



Unveiling IPv6 Scanning Dynamics: A Longitudinal Study Using Large Scale Proactive and Passive IPv6 Telescopes

HAMMAS BIN TANVEER, The University of Iowa, United States

ECHO CHAN, Akamai Technologies, United States and Hong Kong Polytechnic University, China

RICKY K. P. MOK, University of California San Diego, United States

SEBASTIAN KAPPES, Max Planck Institute for Informatics, Germany

PHILIPP RICHTER, Akamai Technologies, United States

OLIVER GASSER, IPinfo, United States

JOHN RONAN, Waterford Institute of Technology, Ireland

ARTHUR BERGER, Akamai Technologies, United States and Massachusetts Institute of Technology, United States

KC CLAFFY, University of California San Diego, United States

We introduce new tools and vantage points to develop and integrate proactive techniques to attract IPv6 scan traffic, thus enabling its analysis. By deploying the largest-ever IPv6 proactive telescope in a production ISP network, we collected over 600M packets of unsolicited traffic from 1.9k Autonomous Systems in 10 months. We characterized the sources of unsolicited traffic, evaluated the effectiveness of five major features across the network stack, and inferred scanners' sources of target addresses and their strategies.

CCS Concepts: • **Security and privacy** → **Firewalls**; *Intrusion detection systems*; • **Networks** → *Network layer protocols*.

Additional Key Words and Phrases: IPv6, Network telescopes, Proactive telescopes, Internet scanning

ACM Reference Format:

Hammam Bin Tanveer, Echo Chan, Ricky K. P. Mok, Sebastian Kappes, Philipp Richter, Oliver Gasser, John Ronan, Arthur Berger, and kc Claffy. 2025. Unveiling IPv6 Scanning Dynamics: A Longitudinal Study Using Large Scale Proactive and Passive IPv6 Telescopes. *Proc. ACM Netw.* 3, CoNEXT3, Article 21 (September 2025), 24 pages. <https://doi.org/10.1145/3749221>

1 Introduction

Internet scanning is a vital tool for researchers and malicious actors alike. Researchers use scanning to characterize network dynamics [12, 20, 41] while malicious actors scan to understand a network's attack surface. Capturing and analyzing scanning traffic allows network operators and researchers

Authors' Contact Information: Hammam Bin Tanveer, The University of Iowa, Iowa City, United States, hammas-tanveer@uiowa.edu; Echo Chan, Akamai Technologies, Cambridge, United States and Hong Kong Polytechnic University, Department of Computing, Hong Kong, China, ecchan@akamai.com; Ricky K. P. Mok, University of California San Diego, CAIDA, La Jolla, United States, cskpmok@caida.org; Sebastian Kappes, Max Planck Institute for Informatics, Saarbrücken, Germany, sebastian.kappes@mpi-inf.mpg.de; Philipp Richter, Akamai Technologies, Cambridge, United States, prichter@akamai.com; Oliver Gasser, IPinfo, Seattle, United States, oliver@ipinfo.io; John Ronan, Waterford Institute of Technology, Waterford, Ireland, john.ronan@waltoninstitute.ie; Arthur Berger, Akamai Technologies, Cambridge, United States and Massachusetts Institute of Technology, Cambridge, United States, awberger@mit.edu; kc Claffy, University of California San Diego, CAIDA, La Jolla, United States, kc@caida.org.



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2834-5509/2025/9-ART21

<https://doi.org/10.1145/3749221>

to study scanner behavior and intent, *e.g.*, exploitation of specific vulnerabilities, and consequently build effective defenses against potential malicious traffic.

The advent of IPv6 has increased the complexity of the Internet scanning ecosystem. Brute-force scanning of the entire IPv6 address space is practically impossible due to the vastness of address space. Thus, IPv6 scanners must intelligently select targets to increase the likelihood of discovering active addresses. This more sophisticated approach, unlike the brute force approach to IPv4 scanning, renders it harder to detect IPv6 scanning activity. In IPv4, scanning detection leverages *darknets*: regions of address space that are inactive, *i.e.*, they neither generate traffic nor host services. Although such dark regions of IPv6 address space are plentiful, they are not as effective for detecting scans, as scanners cannot afford to exhaustively probe such vast inactive address spaces.

This reality leaves us with limited visibility into potentially malicious IPv6 scanning activities, hindering our ability to develop methods to secure and protect IPv6 networks, during a time of continuous growth in both IPv6 usage [29] and attempts to exploit it [3, 6]. The potential threat posed by IPv6 scanning has also been recognized within the IETF, which published Internet drafts on operational considerations for IPv6 blocklisting [27] and a method for operators to specify end-site prefix lengths of their networks to facilitate sensible IPv6 blocklist entries [25].




Fig. 1. Steady increase in IPv6 scan sources per week hitting the CDN. Number of IPv6 addresses (/128s) more than doubled over the two-year window. When aggregated into /64, and /48 subnets, the weekly rate tripled, from ≈ 20 to ≈ 50 -70.

To validate our assumption that IPv6 scanning traffic is a growing threat, we collaborated with a major Content Delivery Network (CDN) to capture unsolicited IPv6 packets reaching a subset of the CDN's servers (230,000 machines in over 700 Autonomous Systems (ASes)) from January 1, 2022 to January 1, 2024. Figures 1 and 2 show the steady growth in weekly counts of scanning source IPs, scanning subnets, and scanning packet volume reaching the CDN over this two-year period.¹ (Figure 13 in Appendix C shows a similar growth in number of ASes sourcing scanning traffic). Not only has scanning traffic volume increased two orders of magnitude but it is also more

¹We define a scan as a source hitting at least 100 IPv6 addresses of the CDN, with a maximum packet inter-arrival time of 3,600 seconds, which allows comparability with earlier work that used this timeout for IPv6 scan detection [52]. This earlier work analyzed the sensitivity of the timeout (3600s, 1800s, 900s) and found only marginal differences, with scan detection rates declining by single-digit percentages under shorter thresholds. Given this trade-off, and for a lower bound of 100 probed targets to declare a scan, a timeout of 1 hour was a reasonable choice for our dataset. The main point of the figures is not the number of scans, but rather its significant increase over time. The prefix aggregation accounts for scanners that choose random source addresses within larger prefixes to evade scan detection [52].




Fig. 2. Weekly scan packets (/64 aggregation), grew 100X, from 10–60M to 1B. In early 2022, scan traffic was often dominated by the most active source(s) (dashed lines), by late 2023 scanning traffic comes from a broad range of sources.

broadly distributed, no longer dominated by one or two scanning sources as it was in early 2022 (see Appendix C for more details).

This growing threat motivates our goal in this study: to develop novel and reproducible techniques to attract, capture, and analyze IPv6 scan traffic. This paper describes our three major contributions: **1. New methods and tools to capture IPv6 scan traffic:** We developed and implemented novel methods to build *proactive telescopes* to attract unsolicited IPv6 scanning traffic (§3). Our approach not only reacts to incoming traffic, but also emits network signals to attract IPv6 scanners. We designed four proactive attraction features: announcing Border Gateway Protocol (BGP) prefixes, registering domain names, issuing Transport Layer Security (TLS) certificates, inserting addresses in IPv6 hitlists. Also, we deployed reactive features—honeypots of different interaction levels to attract IPv6 scanners. *To the best of our knowledge, this work integrates the most comprehensive set of features among the prior work we surveyed (§2.3).*

In collaboration with a regional Internet Service Provider (ISP), we deployed different combinations of features within subnets of their /32 IPv6 address space, namely *honeyprefixes* (§4.2), and compared the results to two passive telescope networks. Our experimental design and deployment strategy allowed us to systematically examine the effect size of individual components in attracting scan activities. We will release the source code of our tools to support reproducibility.

2. Characterization of modern IPv6 scanning: Our proactive techniques induced an increase of three orders of magnitude in unsolicited IPv6 traffic to our previously un-probed honeyprefixes. We stratified deployment of techniques across different subnets to reveal the magnitude and type of IPv6 scanning each technique attracts. Over 90% of observed scanning traffic used ICMP ping, even though the telescopes were responsive to TCP/UDP traffic (§5). We also found that scanning traffic/strategies differed vastly by source ASN; *e.g., Internet Scanner ASNs* mostly sent TCP/UDP packets whereas most scanning traffic from *Hosting/Cloud providers* was ICMP.

3. Implications to IPv6 security: Our findings can help to evaluate and improve the efficacy of network security tools to avoid collateral damage. For example, we found scanners used an entire /30 to send scanning traffic, compared to a /96 for some cloud providers. Awareness of address allocation block size is more critical to safely deploying IPv6 blocklists than in IPv4.

2 Background

We provide an overview of the use of network telescopes in capturing unsolicited network packets from darknets, IPv6 Internet scanning techniques, and complexities introduced by IPv6 address allocation and assignment practices to network security tools/techniques.

2.1 Internet Darknets

Darknets are regions of unused IP address space that do not emit any network traffic; thus, packets in-bound toward the darknet are typically part of unsolicited scanning of the address space. Darknets are used to detect and analyze Internet-wide IPv4 scanning activities, which became pervasive after the release of tools capable of scanning the entire IPv4 address space in minutes [9]. However, the vast size of IPv6 address space makes such brute-force scanning techniques infeasible [13]. Strategic IPv6 target selection becomes critical to maximize the probability of finding responsive addresses. Attracting IPv6 scanning activity to darknets requires simulating network *liveness* by generating signals of network activities.

2.2 IPv6 Internet Scanning Techniques

An obvious approach to scanning IPv6 relies on curated *hitlists* [26, 56, 59, 61, 72]. Researchers have invested in methods to improve IPv6 scanning effectiveness [18, 31, 40, 43, 60, 64, 67, 69], which typically involve two steps: 1) collecting a set of active (seed) IPv6 addresses, 2) generating candidate scanning targets. The first step leverages sources of information containing either active IPv6 addresses (*e.g.*, AAAA records of domains) or hints of address space liveness (*e.g.*, BGP announcements). This step yields either exact addresses to scan or a narrowed search space for discovering previously unobserved IPv6 addresses. The second step uses the data acquired in step 1 to generate (previously unobserved) target addresses with a higher probability of being active. This step often involves using machine learning algorithms to find semantic patterns in observed IPv6 addresses and generating candidate addresses with similar patterns (*e.g.*, [18, 67]).

2.3 Related Work

Early attempts to capture IPv6 scanning traffic using darknets did not achieve significant visibility [19, 23, 32, 39, 54], motivating more creative approaches. Fukuda et al. [24] used *DNS backscatter* to identify widespread scanning activity. Richter et al. [52] passively collected unsolicited network packets at a large-scale commercial CDN, uncovering thousands of weekly scan events originating from dozens of different ASes.

Other methods for capturing IPv6 scanning traffic have *simulated network activity* in darknets to attract scanners. Tanveer et al. [62] launched services from an unused /56 IPv6 prefix to indicate a subnet's "liveness", resulting in an increase of several hundreds of scanning packets per day. Zhao et al. [71] uncovered how scanners utilized DNS-based methods to scan IPv6, *e.g.*, enumerating IPv6 addresses by walking the ip6.arpa zone or finding AAAA records using IPv4 PTR records. They found that most IPv6 scanners they attracted had discovered target IPv6 addresses using IPv4 PTR records. Both Scheitle et al. [57] and Pletinckx et al. [49] showed that exposing a domain name in Certificate Transparency Logs likewise induced traffic from scanners. Egloff et al. [21] uncovered how IPv6 scanners react to BGP prefix announcements.

We build on this previous work to create a more systematic, comprehensive, scalable, and reproducible approach to soliciting IPv6 scanning traffic. We use a larger (/32) network than previous studies and subdivide it into smaller subnets, each employing different combinations of techniques, to isolate the effects of individual variables in our experiment.

3 Methodology

Both prior work and our experiments (§4.1) showed that IPv6 prefixes that emit little to no traffic receive substantially less unsolicited network traffic compared to active prefixes [62, 71]. Our approach leverages both passive as well as *proactive* methods to attract and capture unsolicited IPv6 traffic. The passive approach deploys darknet-based network telescopes to capture unsolicited traffic. We design and implement a novel *proactive telescope*, which not only reacts to incoming traffic but also stimulates Internet scanners that use public data sources to find probing targets.

3.1 Darknet-based Network Telescopes

We deploy darknet-based network telescopes in two ways to establish the baseline for unsolicited traffic. Our first approach monitors the ingress traffic to a dedicated network prefix over time, similar to Cxyz et al. [19, 54]. The second approach captures traffic destined to unused portions of a live network. We leverage the internal routing table of the network's border router to forward ingress traffic destined to these unused subnets. The size of the telescope is dynamic, depending on the current subnet assignment within the network.

3.2 Proactive Network Telescope: Exposing IPv6 Addresses to Public Data Sets

We develop a *proactive network telescope* by 1) increasing the visibility of our network telescope's address space to potential scanning actors by exposing the telescope's IPv6 range in public datasets and 2) emulating live network services by reacting to incoming traffic. We advertise prefixes and addresses used by the telescope to the Internet using multiple network protocols, e.g., BGP, DNS, and TLS, resulting in their inclusion in publicly available datasets, e.g., RIPE Routing Information Service (RIS) [53], DNS zone files [33] etc., thereby signaling evidence of network activity.

Our work improves the deployment of methods proposed in prior research [62, 71] in two ways. First, our proactive telescope utilizes a significantly larger address space — a /32 compared to a /56 in previous studies — which enables us to both capture more unsolicited traffic, and to examine scanning triggers that were previously inaccessible. Second, we conduct a more comprehensive set of active experiments to attract scanning activity e.g., advertising our prefixes via TLS certificates, compared to previous works that primarily relied on like Certificate Transparency (CT) logs [35, 49, 57]. This approach allows us to gain deeper insights into how scanners integrate diverse datasets into their scanning strategies.

Specifically, we conduct the following active experiments in our proactive network telescope.

BGP announcements. We randomly select /48 prefixes from the upper half² of the ISP's /32 and announce them via BGP to create *honeyprefixes*. Then, we verify that the selected prefixes received little to no traffic in the prior month. We choose the /48 as it is the longest prefix that reliably propagates globally [58], allowing us to maximize utilization of the /32 address space.

The use of BGP announcements gave us two advantages over previous works [62, 71]: 1) it signals to scanners that these regions are active (scanners can monitor such announcements using BGP route collectors like RouteViews [65]; and 2) it provides a clear boundary of address space and significantly reduces the search space compared to the covering /32 prefix, increasing the likelihood of scanners discovering active addresses. We conducted all subsequent active experiments within these BGP-announced /48 *honeyprefixes*.

We announce the first subnet in some selected prefixes with a longer prefix length (/49-/64) to examine if scanners can discover prefixes with low visibility in routing tables. We record the time when the routes propagate to public route collectors as the start point for our experiments.

²The ISP requested we use the upper half, and prefix location does not affect generalizability of our study. (Figure 14).

Domain names. We register new domain names in multiple top-level domains (TLDs) and created AAAA records pointing the root domain (*i.e.*, eTLD+1 domain) to random IPs in our honeyprefixes, exposing them to TLD and domain monitoring tools. Scanners that monitor TLD zone files could resolve the domains names to IPs and subsequently probe our honeyprefixes.

Subdomain names. Network operators commonly use subdomain names (*i.e.*, eTLD+2 domains) for various services, such as `www`, `mail`, and `ns`, by convention. We select a total of 374 names listed on at least three of four popular subdomain name lists [10, 14, 42, 50]. We deploy AAAA records to map each subdomain name to a randomly assigned IPv6 address within a honeyprefix.

TLS certificates. As most websites adopt HTTPS, we issue TLS certificates for the root and 50 randomly chosen subdomain names to mimic the presence of web services. CT logs [2] reveal the existence of these subdomains without access to their DNS zone files, allowing scanners to discover subdomains listed in the certificates.

IPv6 hitlist. IPv6 hitlists (§2.3) periodically compile lists of responsive targets using public datasets and active measurement results. Our honeyprefix addresses were thus discovered by and then appeared in such hitlists. We also collaborated with a major hitlist provider to manually add random addresses from two honeyprefixes that would not otherwise be discovered by its process. This allowed us to identify scanners that rely on hitlists as their source of probing targets.

3.3 IPv6-Native Interaction Honeypots

Our proactive telescope engages with incoming IPv6 probes using both low and high interaction honeypots, as scanners may elicit scanning behavior that darknet telescopes cannot observe [30]. We seek to explore the efficacy of such interactions for IPv6 telescopes.

Twinklenet. We designed a lightweight low-interaction multi-protocol honeypot called *Twinklenet* to respond to unsolicited incoming traffic sent to dedicated subnets and/or IPs. Existing open-source honeypots (*e.g.*, AmpPot [36], T-pot [63], and Spoki [30]) support neither IPv6 nor multi-protocol IP aliasing (*i.e.*, handling packets to multiple destination addresses with a single network interface). An alternative approach to enable IP aliasing is to use Network Address Translation (NAT) which does not easily scale in IPv6.

Twinklenet supports IP aliasing for both IPv4 and IPv6 address spaces and responds to four popular protocols (see Appendix D). A single instance of Twinklenet can handle incoming packets toward multiple non-contiguous subnets and IPs. In addition to responding to ICMP pings, Twinklenet can bind to any TCP port of any honeyprefix's subnets and IPs to accept incoming connections. After completing the TCP handshake, and capturing the first data packets sent by the scanner, Twinklenet gracefully closes the connection. For UDP, crafting responses requires parsing protocol-specific queries. Twinklenet supports two popular UDP-based protocols: DNS and NTP. Instead of implementing full services, which attackers may exploit for abuse, Twinklenet replies with error messages to indicate responsiveness to the sender. We plan to make Twinklenet available to the research community upon the acceptance of this paper.

High-Interaction Honeypots. To examine whether scanners behave differently with full-stack systems, we integrated a high-interaction honeypot into our proactive telescope using T-Pot [63], a container-based framework emulating various services. Since each T-Pot instance can only bind to a single IPv4 address, we designed a two-stage setup to enable IPv6 and address aliasing across an entire honeyprefix (see Appendix B, Figure 12).

Traffic to the honeyprefix is redirected via an access router, which maps all packets to the prefix's first address (`::1`) and source port using DNAT. The traffic is then forwarded to a reverse proxy that performs static 6-to-4 translation to T-Pot's IPv4 address and routes it to the appropriate container based on protocol and port. We log DNAT mappings (timestamp, original IPv6 destination, source

port) to recover original destinations in T-Pot logs. The access router also mirrors traffic to a packet capturer, storing it in PCAP format.

3.4 Quantifying effects of controlled experiments

We then measure the impact of our controlled experiments in attracting IPv6 scanning traffic.

Treatment vs. Control Subnets. We divide our /32 address space into two groups: *treatment* subnets (i.e., *honeyprefixes*), and the remaining /48 prefixes (within the covering /32) as control subnets. As the BGP announcement is the initial feature to all honeyprefixes, a control subnet becomes a treatment subnet at the time its BGP announcement is observed by public route collectors.

Quantification Method. We use Bayesian structural time-series models (BSTM) [15] to quantify the effects of our experiments. We use BSTM to construct the counterfactual – what scanning activity would have looked like in the *Treatment* subnet, had the intervention (controlled experiment) not been applied – for each *honeyprefix*. Each counterfactual takes in two inputs; (1) the pre-treatment scanning activity in the treatment subnet, and (2) the scanning activity in the control subnet that received the most scanner attention during the experiment. This approach ensures that we calculate the lower bound of the effect of our experiments.

Using BSTM over traditional Difference-in-Difference method gives us 3 major advantages; First, BSTM enables us to capture complex scanning behavior e.g., sudden bursts of activity, yielding more accurate time series for the controls (prefixes for which we did not deploy features). Second, it does not expect parallel trends i.e., control and treatment to evolve similarly over time. This aspect is important as previous works [62] have shown that external factors e.g., subnet location in the covering prefix, address structure, etc. can also influence scanning activity. Third, BSTM can generate dynamic counterfactuals by tweaking the influence of each input, which results in robust uncertainty quantification.

Calculating Effect Sizes. We compute the Average Effect Size (AES) as the mean of daily difference between scanning activity in the treatment subnet and counterfactual – as shown in Figure 3 – over all days an experiment is active. We focus on two metrics: the AES for traffic volume, denoted as $\Delta_{\mathcal{H}}^{traffic}$, and the AES for the number of unique source ASNs, denoted as $\Delta_{\mathcal{H}}^{ASN}$. These metrics allow us to quantify both the volume of scanning activity and the diversity of scanners attracted by our experiments. To establish the statistical significance of change in scanners/activity, we compute 95% confidence intervals by resampling scanning activity pre and post intervention.




Fig. 3. Counterfactual vs. Observed unique ASNs for a honeyprefix

4 Telescopes and Datasets

We next describe our data capture method and datasets gathered from our three vantage points.

4.1 Passive Network Telescopes

We deployed three geographically and topologically diverse IPv6 network telescopes (Table 1)—one in a transit ISP and two in academic networks—to capture unsolicited traffic to the unused addresses. Over our three overlapping collection windows (10 months, 16 months, and 7 months), we captured over a billion unsolicited packets from 2000 unique ASes targeting more than 150M unique destinations in all of our telescopes combined.

Table 1. Overview of scanning traffic/sources captured and destinations targeted in *NT-A*, *NT-B* and *NT-C*.

Telescope	Address space	Location	Network type	Measurement start-end date	Packets received	Unique traffic sources				Unique dest. targeted		
						/128	/64	/48	ASes	/128	/64	/48
<i>NT-A</i>	/32	Southern Asia	Transit ISP	07/23 - 04/24	654M	259k	190k	138k	1.9k	134M	3.1M	61.5k
<i>NT-B</i>	/48	Ireland	Research	01/23 - 04/24	300k	1.9k	367	354	60	100k	65.5k	1
<i>NT-C</i>	/32	United States	Academic	10/23 - 04/24	250M	57k	26k	24k	507	21M	14.9M	48.8k

NT-A is hosted in an ISP network in Southern Asia with low IPv6 address space utilization. The ISP’s equipment and its customers use the initial five /48s of the APNIC-assigned /32 block. We also collaborate with this ISP to deploy our proactive network telescope (§4.2).

NT-B is an Irish IPv6 research network telescope [11] monitoring incoming traffic to an unused /48 network since June 2022.

NT-C is deployed at a U.S. academic network with a /32 assignment from ARIN. Similar to *NT-A*, it captures all the traffic sent to any unassigned subnets within the address space. The university has assigned the top half (a /33) of the /32 block to equipment and departments on campus.

4.2 Proactive Network Telescope Deployment

We implemented our proactive network telescope in the address space and infrastructure of *NT-A* (Figure 4) as described in §3.2 and §3.3. The access router forwarded ingress traffic destined for unused prefixes in ISP A to server \textcircled{A} , which performed three main functions: 1) Run the BIRD \textcircled{B} Internet routing daemon to announce honeyprefixes in ISP A’s address space, 2) execute Twinklenet \textcircled{C} to respond to incoming traffic based on the experiment’s configuration (Table 2), and 3) capture incoming traffic \textcircled{D} toward unused address space and honeyprefixes. ISP A’s operators register the honeyprefixes on APNIC’s Resource Public Key Infrastructure (RPKI) portal, so that the upstreams and peers accepted and propagated the new routes. To accommodate the high computational and memory demands, we deployed dedicated machines \textcircled{E} to host two T-Pot instances. They reacted to traffic targeting \mathcal{H}_{TPot1} and \mathcal{H}_{TPot2} . In order to map the IPs in T-Pot logs, we collected the NAT table from the DNAT gateway containing the mapping records with timestamps.

4.3 Honeyprefix features and configurations

We describe the implementation of seven features that we experimentally deployed, in multiple combinations, to infer which data sources scanners use to create their target lists. Specifically, we deployed 27 honeyprefixes with different configurations (Table 2) to investigate how Internet scanners discover live hosts and react to different network behaviors.

Table 2. Configuration of honeyprefixes.

Honeyprefixes	BGP	Aliased	ICMP	TCP	UDP	Domain	Subdomain	IPv6 Hitlist *
Description	§4.3.7	§4.3.3	§4.3.4	§4.3.5		§4.3.1 and §4.3.2		§4.3.6
\mathcal{H}_{Alias}	1×/48	✓	●	✗	✗	✗	✗	Aliased
\mathcal{H}_{TCP}	✗ [‡]	✗	●	web, remote	✗	✗	✗	✗
\mathcal{H}_{UDP}	1×/48	✗	●	✗	53, 123	✗	✗	NA, UDP53, ICMP
\mathcal{H}_{Com}	1×/48	✗	✗	web	✗	2×.com	✗	NA, TCP80, TCP443
$\mathcal{H}_{Org/net}$	1×/48	✗	✗	web	✗	1×.org, 1×.net	✓(only .net)	NA, TCP80, TCP443
$\mathcal{H}_{Combined}$	1×/48	✗	●	web, remote	53, 123	1×.net	✓	NA, TCP80, TCP443
\mathcal{H}_{TPot1}	1×/48	✓	●	See Appendix B, Table 5		2×.com	1×.com	Aliased, Manual
\mathcal{H}_{TPot2}	1×/48	✓	●	See Appendix B, Table 5		2×.com	1×.com	Aliased, Manual
$\mathcal{H}_{Specific}$	/49-/64	✗	✗	✗	✗	✗	✗	✗
\mathcal{H}_{BGP}	3×/48	✗	✗	✗	✗	✗	✗	✗

Note: ●/○ represent the entire subnet/specific addresses were responsive to ICMP, respectively.

‡: We configured BIRD to announce the prefix, but the announcement failed to reach the Internet due to a technical problem.

*: Aliased/NA represents the aliased/non-aliased prefixes list, respectively. ICMP, TCP80, TCP443, and UDP53 denote the hitlists that reported at least one IP in the subnet as responsive to the corresponding protocol.

web: 80, 443, 8080, 8443; remote: 22, 23, 2323, 3389.

4.3.6 IPv6 Hitlist. The hitlist measurements [26] discovered some prefixes and IPs in honeyprefixes. The hitlist’s aliased/non-aliased prefix list included all three aliased (\mathcal{H}_{Alias} , \mathcal{H}_{TPot1} , and \mathcal{H}_{TPot2}) and five non-aliased (\mathcal{H}_{TCP} , \mathcal{H}_{UDP} , \mathcal{H}_{Com} , $\mathcal{H}_{Org/net}$, and $\mathcal{H}_{Combined}$) honeyprefixes, respectively. The hitlist also revealed some IPs with open ports on UDP port 53, and TCP port 80 and 443. We collaborated with the hitlist maintainers to manually add two IPs (one at the beginning of the address space, and one random in the honeyprefix) into each hitlist category. In total, we manually insert 40 addresses (20 per honeypot) across 10 hitlist categories. We label this feature as **H**.

4.3.7 BGP only prefixes. We announced 19 BGP only honeyprefixes (3 \mathcal{H}_{BGP} and 16 $\mathcal{H}_{Specific}$). None of them responded to incoming traffic. The three \mathcal{H}_{BGP} prefixes were identical and were announced with a prefix length of /48. We announced a set of 16 $\mathcal{H}_{Specific}$ prefixes with lengths ranging from /49 to /64, with one prefix for each length. We refer to the remaining unused dark address space announced through ISP A’s covering /32 prefix as $\mathcal{H}_{Control}$.

4.4 Metadata and data processing

We map source IPs to ASs and countries using CAIDA’s RouteViews Prefix-to-AS mapping [16] and IPinfo’s geolocation database [34] with datasets collected on the same day as the packet timestamps, ensuring the timeliness of the mapping. We use ASdb [73] to identify the type of AS and Zeek [70] to aggregate packets into flows to facilitate our analysis.

5 Results

We present a comparative analysis of the scanning sources attracted by 3 of our telescopes (§5.1). We conduct a deep dive into scanning traffic collected by our proactive telescope *NT-A* from July 2023 to April 2024. We characterize scanning sources and properties (§5.2), describe our method for setting up our controlled experiments and quantifying their effects (§5.3), and outline the different scanning strategies we observe (§5.4).

5.1 Telescope comparison

To evaluate the efficacy of our proactive telescope (*NT-A*) in attracting scanning traffic, we compare the characteristics of the scan sources across all three telescopes. Overall, *NT-A* accounted for almost 70% of all unsolicited traffic that we captured, with 98.4% of that traffic targeting honeyprefixes. It also attracted scanning traffic from the most diverse set of sources among the three. *NT-C* received

most of the remaining $\sim 30\%$ of the traffic but from a much smaller set of sources. *NT-B* captured only a small fraction of the total traffic, owing to its order of magnitude smaller address space.

We calculated the *Jaccard similarity* to enumerate the overlap of scanning sources between the 3 telescopes. We calculate *Jaccard Similarity* for scan sources at 3 different prefix lengths, /32, /64 and /128, as follow: $\mathcal{J}(NT_y^{agg}, NT_x^{agg}) = \frac{|Sources_{agg}NT_y \cap Sources_{agg}NT_x|}{|Sources_{agg}NT_y \cup Sources_{agg}NT_x|}$. The average Jaccard similarity across prefix aggregation levels of /32, /64 and /128, across all telescope combinations, was ~ 0.1 , indicating that the sets of sources observed at different telescopes were highly distinct. The highest \mathcal{J} of 0.2 was observed between *NT-A* and *NT-C* at a /32 aggregation.

Despite the overall dissimilarity, a few overlapping sources accounted for most of the unsolicited traffic, and targeted the largest number of unique destinations within our telescopes. Comparing source IPs (/128) targeting *NT-A* vs. *NT-B*, the overlapping sources generated 4.3% of unsolicited traffic received by *NT-A*. This fraction increased to 96.3% when we aggregated by a longer prefix length *i.e.*, /64. Common sources between *NT-A* and *NT-C* aggregated by /64 generated 97.3% and 99.2% of the unsolicited traffic received by *NT-A* and *NT-C*, respectively.

NT-A exclusively captured the most exploratory scanners *i.e.*, scanners that were responsible for targeting the most unique destination IPs. For example, overlapping sources (grouped by /64) between *NT-A* and *NT-B* were responsible for most unsolicited traffic (96.8%) in both telescopes, but these sources targeted only 45.1% of all unique destination IPs probed in *NT-A*.

Hosting/cloud providers were responsible for the majority ($> 50\%$) of unsolicited traffic sent toward *NT-A* and *NT-C*, which together received over 99% of all unsolicited traffic across the three telescopes. Notably, *Amazon AWS/Google Cloud Platform* were the top contributors to unsolicited network traffic in *NT-A* and *NT-C*, respectively.

Key Takeaway: This comparison shows that our proactive telescope (*NT-A*) received the most scanning activity from the most distinct sources, many of which were also responsible for the majority of scanning traffic observed in other telescopes. Therefore, in the following sections, we focus on characterizing scanning traffic received by *NT-A*.


5.2 Unsolicited Traffic Sources by network and traffic type

During our controlled experiments, *NT-A* received 654M unsolicited network packets from 259k distinct IPv6 addresses spanning over 1.9k unique ASNs. As the sizes of the subnets that scanners used vary, we aggregated source IP addresses using two common prefix lengths (/48, /64). Table 3 presents the top five ASNs sourcing such traffic and Figure 6 shows a geographic distribution of scan sources (a detailed breakdown of top 20 ASNs is shown in Appendix F). *Amazon-02* and *CERNET* accounted 11 for 80% of all observed unsolicited traffic. While the traffic volumes from these two source ASNs were comparable, they differed significantly in the number of source IP addresses: 44K for *Amazon-02* and only 46 unique source addresses for *CERNET*.

Table 3. Top 5 ASN sources of unsolicited traffic. (See Appendix F, Table 8 for a longer list of top ASNs.)

AS name	ASN	Packet count (share %)	Unique sources		
			/128	/64	/48
AMAZON-02	29014	289M (42.6%)	44k	336	251
CNGI-CERNET	23910	265M (39.0%)	46	4	4
AMAZON-AES	14618	33.7M (4.96%)	11k	25	15
TSINGHUA UNI.	45576	23.8M (3.51%)	5	2	1
HURRICANE	6939	12.7M (1.87%)	3.5k	136	112

Figure 5 shows the type of all 1.9K source ASes with ASdb [73] and statistics on each transport protocol used, proportion of total observed packets, unique destinations, and unique sources. We observed entire IP prefixes and ASes dedicated to IPv6 Internet scanning for various purposes, e.g., the Internet Measurement AS [5]. We manually assigned four such network entities to the *Internet Scanner* category, e.g., *AlphaStrike Labs*, *Shadow Server*.



(a) Unique packet count categorized by AS types. ICMP dominates the traffic, except for *Internet Scanners* which mostly use TCP. (b) Unique source utilized/destination IPs probed by different AS types. Scanners from R&E ASes probe an order of magnitude more target IPs than scanners from other ASes.

Fig. 5. Breakdown of scanner sources by AS type and their proportionate contributions to unique scan sources (/128), unique destinations targeted within our telescope *NT-A* and total number of scanning packets sent broken down by protocol.

Consistent with previous observations [52, 62], hosting/cloud providers generated the most unsolicited traffic, followed by Research/Education (R&E) networks. ICMPv6 was the most common protocol, responsible for 91.6% of all unsolicited traffic. The *Internet scanner* category contained most (90%) of the unique source addresses that sent unsolicited traffic to *NT-A*. These scanners use distributed IP addresses from a covering prefix as large as a /30. The *Internet Scanners* category also presents a different traffic type distribution than other categories, predominantly TCP. R&E networks probed the most distinct target addresses in *NT-A*, accounting for 95% of all destinations.

5.3 Scan Traffic Attraction by Controlled Experiments

We evaluate the effectiveness of our methods to attract IPv6 scanning traffic. We quantify their impact on traffic volume, scanner diversity, and characterize scanner behavior and target scope.

Impact on scan traffic volume. To quantify the increase in scan traffic volume following our controlled experiments, we used $\Delta_{\mathcal{H}}^{traffic}$ as defined in §3. Table 4 shows effect sizes for honeyprefixes averaged over all post intervention days. Our controlled experiments led to a statistically significant increase in scanning activity across all honeyprefixes (Tables 4 and 10). We observed the largest $\Delta_{\mathcal{H}}^{traffic}$ increase in $\mathcal{H}_{Tpot1}^{TLS}$; unsolicited traffic increased by 224k packets/day for this honeyprefix. Figure 10 shows the scanning traffic/scanner effect sizes in $\mathcal{H}_{specific}$ subnets. The $\Delta_{\mathcal{H}}^{traffic}$ sizes in $\mathcal{H}_{specific}$ subnets fall in two distinct classes; $\Delta_{\mathcal{H}}^{traffic}$ was below 10k for 75% of $\mathcal{H}_{specific}$ subnets and above 80k for the rest. The largest $\Delta_{\mathcal{H}}^{traffic}$ in $\mathcal{H}_{specific}$ subnet was observed in a covering honeyprefix from which we announced a /61 IPv6 prefix (§4.3.7); it received over 10M scanning packets in a single day. However, for the most part, scanning traffic in $\mathcal{H}_{specific}$ subnets was sporadic and we observed no correlation between the prefix length we announced and the scanning traffic we observed to that prefix.




Fig. 6. Geographical distribution of /128 scanner source addresses in *NT-A*. We used IPinfo’s geolocation database [34] from April 2024 to infer scanner location from source IPs. Germany is the most prominent source country of scanning due to the vast amount of IP addresses used by *AlphaStrike Labs*.

Table 4. Effect sizes of controlled experiments. Confidence Intervals for traffic rounded to nearest 1000.

Honeyprefix	\mathcal{H}_{bGP}	\mathcal{H}_{Alias}	\mathcal{H}_{TCP}	\mathcal{H}_{UDP}	\mathcal{H}_{com}	$\mathcal{H}_{org/net}$	$\mathcal{H}_{combined}$	$\mathcal{H}_{limit}^{port}$	$\mathcal{H}_{privatist}^{port}$	\mathcal{H}_{TLS}^{port}
$\Delta_{\mathcal{H}}^{traffic}$	3,865	10,670	2,525	112,000	11,903	8,159	11,493	1,115	3,457	224,176
[95% CI]	[4k–3k]	[11k–10k]	[3k–3k]	[113k–110k]	[13k–11k]	[8k–8k]	[12k–11k]	[1k–1k]	[8k–1k]	[236k–213k]
$\Delta_{\mathcal{H}}^{ASN}$	36	25	10	34	35	39	32	26	2	27
[95% CI]	[45–28]	[37–13]	[10–10]	[44–25]	[49–21]	[47–32]	[46–18]	[26–25]	[38–34]	[117–61]

Impact on scanner source diversity. To quantify the increase in scanner source diversity, we use $\Delta_{\mathcal{H}}^{ASN}$ as defined in §3 and shown in Table 4. $\mathcal{H}_{org/net}$ had the largest increase in $\Delta_{\mathcal{H}}^{ASN}$ (39 source ASNs/day), indicating that typical scanners use DNS zone files as seed files to generate scan targets. In contrast, most $\mathcal{H}_{specific}$ had an average $\Delta_{\mathcal{H}}^{ASN}$ of 1 (not listed in Table 4), reflecting the limited scanner diversity observed. This low $\Delta_{\mathcal{H}}^{ASN}$, coupled with sporadic scanning traffic, we attribute to the limited propagation of BGP announcements for $\mathcal{H}_{specific}$. Among the 36 public BGP collectors monitored, announcements from $\mathcal{H}_{specific}$ reached only 5 collectors, compared to an average of 28 collectors for other *honeyprefixes*. This result aligns with the established constraint that /48 is the minimum globally routable IPv6 prefix.

5.3.1 Characterization of Scanner Behavior. To understand scanner behavior, we reviewed scanning traffic/sources per day following our controlled experiments and uncovered IPv6 scanner behaviors.

Changes in scanner attention. We found that the per day difference in scanners/scanning activity fluctuates significantly. Figures 7 and 8 show that scanner attention increases immediately after our controlled experiment begins, marked by the initial BGP announcement for the honeyprefix. This immediate increase lasts different periods of time for different *honeyprefixes*, suggesting that scanner attention depends on the scanning trigger, *i.e.*, which service operates in the *honeyprefix*.

Figure 7 shows that $\Delta_{\mathcal{H}}^{traffic}$ and $\Delta_{\mathcal{H}}^{ASN}$ converge to a stable lower value after 15 days for $\mathcal{H}_{specific}$ and 40 days for $\mathcal{H}_{Tpot1}^{TLS}$. This also explains that despite \mathcal{H}_{com} attracting the most unique scanning source ASes, it received more than an order of magnitude less scanning traffic per day than the interactive honeyprefix \mathcal{H}_{Tpot1} .

Scanner response to additional triggers. We wanted to understand how scanners react to additional triggers inside a *honeyprefix* a month after the initial BGP announcement. To this end, we include two additional triggers (inclusion in IPv6 hitlist and registering TLS certificates) in \mathcal{H}_{Tpot1} (red and blue vertical lines in Figure 7 and Figure 8). Scanning traffic jumped an order of magnitude with each new trigger, despite the number of source ASes remaining similar. This observation reveals that scanners increase the resources expended on a *honeyprefix* if they detect additional activity in a prefix from disparate data sources.

Scanner reaction to trigger retraction. To test whether scanners consistently update their sources for seeding IPv6 targets, we retracted the BGP announcements for two of the remaining \mathcal{H}_{BGP} (not shown here) after 4 months. Within hours, the persistent scanning activity diminished to a negligible level, indicating that IPv6 scanners frequently refresh their data sources.

5.3.2 Characterization of scanner’s scope. We defined IP prefix boundaries for our controlled experiments by announcing a /48 for each of our *honeyprefixes* from within our telescope’s /32 covering prefix. To analyze whether the IPv6 scanners stayed within the address scope of the honeyprefixes, we analyzed the number of /48 prefixes that scanners probed (Figure 9). 95% of scanners probed only two /48 prefixes, 99.92% probed <11 prefixes (excluding $\mathcal{H}_{specific}$ subnets), and 99.97% probed <27 (the total number of honeyprefixes we deployed in our experiment). Only 55 of 191k scanner source IPs scanned beyond this scope, with one probing 61.5k of 65k /48 prefixes. Non-honeyprefix traffic constituted a meager 1.6% of total unsolicited traffic we received in the covering /32 prefix, half of which targeted the first 16 /48 prefixes (xxxx:xxxx:0000::/48 to xxxx:xxxx:000f::/48), with the rest distributed across the remaining subnets. This finding highlights that **1)** /48 is an adequate prefix size to host such controlled experiments and **2)** most IPv6 scanners scan within the prefix address scope of the honeyprefix.

Key Takeaway: Our findings offer foundational guidance for future IPv6 telescope design. A /48 prefix is essential – not only because it is the minimum globally routable prefix, but also because it allows precise attribution of scanning activity; broader covering prefixes can cause collateral scanning spillover. We observe that scanners strictly confine their probing to the announced /48, reinforcing its utility for controlled experiments. Operationally, resource-constrained deployments should prioritize exposing addresses via TLS certificates and deploying both high and low interaction honeypots. These mechanisms significantly amplify scanner engagement, and hosting multiple services within the same /48 can compound visibility—provided the prefix remains consistently announced via BGP.

5.4 Characterizing scanning strategies

The unique characteristics of each honeyprefix suggest which services and protocols were targeted by IPv6 scanners. Figure 11 shows the features that scanning /48 subnets probed. Each *x*-axis label represents a combination of features in nine responsive honeyprefixes that scanners hit. We matched probe packets with honeyprefixes’ features by protocols, (e.g., ICMP and TCP/UDP destination ports), the destination IPs and the probing time. We used the probing time to distinguish features that shared the same protocols and destination addresses. For example, both (sub-)domain names (**S** and **D** in Figure 11) and their corresponding TLS certificates (**s** and **d** in Figure 11) resolve to web services at the same IPs. We labeled scanning attempts targeted to domain names or TLS certificates for probes received before or after certificate issuance, respectively.




Fig. 7. Heatmap of traffic effect size in selected *honeyprefixes*. The figure shows the effect sizes after the initial BGP announcement – \mathcal{I}_{00} – of 3 *honeyprefixes*. We observe an immediate increase in scanning traffic following the announcement and see traffic increase by an order of magnitude for every additional trigger – inclusion in IPv6 hitlist (red line) and registering TLS certificates (blue line).




Fig. 8. Longitudinal view of per-day traffic and ASN effect sizes in selected *honeyprefixes*. This figure illustrates the effect size in terms of unique source ASNs before and after the initial BGP announcement. Unlike traffic volume, the count of unique source ASes remains consistent, indicating that scanners consistently probe the *honeyprefixes* but send less scanning traffic after an initial phase (evident in Figure 7).

ICMP probing. As the entire \mathcal{H}_{TPot1} , \mathcal{H}_{TPot2} , and \mathcal{H}_{Alias} were responsive to ICMP, over 50k unique sources probed addresses in these three *honeyprefixes* *solely* using ICMP pings (**B** in Figure 11a and 11b). About 400 sources probed \mathcal{H}_{UDP} and \mathcal{H}_{RDNS} solely with ICMP, despite both *honeyprefixes* having only 3 addresses responding to ICMP. All these sources targeted only the first address ($: : 1$) of the subnet, although two other addresses were active. Two sources discovered the random address for which we enabled ICMP response in $\mathcal{H}_{Combined}$. They also scanned UDP ports and other non-responsive parts of the subnet (magenta bars in “IU” and “IO”, **A** in Figure 11a).

Domain registration. Many sources used the root AAAA DNS records to compile target lists (i.e., **D**). All *honeyprefixes* with domain names (\mathcal{H}_{Com} , $\mathcal{H}_{Org/net}$, $\mathcal{H}_{Combined}$, \mathcal{H}_{TPot1} , and \mathcal{H}_{TPot2}) showed evidences of scanning targeted to the DNS-mapped IPs (**D** in x -axis, indicated by **C1** in Figure 11a and 11b). These scanners likely learned the names (and corresponding addresses) from




Fig. 9. Number of /48 prefixes targeted by observed scanners. 98.4% of scan traffic we received in our proactive telescope *NT-A*, was directed to one of our *honeyprefixes*, with 99.9% of this traffic targeting *honeyprefixes* other than $\mathcal{H}_{specific}$. That is most IPv6 scanners stay within the prefix address scope of the *honeyprefix*




Fig. 10. $\Delta^{traffic}$ effect size for $\mathcal{H}_{specific}$ subnets. We observe that most (75%) $\mathcal{H}_{specific}$ *honeyprefixes* receive <10k scanning packets and the remaining receive > 80k. This is a result of prefixes longer than /48 not propagating through BGP.

published zone files, as there was a month-long gap between our domain registration (Sept. 19, 2023) and the associated IP appearing on the IPv6 hitlist (Oct. 23, 2023).

Subdomain names and TLS certificates. No sources could detect common subdomains without TLS certificates (*i.e.*, \underline{s} always came with \underline{S} in Figure 11, indicated by **D**). Even for root domain names, ~300 scanners observed by \mathcal{H}_{Com} , $\mathcal{H}_{Org/net}$ only probed the honeyprefixes after the issuance of TLS certificates (**C2** in Figure 11a). Scanners quickly reacted to new TLS certificates. The first scanner from DigitalOcean arrived 7 seconds after certificate issuance, evidence that Certificate Transparency logs [2] were the source of this target list.


IPv6 Hitlist. Manual addition of hitlist entries for \mathcal{H}_{TPot1} and \mathcal{H}_{TPot2} let us isolate the effect of using hitlists. 115 and 111 sources probed addresses in \mathcal{H}_{TPot1} and \mathcal{H}_{TPot2} (and corresponding protocols) specified in the hitlists (blue/orange bars in Figure 11b), respectively. We label scanners with **H** that probed addresses that the hitlist automatically discovered. Although we could not confirm that scanners used the IPv6 hitlist, these IPs saw traffic levels similar to the honeypots. Over 600 sources pinged \mathcal{H}_{UDP} , which had one responsive IP address that we manually inserted into the IPv6 responsive addresses hitlist (**E** in Figure 11b).

Scanning tactics. Most scanners used only ICMP to measure the liveness of the networks. Combining TCP and ICMP probing was a common strategy (**F** in Figure 11a). Scanners probed our honeyprefixes with multiple protocols, particularly for the high-interaction honeypots. Ten scanners probed our honeypots with all ICMP, TCP, and UDP protocols.


Key Takeaway: All features, except subdomains with TLS certificates, effectively attracted scanners. Most scanners use ICMP to test connectivity, but our proactive telescope can detect sophisticated, potentially malicious scanners that integrated multiple data sources to explore live hosts and services in the targeted networks.

6 Discussion

We discuss limitations and future direction of this work.



(a) Many sources discovered our honeyprefixes using domain names and CTLog, not the IPv6 hitlist.



(b) Manually inserted addresses in the TPots reveals that scanners use the IPv6 hitlist to populate target lists.

Fig. 11. Combination of tactics adopted by scanning traffic sources, grouped by /48. Each colored bar shows the number of scanners using the same probing strategy against honeyprefixes. X-axis labels represent combinations of features we introduced (§4.3). O represents scanners probing non-responsive protocols or ports. Red arrows highlight the findings described in §5.4. We exclude results from \mathcal{H}_{TCP} because its /48 prefix was not successfully announced via BGP, leading to only limited probing activity — specifically, traffic from just seven /48s across five ASes.

Generalizability of Our Findings. While our study collected one of the largest known datasets of unsolicited IPv6 traffic, our findings may not generalize to all networks or geographic locations. This limitation has been observed in IPv4 network telescopes, where researchers found that scanning activity can be localized [51], introducing biases across different vantage points [66]. Additionally, certain network types—such as public cloud infrastructures—have been found to attract more diverse and aggressive scanning behavior than others [48]. We expect similar patterns to emerge in IPv6 scanning. Moreover, IPv6 scanning is heavily influenced by the evolution of target generation algorithms, whose effectiveness depends on the seeds and datasets they use [68]. Some scanners may rely on signals not included in our experiments, potentially resulting in false negatives.

Misattribution of Scanning. We took careful steps to accurately attribute unsolicited traffic to each of our experiments—for example, by precisely timing and placing /48 *honeyprefixes* within the

/32 covering prefix. However, our methodology relies on external datasets and mechanisms that may introduce inaccuracies. For instance, our manual inspection revealed misclassifications in ASdb [73]. Another source of errors could be related to routing security. Due to the increasing adoption of RPKI, the upstream provider for *NT-C* would reject BGP announcements of honeyprefixes until the corresponding Route Origin Authorization (ROA) records were successfully registered. As a result, we cannot rule out the possibility that scanners used ROA data, rather than BGP propagation, to infer the presence and location of honeyprefixes.

Implications for Operational Security. Traditional IP blocklists and threat intelligence feeds (e.g., [1, 22]) rely on stable, individual IPs to identify malicious actors. However, in IPv6, network operators assign large prefixes (e.g., /48 to /112) to end-hosts [45], which complicates attribution and mitigation. Overblocking causes collateral damage, while underblocking enables evasion via address rotation. Accurate mapping of prefixes to hosts requires (1) insights into real-world IPv6 allocation and rotation practices [47, 55], and (2) large-scale measurement of scanner behavior. Our study contributes to the latter. For example, we identified scanners operating across /32 prefixes. IETF drafts have offered recommendations for IPv6 blocklisting [27] and to specify end-site prefix lengths of their networks [25]. Our findings can empirically ground such operational guidance as we find evidence of IPv6 scanners utilizing as large as a /29 prefix and as small as a /112 prefix to conduct scanning. Without end-site prefix length specification, developing high accuracy IP/prefix based blocklisting tools will be a complex challenge. We plan to integrate datasets collected and tools developed in this work for (1) creating effective IPv6 defenses e.g., commercially deployed IPv6 only telescopes and (2) helping existing solutions/suggestions to improve [7, 8, 46].

The deployment of IPv4 network telescopes is often constrained by address exhaustion, a limitation that IPv6 does not face. However, existing IPv6-focused platforms that could support such efforts (e.g., Honeydv6 and 6Guard [17, 44]) are outdated or no longer maintained. Our open-source tools enable network operators to deploy IPv6 proactive telescopes within their own networks to gather network-specific threat intelligence. We plan to improve the scalability and reliability of our software and deployment framework, particularly for high-interaction honeypots, to further strengthen the security of IPv6 networks.

Small networks may lack the resources or autonomy to implement all the features we explored. Operators might have limited control over BGP announcements or hold only a single /48 prefix. In such cases, they must balance capability with resource constraints. For instance, rather than deploying multiple honeyprefixes, operators can focus their efforts by placing honeypots, registering domain names, and issuing TLS certificates near the beginning of their assigned address block, where scanners are most likely to initiate probing, to maximize visibility of scanning activities.

7 Conclusion

In this work, we designed and developed a proactive telescope — a reproducible system to capture a broader and more representative range of IPv6 scanner behaviors within an ISP network. Our deployment collected the largest IPv6 telescope dataset to date, in terms of both address space and traffic volume, enabling us to investigate the ecosystem of IPv6 unsolicited traffic. We carefully analyzed the (in)effectiveness of common network features and public datasets that scanners might leverage for IPv6 network discovery. Our longitudinal measurement data further revealed the strategies scanners use to integrate target addresses and explore network services. These insights are valuable for network operators and researchers in developing effective tools to mitigate potential threats in IPv6 networks. We will publish our code and deployment instructions to support reproducibility this work.

Acknowledgments

We thank our shepherd and the anonymous reviewers for their insightful comments. This material is based on research sponsored by the National Science Foundation (NSF) grants NSF OAC-2131987, OAC-2319959, and OAC-2450552. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the funding agencies.

References

- [1] AbuseIPDB. <https://www.abuseipdb.com>. Accessed: 2024-05-16.
- [2] Certificate transparency. <https://certificate.transparency.dev>.
- [3] Crowdsec's new cybersecurity majority report highlights the rise of ipv6 in cybercriminal activities | cybersecurity dive. <https://www.cybersecuritydive.com/press-release/20230727-crowdsecs-new-cybersecurity-majority-report-highlights-the-rise-of-ipv6-in-2/>. (Accessed on 05/16/2024).
- [4] GoDaddy. <https://www.godaddy.com>.
- [5] internet-measurement.com. <https://internet-measurement.com/>. (Accessed on 05/16/2024).
- [6] Ipv6 threats are on the rise as adoption grows - sdxcentral. <https://www.sdxcentral.com/articles/analysis/ipv6-threats-are-on-the-rise-as-adoption-grows/2023/08/>. (Accessed on 05/16/2024).
- [7] The no-compromise solution for attack surface management. [Online; accessed 2025-05-27].
- [8] Sixint: World-leader in ipv6 network intelligence for cyber-security, reconnaissance, and forensics. [Online; accessed 2025-05-27].
- [9] The ZMap Project. <https://zmap.io/>, 2024. Accessed 2024-05-15.
- [10] Assetnote. Commonspeak2. <https://github.com/assetnote/commonspeak2-wordlists/>, 2018.
- [11] Anonymized author(s). Anonymized publication.
- [12] Shehar Bano, Philipp Richter, Mobin Javed, Srikanth Sundaresan, Zakir Durumeric, Steven J Murdoch, Richard Mortier, and Vern Paxson. Scanning the internet for liveness. *ACM SIGCOMM Computer Communication Review*, 48(2):2–9, 2018.
- [13] Robert Beverly, Ramakrishnan Durairajan, David Plonka, and Justin P. Rohrer. In the IP of the Beholder: Strategies for Active IPv6 Topology Discovery. In *Proceedings of ACM Internet Measurement Conference*, 2018.
- [14] bitquark. DNSpop. <https://github.com/bitquark/dnspop>, 2016.
- [15] Kay H Brodersen, Fabian Gallusser, Jim Koehler, Nicolas Remy, and Steven L Scott. Inferring causal impact using bayesian structural time-series models. 2015.
- [16] CAIDA. RouteViews Prefix to AS mappings. https://catalog.caida.org/dataset/routeviews_prefix2as.
- [17] ChristosKon. Github - christoskon/honeyd-ipv6: Git repo for honeyd ipv6. official folder: <https://redmine.cs.uni-potsdam.de/projects/honeydv6/files>. [Online; accessed 2025-04-27].
- [18] Tianyu Cui, Gaopeng Gou, Gang Xiong, Chang Liu, Peipei Fu, and Zhen Li. 6gan: Ipv6 multi-pattern target generation via generative adversarial nets with reinforcement learning. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, pages 1–10, 2021.
- [19] Jakub Czyz, Kyle Lady, Sam G Miller, Michael Bailey, Michael Kallitsis, and Manish Karir. Understanding ipv6 internet background radiation. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 105–118, 2013.
- [20] Zakir Durumeric, Eric Wustrow, and J Alex Halderman. {ZMap}: Fast internet-wide scanning and its security applications. In *22nd USENIX Security Symposium (USENIX Security 13)*, pages 605–620, 2013.
- [21] Isabell Egloff, Raphael Hiesgen, Maynard Koch, Thomas C. Schmidt, and Matthias Wählisch. A Detailed Measurement View on IPv6 Scanners and Their Adaption to BGP Signals. *Proceedings of the ACM on Networking (PACMNET)*, 3(CoNEXT3), September 2025.
- [22] FireHOL. All cybercrime IP feeds. <https://iplists.firehol.org/>, 2025.
- [23] Matthew Ford, Jonathan Stevens, and John Ronan. Initial results from an ipv6 darknet13. In *International Conference on Internet Surveillance and Protection (ICISP'06)*, pages 13–13. IEEE, 2006.
- [24] Kensuke Fukuda and John Heidemann. Who Knocks at the IPv6 Door?: Detecting IPv6 Scanning. In *Proceedings of the Internet Measurement Conference 2018, IMC '18*, pages 231–237, New York, NY, USA, 2018. ACM.
- [25] Oliver Gasser, Randy Bush, Massimo Candela, and Russ Housley. Publishing End-Site Prefix Lengths. <https://datatracker.ietf.org/doc/draft-ietf-opsawg-prefix-lengths/03/>. Active Internet-Draft.
- [26] Oliver Gasser, Quirin Scheitle, Pawel Foremski, Qasim Lone, Maciej Korczynski, Stephen D. Strowes, Luuk Hendriks, and Georg Carle. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *Proceedings of the 2018 Internet Measurement Conference*, New York, NY, USA, 2018. ACM.

- [27] Fernando Gont and Guillermo Gont. Implications of IPv6 Addressing on Security Operations. <https://datatracker.ietf.org/doc/draft-ietf-opsec-ipv6-addressing/00/>. Expired Internet-Draft.
- [28] Google. pcapgo. <https://pkg.go.dev/github.com/google/gopacket/pcapgo>, 2020.
- [29] Google. IPv6 Adoption. <https://www.google.com/intl/en/ipv6/statistics.html>, 2024. Accessed 2024-05-15.
- [30] Raphael Hiesgen, Marcin Nawrocki, Alistair King, Alberto Dainotti, Thomas C. Schmidt, and Matthias Wählisch. Spoki: Unveiling a New Wave of Scanners through a Reactive Network Telescope. In *Proceedings of USENIX Security Symposium*, 2022. Accessed: 2023-2-14.
- [31] Bingnan Hou, Zhiping Cai, Kui Wu, Jinshu Su, and Yinqiao Xiong. 6hit: A reinforcement learning-based approach to target generation for internet-wide ipv6 scanning. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, page 1–10. IEEE Press, 2021.
- [32] Geoff Huston and Mirjam Kuhne. Background radiation in ipv6. *The ISP Column, APNIC*, 2010.
- [33] ICANN. About zone file access. <https://www.icann.org/resources/pages/zfa-2013-06-28-en>, 2013.
- [34] IPinfo. IPinfo Geolocation Database. <https://ipinfo.io/products/ip-geolocation-database>, 2024.
- [35] Brian Kondracki, Johnny So, and Nick Nikiforakis. Uninvited guests: Analyzing the identity and behavior of certificate transparency bots. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 53–70, Boston, MA, August 2022. USENIX Association.
- [36] Lukas Kramer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. AmpPot: Monitoring and Defending Amplification DDoS Attacks. In *Proceedings of the 18th International Symposium on Research in Attacks, Intrusions and Defenses*, 2015.
- [37] Let’s Encrypt. Rate limits. <https://letsencrypt.org/docs/rate-limits/>, October 2021.
- [38] Let’s Encrypt. Challenge types. <https://letsencrypt.org/docs/challenge-types/>, February 2023.
- [39] Chenhuan Liu, Qiankun Liu, Shanshan Hao, CongXiao Bao, and Xing Li. IPv6-Darknet Network Traffic Detection. In Xingming Sun, Xiaorui Zhang, Zhihua Xia, and Elisa Bertino, editors, *Artificial Intelligence and Security - 7th International Conference, IC AIS 2021, Dublin, Ireland, July 19-23, 2021, Proceedings, Part II*, volume 12737 of *Lecture Notes in Computer Science*, pages 231–241. Springer, 2021.
- [40] Zhizhu Liu, Yinqiao Xiong, Xin Liu, Wei Xie, and Peidong Zhu. 6Tree: Efficient dynamic discovery of active addresses in the IPv6 address space. *Comput. Netw.*, 155(C):31–46, may 2019.
- [41] Linda Markowsky and George Markowsky. Scanning for vulnerable devices in the internet of things. In *2015 IEEE 8th International conference on intelligent data acquisition and advanced computing systems: technology and applications (IDAACS)*, volume 1, pages 463–467. IEEE, 2015.
- [42] Daniel Miessler. Combined subdomains. <https://github.com/danielmiessler/SecLists/tree/master/Discovery/DNS>, 2023.
- [43] Austin Murdock, Frank Li, Paul Bramsen, Zakir Durumeric, and Vern Paxson. Target generation for internet-wide ipv6 scanning. In *Proceedings of the 2017 Internet Measurement Conference, IMC ’17*, page 242–253, New York, NY, USA, 2017. Association for Computing Machinery.
- [44] mzweilin. Github - mzweilin/ipv6-attack-detector: Google summer of code 2012 project, supported by the honeynet project organization. [Online; accessed 2025-04-27].
- [45] T. Narten, G. Huston, and L. Roberts. IPv6 address assignment to end sites, March 2011.
- [46] F. Gont SI6 Networks. Implications of ipv6 addressing on security operations, 2 2023. [Online; accessed 2025-04-10].
- [47] Ramakrishna Padmanabhan, John P Rula, Philipp Richter, Stephen D Strowes, and Alberto Dainotti. Dynamips: Analyzing address assignment practices in ipv4 and ipv6. In *Proceedings of the 16th international conference on emerging networking experiments and technologies*, pages 55–70, 2020.
- [48] Eric Pauley, Paul Barford, and Patrick McDaniel. {DScope}: A {Cloud-Native} internet telescope. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 5989–6006, 2023.
- [49] Stijn Pletinckx, Thanh-Dat Nguyen, Tobias Fiebig, Christopher Kruegel, and Giovanni Vigna. Certifiably vulnerable: Using certificate transparency logs for target reconnaissance. In *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*, pages 817–831, 2023.
- [50] rbsec. dnscan. <https://github.com/rbsec/dnscan>, 2022.
- [51] Philipp Richter and Arthur Berger. Scanning the scanners: Sensing the internet from a massively distributed network telescope. In *Proceedings of ACM IMC*, Oct 2019.
- [52] Philipp Richter, Oliver Gasser, and Arthur Berger. Illuminating large-scale IPv6 scanning in the internet. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 410–418, 2022.
- [53] RIPE Network Coordination Centre. Routing information service (RIS). <http://ris.ripe.net/>.
- [54] John Ronan and David Malone. Revisiting and revamping an IPv6 network telescope. In *Proceedings of Irish Signals and Systems Conference (ISSC)*, 2023.
- [55] Erik Rye, Robert Beverly, and Kimberly C Claffy. Follow the scent: Defeating ipv6 prefix rotation privacy. In *Proceedings of the 21st ACM Internet Measurement Conference*, pages 739–752, 2021.

- [56] Erik Rye and Dave Levin. IPv6 Hitlists at Scale: Be Careful What You Wish For. In *Proceedings of the ACM SIGCOMM 2023 Conference*, ACM SIGCOMM '23, page 904–916, New York, NY, USA, 2023. Association for Computing Machinery.
- [57] Quirin Scheitle, Oliver Gasser, Theodor Nolte, Johanna Amann, Lexi Brent, Georg Carle, Ralph Holz, Thomas C. Schmidt, and Matthias Wählisch. The Rise of Certificate Transparency and Its Implications on the Internet Ecosystem. In *Proceedings of the Internet Measurement Conference 2018, IMC '18*, page 343–349, New York, NY, USA, 2018. Association for Computing Machinery.
- [58] Khwaja Zubair Sediqi, Lars Prehn, and Oliver Gasser. Hyper-Specific Prefixes: Gotta Enjoy the Little Things in Interdomain Routing. *ACM SIGCOMM Computer Communication Review*, 52, June 2022.
- [59] Guanglei Song, Jiahai Yang, Lin He, Zhiliang Wang, Guo Li, Chenxin Duan, Yaozhong Liu, and Zhongxiang Sun. AddrMiner: A comprehensive global active IPv6 address discovery system. In *2022 USENIX Annual Technical Conference (USENIX ATC 22)*, pages 309–326, Carlsbad, CA, July 2022. USENIX Association.
- [60] Guanglei Song, Jiahai Yang, Zhiliang Wang, Lin He, Jinlei Lin, Long Pan, Chenxin Duan, and Xiaowen Quan. DET: Enabling Efficient Probing of IPv6 Active Addresses. *IEEE/ACM Trans. Netw.*, 30(4):1629–1643, feb 2022.
- [61] Lion Steger, Liming Kuang, Johannes Zirngibl, Georg Carle, and Oliver Gasser. Target Acquired? Evaluating Target Generation Algorithms for IPv6. In *Proceedings of the Network Traffic Measurement and Analysis Conference (TMA)*, June 2023.
- [62] Hammas Bin Tanveer, Rachee Singh, Paul Pearce, and Rishab Nithyanand. Glowing in the dark: Uncovering {IPv6} address discovery and scanning strategies in the wild. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 6221–6237, 2023.
- [63] Deutsche Telekom Security GmbH. T-Pot - The all in one multi honeypot platform. <https://github.com/telekom-security/tpotce>.
- [64] Johanna Ullrich, Peter Kieseberg, Katharina Krombholz, and Edgar Weippl. On Reconnaissance with IPv6: A Pattern-Based Scanning Approach. In *2015 10th International Conference on Availability, Reliability and Security*, pages 186–192, 2015.
- [65] University of Oregon. RouteViews project. <http://www.routeviews.org/>.
- [66] Gerry Wan, Liz Izhikevich, David Adrian, Katsunari Yoshioka, Ralph Holz, Christian Rossow, and Zakir Durumeric. On the origin of scanning: The impact of location on internet-wide scans. In *Proceedings of the ACM Internet Measurement Conference*, pages 662–679, 2020.
- [67] Grant Williams, Mert Erdemir, Amanda Hsu, Shraddha Bhat, Abhishek Bhaskar, Frank Li, and Paul Pearce. 6Sense: Internet-Wide IPv6 Scanning and its Security Applications. In Davide Balzarotti and Wenyuan Xu, editors, *33rd USENIX Security Symposium, USENIX Security 2024, Philadelphia, PA, USA, August 14-16, 2024*. USENIX Association, 2024.
- [68] Grant Williams and Paul Pearce. Seeds of scanning: Exploring the effects of datasets, methods, and metrics on IPv6 internet scanning. In *Proceedings of the 2024 ACM on Internet Measurement Conference*, IMC '24, pages 295–313. ACM, November 2024.
- [69] Tao Yang, Bingnan Hou, Zhiping Cai, Kui Wu, Tongqing Zhou, and Chengyu Wang. 6Graph: A graph-theoretic approach to address pattern mining for Internet-wide IPv6 scanning. *Comput. Netw.*, 203(C), feb 2022.
- [70] Zeek. An open source network security monitoring tool. <https://zeek.org>, 2024.
- [71] Liang Zhao, Satoru Kobayashi, and Kensuke Fukuda. Exploring the discovery process of fresh ipv6 prefixes: An analysis of scanning behavior in darknet and honeynet. In *Passive and Active Measurement Conference*, pages 95–111, 2024.
- [72] Johannes Zirngibl, Lion Steger, Patrick Sattler, Oliver Gasser, and Georg Carle. Rusty Clusters? Dusting an IPv6 Research Foundation. In *Proceedings of the 2022 Internet Measurement Conference*, New York, NY, USA, 2022. ACM.
- [73] Maya Ziv, Liz Izhikevich, Kimberly Ruth, Katherine Izhikevich, and Zakir Durumeric. ASdb: a system for classifying owners of autonomous systems. In *Proc. ACM IMC*, 2021.

A Ethics

This work does not raise ethical concerns.

B T-Pot Infrastructure




Fig. 12. Overview of IPv6-enabled T-Pot infrastructure.

Table 5. Honeypot containers we deployed in our T-Pot instance and the corresponding ports.

Honeypots	Protocol (destination ports)	\mathcal{H}_{TPot1}	\mathcal{H}_{TPot2}
cowrie	TCP (22-23)	✓	✗
mailoney	TCP (25)	✓	✓
snare	TCP (80)	✓	✓
citrixhoneypot	TCP (443)	✓	✓
ciscoasa	UDP (5000), TCP(8443)	✓	✓
redishoneypot4	TCP (6379)	✓	✗
adbhoney	TCP (5555)	✓	✓
sentrypeer	UDP (5060)	✗	✓
dionaea	TCP (20-21, 42, 81, 135, 443, 445, 1433, 1723, 1883, 3306, 27017), UDP (69)	✓	✓
ddospot	UDP (19, 53, 123, 161, 1900)	✓	✓
compot_kamstrup_382	TCP (1025, 50100)	✗	✓
elasticpot	TCP (9200)	✗	✓
dicompot	TCP (11112)	✗	✓

C Breakdown of Scanner Sources in CDN by Country and Network Type

Also of interest is the country of origin and network type from which the scans are initiated. Table 6 shows the top 20 ASes, ordered by number of packets. For countries, United States and China dominate. Two of the ASes belong to cybersecurity companies. If the scanner has ill-intent, they could likely use cloud service providers or datacenters (though of course others would also be using these platforms) and these platforms are the most popular. Compared to an earlier study by Richter *et. al* [52], covering 15 months starting January 2021, the scan traffic reported here is much more dispersed. In Table 6 the top AS accounted for 18% of the packets across three /64's, whereas earlier, the top three /64's accounted for 87% of the packets.

Table 6. Top 20 source ASes at the CDN by scan packets over the entire measurement window (packets shown for /64 source aggregation).




Fig. 13. Weekly number of source ASes of the IPv6 scans at the CDN.

Rank	AS type	Packets	Scan sources		
			/48s	/64s	/128s
#1	Transit (global)	4.68B (17.6%)	1	3	2745
#2	Datacenter (CN)	4.08B (15.4%)	10	12	45
#3	Cybersecurity (US)	3.74B (14.1%)	7	7	367
#4	Datacenter (US)	3.17B (12.0%)	1	1	11
#5	Cloud (CN)	2.60B (9.8%)	15	17	310
#6	Cloud (CN)	2.42B (9.1%)	6	7	36
#7	Datacenter (CN)	1.72B (6.5%)	2	2	11
#8	Cloud (US/global)	899M (3.4%)	35	43	3312
#9	Cloud (US/global)	833M (3.1%)	4	4	53
#10	Datacenter (CN)	609M (2.3%)	1	1	4
#11	Cloud (US/global)	533M (2.0%)	12	12	2277
#12	Cloud (US/global)	392M (1.5%)	12	19	4475
#13	Cloud (US/global)	360M (1.4%)	22	22	41
#14	Cloud (US/global)	228M (0.9%)	7	7	21
#15	Cybersecurity (US)	91M (0.3%)	2	2	198
#16	Datacenter (CN)	44M (0.2%)	32	138	142
#17	Cloud (US)	28M ($\leq 0.1\%$)	1	1	2
#18	University (CN)	20M ($\leq 0.1\%$)	1	2	2
#19	Datacenter (CA)	14M ($\leq 0.1\%$)	1	1	1
#20	Research (DE)	14M ($\leq 0.1\%$)	1	1	1

Figure 13 shows the weekly number of ASes sending IPv6 scan packets. Like the number of source addresses and packets, the number of scanning ASes grows steadily over time (cf. Section 1).

D Technical Details of Twinklenet

We implemented Twinklenet in Go, using Berkeley Packet Filters (BPF) and PCAPGO [28] to filter, capture, and respond to incoming packets. Twinklenet leverages raw socket to overrides the system’s default TCP/IP stack and routing table, enabling it to return the responses to the sender, emulating live hosts within the honeyprefixes. Table 7 lists the services Twinklenet emulates.

Table 7. Protocols and interactions supported by Twinklenet.

Protocols	Request	Response(s)
ICMP/ICMPv6	ICMP/ICMPv6 Echo request	ICMP/ICMPv6 Echo reply
TCP	TCP SYN to an open port	Complete three-way handshake and close the connection with FIN
	Other TCP packet to an open port	TCP RST
NTP (over UDP)	Any client NTP packet	NTP Kiss-of-Death packet (Reference Identifier=DENY)
DNS (over UDP)	Any DNS query	DNS response with response code SERVFAIL

E Honeyprefix Location in NT-A




Fig. 14. Hilbert map of the /32 IPv6 address space of our telescope in the ISP’s network. We distributed our /48 honeyprefixes randomly across the upper half of unused portion of the /32. The colors of the pixels show different phrases of honeyprefix deployment.

F Sources of Scan Traffic Observed in *NT-A*

Table 8. Top 20 ASN sources of unsolicited traffic in *NT-A* (continued from Table 3).

Rank	AS name	ASN	Packet count (share %)	Unique sources		
				/128	/64	/48
			<i>... (see Table 3)</i>			
#6	BJ-GUANGHUAN-AP	55960	7.1M (1.04%)	695	8	4
#7	ALIBABA-CN-NET	45102	7.0M (1.03%)	17	9	8
#8	INTERNET-MEASUREMENT	211298	6.9M (1.01%)	511	14	14
#9	AKAMAI-AMS	33905	5.5M (0.81%)	3	3	3
#10	PONYNET	53667	4.3M (0.63%)	10	10	4
#11	UONET	3582	3.4M (0.50%)	2	2	2
#12	WESTCLOUDDATA	135629	3.4M (0.50%)	374	3	1
#13	NEXTLAYER	1764	1.9M (0.28%)	11	11	7
#14	CHOOA	20473	1.0M (0.15%)	207	207	54
#15	LEITWERT-RESEARCH	29108	1.0M (0.15%)	11	11	11
#16	BCIX	62193	1.0M (0.15%)	1	1	1
#17	MWN	12816	880K (0.13%)	30	2	2
#18	AKAMAI-ASN1	20940	779K (0.11%)	6	6	6
#19	RICAWEBSERVICES	26832	733K (0.11%)	4	4	2
#20	ORACLE-BMC	31898	676K (0.10%)	42	38	34

Received December 2024; revised June 2025; accepted June 2025