Can Public IP Blocklists Explain Internet Radiation?

Simone Cossaro
University of Trieste
simone.cossaro@studenti.units.it

Damiano Ravalico
University of Trieste
damiano.ravalico@phd.units.it

Rodolfo Vieira Valentim
University of Turin
rodolfo.valentim@unito.it

Martino Trevisan
University of Trieste
martino.trevisan@dia.units.it

Idilio Drago
University of Turin
idilio.drago@unito.it

***Abstract*—Network telescopes (IP addresses hosting no services) are valuable for observing unsolicited Internet traffic from scanners, crawlers, botnets, and misconfigured hosts. This traffic is known as Internet radiation, and its monitoring with telescopes helps in identifying malicious activities. Yet, the deployment of telescopes is expensive. Meanwhile, numerous public blocklists aggregate data from various sources to track IP addresses involved in malicious activity. This raises the question of whether public blocklists already provide sufficient coverage of these actors, thus rendering new network telescopes unnecessary. We address this question by analyzing traffic from four geographically distributed telescopes and dozens of public blocklists over a two-month period. Our findings show that public blocklists include approximately 71% of IP addresses observed in the telescopes. Moreover, telescopes typically observe scanning activities days before they appear in blocklists. We also find that only 4 out of 50 lists contribute the majority of the coverage, while the addresses evading blocklists present more sporadic activity. Our results demonstrate that distributed telescopes remain valuable assets for network security, providing early detection of threats and complementary coverage to public blocklists. These results call for more coordination among telescope operators and blocklist providers to enhance the defense against emerging threats.**

I. Introduction

Network telescopes are IP addresses publicly announced on the Internet but hosting no services [1, 2]. The traffic they receive is thus unsolicited and includes a mixture of packets coming from misconfigured hosts, backscattering (i.e., response traffic from hosts receiving packets with spoofed source addresses) as well as crawlers and network scans. The latter is composed of both legitimate scans performed by security companies [3] as well as malicious ones performed by botnets in the search for vulnerable hosts. This unsolicited traffic is known as *Internet Background Radiation* in the literature [4]. Telescopes are used to capture Internet radiation for multiple network security tasks, both by companies that distribute IP reputation lists [3] and by researchers. Examples of research applications include the study of (i) DDoS attacks [5], (ii) Internet censorship [6], (iii) large-scale scanning [7], and (iv) botnet activities [8].

Traditional telescopes collect Internet radiation traffic from completely dark address spaces [4, 9]. Yet, some recent work has explored alternative telescope designs. Authors of [10] leverage CDN infrastructure to study Internet radiation traffic reaching the distributed replica servers, showing how this traffic differs from traditional telescopes. Attracted by the live CDN nodes, attackers target the nodes with a variety of attacks not observed in classic telescopes. DScope [11] introduces cloud-native telescope deployments, while systems like Spoki [12] and others [13, 14] augment telescopes with the ability to respond to some incoming requests through honeypots.

Regardless of its type, a telescope deployment is expensive. IPv4 addresses are a scarce resource that can hardly be spared for such a monitoring infrastructure.[1] Yet, several works [16, 3, 10] have shown that the information observed from multiple telescopes is complementary. In other words, distributed telescope deployments increase the visibility of ongoing scanning activities. Large deployments however receive a high volume of unsolicited traffic: processing and explaining such Internet radiation traffic becomes a complex task, calling for advanced algorithms and additional measurements for understanding the possible attacks behind the traffic [17, 1].

The ultimate goal of running a telescope is to build IP reputation lists—lists of addresses engaging in particular activities, such as network scans, brute-force attempts, etc. However, multiple Cyber Threat Intelligence (CTI) sources do distribute *public blocklists* already. There are hundreds of lists that aggregate data from various sources, including spam monitoring systems, antivirus and anti-malware software, honeypots as well as private telescopes. These blocklists are often used by network administrators to block traffic and anticipate (or mitigate) cyber-attacks [18].

The challenges and costs for deploying telescopes raise a major question: *Do public blocklists provide sufficient coverage of the scanning activity observed in new telescopes?* Answering this question is instrumental to the deployment of new telescopes. If public blocklists adequately represent the interesting events seen telescope traffic, the burdens and costs of operating new telescopes could be avoided. Conversely, if new telescopes observe a considerable amount of interesting events that are absent from blocklists, they could be used to complement blocklists, reinforcing the call for data exchange and coordination among telescope operators [16, 3, 10].

We thus investigate to what extent IP addresses seen in

Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb) 2025
28 February 2025, San Diego, CA, USA
ISBN 979-8-9894372-9-0
https://dx.doi.org/10.14722/madweb.2025.23009
www.ndss-symposium.org

[1] We ignore IPv6 telescopes, which usually receive little traffic [15].

public blocklists overlap with those scanning telescope IP ranges, completing our previous work [19]. Specifically, we answer the following research questions (RQs):

RQ1 To what extent IP addresses scanning telescope IP ranges are reported in public blocklists?
RQ2 Considering the overlap between blocklists and telescopes, are the reporting of new IP addresses synchronized?
RQ3 How do different blocklists cover the IP addresses scanning telescope ranges?
RQ4 What are the IP addresses seen in telescopes but not in blocklists?

To answer these questions, we collect data from four geographically distributed telescopes, located in Europe and South America, over a two-month period. Simultaneously, we daily gather data from dozens of public blocklists, which vary in size, update frequency and methodology to compose the lists. The telescopes are heterogeneous not only in terms of geographic location but also in terms of IP ranges. They capture a broad spectrum of traffic, receiving packets from approximately $209\,643$ IP addresses per day in median (considering only scanning traffic), with a daily median of $8.3\,\mathrm{GB}$ of traffic. We evaluate the overlap between the IP addresses observed in the telescopes and those reported in the blocklists, considering temporal aspects and regional variations.

Our findings can be summarized as follows:

- Public blocklists only partially cover the scanning traffic seen in telescopes. Daily, $80\,\%$ of the IP addresses seen in telescopes are either in blocklists or are well-known benign scanners. The percentage of telescope traffic that can be explained with those lists ranges from $58\,\%$ to $81\,\%$.
- Telescopes usually observe the activity of IP addresses earlier than they are reported in blocklists. Approximately $5\,\%$ of scanners appear in the blocklists 1-5 days *after* they are observed in the telescopes.
- The vast majority of the coverage is contributed by only 4 out of 50 blocklists.
- The IP addresses found in telescopes but evading blocklists vary quickly and likely belong to compromised user devices or servers.

The remainder of the paper is organized as follows. Section II describes our datasets and methodology. Section III discusses blocklist coverage, while Section IV quantifies the delay of blocklists with respect to telescopes. Then, Section V evaluates the effectiveness and specificity of individual blocklists in enumerating scanners, while in Section VI, we investigate the sources observed in the telescopes that evade blocklists. Finally, Section VII summarizes the related work, and Section VIII concludes the paper.

II. Datasets and methodology

We have collected data for more than 9 weeks from both a distributed telescope infrastructure and public blocklists, specifically from July 5, 2024, to September 10, 2024.[2] Next, we describe our distributed telescope infrastructure and the blocklists used in our analysis.

[2]Due to a malfunction in the collection, data is unavailable for the period between August 25, 2024 and August 28, 2024

TABLE I: Overview of the telescope deployments.

	Location	Subnet Size	Subnets	Reserved IPs	Total IPs
$\mathcal{T}_0$	South America	/19	1	0	8190
$\mathcal{T}_1$	Europe	/24	2	11	497
$\mathcal{T}_2$	Europe	/24	1	2	252
$\mathcal{T}_3$	Europe	/22	1	0	1022

A. Telescopes

We rely on a distributed telescope infrastructure consisting of non-continuous /24 networks deployed across Europe, while South America is represented by a single /19 network. For privacy reasons, we do not disclose specific IP addresses of the telescopes. Table I summarizes the infrastructure in terms of size. The "Reserved IPs" column indicates the number of addresses allocated for other services in each telescope—they are addresses for which we do not register the eventual unsolicited traffic. Network and broadcast addresses are also excluded from the captures. The smallest telescope, $\mathcal{T}_2$, consists of a single /24 subnet with 2 reserved addresses, resulting in a total of 252 free IP addresses. In contrast, the largest telescope, $\mathcal{T}_0$, is a single /19 subnet with no reserved addresses, providing a total of 8190 IP addresses.

These telescopes are continuously monitored using network probes that record all incoming packets. This setup allows us to collect a comprehensive, geographically distributed dataset, enabling the analysis of unsolicited traffic patterns. None of these telescopes have hosted public services in recent years, with the exception of $\mathcal{T}_2$. The $\mathcal{T}_2$ telescope has been previously used for hosting production services. It is known from previous work [16, 5] that scanners more often target IP addresses seen online, and this pattern persists even after the target host goes offline. As such, the traffic observed in this telescope may differ from the others.

Blocklists are expected to report hosts performing malicious activities only, whereas telescopes observe a large number of packets coming from actual *victims* of attacks—the so-called backscattering phenomenon. In these cases, attackers (usually performing DDoS) send large numbers of spoofed packets to victims. When attackers casually pick a telescope address as a spoofed source, the telescope may receive response packets from the victim. These packets must be ignored in our analysis since the IP addresses of these victims are not expected to be present in blocklists.

To filter out these packets, we discard all non-TCP packets and all TCP packets that are not *pure* SYN packets, i.e., any packet with other TCP flags. Moreover, following the approach used in [17, 1], we filter out IP addresses that send fewer than 5 packets per /24 in a telescope, e.g., 20 packets for the /22 telescope. We thus present results considering only the most active IP addresses reaching the telescopes, ignoring any packet compatible with backscattering, even if this filtering approach may drop some malicious scanning traffic. We will show later that this filtering approach does not change our conclusions.

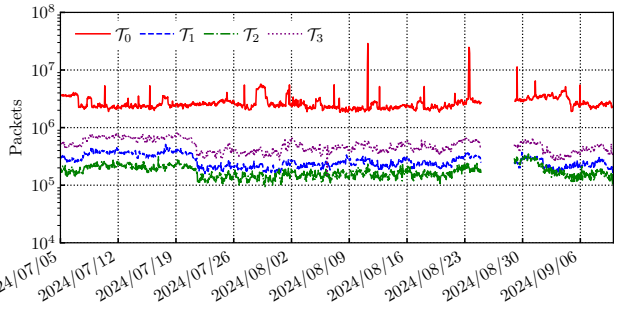

Fig. 1: Number of packets arriving at each telescope.

TABLE II: Overview of the telescope dataset.

Metric	Cumulative	Median daily
Volume of Data	542 GB	8.3 GB
Total Packets	5 673 137 282	84 076 822
Filtered Packets	4 912 714 466	71 456 394
TCP Traffic	92.89 %	92.51 %
UDP Traffic	5.92 %	6.12 %
Unique IP Addresses	4 964 306	209 643
Unique IP Addresses (Filtered)	252 879	22 087

As shown in Figure 1, the traffic volume varies across telescopes and is roughly proportional to the size of the network as reported in previous work [16]. For all four deployments, the figure shows well-known patterns related to unsolicited traffic: a baseline of continuous scan noise and some sporadic peaks of traffic. These episodic events are usually large-scale scans or attacks, often launched from botnets. Interestingly, these peaks are not synchronously seen in all telescopes.

We provide overall statistics on the four telescope datasets in Table II. The telescopes observe a significant traffic volume, totaling 542 GB and 5.6 billion packets. The dataset is largely composed of TCP traffic, i.e., 92.89 % of the total volume, with UDP traffic making up 5.92 %; the remaining traffic is composed of different protocols, e.g., ICMP. In terms of IP addresses, there are nearly 5 million unique addresses, with only 252 879 remaining after applying our filters to drop backscattering. Yet, while we drop approximately 95 % of the total IP addresses, they generate only 12 % of the total traffic.

B. Blocklists

We start from the set of blocklists used by Feal et al. [18], which is a comprehensive previous work. However, due to the evolving nature of online services, some providers are no longer available, calling for an update to the sources. We extend the blocklists with several publicly available ones from the FilterLists[3] online aggregator, a platform that hosts over 300 public lists. These lists span various fields, including malware, phishing sites, ad-blocking, and IP addresses associated with suspicious activity. We here focus on blocklists that provide IPv4 addresses linked to malicious activity such as malware distribution and scans, ignoring all those providing information

[3]https://filterlists.com/

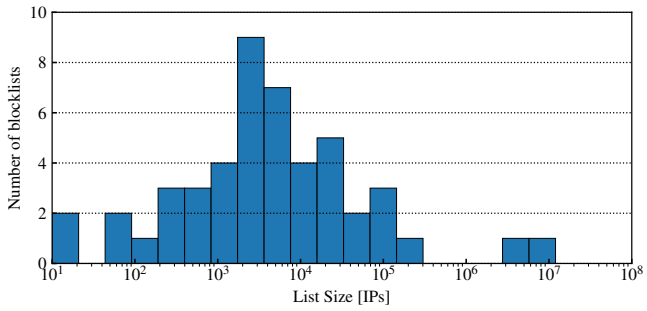

(a) Median size distribution of blocklists.

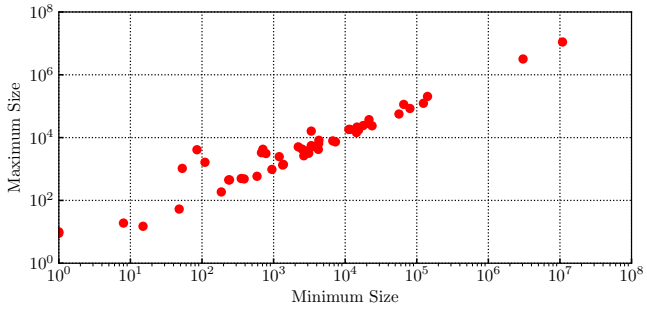

(b) Minimal and maximum daily sizes of each blocklist.

Fig. 2: Size of the 50 selected blocklists (number of IP addresses).

not observed in telescopes, such as hostnames or URLs. The lists are typically formatted as "host files", where each entry corresponds to an individual IP address or an IP range. From these sources, we obtain 50 blocklists providing IP addresses—27 from [18] and 23 from FilterLists. We report the full list in the Appendix. We collect a total of 13.7 GB of blocklist files in our capture period, and they include about 16.5 million unique IP addresses. Due to a change in the FilterLists API, we do not have information for 23 blocklists from July 24, 2024, to August 28, 2024. We will show later that this outage has negligible impact on conclusions as those 23 blocklists usually provide minor coverage of the IP addresses seen in telescopes.

Figure 2 illustrates the size of the selected blocklists in terms of IP addresses observed per day. In the top plot, we show the distribution of the median number of IP addresses reported by each blocklist. The distribution has a bell shape with most lists containing something between 1000 and 10 000 addresses, with some few lists appearing as outliers with less than 10 or more than 1 million IP addresses. Therefore, blocklist sizes vary significantly, with some providers offering small, curated lists, while others provide large datasets that may include less reliable entries. In the bottom plot, we show how the daily sizes of blocklists vary during the period using a scatter plot of the minimum and maximum size of each blocklist. The size of most lists is rather constant, as illustrated by the points along the diagonal. Yet, we see that some lists change substantially, with hundreds of addresses added and/or removed during specific days of our data capture.

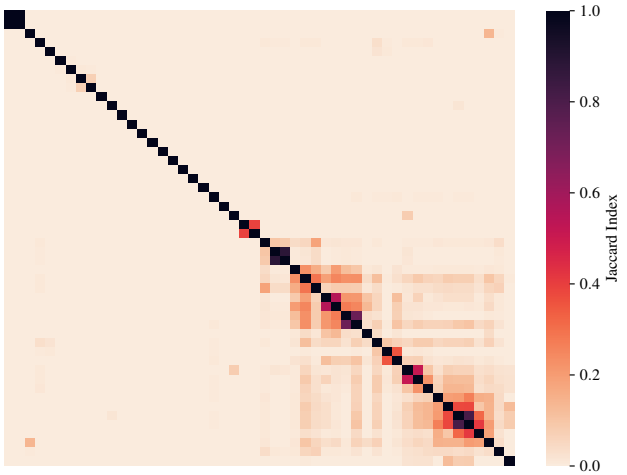

Fig. 3: Daily median pairwise Jaccard index between blocklists.

The smallest list in terms of size is Feodo Tracker IoC, with a median of a single IP address per day and a total of 10 addresses over the data capture period. Spamhaus is the largest blocklist and it includes a median of slightly more than 11 million addresses per day, of which 99.98 % are not present in any other blocklist. This list proved to be highly static, with few additions and rare updates throughout our analysis. Some blocklists, on the other hand, are very dynamic. Considering Nix Spam and IPsum Level 6, for example, approximately one-third of the addresses are renewed each day.

We also analyze the overlap between blocklists by calculating the pairwise Jaccard Index to quantify the similarity between the sets of IP addresses reported by different blocklists in each day of our data capture. We compute the median Jaccard Index over all days for each list pair and show results in Figure 3. The names of the blocklists are omitted to improve visualization.

We observe two major clusters for which the Jaccard index is high, indicating a significant degree of overlap in the IP addresses contained in these lists. Blocklists in these areas may be targeting similar types of threats or employing analogous criteria for the inclusion of addresses. For example, we confirm that the central region in the figure, where high Jaccard indexes are seen, consists of blocklists from multiple providers containing IP addresses associated with malware distribution. These lists are highly similar, thus resulting in high Jaccard indexes. The region with high Jaccard indexes in the bottom right of the figure is formed by a set of blocklists from Dataplane, a provider included in [18]. Dataplane contributes 12 blocklists to the collection, and 5 of these lists are present in this specific cluster. Although the type of activity considered by Dataplane to include addresses in its various blocklists differs, these 5 lists strongly overlap. Finally, the vast majority of blocklist pairs present Jaccard index values close to zero. That is, the lists have no overlap. These lists may focus on different types of malicious activities, which are performed by different threat actors and, consequently, IP addresses.

All these results are consistent with those presented in [18, 20]. We next extend these results by evaluating the extent to which such blocklists can anticipate and explain traffic seen on the telescopes.

C. Benign Scanners

Besides backscattering, telescopes are constantly reached by benign scanners from companies that actively monitor the Internet. Examples include projects like Shodan and Censys, which catalog and index Internet-connected devices for analysis and research. Instead, blocklists should explain the malicious traffic that reaches telescopes. Thus, besides filtering backscattering, we also identify packets from well-known benign scanners.

Studies such as [21, 22, 23, 24] provide some hints on legitimate scanners. Unfortunately, there exists no exhaustive public list of those scanners, and not all of them publish the list of IP addresses they use for scanning. Thus, we leverage the dataset made available by the authors of [22] through the "Acknowledged Scanners" GitLab repository[4] as an initial reference point. This resource catalogs IP addresses linked to scanning activities recognized as non-malicious or neutral, such as academic research projects and security assessment initiatives as well as services like the just-mentioned Censys.

This source offers IP addresses belonging to 41 different scanners, updated on different dates. Updating this data is not straightforward, as many of the scanning services are either fee-based or do not provide information on their IP address ranges. We manually updated 8 out of these 41 lists on June 26, 2024, obtaining in total a snapshot of benign scanners with 12 725 IP addresses. Interestingly, we observe that 3385 acknowledged scanner IP addresses appear in at least one blocklist—even if their activity is supposedly benign. Thus, approximately 27 % of our initial set of benign scanners is present in the blocklists too.

III. Blocklists Coverage

We start by examining the coverage that blocklists provide for telescope traffic, thus answering our **RQ1**. We define *coverage* as the percentage of IP addresses (or percentage of their packets) observed in a given telescope that is also reported in at least one blocklist. Given that many acknowledged scanners are also present in blocklists (see Section II), we include the full set of acknowledged benign scanners in this analysis as an additional list.

We run the analysis on a day-by-day basis as we are interested in temporal patterns. For each day, we use the packets captured by the telescopes on the given day and the snapshot of blocklists downloaded on the same date.

Figure 4a presents the time series of the IP address coverage separately by telescope. The x-axis denotes the date, while the y-axis represents the corresponding coverage. The y-axis is zoomed in for clearer visualization. The coverage in terms of IP addresses is indeed very consistent across telescopes: even daily fluctuations, such as those seen around 26 July 2024, are somewhat replicated across the four deployments. Some of the drops in the series appear to be the beginning of new scanning,

[4]https://gitlab.com/mcollins_at_isi/acknowledged_scanners/

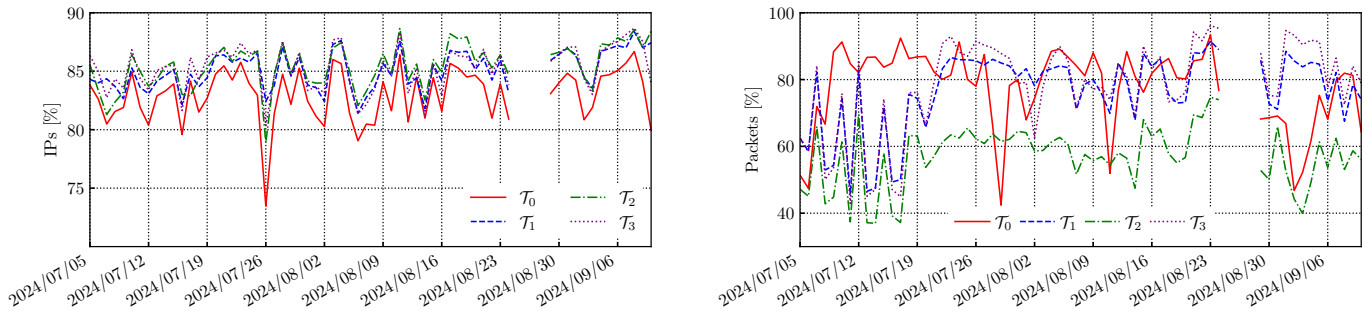

(a) IP

(b) Traffic

Fig. 4: Daily coverage of telescope traffic considering IP addresses present either in blocklists or in our list of acknowledged benign scanners.

which is reported in blocklists with some delay. We will study this phenomenon in the coming sections.

By computing the distribution of the points seen in the figure for each telescope, we find that the median daily coverage across the entire data collection period for $\mathcal{T}_0$, $\mathcal{T}_1$, $\mathcal{T}_2$, and $\mathcal{T}_3$ is $83.57\,\%$, $84.94\,\%$, $85.70\,\%$, and $85.61\,\%$, respectively.

Not shown in the figure, the proportion of IP addresses associated with the acknowledged benign scanners is low and, in median, it varies from $7.38\,\%$ (in $\mathcal{T}_0$) to $11.01\,\%$ in $\mathcal{T}_2$. These differences across telescopes are indeed expected and reflect the fact that different IP address ranges receive traffic from different scanning sources [16]. More interestingly, the low percentage of benign scanning traffic shows that telescopes still capture a large number of attempts by possibly malicious actors too.

Interestingly, we also evaluate the intersection of IP addresses covered by blocklists on different days. Although the daily coverage is on the order of $80\,\%$ to $85\,\%$ (see Figure 4), overall only $71.24\,\%$ of the entire set of IP addresses appear at least once in a blocklist. This happens because there is a high dynamism of addresses evading blocklists. They tend to be different from day to day, whereas a significant share of the ones appearing on blocklists reappear on subsequent days, i.e., blocklists cover the *"usual suspects"* while missing the dynamic sources.

To complete the analysis, we report in Figure 4b the packet coverage, i.e., the share of packets received by the telescopes whose source IP address is in at least one blocklist or in our list of acknowledged scanners. Figure 4b shows the results again separately for the four telescopes. The x-axis represents the date, while the y-axis shows the percentage of packets covered by the lists. Notice the different y-axis limits when compared to Figure 4a.

Here, a much noisier figure emerges. $\mathcal{T}_1$ and $\mathcal{T}_3$ appear to show a similar trend, particularly in the initial two weeks of our monitoring. $\mathcal{T}_2$ shows slightly lower percentages. The coverage in all four telescopes is however very volatile, characterized by frequent peaks and drops. Overall, no clear trends can be marked. When calculating the distribution of the points in the figure, the median coverage for $\mathcal{T}_1$, $\mathcal{T}_2$ and $\mathcal{T}_3$ is $80.48\,\%$, $58.35\,\%$, and $79.68\,\%$, respectively. $\mathcal{T}_0$ has the highest median coverage at $81.21\,\%$.

In a nutshell, while we observe some patterns in terms of IP address coverage, which is around $80\,\%$ to $90\,\%$, packet coverage is a lot noisier. In other words, the blocklists and the list of benign scanners do provide somewhat good coverage of the IP addresses that will be observed scanning a network in a day. They, however, miss a significant percentage of the addresses (up to $20\,\%$), which often are the ones responsible for a large share of the scanning traffic.

IV. Reporting Delays

Next, we investigate to what extent telescopes provide an earlier view of scanning activity with respect to blocklists, thus answering our **RQ2**. This is an important metric as it provides a view of the time networks could remain exposed to malicious scanning if relying solely on blocklists.

We calculate the reporting delays as follows. We use the set of IP addresses contacting the telescope on a specific date as the reference set. Then, we measure the overlap of this reference set with the cumulative set of IP addresses formed by the union of all blocklists (augmented with the benign scanners) in the i days following the reference set, varying i from 0 to 30.[5] Clearly, when $i = 0$, we are measuring the instantaneous coverage, as already reported in Figure 4. We repeat this procedure for 30 different reference days, starting from July 5, 2024, and for each reference day, we compute the overlap over the subsequent i-day period, with i from 0 to 30 for all cases. In the end, we compute the average statistics, considering the 30 reference days. We show results in Figure 5.

Focus first in Figure 5a. We report i on the x-axis, i.e., the number of days after the reference day, during which we accumulate blocklists. The y-axis indicates the coverage as the percentage of addresses observed in telescopes and found in the (accumulated) blocklists. Note again the zoomed y-axis for better visualization. The upward trend in the lines shows that IP addresses initially unlisted in blocklists are gradually included in blocklists in the subsequent days. The coverage

[5]As we download blocklists once a day, our finer granularity for the reporting delays is 1 day.

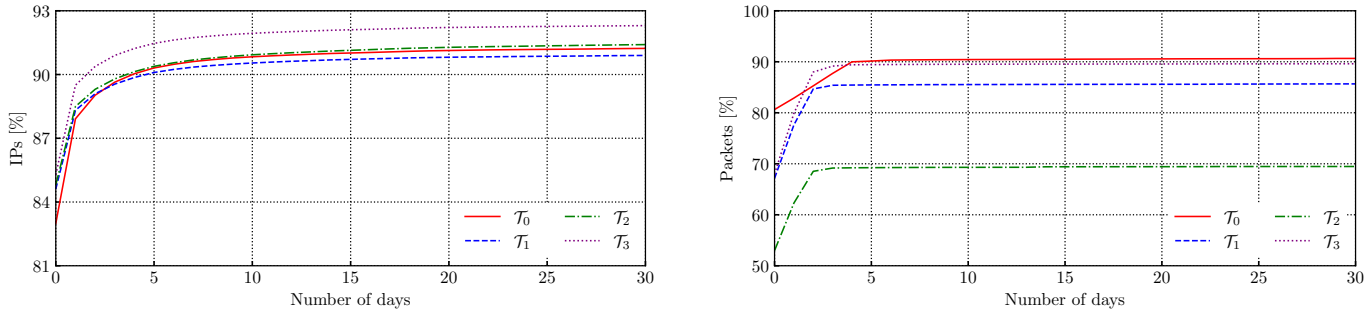

(a) IP addresses

(b) Traffic

Fig. 5: Reporting delay of blocklists with respect to the appearance of the IP addresses in the telescopes.

grows from around 82 % to 85 % on average on the day when IP addresses are observed in the telescopes, to more than 90 % after some days. This result shows that telescopes do observe some addresses earlier, and a significant share of IP addresses are added to blocklists only some days after they first contact the telescopes.

This trend is consistent across different telescopes, irrespective of size, with only $\mathcal{T}_3$ showing a slightly higher percentage. In all cases, coverage increases to around 90 % within 5 days and stabilizes between 91 % to 92 % after about 20 days.[6] Therefore, about 50 % of IP addresses seen by the telescopes but not immediately present in blocklists will appear in the lists within a 20-day window. Notably, the curves never reach 100 %, eventually saturating. Based on manual inspection of the never-reported addresses, we conjecture that this saturation can hardly be linked to benign scanners, which should not be in blocklists. For example, some of the remaining addresses show a behavior compatible with malicious activity, with clear fingerprints and scanning patterns of known botnets such as Mirai.

In Figure 5b we consider how coverage in terms of traffic volume varies according to the reporting delay. We observe an increase of approximately 15 % to 20 % in volume-wise coverage within the first 5 days after IP addresses appear in the telescopes.[7] As commented before, the coverage in terms of traffic varies considerably across telescopes, with $\mathcal{T}_0$ showing the highest values and $\mathcal{T}_2$ the lowest. Notably, all the curves reach saturation sooner compared to those in Figure 5a, practically becoming stable already after 5 days. Apparently, IP addresses generating higher volumes of traffic are included in blocklists more promptly and, consequently, there is a rapid increase in traffic coverage, followed by stagnation, even as more addresses continue to be added to the blocklists.

These large delays suggest that some IP addresses may be included in blocklists when it is already too late, i.e., when their activity has already ceased. To further investigate this behavior, we focus on the *top talkers*, i.e., the IP sources that most frequently contact the telescopes, generating the highest number of packets. For each telescope, we arbitrarily select the top 150 daily most active IP addresses, and then merge all found addresses in a single set, resulting in a total of 2509 unique IP addresses. We then track the activity of these top talkers over the entire data collection period.

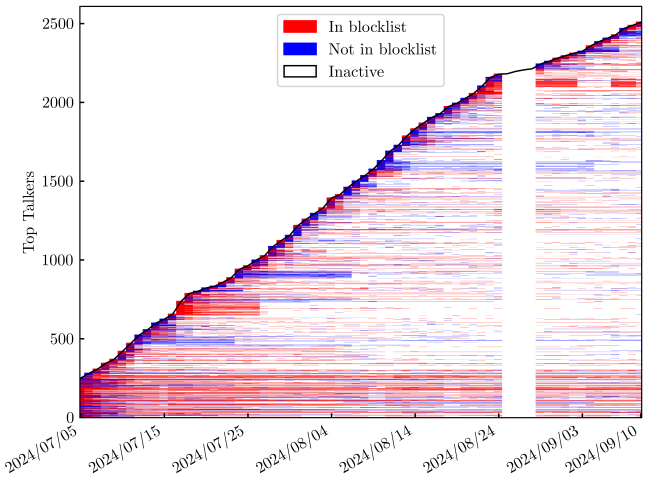

Fig. 6: Activity of the *top talkers* during data collection period.

In Figure 6, we illustrate the full activity of these top talkers. The x-axis represents the date, while the y-axis marks the top talkers sorted by their first appearances in the telescopes. A point is included in the figure when the top talker sends packets to the telescopes on a given day. The color of the point indicates the status of the IP address: red if it is listed in any blocklists, blue if it is not listed. Generally, we observe more density of points around the first moment in which the IP address appears in the telescopes; that is, the addresses are more active when they engage in scanning activity. However, we see that some IP addresses continue to reach the telescopes throughout the entire data collection period. We observe in particular some extremely active sources (around 130) that remain consistently active, contacting the telescope every day throughout the entire data collection period, represented by the lower portion of the figure. Among these examples, only 38 are continuously listed—and never removed—from the blocklists.

[6] $\mathcal{T}_2$ has mean standard deviation of 1.07 %, whereas for $\mathcal{T}_0$, $\mathcal{T}_1$, and $\mathcal{T}_3$, the values are slightly lower, at 0.77 %, 0.69 %, and 0.59 %, respectively.

[7] The mean standard deviations in this case are: $\mathcal{T}_3 = 4.93\,\%$, $\mathcal{T}_0 = 4.55\,\%$, $\mathcal{T}_1 = 3.48\,\%$ and $\mathcal{T}_2 = 3.87\,\%$.

Overall, around 40 % (1014) of the top talkers are never included in any blocklist, and they can be identified by horizontal lines that alternate blue and white. Interestingly, 45 of them are active every day and yet are never included in blocklists. Focusing on the remaining 60 % (1495) of sources that are eventually listed in blocklists at least once, we find that they have a median active period of 4 days. On their first day of appearance, around 53 % (1327) of top talkers are *not* listed in any blocklist, with some of them (313) eventually being included in the blocklists within a median of 4 days. Among those, 288 were also dropped from the blocklists during our capture. This effect is also observed for the remaining 47 % (1182) of sources that are already listed in a blocklist on the first day they contact the telescopes: among these, 1087 are later removed from the blocklists, even if 605 reappear in the blocklists and eventually also in the telescopes.

Overall, a significant number of IP addresses observed in telescopes are reported in blocklists with significant delays, or never reported at all. Among those that eventually get reported, an important share is removed from the blocklists after some time even if they sometimes still appear in telescopes. This suggests that blocklists are not always effective in capturing the most active scanners and that relying solely on them may lead to missing some of the most aggressive scanners.

V. Effectiveness of Blocklists

We now move to **RQ3** and individually evaluate the effectiveness of each blocklist in detecting the scanners that we observe in the telescopes. We are interested in quantifying both (i) the *completeness* of each blocklist in terms of coverage it provides and (ii) its *specificity* in reporting active scanners. We quantify these two aspects by borrowing ideas from the recall and precision metrics commonly employed in machine learning. While leveraging the terminology, we define the metrics in our context as follows:

Recall The portion of IP addresses seen in the telescopes that are included in the given blocklist; it measures the *completeness* (or *coverage*) of the blocklist.

Precision The portion of IP addresses reported in the blocklist that is observed in the telescopes; it measures the *specificity* (or *efficiency*) of the blocklist.

A precision of 100 % would indicate that all IP addresses inserted in the blocklist are active in the telescopes. Achieving 100 % recall would mean that all sources observed in the telescopes are listed on the blocklist.

Figure 7 illustrates the metrics for all blocklists. The x-axis represents the recall of the blocklists (in percentage), while the y-axis represents the precision of the blocklists (also in percentage). Each point corresponds to a blocklist, with the point color indicating the number of IP addresses listed over the entire data collection period. Labels identify blocklists with distinctive characteristics, some of which we will comment on next.

We observe that 58 % (29) of the blocklists have both precision and recall below 20 %, reflecting their minimal contribution to explaining telescope traffic (see also Figure 4). Indeed, at least 80 % of the IP addresses reported in these blocklists show no activity in the telescopes. In some cases,

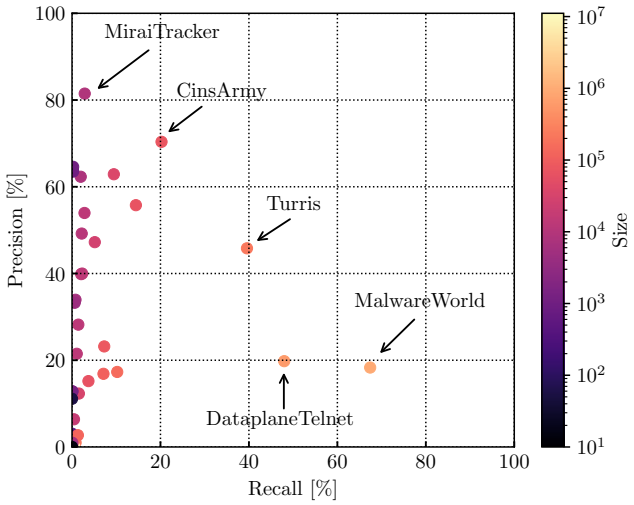

Fig. 7: Recall and precision of blocklists (colors indicate their sizes).

these low metrics are somewhat expected, given the methodologies used to build the blocklists, e.g., spam emails, which may be performed by specific IP addresses. Yet, these results show that using these lists to block malicious network activity is not effective, particularly when considering network scans.

The majority of blocklists are concentrated along the y-axis, indicating very low recall—92 % have a recall below 20 %. However, in some cases, they present a high precision—17 of them have a precision higher than or equal to 20 %. Among these, the MiraiTracker blocklist (IP addresses suspected to belong to the Mirai botnet) achieves the highest precision but has a recall of only 2 %. These blocklists effectively capture entries observed in telescopes because they focus on a single phenomenon that is also observed in telescopes, e.g., the activity of the Mirai botnet. Thus, they may be very effective when the goal is to block such specific events. As many other events are also observed in telescopes (e.g., traffic from multiple other botnets), they offer insufficient coverage of scanning activities or emerging threats in general.

The remaining 8 %—4 blocklists—achieve recall values ranging from 20 % (CinsArmy) to 67 % (MalwareWorld). Interestingly, this higher recall comes with lower precision. In general, the distribution of colors in the figure, representing blocklist sizes, shows that there is an interesting trend between the size of the blocklist and the precision and recall. Larger blocklists generally demonstrate lower precision, whereas smaller blocklists tend to exhibit low recall. There are cases where both precision and recall are exceptionally low; for instance, Spamhaus—the largest blocklist—achieves a precision of 0.01 % and a recall of only 0.5 %.

This pattern appears to reflect a distinction between two types of lists: blocklists with high recall prioritize coverage and comprehensiveness, while those with higher precision focus on quality and accuracy for specific events. Apparently, the former lists achieve higher recall by including IP addresses using relaxed policies and/or aggregating many sources of different nature. Many of these IP addresses are never active

TABLE III: Characterization of IP addresses found in the telescopes but evading blocklists.

		IP addresses	Volume
RDNS	Scanners	188	8.79 %
	Broadband Subscriptions	8915	1.17 %
	Cloud Providers	5410	6.67 %
MaxMind	ChinaNet	10 809	2.09 %
	China Unicom	5249	1.53 %
	DigitalOcean	4129	4.59 %
	Other	37 905	74.98 %
Unknown		133	0.18 %
Total Addresses		72 738	

in the telescopes. Indiscriminately increasing blocklists may result in false positives when they are used to block traffic in production or can pose unnecessary load on middleboxes filtering traffic (e.g., firewalls or routers). Recall, for example, that by considering the union of all 50 blocklists, we also observe 27 % of the IP addresses associated with benign scanners in the blocklists.

All in all, we see a clear distinction between lists that include hundreds of thousands of IP addresses, achieving high recall but low precision, and those including only a few addresses, instead achieving low recall but high precision. While we selected 50 blocklists for our study, only 4 lists provide significant coverage of the telescope traffic (high recall). Some other lists instead provide good precision for some types of events, such as IP addresses associated with specific botnets.

VI. Who are the IP Addresses not In Blocklists?

We conclude by investigating addresses that evade blocklists, i.e., are observed in the telescopes but never appear in any blocklist or list of acknowledged scanners, thus answering our **RQ4**. Overall, out of the 252 879 addresses found with the telescopes, only 72 738 (28.76 %) are never present in any blocklist. Recall from the discussion around Figure 5b that blocklists tend to include the most active IP addresses within few days after they first appear in our telescope. As such, IP addresses that still evade blocklists during the whole data capture generate only 10.54 % of the total traffic. We now provide some insights into those evading the blocklists using two methods: (i) Reverse DNS and (ii) the MaxMind Geo IP database, and we present the results in Table III.

As said, we first use Reverse DNS to obtain a possible Fully Qualified Domain Name (FQDN) for the IP addresses. Notice that not all IP addresses have registered a reverse DNS record (specifically a `PTR` Resource Record). Indeed, we attempt the Reverse DNS resolution for all 72 738 and obtain an FQDN for 32 233 of them. We then inspect those FQDNs using regular expressions to match similar FQDNs, and in case they are informative, we classify the corresponding IP addresses into:

- **Additional benign scanner**: an IP address of a benign scanner not included in the previously used lists, for example, `azpdcs47.stretchoid.com`
- **Broadband Subscription**: a public IP address identifying a subscriber's customer-provided equipment (CPE). They likely represent compromised user devices (or even CPEs), for example `X-X-75-189.shatel.ir`
- **Cloud Provider**: an IP address identifying a node (typically a virtual machine) hosted on a well-known cloud provider. They likely identify compromised machines, for example, `amazonaws.com`

With this method, we can categorize 14 513 IP addresses, since for the remaining 17 720 we are not able to gather sufficient information to characterize them. Note that this method requires manual labeling, and we can only speculate on the nature of an IP address. Achieving certainty on the reasons a given IP address is scanning a telescope requires a different approach, which we leave for future work.

As shown in the first three rows of Table III, using Reverse DNS we discover 188 new IP addresses belonging to benign scanners, mostly belonging to Stretchoid and BynaryEdge organizations. More interestingly, we find 8915 IP addresses belonging to telecommunication companies, e.g., `telecomitalia.it` and `telekom.de`. We also find 5410 IP addresses belonging to cloud providers (e.g., `amazonaws.com` or `googleusercontent.com`).

For the remaining 58 225 IP address that we cannot categorize using Reverse DNS, we employ the MaxMind GeoIP Database[8] to retrieve the organization (i.e., the Autonomous System) and the country associated with IP addresses. Not all addresses appear in the MaxMind GeoIP database, and we gather information for 58 092 of them. Thus, for 133 address we cannot get any information and mark them as "Unknown". Looking at the bottom rows of Table III, we find that Chinese organizations dominate the rank. Specifically, ChinaNet and China Unicom, two Chinese Internet Service Providers are ranked in the first two positions, with 10 809 and 5249 IP addresses, respectively. Digital Ocean, a US cloud provider, is ranked third with 4129 addresses. Overall, considering the country indicated by MaxMind, 20 582 of the sources are from China, followed by India (3770) and the USA (3211). Finally, for 133 "Unknown" neither Revers DNS nor MaxMind can give any insight. Although further investigation might be possible (using other IP intelligence tools or looking at BGP data), this would be out of the scope of this paper.

In the last column of Table III, we report the traffic share due to the different categories of IP addresses, i.e., the percentage of packets originated from IP addresses in each category. Notably, while only a few dozen IP addresses belong to the benign scanner category, they account for up to 8.79 % of the total traffic. Interestingly, nodes hosted in cloud providers target the $\mathcal{T}_2$ telescope in particular. Although it is hard to find the cause, this behavior might be explained by the fact that, as discussed in Section II, the IP range of this telescope was previously used to host legitimate services.

In summary, 72 738 IP addresses never appear in any blocklist. Among these, 14 513 are categorized via Reverse DNS, with most belonging to Cloud Providers or Broadband subscriptions, suggesting potentially compromised machines. Additionally, 188 IP addresses are identified as benign scanners

[8] https://www.maxmind.com/en/geoip-databases

not included in the acknowledged scanner set [22], as detailed in Section II. Finally, using MaxMind, we associate these IP addresses with Autonomous Systems and countries, finding China as the most represented one.

VII. Related Work

We present an overview of the state of the art in research on network telescopes and blocklists within the field of cybersecurity. To the best of our knowledge, no prior work has investigated the relationship between telescopes and blocklists.

A. Network telescopes

With advancements in scanning tools, large-scale IPv4 scanning has become widespread and has been analyzed by several studies. For instance, Durumeric et al. [23] characterize scans and scanners using a 5.5 million IP telescope. More recently, Griffioen et al. [25] analyzed 10 years of TCP port scanning over 45 billion packets using a telescope, noting a 30 times increase from 2015 to 2024, with some organizations scanning the entire IPv4 space and all ports by 2024. Differently, Collins et al. [22] focus on acknowledged scanners using 3 telescope /24 networks, observing predictable IP targeting and distinct port preferences compared to unacknowledged scanners, and provided a scanner list, which we employ (and update) for our study.

Recent studies on telescope traffic reveal its complex patterns and implications for cybersecurity. Research highlights the impact of telescope size and geography on traffic characteristics [16], methods for monitoring with sparse IP blocks [3], and clustering approaches to detect Distributed Reflection Denial of Service (DRDoS) and other attack patterns [26, 5]. Dainotti et al. [27] highlight the importance of telescopes by observing and analyzing the scan of the entire IPv4 address space conducted by Sality botnet through a /8 telescope. Enhanced techniques, such as DarkVec, achieve high accuracy in associating IPs with activity types, even discovering new attack groups [17, 1]. Additionally, graph-based models have proven effective for botnet detection [28].

The substantial volume of traffic generated by telescopes has prompted focused efforts in the research community to classify and characterize such traffic systematically. Fachkha et al. [26] developed an approach based on k-means clustering techniques to infer DRDoS attacks by leveraging data from a /13 telescope. Furthering this work, Jonker et al. [5] proposed a framework that characterizes attack patterns, attack targets, and DDoS Protection Services (DPSs) using telescope traffic. Gioacchini et al. [17] advanced this field with the introduction of DarkVec, a method to associate IP addresses with specific activity types on the telescope; it was able to identify known attack patterns but also facilitated the characterization of new, previously unidentified groups of attackers. In a subsequent study in 2023, the same authors presented an improved and more scalable version of DarkVec [1]. Additionally, Bou-Harb et al. [29] utilized telescope data spanning three years to propose a behavioral model for senders, revealing coordinated activities over time. In this paper, we quantify to what extent off-the-shelf blocklists help in this goal.

B. Blocklists

Several studies have evaluated the effectiveness and challenges of blocklists across internet security, investigating coverage, reactivity, inconsistencies, and overlap.

A comprehensive work has been carried out by Feal et al. [18], who analyzed 2093 blocklists from 69 open providers over 6 months, uncovering significant overlaps, frequent propagation of changes among similar lists, unique overlap patterns, and inconsistencies in labeling and classification processes that hinder content interpretation. In our work, we study the same set of blocklists (restricted to the 27 listing IP addresses), complementing with other 23 by FilterList. We find that the latter set adds a median overlap of less than $2\,\%$, with only $1.73\,\%$ of unique IP addresses on average, as most are already listed or identified as scanners. This confirms significant redundancy among the lists.

Similarly, Umizaki et al. [30] assessed 7 Public Blocklist Providers, highlighting a low update frequency in four lists and outdated entries, with only $0.03\,\%$ valid entries in non-updated blocklists. They also revealed geographic biases and inconsistencies, suggesting the need for standardized blocklist maintenance and improved usability. Again, our study confirms a notable similarity between them.

Ramanathan et al. [20] focus on the dynamism of 157 public blocklists over 11 months, highlighting variations in list size, fragmented information, and recurring malicious IP addresses. They advocate for improved data aggregation and expanded entries to enhance blocklist effectiveness and address co-located threats. In our work, we observe that some lists are often copies and/or use the same method to insert IP addresses, as the discussion of Figure 3 highlights.

Finally, other works explored different angles of the public blocklist ecosystem. Sinha et al. [31] analyzed reputation-based spam blocklists, e.g., NJABL, SORBS, SpamHaus, over 10 days within an academic network of 7000 hosts, finding high false-negative rates across most lists. These lists failed to detect low-volume or short-lived spam sources, suggesting that blocklists alone may not effectively cover the spam landscape. Kuhrer and Holz [32] introduced a system to monitor 49 blocklists tracking servers hosting exploits, malware, and botnets. Over 80 days, they gathered $410\,000$ unique URLs and 2.2 million entries, integrating with DNS and HTTP data for deeper insights into malicious infrastructures like Command and Control servers.

VIII. Conclusions

We evaluated the extent to which public blocklists can explain internet radiation traffic, leveraging 4 geographically distributed telescopes and dozens of public blocklists. Our results showed that, when aggregated, the blocklists cover only a portion of the telescope traffic. Blocklists do include the most active IP addresses observed in the telescopes. However, $28.76\,\%$ of the addresses observed in the telescopes and performing potentially malicious scans are never included in any blocklists, nor are they recognized as legitimate scanners.

To assess the effectiveness of blocklists in covering telescope traffic, we rely on two custom metrics: precision (efficiency) and recall (efficacy). Our findings indicate a clear trend

related to blocklist size: larger blocklists tend to have lower precision, while smaller blocklists typically show reduced recall. Furthermore, we observed that telescopes usually see IP addresses days before they are reported by blocklists. This is true even for the top-talkers, i.e., aggressive sources performing high-volume scans that are likely of major interest to network administrators who need to block attacks in real time.

These findings show that distributed and localized telescopes are still an important asset for network administrators and cybersecurity practitioners. They usually provide earlier information about events and cover more sources targeting a specific network. Thus, they allow for more effective blocking of the scanning activities typically seen in early cyberattack stages. This calls for further cooperation among operators of telescopes to distribute contextualized and real-time information about scans.

As future work, we plan to extend the analysis by considering other types of telescopes, in particular those operating active responders such as honeypots. These responders attract different types of attacks, and we plan to evaluate whether our findings remain valid in those setups. Finally, we plan to engage with some blocklist operators to understand how their lists are formed and eventually set up channels to feed telescope data into the public lists in a timely fashion.

Acknowledgments

The research leading to these results has been funded by the projects ACRE (AI-Based Causality and Reasoning for Deceptive Assets—2022EP2L7H) and COMPACT (Compressed features and representations for network traffic analysis in centralized and edge Internet architectures—2022M2Z728), funded by the European Union - Next Generation EU within the PRIN 2022 program (D.D. 104 - 02/02/2022 Ministero dell'Università e della Ricerca). This manuscript reflects only the authors' views and opinions and the Ministry cannot be considered responsible for them.

References

[1] L. Gioacchini, L. Vassio, M. Mellia, I. Drago, Z. Houidi, and D. Rossi, "i-DarkVec: Incremental Embeddings for Darknet Traffic Analysis," *ACM Transactions on Internet Technology*, vol. 23, no. 3, pp. 1–28, 2023.

[2] M. Kallitsis, R. Prajapati, V. Honavar, D. Wu, and J. Yen, "Detecting and Interpreting Changes in Scanning Behavior in Large Network Telescopes," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 3611–3625, 2022.

[3] W. Harrop and G. Armitage, "Defining and evaluating greynets (sparse darknets)," in *The IEEE Conference on Local Computer Networks 30th Anniversary (LCN'05) l.* IEEE, 2005, pp. 344–350.

[4] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson, "Characteristics of internet background radiation," in *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, 2004, pp. 27–40.

[5] M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto, and A. Dainotti, "Millions of targets under attack: a macroscopic characterization of the dos ecosystem," in *Proceedings of the 2017 Internet Measurement Conference*, 2017, pp. 100–113.

[6] A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé, "Analysis of country-wide internet outages caused by censorship," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, 2011, pp. 1–18.

[7] A. Dainotti, A. King, and K. 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*, ser. BADGERS '12. New York, NY, USA: Association for Computing Machinery, 2012, pp. 13–14.

[8] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis *et al.*, "Understanding the mirai botnet," in *26th USENIX security symposium (USENIX Security 17)*, 2017, pp. 1093–1110.

[9] D. Moore, C. Shannon, G. M. Voelker, and S. Savage, "Network telescopes: Technical report," 2004.

[10] P. Richter and A. Berger, "Scanning the scanners: Sensing the internet from a massively distributed network telescope," in *Proceedings of the Internet Measurement Conference*, 2019, pp. 144–157.

[11] E. Pauley, P. Barford, and P. McDaniel, "DScope: A Cloud-Native internet telescope," in *32nd USENIX Security Symposium (USENIX Security 23)*. Anaheim, CA: USENIX Association, Aug. 2023, pp. 5989–6006.

[12] R. Hiesgen, M. Nawrocki, A. King, A. Dainotti, T. C. Schmidt, and M. Wählisch, "Spoki: Unveiling a new wave of scanners through a reactive network telescope," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 431–448.

[13] L. Izhikevich, M. Tran, M. Kallitsis, A. Fass, and Z. Durumeric, "Cloud watching: Understanding attacks against cloud-hosted services," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, 2023, pp. 313–327.

[14] F. Soro, T. Favale, D. Giordano, I. Drago, T. Rescio, M. Mellia, Z. B. Houidi, and D. Rossi, "Enlightening the darknets: Augmenting darknet visibility with active probes," *IEEE Transactions on Network and Service Management*, vol. 20, no. 4, pp. 5012–5025, 2023.

[15] C. Liu, S. Hao, Q. Liu, C. Bao, and X. Li, "Ipv6-network telescope network traffic overview," in *2021 IEEE 11th International Conference on Electronics Information and Emergency Communication (ICEIEC)*. IEEE, 2021, pp. 1–4.

[16] F. Soro, I. Drago, M. Trevisan, M. Mellia, J. Ceron, and J. J. Santanna, "Are darknets all the same? on darknet visibility for security monitoring," in *2019 IEEE international symposium on local and metropolitan area networks (LANMAN)*. IEEE, 2019, pp. 1–6.

[17] L. Gioacchini, L. Vassio, M. Mellia, I. Drago, Z. B. Houidi, and D. Rossi, "Darkvec: Automatic analysis of darknet traffic with word embeddings," in *Proceedings of the 17th International Conference on emerging Networking EXperiments and Technologies*, 2021, pp. 76–89.

[18] Á. Feal, P. Vallina, J. Gamba, S. Pastrana, A. Nappa, O. Hohlfeld, N. Vallina-Rodriguez, and J. Tapiador, "Blocklist babel: On the transparency and dynamics of open source blocklisting," *IEEE Transactions on Network and Service Management*, vol. 18, no. 2, pp. 1334–1349, 2021.

[19] D. Ravalico, R. Valentim, M. Trevisan, and I. Drago, "Can blocklists explain darknet traffic?" in *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 2024, pp. 1–4.
[20] S. Ramanathan, J. Mirkovic, and M. Yu, "Blag: Improving the accuracy of blacklists," in *NDSS*, 2020.
[21] Z. Durumeric, E. Wustrow, and J. A. Halderman, "ZMap: Fast internet-wide scanning and its security applications," in *22nd USENIX Security Symposium (USENIX Security 13)*, 2013, pp. 605–620.
[22] M. P. Collins, A. Hussain, and S. Schwab, "Identifying and differentiating acknowledged scanners in network traffic," in *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2023, pp. 567–574.
[23] Z. Durumeric, M. Bailey, and J. A. Halderman, "An Internet-Wide view of Internet-Wide scanning," in *23rd USENIX Security Symposium (USENIX Security 14)*, 2014, pp. 65–78.
[24] A. Anand, M. Kallitsis, J. Sippe, and A. Dainotti, "Aggressive internet-wide scanners: Network impact and longitudinal characterization." [Online]. Available: http://arxiv.org/abs/2305.07193
[25] H. Griffioen, G. Koursiounis, G. Smaragdakis, and C. Doerr, "Have you syn me? characterizing ten years of internet scanning," in *Proceedings of the 2024 ACM on Internet Measurement Conference*, 2024, pp. 149–164.
[26] C. Fachkha, E. Bou-Harb, and M. Debbabi, "Inferring distributed reflection denial of service attacks from darknet," *Computer Communications*, vol. 62, pp. 59–71, 2015.
[27] A. Dainotti, A. King, K. Claffy, F. Papale, and A. Pescapé, "Analysis of a"/0" stealth scan from a botnet," in *Proceedings of the 2012 Internet Measurement Conference*, 2012, pp. 1–14.
[28] F. Soro, M. Allegretta, M. Mellia, I. Drago, and L. M. Bertholdo, "Sensing the noise: Uncovering communities in darknet traffic," in *2020 Mediterranean Communication and Computer Networking Conference (MedComNet)*. IEEE, 2020, pp. 1–8.
[29] E. Bou-Harb, M. Debbabi, and C. Assi, "Behavioral analytics for inferring large-scale orchestrated probing events," in *2014 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2014, pp. 506–511.
[30] M. Umizaki, T. Morikawa, A. Fujita, T. Takahashi, T. Lin, and D. Inoue, "Understanding the characteristics of public blocklist providers," *2022 IEEE Symposium on Computers and Communications (ISCC)*, vol. null, pp. 01–07, 2022.
[31] S. Sinha, M. Bailey, and F. Jahanian, "Shades of grey: On the effectiveness of reputation-based "blacklists"," in *2008 3rd International Conference on Malicious and Unwanted Software (MALWARE)*. IEEE, 2008, pp. 57–64.
[32] M. Kuhrer and T. Holz, "An empirical analysis of malware blacklists," *PIK-Praxis der Informationsverarbeitung und Kommunikation*, vol. 35, no. 1, pp. 11–16, 2012.

APPENDIX

TABLE IV: List of all 50 blocklists in alphabetical order.

AlienVault Reputation List	Firehol Level 2
Binary Defense	Firehol Level 3
Blocklist.de	Full Bogons IPv4
CINS Army Bad Guys	gnX Threat Intelligence
DangerRulezSK Brute Force Blocker	Greensnow Blacklisted IPs List
Dataplane DNS CH TXT version.bind	hpHosts EMD (IPs)
Dataplane DNS recursion desidered	hpHosts EXP (IPs)
Dataplane DNS recursion desidered IN ANY	hpHosts FSA (IPs)
Dataplane DNS TCP	Inversion DNSBL
Dataplane IP protocol 41	IPsum Level 4
Dataplane SIP invitation	IPsum Level 5
Dataplane SIP query	IPsum Level 6
Dataplane SIP registration	ISX Solutions Blocklist
Dataplane SMTP data	Maltrail - Parking sites
Dataplane SMTP greeting	Malware World suspicious IPs
Dataplane SSH client connection	Mirai Tracker
Dataplane SSH password authentication	MyIP Blacklist
Dataplane TELNET login	Nix Spam DNSBL
Dataplane VNC RFB	Nordic Filters
EmergingThreats Block IPs	pfBlockerNG - MS-1
EmergingThreats Compromised IPs	pfBlockerNG - MS-3
Feodo Tracker Botnet C2 IOCs	SecLists (Careto IPs by Kaspersky)
Feodo Tracker Botnet C2 IOCs Agressive	Spamhaus DROP v4
Feodo Tracker IP blocklist	Turris greylist
Firehol Level 1	Urlhaus-filter