Birds of a Feather? A Comparative Analysis of DDoS Victimisation by IoT Botnet and Amplification Attacks

Swaathi Vetrivel[1], Arman Noroozian, Daisuke Makita[2,3], Katsunari Yoshioka[3], Michel van Eeten[1], and Carlos H. Gañán[1]

[1] Delft University of Technology
{S.Vetrivel, M.J.G.vanEeten, C.HernandezGanan}@tudelft.nl
[2] National Institute of Information and Communications Technology
{D.Makita@nict.go.jp}
[3] Yokohama National University
{Yoshioka@ynu.ac.jp}

Abstract. The number of Distributed Denial of Service attacks is growing, and the attack vectors are also changing. The advent of IoT botnets like Mirai has challenged the popularity of older techniques like amplification attacks. In this paper, we characterise the consequences of this change for the victimisation pattern of DDoS attacks. We conduct the first empirical comparison of victims of amplification attacks and botnet attacks and draw on the properties of targets outlined in Routine Activity Theory (RAT) to characterise the differences. We analyse the differences in the victimisation patterns at the level of IP addresses and Autonomous Systems. We find differences in the types of networks the victims reside in; botnet attacks are more common against hosting ASes, and its victims are significantly more likely to use dedicated hosting. We also observe that victims of botnet-based attacks tend to be in ASes with larger customer cone sizes. We use a balanced random forest classifier to distinguish the features of the victims of each attack. The model's output confirms our findings and draws out additional geographical differences in the victim distribution. Using the target properties of Value, Inertia, Visibility and Accessibility outlined in RAT, we find that victims of botnet attacks tend to have higher value and visibility and lower inertia than those of amplification attacks. We explain the differences in these patterns of victimisation to the underlying differences in the attack technique. We further use RAT to outline the policy implications of our analysis.

1 Introduction

Any online service faces the threat of being taken offline due to Distributed Denial of Service (DDoS) attacks. These attacks overwhelm the target with spurious requests, exhaust its server capacity and render the service unavailable. The emergence of the DDoS-as-a-service economy made these attacks more accessible and decreased the barriers to entry [28]. Also termed booters, these services enable customers to purchase and launch DDoS attacks against a target of their choice. The downtime due to these attacks imposes significant economic costs in terms of availability, recovery and reputation [9]. It has been estimated that an hour of downtime causes between 61K and 67K USD in loss of revenue [46]. Moreover, in some cases, consumers tended to diversify their providers as a countermeasure against DDoS in the aftermath of DDoS attacks [7]. There has also been a significant negative impact on the stock prices of companies when DDoS attacks disrupt the service to consumers [5].

As is often the case [8], the persistence and growth of DDoS attacks can be better explained by a lack of incentives rather than a lack of technical solutions. The actors best positioned to implement controls for DDoS attacks – the intermediaries – do not have sufficient incentives to do so. For instance, Source Address Validation (SAV) is a technical solution that prevents source address spoofing in a network and therefore limits the prevalence of certain types of DDoS attacks. In an empirical study of 334 ISPs, Lone et al. [36] found that 73% of the ISPs have not fully deployed SAV in their networks. They concluded that the lack of complete adoption is due to a lack of incentives for the network operators. While the cost of the implementation is borne by the operator, the benefits are reaped by the rest of the internet.

This lack of incentives is despite most victims of DDoS attacks residing in broadband access networks, with relatively fewer targets in hosting and enterprise networks [41]. However, this victim distribution is for DDoS attacks that use amplification techniques: sending small spoofed packets that trigger an amplifier – such as a DNS, NTP or SNMP server – to send a much larger response to a victim. These attacks are typically purchased from booters which offer low-powered but accessible and cheap attacks to consumers who use them predominantly against other consumers, such as in the context of gaming. While these low-powered attacks place some additional stress on the broadband networks, the consequences are not severe enough for network operators to take stronger preventive actions like SAV.

However, the last few years have seen the emergence of powerful DDoS attacks from IoT botnets. An attacker infects and takes control of a large number of IoT devices to create a botnet. The collective power of all the devices can then be used to launch attacks with magnitudes higher than earlier techniques. The notorious IoT botnet-powered DDoS attack on 'Dyn' in 2016 was the largest attack seen till then and disrupted services like Reddit, Twitter and CNN. More recently, in August 2022, Imperva reported on a DDoS attack using IoT botnets that had a total of 25.3 billion requests setting a new record for the largest DDoS attack mitigated by them [19]. Moreover, IoT botnet-based DDoS attacks are not only more powerful, but their proliferation is also increasing at an alarming pace. There are reports that current geopolitical tensions and hacktivism have triggered an increase in the proliferation of botnets [39].

What remains unknown, however, is the consequence of this change on the victimisation pattern. Understanding the impact of the change on the victim distribution is important because it might trigger a subsequent change in incentives. Our current knowledge of victim distribution is shaped by high-profile attacks that make the news and industry reports based on limited visibility from their customer networks. On the one hand, we could argue that the change in the attack vector would not make any difference. The operators of the attack infrastructures are not the actors ordering the attacks; their clients are. Thus, one hypothesis would be that the infrastructure is simply a tool, and the actor ordering the attack does not care which tool is used as long as it gets the job done.

On the other hand, IoT botnets enhance the magnitude of DDoS attacks and, at the same time, undermine our current DDoS mitigation techniques, like scrubbing [40]. So another hypothesis would be that this combination of increased attacker power and decreased defence capability would enable the attackers to go after better-defended targets, albeit at a higher cost, thus changing the victimisation patterns. Without a sufficient investigation into the victimisation patterns of IoT botnet-based DDoS attacks, we cannot confirm if the change in attack technique has caused a change in the corresponding targets. Thus, it is essential to study the change in victimisation patterns not only because it is under-explored in literature but also because the findings can help us understand the change in and distribution of the incentives. This would be a necessary step in identifying the changes necessary to law and public policy to better align the incentives.

In this paper, we address this gap. We find out what the change in attack vector means for the victims and how this might reshape incentives to invest in DDoS defence measures. We identify the victims of IoT botnet-based DDoS attacks and compare them to earlier attacks using amplification techniques. We do not know if the victim distribution identified in 2015 by Noorizan et al. [41], still holds both for the amplification attacks since then or for the more recent attack vector based on IoT botnets. Thus, our primary research question is, *'Who are the victims of DDoS attacks using IoT botnets, and how does the victimisation pattern of IoT botnets compare to that of earlier attack techniques?'*.

We answer the question using two existing data sources on DDoS attacks that are, as yet, under-utilised for studying victimisation patterns. First, we collected attack commands sent by the Command and Control servers (C2s) to Mirai bots from Netlab[4]. Next, for the benchmark, we collected amplification attack data; victim IP addresses from amplifier honeypots dubbed AmpPots [32]. Using these data sets, we compare DDoS commands sent by IoT botnets to honeypot data over 15 months (January 2020 to March 2021). We compare network-level features of the target IP addresses, like the type of Autonomous Systems, and host-level features, like the density of domains hosted on the address, to identify victimisation patterns. Using AmpPot data between January 2016 to March 2021, we also study the longitudinal evolution of victims of amplification attacks over the four years. Further, we map the identified features to the four tenets of Value, Inertia, Visibility and Accessibility outlined in Routine Activity Theory

[4] https://netlab.360.com

(RAT) and use the framework to evaluate how incentives might play out in terms of suitable targets and defences. In short, our contributions are as follows:

- We perform the first empirical study outlining the change in the victimisation pattern of DDoS attacks due to IoT botnets. We compare and quantify the differences in victims of IoT botnet-based DDoS attacks and amplification attacks.
- We identify victimisation patterns of DDoS attacks by characterising the networks where the victims reside, i.e., network type, ranking, geo-location and network size. Our results show that botnet attacks are proportionally more common against Hosting ASes. 36.6% of botnet attacks were against hosting ASes compared to only 21.1% of amplification attacks. We also find that the victims of botnet attacks are more likely to use dedicated hosting.
- We also identify the sectors where the targeted victims operate. Victims of IoT botnets are primarily Small and Medium Enterprises. In contrast, online gaming remains the most targeted industry for amplification attacks, accounting for more than a third of the attacks.
- We use a balanced random tree classifier to distinguish the characteristics of victims suffering an IoT botnet attack vs an amplification attack. We identify statistically significant differences in the rankings of the networks where the victims reside – botnet attacks target high-ranked ASes proportionally more. The classifier also outlined geographical differences in the victim distribution. We find that a larger percentage of victims IPs of botnet attacks were in Europe and Africa, while amplification attacks were more prevalent in the Americas and Asia.
- We characterise the differences in the victimisation pattern using the target properties outlined in RAT. We connect the differences identified to the underlying differences between the attack vectors and draw out policy implications.

2 Background and Related Work

DDoS attacks pose a relevant and significant threat in our current digital landscape. In 2020, DDoS attacks grew more than 50% increasing both in complexity and attack volume [50]. The explosion in network traffic due to the changes caused by the COVID-19 pandemic made it easier for attackers to launch DDoS attacks. Since the servers were already under stress due to high traffic volume, it took a relatively lesser effort to overwhelm the servers with requests and take the service offline. According to an industry report [52], in pure numbers, 25% of all attacks in 2020 were targeted at the technology sector, but the corresponding attack size was relatively low. The healthcare sector, on the other hand, suffered the most in terms of average attack size but was amongst the least attacked industry. 2022 and 2021 saw a slight reduction in the percentage of DDoS attacks since 2020, 9.7% and 3.5%, respectively, but the peak bandwidths in 2021 were almost seven times higher than 2020 [21].

The market for DDoS attacks

An important contributing factor to the high volume of DDoS attacks is the low entry barriers to launch one. For the tech-savvy and motivated attackers, there are YouTube tutorials on creating botnets and launching DDoS attacks. For the non-tech-savvy, amplification attacks can be purchased online from the aforementioned DDoS booter services with an ease similar to online shopping. In 2020, these cost a mere $48 for an hour, $134 for a day and $1,000 for a month [50]. Booters also have more in common with e-commerce websites beyond ease of purchase. Musotto and Wall [38] showed that booters are similar to e-commerce websites in terms of how their products, price and customers are differentiated and also noted that the profit margins are not very high.

On the defence side, the threat of DDoS attacks has created a market for DDoS Protection Services. There is a prominent trend towards increased adoption of these services, especially by large web hosters [26]. Such protection is typically classified into proactive and reactive protection [23]. Proactive protection is always on, looking out for potential attacks and, depending on the exact configuration, includes varying levels of packet analysis to determine which packets to block. Reactive protection, on the other hand, only analyses meta-data of network traffic to detect anomalies and traffic mitigation kicks in only when the analysis points to suspicious activity.

Technical studies on DDoS attack: Amplification and botnet based

Technical studies on understanding amplification DDoS attacks have outlined detection mechanisms [27,15,51] while other studies investigate the various protocols that are commonly abused for amplification attacks. They identify commonly used protocols are the UDP-based NTP, LDAP, OpenVPN, ARMS, Ubiquity Discovery Protocol and the like [31] and also observe that there are over 2.5k DDoS attacks in a single day. Kührer et al. [33] reported on the significant diversity in amplifiers used in DDoS attacks and also estimated that TCP handshakes can be abused to cause up to 20x amplification.

Similarly, several studies have contributed to our technical understanding of IoT botnets and specifically of Mirai. A seven-month retrospective analysis of the Mirai botnet [11] studied its emergence, the evolution of its variants, and the competition for vulnerable hosts. Notably, it also pushed forth the understanding that Mirai marks a significant change in the evolutionary development of botnets, both due to the simplicity of its infection vector and its exponential growth. This provided a wake-up call to prioritise the security of IoT devices to prevent such severe DDoS attacks [30]. To that end, Jerkins [24] catalogued vulnerable IoT devices using the same attack vector as Mirai motivating manufacturers to address their poor security practices. Similarly, Rodríguez et al. [45] identified device types and manufacturers of Mirai-infected IoT devices through Web-UI image scans and banner analysis. With regard to the cleanup of Mirai-infected devices, Cetin et al. [13] show that quarantining and notifying infected customers through the ISP has the maximum impact. 92% of infections were remediated within two weeks, and only 5% were reinfected in five months. Others have also pointed to the significant role of broadband ISPs in combating the spread of IoT botnet infections like Mirai [42].

Victims of DDoS attacks

Commercial DDoS protection services claim that any business can be a target for DDoS attacks while available prior research [18,29] on victims places gaming-related services and end hosts at the forefront. Moreover, targets are typically attacked by different types of attacks, and web servers are targeted most often [25]. Other studies on victims were conducted to better understand attacker motives. Abhishta et al., [4] used RAT to analyse the victim properties of 26 DDoS attack events that made the news. They argue that economic reasons are only one of the possible motives and advise companies to monitor the social, political and cultural dimensions of their environment to have a better understanding of the underlying threats. Another study on attacks on Dutch educational institutions [6] lends evidence to this claim. It found a significant correlation between the academic schedules and the attack patterns leading to the conclusion that the attacks were launched by an actor who would have benefited from the disruption to the educational activity.

However, while there is information on the high-profile attacks that make the news, either owing to the target or the severity, there is scarce info on other attacks. It is important to note that while high-profile targets might have protection and redundancy in place to mitigate the severity of the attacks, other businesses might not have sufficient protection in place to prevent even less severe attacks. Moreover, there is no work on distinguishing the victims or targets of DDoS attacks via botnets and amplifiers.

3 Data Sources and Methodology

As mentioned in the Introduction, to conduct this victimisation study we use previously collected data sets on DDoS attacks that have been under utilised to study victimisation patterns. To study the victimisation of IoT botnets, we collected Mirai attack data from the Network Security Research Lab NetLab 360's website[5]. For the comparison to earlier attack, we used amplification attack data, collected through AmpPots [32]. To the raw data obtained from these sources, additional data was added to enable meaningful analysis.

As mentioned in the Introduction, we draw upon Routine Activity Theory (RAT) to analyse the patterns of victimisation observed in both the types of attacks. RAT posits that crime happens at the convergence of space and time where a motivated offender and a suitable target are present in the absence of a capable guardian [16]. Although originally developed for offline crime, RAT has been adapted to the context of online crime [53]. RAT outlines four properties

[5] https://data.netlab.360.com

that affect the target suitability – Value, Inertia, Visibility and Accessibility, often referred to as VIVA. Value is the gains to the attacker from the attack, Inertia is the target's resistance to the attack, Visibility is the degree of exposure of the target to the attacker and Accessibility is the reachability of the target. We map each of the victim attributes analysed to the one of target properties outlined in RAT to study the underlying differences that drive target selection. Figure 1 shows an overview of the data analysis process including the mapping of the attributes to RAT properties.

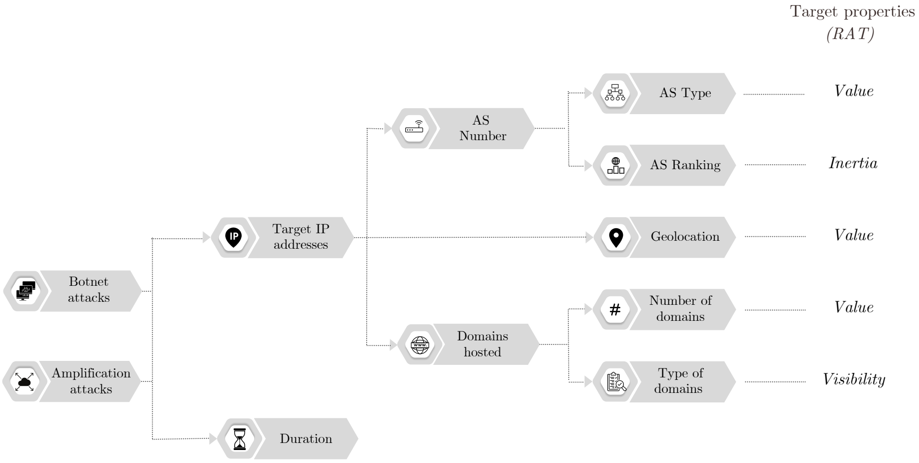

Fig. 1: Overview of attributes analysed and the corresponding mapping to RAT properties

3.1 NetLab Data

The Mirai botnet attack data used in this analysis is sourced from the Network Security Research Lab, Netlab 360. They used analyser programs to heuristically analyse and extract C2 domains or IP information from samples of the Mirai malware. They then track these C2 servers and publish the command information received from the C2s. A detailed explanation of their methods to extract configuration data, attack methods and dictionaries of usernames and passwords from Mirai samples and to classify and track its many variants is provided in [35]. As part of their OpenData Project, till mid March 2021, they released portions of the Mirai attack data thus captured on their website[6]. We scraped and downloaded this attack data set from their website over the collection period between

[6] https://data.netlab.360.com/mirai-c2/

Jan 2020 and March 2021. We scraped and downloaded this attack data set from their website over the collection period between Jan 2020 and March 2021. The data contains a snapshot of commands sent by C2s to Mirai infected devices and includes the time of the attack, the target IP and port and the duration of the attack. The C2 server IP address though available is obfuscated. We found minor inconsistencies in the data, such as port numbers outside the range of 0 to 65535. However, these are likely artifacts of Netlab's data collection methodology or our scraping. Given the low occurrence of such inconsistencies, we omit these data points from our analysis. On average, there were 596 unique target IPs observed each day over the entire period of observation, with a minimum of 11 and a maximum of 1,169 IPs.

3.2 AmpPot Data

The Amplification attack data analysed was collected between Jan 2016 and March 2021 through amplifier honeypots, termed AmpPots and its working is explained in detail in the original paper [32]. In short, these honeypots mimic services commonly abused by attackers for amplification attacks and send back legitimate responses. These services include QotD (17/UDP), CharGen (19/UDP), DNS (53/UDP), NTP (123/UTP), SNMP (161/UDP) and SSDP (1900/UDP). Attackers are thus lured into using these honeypots as amplifiers and data on ongoing attacks, targets and techniques are collected by the AmpPots. These amplifiers are deployed in Japan and depending on the ISP their IPs change every 5 to 30 weeks.

Table 1: Description of AmpPot data over the period of analysis

Time Period	Number of AmpPots	Types of AmpPots
Jan 2016 to May 2017	9	7 proxied and 2 agonostic
June 2017 to March 2018	7	7 proxied
March 2018 to April 2021	19	11 proxied and 8 agnostic

Over the four years, there were differences in the number and type of sensors used which are illustrated in Table 1. Proxied sensors imitate the functionality of the underlying protocol abused by the attackers. They forward the request to internal servers running the abused protocol and send the responses back to the client. Agnostic sensors, on the other hand, reply with a random bytes of response irrespective of the validity of the request. These operate with the assumption that attackers are more concerned about finding hosts that send back large responses than the validity of those responses. However, in this study, we focus on the larger overarching trends in the victims of these amplification attacks rather than the variations due to the differences in the sensors.

In order to separate attacks from scans, an attack is defined as a series of at least 100 consecutive packets where consecutive is defined as less than 60 seconds apart. This is a change from the 3600s and 600s definition used in the earlier papers [32,41] but it allows for analysis at a more granular level. On average, the AmpPots observed 9,475 unique target IPs per day over the entire observation period.

The number of attacks per month for each of the data sets is shown in Figure 2. The size of the AmpPot data set is higher by two orders of magnitude. However, we are comparing the relative proportions of attacks across various aspects and therefore the difference in absolute sizes does not impact the veracity of the results.

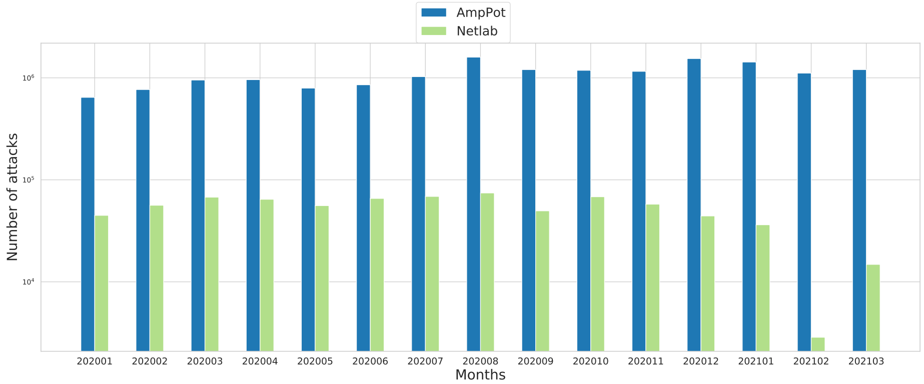

Fig. 2: Attack distribution over the months

4 Methodology

Before describing the methodology used for comparison, we would like to clarify the various terms used in the study. We use the word 'target' or 'target IP' to denote the entity the attacker intended to affect. However, since we can not directly observe the attacker's intention, inline with earlier work [41], we use the term 'victim' or 'victim IP' to refer to the targeted IP address. From both attack data sets we extracted three main attributes for analysis: the duration of the attacks and destination port, both of which were compared directly and the victim IP address, for which additional data was collected for analysis.

Comparison of AS types.

We first compare the types of Autonomous Systems (ASes) that the victims belong to. Since DDoS attacks also impose significant stress on the networks or ASes that the victims reside in, we refer to these ASes as victim ASes. To get the AS that the target belongs to, we looked up the Autonomous System Number (ASN) of the targeted IP using historical BGP routing data obtained from

Routeviews[7]. These files were loaded into the `pyASN` package[8] freely available for python to perform IP to ASN conversion.

Once we obtained the victim ASN, we used a previously built research database for checking the victim AS type. This database was built manually over the years and is organized around ground truth data from an accurate commercial database - Telegeography GlobalComms Database Service [3]. The mapping accurately distinguishes and labels ASes as broadband ISPs, hosting, governmental, mobile ISP, educational and other types of networks.

In addition, we further improve the classification by identifying hosting ASes using the same heuristic as Noroozian et al. [41]. We classify as hosting any AS that was not classified using the database and contains more than 2,700 second level domains (SLDs). To get the count of SLDs per AS, we use a large passive DNS (pDNS) database provided by Farsight Security [2]. The database contains the mapping of IPs and the corresponding domains that they resolved to over our period of observation. We aggregated the IP level domain count based on the corresponding ASN they belong to and obtained the count of SLDs per ASN.

We corroborated our classification of AS types with those from ASDB [54] which classifies AS types using machine learning techniques on data from RIR's WHOIS and Business Intelligence Databases. We found that both our classification and ASDB have similar coverage rate for our data set. For the Netlab data set, ASDB has 19.6% unknowns while our classification has 12.5%. A similar pattern is observed for AmpPot as well. Since the missing ASes are mostly on the tail-end of the frequency distribution – ASes which are not commonly attacked – they do not have a significant impact on our results.

The main difference between the two databases is in the percentage of ISP broadband and hosting ASes. The percentage of ASes classified as ISP broadband is higher in ASDB compared to our database while the percentage of hosting ASes is lower in ASDB. The difference is due to classification of ISP broadband ASes – ASDB classifies as ISP broadband the ASes our classification marks as Hosting. However, since the ISP broadband ASes in our database have been identified using accurate information from Telegeography and have also been manually validated, we consider our database to be more accurate for our data set. Moreover, we have classified the ASes based on the predominant use of the network, either ISP broadband or hosting, while in some cases other types of users might also be present.

Comparison of AS Ranking.

To compare the size and connectivity of the victim ASes in the botnet and amplification attack data sets, we used CAIDA's AS Rankings [1]. These rankings are calculated using customer cone sizes derived from BGP routing data and CAIDA's topological data[9]. An AS's rank is inversely proportional to the size

[7] `http://archive.routeviews.org/`

[8] `https://github.com/hadiasghari/pyasn`

[9] `https://www.caida.org/projects/ark/`

of its customer cone – the sum of its direct and indirect customers. The indirect customers are the customers that can be reached through the ASes that a given AS peers with. These ranks denote both the influence of an AS in the global routing system and its size[10].

Comparison of Victim Location.

We use MaxMind's GeoIP location database [11] to identify the geographical location of the victim IPs in both data sets. The service provides the country an IP address is located in with an accuracy of 99.8%.

Comparison of domains hosted on target IPs.

To compare the domain level attributes, we use the same pDNS database described earlier. We retrieved the domains hosted in the top 100 most common IPs for both sets of data across the entire time period of observation. Since we were interested in domains that are hosted, we restricted ourselves to IPs within hosting ASes. We collected the super set of all the domains hosted on these IPs for each month that they were observed in the data sets.

In addition, to get an estimate of the value of the domains hosted in these IPs, we got the Tranco ranking for these domains. Tranco [12] provides a transparent and reproducible popularity ranking of websites. We then manually analysed these domains to identify their types. Since most of the domains did not have an associated Tranco ranking, we analysed the domains with Tranco ranking separately from others without a corresponding ranking.

Modelling.

In order to better study the differences in the victimisation patterns of the two attack techniques, we constructed a Balanced Random Forest classifier [14]. Random forest is a supervised machine learning algorithm that uses multiple decision trees to arrive at a final class output. Owing to the differences in the size of our data sets, we used a Balanced Random Forest Classifier that is available as part of the imbalanced-learn library [34]. This uses random under sampling of bootstrapped samples to balance the size.

We only selected properties of victims as features because we were interested in the differences in the victimisation patterns across the two data sets. This implies that the duration and port, although they relate to the victim, were not included in the model since they are properties of the attack itself rather than the victims. There were four features input to the classifier. Of these, two were ordinal – the domain count of the victim IP and the the CAIDA ranking of the victim ASN. The other two features were categorical and were one-hot encoded

[10] https://asrank.caida.org/about

[11] https://www.maxmind.com/en/geoip2-country-database

[12] https://tranco-list.eu/

before being input to the model. These were the region, based on the geo-location of the victim IP, and the AS type of the victim ASN. The countries output by the geo-location were grouped into regions for a more concise representation. The duplicate data points were dropped before being input to the model. We ran the model for varying values of number of estimators (the number of trees the model constructs) and maximum depth of the tree and picked the values that had the best accuracy.

RAT properties

As mentioned earlier, we map each of the attribute analysed to one of the four properties of RAT – Value, Inertia, Visibility and Accessibility. We do not use the property of accessibility since all the targets are hosted on the internet and can be reached by any attacker with an internet connection. We mapped the AS level attributes, AS type and ranking, to Value and Inertia respectively. The value gains to an attacker from targeting a specific type of AS, say hosting, will be higher than targeting a broadband customer. The former involves monetary loss from sites hosted on the target IP while the latter will result in lost connection for an individual customer. High ranked ASes will offer higher resistance to the attack and therefore have higher inertia when compared to other low ranked ASes. The location of the target also relates to value. Targets in countries with higher ICT index are more valuable because of the higher dependency of the country on these services. The domain level attributes number of domains and type of domains both relate to visibility. The higher the number of domains that are hosted on a target IP, the higher it's reachability or visibility. Similarly, some types of domains like X are more visible than other types of domains, say Y.

Ethical Considerations.

We adhered to our institution's ethical policy at all times and appropriately handled issues concerning data preservation and data sharing. For the botnet attack data set from Netlab, we notified them about our interest in their data set, and the scraping scripts were designed to minimise the load on their servers. The amplification data set was collected via honeypots. In order for a honey pot to successfully lure attackers, it needs to participate in the attack to a certain degree. However, each honey pot deployment has rate limiting mechanisms to minimise the impact of the participation to a negligible degree.

5 Results

5.1 Distribution of targeted AS types

We first examine differences between amplification and botnet attack victims by comparing differences among the autonomous systems in which victim IPs reside, i.e. by comparing victim ASes. We compared the distribution of victim

AS types identified (as described in section 4) over three dimensions Figure 3 – the percentage of unique IPs, the percentage of attacks and the percentage of unique ASes in each data set. The comparison of the distribution of unique IPs across the AS types shows that the highest percentage of victims in both AmpPot (63.8%) and Netlab (47.1%) data sets are in broadband ISPs. Moreover although the second highest category for both is hosting AS, only 14% of AmpPot victims belong to hosting while about 32.4% of Netlab victims are in hosting ASes.

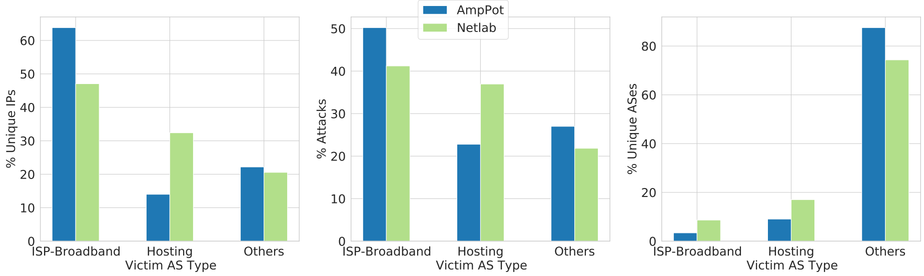

Fig. 3: Comparison of percentage of unique ASes, attacks and IPs across the AS Types

A similar trend is observed in the distribution of unique attacks across the ASes. The most common AS Type for both data sets is ISP broadband (AmpPot - 48.1% and Netlab - 41.7%). The next highest in AmpPot is Others (30.8%) while for Netlab it is hosting (36.6%). However, when it comes to unique ASes the distribution becomes more interesting. The majority for both (AmpPot 87.6% and Netlab 74.3%) is Others but the second most common is hosting for both (AmpPot 9.4% and Netlab 17.6%) and not broadband ISPs (AmpPot 3.4% and Netlab 8.7%). This shows a remarkable concentration of victims in broadband ASes: 64% of victim IPs in the AmpPot data set are from 3.4% of ASes and 47% of victim IPs in Netlab data set are from 8.7% of ASes. Moreover, even when using ASDB for AS Type classification, the proportional distribution of attacks over the AS types is similar.

Further, we see that the distribution of victim ASes across the AS Types has remained relatively stable when compared to earlier work [41]. The ISP broadband AS type still has the highest number of attacks (48%) and hosts the most unique IPs (64%). However, the percentage of attacks in the Others category which includes education, government, gaming, ISP-Mobile and unknowns among others, has increased.

5.2 Rankings of the targeted ASes

Next we compare the differences in the CAIDA AS rankings of the victim ASes. As seen in Figure 4, we observe a difference in the mean ranking of ASes in

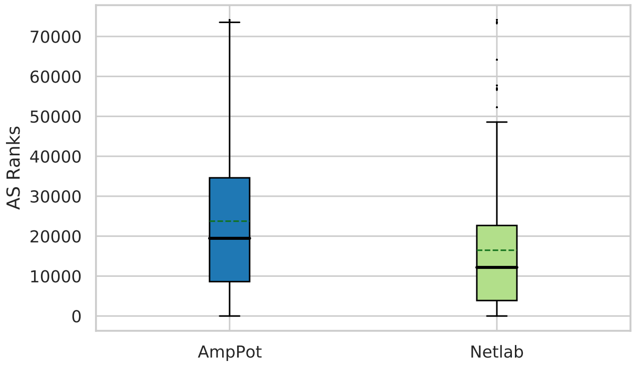

Fig. 4: Comparison of AS ranks between AmpPot and Netlab data sets

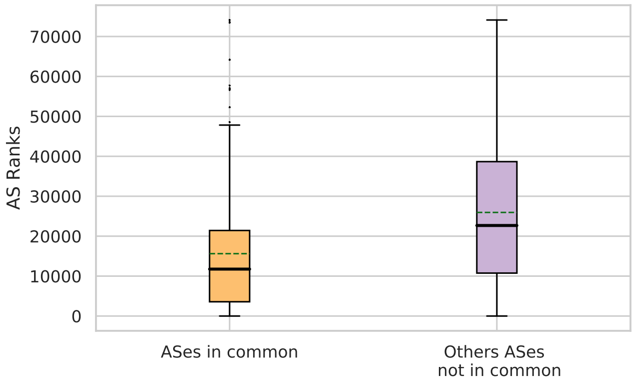

Fig. 5: Comparison of AS ranks between ASes common to both data sets and those unique to each data set

the AmpPot and Netlab data sets. The mean rank of ASes in the AmpPot data set is 23,749.9 while for the Netlab data set it is 16,471.5. In order to check if this difference is significant, we performed the Mann-Whitney U test. The results indicate that there is a statistically significant difference (p<0.001). This indicates that the relative size, connectivity, and therefore influence of ASes in the Netlab dataset are higher than those in the AmpPot dataset.

However, it should be noted that this high ranking (and higher influence) is also seen in the ASes that are common to both AmpPot and Netlab as shown in Figure 5. The Mann-Whitney U test also showed a statistically significant difference in the ranking of these two sets of ASes (p<0.001). We thus see that the larger, more influential ASes contain victims targeted by both amplification and botnet attacks, but the proportion of botnet attack victims in these ASes are significantly more. 94.7% of unique ASes in the Netlab data set are in this common group compared to 23.6% of unique ASes in the AmpPot data set.

5.3 Geographical distribution of the targeted IP addresses

We then compared the geographical distribution of victim IP addresses across our datasets. As mentioned earlier, we used MaxMind's GeoIP[13] location to get the geo-location of the victim IPs.

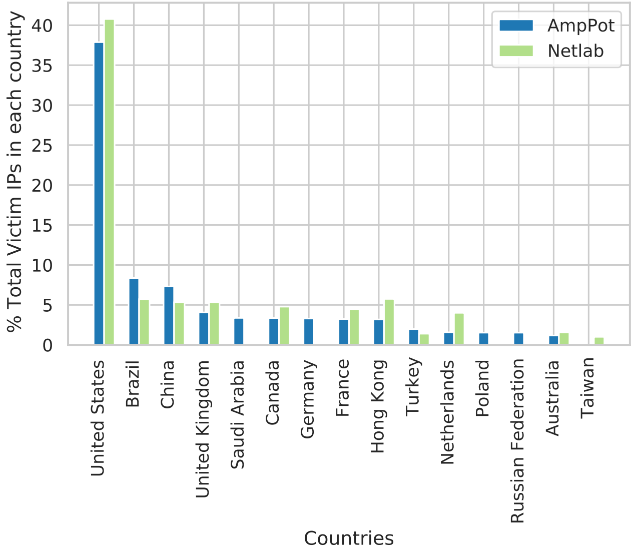

Fig. 6: Countries with more than one percent of victim IPs

Figure 6 shows the top ten countries with the highest percentage of victims. We found that United States has the highest victim IPs in both data sets (Netlab - 44.6% and AmpPot - 37.3%) by a large margin. The next highest country with most victims of botnet attacks is Canada (5.75%) and for amplification it is Brazil (11.6%). Interestingly, in both data sets, China has the third highest number of victim IPs (7.5% in AmpPot and 5.7% in Netlab) and United Kingdom has the fourth highest (3.9% in AmpPot and 5.3% in Netlab). Though Saudi Arabia was the fifth highest country with victim IPs of AmpPot (3.3%), only 0.2% of victim IPs in Netlab were located in Saudi Arabia. The distribution of percentage of attacks across countries with more than one percent of attacks in each data set is shown in Figure 6. Although the graph seems to show an over representation of botnet victims in countries with high GDPs, we did not find any statistically significant correlation. The cross comparison of victim location across our amplification attack data and botnet attack data also suggest minor differences.

[13] https://www.maxmind.com/en/geoip2-country-database

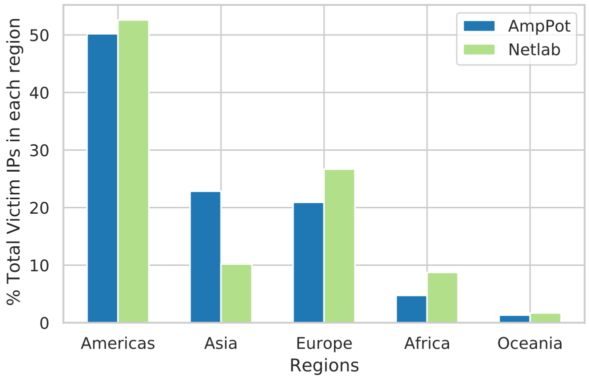

Fig. 7: Percentage of victim IPs across different regions

We then grouped the countries by regions using the Standard area codes provided by the Statistics Division of the UN [14] to study the differences at a more aggregate level. When grouped by regions, the Americas (North America and South America together) rank the highest in both data sets (AmpPot 53.9% and Netlab 52.6%). However, the next highest region with victims of amplification attacks is Asia (21.4%) which has half as many botnet victims (10.2%). The second highest percentage of botnet victims are in Europe (26.7%) which also has 19.4% of amplification victims. We see that the percentage of botnet victims in Africa is twice that of amplification victims. The distribution of the percentage of attacks across all regions is illustrated in Figure 7.

5.4 Comparison of domains hosted in the targeted IPs

Next, we compare victims by examining domain resources hosted behind the attacked victim IPs in our data sets. As described in the Methodology (section 4), we used a passive DNS database to obtain the number of domains hosted on the target IPs in both data sets. We calculated the domain count per IP as the average of the number of domains hosted on the IP through all the months that the IP was seen in the data set. We then compared the domain counts of unique attacks on hosting ASes in each data set. We saw that 77% of attacks against hosting ASes in AmpPot and 40% of attacks in Netlab had a domain count of zero. The cumulative distribution function of the remaining attacks with domain counts greater than zero is shown in Figure 8.

We observe that both the mean and median domain counts of attacks in the AmpPot data set are higher than that of Netlab. The mean and median domain counts for attacks in the AmpPot data set are 26.7 and 10 respectively, while for Netlab they are 2.5 and 2. The significant difference in domain counts illustrated in the box plot in Figure 9 and is also confirmed by the results of the

[14] https://unstats.un.org/unsd/methodology/m49/overview

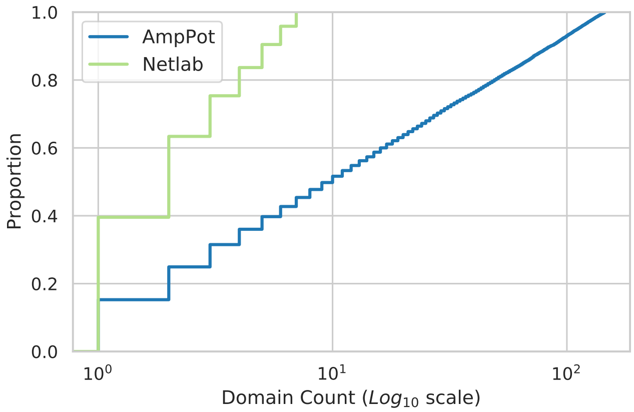

Fig. 8: CDF of domain count of unique attacks in hosting ASes in both data sets

Mann-Whitey U-test (p<0.001). The difference without dropping domain counts of zero is also significant (p<0.001).

5.5 Analysis of domains resolving to top 100 most common IPs.

In order to compare the types of domains hosted on the victim IPs in each data set, we manually analysed them. We extracted the domains hosted on the top 100 most common IPs in each data sets and dropped outlier IPs with significantly larger domain counts (less than 1% of the IPs). This gave us 274 unique domains on Netlab and 418 on AmpPot.

Comparision of Tranco rankings As mentioned earlier, to get a better estimate of the value of these domains, we got their corresponding Tranco ranking[15]. We observed that 87.6% (240) of domains in Netlab and 96% (401) of domains in AmpPot had no associated Tranco ranking. The domains and corresponding ranking in each data set, where available, are presented in Table 2 and Table 3. The average Tranco ranking for the domains unique to AmpPot is 254,141.3 while for Netlab it is 199,796.6. If we use the Tranco ranking of popularity as a proxy for value, we see that the targets of AmpPot have lower value than those of Netlab.

Analysis of domains with Tranco rankings We analysed the subset of domains with an associated Tranco ranking, separately from the rest of the domains. There were six domains in common between the two data sets. These were two hosting/Cloud providers, three Network Service Company Websites

[15] https://tranco-list.eu/

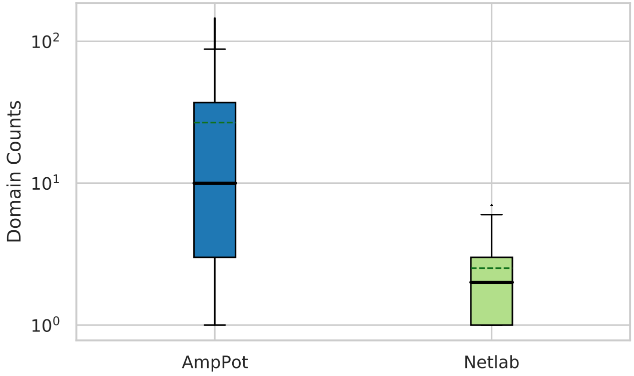

Fig. 9: Box plot comparison of domain counts

(Norton, RIPE, Geolocation API) and two unreachable domains (KKK.com and KKK.bz). Of the rest, AmpPot has four hosting/Cloud provider sites, two African news sites, one UK LGBTQ website and one unknown (17tahun.com). In Netlab however, the remaining domains have a wider classification: nine hosting/cloud provider sites, four gaming related sites, three domains of messaging platforms (discord, telegram and IRC), two pages linking to drugs, two university websites, one each of a Psychic reading website, a News/Information site, an LGBTQ site and a porn site and finally two unknowns.

Table 3: Domains and corresponding Tranco rankings where available - AmpPot

Domain name	Tranco ranking
secureserver.net	1,107
aliyuncs.com	1,930
allafrica.com	4,292
ripe.net	6,674
incapdns.net	8,006
transip.net	226,059
lgbt.foundation	296,282
pro-norton.com	437,159
17tahun.com	741,612
africanews.org	818,292

Analysis of domains without Tranco rankings Most of the domains in the subset of domains without an associated Tranco ranking did not resolve to an

Table 2: Domains and corresponding Tranco rankings where available - Netlab

Domain name	Tranco ranking
avast.com	725
discord.gg	1,242
ovh.com	2,561
ripe.net	6,674
your-server.de	15,249
hetzner.com	25,366
hetzner.de	28,316
acquia-sites.com	38,906
2ksports.com	44,098
psychic-readings-for-free.com	46,716
zbigz.com	91,006
nexus-cdn.com	141,504
unsam.edu.ar	167,213
dathost.net	179,373
dal.net	194,209
softether.net	202,438
verygames.net	223,435
honglingjin.co.uk	251,661
aloneproxy.top	259,587
sexdrug.tech	343,343
fuckarea.biz	345,721
bytebx.com	360,102
sexwax.me	367,564
omgserv.com	392,383
lesbian.com	404,008
iproxies.club	405,449
prick.top	407,477
clouvider.net	647,979

IP address. We therefore ran a check with a domain registration database [16] to get details about the registration status. We found that only 77% and 61 % of domains were still registered as `active` domains respectively in the AmpPot and Netlab datasets. We took a 100 random samples from each of the domains registered as `active` for manual analysis. However, despite being registered as `active` 54% of domains on Netlab and 63% of domains from AmpPot were not accessible. The errors varied from 'Connection refused' (401) to 'Name not resolved'. Of the remaining which were accessible, the distribution is given in Table 4.

We see that the most common type of amplification attacks are on gaming related website in line with earlier research [41]. However, although there exist gaming related victims within Netlab, they are not the most popular. The most common victims of botnet attacks are Small and Medium Enterprises (SMEs).

[16] https://who.is/

Moreover, AmpPot has no domains relating to hosting or Cloud or non-gaming related Servers while Netlab has a fair share of those across both the data sets.

Table 4: Categories of domains in AmpPot and Netlab

Category	**AmpPot**	**Netlab**
Gaming related	26	9
SME	8	18
News/Information	2	-
Network related	1	5
Hosting/Cloud/Server	-	11
Porn/Suspicious	-	3

5.6 Duration of the DDoS attacks

Next, we examine amplification and botnet attack victims differences by analysing and comparing the duration of attacks directed at each victim across data sets. The average duration of an attack in the AmpPot data set is 754.94 seconds, with a median of 164 seconds. However, while the median duration in Netlab is 80 seconds, the average duration is much higher 793,220.66 seconds (about nine days). The high average is undoubtedly driven by a few outliers in the data. These outliers include values like 4,294,967,295 seconds – the maximum value possible in 32 bits (0xFFFFFFFF) – which amounts to about 136 years. We looked at the Mirai source code[17] and found that the attack function returns an error when the duration value is greater than 3600 seconds. We therefore dropped attacks with duration higher than 3600 seconds for the comparison (about 1.1% of total attacks). Since the duration values in the AmpPot data set are obtained through observation of actual attacks via a honeypot, we did not drop any outliers. The longest attack in the AmpPot data set is 2,466,626 seconds (about 28.5 days).

[17] https://github.com/jgamblin/Mirai-Source-Code/blob/master/mirai/cnc/attack.go

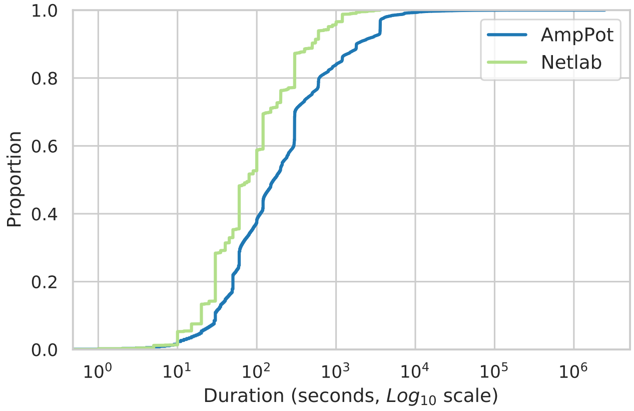

Fig. 10: CDF of duration of attacks

Figure 10 shows the cumulative distribution function for the duration of both data sets. After dropping the outlier durations in Netlab, the average duration is 195.5 seconds and the median is 80 seconds. The results of the Mann-Whitney U-test shows that the differences in the duration are significant ($p<0.001$).

5.7 Modelling

As outlined in the Methodology (Subsection 4), we ran a Balanced Random Classifier model to check for differences in the features of the victim in the two sets. We got the highest accuracy (0.66) with a maximum depth of 8 and 500 trees; the feature weights output by the model are shown in Figure 11.

The results show that highest contributor to the difference between the victims is the domain count of the IPs followed by the ranking of ASes. This is also in line with our analysis which shows significant differences in both the domain counts and the AS ranking. The differences across regions, though lesser by an order of magnitude, is mostly similar to our region analysis. However, where our analysis only showed minute differences in the percentage of victims in the region of Americas, we see that it is the second highest driving factor for differences in region. Upon closer inspection into this divergence using a visual tree interpreter tool, we found that the model uses a higher value for the Americas region as a higher weightage for the AmpPot class. Interestingly, there are only negligible differences amongst the contribution of the different AS types to the classification.

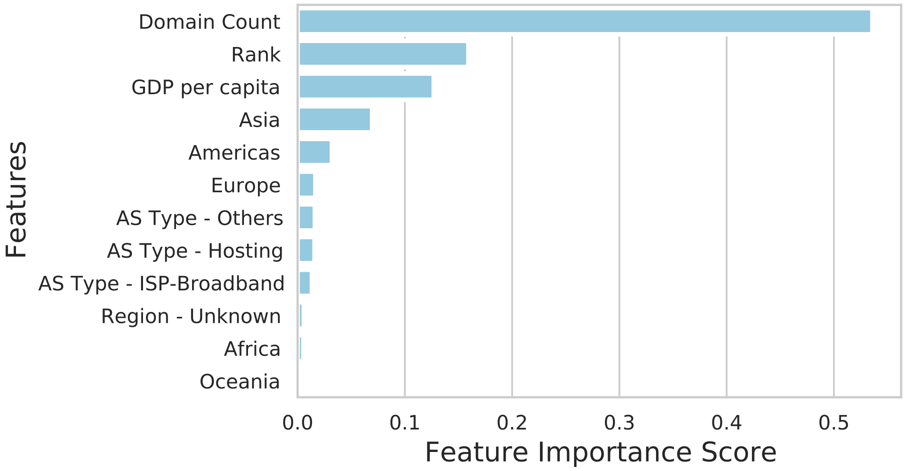

Fig. 11: Random Forest feature importance scores

6 Discussion

Our primary aim with this research was to identify the victimisation pattern for IoT botnets and compare it to the pattern for amplification attacks. We postulated that if there is an underlying difference in the actors deploying these attacks and the corresponding targets, we would see this difference reflected in the victimisation pattern. Conversely, if the attack vector is irrelevant to the actors ordering the attacks, we would not see any difference in the victimisation pattern.

Our results show clear differences in the victimisation pattern within all six attributes analysed. The AS type analysis indicates that a larger percentage of botnet victims reside in hosting ASes. The CAIDA AS ranking shows that relatively more botnet victims are in high-ranking ASes that serve a larger consumer base. The geographical differences show a higher percentage of botnet attacks in countries with higher GDP per capita, albeit with a few exceptions. At the IP level, we observe that victim IPs of botnet attacks within hosting ASes have a higher domain density and also host more popular domains, as indicated by the available Tranco rankings. The classifier highlights the significance of domain count and AS ranking in driving the differences in the victimisation pattern and also draws out other geographical differences. Next, we place these results in the context of the RAT dimensions for an accessible target – Value, Inertia and Visibility. We have not considered accessibility since all the targets are accessed via the internet.

Value Value refers to the gains for the attacker from attacking the target, monetary or otherwise, like status or prestige. Of the six attributes that we analysed, three relate to value – AS type, geolocation and the number of domains. For all three, we find that the higher value of a target correlates with a higher prevalence of botnet attacks on the target. We postulate that higher the value of the target, higher the value to the attacker from attacking the target.

As mentioned earlier, hosting ASes have a higher value from an attacker's perspective than broadband ASes. Hosting ASes charge higher for their service compared to broadband ASes. Therefore, a deterioration in the quality of service due to a DDoS attack on one client will have a higher impact on the revenue for hosting ASes than broadband ASes. Our results show that though broadband ASes are the most common target for both types of attacks, botnet attacks are relatively more common against hosting ASes.

Similarly, countries with higher GDP per capita and ICT development index have a higher dependency on IT services and therefore derive a higher value from them. Although not statistically significant, at a country level, we find a higher occurrence of botnet attacks against victims in the US, UK, Canada, Germany, France and the Netherlands.

Finally, the number of domains hosted on an IP is a clear indicator of value due to the difference in pricing structure between shared hosting and dedicated hosting. Shared hosting, with a larger number of domains per IP, is cost-effective and easy to use, while dedicated hosting with fewer domains per IP offers a more stable and predictable performance. This enhanced performance comes at a higher cost and also requires expertise to set up and maintain. Therefore, dedicated hosting is apt for high-value domains with higher traffic and bandwidth requirements, while shared hosting is an attractive option for personal websites and domains with less traffic. While there is no hard threshold for what counts as shared hosting, prior studies have put it at around ten domains per IP [47,48]. We find that the average number of domains on victim IPs is 10 for botnet attacks and 26 for amplification attacks. Thus, by this metric, AmpPot victim domains are more likely to use shared hosting services, while botnet victims are more likely to have dedicated hosting.

Inertia Inertia refers to the resistance offered by the target to the attacker. This could relate to the size of the targets, attacks on smaller targets are easier to execute than those on larger targets, or the defence capability of the target. In this research, we mapped the CAIDA ranking of the ASes to inertia. We find that lower inertia or higher resistance correlates with a higher percentage of botnet attacks.

Botnet attacks are more prevalent against high-ranking ASes. These high-ranked ASes have larger customer cone sizes and higher revenue compared to lower-ranked ASes. DDoS attacks on these ASes will impose a higher societal cost since a larger number of prefixes are reachable through these ASes. These ASes thus have both the means – due to their higher revenue – and the motive – to decrease the impact of attacks – to invest in DDoS protection.

Visibility Visibility refers to the degree of exposure of the target to the attacker. Our manual analysis and classification of domains resulted in six types: Gaming related, SME websites, News/Information, Network related, Hosting/Cloud/Server and Porn/Suspicious. Of these, we group gaming-related, network related and hosting/cloud/server as low visibility domains because the exact domains are known only to those who have intimate knowledge of these services. For instance, the domain to access the configuration of a hosting or cloud service includes the public domain of the service, say 'aliyuncs.com', but has additional sub-domains like 'susharefile.oss-cn-shenzhen.aliyuncs.com'.

We find that these low visibility domains are more common in amplification attack victims, while high visibility domains like SMEs are more common in botnet attacks with one exception. Two domains related to News/Information are among the amplification attack victims, while there are no domains related to news/information within the botnet attack victims.

The framework of RAT thus helps us understand that there are differences in targets or victims driven by the differences in attack type. Botnet attacks are more common against high-value victims with a lower inertia and higher visibility, while amplification attacks are more common against low-value targets with high inertia and low visibility.

Square pegs and square holes We see that due to the differences in the attack type, each attack type lends itself more suited to a particular target. For instance, botnet attacks are better suited for high-ranked ASes precisely because they might have DDoS prevention and mitigation measures. Evading the defences and launching a successful attack is easier with botnets due to the differences in attack traffic.

Traffic from botnets has legitimate operating system-generated protocol headers, which match the statistical distribution typically observed at the application layer. Due to this similarity with legitimate traffic, machine learning-based mitigation techniques do not achieve high accuracy rates when identifying botnet attack traffic [40,20]. In addition, the diversity of botnets, each with unique characteristics [44], makes detection more difficult. On the other hand, an attacker exploiting amplification vulnerabilities in a protocol uses specific values in certain header fields to trigger an amplified response. These characteristic header field values make it relatively easy to distinguish an attack from legitimate traffic [12].

Moreover, DDoS protection techniques, like DDoS fingerprinting [22] that propose sharing rules derived from fingerprints of attack sources, can be effective against amplification attacks due to the limited number of amplification sources. However, they will not help defend against botnet attacks since the attack sources are diverse and distributed across the entire IPv4 space. Thus, the added difficulty of evading botnet-based DDoS attacks makes them a more attractive option against high-value clients.

The costs of each of the techniques might also explain the higher prominence of botnet attacks against high-value targets. Amplification attack infrastructures are easier to maintain than botnets. The public services that can be abused for amplification are widely available; the attackers only need to cover the cost of scanning for open amplifiers and launching attacks on demand.

On the other hand, IoT botnets need constant renewal since infections of most IoT malware families are non-persistent; a power cycle removes the infection. Even if the bot remains infected, the connection with the C2 servers is often lost within a few days, as C2 addresses are hard-coded in the binaries. So, as soon as the C2 is taken down, the bots are stranded [49]. This makes the upkeep of an IoT botnet onerous and time-consuming, potentially driving up its operators' costs. In contrast, amplification services are widely available and often masquerade as benign stressor websites, and they are not subject to similar take-down efforts.

While we could not find any reliable data specifying the costs for each of the attack techniques, the above-mentioned factors support the conjecture that botnet attacks are priced higher. The superior attack power, higher operational costs and the resultant lower availability might play a role in price differentiation between amplification attacks and botnet attacks. The higher price of botnet attacks would also deter the less motivated attackers, e.g., teenagers wanting to win a game of Minecraft, since cheaper options are available. This, in turn, can also explain the higher prominence of high-value domains in the botnet attack data set. Thus, like square pegs matching square holes, certain types of attacks match best with a certain type of target.

Implications for law and policy. The 2016 study on victims of amplification attacks [41] found that the low price of the attacks attracted behaviour that, like vandalism and file sharing, is strictly speaking illegal but not a profit-driven crime. On the other hand, in our results, we see that botnet-based attacks are more costly to execute and go after more valuable targets than, say, individual gamers and thus causing more economic losses. These are more likely part of a profit-driven cybercrime operation or political action. From a crime prevention and mitigation perspective, the policies for these two types of problems are very different.

For low-level crimes like vandalism and file sharing among consumers, the government usually pursues strategies like awareness campaigns and administrative law mechanisms, like statutory fines. In the past, these strategies have been used to tackle the consumer demand for DDoS attacks. An awareness campaign launched by the UK's anti-cybercrime agency warned youth who searched for DDoS booter services online about the illegality of DDoS attacks [43]. This proved to be effective at keeping new users out of the DDoS markets [17].

For profit-driven crime, on the other hand, they usually approach it via criminal law and law enforcement actions. As botnet techniques further evolve and become more widespread, we can expect to see a higher proportion of these attacks in the DDoS landscape. Thus, more of the mitigation will fall to criminal investigations and disruption efforts rather than consumer-focused interventions.

The framework of RAT further highlights the two main possibilities for the prevention of DDoS attacks. We can decrease the number of suitable targets by limiting the attack power of the IoT botnets. This implies having stronger security measures in our IoT devices and thereby decreasing the size of the botnets. Further, we can also increase the capability of the guardians by improving our ability to defend against IoT botnets.

This can be best done via information sharing, which has shown to be effective in mitigating cybercrime [37]. An anti-DDoS coalition, NoMoreDDoS, established in the Netherlands, enables information sharing and collaboration amongst partners to collectively tackle the threat of DDoS attacks [10]. The partners include government organisations, internet service providers and internet exchanges, among others. This illustrates that network intermediaries like high-ranked ASes, the most common targets of botnet attacks, are uniquely positioned to benefit from and initiate such information sharing. They have the incentive to share information – to minimise the stress on their networks. Further, the size of these ASes would also protect them against a potential negative backlash of being cut off from their peers.

7 Conclusion

Denial of Service attacks are almost as old as the Internet, yet our analyses prove that they show no signs of disappearing. In fact, they are growing in number and diversifying in terms of attack vectors. With the increase of tools to identify services that can be misused to perform amplification attacks, DDoS attacks have been commodified and are accessible to criminals with all kinds of different skills. On top of these amplification services, the appearance of IoT botnets made possible to control millions of devices which in turn are also being used to also launch DDoS attacks. All this together has attracted a great variety of criminals launching DDoS attacks against a wide range of victims.

While this research does not demonstrate a displacement of traditional booter services, the commoditization of IoT botnets may alter the market's technological supply. We did not observe repeated victimisation to increase over the years, meaning that the increase in frequency of DDoS attacks is tied to a growth on the number of victims. However, we see that the newer targets operate in large variety of sectors, with small medium enterprises getting more attacks and gaming services still being a common target.

In the arms race between attackers and defenders, the defence measures must routinely adapt to attack techniques to minimise the impact of an attack. As we observe with our results, the newer techniques, which are tougher to defend against, attract enterprising attackers motivated to cause damage to high-value clients. So, amplification and botnet attacks are not of the same feather and their victims are not flocked together. We see that like square pegs finding square holes, victims who are better able to defend against DDoS attacks are attacked using more robust techniques. This stresses the urgency to adopt stronger legal actions against miscreants launching DDoS attacks.

Acknowlegements

This publication is part of the RAPID project (Grant No. CS.007) financed by the Dutch Research Council (NWO). This research was also supported by the Ministry of Internal Affairs and Communications, Japan (JPJ000254), National Institute of Information and Communications Technology, Japan (05201), and JSPS KAKENHI (21H03444, 21KK0178).

References

1. CAIDA AS Rank. http://as-rank.caida.org/
2. Farsight Security (Jan 2022), https://scout.dnsdb.info/dashboard
3. TeleGeography: GlobalComms Database Service (Jan 2022), https://www2.telegeography.com/en/globalcomms-database-service
4. Abhishta, A., van Heeswijk, W., Junger, M., Nieuwenhuis, L.J., Joosten, R.: Why would we get attacked? An analysis of attacker's aims behind DDoS attacks. J. Wirel. Mob. Networks Ubiquitous Comput. Dependable Appl. **11**(2), 3–22 (2020)
5. Abhishta, A., Joosten, R., Nieuwenhuis, L.J.: Analysing the impact of a DDoS attack announcement on victim stock prices. In: 2017 25th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP). pp. 354–362. IEEE (2017)
6. Abhishta, A., Junger, M., Joosten, R., Nieuwenhuis, L.J.: Victim Routine Influences the Number of DDoS Attacks: Evidence from Dutch Educational Network. In: 2019 IEEE Security and Privacy Workshops (SPW). pp. 242–247 (2019). https://doi.org/10.1109/SPW.2019.00052
7. Abhishta, A., van Rijswijk-Deij, R., Nieuwenhuis, L.J.: Measuring the impact of a successful DDoS attack on the customer behaviour of managed DNS service providers. ACM SIGCOMM Computer Communication Review **48**(5), 70–76 (2019)
8. Anderson, R.: Why information security is hard-an economic perspective. In: Seventeenth Annual Computer Security Applications Conference. pp. 358–365. IEEE (2001)
9. Anderson, R., Barton, C., Böhme, R., Clayton, R., Van Eeten, M.J., Levi, M., Moore, T., Savage, S.: Measuring the cost of cybercrime. The economics of information security and privacy pp. 265–300 (2013)
10. Anti-DDoS-Coalitie: About the coalition (March 2023), https://www.nomoreddos.org/en/about-the-coalition/
11. Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., Cochran, J., Durumeric, Z., Halderman, J.A., Invernizzi, L., Kallitsis, M., et al.: Understanding the mirai botnet. In: 26th {USENIX} security symposium ({USENIX} Security 17). pp. 1093–1110 (2017)
12. Bekeneva, Y., Shipilov, N., Shorov, A.: Investigation of protection mechanisms against DRDoS attacks using a simulation approach. In: Internet of Things, Smart Spaces, and Next Generation Networks and Systems, pp. 316–325. Springer (2016)
13. Çetin, O., Ganán, C., Altena, L., Kasama, T., Inoue, D., Tamiya, K., Tie, Y., Yoshioka, K., Van Eeten, M.: Cleaning Up the Internet of Evil Things: Real-World Evidence on ISP and Consumer Efforts to Remove Mirai. In: NDSS (2019)

14. Chen, C., Liaw, A., Breiman, L.: Using Random Forest to learn imbalanced data. 2004. University of California, Berkeley **110**(1-12), 24 (2004), www.stat.berkeley.edu/users/chenchao/666.pdf
15. Chen, C.C., Chen, Y.R., Lu, W.C., Tsai, S.C., Yang, M.C.: Detecting amplification attacks with software defined networking. In: 2017 IEEE conference on dependable and secure computing. pp. 195–201. IEEE (2017)
16. Cohen, L.E., Felson, M.: Social change and crime rate trends: A routine activity approach. American sociological review pp. 588–608 (1979)
17. Collier, B., Thomas, D.R., Clayton, R., Hutchings, A.: Booting the booters: Evaluating the effects of police interventions in the market for denial-of-service attacks. In: Proceedings of the internet measurement conference. pp. 50–64 (2019)
18. Czyz, J., Kallitsis, M., Gharaibeh, M., Papadopoulos, C., Bailey, M., Karir, M.: Taming the 800 pound gorilla: The rise and decline of NTP DDoS attacks. In: Proceedings of the internet measurement conference. pp. 435–448 (2014)
19. Gabi Stapel, N.K.: Record 25.3 Billion Request Multiplexing DDoS Attack Mitigated by Imperva (September 2022), https://www.imperva.com/blog/record-25-3-billion-request-multiplexing-attack-mitigated-by-imperva/
20. Gupta, B.B., Badve, O.P.: Taxonomy of DoS and DDoS attacks and desirable defense mechanism in a cloud computing environment. Neural Computing and Applications **28**(12), 3655–3682 (2017)
21. Heath, M.: 2023 DDoS Attack Trends (Feb 2023), https://www.f5.com/labs/articles/threat-intelligence/2023-ddos-attack-trends
22. Hesselman, C., Yazdani, R.: DDoS Clearing House for Europe Cross-sector Pilot Demo (Jan 2020), https://www.sidnlabs.nl/downloads/2deJudioEsd0oFWufTXdV9/099fa8c92f7d601e0669bec73b2fa272/NEW-20200123-CONCORDIA-T3.2-demo-review-final.pdf
23. ITSecurityGuru: Proactive vs. Reactive: Which is Better for DDoS Defence? (Jan 2018), https://www.itsecurityguru.org/2018/01/29/proactive-vs-reactive-better-ddos-defence/
24. Jerkins, J.A.: Motivating a market or regulatory solution to IoT insecurity with the Mirai botnet code. In: 2017 IEEE 7th Annual Computing and Communication Workshop and Conference, CCWC 2017. Institute of Electrical and Electronics Engineers Inc. (mar 2017). https://doi.org/10.1109/CCWC.2017.7868464
25. Jonker, M., King, A., Krupp, J., Rossow, C., Sperotto, A., Dainotti, A.: Millions of targets under attack: a macroscopic characterization of the DoS ecosystem. In: Proceedings of the 2017 Internet Measurement Conference. pp. 100–113 (2017)
26. Jonker, M., Sperotto, A., van Rijswijk-Deij, R., Sadre, R., Pras, A.: Measuring the adoption of DDoS protection services. In: Proceedings of the 2016 Internet Measurement Conference. pp. 279–285 (2016)
27. Kambourakis, G., Moschos, T., Geneiatakis, D., Gritzalis, S.: Detecting DNS amplification attacks. In: International workshop on critical information infrastructures security. pp. 185–196. Springer (2007)
28. Karami, M., McCoy, D.: Understanding the emerging threat of ddos-as-a-service. In: 6th {USENIX} Workshop on Large-Scale Exploits and Emergent Threats ({LEET} 13) (2013)
29. Karami, M., Park, Y., McCoy, D.: Stress testing the booters: Understanding and undermining the business of DDoS services. In: Proceedings of the 25th International Conference on World Wide Web. pp. 1033–1043 (2016)
30. Kolias, C., Kambourakis, G., Stavrou, A., Voas, J.: DDoS in the IoT: Mirai and other botnets. Computer **50**(7), 80–84 (2017). https://doi.org/10.1109/MC.2017.201

31. Kopp, D., Dietzel, C., Hohlfeld, O.: DDoS Never Dies? An IXP Perspective on DDoS Amplification Attacks. PAM (2021). https://doi.org/10.1007/978-3-030-72582-2_17
32. Krämer, L., Krupp, J., Makita, D., Nishizoe, T., Koide, T., Koide, T., Yoshioka, K., Rossow, C.: AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. RAID (2015). https://doi.org/10.1007/978-3-319-26362-5_28
33. Kührer, M., Hupperich, T., Rossow, C., Holz, T.: Exit from hell? Reducing the impact of amplification DDoS attacks. USENIX Security Symposium (2014). https://doi.org/null
34. Lemaître, G., Nogueira, F., Aridas, C.K.: Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning
35. Liu, Y., Wang, H.: Tracking mirai variants. Virus Bulletin pp. 1–18 (2018)
36. Lone, Q., Frik, A., Luckie, M., Korczyński, M., van Eeten, M., Ganán, C.: Deployment of source address validation by network operators: a randomized control trial. In: 2022 IEEE Symposium on Security and Privacy (SP). pp. 2361–2378. IEEE (2022)
37. Moore, T., Clayton, R., Anderson, R.: The economics of online crime. Journal of Economic Perspectives **23**(3), 3–20 (2009)
38. Musotto, R., Wall, D.S.: More Amazon than Mafia: analysing a DDoS stresser service as organised cybercrime. Trends in Organized Crime (2020). https://doi.org/10.1007/s12117-020-09397-5
39. NetScout: Botnets Multiply and Level Up (2022), https://www.netscout.com/threatreport/botnets-multiply-and-level-up/
40. Nokia: Nokia Deepfield Network Intelligence Report DDoS in 2021 (Feb 2022), https://onestore.nokia.com/asset/211059?_ga=2.140826161.227459188.1657444403-1091153497.1656679580
41. Noroozian, A., Korczynski, M., Ganan, C., Makita, D., Yoshioka, K., Van Eeten, M.: Who Gets the Boot? Analyzing Victimization by DDoS-as-a-Service. RAID (2016). https://doi.org/10.1007/978-3-319-45719-2_17
42. Noroozian, A., Rodriguez, E.T., Lastdrager, E., Kasama, T., Van Eeten, M., Gañán, C.H.: Can ISPs help mitigate IoT malware? A longitudinal study of broadband ISP security efforts. In: 2021 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 337–352. IEEE (2021)
43. Osborne, C.: NCA launches UK ad campaign to divert kids searching for cybercrime tools (May 2020), https://www.zdnet.com/article/nca-launches-ad-campaign-to-divert-kids-searching-for-cybercrime-tools/
44. Rincón, S.R., Vaton, S., Beugnard, A., Garlatti, S.: Semantics based analysis of botnet activity from heterogeneous data sources. In: 2015 International Wireless Communications and Mobile Computing Conference (IWCMC). pp. 391–396. IEEE (2015)
45. Rodríguez, E., Noroozian, A., van Eeten, M., Gañán, C.: Superspreaders: Quantifying the Role of IoT Manufacturers in Device Infections. In: Workshop on the Economics of Information Security (WEIS) (2021)
46. Sansone, I.: The Damaging Impacts of DDoS Attacks (July 2021), https://www.corero.com/the-damaging-impacts-of-ddos-attacks/
47. Tajalizadehkhoob, S., Korczyński, M., Noroozian, A., Gañán, C., Van Eeten, M.: Apples, oranges and hosting providers: Heterogeneity and security in the hosting market. In: 2016 IEEE/IFIP Network Operations and Management Symposium, NOMS 2016. pp. 289–297. Institute of Electrical and Electronics Engineers (IEEE) (2016)

48. Tajalizadehkhoob, S., Van Goethem, T., Korczyński, M., Noroozian, A., Böhme, R., Moore, T., Joosen, W., van Eeten, M.: Herding vulnerable cats: a statistical approach to disentangle joint responsibility for web security in shared hosting. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. pp. 553–567 (2017)
49. Tanabe, R., Tamai, T., Fujita, A., Isawa, R., Yoshioka, K., Matsumoto, T., Gañán, C., van Eeten, M.: Disposable Botnets: Examining the Anatomy of IoT Botnet Infrastructure. In: Proceedings of the 15th International Conference on Availability, Reliability and Security. ARES '20, Association for Computing Machinery, New York, NY, USA (2020)
50. Toh, A.: Azure DDoS Protection—2020 year in review (Feb 2021), https://azure.microsoft.com/en-us/blog/azure-ddos-protection-2020-year-in-review/
51. Wagner, D., Kopp, D., Wichtlhuber, M., Dietzel, C., Hohlfeld, O., Smaragdakis, G., Feldmann, A.: United We Stand: Collaborative Detection and Mitigation of Amplification DDoS Attacks at Scale. CCS (2021). https://doi.org/10.1145/3460120.3485385
52. Warburton, D.: DDoS Attack Trends for 2020 (May 2021), https://www.f5.com/labs/articles/threat-intelligence/ddos-attack-trends-for-2020
53. Yar, M.: The Novelty of 'Cybercrime' An Assessment in Light of Routine Activity Theory. European Journal of Criminology **2**(4), 407–427 (2005)
54. Ziv, M., Izhikevich, L., Ruth, K., Izhikevich, K., Durumeric, Z.: Asdb: A system for classifying owners of autonomous systems