computers

MDPI

Article

IP Spoofing In and Out of the Public Cloud: From Policy to Practice

Natalija Vlajic *, Mashruf Chowdhury and Marin Litoiu

Department of Electrical Engineering & Computer Science, York University, Toronto, ON M3J 1P3, Canada; mashrufkabir@icloud.com (M.C.); mlitoiu@yorku.ca (M.L.)

* Correspondence: vlajic@cse.yorku.ca

Received: 28 August 2019; Accepted: 25 October 2019; Published: 9 November 2019

Abstract: In recent years, a trend that has been gaining particular popularity among cybercriminals is the use of public Cloud to orchestrate and launch distributed denial of service (DDoS) attacks. One of the suspected catalysts for this trend appears to be the increased tightening of regulations and controls against IP spoofing by world-wide Internet service providers (ISPs). Three main contributions of this paper are (1) For the first time in the research literature, we provide a comprehensive look at a number of possible attacks that involve the transmission of spoofed packets from or towards the virtual private servers hosted by a public Cloud provider. (2) We summarize the key findings of our research on the regulation of IP spoofing in the acceptable-use and term-of-service policies of 35 real-world Cloud providers. The findings reveal that in over 50% of cases, these policies make no explicit mention or prohibition of IP spoofing, thus failing to serve as a potential deterrent. (3) Finally, we describe the results of our experimental study on the actual practical feasibility of IP spoofing involving a select number of real-world Cloud providers. These results show that most of the tested public Cloud providers do a very good job of preventing (potential) hackers from using their virtual private servers to launch spoofed-IP campaigns on third-party targets. However, the same very own virtual private servers of these Cloud providers appear themselves vulnerable to a number of attacks that involve the use of spoofed IP packets and/or could be deployed as packet-reflectors in attacks on third party targets. We hope the paper serves as a call for awareness and action and motivates the public Cloud providers to deploy better techniques for detection and elimination of spoofed IP traffic.

Keywords: IP spoofing; network security; cloud computing; firewalls; computer crime

1. Introduction

IP spoofing—the simple act of modifying an IP packet by replacing its genuine source address with a forged one as illustrated in Figure 1 has long been known as the key precursor for many different forms of cyber attacks and illegitimate online activities, including man-in-the-middle (MitM) attacks, distributed denial of service (DDoS) attacks, ARP and DNS poisoning attacks, spoofed port scanning, etc. What makes IP spoofing particularly challenging for cybersecurity defenders is (a) the fact that no network is entirely immune to attacks that deploy IP spoofing, and (b) there are many readily available tools and programming packages that allow even moderately-skilled hackers to integrate IP spoofing into their attacks, such as Scapy, Hping, Libcrafter, etc.

From a legal standpoint, IP spoofing in itself is not considered a criminal activity. As stated in [1], "Strictly speaking, the simple act of spoofing an identity is not illegal (i.e., no hacking is involved in the commission of the act). It only becomes illegal when a threat of death or violence is involved, or personal data are stolen in order to commit fraud or identity theft". In fact, there are many legitimate reasons/operations that may require the use of packets with forged IP addresses (e.g., network troubleshooting, penetration testing, workload and stress testing, user anonymization, etc.). As a result, until relatively recently, a large number of Internet service providers (ISPs) were turning a blind eye and allowing uninterrupted transmission of all sorts of spoofed IP packets generated by their customers. However, in the past few years, the number of ISPs that have started performing egress filtering on their outbound traffic, by dropping obviously spoofed outgoing IP packets, has grown considerably. This has likely come as a result of the Mutually Agreed Norms for Routing Security (MANRS) initiative [2], which was introduced in 2014 by the Internet Society (ISOC), the parent organization of the Internet Engineering Task Force (IETF) standards body. One of the key recommendations imposed by MANRS to its members is that 'network operators implement anti-spoofing filtering to prevent packets with an incorrect source IP address from entering and leaving their network' [2]. To date, more than a hundred worldwide network operators (ISPs) have agreed to participate in MANRS, each of them representing dozens, hundreds, or in some cases, thousands of network subdomains [2].

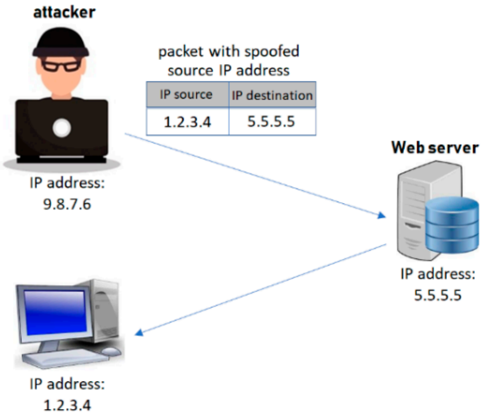

Figure 1. Transmission of spoofed IP packets.

The Center for Applied Internet Data Analytics (CAIDA) [3] is a US-government funded group committed to performing comprehensive measurements and visualization of various type of data pertaining to the global Internet traffic and topology. Among other publicly available outputs of this group, the real-time charts depicting the global as well as country-by-country susceptibility to IP spoofing are particularly interesting. According to these charts (as of June 2019), on average, only 15% to 30% of spoofed traffic is not successfully detected and blocked by worldwide ISPs. This suggests that at present, globally, and on average, ISPs do a reasonably good job of dealing with spoofed IP traffic. However, it is important to note that deviations from this average are quite significant when looking at individual countries. For example, in the United States of America and Canada, the likelihood of IP spoofing successfully taking place is under 5%, while in Bolivia, Uganda, and Zambia, it exceeds 60% (see Figure 2).

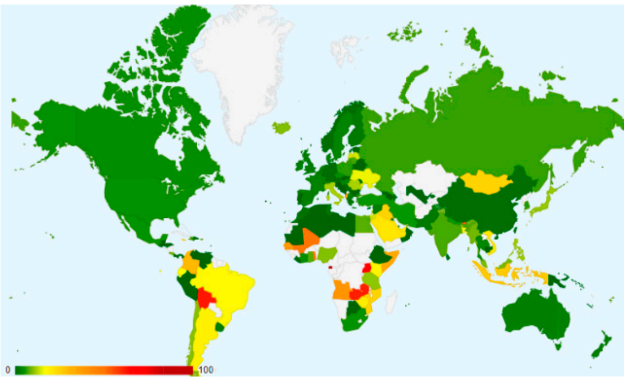

Figure 2. Center for Applied Internet Data Analytics (CAIDA)-produced map that shows the likelihood of IP spoofing successfully taking place within a particular country-level domain [3].

Over the past few years, we have seen increasing evidence that cybercriminals are turning to the public Cloud as an attractive (new) platform for launching a whole range of different attacks. As pointed out in [4], hackers are 'embracing the use of public Cloud services for the same reasons as legitimate organizations: public Cloud services provide flexible, on-demand capacity and resources, and can be provisioned in just a few minutes'. Furthermore, 'by offering bandwidth of between 1 to 10 Gbps, public Cloud service providers enable attack volumes, which can be as much as 1000 times higher than is possible with individual compromised devices, such as home routers or IoT cameras'. According to [5], a quarter of all DDoS attacks conducted in Europe between July 2017 to June 2018 involved the use of Cloud-based botnets. In February 2019, it was revealed that one of the most notorious DDoS-for-hire services had built its operation using a botnet solely consisting of 31 large-capacity Cloud servers [6]. The seriousness of the problem arising from the use of the public Cloud by cybercriminals is also confirmed by the fact that 'abuse and nefarious use of Cloud services' is enlisted as one of the top 12 Cloud security threats by the Cloud Security Alliance (CSA) [7]. All of the above gives reasonable grounds to suspect that the public Cloud may also be that 'new safe haven' from which hackers are more successful in generating spoofed IP traffic, compared to how spoofed traffic has been traditionally generated—by means of compromised desktops and laptops connected to residential ISPs.

While the nature and prevalence of mechanisms deployed by network operators (i.e., ISPs) to deal with spoofed IP traffic are well understood and documented. Unfortunately, there is very little (if any) available information about the state of detection and prevention against IP spoofing in the public Cloud domain. The goal of our recent study, which is the main focus of this paper, was to fill that void. In particular, through our study, we have aimed to (1) provide a comprehensive look at the topic of IP spoofing in the context of public Cloud services and, more specifically to shed light on when and why a hacker would engage in sending spoofed IP packets in and out of the public Cloud, and (2) present real-world data and observations on how a select number of Cloud providers, including Google Cloud, Microsoft Azure, and Amazon Web Services (AWS), actually deal with IP spoofing, both from the policy as well as practical packet-level point of view. According to our knowledge, this is the first reported study that looks at the real-world feasibility of IP spoofing based attacks within the context of public Cloud services and is based on real-world experimentation.

The content of this paper is organized as follows. In Section 2, we provide a brief review of the existing literature pertaining to the problem of IP spoofing and distributed denial of service attacks (DDoS) in the public Cloud. In Section 3, we describe a number of specific attack scenarios involving the transmission of spoofed IP packets from or towards the virtual private servers hosted by a public Cloud provider. In Section 4, we outline the motivation and specific objectives of our real-world experimental study. In Section 5, we present our findings concerning the reference to IP spoofing (or its lack of) in acceptable-use (AU) and term-of-service (ToS) policies of 35 surveyed public Cloud providers. In Section 6, we summarize our experimental results on the feasibility of spoofed-packet transmission *from* or *towards* virtual private servers hosted by a select number of real-world Cloud providers. In Section 7, we close the paper and outline the directions for future work.

2. Related Literature and Significance of This Study

As the popularity and prevalence of public Cloud platforms and services have increased over the last decade, so has the number of research works on the topic of public Cloud security. Some of these works, including [8–13], have offered comprehensive looks at a wide range of security and privacy issues pertaining to network, storage, infrastructure, and software aspects of the public Cloud. Other works have focused on more narrow topics and problems, including the problem of distributed denial of service (DDoS) attacks [14–16] and IP spoofing [17–20] in the public Cloud. Unfortunately, two commonly observed limitations of the existing studies on DDoS and IP spoofing in the public Cloud are (1) They are mostly theoretical works based on hypothetical assumptions. As such, these works provide no insight into the actual real-world feasibility of Cloud-based attacks that rely on IP spoofing, and their results are likely to have very limited practical value, at best. (2) They are largely written from the defender's perspective, so they offer little if any insight into the specific objectives and possible strategies deployed by hackers when conducting Cloud-based DDoS and/or IP spoofing campaigns. Yet, a clear understanding of 'why' and 'how' from the hacker's perspective is always helpful, if not critical, when developing effective real-world defenses.

The work presented in this paper aims to overcome the above-identified limitations of the existing research literature. According to our knowledge, this paper is the first published study that has looked at the problems of cloud-based DDoS and IP spoofing from the offender's perspective, and that has included experimental results involving real-world public Cloud providers. We hope that the discussion and conclusions of this study will provide both cybersecurity researchers and practitioners with a more complete understanding of the addressed problems, and as such, will assist in building an overall safer Internet.

3. IP Spoofing IN and OUT of the Cloud: Possible Scenarios

The first step in understanding the problem addressed in this paper requires us to identify the most likely practical scenarios which involve deliberate transmission of spoofed IP packets by a hacker to or from the virtual private servers (VPSs) (In this paper, the term virtual private server (VPS) is used to refer to a virtual machine (VM) sold as a service by a public Cloud service provider. It is a software-created emulation of a physical server/host with full root access to the VM's operating system of a public Cloud service provider (CSP).) Based on the discussion presented in [14,15], as well as our own knowledge and experience, three such (most likely) scenarios are outlined in Table 1, while more detail descriptions of each scenario are provided in Sections 3.1–3.3.

Table 1. Different attack scenarios in the public cloud involving spoofed IP packets.

Attack Scenario	**adversary leveraging VPS to attack 3rd party (described in Section II.A)**		**adversary attacking hosted VPS** (described in Section II.B)		**adversary using VPS as a packet reflector** (described in Section II.C)	
Variants of the attack	randomly spoofed IP packets	specifically spoofed IP packets	spoofed packet sent from a single machine	spoofed packet sent from multiple machine	spoofed packet sent from a single machine	spoofed packet sent from multiple machine
Flow of Spoofed IP Packets	out of the public Cloud provider's network		into the public Cloud provider's network		into the public Cloud provider's network	

3.1. Adversary Leveraging VPS to Attack 3rd Party

To ensure better fault-tolerance as well as better performance for their users, most public Cloud service providers deploy numerous geographically distributed servers. Moreover, many of these providers also allow their users to explicitly choose by which of these servers they prefer to be hosted. With this in mind, a hacker that resides on an ISP that bans IP spoofing and, despite this ban, is looking for ways to send spoofed IP packets towards some intended target (e.g., during a DDoS attack or spoofed port scanning campaign) could accomplish their objective by adopting the following strategy:

Step (1) Subscribe to a VPS hosting plan with a Cloud service provider with loose Internet traffic regulations towards IP spoofing;
Step (2) Choose to have this VPS specifically hosted by one of the Cloud's servers located in a geographic area (i.e., attached to an ISP) that also has lenient rules about IP spoofing;
Step (3) Use such VPS to run the attack script that generates the intended spoofed packets/traffic.

For an illustration, see Figure 3a,b. In addition note that for the successful transmission of spoofed packets from the given VPS, the conditions of both Step (1) and (2) need to be satisfied. Otherwise, if either the public CSP itself or the ISP(s) hosting the CSP's server(s) block spoofed IP packets, the hacker will not be able to accomplish his objective.

It should also be noted that in the attack scenarios of Figure 3a,b, the spoofed packets are sent 'out' of the Cloud provider's server/network that currently hosts the hacker's VPS, and from the practical point of view these packets could carry either 'randomly' or 'specifically' spoofed source IP addresses. Namely,

(a) The hacker would deploy 'randomly' spoofed source IP addresses (i.e., each generated packet would carry a different randomly generated source IP address) if he is conducting a direct attack on the target machine and is not concerned where the response packets generated by the target machine are going to end up, as shown in Figure 3a. (Here we assume that the spoofed IP packets are actually the 'carriers' of other higher-layer packets, such as ICMP or TCP, which in the majority of cases prompt the receiving machine to generate a response packet. The packet responses generated by the target machine receiving randomly spoofed IP packets are referred to as backscatter [21]).

(b) The hacker would use 'specifically' spoofed source IP addresses (i.e., all generated packets would carry the same forged source IP address) if the actual target of his attack is not the machines receiving the spoofed packets. Instead, the hacker is simply exploiting these 'intermediate' machines as 'packet reflectors', while the true target of his attack is another 'specific' machine. Put another way, the actual target of such an attack is the recipient of the backscatter traffic, as shown in Figure 3b. For more on the difference between direct and reflected attacks, see [14,15,21,22].

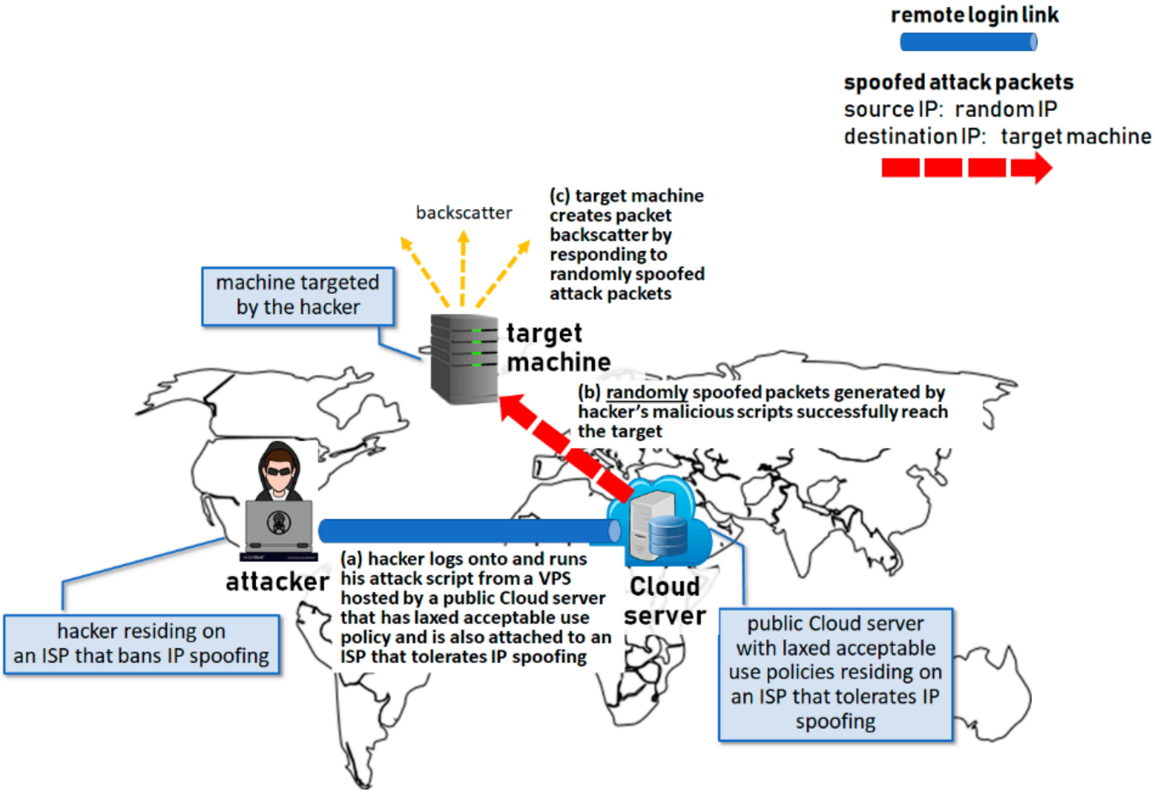
remote login link
spoofed attack packets
source IP: random IP
destination IP: target machine
backscatter
(c) target machine creates packet backscatter by responding to randomly spoofed attack packets
machine targeted by the hacker
target machine
(b) randomly spoofed packets generated by hacker's malicious scripts successfully reach the target
attacker
(a) hacker logs onto and runs his attack script from a VPS hosted by a public Cloud server that has laxed acceptable use policy and is also attached to an ISP that tolerates IP spoofing
Cloud server
hacker residing on an ISP that bans IP spoofing
public Cloud server with laxed acceptable use policies residing on an ISP that tolerates IP spoofing

(**a**)

Figure 3. *Cont.*

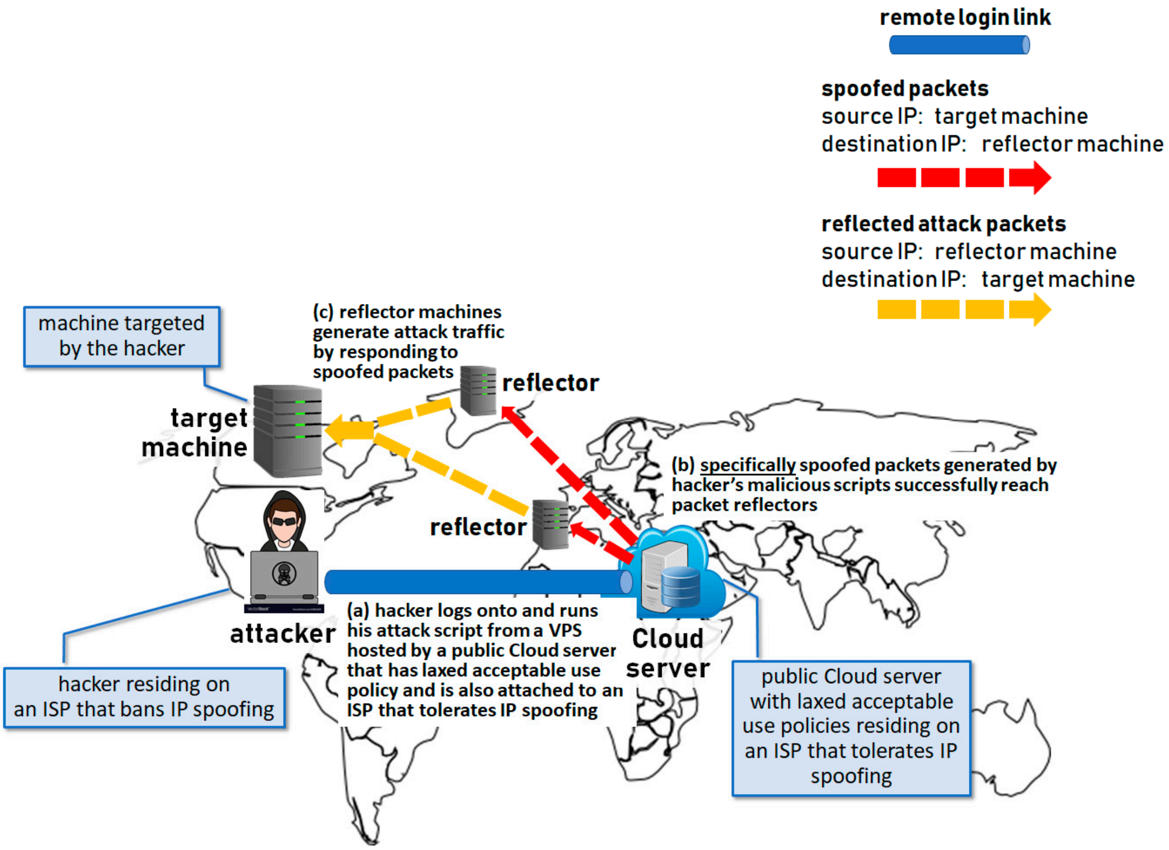
remote login link
spoofed packets
source IP: target machine
destination IP: reflector machine
reflected attack packets
source IP: reflector machine
destination IP: target machine
machine targeted by the hacker
(c) reflector machines generate attack traffic by responding to spoofed packets
reflector
target machine
(b) specifically spoofed packets generated by hacker's malicious scripts successfully reach packet reflectors
reflector
attacker
(a) hacker logs onto and runs his attack script from a VPS hosted by a public Cloud server that has laxed acceptable use policy and is also attached to an ISP that tolerates IP spoofing
Cloud server
hacker residing on an ISP that bans IP spoofing
public Cloud server with laxed acceptable use policies residing on an ISP that tolerates IP spoofing

(**b**)

Figure 3. (**a**) Attacker sending randomly spoofed IP packets from a virtual private server (VPS) hosted by a public Cloud service provider (CSP) in order to conduct a direct attack on the target. (**b**) Attacker sending specifically spoofed IP packets from a VPS hosted by a public CSP in order to conduct a reflection attack on the target.

3.2. Adversary Attacking VPS

DoS and DDoS attacks targeting Cloud infrastructure and services are recognized as one of the top 12 Cloud computing threats [7]. According to [23], 'as service providers place growing importance on the delivery of Cloud-based services to enterprises and consumers, it should come as no surprise that attackers are increasingly targeting these services with DDoS attacks'. The statistics show that the number of DDoS attacks on Cloud-based services nearly doubled between 2016 and 2018 [23], and in many cases, these attacks are conducted by means of spoofed IP packets [14,15]. The illustrations of a DoS and DDoS attack on a VPS of a public CSP using spoofed IP packets are provided in Figure 4a,b, respectively. Recall, in a DoS attack, the attack traffic is generated from one single machine operated by the hacker, while in the case of DDoS attack, the attack traffic is generated by a network of compromised bots/machines.

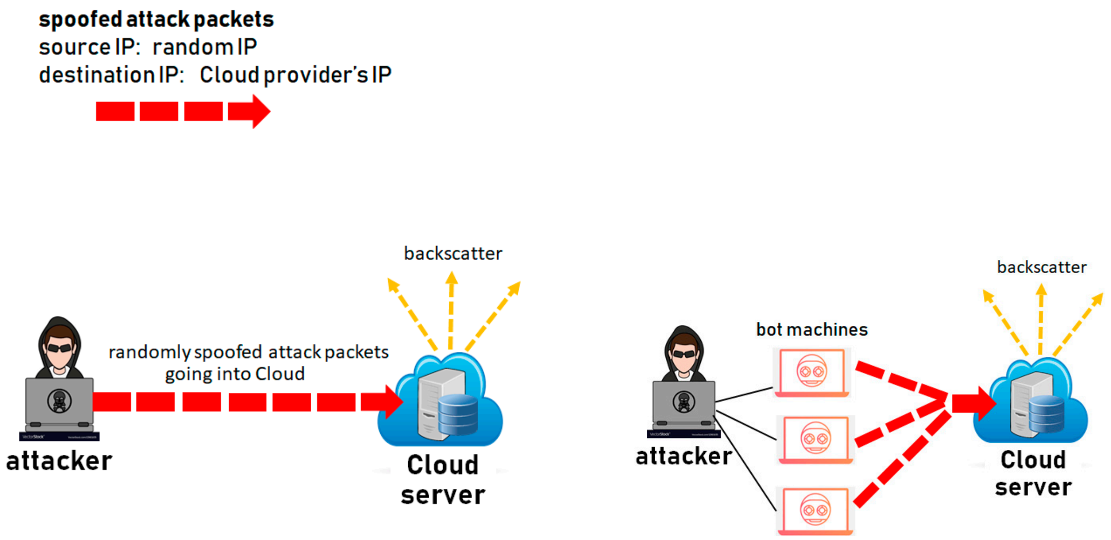

(**a**) (**b**)

Figure 4. (**a**) VPS or a physical server of a public CSP under-going a direct attack involving randomly spoofed IP packets—DoS variant. (**b**) VPS or a physical server of a public CSP under-going a direct attack involving randomly spoofed IP packets—DDoS variants.

It should be noted here that, in contrast to the scenarios of Figure 3 were spoofed IP packets are sent 'out' of the public Cloud, in the attacks of Figure 4 spoofed IP packets are sent 'into' the public Cloud's network. In addition, note that in the DoS variant of this attack, spoofed packets are sent from one single machine (Figure 4a), while in the DDoS variant spoofed packets are generated by a network of attacking/bot machines (Figure 4b).

3.3. Adversary Using VPS as a Packet Reflector

The idea of a reflection attack being initiated from the public Cloud and involving spoofed IP packets has already been discussed in Section 3.1 (see Figure 3b). One should recognize, however, that there is nothing preventing the attacker from actually attempting to deploy the very VPSs of a public CSP as the intermediate packet reflectors in an attack targeting a third-party server/host, see Figure 5. For the attacker, the advantages of this strategy over the use of (e.g.,) standard desktops or IoT devices as packet reflectors, are (i) public Cloud servers generally abound in CPU and bandwidth resources, and are up-and-running 24/7—all of which are very desirable properties of a packet reflector, and (ii) many network operators hesitate to block packets coming from or carrying a source IP address of a well-known public Cloud provider.

An incident that effectively illustrates why blacklisting of a public Cloud provider's IPs should generally be avoided by network operator occurred in April 2018. In an attempt to ban one particular Cloud-based service from operating in Russia, some of this country's main ISPs were ordered to block 1.8 million IP addresses belonging to Amazon and Google Cloud infrastructure. The move to mass-ban so many IP addresses from these two popular Cloud providers had serious secondary repercussions, as it also resulted in unintended disabling of many legitimate Web services (also hosted by those two Cloud providers) in Russia [24]. While for network providers, this story is a cautionary tale about the potential dangers of mass-blocking of IPs affiliated with the public Cloud, for hackers, this story points to additional benefits of deploying public Cloud networks as the launching or reflection points in their attacks.

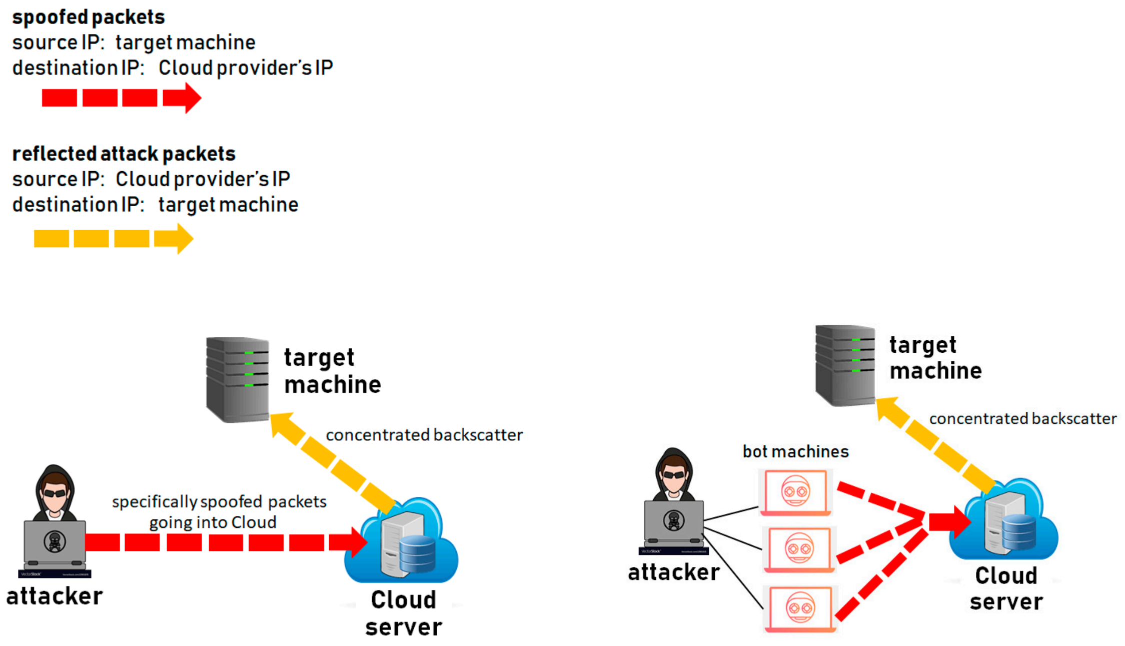

(**a**) (**b**)

Figure 5. (**a**) A VPS or a physical server of a public CSP exploited as a packet reflector in an attack involving specifically spoofed IP packets—DoS. (**b**) A VPS or a physical server of a public CSP exploited as a packet reflector in an attack involving specifically spoofed IP packets—DDoS variants.

4. Motivation and Objectives of Our Study

In the preceding section, we have described three attack scenarios involving the transmission of spoofed packets in and out of the public Cloud, which, based on a number of recent industrial reports as well as research studies, are most likely to take place in practice. Each of these attacks, if successfully executed, can inflict significant damage not only to the attack target but also (where applicable) to the attack's reflection points by making them process and transmit unnecessary traffic.

However, one should recognize that with the rapidly expanding reach of the MANRS initiative, as witnessed over the past few years (see Section 1), the actual real-world feasibility of the attack scenarios described in Section 3 is rather unclear. In other words, while there is a general agreement that the attacks of Section 3 are possible and have sporadically taken place in practice, no one is really able to tell how difficult to execute these attacks are in the present-day Internet.

The goal of the work presented in this paper was to provide a better understanding and explicit answers about the real-world feasibility of the attacks described in Section 3. More specifically, we have structured our work and results around the following three questions:

Question (1) How clearly, if at all, the concept and dangers pertaining to IP spoofing are addressed in the acceptable-use (AU) and terms-of-service (ToS) policies of some of the leading real-world Cloud service providers.

Question (2) How successful an attacker owning a VPS with a real-world public CSP would be in sending spoofed IP packets out of this provider's network to conduct the attacks of Figure 3. Or, looking from the opposite perspective—how effective real-world public CSP is in spotting and blocking outgoing spoofed IP packets generated by their malicious customers.

Question (3) How successful would an attacker be in sending spoofed IP packets to a VPS (IP address) of a real-world public CSP when conducting the attacks of Figures 4 and 5. Or, looking from the opposite perspective—how effective real-world public CSP are in spotting and blocking spoofed IP packets from reaching their servers.

It should be stressed here that our work was only concerned with the 'spoofing' aspect of the attacks in Figures 3–5. That is, the only goal of our study and experimentation was to examine how effective a representative group of real-world Cloud service providers is at detecting and blocking individual instances of spoofed IP packets. Thus, all of our experiments involving these Cloud service providers were based on the transmission of only a few select spoofed-IP-packet probes in and out of these providers' networks. We were not, in any way, interested in testing or challenging the ability of these providers to fend significant floods of spoofed packets, as such an activity could cause serous harm and is considered illegal in most jurisdictions around the world [25].

5. IP Spoofing in AU and ToS Policies of Real-World CSPs

As explained in the preceding section, one of the three main objectives of our study was to survey whether the concept and dangers pertaining to IP spoofing are sufficiently (if at all) addressed in the AU and ToS policies of a representative group of real-world public Cloud service providers. The main reason for this particular objective to be included in our study is a widely known fact that a well-defined AU or ToS policy can have an important deterrence effect on the potential attacker [26]. In other words, the mention and prohibition of a certain activity, in this case, IP spoofing, in the AU or ToS policy of a Cloud service provider could significantly reduce the chances that the users of this service ever engaging in that particular type of activity. Moreover, the mention and prohibition of a certain activity in these policies are also often a good indicator that the provider itself has taken due diligence in implementing effective technical measures aimed at detecting and preventing the given activity [26].

Our study looked at the AU and ToS policies of 35 public Cloud service providers, which included a mix of the top players in the field (e.g., AWS, Google Cloud, and Microsoft Azure) as well as a number of lesser-known public Cloud companies. Our findings, as shown in Table 2, reveal that 19 of 35 (i.e., more than 50%) of the surveyed providers, including Google Cloud and Microsoft Azure, make no direct reference to IP spoofing in their AU and ToS policies, thus failing to serve as potential deterrents. In the AU and ToS policies of the remaining 16 companies, IP spoofing is generally referred to and prohibited either directly using the term 'IP spoofing', or indirectly using the terms 'header forging', 'header falsifying', etc.

Table 2. Reference to IP spoofing in AU and ToS policies of some select public cloud service providers.

Cloud Service Provider (CSP)	IP Spoofing in CSP's AU and ToS Policy
alibabacloud.com	forging TCP/IP header prohibited
arubacloud.com	no explicit mention
aws.amazon.com	forging TCP/IP header prohibited
azure.microsoft.com	no explicit mention
bigrock.dom	no explicit mention
bluehost.com	forging message headers prohibited
buyvm.net	no explicit mention
cloud.google.com	no explicit mention
cloudvps.com	no explicit mention
deltahost.ca	forging TCP/IP header prohibited
digitalocean.com	misleading TCP-IP header prohibited
dostinger.com	packet corruption prohibited
go4hosting.in	no explicit mention
godaddy.com	no explicit mention
hetzner.com	fake source IPs prohibited
hostinginchina.com	IP spoofing prohibited
hostsailor.com	forging TCP/IP header prohibited
hostslayer.com	forging TCP/IP header prohibited
ideastack.com	no explicit mention

Table 2. *Cont.*

Cloud Service Provider (CSP)	IP Spoofing in CSP's AU and ToS Policy
justhosting.ru	no explicit mention
king-servers.com	false message headers prohibited
liquidweb.com	no explicit mention
miran.ru	falsified IP addresses prohibited
parkinhost.com	no explicit mention
ramnode.com	IP spoofing strictly prohibited
sakuraserver.com	no explicit mention
scaleway.com	no explicit mention
seedvps.com	falsifying of packet header prohibited
siteground.com	spoofing prohibited
swiss-vps.com	no explicit mention
unixhost.com	no explicit mention
veesp.com	no explicit mention
vps.net	no explicit mention
vultr.com	no explicit mention
znetliev.com	forging TCP/IP header prohibited

6. Transmission of Spoofed IP Packets To and From Select Public CSPs

6.1. Experimental Framework

As explained in Section 4, two other main objectives of our study were to investigate how effective real-world public Cloud service providers are in preventing spoofed IP packets from being transmitted from and to their respective servers (see Questions 2 and Question 3 in Section 4). To answer these questions, we proceeded by (i) setting up VPSs with a select number of Cloud providers (predominantly the ones from Table 2 whose AU and ToS policies did not explicitly prohibit IP spoofing), and (ii) using these accounts we performed real-world experimentation involving transmission and reception of spoofed IP packets. It should be noted that this second stage of our study did not include all CSPs from Table 2 which did not explicitly prohibit IP spoofing, as some of them demanded expensive monthly payments, payments in Bitcoin, or long-term subscription (3 + months). As a result, we ended up setting up VPS accounts and experimenting with 14 of those providers. (For the actual list of these select providers see Table 4).

The actual real-world experimentation with these select Cloud providers required us to install several network monitoring tools as well as a Python and Scapy environment on each of the 14 VPSs. Further, we developed and used our own custom-coded Python–Scapy scripts for transmission and reception of spoofed IP packets. The actual experiments were grouped into two sets. In Experiment 1, the script sending spoofed IP packets was set to run on a VPS hosted by one of the 14 select Cloud providers, while the script listening for these packets was set to run on our dedicated host at York University, as shown in Figure 6. In Experiment 2, the script sending spoofed IP packets was set to run on a VPS hosted by a Cloud provider that was previously confirmed to allow successful transmission of spoofed IP packets (we refer to it as CSP_1), while the script listening for these packets was set to run on a VPS hosted by one of the remaining 12 select Cloud providers (we refer to it as CSP_2), as shown in Figure 7. Both experiments were performed several times in the period April to July 2019.

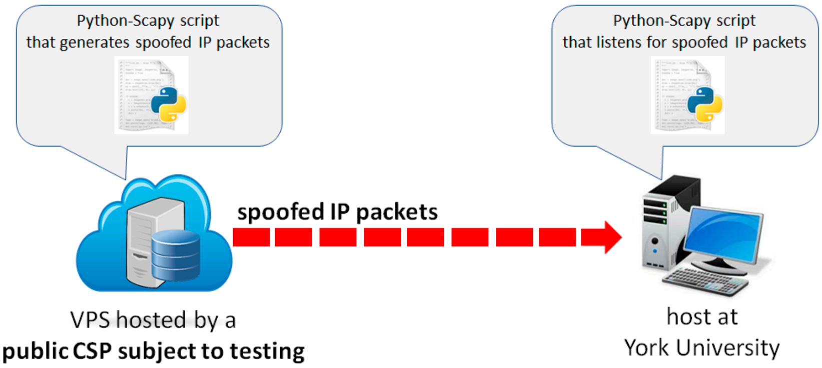
Python-Scapy script
that generates spoofed IP packets
Python-Scapy script
that listens for spoofed IP packets
spoofed IP packets
VPS hosted by a
public CSP subject to testing
host at
York University

Figure 6. Framework of Experiment 1.

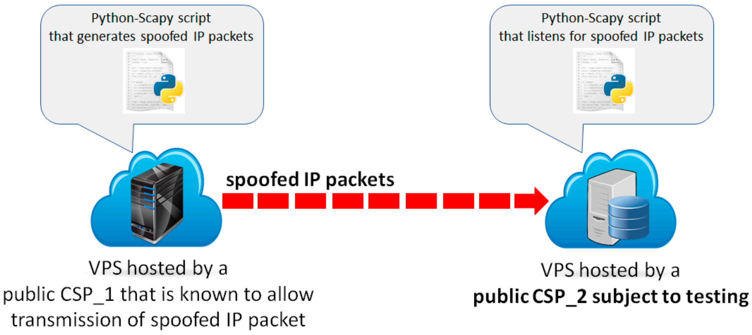
Python-Scapy script
that generates spoofed IP packets
Python-Scapy script
that listens for spoofed IP packets
spoofed IP packets
VPS hosted by a
public CSP_1 that is known to allow
transmission of spoofed IP packet
VPS hosted by a
public CSP_2 subject to testing

Figure 7. Framework of Experiment 2.

It should also be clarified that each experiment involved the transmission of eight different types of spoofed IP packets, as shown in Table 3. The type of packets and the nature of spoofed IP addresses were chosen following the recommendations from [27], and they were specifically designed to allow inference about the presence of some important best-practice filtering methods in the examined Cloud providers' servers and networks. That is, our packet probes were intended to evaluate the Cloud providers' ability not only to deal with spoofed IP packets but also to spot and filter out packets and IP addresses that were/are outright invalid. To understand the reasoning that led to the design of our packet probes, be reminded of the following:

1. 1.2.3.4 is unallocated while 172.16.1.100 is a private IP address [27]. Hence, no packet going into or coming out of the Internet should carry either of these as its source IP address.
2. 8.8.8.8 is the IP address of a well known public Google DNS server, and DNS servers generally do not engage in the transmission of ICMP Echo Requests.
3. A machine or a VPS should not be receiving an ICMP Echo Reply unless this machine/VPS has explicitly sent an initiating ICMP Echo Request.
4. The most widely known passive defense technique against IP spoofing is the so-called TTL-validation [27,28]. Namely, TTL is an 8-bit field in a packet's IP header, and among its other important purposes can be used by the packet's receiving machine to compute the hop-distance to the packet's sending machine. By recording the TTL values of distinct source IP addresses over a period of time, the receiving machine can learn which values are expected from which particular host, and using this knowledge can subsequently be able to identify suspicious/spoofed packets. Now, the common initial values of TTL set by different operating systems are 64 and 128, 256. It can be easily confirmed that the initial TTL value specifically set by the machine(s) corresponding to IP 8.8.8.8 is 64. Consequently, a machine, server or a firewall set to perform TTL-validation on packets arriving from the key Internet servers (including DNS server 8.8.8.8) should be easily able to spot and flag a packet that was sent with the initial TTL = 40 as suspicious. According to [11], the average hop distance in the internet is 15.3 ± 4.2. Thus,

when receiving a packet from the DNS server 8.8.8.8, even the most distant destinations/machines would likely see a TTL ≥ 40. However, a spoofed packet with the initial TTL = 40 would arrive at its destination with the value likely well below 40, which (as previously stated) should be sufficient to flag this packet as anomalous, as per our last packet probe in Table 3.

Table 3. Spoofed IP packet probes used in Experiments 1 and 2.

Type of Packet	Different Spoofed Source IP Addresses Tested
IP packet carrying an ICMP Echo Request	1.2.3.4 172.16.1.100 8.8.8.8 Receiver's own IP (IP of the VPS running the listening script)
IP packets carrying an ICMP Echo Reply	172.16.1.100 8.8.8.8
IP packet with TTL=40 carrying an ICMP Echo Request	172.16.1.100 8.8.8.8

6.2. Experiment 1 Results: Transmission of Spoofed Packets From Select CSPs' Networks

The results of Experiment 1 (outlined in Figure 6) are summarized in Table 4, and they lead to quite an encouraging initial conclusion. Namely, out of the 14 evaluated Cloud service providers, initially (in April 2019), we were able to successfully transmit our spoofed IP packet probes from only three of them. Moreover, in mid-June 2019, during our repeated experimentations, two more CSPs ceased to allow transmission of spoofed IP packets out of their networks—king-servers.com and swiss-vps.com. Thus, presently, it is possible to successfully transmit spoofed IP packets only from one of the evaluated Cloud providers. Practically, this conclusion implies that the real-world feasibility of the attacks from Figure 3 is likely low, as the hacker interested in conducting these attacks would (statistically) face a serious challenge in finding a Cloud provider from which it would be possible to send spoofed IP packets.

Table 4. Results of Experiment 1.

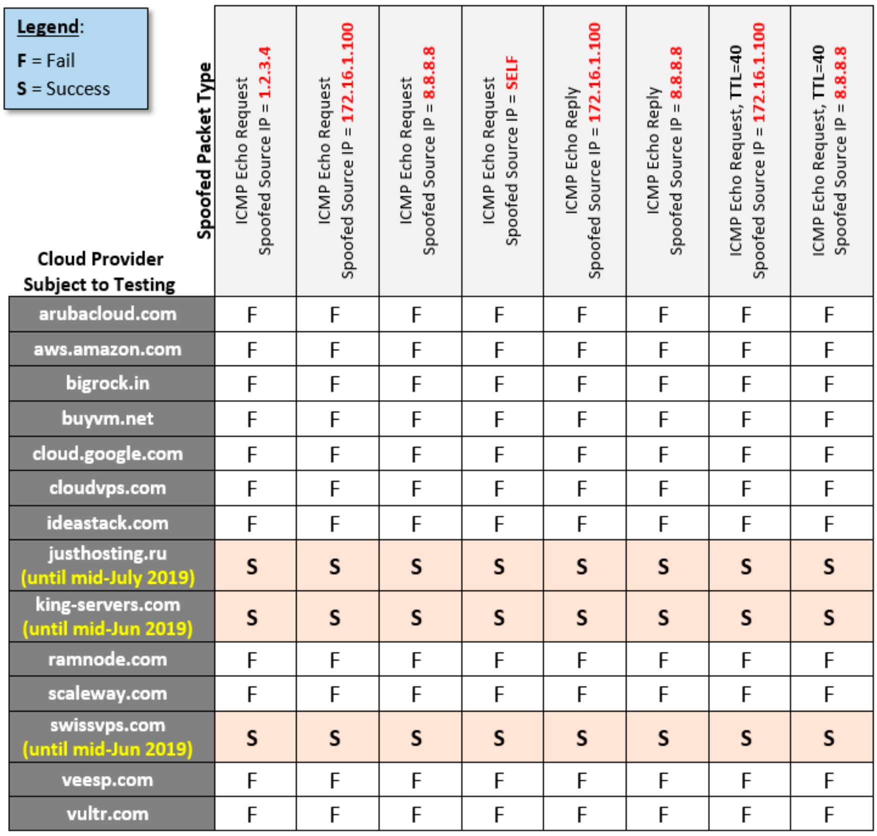

Legend:
F = Fail
S = Success

Spoofed Packet Type / Cloud Provider Subject to Testing	ICMP Echo Request Spoofed Source IP = 1.2.3.4	ICMP Echo Request Spoofed Source IP = 172.16.1.100	ICMP Echo Request Spoofed Source IP = 8.8.8.8	ICMP Echo Request Spoofed Source IP = SELF	ICMP Echo Reply Spoofed Source IP = 172.16.1.100	ICMP Echo Reply Spoofed Source IP = 8.8.8.8	ICMP Echo Request, TTL=40 Spoofed Source IP = 172.16.1.100	ICMP Echo Request, TTL=40 Spoofed Source IP = 8.8.8.8
arubacloud.com	F	F	F	F	F	F	F	F
aws.amazon.com	F	F	F	F	F	F	F	F
bigrock.in	F	F	F	F	F	F	F	F
buyvm.net	F	F	F	F	F	F	F	F
cloud.google.com	F	F	F	F	F	F	F	F
cloudvps.com	F	F	F	F	F	F	F	F
ideastack.com	F	F	F	F	F	F	F	F
justhosting.ru (until mid-July 2019)	S	S	S	S	S	S	S	S
king-servers.com (until mid-Jun 2019)	S	S	S	S	S	S	S	S
ramnode.com	F	F	F	F	F	F	F	F
scaleway.com	F	F	F	F	F	F	F	F
swissvps.com (until mid-Jun 2019)	S	S	S	S	S	S	S	S
veesp.com	F	F	F	F	F	F	F	F
vultr.com	F	F	F	F	F	F	F	F

6.3. Experiment 2 Results: Transmission of Spoofed Packets Into Select CSPs' Networks

The results of Experiment 2 (outlined in Figure 7) are summarized in Table 5, and they appear much more unexpected and potentially concerning than those of Experiment 1. We will analyze these results on a packet-by-packet basis.

(a) Transmission of ICMP Echo Requests with Spoofed Source IP = 1.2.3.4. These spoofed IP packets successfully reached our VPSs on 9 of 12 evaluated Cloud providers. However, as previously explained, 1.2.3.4 is an unassigned IP address, and the ingress firewall of any well-configured Cloud provider should be able to screen for and prevent these packets from entering the respective CSP's network.

(b) Transmission of ICMP Echo Requests with Spoofed Source IP = 172.16.1.100. These spoofed IP packets successfully reached our VPSs on 2 of 12 evaluated Cloud providers. However, recall that 172.16.1.100 is a private IP address, and again a well-configured ingress firewall should be able to screen for and prevent these packets from entering the respective Cloud provider's network.

(c) Transmission of ICMP Echo Requests with Spoofed Source IP = 8.8.8.8. As shown in Table 5, these spoofed IP packets successfully reached our VPSs on 11 of 12 evaluated Cloud providers, and (more importantly) eight of these VPSs automatically generated an ICMP Echo Response (back) to IP = 8.8.8.8—including the VPSs hosted by AWS. We believe this observation is particularly worrisome, for two reasons:

    (c.1) First, as previously explained, DNS servers are generally very unlikely to send/initiate ICMP requests—something a well-configured ingress firewall should be cognizant of. (ICMP requests are typically generated by client machines to probe a server or another client machine).

    (c.2) Moreover, by receiving these requests and responding to them, the servers of the examined Cloud providers could be potentially exploited as reflector points in an attack on the actual DNS server 8.8.8.8, as illustrated in Figure 5. (Given the critical importance of the DNS server 8.8.8.8, any attack on this server can have far-reaching consequences for the entire Internet, and should be prevented by all possible means and by all potentially involved parties).

(d) Transmission of ICMP Echo Requests with Spoofed Source IP = Receiver's Own (Destination) IP. Spoofed IP packets in this category successfully reached our VPSs on 9 of 12 evaluated Cloud providers; moreover, four of them also responded to the given request. This result is also very concerning, for two reasons:

    (d.1) A well-configured ingress firewall should not only be able to screen for and drop any incoming packet that carries the same source and destination IP address but should also be able to flag and drop any packet arriving from the wide-area Internet with an internal IP as its source address. Yet, the firewalls of nine tested Cloud providers failed to do so, including the ones of AWS and Google Cloud.

    (d.2) Machines or servers/VPSs that respond to ICMP requests with their own IP as both the source and destination address are particularly vulnerable to the so-called ICMP Land attack [16]. Our experimentation identified four such VPSs, again including those hosted by Google Cloud and AWS.

(e) Transmission of ICMP Echo Replies with Spoofed Source IP = 172.16.1.100. Spoofed IP packets in this category successfully reached our VPSs on two of the tested Cloud providers. A reason why these probes should have ideally been dropped by all Cloud providers, besides the fact that they carried a private IP address as their source, is the fact that these were entirely unsolicited Echo Replies—something any well-configured stateful firewall should ideally be able to pick on.

(f) Transmission of ICMP Echo Replies with Spoofed Source IP = 8.8.8.8. Spoofed IP packets in this category successfully reached our VPSs on nine of the tested Cloud providers. As in the case of (e), a well-configured stateful firewall should ideally be able to flag these probes as unsolicited/anomalous, and ideally drop them.

(g) Transmission of ICMP Echo Rsequests with Spoofed Source IPs of 172.16.1.100 and 8.8.8.8, and TTL = 40. Spoofed IP packets in these two categories successfully reached our VPSs on 2 and 11 of the tested Cloud providers, respectively. Here, it is particularly striking that all but one tested Cloud provider allowed the (spoofed) IP = 8.8.8.8 requests with clearly invalid TTL values into their network, this even though the TTL-validation technique is widely known and relatively simple to implement (as discussed in Section 6.1).

Table 5. Results of Experiment 2.

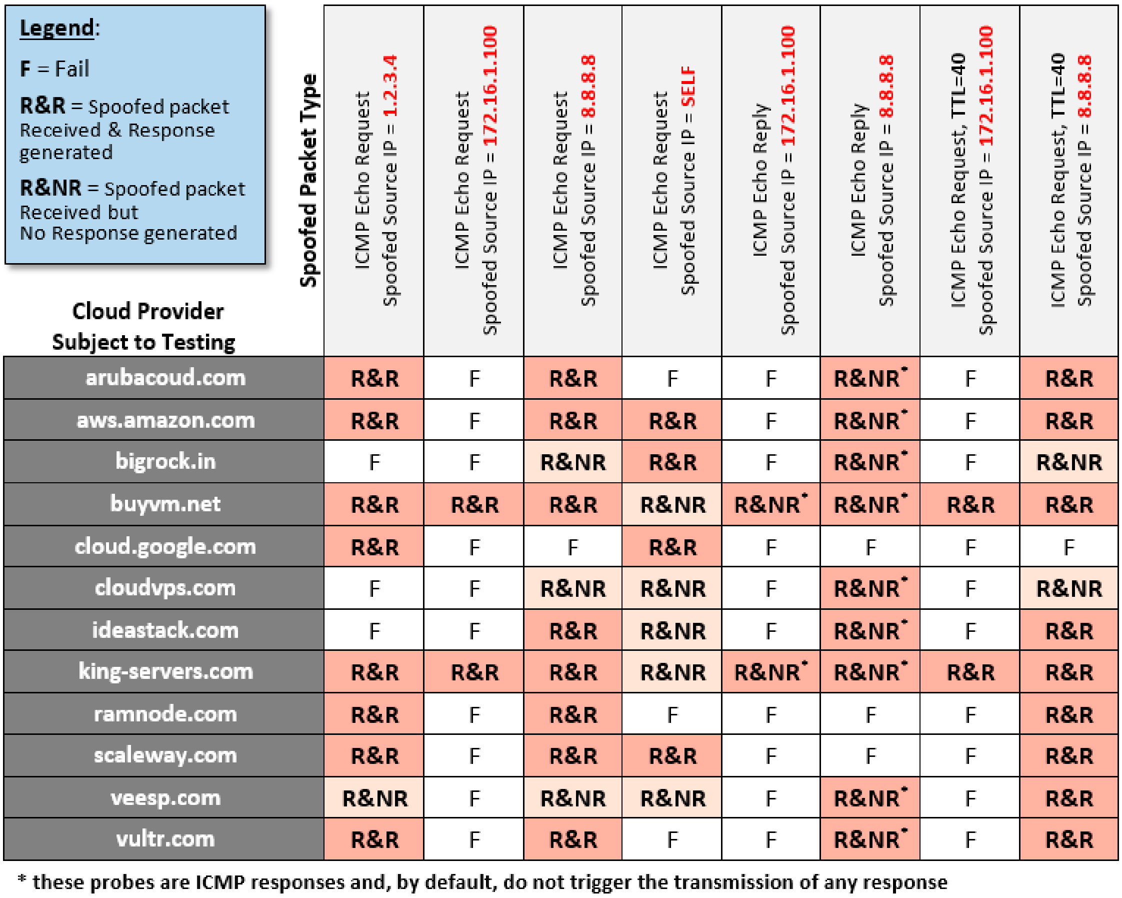

Legend:
F = Fail
R&R = Spoofed packet Received & Response generated
R&NR = Spoofed packet Received but No Response generated

Spoofed Packet Type / Cloud Provider Subject to Testing	ICMP Echo Request Spoofed Source IP = 1.2.3.4	ICMP Echo Request Spoofed Source IP = 172.16.1.100	ICMP Echo Request Spoofed Source IP = 8.8.8.8	ICMP Echo Request Spoofed Source IP = SELF	ICMP Echo Reply Spoofed Source IP = 172.16.1.100	ICMP Echo Reply Spoofed Source IP = 8.8.8.8	ICMP Echo Request, TTL=40 Spoofed Source IP = 172.16.1.100	ICMP Echo Request, TTL=40 Spoofed Source IP = 8.8.8.8
arubacoud.com	R&R	F	R&R	F	F	R&NR*	F	R&R
aws.amazon.com	R&R	F	R&R	R&R	F	R&NR*	F	R&R
bigrock.in	F	F	R&NR	R&R	F	R&NR*	F	R&NR
buyvm.net	R&R	R&R	R&R	R&NR	R&NR*	R&NR*	R&R	R&R
cloud.google.com	R&R	F	F	R&R	F	F	F	F
cloudvps.com	F	F	R&NR	R&NR	F	R&NR*	F	R&NR
ideastack.com	F	F	R&R	R&NR	F	R&NR*	F	R&R
king-servers.com	R&R	R&R	R&R	R&NR	R&NR*	R&NR*	R&R	R&R
ramnode.com	R&R	F	R&R	F	F	F	F	R&R
scaleway.com	R&R	F	R&R	R&R	F	F	F	R&R
veesp.com	R&NR	F	R&NR	R&NR	F	R&NR*	F	R&R
vultr.com	R&R	F	R&R	F	F	R&NR*	F	R&R

*** these probes are ICMP responses and, by default, do not trigger the transmission of any response**

In summary, the results of Experiment 2 have shown that a significant number of tested Cloud providers fail to protect their networks and servers from a range of incoming spoofed packets, yet, from the practical standpoint, each of these packets should be rather easy to detect and drop. This further implies that, in contrast to the attacks from Figure 3, the real-world feasibility of the attacks from Figures 4 and 5 may be rather high and warrants careful consideration by public Cloud providers.

Of course, we again would like to emphasize that our experiments involved the transmission of single packet probes towards the VPSs of the tested Cloud providers, as we did not want to engage in any actual attack-like activity. Consequently, there is a chance that sending large(er) groups of spoofed packets would have produced different outcomes. Put another way, one might argue that some public Cloud providers may be deliberately lenient towards sporadically encountered spoofed IP packets, given their limited threat capabilities; yet, those same Cloud providers could potentially have provisions to step-up their defenses when encountering larger volumes of spoofed packets. While this argument could be true, the problem is (1) there is no easy way of determining whether the real-world Cloud providers actually have such defense-mechanisms in place, so the purpose of our work was to point to a potential vulnerability, and (2) we believe there is no reasonable justification to allow certain undoubtedly problematic categories of spoofed packets (such as incoming packets that carry a local/internal, a private or a non-routable IP address as its source address) into any Cloud provider's network, especially if that Cloud provider is already monitoring the incoming traffic for groups of spoofed packets. Thus, spotting and dropping all 'obviously illegal' packets can be done at no additional cost.

One might also be tempted here to hypothesize here that the observed leniency towards incoming spoofed packets is not an 'oversight' but a 'deliberate choice' by the tested Cloud providers since, after all, spoofed packet can be used for purposes of legitimate network troubleshooting or performance testing (as mentioned in the introduction). However, when evaluating this hypothesis, one should keep in mind that in reality, a valid incoming packet could never originate from an internal, a private or an unassigned IP address (which our work mostly discusses), hence the use of spoofed packets with these specific source IP addresses in the evaluation of a network's performance would have very little (if any) practical sense and justification.

7. Conclusions and Future Work

In recent years we have witnessed a widespread adoption and reliance on public Cloud platforms and services, not only among the general public and companies but also among criminal cyber groups. In this paper, for the first time in the research literature, we have provided a systematic overview of different possible scenarios that involve the transmission of spoofed IP packets for the purposes of a direct attack on or an indirect misuse of the virtual and physical servers of a public Cloud provider. We have also pointed to the fact that, based on our survey of 35 real-world Cloud providers, over 50% of them fail to use their acceptable-use and terms-of-service policies as a potential deterrent against IP spoofing by their customer. Finally, we have presented the results of our experimentation with a number of real-world public Cloud providers. These results are bittersweet: On one side they show that, statistically, the majority of public Cloud providers do a very good job of preventing (potential) hackers from using their VPSs to launch spoofed-IP campaigns on third-party targets, but, on the other hand, they also show that VPSs of a number of public CSP could easily become a target of spoofed IP campaigns themselves, or be used as packet-reflectors in attacks on some critical points in the Internet infrastructure. We hope the findings of this study provide both cybersecurity researchers and practitioners with a more complete understanding of the addressed problems and assist them in building an overall safer Internet.

Our future work will include a more comprehensive study involving a larger number of public Cloud providers, a wider range of spoofed IP probes and a wider range of upper-layer protocols (e.g., we plan to experiment with spoofed IP probes that carry not only ICMP but also TCP or UDP payloads). We also are currently working on a publication that will offer a detailed survey of different techniques that could potentially be used by the public Cloud when dealing with spoofed IP traffic.

Author Contributions: Conceptualization, N.V., Methodology, N.V.; Software, M.C.; Validation, M.C., N.V., M.L.; Formal Analysis, N.V., M.C.; Investigation, M.C.; Resources, M.C., N.V., M.L.; Data Curation, M.C.; Writing-Original Draft Preparation, N.V.; Writing-Review & Editing, N.V., M.C., M.L.; Visualization, N.V., M.C.; Supervision, N.V.; Project Administration, M.C., N.V., M.L.; Funding Acquisition, N.V.

Funding: This research was funded by NSERC CREATE DITA Program.

Conflicts of Interest: Authors declare no conflict of interest.

References

1. Breeden, B.; Lclaughlin, M.; Sindicich, N.; Valentine, A. The C-SAFE Program and the Florida Cyber-Security Manual. Florida Department of Law Enforcement. November 2004. Available online: http://www.secureflorida.org/vendorimages/secureflorida2007/web/C-SAFE/CSAFEcybersecuritymanual.pdf (accessed on 1 November 2019).
2. MANRS–Mutually Agreed Norms for Routing Security. Available online: https://www.manrs.org/ (accessed on 28 October 2019).
3. CAIDA: Center for Applied Internet Data Analysis. Available online: https://www.caida.org/ (accessed on 28 October 2019).
4. ITPRO. Public Cloud Used to Power Supercharged DDoS Attacks. September 2018. Available online: https://www.itpro.co.uk/public-cloud/31884/public-cloud-used-to-power-supercharged-ddos-attacks#gref (accessed on 28 October 2019).
5. Pohle, T. Public Cloud Services Increasingly Exploited to Supercharge DDoS Attacks: New Link11 Research. September 2018. Available online: https://www.link11.com/en/blog/public-cloud-services-increasingly-exploited-to-supercharge-ddos-attacks-new-link11-research/ (accessed on 28 October 2019).
6. Cimpany, C. Operator of Eight DDoS-for-Hire Services Pleads Guilty. February 2018. Available online: https://www.zdnet.com/article/operator-of-eight-ddos-for-hire-services-pleads-guilty/ (accessed on 28 October 2019).
7. Cloud Security Alliance (CSA). The Treacherous 12: Cloud Computing Top Threats in 2016. February 2016. Available online: https://downloads.cloudsecurityalliance.org/assets/research/top-threats/Treacherous-12_Cloud-Computing_Top-Threats.pdf (accessed on 28 October 2019).
8. Singh, S.; Jeong, Y.-S.; Park, J.H. A survey on cloud computing security: Issues, threats, and solutions. *J. Netw. Comput. Appl.* **2016**, *75*, 200–222. [CrossRef]
9. Ali, M.; Khan, S.U.; Vasilakos, A.V. Security in cloud computing: Opportunities and challenges. *J. Inf. Sci.* **2015**, *305*, 357–383. [CrossRef]
10. Subramanian, N.; Jeyaraj, A. Recent security challenges in cloud computing. *J. Comput. Electr. Eng.* **2018**, *71*, 28–42. [CrossRef]
11. Kumar, R.; Goyal, R. On cloud security requirements, threats, vulnerabilities and countermeasures: A survey. *J. Comput. Sci. Rev.* **2019**, *33*, 1–48. [CrossRef]
12. De Donno, M.; Giaretta, A.; Dragoni, N.; Bucchiarone, A.; Mazzara, M. Cyber-Storms Come from Clouds: Security of Cloud Computing in the IoT Era. *MDPI J. Future Internet* **2019**, *11*, 127. [CrossRef]
13. Rathore, S.; Park, J.H. Semi-supervised learning based distributed attack detection framework for IoT. *J. Appl. Soft Comput.* **2018**, *72*, 79–80. [CrossRef]
14. Somani, G.; Gaur, M.S.; Sanghi, D.; Conti, M.; Buyya, R. DDoS attacks in cloud computing: Issues taxonomy, and fugure directions. *Elseiver Comput. Commun. J.* **2017**, *107*, 30–48. [CrossRef]
15. Osanaiye, Q.; Choo, K.-K.R.; Dlodlo, M. Distributed denial of service (DoS) resilience in cloud: Review and conceptual cloud DDoS mitigation framework. *J. Netw. Comput. Appl.* **2016**, *67*, 147–165. [CrossRef]
16. Osanaiye, O.A.; Dlodlo, M. TCP/IP header classification for detecting spoofed DDoS attack in Cloud environment. In Proceedings of the EUROCON, Salamanca, Spain, 8–11 September 2015.
17. Hong, J.B.; Nhlabatsi, A.; Kim, D.S.; Hussein, A.; Fetais, N.; Khan, K.M. Systematic identification of threats in the cloud: A survey. *J. Comput. Netw.* **2018**, *150*, 46–69. [CrossRef]
18. Singh, G.K.; Somani, G. Cross-VM Attacks: Attack Taxonomy, Defence Mechanisms, and New Directions, Part of Versatile Cybersecurity-Advances in Information Security book series (ADIS). *Springer* **2018**, *72*, 257–286.
19. Agrawal, N.; Tapaswi, S. A Lightweight Approach to Detect the Low/High Rate IP Spoofed Cloud DDoS Attacks. In Proceedings of the IEEE International Symposium on Cloud and Service Computing (IEEE SC2), Kanazawa, Japan, 22–25 November 2017.
20. Osanaiye, O.A. Short Paper: IP spoofing detection for preventing DDoS attack in Cloud Computing. In Proceedings of the IEEE International Conference on Intelligence in Next Generation Networks (IEEE ICIN), Paris, France, 17–19 February 2015.
21. Yao, G.; Bi, J.; Vasilakos, A.V. Passive IP Traceback: Disclosing the Locations of IP Spoofers from Path Backscatter. *IEEE Trans. Inf. Forensics Secur.* **2015**, *10*, 471–484. [CrossRef]

22. Chang, R.K.C. Defending against Flooding-Based Distributed Denial-of-Service Attacks: A Tutorial. *IEEE Commun. Mag.* **2002**, *40*, 42–51. [CrossRef]
23. Netscout. 14th Annual Worldwide Infrastructure Security Report. March 2019. Available online: https://www.netscout.com/press-releases/netscout-releases-14th-annual-worldwide-infrastructure (accessed on 28 October 2019).
24. Cimpanu, C. Russia Bans 1.8 Million Amazon and Google IPs in Attempt to Block Telegram April 2018. Available online: https://www.bleepingcomputer.com/news/government/russia-bans-18-million-amazon-and-google-ips-in-attempt-to-block-telegram/ (accessed on 28 October 2019).
25. Valls-Prieto, J. *Handbook of Research on Digital Crime, Cyberspace Security, and Information Assurance*; IGI Global: Hershey, PA, USA, 2014.
26. Herath, T.; Rao, H.R. Protection motivation and deterrence: A framework for security policy. *Eur. J. Inf. Syst.* **2009**, *18*, 106–125. [CrossRef]
27. Beverly, R.; Berger, A.; Hyun, Y. Understanding the Efficacy of Deployed Internet Source Validation Filtering. In Proceedings of the 9th ACM SIGCOMM Conference, Chicago, IL, USA, 4–6 November 2009; pp. 356–369.
28. Wang, H.; Jin, C.; Shin, K.G. Defense against Spoofed IP Traffic Using Hop-Count Filtering. *IEEE/ACM Trans. Netw.* **2007**, *15*, 40–53. [CrossRef]