

Demystifying DDoS as a Service

Ali Zand, Gaspar Modelo-Howard, Alok Tongaonkar, Sung-Ju Lee, Christopher Kruegel, and Giovanni Vigna

The authors present a measurement study of 17 different DaaS providers, in which they analyzed the different techniques used to launch DDoS attacks, as well as the infrastructure leveraged in order to carry out the attacks. Results show a growing market of short-lived providers, where DDoS attacks are available at low cost (tens of dollars) and capable of easily disrupting connections of over 1.4 Gb/s.

ABSTRACT

In recent years, we have observed a resurgence of DDoS attacks. These attacks often exploit vulnerable servers (e.g., DNS and NTP) to produce large amounts of traffic with little effort. However, we have also observed the appearance of application-level DDoS attacks, which leverage corner cases in the logic of an application in order to severely reduce the availability of the provided service. In both cases, these attacks are used to extort a ransom, to hurt a target organization, or to gain some tactical advantage. As it has happened for many of the components in the underground economy, DDoS has been commoditized, and DDoS as a service (DaaS) providers allow paying customers to buy and direct attacks against specific targets. In this article, we present a measurement study of 17 different DaaS providers, in which we analyzed the different techniques used to launch DDoS attacks, as well as the infrastructure leveraged in order to carry out the attacks. Results show a growing market of short-lived providers, where DDoS attacks are available at low cost (tens of dollars) and capable of easily disrupting connections of over 1.4 Gb/s. In our study, particular attention was given to characterize application-level (HTTP) DDoS attacks, which are more difficult to study given the low volume of traffic they generate and the need to study the logic of the application providing the target service.

INTRODUCTION

Distributed denial of service (DDoS) attacks have been a problem on the Internet for more than 15 years. However, the recent increase in the number of DDoS attacks and in the amount of traffic that they generate has attracted the attention of the media, the industry, and the research community alike. This new wave of attacks exploit asymmetries in vulnerable services to generate large amounts of traffic or use large amounts of resources with relatively little effort from the attacker. For example, misconfigured Network Time Protocol (NTP) services can be leveraged to generate gigabytes of data with a simple spoofed request. This generated traffic exhausts the bandwidth available at the target. We call this type of (more traditional) attack an *extensive DDoS*.

However, there is another type of DDoS attack in which the lack of availability of a

resource is due to the fact that a single interaction with the target requires an unusually high amount of resources in order to be processed. For example, on a web site, there might be a search form that, when provided with certain values, might require an extremely large database query that slows the whole website to a crawl. We call this kind of attack an asymmetric application-level or *intensive DDoS*.

While extensive DDoS attacks have been studied for quite a while [1] and some remediation has been provided (e.g., coordinated filtering managed by blacklists, rate limiting, patching of vulnerable services), intensive DDoS attacks have not received the same level of attention. The latter is more difficult to characterize because they often depend on the logic of the application providing the target service. In addition, these attacks do not rely on large volumes of data and therefore can go undetected by volumetric detection mechanisms. Finally, since the attacker communicates with the service following the service protocol, the attacker's requests are similar to a legitimate request and hence more difficult to filter out.

As both extensive and intensive DDoS attacks become an integral part of the efforts of cybercriminals to obtain financial gains (e.g., by blackmailing organizations under attack or by obtaining a tactical advantage in time-sensitive settings), the provision of DDoS service has become commoditized. We now see the rise of DDoS as a service (DaaS) offerings, in which DDoS providers attack a target in exchange for money.

BACKGROUND

In this section we introduce the different types of DDoS attacks available, as well as the basic infrastructure of the DaaS providers, which are the subject of our study.

TYPES OF DDoS ATTACKS

A DDoS attack can be extensive or intensive. An extensive attack relies on high volumes of traffic that by itself is harmless. A malicious actor needs a considerable amount of resources to successfully execute an extensive attack, as it is costly to generate enough traffic volume to impact a large target. Examples of these attacks include SYN flood, UDP flood, reflected Domain Name Service (DNS), and reflected NTP.

In most extensive attacks, miscreants may use a technique called amplification. Leveraging amplification, the attacker continuously abuses a




Figure 1. Infrastructure used by DaaS providers, including the payment platforms employed (phase 1) and the set of resources to launch the selected DDoS attack (phase 2). Intensive attacks predominantly utilize dedicated hosts with high bandwidth.

set of hosts that responds to a request with a considerably larger response that is delivered to the destination of the attacker's choosing. Previous studies have shown that this amplification factor differs according to the used protocol and can be as high as 4670 $\times$. These types of attacks have achieved throughputs as high as 500 Gb/s and affected enterprises with large infrastructures such as Sony PlayStation Network, Cloudflare, and several U.S. banks.

Intensive attacks, on the other hand, target specific weaknesses in a target application. Any request (or request access pattern) that takes a considerably larger amount of resources on the server than the client can be leveraged to perform this attack. These vulnerabilities can be due to problems like memory leaks and long running processes that never free their resources. Most cases of intensive attacks target HTTP servers, given their popularity on the Internet. Examples include submitting data to web forms found on the victim server, at very slow rates (one byte at a time), and opening multiple connections that are kept alive by sending partial packets. These examples have been implemented by the *R-U-Dead-Yet?* (*RUDY*) and *Slowloris* tools [2], respectively. Also worth noting is that intensive attacks only send legit packets, not malformed ones, making the resulting traffic appear legitimate, complicating their detection by security systems.

BASIC SCENARIO FOR A DDoS AS A SERVICE PROVIDERS

The continued rise of DDoS attacks as a way to target the online presence of organizations can be attributed to several factors. One possibility is that these attacks are often conducted through botnets, which often encompass thousands of computers. Pools of vulnerable computers are always available, given the constant discovery of software bugs.

Another possible factor for the rise of DDoS attacks is the commoditization phenomenon that these types of attacks have seen in the last few years. A large number of DaaS providers are avail-

able on the Internet, providing cheap access to both extensive and intensive DDoS attacks. Using a subscription-based model, the providers' fees range between \$2 and \$15 for basic packages. They support different payment mechanisms, ranging from traditional online systems like PayPal to the Bitcoin electronic currency and anonymous payment systems like Paysafecard. The basic packages allow launching attacks for 60–90 s and currently produce attack volume peaking at more than 1.4 Gb/s. More expensive packages are also available, which provide longer attack periods and subscription terms. The same sets of extensive and intensive DDoS attacks are available for all subscription packages.

Figure 1 shows a diagram of the infrastructure used by DaaS providers to offer their *pay*, *point*, and *click* service. The diagram includes the payment platform used (phase 1, *pay*), as well as the components used by the providers to launch a DDoS attack (phase 2, *point* and *click*). As shown in the diagram, intensive attacks are launched using dedicated connections, since only a small set of hosts is required and software needs to be installed to interact with the logic of the web application under attack. Botnets and misconfigured hosts are commonly used when launching the volumetric, extensive attacks.

A common trait found in DaaS providers is the usage of anti-DDoS service providers to protect their web platforms. As many of them claim to be only used to stress test the resources owned by a customer, the providers include DDoS protection mechanisms in their infrastructure.

Given the shady nature of the business, DaaS providers are not particularly dependable services. In our study, we found them to have a short life span (compared to legitimate online services), measured in weeks to months. Of the 17 providers identified and tested, only 7 were functional at the end of our three-month evaluation. Additionally, those providers that were functional delivered an average of only 44 percent of the offered services. We also found several systems provided intermittent service.

Given the shady nature of the business, DaaS providers are not particularly dependable services. In our study, we found them to have a short life span (compared to legitimate online services), measured in weeks to months. Of the 17 providers identified and tested, only 7 were functional at the end of our three-month evaluation.

There are multiple risk factors associated with studying cyber-miscreants. To deal with these factors and to develop the ethical framework for our experiments, we followed the ethical guidelines for computer security research defined in The Menlo Report and consulted previous work where researchers actively interacted with systems or networks used by cyber-miscreants.

DaaS/run	1	2	3	4
APO	2	—	90	2289
BIG	90	415	61	170
DAR	4256	—	—	—
DES	38,194	11,889	20,922	10,727
DIV	—	4	8	—
GRI	20,752	—	—	—
HAZ	—	1	2	1
IDD	—	4	2	64
ION	5	4	4	14,118
IPS	2284	—	—	—
NET	1776	1854	1556	982
POW	2759	3727	3723	—
QUA	8132	—	—	—
RAG	30,505	4018	4	3
RES	8499	—	—	—
TIT	21,609	2274	3501	8238
WRA	7219	6891	11,699	95

Table 1. Traffic generated by each DaaS (MB).

THE DDoS AS A SERVICE LANDSCAPE

METHODOLOGY

We identified 28 different DaaS providers for our study, from visiting multiple hacking sources: forums, blogs, mailings lists, and news sites. A user account was then created on each of the 28 providers. After reviewing the corresponding websites, 17 were determined to be operational. The other 11 failed to provide a working service interface. We later realized that this failure rate is the result of the common short and intermittent life span experienced by DaaS providers (usually weeks to months). For example, 12 out of the 17 providers were available since the start of our investigation, while the other 5 became active later in the process.

Using each of the 17 operational providers, we investigated the DaaS ecosystem from both sides of the attack.

As a DaaS Customer: After registering on the website of each provider, their services were bought for a limited time, selecting the cheapest services available on each website. The prices varied from \$2 to \$15. We studied the different functionalities provided on these websites to help determine how their advertisement, payment systems, and business aspects work. Additionally, our analysis also included a look at their offered attack capabilities.

As a DDoS Victim: We set up a machine to serve as a target of DDoS attacks and ordered each provider to launch the strike against it. The victim machine was an Ubuntu Linux machine with 8 GB of RAM, 1 TB of SSD disk space, dual-

core Intel processor, an optical fiber network connection of 10 Gb/s to the Internet, running an Apache web server with MediaWiki software, and hosting a clone of a university's department website. The machine was connected to the Internet through a dedicated link that allowed isolation of our tests from the rest of the university campus network and prevented it from being negatively affected. We captured all the traffic aimed at our victim machine, its responses, and its internal state during the attacks.

Each DaaS was tested four times over a period of three months, from May to July 2014. In each of the four runs, we tested all the attack types offered by each of the working DaaS and captured all the resulting traffic. At all times during the testing, we ran only one type of attack from a single DaaS. Also, to prevent late packets from one attack from being mixed with the next, we waited for 100 s between consecutive attacks.

ETHICAL CONSIDERATIONS

There are multiple risk factors associated with studying cyber-miscreants. To deal with these factors and to develop the ethical framework for our experiments, we followed the ethical guidelines for computer security research defined in the Menlo Report [3] and consulted previous work where researchers actively interacted with systems or networks used by cyber-miscreants [4, 5].

To reduce the risk of financing possible cyber-miscreants during our experiments, we purchased the cheapest services from the DaaS providers. This meant a single DaaS provider received no more than \$45, as we repeated the experiments three times on the most expensive (\$15) service used.

Another risk factor for studies such as ours is to unwittingly and negatively affect other victims. In this case, the victims can be compromised machines used by the providers to launch the DDoS attacks or other machines and networks on the path of the attack that are affected by the amount of generated traffic. To mitigate the potential risks, our experiments included conditions to restrict the duration and intensity of the attacks, limit the path of the attack traffic, and coordinate the experiments with the system administrators of our campus networks.

As mentioned before, we ran each attack for only 60 s to limit the impact of each attack. In addition, the target machine used to receive the attacks was located on an isolated subnet of our campus network and connected to a dedicated 10 Gb/s link so that the traffic generated during the tests would not affect other subnets (and their hosts) on campus. We also ran all high traffic tests during weekend nights to further reduce impacting network bystanders.

We acquired the campus network administrators' permission to run our tests before proceeding, agreed on a schedule, and established a contingency plan in case an undesirable situation happened. We followed up with the network administrators after each round of experiments and confirmed with them that an experiment had not negatively affected other parts of the campus network before proceeding with the next round.

Finally, it should be mentioned that our research was out of scope of the institution-

al review board (IRB) committee given that the experiments with DaaS providers did not include any type of direct or indirect experiments with human beings.

RESULTS FOR DAAS PROVIDERS

The four test runs generated around 255 GB of traffic and more than 94.1 million packets. The top four protocols (DNS, CHARGEN, Simple Network Management Protocol [SNMP], and NTP) produced 91.3 percent of the total traffic generated. DNS was the top traffic contributor with 71.07 GB, while NTP was the top packet generator with 34.9 million packets. Attacks using HTTP only produced 0.71 GB from 4.72 million packets.

Table 1 shows the amount of traffic generated by each DaaS during a run. Those providers that were not active in a run are shown with a dash (—). Results showed that 10 to 14 DaaS were active in a single run and that traffic generated varied among the different providers. For example, the RAG¹ and DES DaaS generated 30.5 and 38.2 GB each in run 1, while APO and ION only produced 2 and 5 MB. Out of the 47 tests that produced traffic across the four different runs, 26 (55 percent) produced at least 1 GB.

The functionalities provided by different DaaS providers differ greatly in terms of their claimed and actual attack types provided. Table 2 shows the offered attack capabilities of each DaaS. In this table, each row is a type of attack, and each column represents a DaaS. A checkmark (✓) indicates that the feature was offered and indeed worked during the experiments. An (✖) means the feature was offered but did not work for any test run. A blank space means that the feature was not offered.

A total of 28 different attack methods were identified across the 17 DaaS providers under evaluation. Out of these attack methods, 17 were extensive DDoS attacks, 7 were intensive, and 4 never worked. Of these seven intensive attacks, we found that some of the tools used by the providers to launch these attacks targeted different web server implementations. For example, the *Apache Remote Memory Exhaustion* (ARME) tool is only effective against Apache servers, as the name implies, while the Slowloris tool targets Apache, HTTPd, and GoAhead web servers. As observed in our experiments, both tools send partial, legitimate packets to keep connections open and do not generate large volumes of traffic compared to extensive attacks.

Table 3 present the number of completed TCP connections to the victim, the number of unique non-spoofed IP addresses, and the maximum observed throughput for the DaaS producing the largest traffic.

DAAS INFRASTRUCTURE FOR INTENSIVE ATTACKS

To characterize the machines and networks used by the DaaS providers to launch their intensive attacks, we first determined the non-spoofed IP addresses that initiated the attacks. An address was labeled non-spoofed if at least one complete TCP connection was established with our victim server during the test, which provided a lower bound of the actual situation. Among all (intensive and extensive) attack traffic observed, only 0.71 percent was associated with non-spoofed

addresses, an expected result given the usual incognito nature of extensive attacks and the considerably larger traffic they produce.

Using the technique described above, a total of 26,271 non-spoofed IP addresses were identified in all the attacks launched to our victim server and across the five providers that successfully produced the attacks. As shown in Table 4, the number of IP addresses used by a DaaS varied from 35 (TIT) to 21,809 (WRA). The low number of addresses for TIT was a sign of the DaaS soon to go offline, as the service stopped after our second run. WRA, on the other hand, consisted of a large botnet, primarily composed of compromised or misconfigured WordPress web servers. WRA was also the only provider to successfully produce six different types of intensive attacks (GET and POST floods, ARME, Slowloris, RUDY, and XML-RPC pingback) and worked for all four runs.

IP2Location [6] was consulted to determine the geographical information of the IP addresses, their autonomous system number (ASN), and the type of networks to which they were connected. As IP2Location provides various degrees of geolocation accuracy, we limited our analysis to using country and region (state in the United States) information in order to determine the location of addresses. Additionally, we used their classification of subnets and ASNs to label the IP addresses as part of one of the following three types of networks: *broadband/residential*, *commercial hosting providers*, and *other*.

Results show DaaS with different geographical extensions and mixtures of types of machines. The United States and China were the largest sources of machines for the providers, with the United States providing at least 55 percent of the machines in the cases of WRA, DES, and BIG. China was the largest source for RAG and TIT, providing at least 39 percent of the attacking hosts. RAG presented a larger number of countries hosting machines and associated ASNs than BIG, even though they both had similar numbers of IP addresses. 81 percent of the addresses used by RAG were in 10 different countries, and 74.1 percent were connected to broadband networks. In comparison, BIG had 81 percent of its machines located in one country (United States) and 128 addresses (93.3 percent) are connected to networks identified for hosting. Moreover, 85 of those addresses were attributed to a single data center in Arizona. We experienced more effective (able to leave our server unresponsive) and reliable (available through all runs) attacks by using BIG than when launching attacks through RAG, which not surprisingly suggests that machines in hosting networks might be more valuable for DaaS than in those in broadband networks.

After identifying the addresses with at least a complete TCP connection in the intensive attacks, we knew that the attacker's machine either had that IP address, or went through a proxy or VPN using that address. To determine each case, we scanned the IP address actively and also fingerprinted the host passively, as both approaches complement each other. An active scan interacts with the target host by sending a predefined set of packets and determining the type of the host based on its response. As such, this approach allows identifying when a proxy is used. In con-

Our findings show that 81.5 percent of the non-spoofed IP addresses belonged to Linux machines and 12.5 percent to Windows hosts; the rest of the machines were not identified.

The high occurrence of Linux hosts and non-spoofed IP addresses suggests that the DaaS providers depended on machines that use popular OSs, such as dedicated servers and Internet of Things devices, to successfully launch attacks.

¹ Throughout this article, each DaaS provider is referred to by a three-letter code in order to keep its real name anonymous and avoid publicizing its service. For example, a DaaS named GeneralTester could be referred to as GRL.

Attack/DaaS	APO	BIG	DAR	DES	DIV	GRI	HAZ	IDD	ION	IPS	NET	POW	QUA	RAG	RES	TIT	WRA	No. DaaS
Extensive attacks																		
UDP	(x)		✓	✓	(x)	✓	(x)	(x)	(x)	✓	✓			✓		✓		7/12
Home Conn.				✓													(✓)	1/2
XSYN	(x)			✓											(x)		(x)	1/4
SSYN	(x)		(x)	✓					(x)	(x)	✓		✓	✓	(x)		✓	5/10
SSDP			✓										✓		✓			1/1
ESSYN	(x)				(x)			(x)						✓		✓	✓	3/6
ZSSYN																		1/1
NUDP (Net BIOS)				✓														1/1
SUDP (SNMP)		✓		✓	(x)													2/3
Website				✓														1/1
XBOX Live				✓														1/1
DNS					(x)		(x)								✓		✓	2/4
CHARGEN	(x)				(x)			(x)				✓	(x)	✓				2/6
NTP					(✓)									✓	✓		✓	4/5
TCP Amp.																	✓	1/1
RUDP								(x)										1/2
UDPLAG	(x)		✓		(x)			(x)	(x)	✓	✓	(x)	(x)	✓		✓	✓	8/14
Intensive attacks																		
POST				(x)		(x)			(x)		(x)			✓	(x)		✓	2/7
HEAD				(x)		(x)			(x)		(x)			✓	(x)		(x)	1/7
GET				(x)		(x)			(x)		(x)			✓	(x)		✓	2/7
ARME				(x)		(x)			(x)		(x)			✓	(x)		✓	2/7
SLOWLORIS				✓		(x)			(x)		(x)			(x)	(x)	✓	✓	3/8
RUDY			(x)	(x)					(x)	(x)	(x)			(x)	(x)	✓	✓	2/9
XML-RPC		✓	(x)	✓	(x)	(x)					(x)		(x)		(x)		✓	3/9
Not working																		
Source Engine											(x)							0/1
KS												(x)						0/1
Joomla			(x)															0/1
OVH						(x)												0/1
No. Attacks	0/6	2/2	3/7	10/17	0/8	5/12	0/2	0/5	0/9	2/4	4/11	1/3	2/5	10/12	3/12	5/5	12/15	

Table 2. Attack methods offered by each DaaS provider tested.

trast, a passive fingerprinting method observes the traffic originating from the target host and determines its type by looking for patterns that identify a particular operating system or application.

Our findings show that 81.5 percent of the non-spoofed IP addresses belonged to Linux machines and 12.5 percent to Windows hosts; the

rest of the machines were not identified. The high occurrence of Linux hosts and non-spoofed IP addresses suggests that DaaS providers depended on machines that use popular OSs, such as dedicated servers and Internet of Things devices, to successfully launch attacks. In terms of proxies used by the providers, we found that they

DaaS/run	Number of connections/number of unique IP addresses				Max. attack size (Mb/s)/run
	1	2	3	4	
BIG	20,408/127	7076/85	6625/39	2314/50	84.65/2
DES	–/–	–/–	76,483/9409	51/1	690.18/2
RAG	4226/168	1665/168	–/–	–/–	852.49/1
RES	7523/527	–/–	–/–	–/–	1494.05/1
WRA	55,077/459	89,728/271	71,819/278	51,564/21,573	579.84/2

Table 3. Number of connections and unique IP addresses for top traffic generating DaaS per run.

employed proxies in very small numbers, as only 0.76 percent of the non-spoofed addresses were identified as proxies, anonymizing VPN service or TOR exit node. IP2Location also provided information on addresses identified as proxies, validating 92 percent of our results.

Through the four runs of experiments launching intensive attacks, we found few cases of IP address sharing among providers. Most did not share any addresses, and in the cases where they did, it was in very low numbers (1 to 5 addresses). This suggests the appropriation or exclusive control of the machines by each DaaS. WRA was the only exception to this, sharing 5223 addresses with DES, thanks to exploiting a high-risk vulnerability [7] on WordPress servers that was publicly reported during our runs. The vulnerability did not provide a mechanism for attackers to control who could exploit these servers, thus leaving the opportunity for sharing.

Table 5 shows the number of IP addresses reused by BIG and WRA during our experimental runs, as these were the only providers that generated non-spoofed traffic in all four executions. The diagonals in the table show (in bold italic) the total number of IP addresses used by each DaaS in a single run. From our experiments, both providers had to continuously add new machines to their networks, as many of the IP addresses from an attack execution would not be found in the next. As an example, BIG showed 122 addresses in the first run, but only 66 (54 percent) of those would be present in the second run. The attacker needs to constantly find new machines, which is not always trivial. From the second to the third run, BIG went from 82 to 37 IP addresses, and only two of those were new. In the case of WRA, the 21,573 different addresses found in the fourth run correspond to web servers exhibiting the high-risk vulnerability to WordPress, as discussed above.

OPERATIONAL STABILITY

Given the shady nature of their business, DaaS providers are not particularly dependable services. Our study found them to have a short life span (compared to legitimate online services), measured in weeks to months. This was supported by the fact that 11 of the 28 DaaS identified failed to provide any service, while several of the other DaaS briefly disappeared during the different executions. Only seven of the 17 DaaS were functional for all four runs, while four were successfully used in three runs and one DaaS was available in two runs. Additionally, 3 of the 11 providers

that were not working when we first accessed them started working after three months.

13 out of the 17 tested providers claimed to support intensive DDoS attacks, but when we tested them, only five successfully executed one or more types of application layer DDoS attacks. Out of the 17 DaaS providers tested, only 7 were still working after we finished our study.

PAYMENT METHODS

The most popular payment methods used by the DaaS providers were the popular online payment system PayPal and the Bitcoin digital currency. Other methods found included the payment platforms Google Wallet, Paysafecard (which allows anonymous transfers), Payza (transfers using email), and Skrill (focused on low-cost transfers). During the tests, three of the providers had their PayPal accounts deactivated and could not receive money.

DaaS providers offered multiple subscription options for their services at different prices. For 10 providers, a higher price only means a longer period of attack and longer-term subscriptions. In other words, they did not offer additional attack methods or an increase in the intensity of the attacks.

We evaluated GRI, one of the four providers that claimed better throughput and additional methods of attacks, to observe the difference between the cheap and more expensive options. This DaaS was chosen as it offered the most powerful attack, and in terms of throughput, pricing was cheaper than other DaaS (\$50, compared to up to \$300 in the case of RAG), and offered a different class of attack. Results show that the more expensive service gives access to two VIP servers (servers that regular accounts do not have access to) at the same time (and therefore able to execute two concurrent attacks). The amount of traffic generated and the list of offered attacks by each VIP server were not different from its cheap service.

RELATED WORK

Research on the analysis of existing DDoS attack vectors [8–11] has focused on the resources available on the Internet that can be used to launch DDoS attacks. Particularly, researchers have studied the amplification effect produced from using certain network services on the impact from using botnets to create DDoS attacks. Our work complements previous research by providing an unabridged analysis of the new vector available to attackers: application-level, intensive DaaS.

DaaS providers offered multiple subscription options for their services, at different prices. For ten providers, a higher price only means longer period of attack and longer-term subscriptions. In other words, they did not offer additional attack methods or an increase in the intensity of the attacks.

DaaS	Total No. IP addresses	No. countries	No. ASNs	Type of network			No. proxies found	Additional information
				Broadband	Hosting	Other		
BIG	165	20	40	6.7%	93.3%	0.0%	0	U.S. hosts 81.8% of all addresses, while next four countries account for 8.5%
DES	9405	88	1446	11.8%	84.8%	0.4%	11	U.S. hosts 61% of all addresses, followed by 10 countries with more than 100 addresses each
RAG	162	36	84	74.1%	6.8%	19.7%	58	China accounts for 39.5% of all addresses, while Brazil, Indonesia, Rusia, and Guatemala together host 27.16%
TIT	35	10	22	45.7%	48.6%	5.7%	0	China and U.S. host 45% and 22.9%, respectively
WRA	21,809	117	3075	20.12%	79.82%	0.06%	130	U.S. accounts for 55.1% of all addresses, while 19 other countries host at least 140 addresses

Table 4. Geographical distribution of the IP addresses for each of the DaaS providers that generated intensive attacks. The table also includes for each provider: the number of ASNs involved, the type of network to which the addresses were connected, and the number of proxy servers identified.

Run/run	Big				WRA			
	1	2	3	4	1	2	3	4
1	122	66	35	22	426	176	176	157
2	—	82	35	20	—	269	184	163
3	—	—	37	17	—	—	277	170
4	—	—	—	49	—	—	—	21,573

Table 5. Number of non-spoofed IP addresses reused, per run, for BIG and WRA. Values in the diagonal (shown in bold italic) represent the total number of IP addresses used to launch intensive attacks in each run.

Rossow [10] studied several UDP-based services available on the Internet that can be misused for amplification during a DDoS attack, showing that they are numerous and easy to find on the Internet, and providing a byte amplification factor of up to 4670. Kührer *et al.* [9] showed the possibility of using various TCP servers as reflective traffic amplifiers, and measured their possible impact. Czyz *et al.* [8] studied the temporal properties of reflectors, especially from NTP servers, while Rijwijk-Deij *et al.* [11] showed that a byte amplification factor of over 102 is possible by abusing the DNSSEC extensions.

Recent work [12, 13] has also looked at the rising threat of DaaS providers. We consider all previous studies complementary to ours, as they did not analyze the application-level, intensive DDoS attacks that can be launched from these providers, as done in our study. Karami *et al.* [12] only evaluated the infrastructure used for extensive attacks, while Santanna *et al.* [13] limited the study to extensive attacks using the DNS or CHARGEN protocols. Noroozian *et al.* [14] profiled the victims of extensive attacks launched by DaaS providers by using a network of honeypots running open services to launch amplification attacks. The study found that 88 percent of the victims were housed in broadband and hosting ISP networks, while the ICT development and GDP per capita of the host countries also help explain the victimization rate.

CONCLUSIONS

With the goal of demystifying the newly prevalent class of DaaS providers, we identified and studied 28 of these online systems. Given the short life of many of the providers found, we analyzed the behavior of 17 over a period of three months. Results show DaaS providers commonly offer both extensive and intensive DDoS attacks, and over different protocols. Customers only have to spend tens of dollars to have access to the attacks, which we were able to use to launch 1-minute attacks that generated 255 GB of traffic and were able to achieve throughput of 1.4 Gb/s, at a cost of tens of dollars.

In our study, we showed that many of these publicly accessible providers allow users to launch intensive attacks, hence the need to also study this increasingly popular threat. Results show that these providers pose a real threat to web servers on the Internet as they have access to networks of up to tens of thousands of machines to generate traffic that looks inconspicuous but leaves the servers unresponsive.

REFERENCES

- [1] R. Chang, "Defending against Flooding-Based Distributed Denial-Of-Service Attacks: A Tutorial," *IEEE Commun. Mag.*, vol. 40, no. 10, Oct. 2000, pp. 42–51.
- [2] E. Cambiaso *et al.*, "Slow DoS Attacks: Definition and Categorisation," *Int'l. J. Trust Management in Comp. and Commun.*, vol. 1, no. 3–4, Jan. 2013, pp. 300–19.
- [3] D. Dittrich and E. Kenneally, "The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research," U.S. Dept. Homeland Sec., Aug. 2012.
- [4] C. Kanich *et al.*, "Spamalytics: An Empirical Analysis of Spam Marketing Conversion," *Proc. 15th ACM Conf. Comp. Commun. Sec.*, Oct. 2008, pp. 3–14.
- [5] B. Stone-Gross *et al.*, "Your Botnet Is My Botnet: Analysis of a Botnet Takeover," *Proc. 16th ACM Conf. Comp. Commun. Sec.*, Nov. 2009, pp. 635–47.
- [6] IP2Location, commercial IP geolocation databases, Jan. 2015; <http://www.ip2location.com/databases/>, accessed Jan. 5, 2015.
- [7] Symantec, "Security Focus: WordPress Slider Revolution Responsive Plugin 'img' Parameter Arbitrary File Download Vulnerability," July 2014; <http://www.securityfocus.com/bid/68942>, accessed Sept. 13, 2014.
- [8] J. Czyz *et al.*, "Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks," *Proc. ACM SIGCOMM Conf. Internet Measurement*, Nov. 2014, pp. 435–48.

- [9] M. Kührer et al., "Hell of a Handshake: Abusing TCP for Reflective Amplification DDoS Attacks," *Proc. 8th USENIX Wksp. Offensive Technologies*, Aug. 2014.
- [10] C. Rossow, "Amplification Hell: Revisiting Network Protocols DDoS Abuse," *Proc. Network Distrib. Sys. Sec. Symp.*, Feb. 2014.
- [11] R. van Rijswijk-Deij, A. Sperotto, and A. Pras, "DNSSEC and Its Potential for DDoS Attacks," *Proc. ACM SIGCOMM Conf. Internet Measurement*, Nov. 2014, pp. 449–60.
- [12] M. Karami, Y. Park, and D. McCoy, "Stress Testing the Booters: Understanding and Undermining the Business of DDoS Services," *Proc. 25th Int'l. World Wide Web Conf.*, Apr. 2016, pp. 1033–43.
- [13] J. Santanna et al., "Booters: An Analysis of DDoS-as-a-Service Attacks," *Proc. IFIP/IEEE Int'l. Symp. Integrated Network Mgmt.*, May 2015, pp. 243–51.
- [14] A. Noroozian et al., "Who Gets the Boot? Analyzing Victimization by DDoS-as-a-Service," *Proc. Int'l. Symp. Research Attacks, Intrusions, Defenses*, Sept. 2016, pp. 368–89.

BIOGRAPHIES

ALI ZAND (zand@cs.ucsb.edu) received his Ph.D. in 2015 from the University of California Santa Barbara, working on system security research with a focus on cyber situation awareness. His research interests include automatic service dependency detection, automatic asset protection prioritization, botnet C&C signature generation, cyber situation awareness measurement, DDoS attack studies, and social media spam detection.

GASPAR MODELO-HOWARD [SM] (gaspar@acm.org) is a senior principal data scientist in the Center for Advanced Machine Learning at Symantec. His research interest are computer and network security, with a focus on web security, intrusion detection and response, and malware detection. He is also an adjunct professor in computer security at Universidad Tecnológica de Panamá. He is a member of ACM and Usenix.

ALOK TONGAONKAR (alok@redlock.io) is head of Data Science at RedLock. Previously, he was a data scientist director leading the Center for Advanced Data Analytics at Symantec. He has a Ph.D. in computer science from Stony Brook University, New York. His research focuses on application of machine learning and big data technologies for developing innovative security, networking, and mobile app analytic products. He has been granted multiple patents by USPTO. He is a Senior Member of ACM.

SUNG-JU LEE [F] (sjlee@cs.kaist.ac.kr) is an associate professor and an Endowed Chair Professor at the Korea Advanced Institute of Science and Technology (KAIST). He received his Ph.D. in computer science from the University of California, Los Angeles and spent 15 years in the industry in Silicon Valley before joining KAIST. His research interests include computer networks, mobile computing, network security, and HCI. He is a recipient of multiple awards, including the HP CEO Innovation Award and the Test-of-Time Paper Award at ACM WINTER 2016. He is an ACM Distinguished Scientist.

CHRISTOPHER KRUEGEL (chris@cs.ucsb.edu) is a professor in the Computer Science Department at the University of California, Santa Barbara and one of the co-founders of Lastline, Inc., where he serves as the chief scientist. His research interests include most aspects of computer security, with an emphasis on malware analysis, web security, and intrusion detection. He is a recipient of the NSF CAREER Award, MIT Technology Review TR35 Award for young innovators, and IBM Faculty Award.

GIOVANNI VIGNA [SM] (vigna@cs.ucsb.edu) is a professor in the Department of Computer Science at the University of California, Santa Barbara and the CTO at Lastline, Inc. His research interests include malware analysis, vulnerability assessment, the underground economy, binary analysis, web security, and mobile phone security. He leads the Shellphish hacking group, which has participated in more DEF CON CTF competitions than any other group in history. He is a Senior Member of ACM.