

# Evaluating IP Blacklists Effectiveness

Luca Deri  
*ntop*  
Pisa, Italy  
deri@ntop.org

Francesco Fusco  
*IBM Research*  
Zurich, Switzerland  
ffu@zurich.ibm.com

**Abstract**—IP blacklists are widely used to increase network security by preventing communications with peers that have been marked as malicious. There are several commercial offerings as well as several free-of-charge blacklists maintained by volunteers on the web. Despite their wide adoption, the effectiveness of the different IP blacklists in real-world scenarios is still not clear.

In this paper, we conduct a large-scale network monitoring study which provides insightful findings regarding the effectiveness of blacklists. The results collected over several hundred thousand IP hosts belonging to three distinct large production networks highlight that blacklists are often tuned for precision, with the result that many malicious activities, such as scanning, are completely undetected. The proposed instrumentation approach to detect IP scanning and suspicious activities is implemented with home-grown and open-source software. Our tools enable the creation of blacklists without the security risks posed by the deployment of honeypots.

**Index Terms**—IP blacklist, network traffic analysis, host reputation, open-source software.

## I. INTRODUCTION AND MOTIVATION

Reputation systems have been extensively used in network security and network management to maintain networks and networked services secure and reliable. The most commonly used reputation system in network management is blacklisting [34], [48]. A blacklist is an access control mechanism which denies access to selected network resources to peers belonging to a curated list. Blacklists can be compiled for different resources, including IP addresses, domain names, URLs and checksum of files (e.g., binary files corresponding to malware). Blacklists often represent the first line of defence for many networks as they can reduce internal hosts' risk of establishing communications with peers with a bad reputation. For example, several anti-spam filters rely on IP reputation blacklists [10] which store IP addresses and domain names from which it is not recommended to accept emails.

While IP blacklists [16] made the blacklisting approach popular, blacklisting techniques are used at a finer-grained granularity than the IP level, i.e., to block individual network communications marked as suspicious using advanced network fingerprinting mechanisms such as JA3 [2], [33], malicious SSL certificates [22] or IoC (Indicators of Compromise), which combine information such as IP addresses, virus signatures and URLs hashes of malware files [3], [5].

Blacklists are not used only against malware or spam [11], [20], [34], [41], [50] but also for other purposes such as securing networks [31], [41] and blocking advertisements [43] or improving information delivery and security, including

blocking connections from anonymous VPNs or preventing web and security crawlers from scanning a network in search of vulnerabilities that could be potentially used for future attacks [40].

The widespread adoption of IP blacklists has been mostly driven by simplicity and ease of deployment. There are many commercial offerings and several free-of-charge blacklists maintained by volunteers spread across the globe [7], [19], [35]. However, when relying on IP blacklists, one has to consider the inherent limitations of the method [37]. First, blacklists are only effective when maintained in a timely manner [49]. Newly classified malware hosts must be included in the lists, while no longer malicious hosts need to be removed to minimise false positives. Second, blacklists are not equally effective across the planet. In particular, a blacklist built and maintained for a specific region (e.g., North America) is not guaranteed to be effective when deployed in another region (e.g., Europe). Third, blacklists do not necessarily cover the traffic seen in the network where they are deployed.

Since blacklisting approaches have inherent weaknesses, assessing their effectiveness in real-world scenarios is of extreme importance. In recent years, several rigorous studies have been performed to evaluate and compare malware blacklists [18], [27], [46], and domain blacklists when applied to specific application layer services such as email (for spam and phishing) and web traffic [26], [38], [44]. Surprisingly, despite the availability of a large free-of-charge collection and commercial offerings, there are limited studies on the effectiveness of IP blacklists and many important questions remain unanswered. For example, it is unclear if blacklists are equally effective across distinct networks, which false positive rate has to be expected and which amount of malicious traffic goes undetected.

This paper fills this gap and offers an in-depth study of publicly available IP malware blacklists used in large production networks that are different in size, in nature and span multiple geographical locations over the globe. Our goal is to assess the effectiveness of the blacklists by evaluating them against malicious activities, which can be detected with a high degree of confidence using host instrumentation and aggressive heuristics based on passive network monitoring data. In a nutshell, we wanted to evaluate the effectiveness of the blacklist on the *easy cases*, i.e., for hosts that can be detected as malicious with certainty, even using simple mechanisms. Surprisingly, even compared with our conserva-

tive approaches, blacklists can only capture a small fraction of scanning activities, and the recall does not significantly improve when blacklists maintained by distinct parties are combined.

In this paper, we introduce the following contributions. First, we perform a large-scale study evaluating IP blacklists on real-world production networks of more than *hundred thousand* IP hosts belonging to multiple production networks. Second, we describe an effective instrumentation approach to detect IP scanning and suspicious activities toward network servers. Third, we showcase that blacklists are optimized for precision, leaving much of the malicious traffic undetected and a false sense of security. Finally, all the software presented in this paper is home-grown and open-source enabling researchers and network operators to repeat our experiments in their networks.

## II. DATA COLLECTION ARCHITECTURE

To evaluate different IP blacklists, we have selected three networks based in three distinct European locations both using IPv4 and IPv6 addresses:

- An Italian service provider with about 5'000 hosted servers. This site is connected with two 10 Gbit links to the Internet. Network traffic is mirrored at the edge router, converted into flows by a software probe performing DPI (Deep Packet Inspection) and then collected.
- A university located in northern Europe with over 100'000 assigned public IP addresses. In this site, network traffic is exported by three border gateways in NetFlow format [9] and collected at a central location.
- A leading hosting provider located in the Netherlands hosting corporate servers. Each server is monitored using log files (including authentication, web administration, email and TCP/UDP ports monitoring) instead of using raw traffic or flows as in the above two scenarios. These logs contain several invalid login attempts, brute force service attacks, port scans, and data exchange attempts on closed ports. For this reason to the best of our knowledge, no false positives are present, and even if there is a low probability of misconfigurations or human errors they should be considered very minor and will not change this evaluation. Analysis results are presented in Section II-F.

We have done our best to perform our measurements on heterogeneous networks both IPv4 and IPv6 based, located in different European countries. We are aware that a worldwide measurement system would have been desirable, and we are working at that as reported later in the paper. While our monitoring infrastructure has been active for a longer period of time, in this paper, we report results for three weeks of traffic between February 27th and March 19th, 2023. As the monitoring systems are permanent, data is continuously collected and blacklists are updated daily and matched against the traffic. Over the past five months of traffic monitoring, the results were consistent with what is presented in the paper. It is worth remarking that during the Christmas break, the number

of attacks detected decreases significantly compared to non-holiday weeks. For instance, on a mail server, we monitor in this experiment, during the period Dec 27th - Jan 3rd this host received an average of 30k scans/day with a low of 22k on Jan 1st, with respect to 80k scans/day observed in March. A possible interpretation is that behind automated scans there is some sort of human activity that was reduced during holidays, or that perhaps hackers automated the scripts to be less aggressive during holidays.

The main objective of our effort was not to create yet another blacklist to be positioned against the popular ones. Instead, the main goal was to collect the IPs that corresponded to attackers with high confidence in order to evaluate the existing blacklists and understand their strengths and limitations. Our choice to minimize the number of false positives was necessary to estimate:

- the number of attackers' IPs present in our network that *are not* present in public IP blacklists.
- the false positive rate, i.e., how many IP addresses present in blacklists visited the monitored networks without doing an attack, and thus blocking them would have led to a disservice. This is an important metric for deciding if blacklists can be effectively used as a first line of defence, or if they are unreliable and thus unsuitable for this task.

Achieving the goal required analysing the traffic and storing the flow data persistently for the entire period to enable interactive exploration and manual inspection. The matching between the traffic blacklisted and detected was performed daily using the blacklists downloaded on the same day.

### A. Network Flow Analysis

We collect network flows with ntopng [15], a popular open-source network traffic monitoring tool we have developed, and store them in ClickHouse [24], a high-performance columnar database [32]. In addition to providing flow collection capabilities, ntopng implements over hundred distinct rules, which allow the calculation of a *Cyberscore* [13] for every network activity. Cyberscore is a scalar indicator that can be computed by combining multiple rules (called *checks*). Network administrators can configure those checks to identify attackers and the corresponding victims. Since we were interested in minimizing the number of false positives in our setup, we only enabled a handful of simple checks that indicate attacks with very high confidence. The downside of this configuration choice is that some attacks targeting specific hosts might not be considered, but as we will explain later in the section, this is not a limitation. For this reason, ntopng was configured to report alerts produced by remote host scanners, presumably during the pre-attack reconnaissance phase as defined in the Mitre att&ck [42], using a simple algorithm:

- Similar to [12], we monitor RX traffic of unused IPs in the monitored networks, for which we do not expect to observe incoming traffic. Note that in most networks, not all IP addresses are in use; this happens in particular when monitoring research networks that historically have

allocated multiple large (/16 IPv4 or larger) network ranges, and partially used them.

- Remote peers contacting local receive-only hosts using TCP are reported as scanners if and only if they contact at least 128 receive-only hosts in one day. While TCP is bi-directional by design, UDP traffic is ignored in this computation for two reasons. First, not all UDP protocols are always bidirectional (e.g. Syslog and RTP are not). Second, unidirectional flows using bidirectional UDP protocols such as DNS, NTP, and SNMP are not necessarily an indication of a problem. For instance, SNMP GET is a confirmed PDU, whereas SNMP traps are not. Not considering UDP can reduce the number of scanners we detect, but it protects us from possible false positives for the reasons described above.
- In order to avoid web crawlers being classified as scanners, traffic on destination TCP ports 80 and 443, are not considered.

We are aware that a more sophisticated mechanism such as a honeypot could be used as attackers targeting websites (for instance those requesting malicious URLs) or scanners limiting their activities to fewer hosts will not be reported (i.e. false negatives). While honeypots could report richer information (for instance it can be used to grab malware software that attackers will drop on the honeypot host) than our passive approach, deploying a honeypot on an operational network is not always possible as many service providers prohibit the installation of software of this nature. Additionally, this conservative approach is not a problem for the blacklist evaluation because with this setup we have very limited, if any, false positives. In all cases, traffic flows are stored in the database and thus retrieved for later analysis if necessary.

### B. Server Monitoring

In addition to identifying scanning activities with the algorithm described in Section II-A, we monitor log files of core services for the network servers belonging to the networks being monitored. An agent installed on each server monitors both TCP and UDP ports for probing attempts on closed ports, and log files of core services such as SSH, IMAP/SMTP and website management pages. By parsing those logs, the agent can identify intrusion attempts and the IP address from which the attempt has originated. A remote IP is detected as an intruder if logs report multiple unsuccessful access attempts originated at this IP. In addition, similar to what happens with unused hosts in flow monitoring, connections to multiple closed server TCP ports can detect scan attempts, as well as connection attempts to closed ports. The combination of access logs and closed ports monitoring provides evidence of malicious activities, with no false positives, to be computed and used during evaluation together with the scanning detection approach of Section II-A.

### C. IP Blacklist Evaluation

IP blacklists are available in two formats: as a text file that can be used to implement ACLs (Access Control Lists)

on network devices or as cloud services accessible with web REST APIs. As we are interested in evaluating IP blacklists to implement a first level of defence, we need access to them in text format. For this reason, this work focuses on file-based blacklists and uses cloud-based security services to manually verify the results. In particular, we have considered a few popular IP blacklists listed in Table I, each with the average number of IP entries. Although all analysed blacklists are designed to identify security threats, some of them have distinctive features. For instance, Feodo tracks only certain types of malware, and Stratosphere [21] does not track SMTP attacks. These details are relevant for interpreting some of the results presented later in this paper. All blacklists were downloaded for all days of the experiment and evaluated against the observed traffic of the same day. In the rest of this paper, we mainly report only the results for the week of March 13th, with comparable results for all the previous weeks of the experiment. We are unaware of how the lists are created, if they are built in a single geographic region or if they have sources of data spread globally, or what the algorithm is for adding and removing entries. The exception is the Stratosphere Blocklist Generation Project [28], whose algorithm for the Prioritize New (PN) blacklist is documented in [29]. Table I reports the analysed lists and some details such as the average number of entries and change rate (number of lines that change on average in two consecutive days). As reported, the blacklists' length is not homogeneous as it ranges from a few entries of *dshield* to over 14k of Stratosphere's blacklists. All lists report only IPv4 /32 addresses except for *dshield* and *Emerging Threats* which instead also report subnets. While the work reported in this paper applies and has been evaluated on both IPv4 and IPv6 networks, the evaluation discussed below will be limited to IPv4 as we have not been able to find high-quality and actively maintained IPv6 blacklists.

TABLE I  
ANALYSED IP BLACKLISTS

File-based IP blacklists	Entries	IPs	Update Rate (Daily)
<b>Snort IP BlockList</b> [7]	812	812	3%
<b>EmergingThreats (ET)</b> [35]	1'608	16.4 M	2%
<b>Feodo Tracker</b> [1]	184	184	36%
<b>dshield</b> [17]	29	7'936	31%
<b>Stratosphere (PN)</b> [25]	14'518	14'518	9%
<b>AlienVault (AV)</b> [4]	689	609	1%

As shown in Table I Stratosphere IPs and EmergingThreats blacklists are the ones containing the highest numbers of IP addresses. ET is the blacklist containing millions of IP addresses due to the use of IPv4 large prefixes such as /16. All blacklists contain only IPv4 addresses.

While nothing can be said about cloud-based blacklists shown in Table II, for file-based ones, Table I reports the number of entries and the percentage of entries that change daily. Some blacklists update entries very seldom, whereas others are more dynamic. Most lists are updated daily with more changes during the middle of the week rather than on

TABLE II  
ANALYSED CLOUD-BASED IP BLACKLISTS

Cloud-based IP blacklists
VirusTotal [47]
AbuseIP DB [30]
Greynoise [23]

week-ends, others (AlienVault) are updated less frequently. Before comparing blacklists against the collected network data, blacklists were cross-compared to test the intersection of IP addresses among them. The comparison was implemented using a Python script that uses a Patricia tree [39] module to interpret IP addresses.

TABLE III  
BLACKLISTS IP ENTRIES INTERSECTION (CONTAINS/CONTAINED)

List Name	PN	AV	Snort	ET	Feodo	dshield
PN	-	3	13	0	0	0
AV	3	-	0	0	0	0
Snort	13	0	-	1	0	0
ET	978	1	3	-	170	10
Feodo	0	0	170	0	-	0
dshield	1029	0	0	11	0	-

Table III contains the intersection of the lists on March 13th, with similar results (i.e. minor differences in value less than 10%) for the other days of the experiment. The intersection of the lists (i.e. an IP address that is included in the intersection of two lists if it is contained in both of them) is limited. The results are biased for lists that contain subnets as it is unlikely that all subnet IPs are reported in other lists. This is because we assume that a subnet entry is fully included in another list when all the subnet IPs are included, and it explains for instance why in the above table PN (that uses /32 prefixes) has 978 entries contained in ET (that uses larger subnets), but not the opposite as no ET network completely fits inside the PN network list.

#### D. Comparing Flow Data with File-based IP Blacklist

Figure 1 shows the flow distribution of our experiments for the first 7 days of March 2023 at the service provider (SP) and university (U) network.

Table IV reports the dataset size, including the number of IPs with zero cyberscore as defined in II-A.

TABLE IV  
DAILY DATASET SIZE (7 DAYS AVERAGE)

	Service Provider	University Network
Total Flows	153M	194M
Total Flows with Zero Cyberscore	76M	82M
Active Local IPs	49K	79K
Unique Remote Client IPs	1.7M	71K
Remote Scanner IPs	1.8K	3.1K
Remote IPs with Zero Cyberscore	1.3M	65K

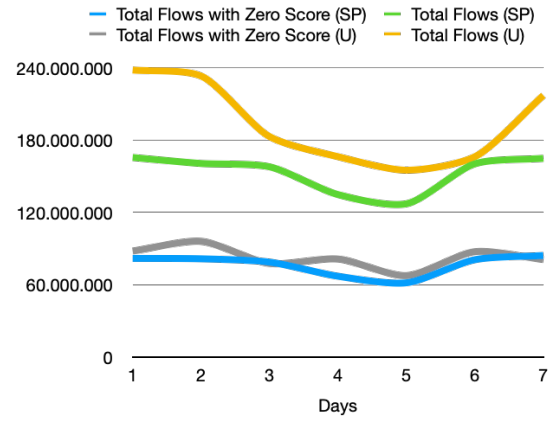


Fig. 1. Flows Distribution (March 1st-7th)

Only 6% scanner IPs of Table IV visited both networks on the same day. These numbers change significantly when monitoring the scanners over two weeks.




Fig. 2. Scanners Propagation (7 Days)

Figure 2 shows that after one day 67% of the university scanners also visited the service provider network. After 13 days this percentage slowly increased to 70%. The interpretation of this data is that most scanners move across networks in less than a day, so promptly updating lists is a key property that a blacklist should have to be effective. In addition, having a large set of monitoring nodes across the Internet helps create a comprehensive blacklist that includes attackers' IPs that soon will visit the network to protect. Neighbourhood Watch is an old concept used by humans, that demonstrated to be effective as an early warning system [45] for Internet monitoring, and thus that should also be considered in blacklists. When some hosts are under scan or attack, likely the attackers will also target other hosts belonging to the same network: if attacker IPs are immediately blacklisted, this could be a very effective early-warning technique to protect the whole network.

Another interesting finding is the average number of unique client IP addresses that visited both networks in a single day. This is probably justified by the fact that even if the service provider network is much smaller, due to the nature of this

business, most hosts are servers, whereas, at the university, it is probably the opposite. This said, about 6.5k unique client IP addresses visited both networks on the same day. Meaning that even if these two networks are different, there is some limited overlap. When looking at these overlapping IPs more in detail, if the last IP address byte is ignored (essentially merging them in /24 networks), we accounted for 20k networks with at least 8 IP addresses, 18k networks with 8...63 IP addresses, and 1.5k networks with at least 64 IP addresses. Randomly selecting IPs from networks with at least 8 IPs and searching them on an IP reputation database [8], most of these networks have a poor reputation or are used by companies that have as core business periodically scan the Internet (e.g., *internet-measurement.com* and *shodan.io*).

In Table V, we report the average number of IP addresses corresponding to scanners that are present in the blacklists. For convenience, we report the percentage of IP addresses present in each blacklist. What appears evident from the results is that the gap between blacklists is substantial. Providers such as AlienVault, Feodo and Snort are ineffective in marking scanner IPs. Stratosphere Prioritize New is the best blacklist to identify scanner IPs, and augmenting it with all the other blacklists improves the detection only marginally (by 1%).

TABLE V  
SCANNER IPS MATCH RATE WITH FILE-BASED IP BLACKLISTS (7 DAYS AVERAGE)

Blacklist	Service Provider	University
Stratosphere (PN)	50.5%	14.6%
EmergingThreats (ET)	13.6%	4.8%
dshield	11.3%	4.5%
AlienVault	0%	0.1%
Snort	0.1%	0%
Feodo	0%	0%
PN+ET+dshield	50.9%	14.7%

We have tried to evaluate the effectiveness of blacklists over time to understand how blacklist updates affect results. For this reason, we have matched the traffic daily (i.e. at midnight of day X we download the blacklist that will match the traffic of the day) against using the blacklist of day X for matching the traffic of days X+1, X+2...X+5 and results of 5 days comparison are reported in Table VI. As expected, using a daily blacklist produces slightly better results than reusing the same blacklist for matching scanner IPs of the following days. However, the slow decay rate is probably surprising, i.e. blacklist of day X is also effective for a couple of days more with no match rate degradation.

TABLE VI  
STRATOSPHERE AP MATCHES (MAR 13-17TH) USING MAR 13TH BLACKLIST (SERVICE PROVIDER)

Match Rate	Mar 13	Mar 14	Mar 15	Mar 16	Mar 17
Daily (%)	49%	53%	51%	51%	51%
Delta (%)	0%	0%	-0.3%	-0.5%	-1.4%

In order to evaluate if blacklists contain false positives, we have extracted the list of IPs whose flows had a maximum cyberscore of 100 (minor issues). We have matched such IPs against all the file-based blacklists and found no overlap, i.e. no flow client IP was listed in blacklists. This experiment reported similar results for the university network. This result is very encouraging as it indicates that blacklists seem to have no false positives with flows having no or minor issues. Said that a more in-depth analysis of flows with a higher score should be performed, the above result encourages the utilization of blacklists for blocking traffic.

In another test, we evaluated the IP blacklist coverage, namely, out of all malicious communications detected by ntopng (including but not limited to suspicious DGA domain contacted, possible exploit, and TLS certificate mismatch) performed by remote peers, what is the (daily) percentage of peers listed in one of the analysed IP blacklists. ntopng can trigger over 100 different types of cybersecurity checks covering both plain-text and encrypted traffic analysis [14]. To avoid false positives and make this experiment reproducible we have enabled only the default checks that trigger alerts based on nDPI security risks (nDPI is an open-source deep packet inspection library on top of which ntopng is built): nDPI inspects the initial flow content and produces a flow risk assessment based the flow packet payload. We have extracted from the alert database the list of remote peers that performed malicious activities from monitored traffic as described in Section II. This has allowed us to identify both network and port scanners listed above, as well as other types of threats reported by nDPI that we have used for validating the blacklists. This assessment is based on the assumption that nDPI has a very low number of false positives if any [6], [36], and to the best of our knowledge this is the case being it is used in many firewall and IPS tools that use it to block traffic and thus that cannot tolerate false positives. In nDPI, the security analysis is based on flow payload inspection (e.g. malformed packet, expired or self-signed TLS certificate, invalid HTTP request) and flow state (e.g. TCP flow with no answer) without enabling checks based on heuristics that might be affected by false positives. We have considered only those hosts that ntopng considers malicious with high confidence (i.e. host cyberscore bigger than 5'000).

TABLE VII  
NTOPNG ALERTED IPS (MATCH %) LISTED IN FILE-BASED IP BLACKLISTS (SERVICE PROVIDER)

Day	PN	AV	Snort	ET	Feodo	dshield
Mar 13	37%	0%	0%	6%	0%	7%
Mar 14	33%	0%	0%	8%	0%	13%
Mar 15	34%	0%	0%	7%	0%	3%
Mar 16	38%	0%	0%	9%	0%	7%
Mar 17	44%	0%	0%	7%	0%	13%
Mar 18	24%	0%	0%	2%	0%	3%
Mar 19	31%	0%	0%	0%	0%	0%
Average 13-19	34%	0%	0%	6%	0%	7%

Table VII shows that in the best case, less than 50% of the

hosts were present in blacklists, which confirms the results reported in Table V that consider only scanners. For those not having a match in the blacklist, we have taken a random subset of 500 hosts (as explained later this is the daily limitation for two of the considered IP blacklists). Over 60% of these IPs are listed in cloud IP blacklists. The outcome is that file-based blacklists can detect only a subset of scanners and attackers, and therefore while they are still useful, file-based IP blacklists are not sufficient to successfully protect a network.

### E. Comparing Flow Data with Cloud-based IP Blacklist

As anticipated earlier in this paper, we have compared cloud-based blacklists with file-based blacklists. Cloud-based services periodically aggregate information from various sources (VirusTotal). Others label an IP as malicious if there is a minimum consensus (AbuseIP), and others (Graynoise) instead use other proprietary and undisclosed techniques. As these services allow a limited number of searches per day (50 hosts/week for GreyNoise, 500/day for VirusTotal and 1000/day for AbuseIP with free plans), we have selected two 500 IPs lists (one observed in the service provider and the other in university network) scanners addresses from a single day of the observation period scanners' IP list and checked the match rate using both file (of the same day) and cloud-based blacklists. For GreyNoise the experiment was limited to 50 hosts due to the above limitations.

TABLE VIII  
CLOUD BASED BLACKLISTS MATCH RATE: 500 IP SCANNERS (SERVICE PROVIDER / UNIVERSITY)

Blacklist	Service Provider		University	
	List.	Cons.	List.	Cons.
VirusTotal	75%	80%	62%	63%
Abuse IP DB	98%	25%	68%	18%
Greynoise (50 IPs)	34%	10%	34%	10%
PN (File-based)	75%	12%	75%	12%
ET (File-based)	36%	4%	36%	4%

The table above shows the results of this experiment, including the match rate of file-based blacklists on the same IPs. Results are divided into two columns: the first one reports a match if the IP is listed, and the second only if there is also a consensus (5 or more matches for VirusTotal, 100% accuracy for Abuse IP, and 'malicious' for GreyNoise). Repeating the experiment on a different day of the observation period produces similar results. It is odd to report that except VirusTotal, all other blacklists perform poorly when matching scanner IPs coming from the university network. According to the above results, with the limitations of the number of queries per day, cloud-based blacklists seem to outperform file-based ones significantly. In conclusion, the nature of cloud-based services prevents them from being used to dynamically configure ACLs in security devices. Therefore, these cloud services are helpful for security analysts but cannot be practically used as the first line of defence as file-based blacklists.

### F. Comparing Server Monitoring Data with IP Blacklist

This section evaluates intrusion attempts on two hosts, one used as an email and the other as a web server. The same service provider hosts both servers, however on two very different networks located in the Netherlands. Only two administrators from specific IP addresses have access to the server, and to avoid interference with collected data, their accesses have not been considered in the analysis using the servers as honeypots. All attempts are recorded by looking at multiple (3+) failures in authentication and application (e.g. email and web server) log files and connection attempts on closed TCP ports. Contrary to the previous experiment where ntopng was used, here we have used ipt\_geofence (<https://github.com/ntop/iptables-geo>) which is deployed on the monitored hosts. This is an open-source tool we have developed and is able to detect scanning attempts, contacts to closed TCP/UDP ports and service attacks looking at system log files (e.g. invalid login reported in the auth.log file). For this reason, we assume that all the above activities represent real scans or non-authorized failed connection attempts not affected by false positives; this is as we take into account only multiple failures on servers used only by the author. Table IX reports the results of a single day of analysis in the paper observation period with similar results recorded on the other days of the week of March 13th.

TABLE IX  
SERVER MONITORING DATA ANALYSIS

Attackers	Web Server	Mail Server
Total number of Attacks In Blacklist	75	450
Also visited Service Provider	73%	42%
Also visited University	14%	11%
	5%	4%

Out of this experiment, these are the main findings:

- Blacklists are more effective for the web server host rather than the mail server, even if the total number of attacks on the mail server is higher. We are not sure why this happens for all blacklists except Stratosphere which is a list built without any email malware feeds. Probably the fact that the two servers are hosted on two totally different networks, and the different intent of blacklists is relevant here.
- Comparing attackers with the list of IPs that visited the same-day service provider and the university network, we observe that only a small minority visited such networks. In this case, as observed before, attackers are visiting the service provider more often than the university. This is a bit odd as the university network has many more hosts than the service provider. Instead, looking at the data over 7 days, we have observed that 70% of the scanners also visited the other network which implies that probably it is just a matter of time. We need to study our data over a longer period of time with more monitored hosts in order better study this fact.

- Another fact worth remarking is that while both hosts are dual stack, 99% of the attackers use IPv4 addresses.

### III. CONCLUSIONS

In this paper, we describe our experience in evaluating free-of-charge IP blacklists over two large production networks and Internet corporate servers. We propose an effective method which can be used to identify attackers with high confidence, with the goal of estimating the coverage of IP blacklist. The method combines passive monitoring with log analysis which is performed directly on production servers using local agents. Our large-scale study shows that IP blacklists provide limited protection against attackers. In particular, top free-to-use blacklists identify about 50% of the scanner attacks. Additionally, we find out that even if different blacklists are not strongly overlapping, i.e. their intersection is low, by combining all the blacklists the coverage can be only marginally improved. The low recall of the malicious IPs can be explained by the small false positive rate, indicating that blacklists are most likely optimized for precision rather than recall. Cloud-based blacklists can help to confirm malicious activities, but the limitation in the API queries prevents these types of blacklists from being used as the first line of defence.

Our methodology in evaluating IP blacklists suggests that detection algorithms tuned for precision (i.e. they include only hosts that performed only major malicious activities in order to minimise false positives, as the blacklists we have evaluated have limited coverage of attacks detected in our experiments) when used in large geographically distributed and always-on monitoring infrastructures might represent a valuable strategy to improve current blacklists.

### IV. FUTURE WORK

This paper monitored traffic in different European locations. A future goal is to extend this work to non-European sites and different network types including residential and mobile networks.

The experiments described in this paper have shown that blacklists are affected by two main limitations: they only detect a subset of network scanners, and their detection rate is not constant across the service provider and the university networks. We believe that deploying more sensors across the Internet, better if in heterogeneous networks, could address the above issues. This is the driving force for creating a feed for publicly available IP malware blacklists using the scanning detection mechanisms developed in this paper (and eventually more sophisticated ones). We are already developing a prototype that we would like to start deploying in the coming months.

Finally, the number of IPs in the blacklists that never incurred malicious activities against the monitored infrastructure was not addressed in this study but represents a topic of interest to better understand the quality of the blacklists in general, and thus a future work item.

### ACKNOWLEDGMENT

The authors would like to thank Joaquin Bogado, and Sebastian Garcia for endless discussions, experiments evaluation and suggestions throughout this research work.

### REFERENCES

- [1] Abuse.ch. Feodo Tracker. [https://feodotracker.abuse.ch/downloads/ipblocklist\\_recommended.txt](https://feodotracker.abuse.ch/downloads/ipblocklist_recommended.txt), 2023.
- [2] Abuse.ch. SSLBL: SSL Blacklists. <https://ssllbl.abuse.ch>, 2023.
- [3] Abuse.ch. Threatfox: sharing indicators of compromise (IOCs). <https://threatfox.abuse.ch>, 2023.
- [4] AlienVault. IP Reputation database. [https://raw.githubusercontent.com/firehol/blocklist-ipsets/master/alienvault\\_reputation.ipset](https://raw.githubusercontent.com/firehol/blocklist-ipsets/master/alienvault_reputation.ipset), 2023.
- [5] Jan Bayer, Yevheniya Nosyk, Olivier Hureau, Simon Fernandez, Ivett Paulovics, Andrzej Duda, and Maciej Korczyński. Study on domain name system (dns) abuse: Technical report. *arXiv preprint arXiv:2212.08879*, 2022.
- [6] Tomasz Bujlow and Valentin Carela-Espanol. Comparison of deep packet inspection (dpi) tools for traffic classification. 2013.
- [7] Cisco. Snort IP Blocklist. <https://snort.org/downloads/ip-block-list>, 2023.
- [8] Cisco. Talos Intelligence. <https://www.talosintelligence.com>, 2023.
- [9] Benoit Claise. Rfc 3954: Cisco systems netflow services export version 9, 2004.
- [10] Duncan Cook, Jacky Hartnett, Kevin Manderson, and Joel Scanlan. Catching spam before it arrives: domain specific dynamic blacklists. In *Proceedings of the 2006 Australasian workshops on Grid computing and e-research-Volume 54*, pages 193–202. Citeseer, 2006.
- [11] Baris Coskun. (un) wisdom of crowds: Accurately spotting malicious ip clusters using not-so-accurate ip blacklists. *IEEE Transactions on Information Forensics and Security*, 12(6):1406–1417, 2017.
- [12] Alberto Dainotti, Alistair King, and Kimberly Claffy. Analysis of internet-wide probing using darknets. In *Proceedings of the 2012 ACM Workshop on Building analysis datasets and gathering experience returns for security*, pages 13–14, 2012.
- [13] Luca Deri and Alfredo Cardigliano. Using cyberscore for network traffic monitoring. In *2022 IEEE International Conference on Cyber Security and Resilience (CSR)*, pages 56–61. IEEE, 2022.
- [14] Luca Deri and Francesco Fusco. Using deep packet inspection in cybertraffic analysis. In *2021 IEEE International Conference on Cyber Security and Resilience (CSR)*, pages 89–94. IEEE, 2021.
- [15] Luca Deri, Maurizio Martinelli, and Alfredo Cardigliano. Realtime high-speed network traffic monitoring using ntopng. In *Proceedings of the 28th USENIX Conference on Large Installation System Administration, LISA'14*, page 69–79, USA, 2014. USENIX Association.
- [16] Christian J Dietrich and Christian Rossow. Empirical research of ip blacklists. In *ISSE 2008 Securing Electronic Business Processes: Highlights of the Information Security Solutions Europe 2008 Conference*, pages 163–171. Springer, 2009.
- [17] Dshield. Top 20 attacking class C subnets over the last three days. [https://raw.githubusercontent.com/firehol/blocklist-ipsets/master/dshield\\_1d.netset](https://raw.githubusercontent.com/firehol/blocklist-ipsets/master/dshield_1d.netset), 2023.
- [18] Mark Felegyhazi, Christian Kreibich, and Vern Paxson. On the potential of proactive domain blacklisting. *LEET*, 10:6–6, 2010.
- [19] FireHOL. All Cybercrime IP Feeds. <http://iplists.firehol.org>, 2023.
- [20] Julien Freudiger, Emiliano De Cristofaro, and Alejandro E Brito. Controlled data sharing for collaborative predictive blacklisting. In *Detection of Intrusions and Malware, and Vulnerability Assessment: 12th International Conference, DIMVA 2015, Milan, Italy, July 9-10, 2015, Proceedings 12*, pages 327–349. Springer, 2015.
- [21] Sebastian Garcia. Modelling the network behaviour of malware to block malicious patterns. the stratosphere project: a behavioural ips. *Virus Bulletin*, pages 1–8, 2015.
- [22] Ibrahim Ghafir, Vaclav Prenosil, Mohammad Hammoudeh, Liangxiu Han, and Umar Raza. Malicious ssl certificate detection: A step towards advanced persistent threat defence. In *Proceedings of the international conference on future networks and distributed systems*, 2017.
- [23] GreyNoise. IP Lookup. <https://viz.greynoise.io/ip/>, 2023.

- [24] Baktagul Imasheva, Nakispekov Azamat, Andrey Sidelkovskiy, and Ainur Sidelkovskaya. The practice of moving to big data on the case of the nosql database, clickhouse. In *Optimization of Complex Systems: Theory, Models, Algorithms and Applications*, pages 820–828. Springer, 2020.
- [25] Stratosphere IPS. Daily Blacklist. [https://mcfp.felk.cvut.cz/publicDatasets/CTU-AIPP-BlackList/Todays-Blacklists/AIP\\_historical\\_blacklist\\_prioritized\\_by\\_newest\\_attackers.csv](https://mcfp.felk.cvut.cz/publicDatasets/CTU-AIPP-BlackList/Todays-Blacklists/AIP_historical_blacklist_prioritized_by_newest_attackers.csv), 2023.
- [26] Marc Kührer and Thorsten Holz. An empirical analysis of malware blacklists. *PIK-Praxis der Informationsverarbeitung und Kommunikation*, 35(1):11–16, 2012.
- [27] Marc Kührer, Christian Rossow, and Thorsten Holz. Paint it black: Evaluating the effectiveness of malware blacklists. In *Research in Attacks, Intrusions and Defenses: 17th International Symposium, RAID 2014, Gothenburg, Sweden, September 17-19, 2014. Proceedings 17*, pages 1–21. Springer, 2014.
- [28] Stratosphere Lab. Blocklist Generation Project. <https://mcfp.felk.cvut.cz/publicDatasets/CTU-AIPP-BlackList>, 2023.
- [29] Stratosphere Research Laboratory. AIP Tool. <https://github.com/stratosphereips/AIP>, 2022.
- [30] Abuse IP Labs. Abuse IP Lists. <https://www.abuseipdb.com/sitemap>, 2023.
- [31] Suhas Mathur, Baris Coskun, and Suhrid Balakrishnan. Detecting hidden enemy lines in ip address space. In *Proceedings of the 2013 New Security Paradigms Workshop*, pages 19–30, 2013.
- [32] Toni Mattis, Johannes Henning, Patrick Rein, Robert Hirschfeld, and Malte Appeltauer. Columnar objects: Improving the performance of analytical applications. In *2015 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!)*, pages 197–210, 2015.
- [33] Pavel Novák and Václav Oujezský. Detection of malicious network traffic behavior using ja3 fingerprints.
- [34] David Oro, Jesus Luna, Toni Felguera, Marc Vilanova, and Jetzabel Serna. Benchmarking ip blacklists for financial botnet detection. In *2010 Sixth International Conference on Information Assurance and Security*, pages 62–67. IEEE, 2010.
- [35] Proofpoint. Emerging Threats fwip rules. <https://https://rules.emergingthreats.net/fwrules/emerging-Block-IPs.txt>, 2023.
- [36] Gregorius Aldo Radityatama, Charles Lim, and Heru Purnomo Ipung. *A Study Of Enterprise Software Support On nDPI*. PhD thesis, Swiss German University, 2017.
- [37] Sivaramakrishnan Ramanathan, Jelena Mirkovic, and Minlan Yu. Blag: Improving the accuracy of blacklists. In *NDSS*, 2020.
- [38] Steve Sheng, Brad Wardman, Gary Warner, Lorrie Cranor, Jason Hong, and Chengshan Zhang. An empirical analysis of phishing blacklists. 2009.
- [39] Joel Sommers. pytricia: an IP address lookup module for Python. <https://github.com/jsommers/pytricia>, 2022.
- [40] Udger s.r.o. LLC. IP Address and DataCenter Lists. <https://udger.com/resources>, 2023.
- [41] Matija Stevanovic, Jens Myrup Pedersen, Alessandro D’Alconzo, and Stefan Ruehrup. A method for identifying compromised clients based on dns traffic analysis. *International Journal of Information Security*, 16:115–132, 2017.
- [42] Blake E Strom, Andy Applebaum, Doug P Miller, Kathryn C Nickels, Adam G Pennington, and Cody B Thomas. Mitre att&ck: Design and philosophy. In *Technical report*. The MITRE Corporation, 2018.
- [43] Abidah Mat Taib, Mohammad Fikri Hanif Ishak, and Nur Khairani Kamarudin. Securing network using raspberry pi by implementing vpn, pi-hole, and ips (vpisec). *International Journal*, 9(1.3), 2020.
- [44] Tran Phuong Thao, Akira Yamada, and Ayumu Kubota. Empirical analysis of domain blacklists. *International Journal On Advances in Security*, 11(1&2):127–137, 2018.
- [45] Northwestern University. Developing a neighborhood watch for the Internet. <https://www.northwestern.edu/newscenter/stories/2008/11/internetwatch.html>, 2008.
- [46] J. van der Velden. Blacklist, do you copy? characterizing information flow in public domain blacklists, January 2020.
- [47] VirusTotal. IP Address Report. <https://www.virustotal.com/gui/home/search>, 2023.
- [48] Chun Wei, Alan Sprague, Gary Warner, and Anthony Skjellum. Identifying new spam domains by hosting ips: improving domain blacklisting. *Proc. of 7th CEAS*, 2010.
- [49] Yinglian Xie, Fang Yu, Kannan Achan, Eliot Gillum, Moises Goldszmidt, and Ted Wobber. How dynamic are ip addresses? In *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 301–312, 2007.
- [50] Asma Zahra and Munam Ali Shah. Iot based ransomware growth rate evaluation and detection using command and control blacklisting. In *2017 23rd international conference on automation and computing (icac)*, pages 1–6. IEEE, 2017.