algorithm assumes the attack traffic to be sent in a streaming fashion, with short inter-arrival times between packets. A longer inter-arrival time then constitutes an indication that packets have been dropped. Figure 1 presents an example of the inter-arrival time distribution for one of the considered attacks, at millisecond resolution. The distribution clearly shows the presence of larger gaps in the inter-arrival time, in this example clustered around 10^2 ms.

Algorithm 1 processes raw traffic traces for an attack and outputs a compensated time-series. For every packet the algorithm determines whether it falls within the current or the next time bin (line 4). In case the packet falls in the current bin, its size in bytes is accounted and the time difference with the previous packet is calculated (line 5–6). We refer to the situation in which the time difference with the previous packet

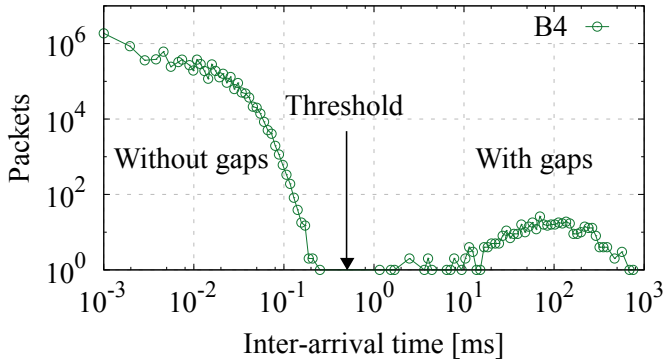


Fig. 1: Example of inter-arrival time distribution.

is larger than the pre-set threshold as a gap. In case a gap is detected, the gap duration is accounted (line 7–8), which is needed for the compensation later on. If the considered packet falls in the next time bin (line 9), the actual compensation is performed (line 10) by dividing the total number of bytes within the bin by the compensated bin size (*i.e.*, the duration of the bin in seconds, subtracted by the total duration of the gaps). Since the missing packets have not been accounted in the total number of bytes per bin (as they are not available in the packet trace) but implicitly in the bin size, as we assume a constant number of packets per bin, only the bin size requires compensation. Finally, after resetting the variables for the next run, the next packet is processed.

The value of `threshold` has been chosen based on the inter-arrival time distribution of each attack. In Figure 1 we have indicated the distribution of the traffic with and without gaps, and the threshold should be chosen such that it discriminates between the left and right part of the distribution. In practice, the value of `threshold` is between 1 and 10 ms. In Section III-A the measured and compensated attacks are shown and analyzed.

III. BOOTER ATTACKS

A. Attack types and volume

Of the 14 Booters from which we purchased attacks, 5 Booters did not perform the UDP-based attacks that we ordered: 3 of those did not send any traffic, and 2 surprisingly generated a handful of TCP packets. The 9 remaining Booters performed as requested, however, and generated more than 250 GB of traffic.

Although there are several types of UDP-based attacks (*e.g.* amplification attacks, based on NTP, SNMP, DNS, and Echo), our measurements only show 7 DNS-based attacks and 2 attacks involving the CharGen protocol. This observation is in line with current trends described in [14] that show DNS and CharGen as two of the most common types of UDP-based attacks.


Both types of attacks (DNS and CharGen) belong to the class of reflection and amplification attacks. These attacks are based on the principle that an attacker sends a relatively small request to a server, crafted with the spoofed IP address of the intended target (reflection), and for which the response is much larger than the request (amplification). For example, in

case of a DNS-based attack, an attacker may send a relatively small DNS query (in the order of 40 – 60 bytes), which may be answered with a large response that can be 4 KB or more in case EDNS0³ is used. In case of CharGen [16], RFC 864 defines that requests to servers should be answered with a randomly-sized reply up to 512 bytes in size. In the next subsection, DNS-based attacks are analyzed, followed by analysis of the CharGen-based attacks.


1) *DNS-based attacks*: Figures 2(a) and 2(b) show the volume of DNS-based attacks measured and compensated, respectively. By analyzing both figures, it is clear that packets have been dropped for attacks with a rate higher than 400 Mbps. The rates of Booters B2, B3, B5, and B7, for which the traffic rate is below 400 Mbps, are barely affected by the compensation algorithm. Booters B1, B4 and B6, on the other hand, show significant gaps for which the algorithm compensates. Our compensated results for all attacks, shown in Figure 2(b), are completely in line with what was measured by SURFnet.

Based on SURFnet’s experience, it is no surprise that some Booters (*e.g.*, B4 and B6) generate attacks with rates of more than 1 Gbps, otherwise schools in the Netherlands would not have been taken offline. More worrying is that all Booters, except B5, generate rates high enough to saturate typical ADSL, ADSL2+ and DOCSIS connections, which are

³The Extension mechanisms for DNS (EDNS0) [15] allow for - among other things - larger DNS responses (than the originally specified 512 bytes), with the most common maximum size configured set to 4 KB.



(a) Measured.



(b) Compensated.

Fig. 2: Traffic rate of DNS-based attacks.




Fig. 3: Packet size distribution (DNS).

used by a large proportion of home Internet users and are also commonly used by small and medium enterprises [17].

Figure 3 shows the distribution of packet sizes in the attacks. The range in packet sizes for all Booters significantly exceeds 512 bytes, the default maximum response size for regular DNS. The CDF even shows that in certain cases (B4) 75% of the distribution is concentrated around values as high as 4 KB. The size of DNS responses is an important factor in an amplification attack, since a large amplification factor – *i.e.* the ratio between the size of the response and that of the request – will lead to an attack that requires less resources on the side of the attacker. By inspecting the packets captured at our measurement point, we found that all attacks make use of EDNS0 [15], which – as mentioned before – allows for responses typically as large as 4 KB. In addition, it should be noted that all the attacks we saw use ANY queries⁴ to achieve maximum amplification. Particularly noteworthy is that all 7 DNS amplification attacks used identical values for certain query parameters. We therefore suspect the script or program used to generate the attacks to be the same or based on a common root source in all 7 cases. Since knowledge of these particular parameters can help greatly in mitigating this particular attack, we will not disclose the specifics in this paper as that could help attackers improve the attack. Figure 3 also shows that Booters that use the same DNS query (see Table I) have a very similar distribution in packet length, such as Booters B1, B2, B3, and Booters B6 and B7. Keeping in mind that the query has almost the same size for all Booters, we conclude that the amplification factor of the former group of Booters is lower than the latter. What finally stands out is that B5 generates the shortest responses.

Table I shows the average rate of each attack, the number of systems involved (misused DNS resolvers) in performing the attacks, the average number of packets per system, and the DNS query used for attacks. The most surprising finding in Table I is that although B5 has the largest set of misused systems (8281 DNS resolvers), the rate of attack was the lowest (6.11 Mbps) of all considered Booters, since the involved resolvers on average sent only 261.8 packets each and the packet size distribution ranges from [20 - 900] bytes. This indicates that the number of hosts involved in an attack is not necessarily an adequate indicator of the attack strength. In fact,


⁴The DNS ANY query is used for retrieving all resource records available for a given domain name.

TABLE I: Details of DNS-based attacks.


Booter	Average rate [Gbps]	Misused systems	Average packets per system	DNS query
B1	0.70	4486	2886.1	root-server.net
B2	0.25	78	116082.5	root-server.net
B3	0.33	54	245169.2	root-server.net
B4	1.19	2970	12327.9	packetdevil.com
B5	0.006	8281	261.8	ddostheinter.net
B6	0.15	7379	1329.2	anonsc.com
B7	0.32	6075	2756.7	anonsc.com

the volume of an attack is a function of the number of systems involved, the number of packets each system sends, and the amplification factor. For example, although B3 relied on a set of misused systems more than 100 times smaller than B7, this Booter generated almost the same volume of traffic as B7. This was possible because the number of packets sent by B3 was 88 times larger than B7.

2) *CharGen-based attacks*: According to several reports [14], [18], DDoS attacks based on CharGen barely appear before 2013, but since then their use has increased significantly. For example, from September to December 2013 Prolexic [14] reports an increase of 92.31%. Figures 4(a) and 4(b) show the rate of CharGen-based attacks measured and compensated, respectively. The traffic rate generated by B9 exceeded our expectations with peaks around 7.0 Gbps, almost 4 times



(a) Measured.



(b) Compensated.

Fig. 4: Traffic rate of CharGen-based attacks.

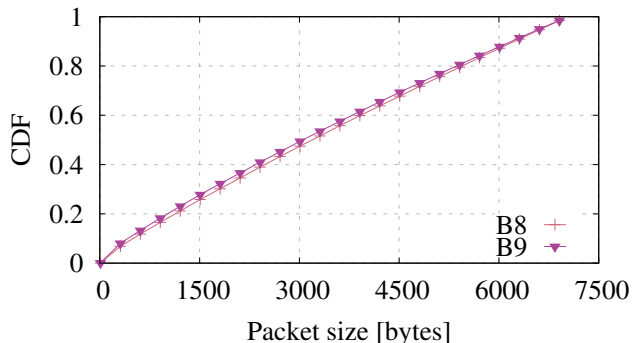


Fig. 6: Packet size distribution (CharGen).

higher than the largest DNS-based attack (B6).

Surprisingly, we notice a large discrepancy between the maximum allowed packet size as per the CharGen protocol specification in RFC 864 [16] (512 bytes) and what we measure. As shown in Figure 6, for both Booters B8 and B9, the size of packets is randomly distributed in the range of [0, 6956] bytes. Therefore, we suspect that the systems involved in the attacks were running a non-RFC-compliant implementation of the CharGen protocol. To verify this, we first examined the misused systems using `nmap`⁵ and we observed that the majority of these systems were running Microsoft Windows.

To verify whether the observed CharGen implementation is specific to MS Windows systems, we installed several recent versions of Microsoft Windows, as well as the reference implementation of the `xinetd`⁶ daemon on Linux (which includes CharGen). When we tested the protocol in our lab environment, our results confirmed those of the live attacks for the implementations on systems running Microsoft Windows. The maximum CharGen packet sizes measured in our lab environment are remarkable: all Microsoft Windows versions from XP up return messages with a random size of [0, 6956] bytes. This confirms that the Windows implementations are non-RFC-compliant. In addition, since CharGen is installed as part of the Simple TCP/IP Services, Windows systems may

⁵<http://nmap.org>

⁶<http://www.xinetd.org>

TABLE II: Details of CharGen-based attacks.

Booter	Average rate [Gbps]	Misused systems	Average packets per system
B8	0.99	281	20491.1
B9	5.48	3779	3514.4

therefore become a powerful base for this type of amplification attack if these services are enabled. Our tests also show that the `xinetd` implementation of CharGen is non-RFC-compliant, although in this case the maximum obtained response size is limited to 1024 bytes, and therefore – on average – 3.4 times smaller than for Microsoft Windows.

Similar to the DNS-based attacks, the traffic rates of the CharGen-based attacks depend on the number of systems involved, the number of packets sent per system, and the implementation of the service on abused systems. In case of Booters B8 and B9, the attacks show a remarkable similarity in the packet size distribution (Figure 6), indicating that both Booters abuse the same type of systems. From Table II, we see that despite the systems controlled by B8 are more aggressive (20491.1 packets/system), B9 has activated a larger set of hosts, which results in this case in a larger attack volume.

B. Geographical distribution of misused systems

We also examined the geographical distribution of the servers abused for the attacks. Since the attacks are reflection-based, the measured source IP addresses are the legitimate addresses of the misused systems and therefore geolocation provides meaningful results. Figure 5(a) and 5(b) show the geographical distributions of the servers involved in the DNS and CharGen attacks, respectively, cumulated for all the measured attacks. The figures also show the top 10 of the most active countries in the considered datasets, in terms of number of misused servers. In case of DNS, the top 10 is dominated by the US, with more than 5.8k hosts, followed by Japan and Germany. This result is not surprising since these countries are among the countries with the highest Internet penetration [17].

More surprising is the distribution of hosts abused for carrying out CharGen-based attacks. In this case, China clearly dominates the top 10, while the US follows with only about a




Fig. 5: Geographical distribution of misused servers.




Fig. 7: Continent breakdown per Booter.

third of the number of hosts in China. It is currently unclear why China dominates the top 10 for CharGen attacks. The predominance of China, however, was already observed in earlier analyses, *e.g.* in [18].

Finally, we investigated the geographical distribution of individual hosts involved in each attack. Figure 7 shows the continent breakdown of hosts misused by each of the 9 tested Booters. For the majority of the Booters, with the exception of B8 and B9 that generated the aforementioned CharGen-based attacks, the majority of misused hosts is located in North America (22-33%) and Europe (31-61%).

C. Potential for future Booter attacks

An interesting aspect to investigate is whether Booters share the same infrastructure for performing attacks. We looked into this issue by comparing the set of misused systems in each of the observed attacks. Since hosts vulnerable to reflection and amplification attacks can potentially be found online by simply scanning, we expect to find an intersection between the set of misused systems.

Table III shows the pairwise intersection between the set of misused systems, calculated as the overlap fraction between the sets of misused host for Booters of B_X and B_Y , expressed as a percentage: $\frac{|B_X \cap B_Y|}{|B_X|} \times 100$. Contrary to our expectations, in almost all cases Booters do not, or only minimally share their attack infrastructure. This indicates that the set of misused hosts in an attack is part of the Booter business model, and we speculate that Booters may employ more advanced techniques in choosing their infrastructure than just harvesting

by scanning. From Table III, we also see that there are some exceptions to this observation. This is the case, for example, for B6, which shares 98.65% of its infrastructure with B7. By looking at the economic aspects we found that the most likely reason for the high intersection is that B6 and B7 are both linked to the same PayPal account, indicating that these Booters share the same owner. This finding also seems to indicate that Booter owners are taking care of offering different products with different prices to attract different customers. Lastly, Table III also shows a small intersection between the host sets used to perform DNS-based and CharGen-based attacks (*e.g.*, B9 correlates with B5, B6 and B7). This same intersection pattern is also described in [19], however without a clear conclusion.

Table III also clearly indicates that Booters have a high potential for future attacks. Booters can easily increase their firepower by using each others infrastructure. Considering that, in the case of DNS-based attacks, our measurement includes 29321 unique misused IP addresses, this could indicate an increase in firepower between 3.5 (B5) and 542 times (B3). For example if Booter B8 that uses 281 systems uses the 3779 systems of Booter B9, then B8 could generate an attack up to 13 times stronger than what we measured, possibly reaching 13 Gbps.

It does not stop here, however, the potential firepower of attacks can be even worse. Recent work by Rossow [20] shows that there are at least 89000 CharGen amplifiers on the Internet at present. If, for instance, Booter B9 would abuse all of these it could increase its firepower by over 23 times, potentially achieving peak attack volumes over 160 Gbps. Kühner et al. [19] describe measurements over a 3 month period in late 2013 and early 2014 that shows a pool of open DNS resolvers between 23 and 25.5 million hosts in size. Assuming the lower bound and that, *e.g.* Booter B6 abuses all available open resolvers, it could reach well over 3000 times the attack volume it achieved in our measurements.

IV. DEFENSE AGAINST BOOTER ATTACKS

Our experiments presented in Section III show that Booters are a threat that should not be taken lightly. Moreover, future attacks can easily be worse than what we have observed so far. In this section, we discuss several options for mitigating the classes of attacks that we have observed in this paper. Furthermore, we highlight our findings related to how Booters protect themselves against the competition by taking preventive measures against DDoS attacks themselves.

TABLE III: Intersection between sets of misused systems by the tested Booters.

\cap	B1	B2	B3	B4	B5	B6	B7	B8	B9
B1	–	0.20	0.20	3.88	0.02	1.07	0.73	0	0
B2	11.54	–	0	0	0	0	0	0	0
B3	16.67	0	–	0	0	1.85	1.85	0	0
B4	5.86	0	0	–	0.20	4.11	1.04	0	0
B5	0.01	0	0	0.07	–	8.38	7.99	0	0.08
B6	0.65	0	0.01	1.65	9.42	–	81.33	0	0.07
B7	0.54	0	0.02	0.51	10.90	98.65	–	0	0.08
B8	0	0	0	0	0	0	0	–	43.06
B9	0	0	0	0	0.18	0.13	0.13	3.20	–




Fig. 8: Percentage of time that 102 Booters are protected, sorted by the year they started to be accessed.

A. Mitigating Botter attacks

CharGen-based attacks are the easiest to be mitigate right away. CharGen uses the fixed UDP port 19, which can simply be filtered at the network border router, or, if this would not be sufficient, by an upstream network provider. It remains questionable in our opinion, however, whether CharGen should be installed on end-systems at all. CharGen seems to be mostly abused for DDoS attacks and have limited benign applications nowadays.

It is much harder to take countermeasures against DNS-based attacks. Blocking DNS traffic, for instance, is impossible as it would prevent end users from properly using the Internet. DNS-based attacks clearly benefit from the large number of open DNS resolvers, often installed by default in Customer Premises Equipment (CPE), such as home routers. An obvious step towards mitigation is therefore to disable such services. Another approach is to rate limit DNS traffic on network edge routers since DNS traffic should rarely be more than a fraction of overall traffic. Note that this may be detrimental to legitimate DNS servers running inside the network that performs the rate limiting so it is no catch all solution.

Reflection and amplification attacks, as described in Section III, rely on the possibility of forging IP addresses that do not belong to the originating network (IP spoofing). Already in 2000, the Internet Engineering Task Force (IETF) proposed a Best Current Practice (BCP38) [21], which strongly encourages network administrators to implement ingress-filtering rules to block all traffic from IP addresses that do not belong to the address space of the originating network. BCP38 constitutes an effective defense mechanism against reflection attacks. However, the data collected in this study has confirmed that many operators do not comply with the best practice. According to the MIT Spoofer project [22], more than 30% of the Autonomous Systems surveyed still allow source address spoofing in February 2014. Therefore, actions to enforce the implementation of BCP38, for example, by including it into peering agreements, ITU regulations or national telecom law, must be more widely taken.

B. What Booters do to be protected

DDoS Protection Services (DPSes) are a popular mitigation method against DDoS attacks. DPSes are online services offered by companies like Versign, Akamai, Prolexic, Incapsula, and CloudFlare, for example. DPSes act as proxies for their

customers, combining the concept of Content Delivery Networks (CDNs) for improving the availability of their customers with advanced traffic filtering to neutralize DDoS attack traffic. In a market where DDoS attacks are the main product, the best way to beat the competition is to isolate them from the Internet. There is evidence that Booters are attacking each another [3]. A possible mitigation approach for Booters, therefore, would be to use DPSes themselves, as a countermeasure against attacks from competing Booters. We investigate if this is indeed the case in this section.

Next to purchasing attacks from 14 Booters, we also tracked where the web sites for a larger list of 102 Booters is hosted. The goal of this is to determine if these Booters make use of DPSes to protect themselves and to track the changes over time. To track where Botter sites are hosted, we use a passive DNS (pDNS) data source, called DNSDB provided by Farsight Security⁷. These data sources have information from the DNS over time about the mapping of host names (in our case Botter host names) to IP addresses. Based on these IP addresses we can determine whether or not the Botter is inside a DPS. For more information about pDNS see [23].

Figure 8 shows the percentage of the 102 Booters we tracked that is protected by a DPS. Although there are several DPSes available on the Internet, only Incapsula and CloudFlare are used by Booters in our dataset. A possible explanation for this penchant for Incapsula [24] and CloudFlare [1] may be that only these two offer the option of a free subscription. However, we are also aware of the fact that the DDoS attack mitigation mechanisms of these DPSes are not included in the free subscription, leading us to conclude that Booters may subscribe to more advanced protection plans.

Figure 8 also shows that 89% (53 out of 59) of Booters that started their activities between 2012 and 2014 were protected by DPSes for part of their lifetime. On average, Booters spend an increasingly large fraction of their lifetime protected by DPSes, as indicated in Table IV. New Booters that started operations in 2014 are all protected by a DPS. The more long-lived Booters, like the ones that became active in 2010 and 2011, also make use of DPSes. However, we found that these Booters mostly started subscribing to protection services from 2011. This leads us to believe the trend of using DPSes in the Botter market started in 2011.

An in-depth analysis of the functionality of DPSes is outside the scope of this paper. We note here, however, that

⁷<https://www.dnsdb.info/>

TABLE IV: Average fraction of time in DPSes, per year.

Year	CloudFlare	Incapsula	Unprotected
2010	22.41	0.09	77.51
2011	9.07	0.08	90.85
2012	43.55	2.75	53.70
2013	58.18	2.44	39.37
2014	100	0	0

DPSes could play a major role in mitigating Booters. By acting as a proxy, DPSes are able to access information specific to Booters, such as the real IP address of the Booter, in addition to information about customers accessing the service or the Booter owners. Attack parameters such as the target IP address could also be used to preempt attacks. More research is needed to understand which type of information these services are able to access de facto and which type of mechanisms they offer against different types of DDoS attacks.

V. CONCLUSIONS

To the best of our knowledge, this paper is the first to present a characterization of a large set of Booter attacks. We assessed the characteristics of UDP-based attacks carried out by 14 distinct Booters, which were active around the end of 2013. To better understand the Booter market, and in particular how they use DDoS Protection Services, we analyzed a total of 102 Booters, using (historical) DNS data from North America, covering a period since 2010. Finally, we discussed ways to protect against Booter attacks.

In the second half of 2013 we bought attack packages from 14 Booters to attack our own infrastructure and analyze the traffic generated during these attacks (see Section III). We measured primarily UDP-based reflection and amplification attacks, using DNS and CharGen. While DNS amplification attacks are well-known, CharGen attacks are relatively new and since 2013 have rapidly been gaining popularity. The DNS-based attacks, for which we had to pay only a few dollars, showed traffic peaks of up to 1.6 Gbps, whereas the CharGen attacks even showed peaks around 7.0 Gbps.

A surprising conclusion regarding the attack sources is that Booters do not (yet?) use the same (DNS and CharGen) hosts to amplify their attacks. This means that attacks might become much stronger once a single Booter will be able to exploit *all* systems currently used for amplification. Even worse, if we take into account the results of Internet-wide-scan projects, the number of available systems that can be abused might be almost 4000 times higher than the number of systems that we saw in the most powerful DNS-based attack on our own infrastructure.

When we look at the Booter market, we conclude that there is apparently cut-throat competition. Since 2010, the number of Booters that protect themselves against the same types of attack that they sell themselves by making use of DDoS Protection Services (DPSes) has grown dramatically, such that every Booter we know to be active in early 2014 was behind a DPS. This does open up avenues for future work in combating Booters; if DPSes can be compelled to collaborate, characterizing who runs Booters and what their internal infrastructure looks like becomes a lot easier.

Finally we discussed several countermeasures against Booter attacks (see Section IV). We consider ingress filtering, as described in BCP38, to be an effective solution against reflection-based DDoS attacks. However, the attacks collected in this study have once again confirmed that many operators do not implement BCP38. In addition, operators should filter CharGen traffic, since there are no good reasons why CharGen should be accessible outside a local environment.

ACKNOWLEDGMENTS

This work was funded by the Network of Excellence project FLAMINGO (ICT-318488), which is supported by the European Commission under its Seventh Framework Programme. Special thanks go to the teams of SURFnet, SURFcert, CERT-UT, and Farsight Security for assisting us in the experiments and providing valuable information. We also thanks to reviewers from the Internet measurement community for valuable advices on how to improve our research.

REFERENCES

- [1] M. Prince, "Technical Details Behind a 400Gbps NTP Amplification DDoS Attack," <http://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack>.
- [2] —, "The DDoS That Almost Broke the Internet," <http://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet>.
- [3] M. Karami and D. McCoy, "Understanding the Emerging Threat of DDoS-as-a-Service," in *6th USENIX Workshop on Large-Scale Exploits and Emergent Threats*, ser. LEET'13, 2013.
- [4] Krebs, B., "DDoS Services Advertise Openly, Take PayPal; The Obscurest Epoch is Today; The World Has No Room For Cowards," <http://krebsonsecurity.com/tag/booter-tw/>.
- [5] J. Mirkovic and P. Reiher, "A taxonomy of ddos attack and ddos defense mechanisms," *ACM SIGCOMM Computer Communication Review*, 2004.
- [6] T. Peng, C. Leckie, and K. Ramamohanarao, "Survey of network-based defense mechanisms countering the dos and ddos problems," *ACM Computing Surveys*, 2007.
- [7] A. Srivastava, B. Gupta, A. Tyagi, A. Sharma, and A. Mishra, "A recent survey on DDoS attacks and defense mechanisms," *Advances in Parallel Distributed Computing*, 2011.
- [8] S. Zargar, J. Joshi, and D. Tipper, "A survey of defense mechanisms against distributed denial of service (ddos) flooding attacks," *Communications Surveys & Tutorials, IEEE*, 2013.
- [9] J. J. Santanna, R. Durban, A. Sperotto, and A. Pras, "Inside Booters: An Analysis on Operational Databases," in *14th IFIP/IEEE International Symposium on Integrated Network Management (accepted)*, ser. IM 2015, 2015.
- [10] M. Karami and D. McCoy, "Rent to Pwn: Analyzing Commodity Booter DDoS Services," *USENIX*, 2013.
- [11] D. Manky, "Cybercrime as a service: a very modern business," *Computer Fraud & Security*, 2013.
- [12] A. Sood and R. Enbody, "Crimeware-as-a-service - A survey of commoditized crimeware in the underground market," *International Journal of Critical Infrastructure Protection*, 2013.
- [13] J. J. Santanna and A. Sperotto, "Characterizing and Mitigating The DDoS-as-a-Service Phenomenon," in *8th International Conference on Autonomous Infrastructure, Management and Security*, ser. AIMS'14, 2014.
- [14] Prolexic, "Prolexic Quarterly Global DDoS Attack Report (Q1 2014)," <http://www.prolexic.com/knowledge-center-ddos-attack-report-2014-q1.html>.
- [15] J. Damas, M. Graff, and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))," RFC 6891, 2013.
- [16] J. Postel, "Character Generator Protocol," RFC 689, 1983.

- [17] Akamai, "The State of the Internet (Q3 2013)," http://www.akamai.com/dl/akamai/akamai-soti-q313.pdf?WT.mc_id=soti_Q313.
- [18] Prolexic, "Prolexic Quarterly Global DDoS Attack Report (Q3 2013)," <https://www.prolexic.com/knowledge-center-ddos-attack-report-2013-q4.html>.
- [19] M. Kühner, T. Hupperich, C. Rossow, and T. Holz, "Exit from Hell? Reducing the Impact of Amplification DDoS Attacks," in *23rd USENIX Security Symposium*, ser. USENIX'14, 2014.
- [20] C. Rossow, "Amplification Hell: Revisiting Network Protocols for DDoS Abuse," in *Network and Distributed Systems Security Symposium*, ser. NDSS'14, 2014.
- [21] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," BCP 38, 2000.
- [22] R. Beverly, A. Berger, Y. Hyun, and K. Claffy, "Understanding the Efficacy of Deployed Internet Source Address Validation Filtering," in *9th ACM SIGCOMM/USENIX Internet Measurement Conference*, ser. IMC'09, 2009.
- [23] CISCO, "Tracking Malicious Activity with Passive DNS Query Monitoring," <https://blogs.cisco.com/security/tracking-malicious-activity-with-passive-dns-query-monitoring/>.
- [24] Incapsula, "Protect and Accelerate Your Online Service," <http://www.incapsula.com/pricing-and-plans.html>.