



Delft University of Technology

Quantifying the Spectrum of Denial-of-Service Attacks through Internet Backscatter

Blenn, Norbert; Ghiëtte, Vincent; Doerr, Christian

DOI

[10.1145/3098954.3098985](https://doi.org/10.1145/3098954.3098985)

Publication date

2017

Document Version

Final published version

Published in

ARES 2017

Citation (APA)

Blenn, N., Ghiëtte, V., & Doerr, C. (2017). Quantifying the Spectrum of Denial-of-Service Attacks through Internet Backscatter. In *ARES 2017: Proceedings of the 12th International Conference on Availability, Reliability and Security* (pp. 1-10). Article 21 Association for Computing Machinery (ACM).
<https://doi.org/10.1145/3098954.3098985>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Quantifying the Spectrum of Denial-of-Service Attacks through Internet Backscatter

Norbert Blenn, Vincent Ghiëtto and Christian Doerr

Delft University of Technology, Cyber Security Group

Mekelweg 4

2628CD Delft

n.blenn@tudelft.nl, v.d.h.ghietto@student.tudelft.nl, c.doerr@tudelft.nl

ABSTRACT

Denial of Service (DoS) attacks are a major threat currently observable in computer networks and especially the Internet. In such an attack a malicious party tries to either break a service, running on a server, or exhaust the capacity or bandwidth of the victim to hinder customers to effectively use the service. Recent reports show that the total number of Distributed Denial of Service (DDoS) attacks is steadily growing with “mega-attacks” peaking at hundreds of gigabit/s (Gbps).

In this paper, we will provide a quantification of DDoS attacks in size and duration beyond these outliers reported in the media. We find that these mega attacks do exist, but the bulk of attacks is in practice only a fraction of these frequently reported values. We further show that it is feasible to collect meaningful backscatter traces using surprisingly small telescopes, thereby enabling a broader audience to perform attack intelligence research.

KEYWORDS

denial-of-service, backscatter, telescope, threat intelligence

ACM Reference format:

Norbert Blenn, Vincent Ghiëtto and Christian Doerr. 2017. Quantifying the Spectrum of Denial-of-Service Attacks through Internet Backscatter. In *Proceedings of ARES '17, Reggio Calabria, Italy, August 29-September 01, 2017*, 10 pages.

<https://doi.org/10.1145/3098954.3098985>

1 INTRODUCTION

Denial of Service (DoS) attacks are a major threat currently observable in computer networks and especially the Internet. In such an attack a malicious party tries to either break a service, running on a server, or exhaust the capacity or bandwidth of the victim to hinder customers to effectively use the service. Current research [2] shows that the total number of Distributed Denial of Service (DDoS) is steadily growing with “mega-attacks” peaking at several hundred of gigabits (Gbps), with recent attacks from the Mirai botnet exceeding volumes of 1 Terabit per second.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ARES '17, August 29-September 01, 2017, Reggio Calabria, Italy

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5257-4/17/08...\$15.00

<https://doi.org/10.1145/3098954.3098985>

The bulk of these reports have come from security companies, network operators, and vendors of mitigation solutions, which raises the question about the representativeness of these examples within the entire spectrum of denial-of-service attacks today. While there is an obvious media-selection bias – only attacks citing massive numbers and which exceed the status quo are “news-worthy” –, these vendors have an implicit interest to identify DDoS as a significant problem. Finally, from the perspective of a network operator which will see traffic volume in the hundreds of gigabit and Terabit range across its peering points, PoP and IXP links, small denial-of-service attacks might not make a large enough difference to be systematically recorded and find their way into statistics.

Two types of attacks are prevalent in the current Internet. In one, an attacker sends a huge amount of requests to a server in order to exhaust processing capabilities, as the server would spawn a new process to execute each request. This way, certain capabilities of a server might get exhausted in terms of CPU, memory usage or the total amount of concurrent connections. In a second type, an attacker would actually target a server’s connection by sending an amount of packets that saturate the victims’ connection. This way, legitimate traffic is not able to reach the server anymore and will be dropped eventually.

As an attacker has an intrinsic motivation not to reveal its own IP address, in order to stay anonymous the attacker can send the attack traffic with (randomly) spoofed source IP addresses. A server that receives such a message will, as long as it is possible to process requests, send an answer to that spoofed IP address as shown in figure 1. When observing a fraction of the IPv4, this so-called backscatter – which is delivered to a large chunk of unused IP addresses, a so-called network telescope – allows us to identify and quantify ongoing (D)DoS attacks in real-time, as well as estimate their size and impact.

Using backscatter data from a /15-sized network telescope that was operated over several years, we address the above described gap, and derive an empirical quantification of the spectrum of denial-of-service attacks in today’s Internet. While backscatter can provide important insights into the type and magnitude of ongoing attacks, publicly available traces from for example CAIDA – onto which still a sizable chunk of backscatter attack research is being performed –, are by now 8 years old. During this time, the Internet as a whole, but also the methods used by attackers and defenders has evolved significantly, thus making new collections and analysis of today’s backscatter necessary. Our paper makes four major contributions:




Figure 1: An attacker (A) sends packets (requests) to the victim (B), spoofing random IPs (C and D). The server of the victim (B) will reply to these requests, as long as possible, sending answers (backscatter) to the spoofed IPs.

- We analytically show that it is possible to obtain statistically significant backscatter traces using comparatively small network telescopes. This makes research on Internet attacks possible for a wider audience.
- We provide a quantification of today’s DDoS attacks, and find that while very large attacks exist they are the absolute exception. Most attacks are very small and short in nature, and UDP – which has been largely ignored in existing backscatter work – has a significant attack surface in practice.
- We demonstrate that the effect of attacks on services and service outage times may be measured from ICMP errors in the backscatter.
- We introduce a method to test the random spoofing of source IP addresses by the attacker, who aims to avoid attribution.

This paper is organized as follows. Section 2 describes the related work and existing findings. Section 3 investigates the size of network telescopes required to obtain a representative insight into attacks on the Internet using backscatter. Section 4 describes our data collection process and introduces the used data set as well as the strategies used to identify backscatter. Section 5 presents results on attack sizes and duration, targeted services, outage durations, as well as the test for random source spoofing. Section 6 concludes our work.

2 RELATED WORK

The concept of backscatter was probably first used in 2006 by Moore *et al.* [8] who introduced the technique of backscatter analysis to provide an estimate of worldwide denial-of-service activity. The authors investigated flooding attacks, in which the victim’s CPU, memory or network resources become exhausted by sending a high number of requests to the victim’s server. The most popular tools, used for distributed attacking, were found to select source addresses at random. In order to classify attacks, Moore *et al.* created “flows”

based on packets arriving from individual IPs in which all packets with an inter-arrival time of less than 5 minutes were combined. Additionally, all flows with less than 100 packets, a flow duration of less than 60 seconds and flows in which data was only sent to one destination IP, were removed from the data set in order to address misconfigurations and other effects, causing packets to reach their monitored address range which had the size of $/8$ (2^{24} distinct IPs).

In 2006, Mao *et al.* [7] compared backscatter analysis from a mainly unused /8 network with direct flow-level measurements executed at a tier-1 ISP. The findings suggest that random address spoofing is only used to a limited extent, implying that backscatter based measurement techniques may not be sufficient. Within the tier-1 ISP, most denial-of-service attacks were probably executed from botnets which did not make use of IP spoofing, as the actual attacker is hidden behind the command and control structure of the botnet. In order to compare the results, the same method of creating flows, given by Moore *et al.* was implemented. An additional constraint was that flows, originated from the direct measurement within the tier-1 ISP, had to consist of a significant amount of packets of the same type. That means that more than 95% of packets had to be of type UDP and originate from a large number of source IPs to classify a potential UDP flooding attack. Likewise, to identify ICMP flooding attacks, more than 95% of packets in a flow had to be of type ICMP and for TCP attacks, more than 90% of all traffic had to be of type TCP having only a single flag set. The results showed, that a high number (90%) of all targets were small businesses or end-users. It is also shown, that most attacks last for less than an hour and that packet rates can potentially be as high as tens of thousands per second with some even reaching one million packets per second.

In 2010, Wustrow *et al.* [11] aimed to analyze the “Internet background radiation” by monitoring network data arriving at multiple /8 networks. In their work the influence of scanning activities and misconfigurations was, in contrast to the earlier mentioned publications, considered additionally to denial-of-service backscatter. Wustrow *et al.* define all TCP packets where only the SYN flag is set, as scans, and packets having the SYN+ACK, RST, RST+ACK or ACK flag set as backscatter. The remaining traffic is considered to arise from misconfigurations. In the longitudinal study, in which data from 2006 until 2010 was analyzed, a possible increase in scanning activities was found, where at the same time the amount of traffic attributed to denial-of-service attacks decreased. The main parts of “pollution” in their data set arose from misconfigurations, like invalid announcements of servers (like BitTorrent servers) and programming errors within certain networked devices/routers.

Durumeric *et al.* [4], also report an increase in terms of Internet scanning activities. Especially through the availability of fast network scanning tools like Zmap or Masscan, that are able to scan the whole IPv4 range in only a few minutes. By observing traffic arriving at a large darknet of 5.5 million addresses, most scans are intended to discover services like ssh, http or mysql servers, but numerous scans are also intended towards the discovery of denial-of-service amplification services like NTP, DNS, Chargen or Quote of the day (for a complete list, please see Rossow [10]).

A recent study by Kr amer *et al.* [6] used a /16 network and employed honeypots for amplification attacks called AmpPots. These

honeypots were rate limited in order to limit the abuse potential in actual DDoS attacks by not sending more than 10 packets per minute. Interestingly, by just placing these honeypots at different locations within the IPv4 range, attackers found and tried to employ them rather quickly, a fact that can be attributed to large scanning activities for amplification services. Krämer *et al.* showed that it is possible to passively monitor amplification attacks through placing and operating DoS amplification honeypots.

In our work we combined the knowledge about IPv4 scans in order to estimate the size and ongoing threat arising from denial-of-service attacks.

3 NETWORK TELESCOPE AND THE ESTIMATION OF INTERNET ATTACKS

As discussed in the related work, the bulk of previous work in academia on Internet backscatter analysis is centered around a few datasets such as the CAIDA backscatter dataset from 2008 [1], and is still used by publications appearing today. With the CAIDA measurement turning now 9 years old, we can however expect many traffic characteristics and attacks on the Internet to have significantly changed since then. Several methods such as smurf attacks are no longer widely used, while new types of threats such as NTP amplification attacks have become mainstream. With advances in network technology, shifts in applications and access vectors such as the rise of mobile platforms, introduced host and network defense strategies and evolving attack methods, and finally also questions that have been raised about the integrity of the CAIDA dataset [3], it might be time to resample today's backscatter to understand the evolution over the past decade.

One of the reasons for the popularity for the CAIDA dataset is the easy availability of the data and the significant size of the network telescope with which they were collected. This however naturally poses the question that if new datasets were to be collected today, how large of a telescope would be needed to arrive at a reliable snapshot, from which estimations of attack types and sizes on today's Internet can be made and statistically significant conclusions be drawn. This evaluation is the focus of this section.

Consider a network telescope consisting of k IP addresses out of the total 2^{32} IPv4 address space. Any IP connected to the Internet will normally receive three types of traffic as shown in figure 2: traffic in response to for example a previous outgoing user request, network scans from parties in the Internet probing the host for any open ports, and backscatter from attacks that arrives because the attacker has spoofed the source IP of the attack packets to complicate attribution and attack mitigation. As the telescope would normally be built using dark IPs, no user traffic would be seen in such a deployment. Network scans can be eliminated from the incoming traffic using the techniques discussed in section 4, leaving only the echo from attacks.

In non-amplification attacks, the adversary will typically spoof the IP address of the request packet sent to the attacked host or network as shown in figure 1. This will disguise the actual origin of the traffic to the victim and also render mitigation techniques such as IP filtering useless. The adversary may follow two options, generate the IP addresses randomly or – less commonly – follow an algorithm such as a generator in multiplicative group modulo p .




Figure 2: Incoming data on any IP address will be a mix of requested user-related data traffic, scans performed on the local host or network as well as backscatter from attacks performed against remote parties.

The former approach equals a drawing with replacement and the repeated Bernoulli trial over the course of an attack will lead to a binomial distribution of packets observed across IPs at the observatory, while the latter one will show up as a uniform distribution. We identify the IP spoofing algorithm using a statistical test for each incoming attack as discussed below.

Suppose a number of o packets was received at the k observed IP addresses over a given timespan. This will mean that the overall size of the attack directed against the third party μ_O may be approximated as

$$\mu_O = o \cdot \frac{2^{32}}{k},$$

which is the estimated average number of packets generated in the attack. In the typical case of masking by IP address randomization, a smaller network telescope would however be more sensitive to noise and have a higher likelihood of over- or underestimating the total size of the attack. Assuming a 5% error margin for an erroneous attack size estimation, the bounds of the potential estimation error are obtained by finding the binomial distribution with mean μ_U for which the probability to obtain μ_O or fewer packets is 2.5%. Together with the similarly obtained lower distribution with mean μ_L a confidence interval for the attack size can be obtained as shown in figure 3.




Figure 3: Estimation of the lower and upper bound of the average number of attack packets.

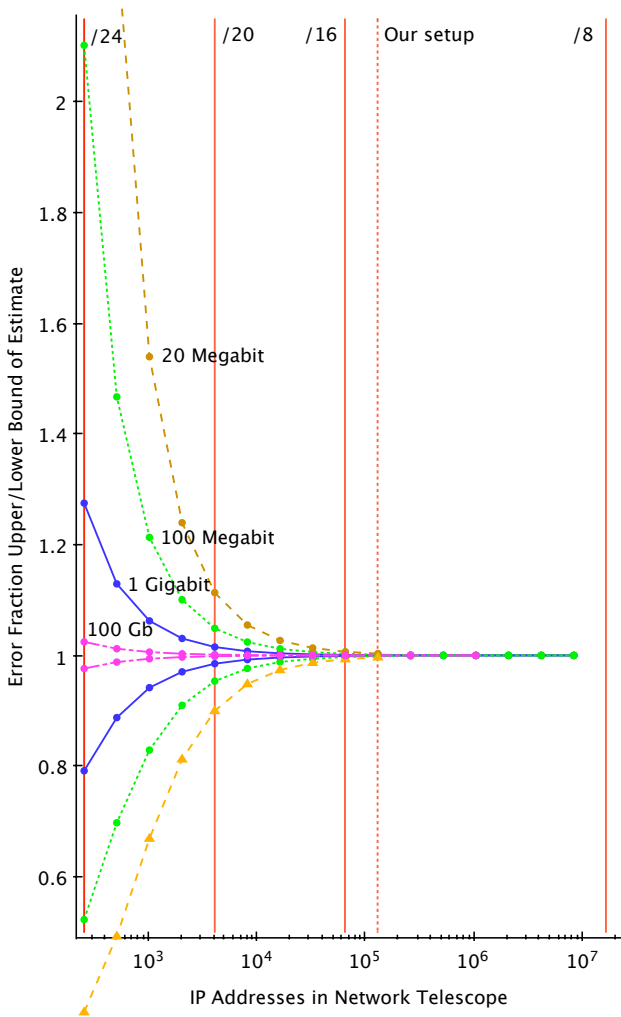


Figure 4: Estimation error of attack sizes based on backscatter as a function of telescope size.

Figure 4 shows the attack size over- and underestimation as a fraction of the actual size as a function of the number of packets sent in the attack and the IP addresses available in the telescope available for observation. As can be seen in the graph for a packet flood with the equivalent of 100 Mbps (assuming minimal packet size), a network telescope across 512 dark IP addresses would place an estimate of the attack somewhere between 69 to 146 Mbps. This error margin decreases with the increase in observed IP addresses, so that at 2^{10} available IPs the bounds have shrunk to 82 to 121 Mbps. From the figure a number of insights may be drawn:

- Research on large attack sizes beyond 10 Gigabit / s can be effectively facilitated even through very small installations. With a /24 subnet the error margin for medium-size attacks is below 10%, making data collection and analysis of backscatter possible for most research groups.

- Size does matter – especially for smaller attack volumes. At 128 IP addresses and a 20 Mbit/s attack volume, the estimation from the telescope may be almost half an order of magnitude larger than the actual attack volume. Since we find that 50% of backscattering attacks are smaller than 20 Mbps, significantly more IP addresses are needed to facilitate analysis on the entire spectrum of protocol-based and volumetric attacks. A /19 network or 8192 IP addresses is the smallest telescope size that brings the error margin for small scale attacks below 10%.
- Network telescopes beyond a /17 network are very insensitive to sampling biases. With 32768 observed IPs error margins below 2% can be realized even for 20 Mbps-sized attacks.

The statistical analysis points out that it is actually possible to a wide(r) audience to setup their own network telescope and obtain backscatter observations that lead to statistically significant and meaningful insights. With this validation in place, the remainder of this paper will discuss analysis of attacks and classification of data obtained through two /16 telescope blocks.

4 DATA COLLECTION

The data used in this paper arises from a darknet of the TU Delft in which all non-routable packets directed to two nearly unused /16 networks over a period of 26 months is collected. On a typical day approximately 15 GB of scan traffic and backscatter is received, yielding a corpus of more than 7 TB. The analyses in this paper are based on a subset of this trace between March 25, 2015 and June 20, 2015. This evaluation dataset of 1.3 TB contains approximately 3.532 billion received packets. While previous work on only Internet backscatter has largely excluded UDP traffic and [3] remarked the absence of UDP in some public traces, we do find that a non-trivial amount actually uses UDP - in the above trace from 2015, 74.58% of frames are TCP, 22.37% UDP.

As the packets arriving in this sinkhole are sent to IPv4 IPs behind which no hosts exist, the data arriving in the sinkhole therefore arises from either (a) IPv4 scans, (b) IP misconfigurations, or (c) backscatter. This in turn also eliminates any user privacy concerns in the data collection for our research, as the IP addresses are not assigned to hosts and therefore no user communication could be accidentally recorded. In order to identify backscatter from scans or misconfigurations, we dissected traffic on a protocol basis, and if necessary also parsed the TCP/UDP/ICMP payload for select application protocols.

4.1 Identifying TCP backscatter

Unlike other protocols, TCP contains header information necessary for session management and that signals the state of the ongoing connection. Initial packets sent from a client to a server contain a SYN packet, to which the server responds with a SYN+ACK, triggering the client to complete the handshake with an ACK flag. We refer to [5] for a description of the TCP finite state machine and other flags triggering transitions. (We note that although some tools and tutorials extensively describe exotic flag combinations in scans and attacks, such as Xmas or NULL packets, we find little evidence of these happening in practice, with all flags other than

SYN/ACK/RST combined accounting for only a fraction of 1% of all observed packets.)

Based on this background of the TCP protocol, and following established practices, it is thus possible to separate scan from backscatter through the SYN+ACK and RST flag only. Network scans or attack traffic from clients to a server will feature the SYN flag (88% of packet in our dataset) – a server will never send SYNs –, backscatter from attacks must thus contain a SYN+ACK (9%) in response at an open port and RST (1%) for a closed port.

4.2 Characterizing attacks through ICMP backscatter

A service protocol of IP, ICMP serves many network-related functions. ICMP packets signal the status of networks and hosts to remote locations, can be used for diagnosing purposes (“ping”) and also carry back error messages when delivery failures occurred. There are a total of 48 ICMP messages in use today, which are categorized by ICMP type, each further subdivided into message codes. The type of ICMP identifies the reason for the packet being sent, for example type 8 being the well-known ping request (which is replied to by a type 0), type 11 being TTL exceeded and type 3 destination unreachable notices. An ICMP type 3 message such as destination unreachable is further split up into the exact cause of non-delivery, such as the network being unreachable (type 3 - code 0) or the host being administratively prohibited from communicating with this source (type 3 - code 9).

While 70% of ICMP packets are echo requests and thus remote network scans, the remaining 30% are comprised of destination unreachable replies, time exceeded, and replies to ping requests. Table 1 breaks down the type of ICMP packets recorded among the 3.5 billion received packets, table 2 breaks type 3 packets – which are the most relevant for characterizing denial-of-service attacks – further down in the concrete status codes. As attack packets at some point no longer reach the victim host, these ICMP packets are returned as backscatter to our telescope. We find that DDoS attempts by ping are a rarity, but also notice that TTL exceeded packets – which can in theory be used to saturate and thus execute a DDoS on routers – occur nearly 6 times more often.

Three quarters of ICMP backscatter are type 3 destination unreachable notices, which are important as they carry additional response codes. To the best of our knowledge, ICMP backscatter response codes have not systematically been mined before, even though in combination with other backscatter data allow a surprisingly detailed peek at the mechanics of individual attacks. As we will show exemplarily in section 5, type 3 packets may be used to identify and characterize attacks.

Table 1: Breakdown of ICMP Packet Types

ICMP Type	Percent Packets
Echo	70%
Destination unreachable	24%
Time exceeded	6%
Echo reply	1%
Fragmentation needed and DF was set	< 1%

Table 2: Breakdown of ICMP Type 3 Packets (Destination Unreachable)

ICMP Type	Percent Packets
Port unreachable	89%
Communication with host is prohibited	4%
Protocol unreachable	3%
Net unreachable	2%
Host unreachable	1%
Others	each < 1%

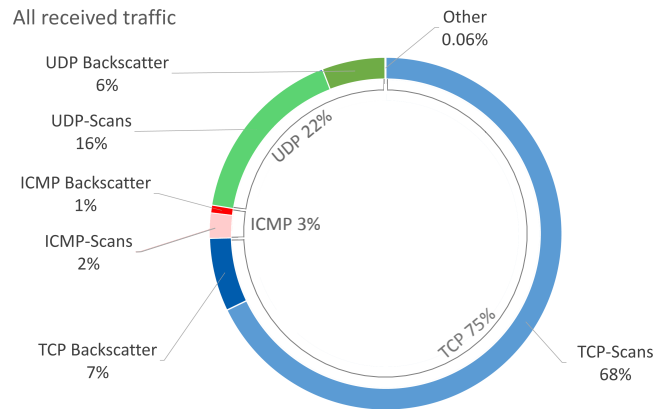


Figure 5: Fraction of packets identified as backscatter received in our network monitor over all received packets.

4.3 Identifying UDP backscatter

Another type of packet that has unjustifiably received little attention in network backscatter analysis today is UDP. UDP is relevant, as first about 22% of all traffic is actually UDP, and second UDP packets typically carry other types of traffic such as DNS or NTP, which are those that are predominantly abused for their high amplification ratios (together with UDP’s handshake-less nature) in today’s denial-of-service attacks.

One of the reasons we believe why there is almost no work on UDP backscatter, is the notorious difficulty to classify UDP packets as backscatter. As the protocol is stateless and only has a minimalistic header, it cannot be judged from the header fields alone whether a packet is actually a request or a response, other than in TCP where the presence of a SYN and/or an ACK clearly differentiates the two at the transport layer protocol level. Understanding UDP thus requires an understanding of the UDP packet, as more than a fifth of frames should not be ignored, we are classifying UDP into requests and responses based on a two-step process. In a first step we filter for all packets received from standardized UDP source ports as listed in RFC-1340 [9]. Additionally we added and included the 20 most frequent UDP source ports in the dataset into the analysis. In a second step, protocol parsers were written for key application protocols that analyze the content of the UDP payload to determine the nature of the packet. An example of this is DNS, where if the header passes a consistency check and bit 17 is set to 1 the packet therefore constitutes a stateless reply.

Using the methods discussed above, approximately 9% of all TCP traffic, 30% of ICMP and 27% of UDP could be identified as backscatter as shown in figure 5.

5 AN ANALYSIS OF ATTACKS

Based on the identification of backscatter as introduced in the last section, this part will report on the types and sizes of attacks we have observed through backscatter analysis. For the duration of approximately three months, a total of 4.688 million attacks were identified, targeting 2.21 million IP addresses. DDoS attacks were targeted against 35% of all autonomous systems in the Internet, affecting nearly all countries in the world.

In this section, we address four main questions on the characteristics of current attacks in the Internet: First, given the media reports on extremely large attack volumes, we quantify the entire spectrum of ongoing attack sizes. Second, we investigate which ports and services are commonly targeted. Third, we demonstrate that it is possible to quantify downtimes of services based on backscatter and estimate the amount of floods a particular service can sustain before collapsing under the load. And fourth, we introduce the notion of statistical testing on the source IP addresses spoofed in an attack to show that an attack has randomly selected source IPs to hide its tracks.

5.1 Attack Characteristics

As reported in the introduction, record-breaking attack volumes are reported by ISPs and service providers specializing in attack mitigation on a regular basis and circulated in the media. While it is evident that such reports are subject to a selection bias, this however naturally triggers the question how much of an outlier DDoS attacks of several hundred Gbps actually are.

For all 4.688 million attacks observed between March and June of 2015, we computed the likely total attack size as well as their confidence intervals as discussed in section 3. Figure 6 shows a plot of the cumulative density function for the estimated total number of packets per second (pps) in the attack. While there do exist a number of large and very large attack volumes exceeding 10^9 packets per second, the majority of DoS attacks only encompasses around 10^4 pps.

To further put this into context, let’s assume that the attacker aims to maximize the number of packets sent in an attack, e.g., to overwhelm the host at the protocol level. With a minimum packet size of 52 bytes for UDP and 64 bytes for TCP on the wire, this means approximately 90% of all observed DDoS attacks had an average traffic volume of less than 10 Mbps. Given previously reported large scale attack sizes, this finding seems surprising at first, but straight forward on second thought.

In the spectrum of attacks on services we need to differentiate between protocol-based and volume-based attacks as shown in figure 2, i.e., those that intend to exceed a block of finite resources on the host – in case of TCP typically the size of the transmission control block (TCB) – or saturate the uplink of the host or network providing the service. It is also logical to assume that an attacker will be using attack resources conservatively, i.e., only utilize as much as necessary (with a safety margin) in order to be successful. For a protocol-based attack such as TCP SYN floods trying to exhaust

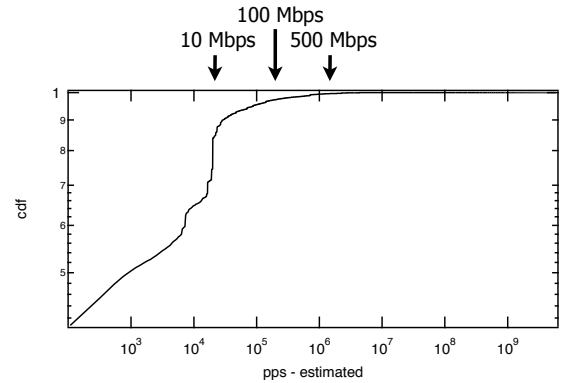


Figure 6: Cumulative density function of attack sizes.




Figure 7: Cumulative density function of attack duration.

a host’s available connection pool, 10^4 attempted connections per second is likely sufficient to impair most service-providing hosts on the Internet, unless employing defensive means such as TCP SYN cookies. Only in the upper 5% of attacks shown in figure 6 do we see the second regime of volume-based attacks that aim to saturate the connection itself, with an estimate 100 Mbps volume or larger applying to the top 5 % and 1 Gbps or larger flood to the top 1%. As the user base of Internet services follows a power law distribution, the vast majority of DoS attacks actually goes unnoticed to the bulk of users, and those targeting the top 0.01% which would be noticed widely also require a significantly-sized flood.

When looking at the duration of DDoS attacks as shown in figure 7, a fundamentally different picture emerges. The duration of DDoS attacks exhibits a long tail, with 10% continuing for at least one day and approximately 2% exceeding one month. Note also the sharp peak of DDoS attack durations at 20 seconds, and a smaller one at 60 seconds. These intervals match typical attack durations provided by DDoS services, which sometimes also offer free attack demonstrations at these durations (showing increasing professionalization with 24 hour customer service and money back guarantees).

Figure 8 combines the data from figures 6 and 7 in a surface plot, showing the number of observed attacks as a function of the duration of the attack and the number of packets received in the

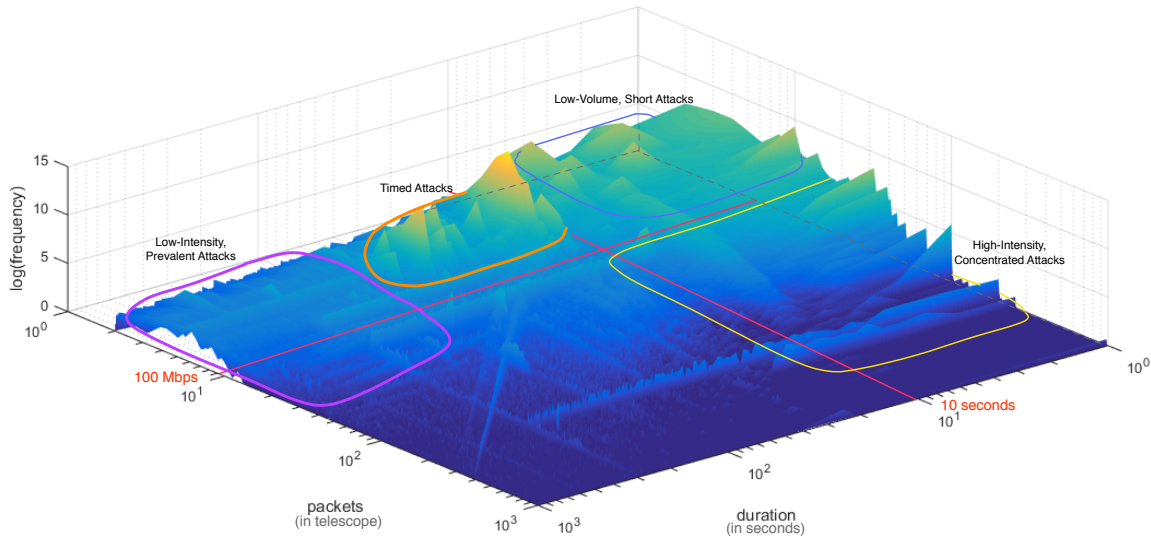


Figure 8: Number of observed attack instances as a function of attack volume and duration.

telescope (without the scaling towards the estimated total attack size in figure 6) all in logarithmic scale. Two lines indicating 100 Mbps attack volume and 10 second attack duration are added as guidelines. The spectrum of attacks can be classified in four different regimes: (1) In the center of the graph and with the highest peak (indicated in orange), we find attacks of fixed duration such as the distinctive 20 seconds, 60 seconds already indicated in the previous figure, but at a logarithmic scale we can also locate smaller peaks around 30, 40 and 90 seconds. It is noteworthy that although being concentrated in time, these attacks exhibit a significant variation in volume per campaign. (2) Only a small fraction of attacks becomes prevalent, which can be found in the left part of the figure, indicated in violet. We find that if an attack is launched continuously at a target, the attack intensity is low, with the mode around 50 Mbps. Long duration attacks beyond 200 Mbps are virtually absent. (3) High volume attacks do exist however, as shown in the right part of the figure in yellow. A significant share of these high intensity attacks routinely exceed Gigabit speed, but are only very short in nature - they typically do not last longer than 10 seconds. While running at a high speed, they are less visible in an aggregate plot such as the one shown in figure 6. (4) About a third of all campaigns are both short and low in intensity, as indicated in blue at the back of the figure. Here we find the bulk of activity at intensity levels below 5 Mbps.

The fact that we can identify a number of different regimes raises the question of the types of attackers behind these DDoS campaigns and their underlying motivation, which to this date is not well researched and understood yet. Differences between the “low-volume, short attack” and “high-intensity, concentrated attack” classes can be explained by the resources available to the initiating party. Even when working at high amplification factors, a flood at the Gigabit/second scale requires a sizeable uplink to inject packets

from, more than the typical residential customer would have. This is even more relevant in case of a generic TCP SYN flood, where the request/response traffic ratio is essentially 1. Thus, the difference between these two regimes would indicate a home-brewn operation in contrast to a professionally run and hosted/distributed service. The presence of such services is also hinted at from the concise attack durations, matching common “order” intervals of DDoS services. Noteworthy is however the absence of high-intensity long-lived attacks in the longitudinal analysis. If we take the assumption that the perpetrators are economically-thinking actors who choose and dimension attacks to fulfill their objective (as always going “all in” will waste capacity that could be sold otherwise, reveals their maximum DDoS capabilities, and in case of a botnet also makes detection of the infected end hosts more likely), we might speculate that there is simply no significant market for these attacks. On the one hand, the attacker’s capability to take down a service or organization can be proven by a short peak in traffic to convince the victim to pay up, on the other hand, as we will show in section 5.3, low attack volumes are typically sufficient to trip a targeted service and thus also achieve the objective of the attacker.

5.2 Targeted Services

Figure 9 shows a frequency distribution of the ports at which attack packets were directed at the victim’s host for both TCP and UDP for the privileged ports 1 through 950. Note that the y-axis is in logarithmic scale - attacks on port 80 by far dominate any other target and almost 34 times more volume is directed at port 80 than the second most frequency targeted TCP port 443.

While previous backscatter studies have largely excluded UDP, we do see significant attack activity on select ports such as port 53, where UDP traffic is about 1.5 orders of magnitude more than TCP-based attacks. Furthermore, 5 out of the 10 most targeted




Figure 9: Histogram depicting the number of packets received from source ports 1 - 950 in our network monitor.

port/protocol combinations involve UDP. Also the recent rise of amplification attacks such as using NTP – the peak at port 123 – is clearly visible in the backscatter. The data also reflects recent trends in DDoS techniques, such as the switch from NTP-based to SSDP-based attacks, as there are simply many more commodity routers that have implemented the Simple Service Discovery Protocol and are not easily upgraded as a population.

Aside from well known and rising vectors, we note that a significant chunk of attacks is directed against gaming servers. Within TCP, Minecraft is the third most targeted service, with approximately 42% of the attack volume directed against HTTPs and in fact more attack packets than sent against DNS on TCP.

5.3 Quantifying “Fallover” Levels

When a TCP port receives a SYN packet, it will respond with a SYN+ACK if open, and a RST if closed. In case the operating system is unable to process the request (for example due to an overload from a SYN flood), an ICMP destination unreachable frame will be generated unless this feature is suppressed. If such an ICMP message of type 3 is sent, the original packet sent to the server is typically conveniently included in the response, which allows us to inspect the concept of the original attack packets.

Figure 10 shows an example of a successful DDoS attack on a service which finally collapses and shows an initial 11 second outage. Within this outage time, the service partially recovers but with the still ongoing attack will from now on be in a period of intermittent service availability. The screenshot of the ICMP type 3 packet shows the original SYN frame sent to port 80 on the affected IP address.

Aside from enabling a quantification of attack types – such as the Kaminsky attack on DNS, which is beyond the scope of this paper – this observable switch from SYN+ACKs to ICMP type 3 in the backscatter data allows an estimation of the size of packet flood the service was able to cope with before collapsing under the load. Figure 11 shows the distribution of estimated DDoS sizes in packets per second (pps) under which remote hosts collapsed. The distribution is long-tailed with resilient (but not resilient enough) services that sustain floods well beyond 100 Mbps. A very high

number of services is however not well provisioned and become unavailable well below 1 Mbps, connected with (self-hosted) game servers such as the Minecraft hosts mentioned above.

These insights about the relatively low packet flood levels at which a host “falls over” aligns with the earlier finding that in the Internet as a whole DDoS attack sizes are comparatively small, most likely because they don’t need to be larger and as a significant number of services are not that well provisioned as we would expect from our first hand experience with major Internet platforms, service providers and web applications.

Together with the findings presented in section 5.1, these results are somewhat sobering. They clearly demonstrate that the bulk of denial-of-service attacks is actually surprisingly small, 95% of attacks never exceed 100 Mbps, and an attack beyond 1 Gbps is extremely rare when looking at the entire spectrum of attacks. Adversaries have however the goal of bringing down a service as part of a given attack and the data (and experience) shows that botnets and DDoS services have the capability to run attacks in the tens and hundreds of Gigabits. The reason that they do not regularly attack at such volumes probably lies in the fact that they do not have to; as low volumes are sufficient, why go “all-in”? While for example protocol-based attacks such as TCP SYN floods can be easily mitigated by modern operating systems through means such as SYN cookies, figure 11 seem to indicate that in the field a very large number of services do not deploy or enable such countermeasures.

5.4 Source IP Spoofing and Attack Attribution

When performing a DoS, the attacker will try to hide his tracks and make the mitigation for the victim host more difficult by including other source IP addresses in the request packets than his own. For this the adversary has two ways, randomly pick – typically from the entire IPv4 space to maximize the difficulties for the defender –, or enter a set of predefined IP addresses as the source to attribute the attack to some third party. This makes attribution of denial-of-service attacks extremely difficult, since even if the actual hosts that are injecting attack packets have been located through back tracing of attack traffic, these tend to be compromised hosts themselves –

No.	Time	Source	sreport	Destination	dstport	Protocol	Length	Info
231666	2015-07-03 05:09:13.865169000	183.61.182.27	80	130.161.232.187	57984	TCP	60	80-57984 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
232367	2015-07-03 05:09:14.366246000	183.61.182.27	80	145.94.133.5	36737	TCP	60	80-36737 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
235862	2015-07-03 05:09:16.459540000	183.61.182.27	80	131.180.219.12	38060	TCP	60	80-38060 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
236336	2015-07-03 05:09:16.731608000	183.61.182.27	80	130.161.182.205	39208	TCP	60	80-39208 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
241710	2015-07-03 05:09:20.149237000	183.61.182.27	80	130.161.125.137	8858	TCP	60	80-8858 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
244821	2015-07-03 05:09:22.195657000	183.61.182.27	80	145.94.102.106	55611	TCP	60	80-55611 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
252422	2015-07-03 05:09:27.440474000	183.61.182.27	80	145.94.0.207	54844	TCP	60	80-54844 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
254181	2015-07-03 05:09:28.564546000	183.61.182.27	80	130.161.208.162	54253	TCP	60	80-54253 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
264706	2015-07-03 05:09:34.976662000	183.61.182.27	50083	145.94.223.216	80	ICMP	590	Destination unreachable (Port unreachable)
269627	2015-07-03 05:09:38.073324000	183.61.182.27	61341	145.94.132.55	80	ICMP	590	Destination unreachable (Port unreachable)
273127	2015-07-03 05:09:40.222578000	183.61.182.27	24146	145.94.65.55	80	ICMP	590	Destination unreachable (Port unreachable)
274610	2015-07-03 05:09:41.225098000	183.61.182.27	26881	130.161.240.66	80	ICMP	590	Destination unreachable (Port unreachable)
277058	2015-07-03 05:09:42.713824000	183.61.182.27	56591	131.180.254.182	80	ICMP	590	Destination unreachable (Port unreachable)
277086	2015-07-03 05:09:42.734640000	183.61.182.27	718	145.94.108.86	80	ICMP	590	Destination unreachable (Port unreachable)
279579	2015-07-03 05:09:44.234599000	183.61.182.27	10250	130.161.249.1	80	ICMP	590	Destination unreachable (Port unreachable)
280570	2015-07-03 05:09:44.898170000	183.61.182.27	5054	131.180.172.236	80	ICMP	590	Destination unreachable (Port unreachable)
281110	2015-07-03 05:09:45.286588000	183.61.182.27	3610	131.180.8.231	80	ICMP	590	Destination unreachable (Port unreachable)
281729	2015-07-03 05:09:45.749734000	183.61.182.27	80	130.161.199.135	44760	TCP	60	80-44760 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
282730	2015-07-03 05:09:46.399579000	183.61.182.27	80	130.161.242.156	49923	TCP	60	80-49923 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
283696	2015-07-03 05:09:46.837367000	183.61.182.27	80	130.161.199.135	44760	TCP	60	[TCP Retransmission] 80-44760 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
284526	2015-07-03 05:09:47.316086000	183.61.182.27	80	145.94.92.137	45165	TCP	60	80-45165 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
285400	2015-07-03 05:09:47.889853000	183.61.182.27	80	130.161.185.168	52449	TCP	60	80-52449 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
286007	2015-07-03 05:09:48.289702000	183.61.182.27	80	145.94.113.41	20610	TCP	60	80-20610 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
286758	2015-07-03 05:09:48.745811000	183.61.182.27	80	131.180.254.36	19215	TCP	60	80-19215 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
286845	2015-07-03 05:09:48.784206000	183.61.182.27	80	130.161.199.135	44760	TCP	60	[TCP Retransmission] 80-44760 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
287534	2015-07-03 05:09:49.113313000	183.61.182.27	80	131.180.254.36	19215	TCP	60	[TCP Retransmission] 80-19215 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460
288649	2015-07-03 05:09:49.627376000	183.61.182.27	80	145.94.107.103	36476	TCP	60	80-36476 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460

Figure 10: Example of an successful attack on a system with intermittent service availability.




Figure 11: Distribution of estimated DDoS sizes in packets per second (pps) in which the remote service collapsed.

such as bots rented for the attack -, while the actual perpetrator will hide in yet another step in the chain.

A for the operational mitigation of the attack actually far more relevant question centers around the IP addresses that are spoofed in the attack. The most difficult to mitigate attack is one where the adversary is injecting random source addresses in the request packets. If the adversary is spoofing a few, specific IP addresses in an attack, for example those of a competitor, the attack would be easier to mitigate by the recipient of the DDoS traffic, but lay the blame for the attack on someone else - and in this case this alleged specific origin network or organization would be the actual victim.

Whether or not the adversary is randomly sampling IP addresses from the entire Internet or a large chunk of the IP address space and thus attributed the attack to everyone can actually be efficiently tested on a per-attack basis using the incoming backscatter. This "source origin test" is possible as the number of packets would follow a binomial distribution, which given a large enough telescope to obtain statistically significant results can be empirically verified. In case of random origin spoofing, the source addresses of each packet would drawn with equal probability, the probability to observe this packet at a particular IP address belonging to the network telescope equals $p = \frac{1}{2^{32}}$. With a total of n packets sent in the attack the probability of receiving k backscatter packets at any one IP

of the network telescope follows a Binomial distribution $B(n, p)$. For network telescopes with any predictive power (see section 3), we can approximate the binomial distribution $B(n, p)$ through the simpler Gaussian distribution $N(np, np(1 - p))$. If the number of packages received at all IPs belonging to the telescope does indeed follow this distribution, we may conclude that most likely this attack did not attribute the attack specifically to one network but indeed injected packets from random source IP addresses.




Figure 12: Probability of receiving P packets at one IP in our network monitor from one source IP, fitted by a Gaussian distribution.

Figure 12 shows this evaluation for the case of an actual campaign, with the red line indicating the frequency of the total number of packets received by a destination IP in the telescope belonging to this specific attack. Extrapolating for the total sum of backscatter and the size of the telescope, we can obtain an estimation of the attack size. These values match the Gaussian with $\mu = 69.9$ and $\sigma = 8.4$, therefore this particular attack used a random spoofing.

6 CONCLUSION

In this paper we have made the case that since the collection of major publicly available backscatter traces, much has changed in the Internet and it is time to re-collect and re-do such measurements,

leading to new insights about current threats and the developments of techniques over time.

We find that although new records in denial-of-service attack sizes are reported in regular intervals in the media, the vast majority of DDoS attacks are tiny, and could be defended with readily available technology and software. We demonstrate that it is feasible to collect sufficient amounts of data to arrive at statistically sound estimations of ongoing attacks on the Internet using comparatively small network telescopes. We furthermore show that it is possible to mine a broad array of statistics from ongoing attacks using backscatter, such as attack size and duration distributions, statistically check the attribution of the attack, and remotely determine the state of services and under which load a particular host collapses.

REFERENCES

- [1] The ucsd caida backscatter dataset - 2008. [Online]. Available: http://www.caida.org/data/passive/backscatter_dataset.xml
- [2] "Q2 2015 State of the Internet - Security Report," 2015. [Online]. Available: <https://www.stateoftheinternet.com/resources-cloud-security-2015-q2-web-security-report.html>
- [3] E. Balkanli and A. N. Zincir-Heywood, "On the analysis of backscatter traffic," *8th IEEE Workshop on Network Measurements*, 2014.
- [4] Z. Durumeric, M. Bailey, and J. A. Halderman, "An internet-wide view of internet-wide scanning," in *23rd USENIX Security Symposium (USENIX Security 14)*. San Diego, CA: USENIX Association, Aug. 2014, pp. 65-78. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/durumeric>
- [5] Information Sciences Institute, "Transmission control protocol (rfc793)," IETF, Tech. Rep., 1981.
- [6] L. Kr amer, J. Krupp, D. Makita, T. Nishizoe, T. Koide, K. Yoshioka, and C. Rossow, "AmpPot: Monitoring and Defending Amplification DDoS Attacks," in *Proceedings of the 18th International Symposium on Research in Attacks, Intrusions and Defenses*, November 2015.
- [7] Z. M. Mao, V. Sekar, O. Spatscheck, J. van der Merwe, and R. Vasudevan, "Analyzing large ddos attacks using multiple data sources," in *Proceedings of the 2006 SIGCOMM Workshop on Large-scale Attack Defense*, ser. LSAD '06. New York, NY, USA: ACM, 2006, pp. 161-168. [Online]. Available: <http://doi.acm.org/10.1145/1162666.1162675>
- [8] D. Moore, C. Shannon, D. J. Brown, G. M. Voelker, and S. Savage, "Inferring internet denial-of-service activity," *ACM Trans. Comput. Syst.*, vol. 24, no. 2, pp. 115-139, May 2006. [Online]. Available: <http://doi.acm.org/10.1145/1132026.1132027>
- [9] J. Reynolds and J. Postel, "Rfc1340: Assigned numbers," in *IETF*, 1992.
- [10] C. Rossow, "Amplification Hell: Revisiting Network Protocols for DDoS Abuse," in *Proceedings of the 2014 Network and Distributed System Security (NDSS) Symposium*, February 2014.
- [11] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Huston, "Internet background radiation revisited," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '10. New York, NY, USA: ACM, 2010, pp. 62-74. [Online]. Available: <http://doi.acm.org/10.1145/1879141.1879149>