Recommended Paper

A Comprehensive Measurement of Cloud Service Abuse

Naoki Fukushi[1,†1,a)] Daiki Chiba[2,b)] Mitsuaki Akiyama[2,c)] Masato Uchida[1,d)]

Received: May 7, 2020, Accepted: October 6, 2020

Abstract: Cloud services are maliciously used as an infrastructure for cyber-attacks. In a cloud service, the assigned Internet Protocol (IP) address for a server is owned by the cloud service provider. When the server is shut down, the assigned IP address is freed for reuse and assigned to another server in the same cloud service. Cyber-attackers abusing cloud services in this way therefore pose a serious risk since legitimate service providers, developers, and end users may be mistakenly blacklisted which lowers the image and hurts the reputation of the service. In this study, we conducted a large-scale measurement of cloud service abuse by using blacklisted IP addresses. Our analysis of four cloud services over 154 days using 39 blacklists revealed that a total of 61,060 IP addresses from these cloud service providers were blacklisted, approximately 14,000 IP addresses continue to be blacklisted, and approximately 5% are replaced daily. Moreover, our study revealed trends in attacks that abuse cloud services with respect to attack type, region, duration, and anti-abuse countermeasures. Finally, we discuss recommendations for cloud service users, cloud service providers, and blacklist providers.

Keywords: cloud, IP address, blacklist

1. Introduction

Cloud services are widely used on the Internet and have currently established themselves as a new infrastructure. Cloud services are designed to provide users with a required amount of computing resources, such as servers, storage, and applications which are owned by the cloud service provider. Due to the considerable convenience they provide, the global market for cloud services continues to expand rapidly. It has been shown that the size of the global market for cloud services increased by 37% to $27.5 billion in the third quarter of 2019 from $20.2 billion in the previous year [2]. However, cloud services can also be abused as an infrastructure for cyber-attacks. For example, a large number of cloud servers were used in a brute-force attack to hijack Instagram accounts [3], wherein the servers on the cloud were used as a command and control (C&C) server [4]. In this paper, we refer to cyber-attacks that abuse cloud services as "cloud service abuse."

Currently, service providers, developers, and end users are relying on the cloud as an infrastructure for their services and businesses. In fact, according to the Cisco's report [5], global cloud Internet Protocol (IP) traffic will account for 95% of the total data center traffic by 2021. Besides, Flexera's survey that focused on technical professionals representing organizations of various sizes spanning numerous industries, revealed that 94% of them use cloud services [6]. Therefore, cloud service abuse poses significant risks for both legitimate cloud service users and cloud service providers. For example, in the cloud-based email-sending service Amazon Simple Email Service (SES), a legitimate user could not send emails because the IP address assigned to the user was blacklisted [7]. This is because the cloud service provider reused the IP address among multiple users within a short time and some of the users were involved in cyber-attacks. Thus, there is a risk that legitimate service providers, developers, and end users will be subjected to various restrictions if IP addresses that have been blacklisted for malicious activity in the past are assigned to the servers that they use. There is also the resultant risk of damage to the reputation of the cloud service provider.

To the best of our knowledge, this study represents the first large-scale analysis of cloud service abuse to determine the actual situation and effective countermeasures. To conduct this analysis, we required a method for the large-scale observation of cloud service abuse involving various types of cyber-attacks. To perform this, we focused on using different types of blacklists in combination. We then conducted a large-scale analysis of cloud service abuse of four large cloud service providers: Amazon Web Service Elastic Compute Cloud (AWS), Microsoft Azure (Azure), Google Cloud Platform (GCP), and Oracle Cloud (Oracle). In our analysis using 39 blacklists for a duration of 154 days, a total of 61,060 blacklisted IP addresses from these cloud service providers were observed.

We also discovered five different aspects of cloud service abuse: (1) changes in the number of blacklisted cloud IP ad-

[1] Waseda University, Shinjuku, Tokyo 169–8555, Japan
[2] NTT Secure Platform Laboratories, Musashino, Tokyo 180–8585, Japan
[†1] Presently with NTT Secure Platform Laboratories
[a)] naoki.fukushi.kz@hco.ntt.co.jp
[b)] daiki.chiba@ieee.org
[c)] akiyama@ieee.org
[d)] m.uchida@waseda.jp

This paper is an extended version of a paper published in IEEE Conference on Communications and Network Security (CNS) 2020 [1].
The preliminary version of this paper was published at anti Malware engineering WorkShop (MWS 2019), October 2019. The paper was recommended to be submitted to Journal of Information Processing (JIP) by the program chair of MWS 2019.

dresses over time, (2) types of attacks involved in cloud service abuse, (3) trends regarding IP address regions, (4) differences in on-list duration of the blacklisted IP addresses depending on the attack type and cloud service provider, and (5) the status of the deregistration of blacklisted cloud IP addresses. These findings suggest that cloud service providers need to detect abuse of their services early and take appropriate countermeasures.

The contributions of this study are as follows.

- We conducted the first large-scale analysis of cloud service abuse and revealed the actual situation.
- We proposed an observational method for cloud service abuse using diverse blacklists.

The rest of this paper is organized as follows. We explain cloud services and blacklists in Section 2. In Section 3, we describe our observation method for cloud service abuse using diverse blacklists and the viewpoints of our measurement. We discuss the results from our measurement in Section 4 and the recommendations for cloud service users, providers, and blacklist providers and the limitations of our study in Section 5. In Section 6, we summarize related work. Finally, we conclude our paper in Section 7.

2. Cloud Services and Blacklists

Figure 1 shows the relationship between cloud service users (e.g., attackers, legitimate service providers, developers, and end users), cloud service providers, and blacklist providers in situations where cloud service abuse occurs. The threat model in this paper is as follows. When using cloud services, (1) the IP address of the cloud service provider is assigned to the server. In this situation, (2) an attacker abuses the cloud service to carry out cyber-attacks. (3) Blacklist providers observe these attacks. Then, (4) they add the IP address of the cloud service provider to their blacklists. When the attacker stops using the cloud service, (5) the IP address assigned to the server is returned to the cloud service provider. Because IP addresses are reused among users in cloud services, (6) the blacklisted IP address may be assigned to a server used by legitimate service providers, developers, and end users. (7) This results in false restrictions, which means that legitimate service providers, developers, and end users cannot communicate normally because their IP addresses are blacklisted. In this section, we describe the cloud services (Section 2.1) and the blacklists (Section 2.2).

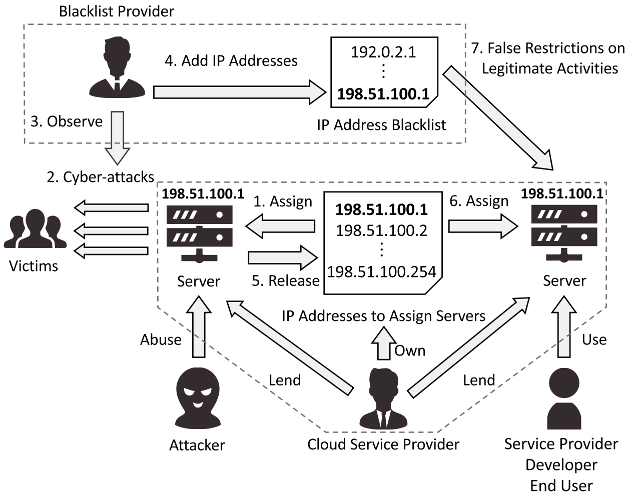

Fig. 1 Stakeholders of cloud service abuse.

2.1 Cloud Services

Cloud services are available with lower initial investment, maintenance, and operational costs to users compared to physical hardware. In general, cloud service providers have data centers in multiple regions around the world. When using cloud services, the users can select the region where they want the computing resources they are using to be located. The prices and range of IP addresses assigned to the server vary depending on the selected region. IP addresses in cloud services are reused between users. Thus, the IP address assigned to the server is freed for reuse when the server is shut down and is assigned to another server.

There are numerous cloud service providers. In this study, we focused on AWS, Azure, GCP, and Oracle for our measurements. There are two reasons for selecting these four cloud services. First, these are typical/popular cloud services and are considered to be abused by more attackers. Second, in these cloud services, the range of public IP addresses assigned to the server is available to the public [8], [9], [10], [11]. This allowed us to determine whether a given IP address was used by these cloud services.

2.2 Blacklists

A blacklist is a list of IP addresses that have been found to be involved in malicious activity. Blacklists are used to identify communications that use the blacklisted IP addresses as their source or destination. Blacklists are not static but are updated regularly by the blacklist provider. However, the update time and intervals differ for each blacklist. In addition, the period for which a blacklisted IP address remains on the blacklist differs depending on the blacklist. This is because blacklist providers have different respective listing policies. We can obtain the policy information when the blacklist provider offers it along with the blacklist. To collect as many blacklisted IP addresses as possible when multiple listing policies are offered for the same blacklist, we acquired the blacklist with the longest one (up to 30 days).

Some blacklists allow third parties to apply for the deregistration of blacklisted IP addresses. There are two main reasons why this is possible. One is that the IP addresses of users who did not originally perform the malicious activities might be mistakenly blacklisted. Another reason is that the blacklisted IP address may have already stopped performing malicious activities due to subsequent countermeasures.

3. Measurement Method

We conducted a large-scale analysis of cloud service abuse of four popular cloud services: AWS, Azure, GCP, and Oracle. For this analysis, it was necessary to have an environment where large-scale, continuous, and direct observation of cloud service abuse involving various types of attacks was possible. However, it is extremely difficult to prepare such an environment. For example, honeypots and darknets are not suitable for observing cloud service abuse. Observation using only darknets or honeypots may result in a biased observation range, or observable attack types

Table 1 Provider name, number of blacklists from each provider, total number of unique blacklisted IP addresses, and listing policy information.

#	Provider Name	# Lists	# IP Addrs	Policy	#	Provider Name	# Lists	# IP Addrs	Policy
1	Badips	11	981,266	30 days	12	Feodo	1	2,061	-
2	Fail2ban	7	473,875	2 days	13	Haley	1	62,920	-
3	ProjectHoneyPot	3	12,537	30 days	14	LashBack	1	1,492,213	-
4	AlienVault	1	408,943	-	15	MyIP	1	7,017	10 days
5	Bambenek	1	4,632	-	16	NixSpam	1	493,669	12 hours
6	BinaryDefense	1	35,472	-	17	Nothink	1	32,560	-
7	BotScout	1	85,201	30 days	18	Sblam	1	51,861	-
8	CleanTalk	1	520,401	30 days	19	StopForumSpam	1	287,642	30 days
9	CyberCrime	1	2,064	-	20	Talos	1	3,109	-
10	Dangerrulez	1	5,336	30 days	21	VoIPBL	1	95,625	-
11	DShield	1	52,094	30 days					

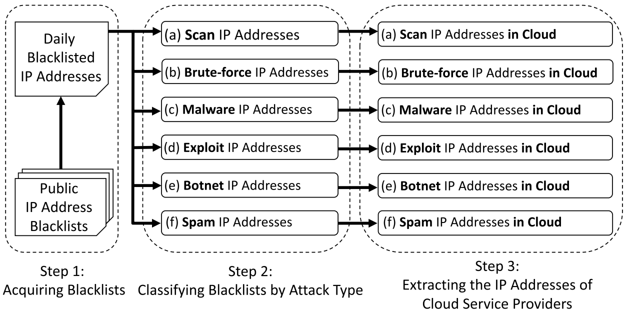
Daily Blacklisted IP Addresses
Public IP Address Blacklists
(a) Scan IP Addresses
(b) Brute-force IP Addresses
(c) Malware IP Addresses
(d) Exploit IP Addresses
(e) Botnet IP Addresses
(f) Spam IP Addresses
(a) Scan IP Addresses in Cloud
(b) Brute-force IP Addresses in Cloud
(c) Malware IP Addresses in Cloud
(d) Exploit IP Addresses in Cloud
(e) Botnet IP Addresses in Cloud
(f) Spam IP Addresses in Cloud
Step 1: Acquiring Blacklists
Step 2: Classifying Blacklists by Attack Type
Step 3: Extracting the IP Addresses of Cloud Service Providers

Fig. 2 Flow of our observation method for cloud service abuse.

Table 2 Attack type, number of corresponding blacklists, and total number of unique blacklisted IP addresses.

Attack Type	# Lists	# IP Addrs
(a) Scan	2	86,572 (2.0%)
(b) Brute-force	9	546,881 (12%)
(c) Malware	3	8,757 (0.20%)
(d) Exploit	10	381,743 (8.6%)
(e) Botnet	6	192,740 (4.4%)
(f) Spam	9	3,213,420 (73%)
Total	39	4,430,113 (100%)

may be limited to brute-force attacks and scan attacks.

In this study, we focused on IP address blacklists to solve this problem. By integrating different types of blacklists, both the observation range and observable attack types were increased. In addition, since blacklists are updated regularly, it is easy to conduct a time series analysis. Conventional studies on blacklists [12], [13], [14] have focused on the accuracy and characteristics of blacklists themselves. We used multiple blacklists as the most practical method for broadly observing cloud service abuse without direct observation. This is a significant difference compared to previous studies on blacklists and also signifies the technical importance of our measurement method.

3.1 Observation of Cloud Service Abuse

This section describes the observation method for cloud service abuse using blacklists. **Figure 2** shows the flow of our observation method. The method consists of three steps: acquiring blacklists, classifying blacklists by attack type, and extracting the IP addresses of cloud service providers.

Step 1: Acquiring Blacklists. A total of 39 public blacklists were acquired from 21 different blacklist providers once per day at the same time. After considering results from investigating the update frequency of each blacklist provider, we decided to acquire these blacklists once a day. The acquisition period lasted 154 days from June 30, 2019 to November 30, 2019. The blacklists acquired in this study were also used in Refs. [14], [15], which compared and analyzed the characteristics of multiple IP address blacklists. We selected blacklists that were continuously updated during the entire acquisition period.

Table 1 summarizes the names of the blacklist providers, the number of blacklists from each provider, the total number of unique blacklisted IP addresses, and listing policy information found from the acquired blacklists. Not all blacklists make policy information available to the public. In Table 1, if policy information is not available, it is shown as "-". Note that three blacklist providers created multiple blacklists, whereas the remaining providers each created one blacklist.

Step 2: Classifying Blacklists by Attack Type. Next, we describe the procedure for classifying the acquired blacklists according to attack type. The 39 acquired blacklists were classified into 6 attack types based on explanations from the blacklist providers: (a) Scan, (b) Brute-force, (c) Malware, (d) Exploit, (e) Botnet, and (f) Spam. These six attack types are defined in Ref. [14]. This categorization enabled us to conduct an analysis that considers what type of attack each blacklisted IP address performed. Here, hosts performing port or vulnerability scans are classified as (a) Scan; hosts making brute-force login attempts are classified as (b) Brute-force; malware C&C and distribution servers are classified as (c) Malware; hosts attempting to remotely exploit vulnerabilities are classified as (d) Exploit; compromised hosts belonging to a botnet are classified as (e) Botnet; and hosts sending spam are classified as (f) Spam. For example, a blacklist with the explanation "IPs launching SSH dictionary attacks" is classified as (b) Brute-force, whereas a blacklist with the explanation "IP addresses that sent spam" is classified as (f) Spam. **Table 2** lists the number of blacklists classified into each attack type and the total number of unique blacklisted IP addresses over the observation period from all classified blacklists for each attack type.

Step 3: Extracting the IP Addresses of Cloud Service Providers. Finally, we describe the procedure for extracting the IP addresses of cloud service providers from the blacklisted IP addresses. As explained in Section 2.1, each of the four cloud services selected for this study releases a range of public IP addresses that can be assigned to servers lent to users [8], [9], [10], [11]. We compared and matched the blacklisted IP addresses and the above range of public IP addresses. As a result, the blacklisted IP addresses of the cloud service providers were extracted.

By observing the blacklisted cloud IP addresses, we were able to indirectly observe the occurrence of cloud service abuse. In this study, we investigated only IPv4 addresses because the blacklists acquired in this study contained only a very small number of IPv6 addresses. The total number of unique blacklisted IPv6 addresses was 987 which was approximately 0.025% of the total number of unique blacklisted IPv4 and IPv6 addresses.

3.2 Viewpoints of Analysis

We analyze trends in cloud service abuse cases using blacklists based on the following technique. We integrated the 39 blacklists acquired on the same day into one list. The IP addresses in the integrated list were blacklisted in at least one of 39 blacklists that was used for integration. Since the blacklists were acquired over 154 days, 154 integrated lists were created. Using these 154 integrated lists, we investigated when and which cloud IP address was blacklisted and the attack type of the blacklists in which it was registered. This integrated list is useful for identifying the actual situation of cloud service abuse.

We analyzed cloud service abuse from the following seven viewpoints.

- How many IP addresses are abused per day and how do they change over time? (Section 4.1)
- Is there a difference in the number of abused IP addresses depending on the cloud service provider and type of attack? (Section 4.1 and Section 4.2)
- Are there any regional characteristics in the abused IP addresses? (Section 4.3)
- How long do abused IP addresses stay on the blacklist? (Section 4.4)
- Can we observe applications from cloud service providers or users for the deletion of IP addresses from blacklists? (Section 4.5)
- Did the IP addresses registered in the blacklists actually perform cyber-attacks? (Section 4.6)
- What positive effects did the integration of multiple blacklists have on the observation of cloud service abuse? (Section 4.7)

4. Measurement Results

4.1 Number and Changes of Blacklisted IP Addresses

For each cloud service provider, we investigated the number of IP addresses registered in at least one of the 39 acquired blacklists. On average, 7,335 AWS, 3,108 Azure, 3,416 GCP, and 330 Oracle IP addresses were blacklisted per day during the observation period. Moreover, on average, 401 AWS, 110 Azure, 157 GCP, and 12 Oracle IP addresses were replaced per day. In other words, approximately 14,000 cloud IP addresses were blacklisted and about 5% were replaced per day. **Figure 3** shows a graph of date versus the number of cloud IP addresses registered in the blacklists acquired on that date. The change in the number of blacklisted IP addresses over time differs for each cloud service provider. The number of blacklisted AWS IP addresses showed roughly two gradual decreases and one gradual increase from the observation start date of June 30 to the end of October, and then increased significantly after October. Additionally, throughout the observation period, the number of blacklisted Azure IP addresses increased slowly, the number of blacklisted GCP IP addresses changed slightly, and the number of blacklisted Oracle IP addresses remained almost unchanged. To clarify the reasons for the increase or decrease in the number of blacklisted IP addresses, we need to know the full extent of the attacker's behavior. This is impractical and the lack of further clarification is a limitation on our work. On the other hand, both cloud service providers and users should be more cautious while the number of blacklisted IP addresses is increasing such as from the end of October 2019 for AWS. The total number of unique blacklisted IP addresses over the observation period was 37,500 for AWS, 8,806 for Azure, 13,660 for GCP, and 1,094 for Oracle.

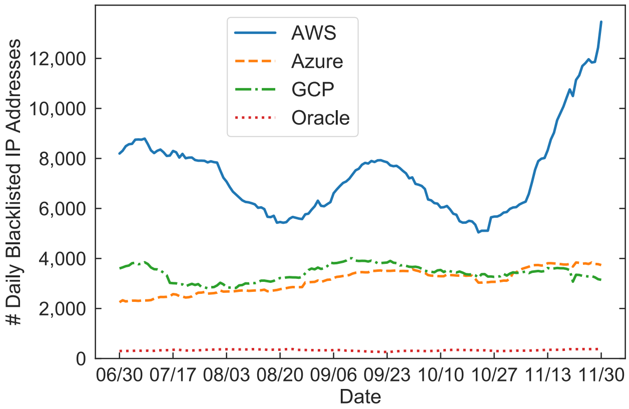

Fig. 3 Change in the number of daily blacklisted IP addresses.

4.2 Type of Attack

In Section 4.1, we did not consider the type of attack. In this section, we conducted an analysis of blacklisted IP addresses based on the type of attack. We investigated the number of unique blacklisted IP addresses for each attack type and the proportion of these numbers with respect to the total number for all attack types. The results are listed in **Table 3**. Several IP addresses from each cloud service provider were registered in the blacklists of (b) Brute-force and (f) Spam, which accounted for 71% to 85% of the total blacklisted IP addresses. This suggests that brute-force attacks and spam-sending are the most common attacks in these cloud services. This trend follows the percentage of blacklisted IP addresses by attack type shown in Table 2. In contrast, each cloud service provider differed in terms of the proportions of each attack type. For example, in Azure, the proportion of IP addresses used for (b) Brute-force is larger than for other cloud services. However, for (f) Spam it is smaller than for the other cloud services.

The total number of blacklisted IP addresses in Table 3 does not coincide with the total number of unique blacklisted IP addresses shown in Section 4.1. This number mismatch occurs because some IP addresses were registered in multiple attack types of blacklists. These IP addresses are considered to be associated with multiple types of cyber-attacks. The number of these IP addresses was 2,467 for AWS, 1,682 for Azure, 2,049 for GCP, and 206 for Oracle. Several such IP addresses were registered in (f) Spam blacklists and other blacklists such as (d) Exploit and (b) Brute-force. This fact suggests that cloud service abuse involv-

Table 3 Number of blacklisted IP addresses for each attack type.

Attack Type	#	AWS	Azure	GCP	Oracle
(a) Scan	2	475 (1.2%)	182 (1.7%)	297 (1.8%)	26 (2.0%)
(b) Brute-force	9	10,458 (26%)	3,634 (34%)	4,320 (27%)	368 (28%)
(c) Malware	3	301 (0.75%)	103 (0.96%)	95 (0.59%)	1 (0.076%)
(d) Exploit	10	3,499 (8.7%)	2,137 (20%)	2,167 (13%)	253 (19%)
(e) Botnet	6	1,907 (4.7%)	675 (6.3%)	1141 (7.1%)	16 (1.2%)
(f) Spam	9	23,607 (59%)	4,001 (37%)	8,046 (50%)	660 (50%)
Total	39	40,247 (100%)	10,732 (100%)	16,066 (100%)	1,324 (100%)

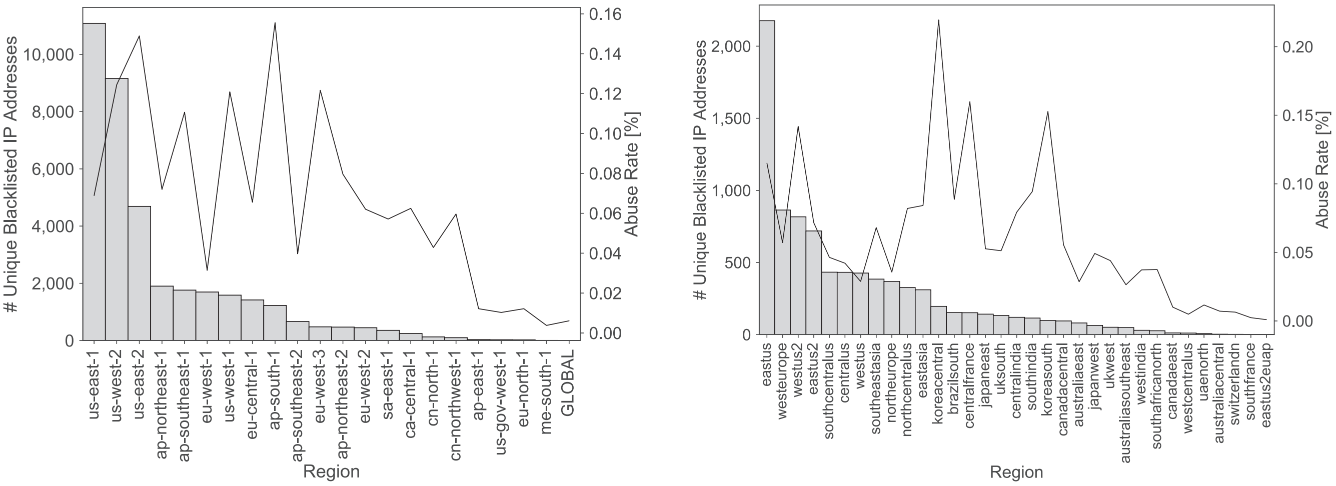

(a) AWS

(b) Azure

Fig. 4 Number of unique blacklisted IP addresses and abuse rate for each region.

ing multiple types of cyber-attacks tends to exploit vulnerabilities remotely or perform brute-force attacks in addition to sending spam.

4.3 IP Address Regions

In this section, we investigate the regions of the blacklisted IP addresses from the cloud service providers. We selected AWS and Azure as the cloud service providers to use for analyzing regions because these two providers publish information linking IP addresses and regions and have a sufficient number of blacklisted IP addresses and regions. We investigated 37,500 AWS and 8,806 Azure IP addresses which was the total number of unique blacklisted IP addresses over the observation period clarified in Section 4.1. In addition to the number of blacklisted IP addresses for each region, we investigated the proportion of these numbers to the total number of IP addresses of each region. In this paper, this proportion is called the IP address "abuse rate" for each region. The total number of IP addresses in each region was derived based on the following procedure. First, we expanded the IP address range (CIDR notation) of each region. We then calculated the total number of expanded IP addresses.

Figure 4 shows the graphs of region versus the number of unique blacklisted IP addresses and the abuse rate for AWS and Azure. The regions on the horizontal axis in Fig. 4 had at least one unique blacklisted IP address. Both AWS and Azure have a bias in the number of unique blacklisted IP addresses among all regions. In AWS, US regions, except for *us-west-1*, are available at low prices [16]. Therefore the attackers probably selected inexpensive regions the same as done by legitimate users. Similarly, in Azure, there are several blacklisted IP addresses in US regions with low prices. However, the number of blacklisted IP addresses associated with *westeurope*, which has a relatively high price, is also large or namely the regions are not selected based only on price. Here, we referred to Ref. [17] for the price in each region in Azure.

Similar to the number of unique blacklisted IP addresses, there is also a bias in the abuse rate of each region. We can see a relatively strong positive correlation between the number of unique blacklisted IP addresses and the abuse rate in both AWS and Azure. In contrast, some regions have high abuse rates relative to the number of unique blacklisted IP addresses, for example, *ap-south-1* (India) and *eu-west-3* (France) for AWS and *centralfrance* (France), *koreacentral*, and *koreasouth* (Korea) for Azure. In these regions, the number of unique blacklisted IP addresses is small, however, the risk of these IP addresses being assigned to a legitimate user is high.

4.4 Probability that IP Address Continues to be Blacklisted

In this section, we define the probability that an IP address will continue to be on a blacklist for M days after first being listed as "on-list duration probability for M days." In our analysis, for example, if an IP address is blacklisted for one day and then blacklisted again 30 days later for one day, we assess this as two different IP addresses are blacklisted for one day each. In this example, for the type of attack involving this IP address, we observed a tendency towards a low on-list duration probability. We analyzed the on-list duration probability using the Kaplan–Meier method [18]. In this study, the observation period was limited to 154 days. Using the Kaplan–Meier method, we can calculate the on-list duration probability even under this condition. We calcu-

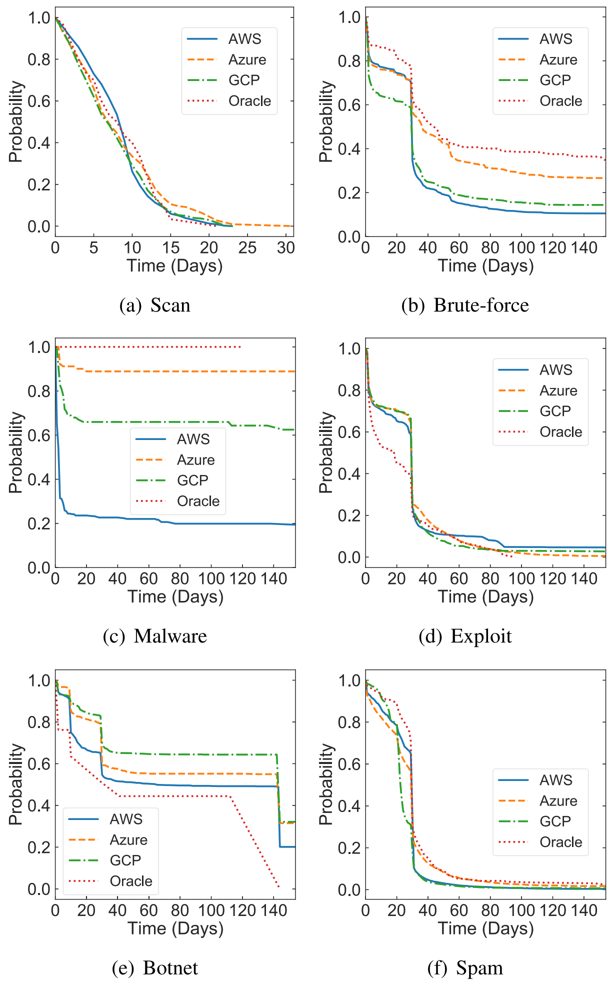

(a) Scan

(b) Brute-force

(c) Malware

(d) Exploit

(e) Botnet

(f) Spam

Fig. 5 On-list duration probability of blacklisted IP addresses for each attack type and cloud service provider.

lated the on-list duration probability for each type of cyber-attack and for each cloud service provider.

The six graphs in **Fig. 5** show the results of the Kaplan–Meier analysis. For each graph, the vertical axis indicates the on-list duration probability (complementary cumulative distribution) and the horizontal axis indicates the elapsed days since the IP address was first blacklisted. If the on-list duration probability remains high as time passes, the cloud IP addresses continued to be blacklisted for a long time, which means that the attacker continuously used them for cyber-attacks or that there were no applications to deregister them once they were no longer used for cyber-attacks. In any case, these cloud service providers need to take adequate countermeasures to deal with abuse of their services.

From Fig. 5, we can confirm that the on-list duration probability is different for each type of cyber-attack and each cloud service provider. For example, the on-list duration probability of (a) Scan decreases steadily and reaches almost 0 after 20 days. In contrast, the on-list duration probability of (c) Malware decreases in the first few days but remains high with only a little decline from then onwards, and that of (e) Botnet decreases slowly and converges to a lower limit in the range of 0.2 to 0.4 except for Oracle, and those for (b) Brute-force, (d) Exploit, and (f) Spam decrease significantly after 30 days. Moreover, for (b) Brute-force, AWS and GCP have lower on-list duration probabilities than Azure and Oracle, and AWS for (c) Malware and Oracle for (e) Botnet have a lower on-list duration probability compared with the other cloud service providers. The early decline in the on-list duration probability indicates that the blacklists had a short policy-defined period. Another point is that countermeasures of cloud service providers may have been successful, or attackers may have changed the IP address for misuse or attacks within a short period of time. To understand this, we need both the internal countermeasure status of cloud service providers and information on the attackers. The lack of access to this information is a limitation on our work.

Table 4 Number of blacklisted IP addresses removed before the policy-defined period.

	AWS	Azure	GCP	Oracle	Total
# IP Addresses	2	512	2	1	517

Here, we analyzed why the on-list duration probability significantly decreases for (b) Brute-force, (d) Exploit, and (f) Spam for elapsed times greater than 30 days. Several of the blacklist providers with IP addresses classified into these attack types specify a policy wherein IP addresses remain on the blacklist for at least 30 days after the final observation of an attack. If the attack is observed again M ($\leq$ 30) days after the first blacklisting, the blacklisting period will continue for at least $M + 30$ days. In other words, cyber-attacks were not observed after the first blacklisting, and several IP addresses were immediately removed from the blacklists once the blacklisting period exceeded 30 days according to this policy.

4.5 Deregistration of Blacklisted Cloud IP Addresses

We also investigated the deregistration of blacklisted cloud IP addresses. A legitimate user or cloud service provider can apply for the deregistration of blacklisted cloud IP addresses that are no longer involved in cloud service abuse. By investigating the status of deregistration of blacklisted cloud IP addresses, we can estimate the current status of countermeasures against cloud service abuse. If the blacklisted IP address is removed before the policy-defined period (e.g., 30 days), we assumed that there was an application for deregistration from a third party. As a result of the investigation, we identified the cloud IP addresses that were assumed to be removed following an application for deregistration. Specifically, for a certain blacklist of (f) Spam with a policy of on-list duration of 30 days, 517 IP addresses were deleted within 30 days. The number of these IP addresses for each cloud service provider is shown in **Table 4**. This is less than 5% of the total number of unique cloud IP addresses on the blacklist. For other blacklists, we could not observe IP addresses that were deleted within the policy-defined period. This means that most of the blacklisted IP addresses were continuously used for cyber-attacks or that there was no application for the deregistration of blacklisted IP addresses that were no longer being used for cyber-attacks.

4.6 Observation of Cloud Service Abuse Using Darknet

To evaluate the reliability of the results of our analysis, it was necessary to verify whether the blacklisted cloud IP addresses were actually involved in cyber-attacks. To this end, we needed to observe cloud service abuse by using methods other than blacklists. In this study, we used the darknet. The darknet is a space of reachable and unused IP addresses on the Internet, and almost all packets arriving at the darknet can be regarded as malicious [19]. In this section, we discuss the source IP addresses of packets arriving at the darknet and investigate how many of them are cloud IP addresses. We focus on (a) Scan, which is the only attack type that the darknet can observe.

The darknet used in this study was the UCSD network telescope [19]. This darknet is a /8 network that has over 16 million IP addresses. We investigated the packets arriving at this darknet for 93 days from August 7 to November 7, 2019. The number of unique packet source IP addresses observed during this period was 200,658,917. These IP addresses included 276,297 AWS, 37,444 Azure, 46,285 GCP, and 5,605 Oracle IP addresses.

Next, we compared these IP addresses and the unique IP addresses registered in blacklists classified as (a) Scan during the same period. By way of this matching, we verified whether the blacklisted IP addresses actually conducted the scan. The matching results are shown in **Table 5**, where coverage refers to the proportion of blacklisted IP addresses that were observed in the darknet. The results show that the lowest coverage was 85% for Azure. In other words, most of the blacklisted cloud IP addresses were observed in the darknet during the same period. This shows that a multitude of cloud IP addresses that were registered in the blacklists of (a) Scan were actually involved in cyber-attacks.

Table 5 Number and coverage of blacklisted IP addresses observed in the darknet.

	AWS	Azure	GCP	Oracle
# (a) Scan Blacklisted IP Addresses	394	137	230	22
# Observed in Darknet	348	117	208	22
Coverage	88%	85%	90%	100%

4.7 Effectiveness of Our Observation Method Integrating Multiple Blacklists

In this section, we evaluate the effectiveness of integrating multiple blacklists as our observation method for cloud service abuse. There are two viewpoints for this evaluation. One is whether we could observe the attack more broadly. Another is whether we could observe the attack earlier. We define an IP address which was registered only in a single blacklist as “single-blacklisted IP address” and that which was registered in two or more blacklists as “multiple-blacklisted IP address.” We use single-blacklisted IP addresses for the former viewpoint and multiple-blacklisted IP addresses for the latter. The larger the number of single-blacklisted IP addresses, the more broadly we could observe cloud service abuse by integrating multiple blacklists since each blacklist had a different observation range. Moreover, we examined the first registered date of each blacklist for a multiple-blacklisted IP address. By investigating the time lag between the fastest blacklist and the slowest blacklist, we can determine how many days earlier we can observe that IP address compared to when using a single blacklist.

As shown in Section 1, the total number of unique blacklisted cloud IP addresses was 61,060. Among these, we revealed that the number of single-blacklisted IP addresses was 48,040 (79%) and that of multiple-blacklisted IP addresses was 13,020 (21%). A large number of single-blacklisted IP addresses indicate that cloud service abuse can be observed broadly by integrating multiple blacklists. The multiple-blacklisted IP addresses are also important for early detection of attacks. In fact, we found that we could observe an attack 12 days earlier on average compared with only using a single blacklist.

In contrast, the above analysis does not consider the characteristics of each blacklist. More specifically, only several large blacklists may have made the above contributions because the small blacklist was a subset of the large blacklist. To confirm this, we conducted a more detailed analysis focusing on each blacklist. For each blacklist, we defined the proportion of the number of single-blacklisted IP addresses to the total number of unique blacklisted cloud IP addresses as the “single-blacklisted IP address rate.” In addition, the proportion of the number of multiple-blacklisted IP addresses which registered faster than other blacklists is defined as the “fastest-blacklisted IP address rate.” **Table 6** summarizes the name, total number of unique blacklisted

Table 6 Single-blacklisted IP address rate and fastest-blacklisted IP address rate for each blacklist.

#	Blacklist Name	# IP Addrs	Single	Fastest	#	Blacklist Name	# IP Addrs	Single	Fastest
1	Badips Brute-force	7,099	48%	58%	21	ProjectHoneyPot Spammers	38	92%	100%
2	Badips DDos	20	35%	54%	22	AlienVault	6,407	77%	53%
3	Badips DNS	372	98%	89%	23	Bambenek	281	94%	29%
4	Badips FTP	104	36%	52%	24	BinaryDefense	980	52%	27%
5	Badips HTTP	4,307	45%	72%	25	BotScout	912	38%	51%
6	Badips Mail	5,456	63%	60%	26	CleanTalk	12,131	83%	55%
7	Badips RFI	36	25%	63%	27	CyberCrime	219	67%	58%
8	Badips SQL	1,666	89%	47%	28	Dangerrulez	125	0.80%	33%
9	Badips SSH	10,273	37%	80%	29	DShield	0	-	-
10	Badips VoIP	68	22%	70%	30	Feodo	0	-	-
11	Badips XML	10	20%	63%	31	Haley	1,364	35%	70%
12	Fail2ban Bots	212	38%	23%	32	LashBack	9,964	95%	50%
13	Fail2ban Brute-force	2,128	31%	40%	33	MyIP	616	75%	50%
14	Fail2ban FTP	55	40%	30%	34	NixSpam	1,730	71%	43%
15	Fail2ban IMAP	794	0.13%	58%	35	Nothink	504	39%	25%
16	Fail2ban Mail	1,578	31%	51%	36	Sblam	263	27%	40%
17	Fail2ban SIP	454	11%	8.1%	37	StopForumSpam	1,674	55%	51%
18	Fail2ban SSH	5,594	18%	38%	38	Talos	19	58%	88%
19	ProjectHoneyPot Commenters	205	94%	46%	39	VoIPBL	1,986	90%	68%
20	ProjectHoneyPot Dictionary	2	100%	-					

cloud IP addresses, the single-blacklisted IP address rate, and the fastest-blacklisted IP address rate for each blacklist. The single-blacklisted IP address rate can be defined when a blacklist registered at least one IP address. Similarly, the fastest-blacklisted IP address rate can be defined when the blacklist contained at least one multiple-blacklisted IP address. In Table 6, both the single-blacklisted IP address rate and the fastest-blacklisted IP address rate of 29: DShield and 30: Feodo, and the fastest-blacklisted IP address rate of 20: ProjectHoneyPot Dictionary cannot be defined and are shown as "-."

The average values of the single-blacklisted IP address rate and the fastest-blacklisted IP address rate were 52% and 53%, respectively. Here, from Table 6, it can be observed that the single-blacklisted IP address rate or the fastest-blacklisted IP address rate of the blacklist with a large number of blacklisted cloud IP addresses may not always be high. (For example, the single-blacklisted IP address rate of 9: Badips SSH and the fastest-blacklisted IP address rate of 32: LashBack.) In contrast, the single-blacklisted IP address rate or fastest-blacklisted IP address rate of the blacklist with a small number of blacklisted cloud IP addresses may be high. (For example, both the single-blacklisted IP address rate and the fastest-blacklisted IP address rate of 21: ProjectHoneyPot Spammers.) In other words, the size of the blacklist does not necessarily have to be large in order to increase the observation range of attack and to observe the attack earlier. Therefore, it can be said that even a blacklist with a small number of blacklisted cloud IP addresses is worth integrating.

5. Discussion and Limitations

5.1 Discussion

From the discussion in Section 4, we confirmed that the seriousness of cloud service abuse cannot be ignored and that effective countermeasures are necessary. Based on our findings, we believe that the risk of cloud service abuse can be mitigated. In this section, we make the following recommendations for cloud service users, cloud service providers, and blacklist providers. Note that we do not discuss how to prevent cloud service abuse *in advance* in this paper. Instead, we discuss what each stakeholder should do after cloud service abuse has occurred.

Cloud Service Users. When using a cloud service, users should check whether the assigned IP address is blacklisted. If the assigned IP address was previously used for cloud service abuse, the IP address could still be blacklisted and communication might be blocked. For example, using the web service IPVoid [20], we can compare the IP address against a total of over 100 blacklists just by entering the IP address we want to check into the browser. If the IP address is blacklisted, users can take measures such as receiving a different IP address by stopping and restarting the server.

Cloud Service Providers. In situations where cloud service abuse occurs constantly, if an abused IP address is freed for reuse and immediately assigned to another user, then various restrictions are placed on that user's service. Therefore, cloud service providers should detect and manage cloud service abuse at an early stage to minimize cyber-attacks using their services and to maximize their availability. To achieve this, cloud service providers can use the measurement method conducted in this study or namely, collect multiple IP address blacklists, find their blacklisted IP addresses, and take action to warn or suspend the corresponding malicious user/users.

Blacklist Providers. Blacklist providers need to reduce false positives as much as possible when creating blacklists. In Section 4.4, we showed that several IP addresses continue to be blacklisted for 30 days according to policy but are considered to have conducted no attacks after the first blacklisting. Because the IP addresses of cloud service providers are reused among users, it is not desirable to keep such IP addresses blacklisted for a long time. Using the measurement method in this study, blacklist providers can identify cloud IP addresses and treat them differently to other IP addresses. For example, when an attack from a cloud IP address is observed, not only blacklisting but also informing the cloud service provider may lead to a faster response. Reports of abuse can be made through channels possessed by the cloud service providers [21], [22], [23].

5.2 Limitations

There are some limitations to this study. One is that the observation range of cloud service abuse is limited to the observation range of the blacklist providers. As shown in Section 4.6, there are a large number of IP addresses that were observed in the darknet but were not blacklisted. Because it is impossible to observe all possible attacks on the Internet, it is not easy to completely solve this limitation. However, in the future, we will conduct an analysis that is as comprehensive as possible by increasing the number of acquired blacklists and the days of investigation. There is also the limitation that the accuracy of the measurement method used in this study depends on the accuracy of the information provided by the blacklist provider. However, in Section 4.6, we confirmed that highly accurate information was provided by at least those blacklists classified as (a) Scan. Further in Section 4.7, we showed that 21% of the total unique blacklisted IP addresses were registered in two or more blacklists. Observing multiple-blacklisted IP addresses increases the certainty that they are actually performing malicious activities meaning that integrating more blacklists will improve the accuracy accordingly. Moreover, attackers may be abusing cloud services called bulletproof hosting services, which allow users to host any content. In fact, many C&C servers are said to be hosted by bulletproof hosting service providers [24], [25]. Since we have not focused on such cloud services in this work, the countermeasures proposed in Section 5.1 do not take into account the abuse of these services. We will expand the scope of our measurement to include such services in future work.

6. Related Work

IP addresses are the most basic and essential identifier on the Internet; therefore, numerous related studies have been conducted. In this section, we summarize related works that are roughly divided into studies that focused on malicious IP addresses used for attacks and studies that focused on changes in the IP addresses themselves.

Malicious IP Addresses. Ramachandran et al. [26] analyzed the characteristics of the source IP addresses of a large amount of spam mail collected by spam traps from 2004 to 2005 and showed that the source addresses of spam mail were biased to a specific IP address range. Moreover, they showed that commercial IP blacklists for spam mail countermeasures cannot identify more than 30% of spam mail source IP addresses, and such source IP addresses are not blacklisted for more than one month [12]. Metcalf et al. [13] showed that there are several unique malicious IP addresses for each blacklist, and that there is less duplication of IP addresses among blacklists by collecting and investigating multiple blacklists that registered the malicious IP addresses. Zhao et al. [27] analyzed the characteristics in changes of malicious activity over the ten years from 2007 to 2017 using 22 types of IP address blacklists acquired from Wayback Machine. Mi et al. [28] conducted the first study on RESIP (Residential IP) proxy services, which revealed that the attacker used these IP addresses for malicious activities such as sending spam or hosting malicious sites. In 2019, Li et al. [14] collected a large number of public and commercial IP address blacklists and proposed objective evaluation indicators. Based on these indicators, they proved that the current IP blacklist was still insufficient for protecting users and organizations.

Our study is based on multiple malicious IP addresses or IP blacklists, which is similar to the above studies. However, there are two major differences. First, in our study, the investigated type of attack was not limited to spam emails but instead considers the more general current trends of cyber-attacks. Second, we focused on the IP addresses of cloud service providers and revealed the actual status of abuse specific to cloud services.

Changes in IP Addresses. Liu et al. [29] defined a DNS record that remains even though the DNS record reference resource (e.g., domain name or IP address) has not been used and is invalidated as a dangling DNS record (Dare). This was the first study to identify the security risks of Dare. They clarified that when a Dare reference destination is an IP address from a cloud service, a third party can obtain it after it is released. Borgolte et al. [30] showed that obtaining these cloud IP addresses is easy in terms of time and cost, which poses the risk of the attacker hijacking the domain name and issuing an SSL certificate. Pariwono et al. [31] verified the same problem as Refs. [29] and [30] by focusing on the domain name and IP addresses referenced from an Android application and revealed its risks. Nakamori et al. [32] emphasized that the same IP address is not always assigned to the same user when a dynamic IP address is assigned by an ISP or when an IP address is assigned by a cloud service provider. They proposed a method to identify such changeable IP address regions from the continuity of PTR records.

Considering the nature and actual status of changes to the owners of cloud IP addresses or dynamic IP addresses, as shown in the above studies, we used 39 blacklists and analyzed any changes in malicious cloud IP addresses. We clarified the characteristics of malicious cloud IP addresses that continue to be blacklisted and the attack trends unique to cloud services for the first time.

7. Conclusion

In this paper, we conducted the first large-scale analysis of cloud service abuse. The main idea of our study was to use diverse blacklists for observing cloud service abuse without direct observation. Our analysis of four typical/popular cloud services using 39 blacklists over 154 days revealed the actual status of cloud service abuse: changes in the number of blacklisted cloud IP addresses over time, the types of attacks, trends regarding IP address regions, on-list duration of the blacklisted IP addresses, and the status of deregistration. Moreover, we showed the effectiveness of our observation method. The findings of this study provide a foothold for cloud service users, cloud service providers, and blacklist providers to start taking effective countermeasures to deal with abuse of cloud services.

Acknowledgments This work was supported in part by the Japan Society for the Promotion of Science through Grants-in-Aid for Scientific Research (C) (20K11800).

References

[1] Fukushi, N., Chiba, D., Akiyama, M. and Uchida, M.: A Large-scale Analysis of Cloud Service Abuse, *Proc. IEEE CNS* (2020).

[2] Canalys: Global cloud market up 37%, with channels creating new growth engine (online), available from ⟨https://www.canalys.com/newsroom/global-cloud-market-Q3-2019⟩ (accessed 2020-05-07).

[3] Hot for Security: How any Instagram account could be hacked in less than 10 minutes (online), available from ⟨https://hotforsecurity.bitdefender.com/blog/how-any-instagram-account-could-be-hacked-in-less-than-10-minutes-21409.html⟩ (accessed 2020-05-07).

[4] Cybers Guards: Hackers abuse Microsoft Azure to use malware and evasion technology on C2 servers (online), available from ⟨https://cybersguards.com/hackers-abuse-microsoft-azure-to-use-malware-and-evasion-technology-on-c2-servers⟩ (accessed 2020-05-07).

[5] Cisco: Cisco Global Cloud Index: Forecast and Methodology, 2016-2021 White Paper (online), available from ⟨https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html⟩ (accessed 2020-05-07).

[6] Flexera: RIGHTSCALE 2019 STATE OF THE CLOUD REPORT FROM FLEXERA (online), available from ⟨https://resources.flexera.com/web/media/documents/rightscale-2019-state-of-the-cloud-report-from-flexera.pdf⟩ (accessed 2020-05-07).

[7] Amazon Web Services, Inc.: Barracuda blocking email from SES (online), available from ⟨https://forums.aws.amazon.com/thread.jspa?messageID=897282⟩ (accessed 2020-05-07).

[8] Amazon Web Services, Inc.: AWS IP Address Ranges (online), available from ⟨https://docs.aws.amazon.com/general/latest/gr/aws-ip-ranges.html⟩ (accessed 2020-05-07).

[9] Microsoft: Azure IP Ranges and Service Tags - Public Cloud (online), available from ⟨https://www.microsoft.com/en-us/download/details.aspx?id=56519⟩ (accessed 2020-05-07).

[10] Google Cloud: Google Compute Engine FAQ (online), available from ⟨https://cloud.google.com/compute/docs/faq?hl=en⟩ (accessed 2020-05-07).

[11] Oracle Cloud: IP Address Ranges (online), available from ⟨https://docs.cloud.oracle.com/en-us/iaas/Content/General/Concepts/addressranges.htm⟩ (accessed 2020-05-07).

[12] Ramachandran, A., Feamster, N. and Vempala, S.: Filtering spam with behavioral blacklisting, *Proc. ACM CCS*, pp.342–351 (2007).

[13] Metcalf, L. and Spring, J.M.: Blacklist Ecosystem Analysis: Spanning Jan 2012 to Jun 2014, *Proc. ACM WISCS*, pp.13–22 (2015).

[14] Li, V.G., Dunn, M., Pearce, P., McCoy, D., Voelker, G.M. and Savage, S.: Reading the Tea leaves: A Comparative Analysis of Threat Intelligence, *Proc. USENIX Security*, pp.851–867 (2019).

[15] FireHOL: All Cybercrime IP Feeds, FireHOL (online), available from ⟨https://iplists.firehol.org/⟩ (accessed 2020-05-07).

[16] Concurrency Labs: Save yourself a lot of pain (and money) by choosing your AWS Region wisely (online), available from ⟨https://www.concurrencylabs.com/blog/choose-your-aws-region-wisely/⟩ (accessed 2020-05-07).

[17] Azure VM Price: Azure VM Comparison (online), available from ⟨https://azureprice.net/⟩ (accessed 2020-05-07).

[18] Kaplan, E.L. and Meier, P.: Nonparametric estimation from incomplete observations, *Journal of the American Statistical Association*, Vol.53, No.282, pp.457–481 (1958).

[19] CAIDA: The UCSD Network Telescope (online), available from ⟨https://www.caida.org/projects/network_telescope/⟩ (accessed 2020-

05-07).
[20] IPVoid (online), available from ⟨https://www.ipvoid.com/⟩ (accessed 2020-05-07).
[21] Amazon Web Services, Inc.: How do I report abuse of AWS resources? (online), available from ⟨https://aws.amazon.com/premiumsupport/knowledge-center/report-aws-abuse/?nc1=h_ls⟩ (accessed 2020-08-24).
[22] Microsoft: Submit Abuse Report (CERT) (online), available from ⟨https://portal.msrc.microsoft.com/en-us/engage/cars⟩ (accessed 2020-08-24).
[23] Google Developers: Report suspected abuse on Google Cloud Platform (online), available from ⟨https://support.google.com/code/contact/cloud_platform_report?hl=en⟩ (accessed 2020-08-24).
[24] Khattak, S., Ramay, N.R., Khan, K.R., Syed, A.A. and Khayam, S.A.: A taxonomy of botnet behavior, detection, and defense, *IEEE Communications Surveys & Tutorials*, Vol.16, No.2, pp.898–924 (2013).
[25] Goncharov, M.: Criminal Hideouts for Lease: Bulletproof Hosting Services, Trend Micro (online), available from ⟨https://www.trendmicro.no/media/wp/wp-criminal-hideouts-for-lease-en.pdf⟩ (accessed 2020-08-24).
[26] Ramachandran, A. and Feamster, N.: Understanding the network-level behavior of spammers, *Proc. ACM SIGCOMM*, pp.291–302 (2006).
[27] Zhao, B.Z.H., Ikram, M., Asghar, H.J., Kaafar, M.A., Chaabane, A. and Thilakarathna, K.: A Decade of Mal-Activity Reporting: A Retrospective Analysis of Internet Malicious Activity Blacklists, *Proc. ACM AsiaCCS*, pp.193–205 (2019).
[28] Mi, X., Feng, X., Liao, X., Liu, B., Wang, X., Qian, F., Li, Z., Alrwais, S., Sun, L. and Liu, Y.: Resident evil: Understanding residential IP proxy as a dark service, *Proc. IEEE S&P*, pp.1185–1201 (2019).
[29] Liu, D., Hao, S. and Wang, H.: All Your DNS Records Point to Us: Understanding the Security Threats of Dangling DNS Records, *Proc. ACM CCS*, pp.1414–1425 (2016).
[30] Borgolte, K., Fiebig, T., Hao, S., Kruegel, C. and Vigna, G.: Cloud Strife: Mitigating the Security Risks of Domain-Validated Certificates, *Proc. NDSS* (online), DOI: 10.14722/ndss.2018.23327 (2018).
[31] Pariwono, E., Chiba, D., Akiyama, M. and Mori, T.: Don't throw me away: Threats Caused by the Abandoned Internet Resources Used by Android Apps, *Proc. ACM AsiaCCS*, pp.147–158 (2018).
[32] Nakamori, T., Chiba, D., Akiyama, M. and Goto, S.: Detecting Dynamic IP Addresses and Cloud Blocks Using the Sequential Characteristics of PTR Records, *Journal of Information Processing*, Vol.27, pp.525–535 (2019).

Editor's Recommendation

This paper provides a large-scale measurement of cloud service abuse using blacklisted IP addresses. This study analyzes actual cases and reveals trends in attacks with respect to attack type, region, duration, and anti-abuse actions. These findings are novel as well as practical. The paper gives insights to readers in this research field and thus is selected as a recommended paper.

(Program chair of MWS 2019 Mamoru Mimura)

Naoki Fukushi received his B.E. and M.E. degrees in Computer Science from Waseda University in 2018 and 2020. His research interest covers Cyber Security. He is now with NTT Secure Platform Laboratories, Tokyo, Japan.

Daiki Chiba is currently a researcher at NTT Secure Platform Laboratories, Tokyo, Japan. He received his B.E., M.E., and Ph.D. degrees in computer science from Waseda University in 2011, 2013, and 2017. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 2013, he has been engaged in research on cyber security through data analysis. He won the Research Award from the IEICE Technical Committee on Information and Communication System Security in 2016, 2018, and 2019 and the Best Paper Award from the IEICE Communications Society in 2017. He is a member of IEEE and IEICE.

Mitsuaki Akiyama received his M.E. and Ph.D. degrees in information science from Nara Institute of Science and Technology, Japan in 2007 and 2013. Since joining Nippon Telegraph and Telephone Corporation (NTT) in 2007, he has been engaged in research and development on cybersecurity. He is currently a Senior Distinguished Researcher with the Cyber Security Project of NTT Secure Platform Laboratories. His research interests include cybersecurity measurement, offensive security, and usable security and privacy. He is a member of the IEEE, IPSJ, and IEICE.

Masato Uchida is currently a Professor in the Department of Computer Science and Engineering, School of Fundamental Science and Engineering, Waseda University, Tokyo Japan. He received his B.E., M.E. and Ph.D. degrees from Hokkaido University, Hokkaido, Japan, in 1999, 2001, and 2005, respectively. In April 2001, he joined NTT Service Integration Laboratories, Tokyo, Japan. From August 2005 to March 2012, he has been an Associate Professor in the Network Design Research Center, Kyushu Institute of Technology, Fukuoka, Japan. From April 2012 to March 2015, he has been an Associate Professor in the Department of Electrical, Electronics and Computer Engineering, Faculty of Engineering, Chiba Institute of Technology, Chiba, Japan. From April 2015 to March 2016, he has been a Professor in the same department. From April 2016 to March 2017, he has been a Professor in the Department of Information and Communication Systems Engineering of the same university. He has been engaged in research on data science and machine learning with application to various fields in computer science including information networking and information security. He is a member of the IEEE, ACM, IPSJ, and IEICE.