International Journal of Network Management

WILEY

SPECIAL ISSUE PAPER

Beyond Size: Investigating the Impact of Scaled-Down Network Telescopes on Threat Detection

Arthur Vinícius Cunha Camargo | Lisandro Granville | Leandro M. Bertholdo

Instituto de Informática, Universidade Federal do Rio Grande do Sul, Porto Alegre, Brazil

Correspondence: Arthur Camargo (avccamargo@inf.ufrgs.br)

Received: 2 September 2024 | **Revised:** 30 January 2025 | **Accepted:** 5 March 2025

Keywords: cybersecurity | darknet | network telescopes | security service | service management

ABSTRACT

Cyber threat intelligence relies on network telescopes to detect attacks and emerging threats, traditionally utilizing a substantial portion of the IPv4 address space. However, the escalating scarcity and value of this resource force companies and research centers to grapple with the challenge of repurposing their address spaces, potentially impacting cybersecurity effectiveness and hindering research efforts. In this article, we first investigate the historical usage of IPv4 address space in network telescopes and the current reduction trend in several initiatives. Then, we examine the impact of reducing the allocated space on the ability of these systems to identify attackers and attack campaigns. We explore two network telescopes with the intention of assessing the impact of this reduction by quantifying the losses in several ways. Our findings reveal that even halving the allocated space for a network telescope may still permit the detection of 80% of unique cyberattack sources and the address allocation schema has little to no influence on this detection. We also found that most of the proportions and patterns remain present, albeit with reduced intensity.

1 | Introduction

Network telescopes, also known as "darknets" [1], capture and record unsolicited Internet traffic directed at unused, globally routed IP address space. While network telescopes have been utilized for years, they remain essential tools for detecting and studying cyber threats and global events.

The main application of network telescopes includes monitoring and analyzing Internet traffic, helping cybersecurity experts identify new threats, attack patterns, and understand the behavior of potential attackers. They have been used for a long time to observe and better understand cyberattacks on an Internet scale, such as botnets propagation [2], distributed denial of service (DDoS) [1, 3], and network scan campaigns [4, 5]. The collected data allow us to provide a myriad of insights into the malicious, unwanted, and unexpected behavior of cyberattacks.

Network telescopes typically collect large volumes of data. For example, a medium sensor with 65,536 IPv4 addresses (/16) can generate between 10 GB and 1 TB of data daily. Analyzing this massive amount of data is challenging, but the use of artificial intelligence (AI) techniques is transforming this scenario [6]. Advanced IA methods significantly improve one's ability to understand and interpret the data collected by those sensors, thereby providing deeper insights into emerging cyber threats.

Despite their many benefits, network telescopes face an inherent challenge: IPv4 address space has become a scarce and expensive resource. For example, an address space /19 (8,192 IP addresses) is rated between US$ 357,990 to US$ 395,673 [7]. On the other hand, IPv6 Telescopes still have limited value in detecting attacks, with little traffic, most of which is generated by research IPv6 mapping efforts [8]. In this context, companies, research networks, and universities still rely solely on network

Abbreviations: AI, artificial intelligence; ASes, autonomous systems; DDoS, distributed denial of service; DoS, denial of service; IBR, Internet background radiation.

telescopes using IPv4 as their primary source of information. However, they now face growing pressure to release the IPv4 address space used in these sensors for other purposes, or even to sell or rent these addresses.

Recent studies, such as the exploration of dynamic telescopes in cloud environments [9], have highlighted financial challenges arising from the increasing cost of IPv4 address space. Beginning January 1, 2024, major cloud providers like AWS and Google Cloud will impose new charges for IPv4 usage [10, 11]. These costs pose significant obstacles to ongoing and future telescope research in cloud platforms, particularly by limiting access to affordable IPv4 resources.

This paper builds on previously published work [12], and the sections until 4.3 retain similar wordings for clarity and context. Our previous work presents a dual-fold investigation aimed at understanding the current state of network telescopes. Its first contribution involves a meticulous literature review to discern the presence and magnitude of a potential decline in the size of network telescopes. Its second contribution is an exploration of our Network Telescope to investigate the consequences of reducing its address space. The main objective is to experiment with different approaches for allocating the limited IPv4 address resource and to understand the impact of this reduction on its effectiveness in threat detection and intelligence.

Our analysis involves the sampling of two network telescopes, exploring the effects of reduced IPv4 address space on traffic volume and unique source detection. Building from that work, and after sampling the telescope, we examine its behavior in relation to different protocols and attack scenarios in this work, contributing to a better understanding of Network Telescope dynamics and offering insights for addressing challenges posed by IPv4 scarcity in future network telescope development.

This paper is organized as follows: In Section 2, we provide a literature review addressing key concepts and new approaches for network telescopes. In Section 3, we detail the various methodologies applied throughout this paper, including those used to assess the impact of reducing IPv4 address space. In Section 4, we present the results obtained from each methodology, such as the review of other Network Telescope initiatives, sampling methods to collect unique sources, the behavior of different protocols when the telescope's size is halved, and how much of the TCP, UDP, and ICMP attacks can still be captured after this reduction. We also show our results on investigating attack campaigns and time series events. Finally, in Section 5, we summarize our conclusions and give insights about future works.

2 | Background, Definitions, and Related Work

A Network Telescope, or "darknet," is a technique that allows us to capture traffic destined for unused IP address space on the Internet. Due to the unadvertised nature of these "dark" spaces, all the traffic received by this infrastructure is unsolicited, as no real service is assigned to this address space. This makes the traffic very likely to be malicious, with a very low number of false positives. The term "Network Telescope" is also known by various alternative names, such as darkspace, darknet, or blackhole [13]. Throughout this paper, we consistently use the term "Network Telescope" to maintain terminological consistency.

Those sensors normally operate by designating a block of unused IP addresses to monitor unsolicited traffic. This IP block must be allocated exclusively for monitoring purposes to avoid mixing legitimate user data and malicious traffic. Additionally, a router or switch is required to forward packets destined for this IP space to a data capture system. Finally, the captured packets can then be stored and analyzed for further exploration. There are various methods to set up and deploy a network telescope, but configuring the network and ensuring sufficient storage can be challenging. Bailey et al. [14] describe some of the methods and requirements to deploy sensors.

An example of deployment can be seen in Figure 1, where the router statically routes an entire block to the telescope sensor. That way, although very simple, requires that all the blocks be specifically destined for monitoring. That method was the one utilized to monitor our network telescope in our experiments.

Recent advancements in data analysis, automation techniques, and artificial intelligence have leveraged the usability of large datasets generated by network telescopes. This enhancement aims to boost their effectiveness in cyber threat analysis, resulting in increased interest in network telescopes in recent years. Some research showcasing this evolution can be found in Soro et al. [15] and Cabana et al. [5].

Soro et al. [15] introduced community detection algorithms applied to represent network telescope activity as a graph, grouping hosts infected by a botnet that is actively scanning the network in search of vulnerable services. Cabana et al. [5] utilized a combination of network telescope traffic analysis and artificial intelligence to analyze reconnaissance attack campaigns against industrial control systems, allowing an automatic determination of the threat level associated with each campaign.

Network telescope sensors have various applications, including the analysis of Internet scan campaigns [4], locating botnets [16], observing the proliferation of Internet worms [17], and the analysis of Internet Backscattering Radiation (IBR) [18]. These sensors are valuable for detecting and studying such threats, how they spread across the Internet and how attackers select their targets. However, [19] shows that the size of blocks allocated for this purpose affects its capabilities.

An alternative for expanding network telescope size, or beginning new initiatives, is IPv6. However, current challenges in

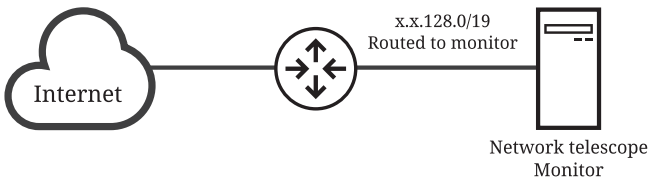

FIGURE 1 | An example of network telescope deployment is routing the spare addresses to the telescope monitor, also known as the sensor.

effectively scanning the IPv6 address space often lead attackers to prefer IPv4. IPv6 scans perceived from the viewpoint of IPv6 network telescopes are infrequent and exhibit significantly different characteristics compared to the more common IPv4 scans [20], most times just capturing address discovery campaigns and mapping efforts [8]. There has been some research made by RIPE [21] using a very large address space, although, as expected, it only received very few scans compared to IPv4. These findings suggest that IPv4 network telescopes will continue to play a significant role in the foreseeable future.

Additionally, despite the existence of large, well-known global network telescopes like CAIDA [22] and Merit [23], many organizations face constraints in allocating addresses for threat intelligence. Deploying and maintaining a network telescope is often hindered by the scarcity of unused IP address space, making it a significant entry-level barrier. Our research has noted a limited examination of the resources invested by companies and researchers in network telescopes. Thus, in Section 4.1, we review existing literature to identify utilized network telescopes in research.

Several telescopes, including CAIDA and Merit, have reduced their size over the years. One of the earliest approaches to employing a smaller address space for a network telescope while preserving the benefits for cyber threat intelligence was proposed by Harrop [24]. The authors suggested sparsely populating the sensors between actively used address spaces, coining their approach as "greynets"—a mixture of unused address space within specific subnets. However, this is not an option when you relinquish a part of the address space, as you lose ownership.

Following the subject of comparing or reducing of address space usage of network telescopes, there are other authors who explore that topic.

Pemberton et al. [25] explores sampling Network Telescopes and focus on the arrival density of backscatter radiation using a /16 network. They suggest four schemes to slice its address space, horizontal, a contiguous /24 block, vertical, 1 address of each /24 and Random-30 and Random-256, that, respectively, chooses random addresses from 30 and 256 /24s, respectively. Their work concludes that deploying a number of random /32 networks across the telescope is the best way to predict backscatter radiation activity. We considered their study when analyzing possible sampling methodologies to adopt in Section 3.2.

Chindipha et al. [26] compares how different subnets behave on the matter of collecting IBR. The article computes the overlap of unique sources IPs across the whole network telescope. It concludes that lower /24 subnets do receive more unique origin hosts that the others, noting that 67% of the sources does only scan one IP of the sensor. In Section 3.3, we analyze two network telescope ability to identify unique sources before and after a reduction in size.

Soro et al. [27] also compare 3 network telescopes with different sizes, /19, /15, and three /24, to assess how the size of each sensor influence the efficiency and detection of different types of events. The objective was to assess how the size of each telescope influences the efficiency and detection of different types of events. The study presented evidence that the sources of traffic significantly vary based on the IP range and the size of the network telescope The authors analyzed one week of data, aggregating it around autonomous systems (ASes). They demonstrated that reducing the Network Telescope by half minimally affects the visibility of network scans but results in different behavior in backscatter analysis when considering ASes. In our work, we used similar metrics and added others to assess the impact of reducing the telescope size on event detection efficiency.

Recent developments addressing the scarcity of IPv4 address space for constructing new network telescopes propose other approaches while upholding their primary goals for cyber intelligence.

The "cloud-native Internet telescope" [9] suggests deploying short-lived telescopes on virtual machines within a cloud provider, leasing IPv4 address space, and releasing it after use. Their results indicate that optimal price performance per IPv4 is achieved in 8 min, and 90% of the steady-state traffic to a given IP address, compared with a regular network telescope, can be observed after only 72 min. In autorefsec:campaign-eval, we evaluate the ability to detect events over time.

Another approach, called "meta-telescope" [28], proposes identifying "unlikely to be used" address space in central points of the Internet, also known as Internet Exchanges, and capturing unsolicited traffic to this address space. In their research, they were able to capture unsolicited traffic for more than 350k /24 blocks in over 7k ASes.

Our contribution aims to investigate the normally used range of addresses utilized in telescopes across different organizations. To the best of our knowledge, we have not found any work related to identifying how much IP resources companies and researchers have spent on network telescopes Additionally, we explored how reducing the number of addresses impacts its threat intelligence. Unlike Pemberton [25] and Chindipha [26], we do not focus solely on collecting backscatter radiation data (TCP-ACKs coming from possible DDoS attack victims). Instead, we are primarily concerned with the number of attacks, requests, and unique sources that can be gathered, which has not been done previously.

3 | Methodologies

In this section, we describe all methodologies used in this paper to accomplish our two goals: measuring the current trends in telescope sizes and assessing the impact of their reduction. In the first part, we detail our process of reviewing the literature to identify all known network telescopes, along with their deployment characteristics, address space usage size, and relevant research results based on each infrastructure. Second, we explain the methodologies used to explore sampling strategies and analyze the impact of reducing the IPv4 address space in an existing network telescope from various perspectives, including the number of unique sources, IDS alerts, protocol attributes, and time series analysis. Lastly, we detail how we evaluated the impact of the reduction. This section delves into the methodologies employed in the three cases.

3.1 | Network Telescopes Over Time

To understand the current landscape of IP address utilization in network telescopes, we conducted a literature review. Specifically, we gathered information on the types of sensors being used for research or production purposes, aiming to present a fresh perspective on their deployment trends over time. Additionally, we intend to show the status and information gathered by the projects and its approaches.

To achieve this objective, we selected data from research papers and significant projects related to network telescopes. Given our primary focus on studying address space usage, certain key characteristics are deemed essential for each subject under review. These include the number of IPv4 addresses utilized, the date of deployment, and the primary objectives pursued by the authors.

To make the data more accurate, we only select survey, reviews, essays, databases, and papers related to network telescopes that are at least from year 2000. Non-relevant works were not selected (e.g., white paper, experimental studies, and reports lacking the information we are intending to collect). To be included, the documents must prove to be informative and descriptive and meet one of the following criteria: (i) published by a respected organization with strong scientific endeavor or (ii) published in influential journals and conferences. In addition to those requirements, it is important that the source expresses the steps used to deploy and maintain the Network Telescope, as well as the results that were achieved during its activity.

To collect high-quality research and studies, we conducted a manual search for keywords related to each topic on Google Scholar. We used keywords such as "darknets," "Network Telescope," or "network blackhole" in our search. After reading the abstract of each article, we excluded those that were not eligible. Then, were selected relevant studies for further reading while discarding off-topic articles and papers. In Section 4.1, we present the results from applying that methodology.

3.2 | Defining a Sampling Strategy

To explore the implications of reducing the address space utilized in a network telescope and to understand whether some parts of the address space are more prone to receiving higher traffic, we conducted a comprehensive analysis using one month of data collected from two different sources at different times: the Japanese NICTER Darknet in 2018 [29] and the Darknet-BR, the telescope utilized by Soro et al. in 2019 [27]. For Darknet-BR, we gathered a more recent, complete, and unanonymized dataset from December 2023 to January 2024, as the previous analysis by Soro et al. was considered sufficiently complete and a useful and complementary analysis for comparison with our new dataset.

From these datasets, we inspected and compared the number of requests each IP address received in both telescopes. Furthermore, we analyze sampling strategies, taking into consideration the number of unique sources that could be detected by the telescopes, by calculating the expected value for each arrangement of addresses. We show the results of this methodology in Section 4.2.

Datasets Details

The Network Incident Analysis Center for Tactical Emergency Response (NICTER) sensors are an integration of large-scale network monitoring for the analysis of cyber threats, such as botnets or DDoS attacks. Its dataset encompasses information from eight sensors distributed worldwide, covering networks ranging from /20 to /17. The dataset archives 1 month of data from October 2018 and includes only TCP SYN packets.

The Darknet-BR is a Brazilian Network Telescope operating on a /19 IPv4 prefix over which we have full control, enabling us to perform more in-depth analysis on the captured packets. We used a dataset from December 14, 2023, to January 14, 2024, in our analysis. Table 1 shows the period during which each dataset was collected, the daily volume of collected data, and the address space size of each sensor.

Choosing a Sampling Strategy

One way that organizations have to mitigate the problems of IPv4 exhaustion and scarcity is to reduce their telescopes. To do that, it faces the dilemma of how to divide the blocks. Said that, using the Darknet-BR, we considered two possible alternatives: (1) a reduction from a /19 (8192 IPs) to a contiguous /20 (4096 IPs), splitting it in half and then evaluating if there are differences between the first and second /20, or (2) reducing it to several /24s (256 IPs) by adopting a sample strategy proposed by [25]. In this paper, we consider both solutions, being the first solution (1) a more simple and straightforward approach, and the second (2) one a more complex but still valid approach.

For the first study, we consider just the horizontal sampling from Pemberton et at. [25], since vertical sampling we consider operationally unfeasible—when delegating a prefix we lose control over the that range. In horizontal sampling, we select half of the /24 blocks for the telescope, while the others will be assigned for users. As for the second study, we considered selecting alternating /24 blocks, in a way that they are equally spaced, trying to cover a larger area of addresses. Additionally, we observed individual blocks at specific locations within the network telescope, such as the beginning, end, and middle, exploring potential correlations with the findings of Harrop [24] and Chindipha [26].

Distribution of Request per IP on Telescope

For the purpose of further analyzing the relation of address space and threat detection abilities, we analyze the number of requests received by each IP address in the network telescopes.

As a way to confirm our findings, we applied our method to another network telescope sensor. We selected datasets from NICTER sensors, which cover several IP spaces. After analyzing all of these sensors and noticing they present similar behavior, we chose just one for further comparison (sensor E), as this sensor has the same size as the Darknet-BR telescope. It is worth mentioning that we only analyzed the number of scan

TABLE 1 | Network telescope datasets from NICT and Darknet-BR.

Sensor	NICTER-A	NICTER-B	NICTER-C	NICTER-D	NICTER-E	NICTER-F	NICTER-G	NICTER-H	Darknet-BR
IPv4 address space size	/17	/18	/20	/20	/19	/18	/21	/21	/19
Volume per day	10 GB	6 GB	2 GB	2 GB	3 GB	6 GB	0.8 GB	0.80 GB	3 GB

events (TCP-SYN) and unique attacker sources in this section. We did not consider UDP or ICMP data to keep the comparison consistent—NICTER-E only provided TCP-SYN data.

Impact of Recognition of Unique Sources

In order to estimate the number of unique sources expected for allocating different sizes of network telescopes, we use a probabilistic approach based on the number of different destinations each source has. Considering N the set of addresses in the network telescope, and n its size, and S a subset of N in a way that $S \subseteq N$, and being s its size. Naming K as the set of source IPs captured by the telescope and the probability of k being observed by the smaller version S considering a uniform distribution is given by P_k in Equation (1). This formula is simply the complement of the probability of not observing any k in all the combinations of subsets S of N.

In addition, we grouped the number of IPs k that target the same number of T_k together as G_{tk}, as they provide roughly the same information for our model. That way, we are able to deduct the expected number of unique sources that the reduced version of the sensor will capture utilizing the formula of expected value—see $E[S]$ in Equation (2).

$$P_k = 1 - \frac{\binom{n-T_k}{s}}{\binom{n}{s}} \tag{1}$$

$$E[S] = \sum_{k \in K} P_k \times G_{tk} \tag{2}$$

As we first focus on gathering unique sources, we use the complementary probability of not detecting any attacks. In this way, the only parameter related to the reduced network telescope is its size, $|S|$, meaning that the formula is independent of the specific addresses selected in the smaller version.

3.3 | Network Telescope Reduction Evaluation

Next, we evaluate our proposed strategy—using our telescope—by analyzing the impact on different protocols, including TCP, UDP, and ICMP. Additionally, to further improve our understanding of how the reduction affects cyber threat detection, we utilized rules from the widely used Intrusion Detection System (IDS) as a ground truth for events. Finally, we investigate how some UDP campaigns (e.g., UPnP and NTPxd) are affected by the reduction in telescope size while considering the time of the attacks.

Impact on Protocol Types

It is important to determine whether the information lost due to the reduction of the telescopes affects the types and characteristics of the traffic. If so, this would indicate that valuable information is being lost as a result of the reduction.

To measure the impact of data loss in a reduced telescope, we collect information about the most commonly received protocols and how they are distributed within the full or half portion of the telescope. In this analysis, we only consider Darknet-BR, as NICTER does not provide information beyond TCP-SYN packets.

We analyze whether the distribution of protocols remains unaltered and whether the accessed ports are still proportional, even when using a telescope of half the size. We categorize the protocols based on their flags and content attributes. TCP traffic is subdivided according to its various flags, UDP traffic is classified based on its content using Wireshark's [30] protocol classifications and ICMP traffic is analyzed by its most commonly used types and codes. Additionally, we present the most common ICMP types, and TCP and UDP ports along with their intended and typical usage, and establish a "weight" based on the average between the number of packets and the volume of traffic received by each protocol usage. The results of this step are demonstrated in Section 4.3.

Impact on Attacks Recognition

To understand the types of attacks being attempted, we utilize the IDS alert system from Snort-2.9 [31] and complement the already built-in rules (3136 rules) with additional community-driven lists (1373 rules), resulting in a total of 4510 rules [32]. The Snort tool runs checks using multiple regular expressions and takes actions according to the results. This approach allows us to identify which threats and possible exploits are being attempted in our network without manually inspecting all the packets. The results from this methodology are presented in Section 4.4.

Finally, we also explore how the traffic behaves in relation to the timeline of events, a known characteristic of attack campaigns. We gather information about the time that it takes to a source to appear again, that way we can understand some scanning and automatic behavior. We also observe traffic spikes on common alerts for UPnP, that can expose local devices to the Internet [33], and NTPxd as they are very common attacks and shows insistent but short campaigns. The results of this step are shown in Section 4.5.

4 | Results

Here, we present the results obtained by applying the previously described methodology. In Section 4.1, we discuss our findings on the reduction of Network Telescope initiatives over time. We present our results on the sampling strategy aimed at reducing the Telescope size in Section 4.2. We also analyze the impact of reducing the Telescope size in terms of protocol distribution in Section 4.3, cyber threat analysis in Section 4.4, and the effectiveness of detecting attack campaigns in Section 4.5.

4.1 | The Decline in Network Telescopes' Address Space: A Temporal Analysis

After conducting our paper review on network telescopes, we summarize the main telescopes initiatives we identified in Table 2. It is important to note that network telescopes in the cybersecurity industry typically do not publish their data, address space size, or even their existence to safeguard the secrecy of their initiatives. Consequently, we could not gather much information about those environments.

Additionally, most Network Telescope initiatives referred in research also does not disclose the address space they use. They justify this approach to avoid *adversary traffic*–when an attacker avoids scanning or using the network telescope address to avoid being identified. We have listed the address spaces that we could verify.

From the paper review we observed that the majority of significant network telescopes emerged between 2000 and 2007, a period characterized by fewer issues related to IPv4 allocation. The onset of IPv4 address exhaustion was first announced by the Regional Internet Registry (RIR) in Asia in 2011, followed by announcements from other RIRs in subsequent years [55–58].

The IPv4 exhaustion resulted in the absence of new relevant network telescope initiatives and even a reduction in existing network telescopes in recent years. For instance, UCSD-NT/CAIDA, which is part of the US Protected Repository for the Defense of Infrastructure against Cyber Threats (PREDICT) program and its successor, the Information Marketplace for Policy and Analysis of Cyber-risk & Trust (IMPACT) [59], saw a reduction in its IPv4 address space from a /8 to a /9 plus /10 in 2019.

In Figure 2, we can better visualize the decline in the utilization of public IPv4 addresses for network telescopes over the years. The figure is built from a more comprehensive list of publications we investigated, encompassing 28 network telescopes (blue crosses). From this graph, is reasonable to infer an initial reduction around 2010, correlated with the depletion of IPv4 address in RIRs. The second reduction, around 2018 may be linked to the escalating prices of IPv4 in the market. Notably, during this period, portions of addresses from large network telescopes shifted to major companies such as Google and Amazon.

The reduction trend become more evident when we visualize the smoothed tendency over a two-year time window (red line). A linear regression (dashed black line) for this trend is also influenced by the deallocation of IPv4 address space for network telescopes after 2018—as observed in the case of UCSD-NT/CAIDA.

The indicated trend corroborates the difficulty of maintaining the address space active by various initiatives. As a result, new research efforts such as *meta-telescope* [28] and *DScope* [9] aim to explore new ways to deploy temporary telescopes.

4.2 | Effects of Address Space Reduction on Capturing Unique Sources in Network Telescopes

Here, we present our results on assessing the impacts of reducing the size of a network telescope in terms of collecting unique sources and categorizing the types or services that were likely targeted.

We consider TCP and UDP from the transport layer and ICMP from the network layer, using data collected solely from the Darknet-BR dataset, as NICTER does not provide that information. The goal is to verify whether the reduction in sensor size affects the types of attacks received.

We then explored the datasets and analyzed possible strategies to overcome problems related to the scarcity of IPv4 address space. Initially, we examined the distribution of requests per destination IP across all the datasets we listed.

In Figure 3a,b, we can visualize the distribution of requests received per IP address in both telescopes in a period of 31 days (one month). Here, we depict Darknet-BR for the period of Dec/2023 and NICTER-E for Oct/2018. It is noteworthy that other NICTER sensors have shown a similar behavior when compared with NICTER-E. In all datasets, the telnet scans (TCP/23) were the primary target, as we further explain in Section 4.4.1 for Darknet-BR, although that is also true for NICTER-E.

In Figure 3a,b, we show the distribution of requests received per IP address across two telescopes over a 31-day period (1 month). Specifically, we present data for Darknet-BR covering December 2023 and for NICTER-E from October 2018. It is important to note that other NICTER sensors have exhibited similar patterns to those seen in NICTER-E. Across all datasets, Telnet scans (TCP/23) were the primary target, as detailed in Section 4.4.1 for Darknet-BR, a trend that also holds true for NICTER-E when we analyzed it.

Applying the methods previsouly described, we calculated the expected value of unique sources for all sizes of S for both sensors (Darknet-BR and NICTER-E). Figures 4 and 5 show the percentage of the unique sources that would be visible when we project a reduction in the number of addresses being utilized by the telescope. The figure illustrate the projection of the percentage from the original telescope that we can reach (Y-axis) by the number of hosts needed (X-axis). This estimation allows us to assess how the reduction will impact the capture of unique sources by providing information on the origins already captured and the desired size of the new telescope.

TABLE 2 | Summary of network telescope projects, referred from the year 2000 to the present, ordered by the largest available address space.

IPv4 addresses	Year	Name	Comments
50,331,648	2010	APNIC/ARIN	APNIC and ARIN collaborated on IBR research utilizing unallocated addresses 1/8, 50/8, and 107/8. This telescope had a lifespan of 1 week in 2006 [34].
17,048,576	2001	Internet motion sensor	Arbor Networks and the University of Michigan project deploys sensors in diverse locations to enhance the diversity, sparsity, and size of a network telescope. The IMS initiative seems ending in 2004 and spanning into Merit telescope [19].
16,777,216	2005	MERIT	Merit network telescope used the 35/8 address from 2005 to 2018. After this date, the Michigan University formalized the Orion telescope with a smaller address space [23].
16,777,216	2001	UCSD-CAIDA	The UCSD network telescope, a project from the University of San Diego/US was built on the globally routed 44/8 prefix (former AMPRNet) from 2001 to 2019 [22].
12,582,912	2019	UCSD-CAIDA	The UCSD Network Telescope reduced it size from a /8 to a /9 and /10 network [22].
~ 2,000,000	2012	SWITCH	Collect data from the address space from multiple networks across Switzerland [35].
626,944	2004	Team Cymru	Multiple sensors deployed by the company Team Cymru [36].
524,288	2014	Farsight	Farsight's network telescope, now part of DomainTools, offers data through subscription [37].
475,136	2018	ORION-MERIT	Michigan State University's project, known as the Observatory for Cyber-Risk Insights and Outages of Networks, focus on Internet backscatter radiation. Designed and engineered with support from the US-NSF, it consists of 1856 /24 subnets [23].
270,000	2005	NICTER	The Japanese organization NICTER (Network Incident Analysis Center for Tactical Emergency Response) integrates its network telescope for large-scale monitoring and analysis of cyber threats, including botnets and DDoS [38].
178,000	2018	MIT-Akamai	The first network telescope built over a content delivery network (CDN) infrastructure. It is composed of two IPs on each of the 89,000 Akamai servers across the globe [39].
131,072	2018	NL-Darknet	Network telescope maintained by SurfNET in the Netherlands [27].
65,636	2019	HEAnet	Ireland's national education and research network telescope [40].
65,536	2004	IUCC/IDC	The Israel InterUniversity Computation Center (IUCC) network telescope [41].
65,536	2006	Anonymous	The University of Wellington, NZ, utilizes an undisclosed /16 Network Telescope to test various address sampling strategies for measuring arrival density [25].
65,536	2021	Anonymous	An undisclosed enterprise network telescope identifies a specific stateless-scanning malware, and a response is forged to slow down the malware's propagation, deceiving botnet scanners. The research is being conducted in Germany [42].
8192	2018	BR-Darknet	A /19 network telescope in Brazil (used in this paper) [27].
4096	2017	JP-Darknet	Another /20 network telescope hosted in Japan [43–45].
4096	2010	INRIA	French Telescope at INRIA's High Security Laboratory [43–46].
765	2017	IT-Darknet	Italian network telescope [27]
512	2009	Rhodes University	The first known network telescope in the Afrinic Region [47].
512	2013	KISTI	Science and Technology Security Center South Korea (KISTI) does provide 2 sensors with a /24 mask [48].
256	2006	—	After 2006, numerous minor initiatives deployed temporary network telescopes with short lifespans (1–2 years), ranging from /28 to /24, for specific research [49–54].

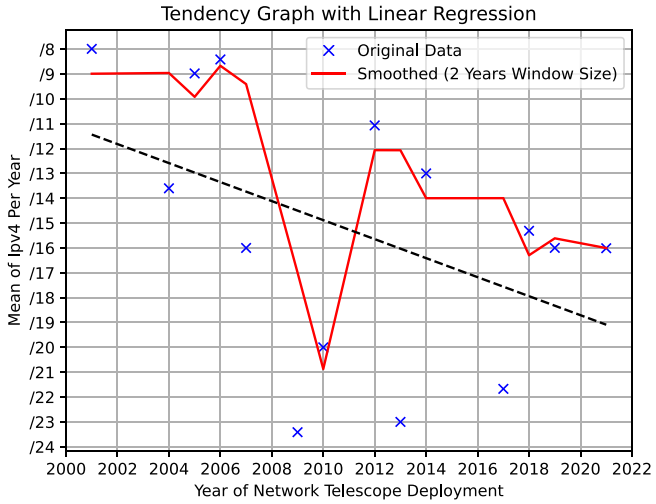

FIGURE 2 | Tendency graph of IPv4 usage on network telescopes in each year, considering the first deployment of each one. Here, we omitted data from initiatives who reduced the size of their telescope (i.e., UCSD and MERIT).

When we compare the projected results for both telescopes, some similarities become evident. Both telescopes expect to capture more than 80% of unique sources when the address space is reduced by half, and more than 60% when reduced to a quarter. Our main observation here is that exists sources that tend to scan more than one address in both telescopes.

Another observation is related to the incline of the curve in the two graphs. In this regard, NICTER-E shows slower growth in the beginning. This result is attributed to different attack methods, with NICTER-E registering most scans at lower addresses in each /24 (scans for routers). The takeaway here is that reducing the address space too much may impact in our ability to detect certain types or methods of scans.

Additionally to the computation of the estimations equation, we also relied on different sampling approaches to make those results more concrete. Here, we arranged the Darknet-BR dataset in 4 different subnets allocation schemata (e.g., one /20 or sets of /24), and gathered the real number of unique sources that would be perceived in each context.

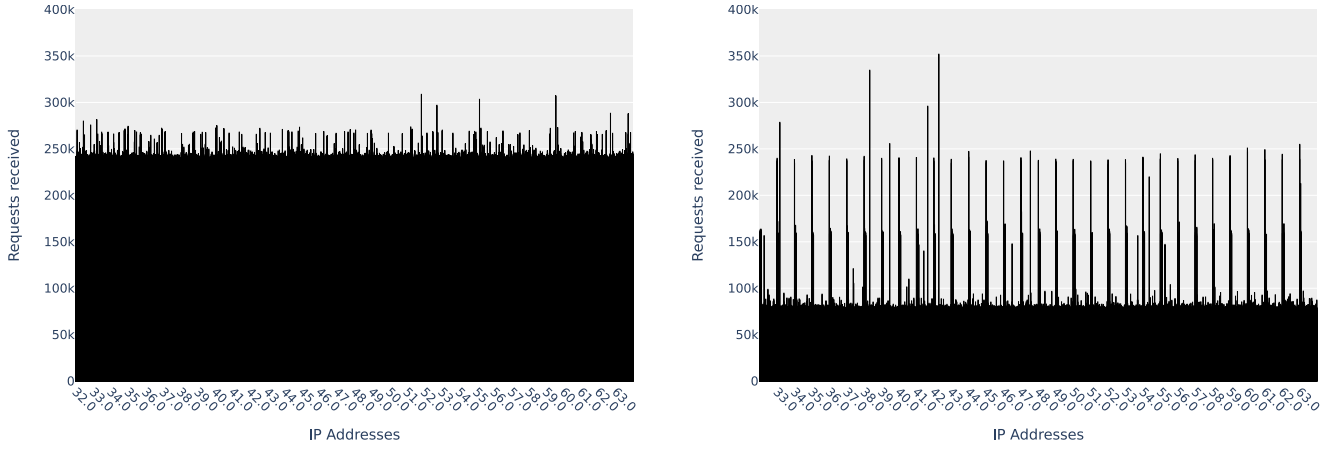

(a) Number of requests per IP, Darknet-BR.

(b) Number of requests per IP, NICTER-E.

FIGURE 3 | Distribution of received scans (TCP-SYN) per IP address in two different network telescopes. Both using a /19 block in 1 month.

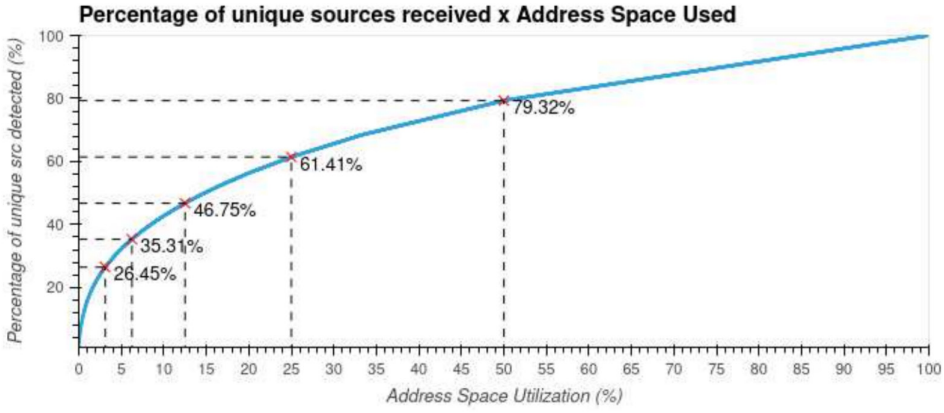

FIGURE 4 | Darknet-BR /19 percentage of unique sources received per IP used.

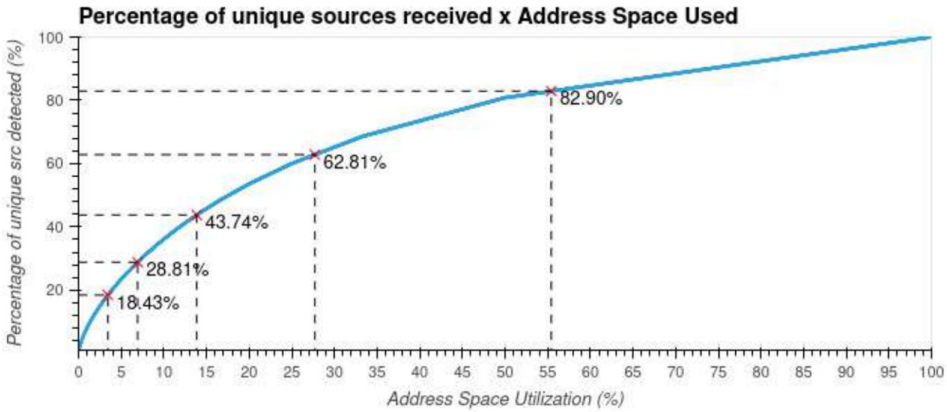

FIGURE 5 | NICTER-E /19 percentage of unique sources received per IP used.

TABLE 3 | Number of unique sources and requests seem by different sampling methods in the Darknet-BR telescope.

Method	Unique sources (%)	Number requests (%)
Total	100.00	100.00
Low /20	80.30	50.03
High /20	80.26	49.97
Even /24 allocation	80.26	50.01
Odd /24 allocation	80.39	49.99

Note: Considering half of the sensor, observe 4096 IP addresses.

TABLE 4 | Number of unique sources and requests seem by different sampling methods in the NICTER-E telescope.

Method	Unique sources (%)	Number requests (%)
Total	100.00	100.00
Low /20	74.73	50.01
High /20	74.46	49.98
Even /24 allocation	74.74	49.97
Odd /24 allocation	74.46	50.02

Note: Considering half of the sensor, observe 4019 IP addresses.

In Tables 3 and 4, we present the results of reducing the size of Darknet-BR and NICTER by half, that is, from 8k addresses to 4k addresses. The figure illustrates the impact on unique sources and the number of requests observed for each applied sampling method. Our results from testing four different allocation—Low /20, High /20, Even /24 allocation, Odd /24 allocation—over both datasets do not show a significant difference, less than 0.1% for identifying unique scan sources or requests.

For Darknet-BR, this phenomenon can be explained by the uniform distribution present in the requests per IP, meaning we did not find any particular set of IPs that are more targeted than the other. On the other hand, the NICTER sensor does still have little differences because the overall distribution inside a /24 is still very much the same between all the blocks, which can also explain the little difference between the allocation methods.

To better understand this behavior, we also examined several individual /24 blocks at the beginning, end, and middle of the /20 block to ensure that no specific block was significantly different from the others. This was true for both datasets and showed very little discrepancy between the number of unique sources and the number of requests, as expected.

After conducting this experiment, we found that it is still possible to identify 80% of the unique IP sources when reducing the number of IPs by half, as a portion of the IP sources appears more than once in each telescope. Moreover, there is no significant discrepancy based on how the blocks are split, as confirmed by experimenting with both the first and second halves and by splitting them into multiple /24 blocks, which resulted in a difference of less than 0.1%. So we decided to explore further the more simple approach of using two /20 prefixes.

4.3 | Effects of Network Telescope Shrinkage on TCP, UDP, and ICMP Traffic Distribution

In the previous section, we identified that our reduced telescope is capable of capturing around 80% of the unique IPv4 addresses, rather than the initially expected 50%—a simplistic reasoning. In this section, we conduct an evaluation of packet counts and traffic volume losses considering the reduction of the Darknet-BR telescope in half (higher /20 prefix and lower /20 prefix). Here, we expand our initial traffic distribution per IP to also analyze the protocol distribution, measuring the most commonly received protocols and examining how they are distributed across the chosen sensor sampling. This analysis aims to determine if one-half of the address space has a preference for receiving traffic over the other as well as exploring how the attacks prefer each protocol.

We analyzed the data reduction by comparing TCP, UDP, and ICMP. We found that the distribution was similar to the full dataset from Darknet-BR. The majority of the scans still attempted to establish a TCP connection, making it the most significant

type of traffic in our sample, accounting for more than 94% of all received packets.

The breakdown of each TCP type of request and flags is shown in Table 5. The presence of a few TCP-ACK packets further underscores the importance of a large telescope for identifying backscatter radiation and reflection victims. Further observation shows that we are receiving packets with the reserved flags on (R1, R2, and R3) which can be explained by some fingerprinting tactics or Xmas tree attacks [60], where the source tries to set all the flags to one in order to provoke bugs and determine implementation subtleties.

Moreover, even with the majority of TCP-SYN packets, it is important to note that other types of packets still contain useful information about different types of attacks. For example, several vulnerable and exploitable services run on top of UDP. Notable examples include NTP, Memcached, and DNS, which are well-known vectors for reflection and amplification attacks, as explored in Table 6.

Most of the UDP data comes from payloads encapsulated solely within the transport protocol. These payloads contain custom data, which can serve as vectors for attacks. In our exploration of the payloads, we found some TSource Engine queries, HTTP headers, stuffed characters, and characters followed by null bytes. Additionally, we observed that much of the traffic lacks a payload and uses the UDP protocol solely for delivery—a tactic commonly employed in port scanning.

Other protocols are well-known for being the source of reflection and amplification DoS attacks, such as NTP, DNS, SNMP, SSDP, CHARGEN, QUAKE3, STEAM_IHS_DISCOVERY [61], MDNS, RIP, RPC, CLDAP, TFTP, and MEMCACHED [62]. Even if they are not as numerous and represent less than 5% of the traffic, collecting information about those attacks can be valuable in the context of network telescopes as they are capable of revealing global events and can also contain victims' addresses as their source IP.

ICMP is divided between its types and codes, which define the purpose of each packet. In Table 7, we can observe that most of the ICMP traffic consists of Echo requests, which are normally utilized as a reconnaissance tool, and can also be utilized as a flood mechanism. Furthermore, we can also observe that a lot of port unreachable appear in second place, which means that some datagrams could not be received. Collecting that combination of type and code is interesting in the context of network telescopes, as it indicates a probable spoofed traffic and also contains 64 bits of the original data, giving a hint of what was the intent of the packet.

Reducing the size of the telescope does not significantly alter the proportions of attack types, which implies that most of the traffic follows a uniform distribution across the sensor. As shown in the tables, the difference between each side of the telescope (high and low) is no more than 3%, indicating minimal imbalance. The overall proportion remains close to 50%, as expected. Most of the observed imbalances stem from categories containing very few packets.

4.4 | Assessment of Targeted Services and Attack Distribution Across Protocols

4.4.1 | Detailed Exploration of TCP Attack Vectors and Targeted Services

As there is not much information that can be gathered inside only the TCP-SYN packet, we also do explore the destination port that are targeted. There is not way to be sure if the requester

TABLE 5 | TCP flags distribution.

Flags	TCP packet share (%)	Packet (%) (half-low)	Packet (%) (half-high)	Volume (%) (half-low)	Volume (%) (half-high)	Weighted (%) (half-low)	Weighted (%) (half-high)
SYN	98.17%	50.01%	49.99%	50.01%	49.99%	49.10%	49.07%
ACK	1.00%	50.52%	49.48%	50.34%	49.66%	0.51%	0.50%
CON	0.28%	49.87%	50.13%	49.87%	50.13%	0.14%	0.14%
ECHO	0.27%	49.90%	50.10%	49.90%	50.10%	0.13%	0.14%
RES	0.21%	49.88%	50.12%	49.88%	50.12%	0.11%	0.11%
PUSH	0.03%	49.91%	50.09%	49.77%	50.23%	0.01%	0.01%
R2	0.01%	49.27%	50.73%	49.19%	50.81%	0.01%	0.01%
FIN	0.01%	49.61%	50.39%	49.60%	50.40%	0.00%	0.00%
URG	0.01%	50.10%	49.90%	50.10%	49.90%	0.00%	0.00%
NONCE	0.00%	45.69%	54.31%	36.62%	63.38%	0.00%	0.00%
R3	0.00%	46.70%	53.30%	32.53%	67.47%	0.00%	0.00%
R1	0.00%	43.21%	56.79%	31.84%	68.16%	0.00%	0.00%
Total						**50.02%**	49.98%

Note: TCP accounts for 94.09% of packets and 87.18% of volume. Green cells show the highest number of requests or traffic volume.

TABLE 6 | UDP protocols distribution.

UDP protocols	UDP protocols packet share (%)	Packet (%) (half-low)	Packet (%) (half-high)	Volume (%) (half-low)	Volume (%) (half-high)	Weighted (%) (half-low)	Weighted (%) (half-high)
DATA	48.78%	49.44%	50.56%	45.58%	54.42%	24.11%	24.66%
UDP ONLY	6.98%	49.99%	50.01%	49.99%	50.01%	3.49%	3.49%
NTP	4.69%	50.09%	49.91%	50.03%	49.97%	2.35%	2.34%
DNS	4.23%	50.02%	49.98%	50.01%	49.99%	2.12%	2.12%
MDNS	3.60%	49.99%	50.01%	50.13%	49.87%	1.80%	1.80%
SIP	3.46%	49.94%	50.06%	49.94%	50.06%	1.73%	1.73%
MEMCACHE	2.65%	50.09%	49.91%	50.06%	49.94%	1.33%	1.32%
SNMP	2.30%	50.07%	49.93%	50.08%	49.92%	1.15%	1.15%
ISAKMP	2.22%	50.22%	49.78%	50.29%	49.71%	1.11%	1.10%
SSDP	1.77%	50.03%	49.97%	50.03%	49.97%	0.89%	0.88%
RPC	1.53%	49.93%	50.07%	49.92%	50.08%	0.76%	0.77%
COAP	1.52%	50.01%	49.99%	50.01%	49.99%	0.76%	0.76%
NBNS	1.11%	49.94%	50.06%	49.94%	50.06%	0.55%	0.55%
OPENVPN	1.09%	49.99%	50.01%	49.99%	50.01%	0.55%	0.55%
CLDAP	1.06%	50.06%	49.94%	50.07%	49.93%	0.53%	0.53%
Other protocols	1.00%	48.17%	51.83%	48.55%	51.45%	0.48%	0.52%
DTLS	0.84%	49.96%	50.04%	49.98%	50.02%	0.42%	0.42%
SRVLOC	0.81%	50.64%	49.36%	50.72%	49.28%	0.41%	0.40%
PORTCONTROL	0.76%	50.28%	49.72%	50.28%	49.72%	0.38%	0.38%
RIP	0.70%	50.16%	49.84%	50.17%	49.83%	0.35%	0.35%
CHARGEN	0.70%	50.37%	49.63%	50.36%	49.64%	0.35%	0.35%
BVLC:BACNET:BACAPP	0.69%	50.01%	49.99%	50.01%	49.99%	0.34%	0.34%
RDPUDP	0.58%	49.96%	50.04%	49.87%	50.13%	0.29%	0.29%
STUN	0.56%	49.97%	50.03%	49.98%	50.02%	0.28%	0.28%
UDPENCAP:ESP	0.47%	50.02%	49.98%	49.93%	50.07%	0.24%	0.24%
XDMCP	0.47%	49.91%	50.09%	49.91%	50.09%	0.23%	0.23%
GQUIC	0.44%	49.97%	50.03%	49.97%	50.03%	0.22%	0.22%
RMCP:IPMI_ SESSION:DATA	0.44%	49.95%	50.05%	49.95%	50.05%	0.22%	0.22%
TFTP	0.43%	50.02%	49.98%	50.02%	49.98%	0.22%	0.22%
L2TP	0.38%	49.96%	50.04%	49.96%	50.04%	0.19%	0.19%
RADIUS	0.36%	50.03%	49.97%	50.04%	49.96%	0.18%	0.18%
ECHO	0.35%	49.75%	50.25%	49.76%	50.24%	0.17%	0.18%
QUAKE3	0.33%	49.99%	50.01%	50.00%	50.00%	0.16%	0.16%
OMRON	0.27%	50.43%	49.57%	50.43%	49.57%	0.14%	0.13%
GTP	0.25%	50.63%	49.37%	50.63%	49.37%	0.13%	0.13%

(Continues)

TABLE 6 | (Continued)

UDP protocols	UDP protocols packet share (%)	Packet (%) (half-low)	Packet (%) (half-high)	Volume (%) (half-low)	Volume (%) (half-high)	Weighted (%) (half-low)	Weighted (%) (half-high)
HART_IP	0.22%	50.02%	49.98%	50.02%	49.98%	0.11%	0.11%
ECATF:ECAT	0.21%	49.98%	50.02%	49.98%	50.02%	0.10%	0.10%
KERBEROS	0.20%	49.86%	50.14%	49.85%	50.15%	0.10%	0.10%
ICMP:IP:UDP:DATA	0.19%	49.82%	50.18%	49.74%	50.26%	0.09%	0.10%
GTP:GTPV2	0.19%	49.95%	50.05%	49.95%	50.05%	0.09%	0.09%
GRE:DATA	0.19%	49.86%	50.14%	49.86%	50.14%	0.09%	0.09%
BFD	0.17%	50.53%	49.47%	50.67%	49.33%	0.08%	0.08%
DAYTIME	0.15%	49.84%	50.16%	49.84%	50.16%	0.07%	0.07%
GRE:IP:UDP:DATA	0.13%	49.91%	50.09%	49.91%	50.09%	0.07%	0.07%
GRE:UDP ONLY	0.10%	49.94%	50.06%	49.92%	50.08%	0.05%	0.05%
KIP	0.10%	49.20%	50.80%	49.20%	50.80%	0.05%	0.05%
STEAM_IHS_ DISCOVERY	0.09%	49.70%	50.30%	49.69%	50.31%	0.05%	0.05%
RPC:DATA	0.09%	49.93%	50.07%	49.93%	50.07%	0.04%	0.04%
ICMP:IP:UDP	0.07%	49.93%	50.07%	49.44%	50.56%	0.04%	0.04%
TIME	0.07%	49.17%	50.83%	49.26%	50.74%	0.03%	0.04%
Total						48.81%	**51.19%**

Note: UDP accounts for 5.07% of packets and 10.73% of the volume. Green cells show the highest number of requests or traffic volume.

is really trying to exploit that specific service based on the destination port that is being targeted, but it provides a good guess of intention.

The graph (Figure 6) shows the top 50 ports by the percentage of requests in a bar graph. It contains two bars. The larger bar represents the baseline scenario, where 100% of the telescope is being utilized. The smaller, overlayed, and highlighted bar represents the reduced sensor scenario, where only 50% of the IPs are utilized.

Halving the sensor still does not change the proportion of those port analysis, which means that they also tend to be uniformly distributed between the addresses. It is possible to still gather this proportion using a smaller telescope and use that information to understand which techniques are more prevalent than the others. That enable administrators to understand the most targeted services in their networks and create better custom defenses.

It is clear that most of the attacks targets remote access services like telnet, ssh, and RDS, more critical infrastructure like SQL databases and Microsoft Directory Service and some popular services like http and https. Those services normally does grantee a great impact if successfully accessed, as they are also services that contains known vulnerabilities that can be remotely exploited. Most of the data indicates that Telnet is also related to IoT botnet campaigns that brute force password of weak devices like IP Cameras, DVRs, and consumer routers [63]. No conventional ports like 81, 8888, and 2323 are also normally related to http and telnet, appearing as popular targets [64] and also represent types of remote access attempts.

4.4.2 | Detailed Analysis of UDP Attack Patterns and Service Risks

The UDP protocol, as being a connection-less protocol, does tend to contain some wildly exploited application attacks with NTP, SSDP, CLAP, DNS, and others [65]. In that sense, we also explored the ports being destined and the listed the probable service being targeted.

As it is presented in Figure 7, there are also the patterns of receiving only half of the requests in the top 50 ports, which means that the traffic port are distributed between the IP addresses uniformly and that reducing the size of the telescope will wield the same kind of attacks, although, with less quantity. Differently from TCP, we observe that there are a smaller number of ports that are responsible for most of the traffic. While 74.66% of the packets target the top 50 ports in UDP, in TCP that number is 27.32%. That can be explained by the fact that a lot of the TCP traffic consists of scans without significant data with the objective to identify open ports and services, while in UDP permits a more payload based attacks where the first request is already trying to exploit possible vulnerabilities and misconfigurations.

TABLE 7 | ICMP telescope distribution.

Description	Type	Code	ICMP share (%)	Packet (%) (half-low)	Packet (%) (half-high)	Volume (%) (half-low)	Volume (%) (half-high)	Weighted (%) (half-low)	Weighted (%) (half-high)
Echo request	8	0	96.55%	51.27%	48.73%	51.81%	48.19%	49.50%	47.05%
Port unreachable	3	3	1.35%	49.87%	50.13%	49.71%	50.29%	0.67%	0.68%
Redirect	5	1	0.82%	49.85%	50.15%	49.46%	50.54%	0.41%	0.41%
Echo reply	0	0	0.53%	51.98%	48.02%	52.54%	47.46%	0.28%	0.26%
TTL exceeded	11	0	0.30%	49.94%	50.06%	50.13%	49.87%	0.15%	0.15%
Timestamp	13	0	0.20%	49.79%	50.21%	49.79%	50.21%	0.10%	0.10%
Fragmentation required	3	4	0.17%	49.87%	50.13%	49.77%	50.23%	0.09%	0.09%
Host prohibited	3	10	0.03%	48.91%	51.09%	48.28%	51.72%	0.01%	0.01%
Host unreachable	3	1	0.02%	50.22%	49.78%	51.55%	48.45%	0.01%	0.01%
Protocol unreachable	3	2	0.01%	31.77%	68.23%	39.63%	60.37%	0.00%	0.01%
Communication prohibited	3	13	0.01%	49.63%	50.37%	50.98%	49.02%	0.00%	0.00%
Network unreachable	3	0	0.01%	47.23%	52.77%	47.62%	52.38%	0.00%	0.00%
Unknown	233	734	0.00%	27.12%	72.88%	64.13%	35.87%	0.00%	0.00%
Fragment time exceeded	11	1	0.00%	50.00%	50.00%	50.00%	50.00%	0.00%	0.00%
Bad IP header	12	0	0.00%	66.67%	33.33%	66.67%	33.33%	0.00%	0.00%
Total								**51.48%**	48.52%

Note: ICMP accounts for 0.7% of packets and 0.8% of volume. Green cells show the highest number of requests or traffic volume.

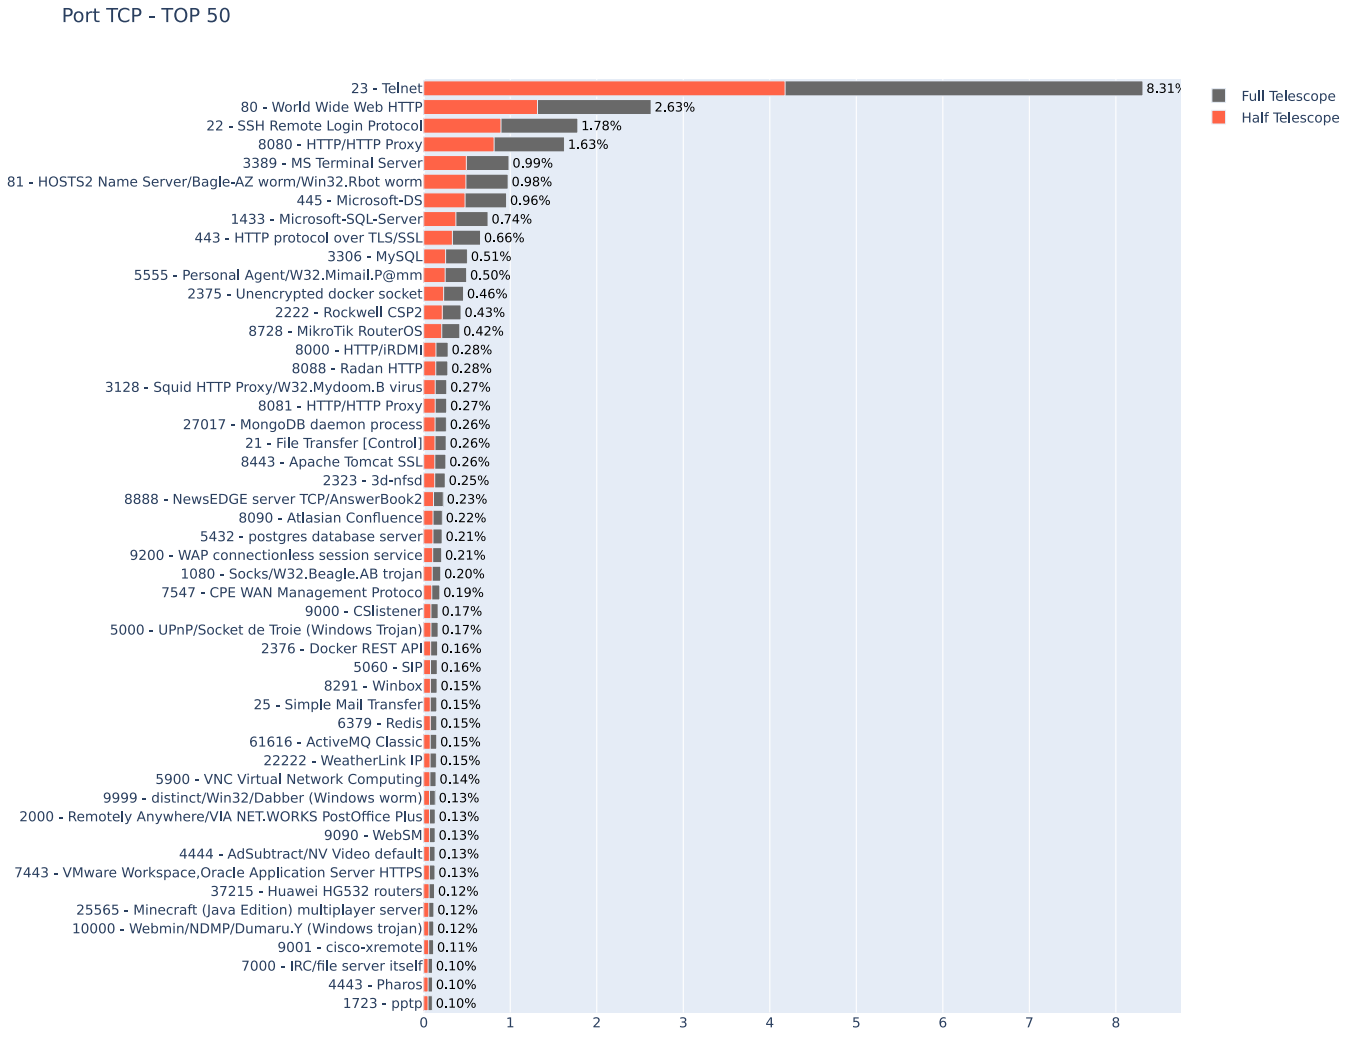

FIGURE 6 | TCP ports percentage.

It is worth noting that some unassigned ports appears that do not have an attributed purpose and that does not follow an uniform distribution along the destination IP; in Figure 7, we can see that port 14580 does not appear on the reduced version of the telescope while all of the packets that target port 46235 appear in the shrunk version. Although there exists this discrepancy, most of those ports are not reserved or default any type of service, which implies that they are explored randomly for scanning purposes.

4.4.3 | Detailed Analysis of ICMP Attack Vectors and Distribution

In this section, we explore the ICMP data received by the network telescope. The literature provides multiple ICMP related attacks; the most common are the Ping Flood attack, ICMP tunneling, and Redirect attacks [66, 67]. Those attacks differ in means and objectives; a ping flood attack tries to compromise thethe system's availability by sending an excessive amount of ping packets (Echo Requests) that can saturate the network and the computation capabilities of servers. On the other hand, an ICMP tunneling circumvents defenses by encapsulating all the traffic using ICMP packets. Redirect attacks try to manipulate the routing table of hosts with the means to eavesdrop on the connection using man-in-the-middle (MITM) approaches.

In order to identify that attacks into our network telescopes, we tried to find common patterns that normally appears on those requests and packets that are related to each type of attack.

First of all, the type and code of the packets are analyzed to understand its distribution as can be observed in Table 7. Most of the packets consists of echo requests that can be explained as pings that are frequently utilized to know if a given host is up on the Internet. We do not think that all of those need to be considered being malicious as a lot of research institutions do scan the internet for data collection. A lot of those organizations do identify their packets in the payload, but we were able to identify only the Ripe Atlas.

As we observed some of the content of the packets, some known attacks become more apparent and obvious. We have found some data containing payloads of destination ports unreachable with seemingly malicious payloads; as ICMP includes data about the packets in those cases, it proves to be an good intelligence gathering mechanism for the telescope. That information

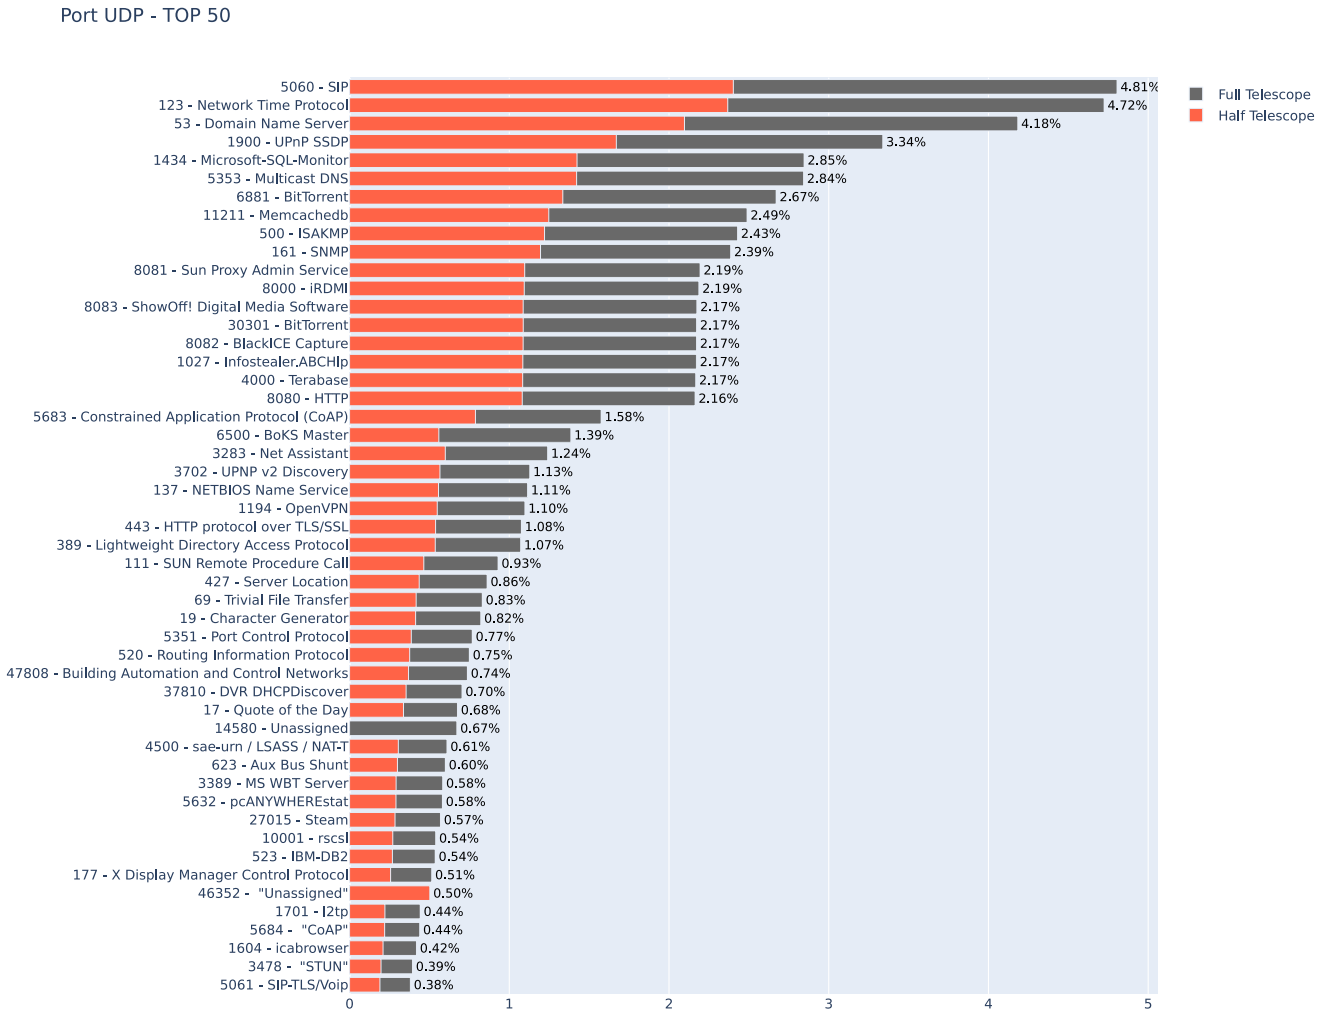

FIGURE 7 | UDP ports percentage.

can give some clue about victims of attacks and tell more about the whole Internet landscape. We have found strings related to *TSource Engine Queries* that are related to distributed reflective denial of service (DRDoS) that utilizes Valve software protocol [61, 62]. Researchers were able to accomplish amplification factors of 33 times in 2013 using that type of attack on gameservers [68]. Additionally, we found a lot of Bittorrent traffic within those packets, which can be related to spoofed traffic.

We also leveraged a list of the protocols that are encapsulated within ICMP, which can give provide some factors (Table 8). That table shows that protocol using wireshark layers notation, where it becomes explicit how the packeets are encapsulated; for example, icmp:ip:udp:data means that there is an icmp packet that encapsulates the ip protocol that encapsulates udp and that last one contains bytes of data.

We observe that the top 5 protocols found by Wireshark does correspond to more than 99% of the packets and that the majority of those contains some kind of data payload. The content of each payload varies between having a stuffed payload like *1234567* or *abcdefh*, which show that not all the data tries to be disruptive in the first glace, although a great part of the payloads that also contains UDP holds some more sophisticated data like Bittorrent headers and The TSource Engine Queries.

TABLE 8 | ICMP protocols.

Protocol	Count	Percentage (%)
icmp:data	16,299,429	90.049139
icmp only	1,294,107	7.149528
icmp:ip:udp:data	219,374	1.211971
icmp:ip:udp	85,717	0.473559
icmp:ip:icmp	85,145	0.470399

4.5 | Time-Based Detection and Analysis of Network Attack Campaigns

In this section, we observe how the time influence the prevalence of the requests and how long does it need to the same source to appear again. There are multiple factors that can influence the number of requests received by time, the discovery of new vulnerabilities, big international campaigns, and even recurring scans that are programmed in a fixed hour. There are also some known organizations that make periodic scans across the internet in order to identify open ports and services held by IPv4 addresses (Table 9).

TABLE 9 | Mean seconds until source reappearance.

	Reappearance time (seconds)
Mean	42,123.03
Standard deviation	144,330.77
Min value	0.000003
25%	8.74
50%	970.04
75%	12,883.12
Max value	2,680,215

Taking in consideration a smaller time, and a prevalent type of attack, we can observe shorted-lived campaigns that have a short duration. Figure 8 shows attacks from a well-known NTPxd overflow exploit attempt and the prevalent UPnP protocol being targeted frequently by bursts of traffic. The orange lines represent the data that would be collected have we reduced the telescope, which shows that the majority of those attacks can still be perceived by a smaller telescope.

For further approval that we can still gather qualitative information from reducing our sensors, we also found that pattern appears for all the alerts across the experiment's lifespan. Figure 9 shows the sum of the count of all the alerts that were collected

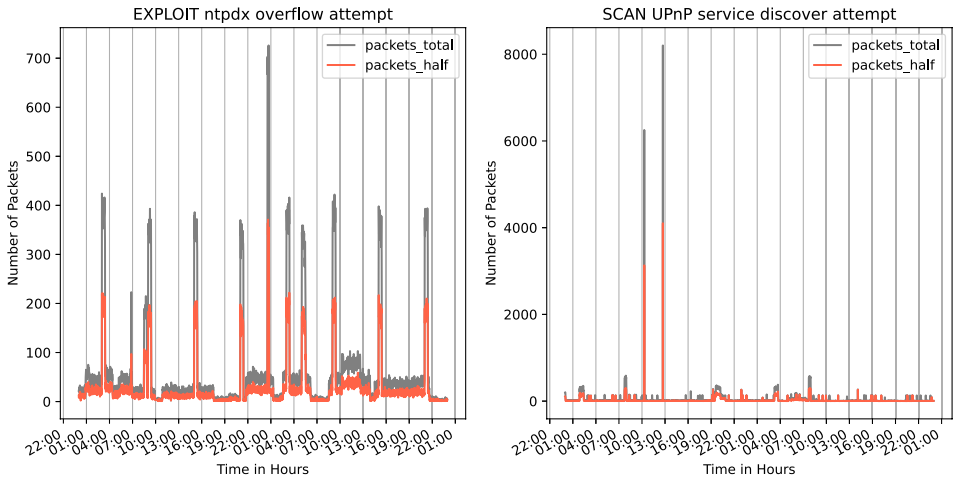

FIGURE 8 | Diary requests received trying to Scan UPnP and to exploit NTPxd.

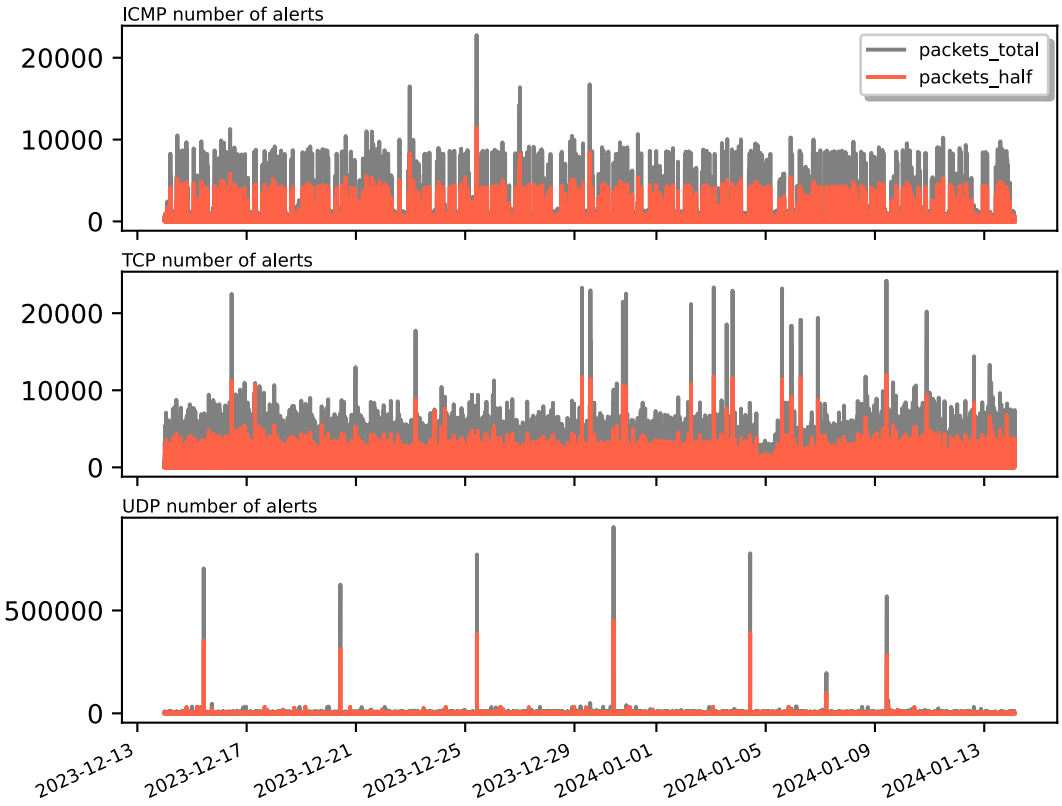

FIGURE 9 | Alerts received across 30 days of telescope activity. Considering TCP, UDP, and ICMP.

using Snort-2. It makes clear that it is still possible to identify similar patterns in the halved telescope although having less requests to work with.

Those different protocols present different types of alerts and what type of attacks were collected because of that. TCP rules often presents some scan fingerprinting regular expressions and some flag based type of identification (e.g., xmas attacks). Also, UDP rules are more payload based and can identify more types of real exploits and attacks, as well as tentative of scanning. ICMP rules focus mostly on gathering information about the types and codes and a bit of the payload in order to fingerprint the technology stack being utilized.

5 | Conclusions and Future Works

In this work, we review the literature to gather information about all the network telescopes that have been deployed in the last 23 years. We identified just 28 distinct initiatives, where we believe 18 are still active today. The reduction in the IPv4 address space for network telescopes is a trend, with the larger telescopes yielding part of their space for other organizations.

A second contribution of our work is the analysis of two network telescopes with the objective of better understanding the effects of reducing their address space. We achieve this by estimating the number of unique sources each one would receive in the case of reduction using an expected value formula and applying sampling techniques. Our findings indicate that reducing a /19 telescope in half will still maintain more than 80% of its total number of unique sources detected, although it would perceive only 50% of the number of requests. Additionally, the addressing schema adopted in the reduction, such as splitting into /24 or just selecting one /20, has a low influence of less than 0.1%.

As a third contribution we also explored how different aspects of received packets are affected by the telescope size and found that most of the proportions and patterns present in bigger telescopes also appears in small sized ones, although with less intensity and volume. We accomplish that by making analysis of protocols (e.g., TCP flags, UDP application protocols, ICMP types and code), using IDS alerts for measuring real threats and by making a time series analysis of our 30-day capture.

We believe that the problem of IPv4 scarcity will increase even more in the next few years and that Regional Internet Registries face a great challenge in managing the allocation of this still necessary resource. This greatly impacts small and new organizations that tend not to have addresses and need to wait almost 2 years to receive /22 blocks [69]. This has resulted in the popularization of IP brokers who treat addresses as a commodity, invariably increasing the costs of having an Internet business. In this context, the usage of network telescopes, although a very important cybersecurity tool, underuses addresses that could be repurposed for more active services. In that matter, we support that it is possible to reduce the telescopes and still have some threat detection and a way to gather background information.

For future works and exploration, there are some possible works that was not explored in this article for being out of the scope. IPv6 as explained in Section 1 pose great challenges for scanners and tend to be more scarce in terms of background noise. That normally makes harder to gather interesting information on attacks and events. Developing new techniques for attracting scanners and threats to the IPv6 is a need for exploring network telescopes inside that space.

References

1. D. Moore, C. Shannon, G. M. Voelker, et al., "Network Telescopes: Technical Report," (2004), https://escholarship.org/uc/item/1405b1bz.

2. M. Antonakakis, T. April, M. Bailey, et al., "Understanding the Mirai Botnet," in *26th USENIX Security Symposium (USENIX Security 17)*, (2017): 1093–1110.

3. M. Jonker, A. King, J. Krupp, et al., "Millions of Targets Under Attack: A Macroscopic Characterization of the DoS Ecosystem," in *Proceedings of the 2017 Internet Measurement Conference*, IMC '17, (Association for Computing Machinery, 2017): 100–113.

4. 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): 144–157.

5. O. Cabana, A. M. Youssef, M. Debbabi, et al., "Threat Intelligence Generation Using Network Telescope Data for Industrial Control Systems," *IEEE Transactions on Information Forensics and Security* (2021).

6. E. d'Andréa, J. François, O. Festor, et al., "Multi-Label Classification of Hosts Observed Through a Darknet," in *NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium*, (IEEE, 2023): 1–6.

7. IPv4 Global, "November 2023 Sales Report," (2023), https://ipv4.global/reports/november-2023-sales-report/. Accessed January 30, 2025.

8. J. Ronan and D. Malone, "Revisiting and Revamping an IPv6 Network Telescope," in *2023 34th Irish Signals and Systems Conference (ISSC)*, (IEEE, 2023): 1–6.

9. E. Pauley, P. Barford, and P. McDaniel, "DScope: A Cloud-Native Internet Telescope," in *32nd USENIX Security Symposium (USENIX Security 23)*, (USENIX Association, 2023): 5989–6006.

10. A. Huides, A. Santhanam, and M. Lehwess, "Identify and Optimize Public IPv4 Address Usage on AWS | Networking & Content Delivery," (2023), https://aws.amazon.com/blogs/networking-and-content-delivery/identify-and-optimize-public-ipv4-address-usage-on-aws/. Accessed January 30, 2025.

11. Google Cloud, "Google Cloud Virtual Private Cloud (VPC) Pricing," (2023), https://cloud.google.com/vpc/pricing-announce-external-ips?hl%3Dpt-br. Accessed January 30, 2025.

12. A. Camargo, L. Bertholdo, and L. Granville, "Less Is More? Exploring the Impact of Scaled-Down Network Telescopes on Security and Research," in *Anais do xlii simpósio brasileiro de redes de computadores e sistemas distribuídos*, (SBC. SBRC, 2024): 1050–1063.

13. C. Fachkha and M. Debbabi, "Darknet as a Source of Cyber Intelligence: Survey, Taxonomy, and Characterization," *IEEE Communications Surveys & Tutorials* (2016), https://doi.org/10.1109/comst.2015.2497690.

14. M. Bailey, E. Cooke, F. Jahanian, et al., "Practical Darknet Measurement," in *2006 40th Annual Conference on Information Sciences and Systems*, (IEEE, 2006): 1496–1501.

15. F. Soro, M. Allegretta, M. Mellia, et al., "Sensing the Noise: Uncovering Communities in Darknet Traffic," in *2020 Mediterranean Communication and Computer Networking Conference (MEDCOMNET)*, (2020): 1–8.

16. E. Le Malécot and D. Inoue, "The Carna Botnet Through the Lens of a Network Telescope," in *Foundations and practice of security: 6th International Symposium, FPS 2013, La Rochelle, France, October 21–22, 2013, Revised Selected Papers*, (Springer, 2014): 426–441.

17. U. Harder, M. W. Johnson, J. T. Bradley, et al., "Observing Internet Worm and Virus Attacks With a Small Network Telescope," *Electronic Notes in Theoretical Computer Science* 151, no. 3 (2006): 47–59.

18. E. Balkanli and A. N. Zincir-Heywood, "On the Analysis of Backscatter Traffic," in *39th Annual IEEE Conference on Local Computer Networks Workshops*, (IEEE, 2014): 671–678.

19. E. Cooke, M. Bailey, D. Watson, et al., "The Internet Motion Sensor: A Distributed Global Scoped Internet Threat Monitoring System," (2004), Technical Report CSE-TR-491-04.

20. P. Richter, O. Gasser, and A. Berger, "Illuminating Large-Scale IPv6 Scanning in the Internet," in *Proceedings of the 22nd ACM Internet Measurement Conference*, IMC '22, (Association for Computing Machinery, 2022): 410–418.

21. S. D. Strowes, E. Aben, R. Wilhelm, et al., "Debogonising 2a10::/12: Analysis of One Week's Visibility of a New/12," (2020).

22. CAIDA, "Historical and Near-Real-Time UCSD Network Telescope Traffic Dataset," (2024), https://www.caida.org/catalog/datasets/telescope-near-real-time_dataset. Accessed January 30, 2025.

23. M. Network, "Orion Network Telescope—Merit," (2024), https://www.merit.edu/initiatives/orion-network-telescope/. Accessed January 30, 2025.

24. 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): 344–350.

25. D. Pemberton, P. Komisarczuk, and I. Welch, "Internet Background Radiation Arrival Density and Network Telescope Sampling Strategies," in *2007 Australasian Telecommunication Networks and Applications Conference*, (2007): 246–252.

26. S. D. Chindipha, B. Irwin, and A. Herbert, "Effectiveness of Sampling a Small Sized Network Telescope in Internet Background Radiation Data Collection," in *Southern Africa Telecommunication Networks and Applications Conference (SATNAC)*, (2018).

27. F. Soro, I. Drago, M. Trevisan, et al., "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): 1–6. ISSN: 1944-0375.

28. D. Wagner, S. A. Ranadive, H. Griffioen, et al., "How to Operate a Meta-Telescope in Your Spare Time," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, IMC '23, (Association for Computing Machinery, 2023): 328–343.

29. C. Han, J. Takeuchi, T. Takahashi, et al., "Dark-Tracer: Early Detection Framework for Malware Activity Based on Anomalous Spatiotemporal Patterns," *IEEE Access* 10 (2022): 13038–13058.

30. L. Bock, *Learn Wireshark: A Definitive Guide to Expertly Analyzing Protocols and Troubleshooting Networks Using Wireshark* (Packt Publishing, 2022).

31. snort.org, "Snort–Network Intrusion Detection & Prevention System," https://www.snort.org/. Accessed August 21, 2024.

32. snort.org, "Snort—Network Intrusion Detection Community Rules," https://www.snort.org/downloads#rules. Accessed January 20, 2025.

33. G. Kayas, M. Hossain, J. Payton, et al., "An Overview of UPNP-Based IoT Security: Threats, Vulnerabilities, and Prospective Solutions," in *2020 11th IEEE Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*, (IEEE, 2020): 452–0460.

34. E. Wustrow, M. Karir, M. Bailey, et al., "Internet Background Radiation Revisited," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, (ACM, 2010): 62–74.

35. "SWITCH," https://www.switch.ch/.

36. T. Cymru, "Team Cymru Darknet Project," https://www.team-cymru.com/. Accessed July 10, 2023.

37. "Farsight Security, Cyber Security Intelligence Solutions," https://www.farsightsecurity.com/.

38. "Nicterweb—Dark Net Observation | National Institute of Information and Communications Technology Cybersecurity Laboratory," https://www.nicter.jp/en.

39. P. Richter and A. Berger, "Scanning the Scanners: Sensing the Internet From a Massively Distributed Network Telescope," in *Proceedings of ACM IMC 2019*, (ACM, 2019).

40. J. O'Hara, "Cloud-Based Network Telescope for Internet Background Radiation Collection" (PhD Thesis, 2019).

41. "The IUCC/IDC Internet Telescope," (2024), https://nocvm.iucc.ac.il/research/telescope/.

42. H. Griffioen and C. Doerr, "Could You Clean Up the Internet With a Pit of Tar? Investigating Tarpit Feasibility on Internet Worms," in *2023 IEEE Symposium on Security and Privacy (SP)*, (IEEE, 2023): 2551–2565. ISSN: 2375-1207.

43. S. Lagraa, Y. Chen, and J. François, "Deep Mining Port Scans From Darknet," *International Journal of Network Management* 29, no. 3 (2019): e2065, https://onlinelibrary.wiley.com/doi/pdf/10.1002/nem.2065.

44. M. Zakroum, J. François, M. Ghogho, et al., "Self-Supervised Latent Representations of Network Flows and Application to Darknet Traffic Classification," *IEEE Access* 11 (2023): 90749–90765.

45. E. d'Andréa, J. François, O. Festor, et al., "Multi-Label Classification of Hosts Observed Through a Darknet," in *NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium*, (IEEE, 2023): 1–6. ISSN: 2374-9709.

46. A. Houmz, G. Mezzour, K. Zkik, et al., "Detecting the Impact of Software Vulnerability on Attacks: A Case Study of Network Telescope Scans," *Journal of Network and Computer Applications* 195 (2021): 103230, https://doi.org/10.1016/j.jnca.2021.103230.

47. B. Irwin, "A Framework for the Application of Network Telescope Sensors in a Global IP Network" (PhD Thesis, 2011).

48. F. Gadhia, J. Choi, B. Cho, et al., "Comparative Analysis of Darknet Traffic Characteristics Between Darknet Sensors," in *2015 17th International Conference on Advanced Communication technology (ICACT)*, (IEEE, 2015): 59–64. ISSN: 1738-9445.

49. E. Ahmed, A. Clark, and G. Mohay, "A Novel Sliding Window Based Change Detection Algorithm for Asymmetric Traffic," in *2008 IFIP International Conference on Network and Parallel Computing*, (IFIP International, 2008): 168–175.

50. E. Ahmed, "Monitoring and Analysis of Internet Traffic Targeting Unused Address Spaces" (PhD Thesis, 2010).

51. Y. Feng, Y. Hori, K. Sakurai, et al., "A Behavior-Based Method for Detecting Distributed Scan Attacks in Darknets," *Journal of Information Processing* 21, no. 3 (2013): 527–538.

52. R. Niranjana, V. A. Kumar, and S. Sheen, "Darknet Traffic Analysis and Classification Using Numerical AGM and Mean Shift Clustering Algorithm," *SN Computer Science* 1, no. 1 (2019): 16, https://doi.org/10.1007/s42979-019-0016-x.

53. D. T. Eze, D. L. Speakman, and D. C. Onwubiko, *ECCWS 2020 19th European Conference on Cyber Warfare and Security* (Academic Conferences and Publishing Limited, 2020), Google-Books-ID: lB4EEAAAQBAJ.

54. M. Kori, V. A. Kumar, R. Pachouri, et al., "Quantitative and Qualitative Evaluation of TCP Target Ports Through Active Network Telescope," *International Journal of Information Technology* (2024): 1–15.

55. APNIC, "APIC IPv4 Exhaustion," (2024), https://www.apnic.net/manage-ip/ipv4-exhaustion/. Accessed January 30, 2025.

56. RIPE, "What Is IPv4 Run Out?," (2024), https://www.ripe.net/manage-ips-and-asns/ipv4/ipv4-run-out/. Accessed January 30, 2025.

57. LACNIC, "Estadísticas de asignación de lacnic," (2024), https://www.lacnic.net/999/1/lacnic/. Accessed January 30, 2025.

58. AFRINIC, "AFRINIC IPv4 Exhaustion Statistics," (2024), https://stats.afrinic.net/ipv4/exhaustion/ipv4_available. Accessed January 30, 2025.

59. CAIDA, "Supporting Research and Development of Security Technologies Through Network and Security Data Collection," (2018), https://apps.dtic.mil/sti/trecms/pdf/AD1054333.pdf. Accessed on January 21, 2024.

60. G. Kambourakis, M. Anagnostopoulos, W. Meng, et al., *Botnets: Architectures, Countermeasures, and Challenges* (CRC Press, 2019).

61. C. Rossow, "Amplification Hell: Revisiting Network Protocols for DDoS Abuse," in *NDSS*, (ISOC, 2014): 1–15.

62. S. Ismail, H. R. Hassen, M. Just, et al., "A Review of Amplification-Based Distributed Denial of Service Attacks and Their Mitigation," *Computers & Security* 109 (2021): 102380.

63. M. Antonakakis, T. April, M. Bailey, et al., "Understanding the Mirai Botnet," in *26th USENIX Security Symposium (USENIX Security 17)*, (USENIX Association, 2017): 1093–1110.

64. S. Torabi, E. Bou-Harb, C. Assi, et al., "Inferring and Investigating IoT-GENERATED SCANNING Campaigns Targeting a Large Network Telescope," *IEEE Transactions on Dependable and Secure Computing* 19, no. 1 (2022): 402–418.

65. T. Heinrich, R. R. Obelheiro, and C. A. Maziero, "New Kids on the DRDoS Block: Characterizing Multiprotocol and Carpet Bombing Attacks," in *International Conference on Passive and Active Network Measurement*, (Springer International Publishing, 2021): 269–283.

66. X. Feng, Q. Li, K. Sun, et al., "Off-Path Network Traffic Manipulation via Revitalized ICMP Redirect Attacks," in *31st USENIX Security Symposium (USENIX Security 22)*, (USENIX Association, 2022): 2619–2636.

67. B. Bouyeddou, B. Kadri, F. Harrou, et al., "DDos-Attacks Detection Using an Efficient Measurement-Based Statistical Mechanism," *Engineering Science and Technology, an International Journal* 23, no. 4 (2020): 870–878.

68. A. N. Blanco, "Amplification DDoS Attacks With Games Servers From the Perspective of Both the Attacker and the Defender," in *Grehack conference*, (GreHack, 2013): 1–12.

69. ARIN, "ARIN Waiting List Status Report," (2025), https://www.arin.net/resources/guide/ipv4/waiting_list/. Accessed January 30, 2025.