# A Detailed Measurement View on IPv6 Scanners and Their Adaption to BGP Signals

ISABELL EGLOFF, HAW Hamburg, Germany
RAPHAEL HIESGEN, HAW Hamburg, Germany
MAYNARD KOCH, TU Dresden, Germany
THOMAS C. SCHMIDT, HAW Hamburg, Germany
MATTHIAS WÄHLISCH, TU Dresden, Germany

Scanners are daily visitors of public IPv4 hosts. Scanning IPv6 nodes successfully is still a challenge, which an increasing crowd of actors tries to master. In this paper, we analyze IPv6 scanning under various network conditions to disclose the impact on scanning. We deploy four IPv6 network telescopes, including a reactive /48 telescope and a proactive /32 telescope that is periodically reconfigured by changing BGP announcements. We provide a longitudinal study of eleven months and classify the observed scanners w.r.t. their temporal behavior, their target and network selection strategies, as well as their individual tools, fingerprints, and correlations across categories. We find that silent subnets of larger covering prefixes remain mainly invisible, whereas BGP prefix announcements quickly attract attention by scanners. Based on our findings, we derive operational guidance on how to deploy network telescopes to increase visibility for IPv6 scanners and understand corresponding biases.

CCS Concepts: • **Networks** → **Network measurement**; **Network experimentation**; **Public Internet**; *Network layer protocols*; **Network security**.

Additional Key Words and Phrases: network telescopes, darknet, IPv6, BGP, measurement instrumentation

**ACM Reference Format:**
Isabell Egloff, Raphael Hiesgen, Maynard Koch, Thomas C. Schmidt, and Matthias Wählisch. 2025. A Detailed Measurement View on IPv6 Scanners and Their Adaption to BGP Signals. *Proc. ACM Netw.* 3, CoNEXT3, Article 15 (September 2025), 23 pages. https://doi.org/10.1145/3749215

## 1 Introduction

Scanning is prevalent on the Internet [1, 15]. Researchers, commercial services, as well as malicious actors explore the IPv4 address spaces regularly and with high intensity. Many of these build their tooling on stateless scanning [16, 26], which allows for traversing all IPv4 addresses within less than an hour. The huge IPv6 address space renders this impossible. It is thus vital for research, operations, and security to develop methods for observing and understanding the emerging ecosystem of IPv6 scanners.

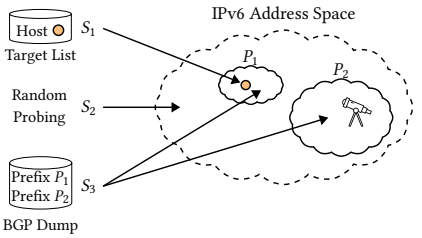


Fig. 1. IPv6 address probing methods.

Authors' Contact Information: Isabell Egloff, HAW Hamburg, Hamburg, Germany, isabell.egloff@haw-hamburg.de; Raphael Hiesgen, HAW Hamburg, Hamburg, Germany, raphael.hiesgen@haw-hamburg.de; Maynard Koch, TU Dresden, Dresden, Germany, maynard.koch@tu-dresden.de; Thomas C. Schmidt, HAW Hamburg, Hamburg, Germany, t.schmidt@haw-hamburg.de; Matthias Wählisch, TU Dresden, Dresden, Germany, m.waehlisch@tu-dresden.de.


ACM 2834-5509/2025/9-ART15
https://doi.org/10.1145/3749215

Scanning IPv6 nodes means mastering the challenge of finding them. Different strategies have been developed (see Figure 1), more are expected to emerge. Simple random probing, such as $S_2$ in Figure 1, does not scale for an Interface Identifier (IID) space of $2^{64}$ addresses per subnet. Nevertheless, random probing can be successful if its distribution accounts for IPv6-inherent address structures. Scanners can also execute along a target list ($S_1$), which may be taken from some static hitlist or generated algorithmically on the fly. Scanners can also react to external events ($S_3$), such as the publication of DNS records or—more network-centric—the announcement of new prefixes via BGP.

Measuring the IPv6 scanning landscape and obtaining a meaningful picture is about as challenging as the scanning itself. An operator of a public IPv6 network without any further attractors may experience close to no scan visits. This challenges the operation of network telescopes (and honeypots), which are the traditional observatories of scanners. Deploying specific attractors to influence scanning will increase packet reception, but also introduces biases that misguide conclusions.

Internet scanning regularly precedes malicious attacks. Measuring malicious scanners may therefore serve as a quantitative assessment of the current threat landscape [25, 27, 33]. Accurate measurements and quantitative assessments, however, are intricate. Distributed honeypots have shown to span a very selective view on the Internet [43]. Operating an IPv4 telescope and interpreting its data correctly likewise exhibits high complexity [40]. Soundly assessing the limitations and biases of our measurement instruments is essential but difficult. In IPv6, these difficulties exponentiate, since any scope of perception remains tiny compared to the huge address space. In view of this, the importance of understanding the interplay between network characteristics and scanner reactions becomes even more evident.

In this paper, we explore the reactivity of scanners to network embeddings of telescopes by performing a large measurement study, in which we deploy four telescopes with different network attachments for eleven months. With this setup, we want to study the effective adaptation of scan traffic to these telescopes, in particular to changes in BGP. After summarizing the state of the art on IPv6 address characteristics and IPv6 scan studies (§2), we present our main contributions:

(1) In controlled experiments, we observe IPv6 scanning during 11 months from four network telescopes with contrasting network properties. This includes a proactive telescope that varies visibility in BGP with up to 17 IPv6 prefixes (§3). We capture and analyze 51M packets (§4).
(2) We derive a taxonomy and classify scanners following a thorough analysis of their temporal behavior, of their target selection strategies, and their origins (§5).
(3) We analyze and compare all four telescopes during a 12 week observation period, revealing unique traffic patterns in terms of packet intensity, sources, and scanning strategies (§6).
(4) We analyze the reactions of scanners to our 8 months of changing BGP signals and find a strong correlation with packet arrival (+286%). 70% of all scanners were observed only once, 9% target uniformly all of the announced prefixes and account for 63% of all packets (§7).
(5) We derive operational guidance on the visibility and expected biases of IPv6 telescopes (§8).

## 2 Background and Related Work

The huge IPv6 address space challenges Internet measurements. Discovering active endpoints is in focus of of active research, industry projects, and malicious actors. Notably, TU Munich maintains a comprehensive hitlist of responsive IPv6 hosts [20, 70], which estimates a lower bound of all active IPv6 addresses. While scanners actively explore the address space, incoming scan activities can be observed through telescopes to analyze scanners. Table 1 summarizes related scientific work.

**IPv6 address characteristics.** Manual configuration often leads to four easy-to-predict and memorable patterns [22]. (*i*) Low-byte addresses use Interface Identifier (IID) bytes of zero except

Table 1. Overview of related work that observes IPv6 scanners. In contrast to prior work, we perform an active and controlled experiment on the control plane and deploy several prefixes of different characteristics for comparison. *Passive* telescopes originate no packets. *Traceable* telescopes originate or receive traffic controlled by the authors. *Active* telescopes react to connection attempts. Duration is some weeks (w) or months (m).

| | Time | | Telescope | | | | Announced | Application | Packets |
|---|---|---|---|---|---|---|---|---|---|
| Publication | Year | Duration | Size | Passive | Traceable | Active | Prefix Sizes | Attractors | Received |
| Ford *et al.* [17] | '04-'06 | 16m | /48 | ✓ | ✘ | ✘ | /48 | | 12 |
| Czyz *et al.* [14] | '12-'13 | 3m | 5×/12 | ✓ | ✘ | ✘ | 5× /12 | | 209M |
| Fukuda *et al.* [19] | '17-'18 | 9m | /37 | ✓ | ✘ | ✘ | /37 | | 15k |
| Strowes *et al.* [55] | '20 | 1w | /12 | ✘ | ✓ | ✘ | /12, 4×/32, 4×/48 | | 6.5M |
| Liu *et al.* [36] | '19-'21 | 4m, 3w, 3w | /20 | ✓ | ✘ | ✘ | /20 | | 2.9M |
| Tanveer *et al.* [57] | '21 | 48w | /56, 24×/64 | ✘ | ✓ | ✘ | (unclear) | Web, DNS, NTP, TOR | 14.6M |
| Richter *et al.* [47] | '21-'22 | 14m | (CDN logs) | ✘ | ✘ | ✘ | n/a | | 2.04B |
| Ronan *et al.* [49] | '22 | 6m | /48 | ✓ | ✘ | ✘ | /48 | | 5.13k |
| Zhao *et al.* [69] | '23 | 6m | /56, 12×/64 | ✘ | ✘ | ✓ | /48 | DNS exposure | 33M |
| This work | '23-'24 | 11m | /32, 3×/48 | ✓ | ✓ | ✓ | /29, /32–/48, /48 | Productive subnet | 51M |

for the least significant byte, *e.g.,* `2001:db8::1`, (*ii*) service port addresses embed the port of a running service in the IID, *e.g.,* `2001:db8::443` for HTTPS, (*iii*) IPv4-based addresses embed the IPv4 address of the network interface in the IID, *e.g.,* `2001:db8::192.0.0.1`, and (*iv*) wordy addresses include semantic hints such as `2001:db8::cafe`.

Stateless address autoconfiguration [60] maps MAC addresses and thus also follows predictable patterns in its assignment strategy if not used with a privacy extension [42]. One strategy is to scan these address ranges first, as the density of active hosts is expected to be higher. When observing scanners in this work, we analyze target address patterns to infer their scanning strategies.

**IPv6 scanning strategies.** Early drafts of RFC 7707 [22] dating back to 2012 already described active IPv6 address discovery strategies. Ullrich *et al.* [61] were the first to publish on pattern-based scanning techniques, followed by Foremski *et al.* [18] who probabilistically modeled IPv6 addresses. Starting in 2019, multiple studies appeared on target generation algorithms (TGA) [9, 11–13, 30, 31, 37–39, 52, 66–68], which can be roughly divided into (*i*) static and (*ii*) dynamic TGAs [53]. While static algorithms only generate potential candidates for scanning based on a fixed training set, dynamic TGAs adjust their training set by evaluating the activity of generated addresses immediately through active scanning. In our scanning analysis, we use the order of selected target addresses to infer scanning strategies.

**Passive network telescopes.** Ford *et al.* [17] were the first to analyze IPv6 scanners by examining the background radiation of a /48 prefix in 2005 and 2006. They only received 12 packets. Seventeen years later, Ronan *et al.* [49] re-examined the same /48 prefix. They collected over 5k packets in 6 months. Most sources sent ICMPv6, about a fifth TCP, and a few UDP. Scanners generally iterated the network prefix systematically. In 2021, Liu *et al.* [36] investigated the background radiation of a previously unused /20 IPv6 prefix and collected more than 2.9M packets during roughly 6 months. The authors observed similar shares of transport protocols. 95% of the packets originating from only 10 sources. Our measurements also use passive telescopes but of controlled, variable BGP exposure.

**Temporary network telescopes.** Czyz *et al.* [14] analyzed IPv6 Internet background radiation (IBR) by announcing five /12 address blocks assigned to each of the five RIRs—the covering prefixes of the full address space allocated by LIRs. This experiment extracted darknet traffic, *i.e.,* packets to the subnets (of the /12) that were never allocated nor routed. This darknet received only 5% (209M) of all packets, as most packets were sent to address sub-spaces that was previously allocated by others. Strowes *et al.* [55] analyzed traffic to a newly announced /12 prefix for one week to check visibility

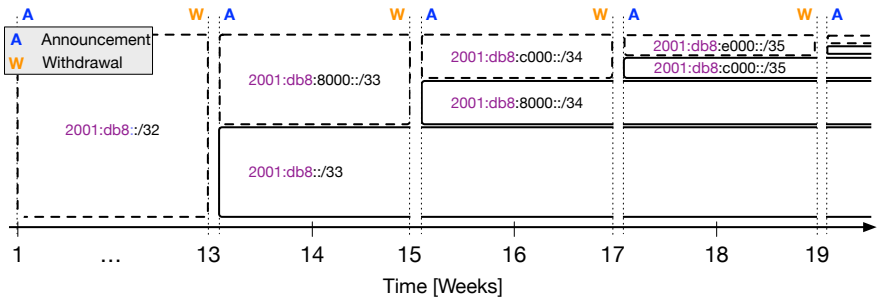


Fig. 2. After an initial baseline observation phase (week 1-13), we recursively split one prefix into two more-specific prefixes every two weeks (dotted vertical lines) until we announce 17 prefixes, our most-specific prefix is /48. (`2001:db8::/32` is not the prefix we announce but reserved for documentation.)

and reachability in IPv6 routing. 4× /32 and 4× /48 prefixes of a covering /29 were announced separately and actively probed with ICMPv6 packets, which generated most of the received traffic. In contrast, we focus on the visibility of BGP announcements for scanners at network telescopes.

**Server logs.** Richter *et al.* [47] analyzed IPv6 scanning activities using the firewall logs of a large CDN but exclude ICMP traffic and TCP packets to ports 80 and 443. Two sources originated 70% of all traffic. 75% of scan sources (aggregated by /64 source prefix) only focused on addresses discoverable via DNS. While authors often used /64 source address aggregation, they concede that aggregation is usually case-specific. Our grouping of scanner activities relies on scan sessions, which showed to be a stable measure under aggregation.

**Attracting scanners on the application layer.** Tanveer *et al.* [57] examined how IPv6 host activities influence the behavior of IPv6 scanners in a previously unused /56 subnet. Each of the six traffic classes was deployed in four randomly selected /64 subnets. Publicly visible active services (NTP, Tor, DNS zones) attracted significantly more scan traffic than client activities from the telescopes (web crawls and DNS probes). Most probes targeted low-byte addresses and random IIDs. Scanners focused on a few subnets or scanned across all 256 subnets.

Zhao *et al.* [69] investigated the effect of address exposure via DNS using a previously unused /56 network. They published addresses via four DNS methods (*e.g.,* PTR record for random addresses). Each method was deployed once in a /64 darknet and once in a /64 honeynet. Associating IPv6 addresses with domains that have IPv4 PTR records attracted over 99.99% of the scans. In honeynets, scanners were less focused, targeting exposed addresses and unexposed addresses in the vicinity more equally. Low-byte scans comprised only 0.03% of all scans.

Parallel to our work, Tanveer *et al.* [56] deployed a proactive network telescope with multiple application layer attractors to evaluate the effectiveness of these features across the network stack, and to infer the strategy of scanners.

In contrast, this work focuses on the network conditions of proactive and reactive IPv6 telescopes and analyzes how scanners react. We do not proactively attract scanners via specific application services but include a prefix in our analysis that hosts an actively operating subnet. These controlled properties of our four telescopes allow for selectively characterizing scanner behavior as well as the corresponding bias envisioned in the different telescopes.

## 3 Measurement Method and Setup

In this section, we introduce our measurement infrastructure consisting of four different network telescopes (T1–T4). Each telescope has specific properties that reflect different aspects of network

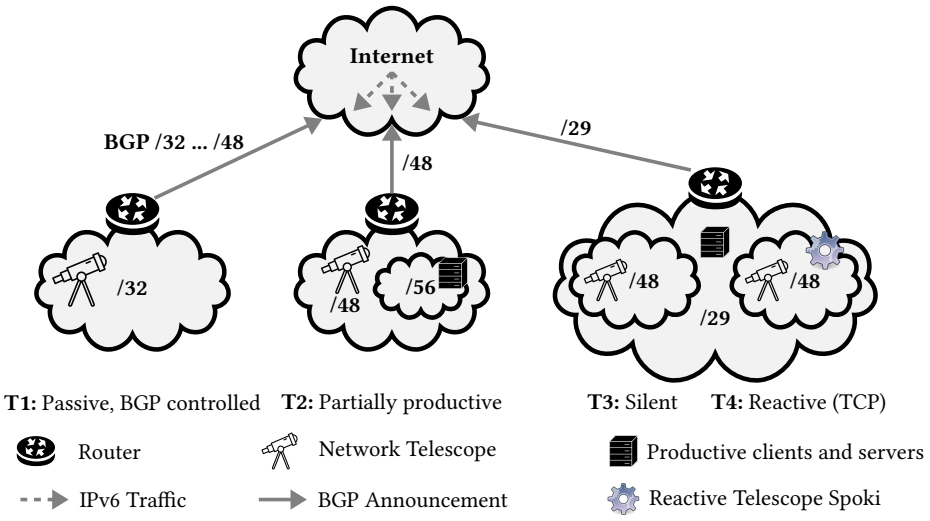


Fig. 3. Our experiment setup, showing all four network telescopes, their sizes, and announced prefixes in BGP.

embedding (see Figure 3). We define our network experiments along with our scanner detection method.

### 3.1 Network Telescopes

Scanning success in the IPv6 address space requires strategy and guidance, reasonably based on external observations and measurements. Hitlists and target generation algorithms can help to reduce the search space, but scanners may pursue different strategies, which we want to discover.

**T1: BGP controlled /32 – /48.** This untainted /32 IPv6 prefix was first announced at the beginning of our experiments. All addresses in this prefix are passive, and neither is assigned to an endpoint nor to other network services such as the DNS. In the first 12 weeks of our experiment, this prefix was kept stable. During this initial observation period, we collected baseline data. Thereafter, we used this prefix to study the influence of BGP announcements, announced prefix sizes, as well as the number and relative position of (sub-)prefixes on scan activities. For this active, controlled experiment we carefully selected the duration of the announcements, as well as the number and size(s) of the announced prefix(es), to avoid interference with scanner behavior in any other way.

To identify a reasonable announcement period, we explore the temporal correlation of scanners with a new BGP announcement during our initial observation. We do not find convergence, but after two weeks the number of scan sources from new prefixes reduces notably (see Figure 4). Thereby, we decide on a two-weeks prefix split interval as a trade-off between overall experiment duration and time to observe how scanners react to our changes.

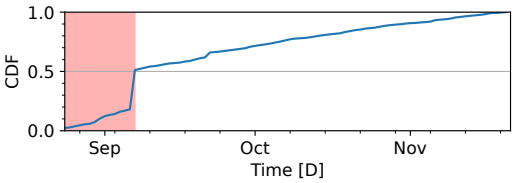


Fig. 4. Number of new prefixes hosting scan sources discovered during our initial 12-week observation period. Two weeks (red area) mark the trade-off to learn enough without running an announcement too long.

Every two weeks, we withdraw our prefixes for one day. On the day after, we announce a set of new prefixes that contains (*i*) a new pair of prefixes created by splitting one prefix into two more-specific prefixes of equal size and (*ii*) all previously announced prefixes except the covering

prefix. Since low-byte addresses tend to attract more traffic, we chose to split the most-specific prefix that does not contain the respective low-byte address if possible (*e.g.,* the low-byte address of `2001:db8::/32` is `2001:db8::1`). This creates two new prefixes with low-byte addresses that do not (byte-wise) match the low-byte addresses of the previously announced prefixes. The number of prefixes we announce increases by one at each interval, including two previously unannounced prefixes. Figure 2 visualizes our asymmetric prefix splitting process.

**T2: Partially productive /48.** This /48 prefix has been continuously announced for 13 years and a /56 subnet (not separately announced) is in productive use since then. The productive subnet contains web servers, end hosts, and IoT devices, several of which with persistent DNS entries; traffic from or to the /56 subnet is excluded from our measurements. In addition, one address within the /48 prefix has a DNS entry outside the active /56 sub-prefix. This name co-exists in IPv4 and is part of the CISCO Umbrella popularity list [10].

**T3: Silent /48.** We borrowed this /48 network, which is not separately announced in BGP, as part of a larger /29 prefix. This network neither hosts services nor active clients but is entirely silent.

**T4: Reactive /48.** This /48 is part of the same /29 covering prefix as T3. In contrast to T3, T4 answers to TCP SYN packets. T4 deploys the reactive network telescope Spoki [26], which interacts generically on the transport layer. Spoki accepts TCP connections and reacts to SYN scanning for analyzing scanner behavior. First results from Spoki can be found in [64].

### 3.2 Experiment Setup

We run Free Range Routing (FRR) [46] software on Linux to connect our autonomous system to an IXP and to upstream providers. BGP announcements for T2—T4 remain stable during our measurement period. For T1, the bi-weekly (re-)configuration is performed automatically to avoid errors. We confirm visibility of our announcements via a looking glass [59] and RIPEstat [44].

**Route6 object and ROAs.** Route(6) objects are information records that publish peering relations in the RIR database and are often used in public peering, occasionally also by upstream providers to validate that routes received from their peers are legitimate. To assess the impact of route objects on scanners, we first omitted this for our initial /32 prefix (T1). This did not impair the visibility of our prefix via the upstream. Four months after its first announcements, we created a route object for the non-split /33 prefix. Creating the route object had no noticeable effect on scanners. We did not add an RPKI ROA [34], since this would neither enhance reachability nor visibility of our prefixes given that prefixes validated as *not found* are not filtered.

**Presence in the TUM hitlist.** The TUM hitlist service [20, 70] is the most popular IPv6 address collection of active hosts. It comprises lists of responsive addresses, aliased prefixes, and non-aliased prefixes. If present on the TUM hitlist, the prefixes from our telescopes T1 to T4 could see enhanced visibility to scanners. At the beginning of our experiment, T1 was absent but T2 as well as the covering /29 prefix of T3 and T4 were already listed on the hitlist (non-aliased prefixes). We first observed the /32 prefix of T1 on the same hitlist on August 29, 2023—5 days after its announcement. We report on the presence of new prefixes from T1 on the hitlist in Section 7. It is noteworthy that T4—even though responsive from every address—never appeared on the aliased prefix list.

### 3.3 Scanner, Sources, and Sessions

We now specify terminology concerning scanners and explain our methods of identification.

**Scanner and scan sources.** A scanner may send packets from a single source, or from locally as well as globally distributed addresses. Identifying globally situated scanning entities is a complex task and we leave this for future work. Instead, we focus on localizable scan sources in this work.

A localizable *scan source* is an individual IPv6 address or an aggregation of addresses from some network. Inspecting addresses individually, *i.e.,* a /128, is the most fine-grained view. Aggregating sources based on network addresses (/64) can help to reveal scanners that change addresses within their subnet. Aside from local subnets, sources may be also aggregated on a prefix level, *e.g.,* in a /48 [54]. Aggregating sources in larger prefixes than /48 may likely combine unrelated scan sources, especially for hosting networks. Most related work [14, 36, 47, 49, 55, 69] identifies scanners by addresses (/128), followed by aggregated /64 network prefixes [47, 55]. Richter *et al.* [47] additionally evaluated /48 scan sources.

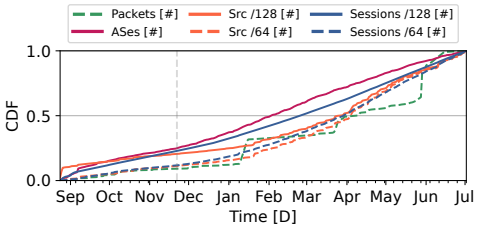


Fig. 5. Relative number of packets, ASes, and sources (/128, /64) learned during our measurements aggregated over all 4 telescopes. The vertical line indicates the end of our initial observation period.

We examine the influence of different aggregation levels on our telescopes, see Figure 5. We find divergence between the aggregation levels /128 and /64. For this reason, we proceed by analyzing both, full source addresses (/128) and /64 aggregation. We report in detail on the scan sources and their origin in Section 7.

**Scan sessions.** We define a *scan session* as a sequence of consecutive packets from a single source, in which the inter-arrival time between two subsequent packets remains below a timeout value $T$. Benson *et al.* [3] used a timeout value of 5 minutes for IPv4—the time needed to traverse the entire address space. IPv6 requires an adjustment of this value to cope with scanners traversing randomly through large subnets. Richter *et al.* [47] as well as Zhao *et al.* [69] selected one hour as session timeout to trade-off between too loose packet aggregation and too long sessions, which we adopt. Nevertheless, we do not apply any other constraints, such as a minimum number of packets or targets per session, since we want to consider all packets of all sources.

In our subsequent analyses, we will focus on scan sessions instead of packets, as sessions are more expressive and consolidate heavy hitters. Also, sessions remain relatively stable for the different source aggregation levels (/128 and /64). Similar to the scan sources, the number of scan sessions for full addresses has a more pronounced increase over /64 aggregation (shown in Figure 5). We report on the correlation between the high number of /128 scan sources and sessions in Section 7.

## 4 Overview of Data Corpus

We now introduce the data corpus collected during our experiments. After a brief overview of the traffic observed during the initial observation period, we present an aggregated view of the full observation period regarding target address types, protocol, and port usage. We use the free MaxMind GeoIP database [41] to geolocate scan sources and the freely available RouteViews dataset [62] for IPv6 address to ASN mapping.

### 4.1 Initial Observation Period

We captured about 4.6M packets during the first 12 weeks across all telescopes. These packets originate from 7.6k source addresses (3k /64 subnets) and can be aggregated into 168k (/128) or 19k (/64) sessions, respectively. We observe scan sources in 666 ASes distributed over 89 different countries.

### 4.2 Full Observation Period

In total, we captured over 51M packets across all telescopes, originating from 36k source addresses (26k /64 subnets). We can assign packets to 754k (/128) and 151k (/64) sessions, respectively. Scan sources originate from 2k ASes in 127 different countries. All CDFs grow smoothly over the whole time except for the number of packets, which discontinuously grows due to some heavy hitters sending large amounts of packets within short periods of time, see Figure 5 and Section 6 for details.

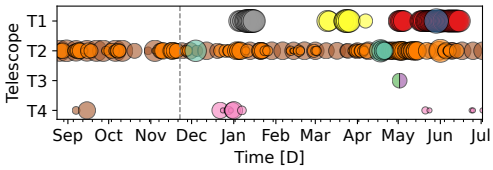


Fig. 6. Daily activity from heavy hitters at the four telescopes. Bubble sizes signify packet counts and colors represent scan sources (same color means same scan source).

**Heavy hitters.** A few individual sources stand out by contributing more than 10% of the scan packets at one telescope. We found ten such heavy hitters at the four telescopes, four in T1, three in T2, two in T3, and two in T4 (one source is identified as a heavy hitter in T2 and T4). Most heavy hitters send large amounts of packets in very few sessions during a short time span. Two heavy hitters in T2, however, send packets repeatedly over the complete observation period (see Figure 6), one of which has an RDNS entry pointing to the 6Sense scanning campaign [63]. Only three heavy hitters have an RDNS entry but we can attribute seven out of ten heavy hitters to a research context. Three of four heavy hitters in T1 are located in networks of hosting providers, one of them claims to be a 'bullet-proof' hosting provider, indicating that its scan campaign might be of malicious intent. We do not exclude heavy hitters from our analysis as they do not dominate our session-centric statistics. Even though heavy hitters account for 73% of all packets, they only contribute 0.04% of all sessions.

**Target address types.** To gain a better understanding of the selected destination addresses, we categorize addresses using the *addr6* tool of the *IPv6Toolkit* [51]. The address types are defined according to the specifications of RFC 7707 [22]. Additionally, we identify destination addresses ending with `::0` as Subnet-Router anycast addresses according to RFC 4291 [28]. We observe most targeted addresses to be randomized (64%), followed by low-byte (23%), and pattern-byte addresses (6%), see Table 3. While randomized addresses comprise most packets, they only account for 6% of the scanners. 90% of all scanners target at least one low-byte address, showing a clear trend in the scanning strategy. We analyze these results in detail in Sections 6 and 7.

Table 2. Packets, sessions, and sources per transport protocol. Shares of sessions and sources may exceed 100% due to multi-protocol scanners.

| | Packets | | Sessions /128 | | Sources /128 | |
|---|---|---|---|---|---|---|
| Protocol | [#] | [%] | [#] | [%] | [#] | [%] |
| ICMPv6 | 33,889,898 | 66.2 | 132,816 | 20.1 | 20,373 | 56.5 |
| UDP | 11,967,255 | 23.4 | 36,780 | 5.6 | 7113 | 19.7 |
| TCP | 5,372,494 | 10.5 | 614,223 | 92.8 | 19,977 | 55.4 |

**Protocol usage.** We group the packets per transport protocol. Most scanners use ICMPv6 as the primary protocol. We observe 33.9M ICMPv6 packets, 12M UDP packets, and 5.4M TCP packets. 85% of the UDP packets are DNS requests from a single scanner. While only 10.5% of the total packets are TCP, they originate from 55.4% of the scan sources and 92.8% of the sessions, implying that TCP is scanned quite often but with small amounts of packets (see Table 2). In contrast, UDP is probed in only 5.6% of the sessions but accounts for 23.4% of the packets, which shows the impact of heavy hitters on the telescope traffic.

**Target ports.** We count all targeted destination ports once per session in which they occur. We show the top 5 ports for TCP and UDP packets in Table 4. Since some scanners rotate their source interface IDs per destination port during vertical scans, we aggregate sessions by /64 subnets for

Table 3. Distribution of target types. Shares may exceed 100% due to probing multiple address types.

| | Packets | | Sources /128 | |
|---|---|---|---|---|
| Address Type | [#] | [%] | [#] | [%] |
| randomized | 32,911,060 | 64.24 | 2101 | 5.83 |
| low-byte | 11,828,733 | 23.09 | 32350 | 89.71 |
| pattern-bytes | 3,054,847 | 5.96 | 570 | 1.58 |
| embedded-ipv4 | 2,026,762 | 3.96 | 547 | 1.52 |
| subnet-anycast | 1,173,137 | 2.29 | 1476 | 4.09 |
| embedded-port | 138,656 | 0.27 | 80 | 0.22 |
| ieee-derived | 96,627 | 0.19 | 26 | 0.07 |
| isatap | 217 | ≤ 0.01 | 2 | ≤ 0.01 |

Table 4. Top 5 ports targeted by sessions on /64 source aggregation level. Shares add up to >100% as multiple ports may be addressed in one session.

| | TCP | | | UDP | | |
|---|---|---|---|---|---|---|
| Rank | Port | [#] | [%] | Port | [#] | [%] |
| #1 | 80 | 48,070 | 87.2 | Traceroute[1] | 7067 | 71.4 |
| #2 | 443 | 16,223 | 29.4 | 53 | 1945 | 19.7 |
| #3 | 21 | 2592 | 4.7 | 161 | 1718 | 17.4 |
| #4 | 8080 | 2172 | 3.9 | 500 | 1710 | 17.3 |
| #5 | 22 | 1849 | 3.4 | 123 | 1669 | 16.9 |

[1]Default port range of traceroute [33434, 33523].

this analysis. Most sessions include TCP packets with destination port 80 (HTTP, 87%) and 443 (HTTPS, 29%). The remaining top 5 ports received comparatively little attention (3%-5%). In total, 1,335 TCP ports were hit at least once. Additionally, each of the top 72 ports were targeted in at least 1k sessions, which shows that apart from port 80 and 443 some scanners cover a broad port range during their scans. A different picture emerges for UDP. 71% of all UDP sessions targeted a port from the standard traceroute range. The remaining top 5 ports belong to DNS, SNMP, ISAKMP, and NTP, which are observed in a similar number of sessions. In total, 91 UDP ports (aggregating all traceroute ports) were visited at least once.

**Key observations.** Scan traffic arrived continuously in our telescopes during the eleven months of experiments. ICMPv6 is the dominating protocol. Low-byte address scanning was the most popular strategy, but the majority of packets targeted randomized addresses. HTTP ports (TCP 80 and 443) are probed frequently while UDP packets predominantly target common traceroute and DNS ports.

## 5 Taxonomy of Scanning Behavior and Public Scan Tools

Next we introduce our classification of IPv6 scanners. We categorize *temporal behavior* of scanners, *network selection*, *i.e.,* properties of target networks, and *addresses selection*, *i.e.,* properties of the IPv6 target interface IDs. In addition, we classify IPv6 scan tools based on their signatures left in payloads and RDNS entries. These classifications are applied to scan sessions (*cf.,* Section 3.3).

### 5.1 Classifying Temporal Behavior

The temporal behavior may get influenced by internal schedules and external events. We classify scanners into three basic categories as visualized in Figure 7.

**One-off** scanners perform a single scan session and then disappear for the remaining measurement. Two edge cases are captured in this category: (*i*) scanners that perform slow, long-lasting scans and thus maintain a single scan session and (*ii*) scanners that alter source addresses repeatedly during a session, provided we can attribute the individual sessions.

**Periodic** scanners perform scan sessions in periodic intervals. The period can vary between scanners from hours to months but scanners must appear more than twice and show a stable period between scans.

**Intermittent** scanners perform multiple scan sessions but do not show a periodic pattern. Specifically, such scanners must appear with at least two scan sessions in our dataset but show no detectable period between scans.

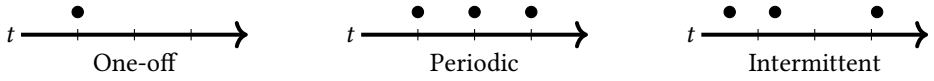


Fig. 7. Classification of temporal scanner behavior.

**Classification method.** Scanners fall into exactly one of the three exclusive categories. We identify *periodic* scanning based on a period detection by autocorrelation [7], which leaves those recurrent scanners as *intermittent* for which no period is detectable.

### 5.2 Classifying Network Selection

Scanners implement a strategy for probing IPv6 network ranges, which they can reach as soon as their prefixes are announced in BGP. We are interested in the target selection of network addresses by scanners with a particular focus on scanner behavior related to BGP prefix announcements. A scanner follows either a single-prefix or a multiple-prefix scanning strategy, and then scans the selected networks independent or dependent of their size or in a mix of both. In our setup, we can only distinguish between network-size independent and dependent scanning for multiple-prefix scanners as our view point is the context of arriving probe packets.

**Single-prefix scanning** A scanner that only scans one announced prefix during each period of announcement in BGP. The chosen (arbitrary) prefix may vary between periods.

**Network-size independent** A scanner hits networks of different sizes with (roughly) the same number of scan sessions. This behavior becomes particularly visible in our BGP experiments (T1) since all but two of the prefixes differ in size by a factor of at least two. For example, a scanner that continuously traverses the entire address space belongs in this category as each prefix will see the same number of sessions. Since our classification is based on sessions, it does not matter how many packets are sent into each prefix.

**Network-size dependent** A scanner varies the number of sessions based on the network size. In this category, we expect scanners to target smaller networks and more-specific prefixes with significantly fewer scan sessions as they contain fewer addresses. Examples are course-grained scans that more likely hit less-specific prefixes. Also, scanners may direct sessions toward fixed-size subnets. This allows scanners to probe a specific share of addresses per network and add sessions to increase the coverage of less-specific prefixes.

**Inconsistent** A scanner that changes its behavior during announcement periods shows inconsistent behavior and does not clearly fit into any of the categories above.

**Classification method.** For T2–T4, a scanner can only be classified as single-prefix scanner. For T1, we determine the network selection strategy separately for each announcement cycle, *i.e.,* the two weeks during which we announce a specific set of prefixes in T1. Our classification is based on the density-based clustering algorithm DBSCAN.

### 5.3 Classifying Address Selection

For a given network prefix, a scanner needs to select endpoint addresses (notably interface IDs) for probing. We categorize the target selection strategy of a scanner per scan session in three groups.

**Structured** A scanner selects addresses with a detectable pattern (or strong tendency) towards a specific address structure. This includes a focus on low-byte addresses, *i.e.,* choosing to probe the `::x` addresses of announced prefixes, following other well-known IPv6 address structures (*cf.,* Section 2), or algorithmically subdividing and iterating prefixes.

Table 5. Comparison of network telescopes before the split period.

(a) Sources, ASes, destination addresses, and total number of packets. T2 receives 14% more packets than T1, and is probed by 380% more /128 sources.

| | T1 | T2 | T3 | T4 |
|---|---|---|---|---|
| /128 Source addr. | 1386 | 6611 | 7 | 253 |
| /64 Source addr. | 1199 | 2113 | 6 | 251 |
| ASN | 418 | 478 | 6 | 9 |
| Destination addr. | 796,443 | 714,169 | 20 | 1817 |
| Packets | 2,161,354 | 2,464,417 | 43 | 3416 |

(b) Distinct sources per transport protocol. Percentages add to more than 100% as a single source can probe multiple protocols.

| | T1 | | T2 | | T3 | | T4 | |
|---|---|---|---|---|---|---|---|---|
| Protocol | [#] | [%] | [#] | [%] | [#] | [%] | [#] | [%] |
| ICMPv6 | 1110 | 80.1 | 4112 | 62.2 | 7 | 100 | 246 | 97.2 |
| TCP | 40 | 2.9 | 5311 | 80.3 | 0 | 0 | 6 | 2.4 |
| UDP | 266 | 19.2 | 1768 | 26.7 | 0 | 0 | 1 | 0.4 |

**Random** We examine the randomness of selected addresses in scan sessions using a statistical test suite, which differs from previous studies by Richter *et al.* [47].

**Unknown** If no detectable pattern is visible, we classify address selections as unknown.

**Classification method.** We test for address structures using the *addr6* tool of the *IPv6Toolkit* [51] and for randomness using the frequency test from the NIST test suite [2]. For statistical testing, we select sessions of at least 100 packets from our dataset, since the test requires a minimum of bits as input. The test calculates a p-value between 0 (non-random) and 1 (random). We consider scans with a p-value of at least 0.01 to use randomly generated addresses, *i.e.,* significance level of $\alpha = 0.01$. If target sets neither show structure, nor randomness we classify it as unknown.

### 5.4 Scan Tools and Systems

Probes sent by scan tools can carry tool-specific payloads. We cluster the hex-byte representation using DBSCAN, a density-based clustering approach. Then, we analyze the payload features of each cluster manually by matching them against common publicly measurement tools and systems (*e.g., RIPE Atlas probes* [48], *Yarrp* [4], and *traceroute*). To consider less common tools, we also search *GitHub* repositories and publications on measurement tools. In addition to payload analysis, we execute reverse DNS queries for each scan source to gain more detailed information on who is scanning our telescopes. We label each cluster according to our tool and DNS analysis. If no reverse DNS entry exists and no public tool is found, we categorize them using other payload characteristics *e.g.,* random bytes, or by additional patterns in the scan behavior *e.g.,* address rotation.

## 6 Scanners at the Network Telescopes During the Initial Observation Period

In this section, we analyze the network traffic that each of the four telescopes attracts during the initial observation period. We will present the results during our BGP experiment later in Section 7.

**Network traffic.** We notice in Table 5(a) that the two network telescopes with prefixes separately announced in BGP (T1 and T2) receive 4 to 6 orders of magnitude more traffic than those embedded in a covering BGP announcement (T3 and T4). The reactive network telescope T4, on the other hand, receives two orders of magnitude more traffic than the completely silent T3. These observations motivate questions about how BGP announcements impact scan traffic.

Figure 8(a) depicts longer and higher traffic peaks for T2, which emerge from scanners specifically targeting the only address in T2 for which a DNS entry exists. 50% of all observed scanners exclusively target this address, leading to twice as many /64 sources visible in T2 than in T1. In parallel, T2 scanners exhibit a unique characteristic in source address utilization. Table 5(a) shows little difference in the two source aggregation levels (/64, /128) for T1, T3, and T4, whereas T2 sees three times as many /128 scan sources as /64, because T2 attracts scanners that use address

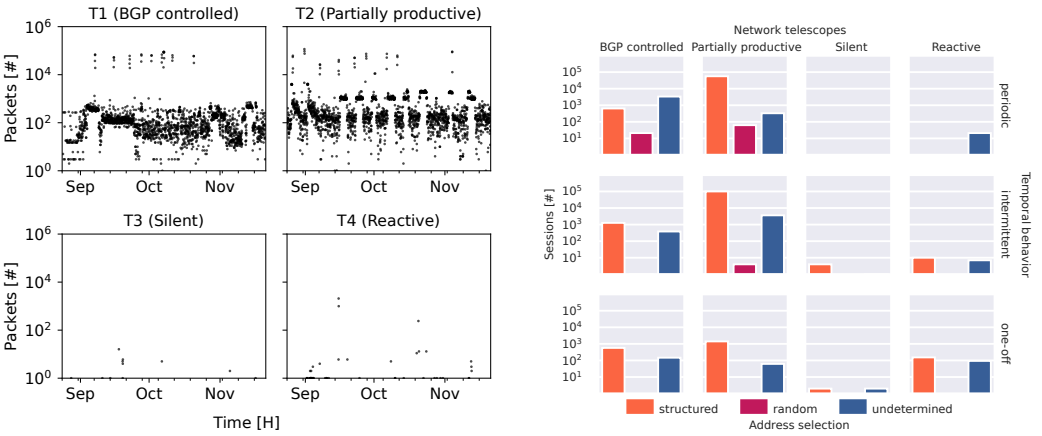


(a) Network traffic per hour across the four telescopes.

(b) Classification of scanners according to our taxonomy proposed in Section 5.

Fig. 8. Scan traffic and scanning behavior during the initial 12 weeks.

rotation within a /64 subnet. The partially productive telescope (T2) receives 14% more packets and is probed by 380% more individual sources than our BGP controlled T1. In contrast, scanners probe 12% more distinct targets in T1 compared to T2. Furthermore, the majority of individual scan sources probes ICMPv6 for T1 (80%), T3 (100%), and T4 (97%). For T2, however, the prevalent probed protocol is TCP (80%), followed by ICMPv6 (62%), as shown in Table 5(b). The differences are mainly caused by the DNS attractor and scanners rotating their source addresses within a /64 subnet, which we exclusively observe for T2.

**Behavior of scanners.** Figure 10 displays the evolution of session intensities per telescope and week. Results are rather stable for T1 and T2, while sporadic for T3 and T4. The single October peak for T4 is due to a single scanning campaign.

We next categorize scanners and their sessions according to our taxonomy in Section 5 w.r.t. (*i*) temporal behavior and (*ii*) address selection. Figure 8(b) summarizes the classification for all telescopes. For each telescope (columns), we separate sessions by the temporal behavior of the scanner (rows), and plot bars representing the number of sessions per address selection strategy. Most scanners return (intermittent: 41% or periodic: 29%) and follow a *structured* selection strategy. T1 and T2 appear similar in their overall distribution. In contrast to T4, scan sessions from *one-off* scanners are relatively less frequent in T1 and T2 where *intermittent* and *periodic* are more common. For T3 and T4, which have fewer sessions overall, *structured* address selection is the only identifiable strategy. For T3 and T4, none of the sessions are classified as *random*.

**Correlation between the four telescopes.** The initial observation period allows for a brief comparison between all telescopes before exploring the effect of more specific BGP announcements in the following. At T4, only a small fraction of ASNs is exclusively observed, most of them overlapping with T1-T3. At T3, all source ASNs can also be observed in T1, T2, and T4. Around half of all ASNs in T1 and T2 overlap, which shows a partial similarity between these telescopes. Figure 9(a) lays out detailed ASN observations in an UpSet plot. The bar graph on the left shows the non-exclusive shares observed in each telescope. In contrast, the top bars show the exclusive combinations of telescopes that observed a given share of ASNs.

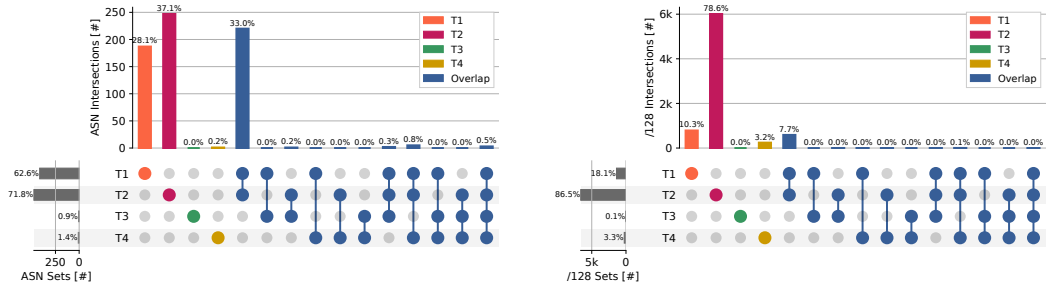


(a) ASN (b) /128

Fig. 9. Intersections of originating autonomous systems and scan sources (/128) between all telescopes.

A detailed analysis of the /128 scan sources, presented in Figure 9(b), reveals that a dominant fraction (≈ 90%) of the sources exclusively scans a single telescope. We infer that these differently configured telescopes attract different scanners. Even if we exclude all sources that exclusively target the address with a DNS name, diverse traffic remains at all telescopes.

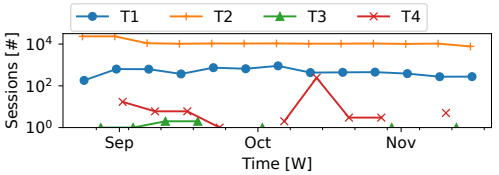


Fig. 10. Weekly scan sessions at the four telescopes.

**Key observations.** Telescopes that are subnets of covering prefixes in BGP receive significantly less traffic than telescopes that coincide with in BGP announced prefixes. Sessions are a stable measure of scanner behavior, but behavior appears to vary with the kind of scanner attractors. Most scanners follow a structured address selection strategy and return. Heavy hitters are either research scanners or malicious.

## 7 Adaption of Scanners to BGP Signals

In this section, we report on the detailed results of our BGP experiment. We discuss how traffic is attracted by prefix advertisements and how scanners react to BGP signals. We also compare the activities in our BGP controlled telescope (T1) against those in the other telescopes.

### 7.1 BGP-controlled Telescope Activity and Impact on the Behavior of Scanners

**Attracting traffic.** During our BGP experiment, T1 observed a weekly increase in the average number of observed scan sources by 275% and 555% in the average number of scan sessions. Activities remain stable in all other telescopes (see Figure 12). Announcements of the two new more-specific prefixes attract significantly more traffic than to the single announcement of the companion prefix. Figure 11 illustrates the growth in sessions per most-specific prefix. As long as subnets remain silent within a larger covering prefix, they hardly attract any attention, *e.g.*, /48 subnets receive only 0.4% of the total number of sessions during the first two weeks of the experiment. Only after subnets transform to prefixes, they receive significant traffic. The overall number of packets arriving in the iteratively split /33 segment exceeds packet counts for the stable companion /33 by +286%. For the final announcement period, the /48 prefixes observed 15.7% of all sessions (*i.e.*, increased by 39×).

**Network selection of scanners.** We observe from Table 6 that 90% of all scan sources probe addresses from a single prefix, 9% scan independent of network size, less than one percent show inconsistent or network-size dependent behavior. Single-prefix scanners originate mainly from

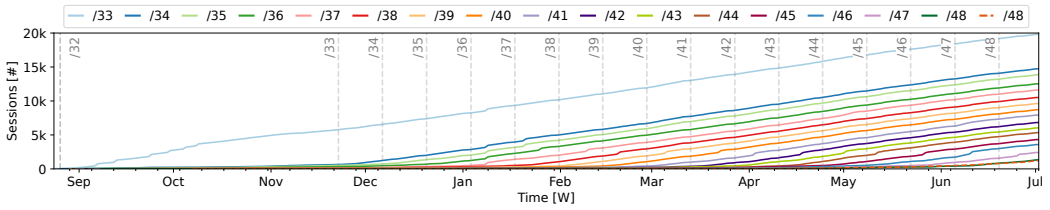


Fig. 11. Cumulative number of scan sessions per target prefix. Addresses receive significantly more attention by scanners when more specific prefixes are announced in BGP.

two entities (*i*) RIPE Atlas (57%) topology measurements, and (*ii*) Alpha Strike Labs [21] (36%), a cybersecurity company. These scanners account for 20% of the sessions.

With more than one thousand individual scan sources and 31% of the total sessions, a large portion selects target networks independent of the actual prefix size, *i.e.,* during an announcement period they cover each of our prefixes with a roughly equal number of scan sessions. Since the sizes of our BGP prefixes differ largely these scanners either aim at targeting every prefix during each session or they scan the BGP space with fine granularity, hitting our prefixes as a side effect of their strategy. We observe inconsistent behavior for only 64 scan sources, but these execute almost 50% of all sessions. In most cases, inconsistent behavior tends to become network-size independent towards the end of the experiment. In the beginning, however, these scanners probe the larger prefixes more often.

Only 24 scan sources (0.2%) probe networks depending on their size, *i.e.,* larger prefixes receive more traffic. A minimal telescope setup with a single /48 announcement would miss those who exclusively probe larger prefixes, such as /32. In contrast, we expect scanners that select their target networks independent of size to be observable in all network telescopes with a prefix visible in BGP. Even smaller telescopes should receive a notable amount of their scan traffic.

**Address selection of scanners** classifies into the categories (*i*) structured, (*ii*) random, and (*iii*) unknown according to our taxonomy (*cf.,* Section 5). Figure 13 illustrates samples of a structured and a random selection session. For each scan session, we show all targeted addresses in hexadecimal representation (y-axis) from low nibble (top) to high nibble (bottom) and rank them by time of arrival. Our telescope prefix is concealed in gray. Figure 13(a) illustrates a scan session by AS132203 with 151k packets. While most nibbles across the figure are zero, periods with frequently changing characters are visible as stripes. By sorting the addresses lexicographically, Figure 14 visualizes a clear structure of the traversal through the network. Although blocks of zeros become more pronounced, iterative traversal is now visible in stripes of shifting color (from `0` to `f`). After roughly two thirds of the packets the pattern shifts and shows a tree like structure from traversing into the 'leaves' of some subnets. Figure 13(b) depicts a scan session by AS53667 with 113k packets. Nibbles 11 and 12 show a structured iteration through our subnets, but there is no pattern or structure apparent in the remaining segments. This observation suggests a random generation of the last 80 bits of the

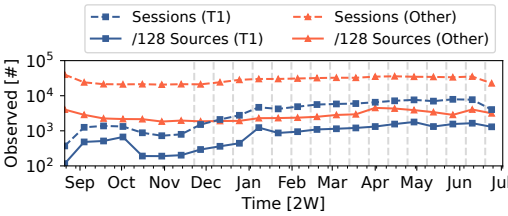


Fig. 12. Bi-weekly evolution of sessions and sources in the BGP controlled (T1) and the aggregated remaining telescopes. Different from the other telescopes, T1 shows an increase in sessions and sources during the experiment. The vertical lines indicate the prefix splits.

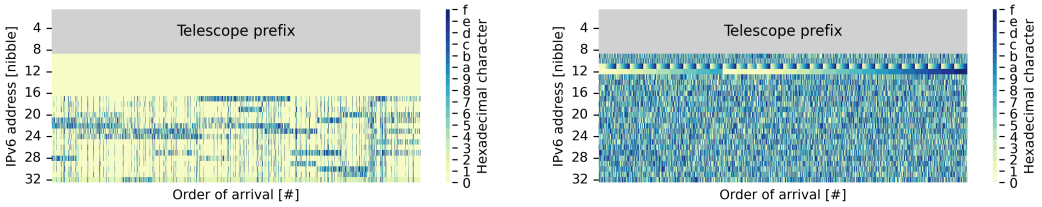


(a) Structured (b) Random

Fig. 13. Structured and randomized target address generation observed in two scan sessions. Sequentially targeted destination addresses are represented in hex using a different color for each digit value.

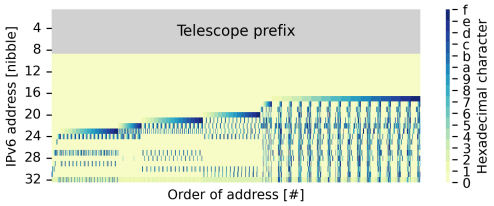


Fig. 14. Addresses from Fig. 13(a) numerically sorted.

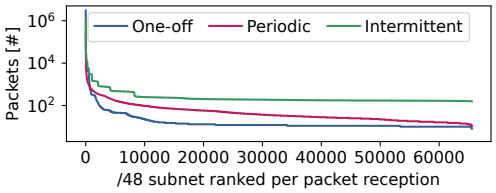


Fig. 15. Packets per scanner type across /48 subnets.

target addresses. Appendix C presents a randomness assessment of the IID and subnet parts, using the test suite from the National Institute of Standards and Technology (NIST) [2].

The prevalent probing strategy within scan sessions is for structured addresses (see Figure 16). Still, many sessions randomly traverse the address space, especially those from periodic scanners. While random probing can be useful to detect aliased prefixes, it is commonly used in topology measurements to detect routers [4, 6, 29, 35, 50, 65] as it is likely to hit unassigned addresses or networks, which triggers ICMP error messages from on-path routers and reveals network topology.

**Temporal behavior of scanners.** Our classification of temporal behavior as summarized in Table 6 shows that the vast majority of scanners appear only once (69.7%), which may relate to our limited observation period between the biweekly announcements. Intermittent and periodic scanners remain with a similar share around 15%. Scanners of periodic behavior (14.8%) comprise the majority of sessions (72.8%) due to short periods. One-off sessions contain fewer packets and are rather short, *i.e.,* always less than 4 hours (top 10 $\geq$ 1 hour each). These scanners may act very coarse-grained or use a distributed scan infrastructure, which requires little work from individual sources. The latter could show up as comparatively many associated scanning source addresses with fewer packets each. In our dataset, we detect such behavior from RIPE Atlas probes and Alpha Strike Labs. In contrast, scan sessions from periodic and intermittent scanners tend to be much longer. The longest session from a periodic scanner spans over 311 hours (top 15 $\geq$ 78 hours each); for intermittent, the longest session spans 316 hours (top 15 $\geq$ 6 hours each).

Beyond session schedules, the coverage of the /48 subnets in our /32 telescope prefix is of interest. We observe two kinds of behavior, (*i*) probing a wide range of /48 subnets with only a few packets each, or (*ii*) focusing on a few subnets and probing these in more depth. Figure 15 ranks all subnets from highest to lowest packet reception. Intermittent scanners probe the majority of subnets more evenly, while one-off scanners focus on a few selected subnets. Periodic scanners probe a wider

Table 6. Classification according to our taxonomy in Section 5, combining internal and external schedules during the split period.

| | Scanners | | Sessions | |
|---|---|---|---|---|
| Classification | [#] | [%] | [#] | [%] |
| *Temporal behavior* | | | | |
| One-off | 8244 | 69.71 | 8244 | 8.95 |
| Intermittent | 1832 | 15.49 | 16,842 | 18.28 |
| Periodic | 1750 | 14.80 | 67,067 | 72.78 |
| *Network selection* | | | | |
| Single-prefix scanning | 10,703 | 90.50 | 17,939 | 19.47 |
| Network-size independent | 1035 | 8.75 | 28,433 | 30.85 |
| Inconsistent behavior | 64 | 0.55 | 44,294 | 48.07 |
| Network-size dependent | 24 | 0.20 | 1487 | 1.61 |


Fig. 16. Classification of scanners observed in T1 during the split period according to our taxonomy.

range of subnets with more packets but tend to visit subnets selectively. Periodic and intermittent scanners focus on the `0000` subnet while one-off scanners probe the `e000` subnet the most.

**Cross-category behavior.** All three temporal behaviors show specific patterns w.r.t. the target network and address selection. One-off sessions are highly biased towards single-prefix scans (95%) and often follow a structured target selection approach. In contrast, sessions of intermittent scanners predominantly show inconsistent behavior (50%) followed by single-prefix scanning (32%). While we also observe most periodic sessions as inconsistent (54%), 39% scan independent of network-size. Note that in terms of individual scan sources we observe inconsistent behavior for only 13 intermittent (0.7%), and 51 periodic scanners (2.9%). The results of address selection, temporal behavior, and network target selection provide different perspectives on the observed scanning behavior. Our experiment clearly reveals a trend towards structured, BGP-aware scanning.

### 7.2 Scan Tools and Scanning Sources

We analyze all scan sources and packet payloads (if present) for patterns that allow for a better understanding of the scanners that visit our telescope, including the tools they use.

**Public tools used by scan sources.** 17M (40%) of 43M captured packets contain a payload. These packets are sent by 11,001 (93%) scan sources and cover $\approx 70K$ (76%) of all sessions. We analyze the payloads for fingerprints and map them to public tools. We were able to identify eight public tools (see Table 7), mostly traceroute-type implementations for IPv6 topology and periphery measurements. Interestingly, we observe scanners using the *Htrace6* tool in December 2023, even though the code was first published in late January 2024. The traceroute-type tool *Yarrp6* is used by 22 distinct scan sources, all appear periodically. It is the only open source tool which we observe regularly over the complete observation period. A large fraction of all scan sources (55%) are RIPE Atlas probes,

Table 7. Overview of identified scan tools. RIPE Atlas probes account for 55% of all sources we observe at T1.

| | Scanners | | Sessions | |
|---|---|---|---|---|
| Scan Tool | [#] | [%] | [#] | [%] |
| RIPEAtlasProbe [45] | 6483 | 54.82 | 11,859 | 12.87 |
| Yarrp6 [5] | 22 | 0.19 | 562 | 0.61 |
| Traceroute [32] | 19 | 0.16 | 163 | 0.18 |
| Htrace6 [24] | 9 | 0.08 | 16 | 0.02 |
| 6Seeks [58] | 5 | 0.04 | 17 | 0.02 |
| 6Scan [23] | 3 | 0.03 | 17 | 0.02 |
| CAIDA Ark [8] | 2 | 0.02 | 2019 | 2.19 |

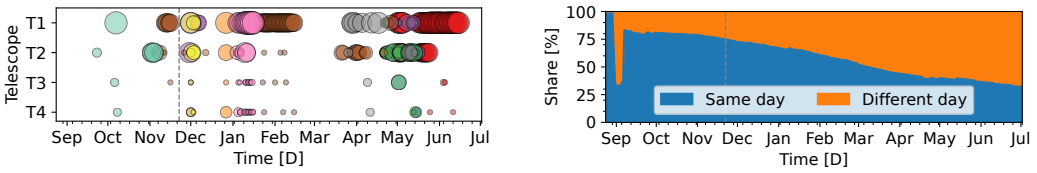


(a) Daily activity from scan sources across all four telescopes. Bubble sizes signify packet counts and colors represent scan sources (same color means same scan source).

(b) Cumulative share of sources observed in T1 and T2 on the same or different days.

Fig. 17. Analysis of overlapping sources. The dashed line marks the start of our active experiment.

almost exclusively identified as one-off scanners. They always target the ::1 addresses in each prefix.

**Context of scan sources.** While the majority of scan sources (96%) originate from hosting or ISP networks (see Table 8), most of them can be associated with a research context – 97% of the sources in ISP networks and 22% in the hosting networks belong to RIPE Atlas probes. 58% of sources in the hosting category belong to a single company (Alpha Strike Labs) that commercializes research scanning. Heavy hitters contribute a significant portion of all packets, but only account for a negligible fraction of sessions (*cf.,* §4). Table 8 shows that three out of four heavy hitters are located in hosting networks, without research context, thus indicating malicious activity.

Our findings emphasize that research scanners probe the IPv6 address space more often and regularly, making them a good reference point for telescope setup as they predominantly act BGP-aware. Furthermore, we find 18 scan sources to live-monitor BGP announcements, since we reliably observe them within 30 minutes after a new BGP announcement.

**Source overlap across telescopes.** Only one /128 scan source, originating from a hosting network, probed all four telescopes during the initial 12 weeks, at the end of September and early October. These probes matched the signature of "Yarrp6". The same scan source returned to T2 in November, but with a different signature. Either the same entity ran a new scan campaign or a different entity incidentally used the same IP address.

Table 8. Network types of scan sources. Most sessions (76%) and sources (96%) originate from hosting and ISP networks. Three out of four heavy hitters (Hit.) are located in hosting networks.

| | Scanners | | Sessions | | Packets | |
|---|---|---|---|---|---|---|
| Network | [#] | [%] | [#] | [%] | [#] | [%] |
| Hosting | 6624 | 56.01 | 23,682 | 25.70 | 28,371,475 | 65.06 |
| w/o Hit. | 6621 | 55.99 | 23,674 | 25.69 | 4,496,454 | 10.31 |
| ISP | 4681 | 39.58 | 46,864 | 50.85 | 1,478,591 | 3.39 |
| Education | 245 | 2.07 | 17,634 | 19.14 | 13,629,270 | 31.25 |
| w/o Hit. | 244 | 2.06 | 17,627 | 19.13 | 4,375,030 | 10.03 |
| Business | 194 | 1.64 | 2259 | 2.45 | 71,689 | 0.16 |
| Government | 6 | 0.05 | 7 | 0.01 | 96 | 0.00 |
| Unknown | 76 | 0.64 | 1707 | 1.85 | 58,527 | 0.13 |

Nine other /128 scan sources were observed at every telescope during the entire measurement period. Figure 17(a) visualizes the daily activities, marking each scan source with a distinct color. Larger markers signify more packets on a given day. For each source, T1 and T2 together received about 98% of the packets, although in most cases one of them was focused (≥90%). While most scan campaigns were localized in time, some scan sources showed up repeatedly across a subset of the telescopes. As an example, in mid-November one source was observed in T1, T2, and T3, but not in T4. Overall, six of these ten scan sources belong to four different hosters. The remaining four scan sources are located in three research networks. Two of them were both active in May, largely overlapping in time, and were likely part of the same scan campaign.

Limiting the analysis to our separately announced prefixes (T1 and T2) reveals an overlap of 2,169 addresses in total or 598 addresses during the initial period. Figure 17(b) shows the cumulative share of scan sources observed on the same day in blue and on different days in orange. After few addresses were observed on the first days, the dip in September followed by a sudden rise was caused by 254 new sources scanning on the same day. Roughly 75% of scan sources were observed in both telescopes on the same day during the initial period. Attracting scanners by our active experiment only in T1 leads to a drifting apart between the telescopes as the share declines to approximately 30% over the course of our active experiment.

**Correlation with the TUM hitlist.** While our prefixes appear on the TUM hitlist within a few days after the announcement, there is no noticeable impact on the number of sessions or traffic, implying that the prefixes on the TUM hitlist are rarely used by BGP-reactive scanners.

**Key observations.** Announcing more specific prefixes in BGP attracts significantly more traffic (+286%) and scanners (+275%), which in contrast to stable prefixes prefer one-time visits (69.7%) and single-prefix scans. A large portion of scanners probes prefixes independent of its size and applies a structured target selection. The majority of scanners can be tied to its origin or a public tool.

## 8 Discussion, Conclusion, Outlook

In this paper, we explored the behavior of IPv6 scanners from the perspective of four network telescopes with contrasting properties. Our telescopes were passive, traceable, or (re-)active. To specifically study scanners that are triggered by BGP signals, we introduced and utilized a controlled, active measurement method that relies on IPv6 prefixes of varying sizes announced proactively in BGP.

**What should an IPv6 telescope operator consider?** Our findings include the following practical implications for operating network telescopes. (*i*) Visibility of a network largely increases if its prefix is individually announced in BGP instead of being a subnet that is only part of a covering prefix; (*ii*) the size of an IPv6 prefix is of lower relevance for a network telescope than the number of individually announced prefixes; (*iii*) different attractors, *e.g.,* DNS versus BGP, draw different kinds of scanners; (*iv*) active network services draw scanners to neighboring address space; (*v*) structured target addresses are preferred by many scanners.

**Are observations in telescopes unbiased?** No. Scanners seem to contact telescopes following external triggers from the network or the application side, which in turn means that triggers attract only those scanners that react to them. We are aware of the different specific biases related to our network telescopes and emphasize our converse finding: We measure the effects of network triggers and show, how and which scanners react to them, *i.e.,* we quantify the biasing factors.

**Are IPv6 telescopes suitable to monitor DDoS?** No. Telescopes commonly monitor DDoS by capturing the backscatter from randomly spoofed attack traffic. It is very unlikely to capture packets with randomly selected IPv6 destination addresses in a telescope. Researchers and security experts will need to identify new ways to assess DDoS—without relying on IPv6 background radiation.

**What next?** This work gives rise to the following research directions. (*i*) Future measurements and analyses shall quantify the effect of further triggers that attract traffic to IPv6 network telescopes. (*ii*) With various, compatible trigger measurements at hand, a correlation analysis should enable a more realistic, quantitative assessment of the biases inherited from IPv6 for network telescopes. (*iii*) New methods for Internet observatories are needed to capture IPv6 background radiation.

## Acknowledgments

This work was partly supported by the Federal Ministry of Research, Technology and Space (BMFTR) within the projects IPv6Explorer (16KIS1815) and AI.Auto-Immune (16KIS2332K and 16KIS2333).

## References


[1] Mark Allman, Vern Paxson, and Jeff Terrell. 2007. A Brief History of Scanning. In *Proc. of the ACM IMC*. ACM, New York, NY, USA, 77–82.

[2] Lawrence Bassham, Andrew Rukhin, Juan Soto, James Nechvatal, Miles Smid, Elaine Barker, Stefan Leigh, Mark Levenson, Mark Vangel, David Banks, Nathanael Alan Heckert, James Dray, and San Vo. 2010. *A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications*. Special Publication NIST SP 800-22. National Institute of Standards & Technology, Gaithersburg, MD, US.

[3] Karyn Benson, Alberto Dainotti, kc claffy, Alex C. Snoeren, and Michael Kallitsis. 2015. Leveraging Internet Background Radiation for Opportunistic Network Analysis. In *Proceedings of the 2015 Internet Measurement Conference* (Tokyo, Japan) *(IMC '15)*. ACM, New York, NY, USA, 423–436.

[4] Robert Beverly. 2016. Yarrp'ing the Internet: Randomized High-Speed Active Topology Discovery. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 413–420. https://doi.org/10.1145/2987443.2987479

[5] Robert Beverly. 2025. YARRP. https://github.com/cmand/yarrp

[6] Robert Beverly, Ramakrishnan Durairajan, David Plonka, and Justin P. Rohrer. 2018. In the IP of the Beholder: Strategies for Active IPv6 Topology Discovery. In *Proc. of ACM IMC (IMC '18)*. ACM, New York, NY, USA, 308–321. https://doi.org/10.1145/3278532.3278559

[7] Tim Breitenbach, Bartosz Wilkusz, Lauritz Rasbach, and Patrick Jahnke. 2023. On a Method for Detecting Periods and Repeating Patterns in Time Series Data with Autocorrelation and Function Approximation. *Pattern Recognition* 138 (2023), 1–22.

[8] CAIDA. 2007. Archipelago Project. https://www.caida.org/projects/ark/

[9] Xue Chen, Weiwei Shi, Jieyan Liu, Mengshu Hou, and Yujun Li. 2023. 6Community: An Active IPv6 Address Detection Method Based On Community Discovery Algorithm. In *ADMIT 2023*. ACM, New York, NY, USA, 114–119. https://doi.org/10.1145/3625403.3625426

[10] Cisco. 2025. Cisco Umbrella Popularity List. https://s3-us-west-1.amazonaws.com/umbrella-static/index.html

[11] Tianyu Cui, Gaopeng Gou, and Gang Xiong. 2020. 6GCVAE: Gated Convolutional Variational Autoencoder For IPv6 Target Generation. In *2020 24th PAKDD*. Springer, Springer International Publishing, Cham, 609–622.

[12] Tianyu Cui, Gaopeng Gou, Gang Xiong, Chang Liu, Peipei Fu, and Zhen Li. 2021. 6GAN: IPv6 Multi-Pattern Target Generation Via Generative Adversarial Nets With Reinforcement Learning. In *IEEE INFOCOM 2021*. IEEE, IEEE, Piscataway, NJ, USA, 1–10.

[13] Tianyu Cui, Gang Xiong, Gaopeng Gou, Junzheng Shi, and Wei Xia. 2020. 6VECLM: Language Modeling In Vector Space For IPv6 Target Generation. In *ECML PKDD 2020*. Springer, Springer International Publishing, Cham, 192–207.

[14] Jakub Czyz, Kyle Lady, Sam G. Miller, Michael Bailey, Michael Kallitsis, and Manish Karir. 2013. Understanding IPv6 internet background radiation. In *Proc. of the ACM IMC* (Barcelona, Spain). ACM, New York, NY, USA, 105–118. https://doi.org/10.1145/2504730.2504732

[15] Zakir Durumeric, Michael Bailey, and J. Alex Halderman. 2014. An Internet-Wide View of Internet-Wide Scanning. In *Proc. of the 23rd USENIX Conference on Security Symposium* (San Diego, CA). USENIX Assoc., Berkeley, CA, USA, 65–78.

[16] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-Wide Scanning and its Security Applications. In *Proc. of the 22nd USENIX Security Symposium*. USENIX Assoc., Berkeley, CA, USA, 605–620.

[17] Mat Ford, J. Stevens, and John Ronan. 2006. Initial Results from an IPv6 Darknet. In *Proc. of the ICISP*. IEEE, Piscataway, NJ, USA, 13–13. https://doi.org/10.1109/ICISP.2006.14

[18] Pawel Foremski, David Plonka, and Arthur Berger. 2016. Entropy/IP: Uncovering Structure in IPv6 Addresses. In *Proc. of the ACM IMC* (Santa Monica, California, USA) *(IMC '16)*. ACM, New York, NY, USA, 167–181. https://doi.org/10.1145/2987443.2987445

[19] Kensuke Fukuda and John Heidemann. 2018. Who Knocks at the IPv6 Door? Detecting IPv6 Scanning. In *Proc. of IMC* (Boston, MA, USA) *(IMC '18)*. ACM, New York, NY, USA, 231–237. https://doi.org/10.1145/3278532.3278553

[20] Oliver Gasser, Quirin Scheitle, Sebastian Gebhard, and Georg Carle. 2016. Scanning the IPv6 Internet: Towards a Comprehensive Hitlist. In *Proc. of TMA* (Louvain La Neuve, Belgium). IFIP, Laxenburg, MD, Austria, 1–8.

[21] Alpha Strike Labs GmbH. 2017. Alpha Strike Labs. website. https://www.alphastrike.io/en/

[22] F. Gont and T. Chown. 2016. *Network Reconnaissance in IPv6 Networks*. RFC 7707. IETF. https://doi.org/10.17487/RFC7707

[23] hbn1987. 2025. 6Scan. https://github.com/hbn1987/6Scan

[24] hbn1987. 2025. Htrace6. https://github.com/hbn1987/6Scan/tree/master/Htrace6

[25] Raphael Hiesgen, Marcin Nawrocki, Marinho Barcellos, Daniel Kopp, Oliver Hohlfeld, Echo Chan, Roland Dobbins, Christian Doerr, Christian Rossow, Daniel R. Thomas, Mattijs Jonker, Ricky Mok, Xiapu Luo, John Kristoff, Thomas C. Schmidt, Matthias Wählisch, and KC Claffy. 2024. The Age of DDoScovery: An Empirical Comparison of Industry and Academic DDoS Assessments. In *Proc. of ACM Internet Measurement Conference (IMC)*. ACM, New York, 259–279. https://doi.org/10.1145/3646547.3688451

[26] Raphael Hiesgen, Marcin Nawrocki, Alistair King, Alberto Dainotti, Thomas C. Schmidt, and Matthias Wählisch. 2022. Spoki: Unveiling a New Wave of Scanners through a Reactive Network Telescope. In *Proc. of 31st USENIX Security Symposium*. USENIX Association, Berkeley, CA, USA, 431–448. https://www.usenix.org/system/files/sec22-hiesgen.pdf

[27] Raphael Hiesgen, Marcin Nawrocki, Thomas C. Schmidt, and Matthias Wählisch. 2024. The Log4j Incident: A Comprehensive Measurement Study of a Critical Vulnerability. *IEEE Transactions on Network and Service Management (TNSM)* 21, 6 (December 2024), 5921–5934. https://doi.org/10.1109/TNSM.2024.3440188

[28] R. Hinden and S. Deering. 2006. *IP Version 6 Addressing Architecture*. RFC 4291. IETF. https://doi.org/10.17487/RFC4291

[29] Florian Holzbauer, Markus Maier, and Johanna Ullrich. 2024. Destination Reachable: What ICMPv6 Error Messages Reveal About Their Sources. In *Proc. of ACM IMC* (Madrid, Spain). ACM, New York, NY, USA, 1–15. https://doi.org/10.1145/3646547.3688420

[30] Bingnan Hou, Zhiping Cai, Kui Wu, Jinshu Su, and Yinqiao Xiong. 2021. 6Hit: A Reinforcement Learning-Based Approach To Target Generation For Internet-Wide IPv6 Scanning. In *IEEE INFOCOM 2021*. IEEE, IEEE, Piscataway, NJ, USA, 1–10.

[31] Bingnan Hou, Zhiping Cai, Kui Wu, Tao Yang, and Tongqing Zhou. 2023. 6Scan: A High-Efficiency Dynamic Internet-Wide IPv6 Scanner With Regional Encoding. *IEEE/ACM Transactions on Networking* 31, 4 (2023), 1870–1885.

[32] Van Jacobson. 1987. traceroute. https://linux.die.net/man/8/traceroute6

[33] Jonas Kaspereit, Gurur Öndarö, Gustavo Luvizotto Cesar, Simon Ebbers, Fabian Ising, Christoph Saatjohann, Mattijs Jonker, Ralph Holz, and Sebastian Schinzel. 2024. LanDscAPe: Exploring LDAP Weaknesses and Data Leaks at Internet Scale. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Philadelphia, PA, 1225–1242. https://www.usenix.org/conference/usenixsecurity24/presentation/kaspereit

[34] M. Lepinski and S. Kent. 2012. *An Infrastructure to Support Secure Internet Routing*. RFC 6480. IETF. https://doi.org/10.17487/RFC6480

[35] Xiang Li, Baojun Liu, Xiaofeng Zheng, Haixin Duan, Qi Li, and Youjun Huang. 2021. Fast IPv6 Network Periphery Discovery And Security Implications. In *IEEE DSN 2021*. IEEE, IEEE, Washington, DC, USA, 88–100.

[36] ChenHuan Liu, ShanShan Hao, QianKun Liu, CongXiao Bao, and Xing Li. 2021. IPv6-Network Telescope Network Traffic Overview. In *2021 IEEE 11th ICEIEC*. IEEE, Piscataway, NJ, USA, 1–4. https://doi.org/10.1109/ICEIEC51955.2021.9463724

[37] Ning Liu, Chunbo Jia, Bingnan Hou, Changsheng Hou, Yingwen Chen, and Zhiping Cai. 2023. 6Search: A Reinforcement Learning-Based Traceroute Approach For Efficient IPv6 Topology Discovery. *Computer Networks* 235 (2023), 1–10.

[38] Qiankun Liu and Xing Li. 2023. 6Former: Transformer-Based IPv6 Address Generation. In *ISCC 2023*. IEEE, IEEE, Piscataway, NJ, USA, 1142–1148.

[39] Zhizhu Liu, Yinqiao Xiong, Xin Liu, Wei Xie, and Peidong Zhu. 2019. 6Tree: Efficient Dynamic Discovery Of Active Addresses In The IPv6 Address Space. *Computer Networks* 155 (2019), 31–46.

[40] Alexander Männel, Jonas Mücke, kc Claffy, Max Gao, Ricky K. P. Mok, Marcin Nawrocki, Thomas C. Schmidt, and Matthias Wählisch. 2025. Lessons Learned from Operating a Large Network Telescope. In *Proc. of ACM Special Interest Group on Data Communication (SIGCOMM)*. ACM, New York.

[41] MaxMind, Inc. 2025. MaxMind – GeoLite Country. https://dev.maxmind.com/geoip/geolite2-free-geolocation-data

[42] T. Narten, R. Draves, and S. Krishnan. 2007. *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*. RFC 4941. IETF. https://doi.org/10.17487/RFC4941

[43] Marcin Nawrocki, John Kristoff, Chris Kanich, Raphael Hiesgen, Thomas C. Schmidt, and Matthias Wählisch. 2023. SoK: A Data-driven View on Methods to Detect Reflective Amplification DDoS Attacks Using Honeypots. In *Proc. of IEEE Euro Security & Privacy* (Delft, Netherlands). IEEE, 576–591. https://doi.org/10.1109/EuroSP57164.2023.00041

[44] RIPE NCC. 2024. RIPEstat. https://stat.ripe.net/ui2013/

[45] RIPE NCC. 2025. RIPE Atlas software probe. https://github.com/RIPE-NCC/ripe-atlas-software-probe

[46] FRRouting Project. 2025. FRR. https://frrouting.org

[47] Philipp Richter, Oliver Gasser, and Arthur Berger. 2022. Illuminating Large-Scale IPv6 Scanning in the Internet. In *Proc. of the ACM IMC*. ACM, New York, NY, USA, 410–418. https://doi.org/10.1145/3517745.3561452

[48] RIPE NCC. 2010. What is RIPE Atlas? https://atlas.ripe.net/about/

[49] John Ronan and David Malone. 2023. Revisiting and Revamping an IPv6 Network Telescope. In *2023 34th ISSC*. IEEE, Piscataway, NJ, USA, 1–6. https://doi.org/10.1109/ISSC59246.2023.10162033

[50] Erik C Rye and Robert Beverly. 2020. Discovering The IPv6 Network Periphery. In *Proc. of PAM Conf.* Springer, Springer, Berlin Heidelberg, 3–18.

[51] SI6 Networks. 2016. IPv6 Toolkit. https://www.si6networks.com/research/tools/ipv6toolkit/.

[52] Guanglei Song, Jiahai Yang, Zhiliang Wang, Lin He, Jinlei Lin, Long Pan, Chenxin Duan, and Xiaowen Quan. 2022. DET: Enabling Efficient Probing Of IPv6 Active Addresses. *IEEE/ACM Transactions on Networking* 30, 4 (2022), 1629–1643.

[53] Lion Steger, Liming Kuang, Johannes Zirngibl, Georg Carle, and Oliver Gasser. 2023. Target Acquired? Evaluating Target Generation Algorithms for IPv6. In *Proc. of TMA*. IEEE, Piscataway, NJ, USA, 1–10.

[54] Stephen Strowes. 2019. *Visibility of IPv4 and IPv6 Prefix Lengths in 2019*. RIPE Labs article. RIPE. https://labs.ripe.net/author/stephen_strowes/visibility-of-ipv4-and-ipv6-prefix-lengths-in-2019/
[55] Stephen D Strowes, René Wilhelm, Florian Obser, Riccardo Stagni, Agustín Formoso, and Emile Aben. 2020. Debogonising 2a10::/12 Analysis of one week's visibility of a new /12. In *Proc. of TMA*. IFIP, Laxenburg, MD, Austria, 1–9.
[56] Hammas 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. *Proceedings of the ACM on Networking* 3, CoNEXT3 (September 2025). https://doi.org/10.1145/3749221
[57] Hammas Bin Tanveer, Rachee Singh, Paul Pearce, and Rishab Nithyanand. 2023. Glowing in the Dark: Uncovering IPv6 Address Discovery and Scanning Strategies in the Wild. In *Proc. of the USENIX Security Symposium*. USENIX Association, Anaheim, CA, 6221–6237. https://www.usenix.org/conference/usenixsecurity23/presentation/bin-tanveer
[58] Yang Tao. 2025. 6Seeks. https://github.com/6Seeks/6Seeks
[59] Telia. 2024. Telia Looking Glass. https://lg.telia.net/
[60] S. Thomson, T. Narten, and T. Jinmei. 2007. *IPv6 Stateless Address Autoconfiguration*. RFC 4862. IETF. https://doi.org/10.17487/RFC4862
[61] Johanna Ullrich, Peter Kieseberg, Katharina Krombholz, and Edgar Weippl. 2015. On Reconnaissance with IPv6: A Pattern-Based Scanning Approach. In *10th International Conf on Availability, Reliability and Security*. IEEE, Piscataway, NJ, USA, 186–192.
[62] University of Oregon. 2017. Route Views Project. http://www.routeviews.org/.
[63] Grant Williams, Mert Erdemir, Amanda Hsu, Shraddha Bhat, Abhishek Bhaskar, Frank Li, and Paul Pearce. 2024. 6Sense: Internet-Wide IPv6 Scanning and its Security Applications. In *Proc. of the 33rd USENIX Security Symposium*. USENIX Association, San Diego, CA, USA.
[64] Yue Xin, Maynard Koch, Isabell Egloff, Raphael Hiesgen, Thomas C. Schmidt, and Matthias Wählisch. 2025. POSTER: Two-Phase Scanning in IPv6 - First Observations from a Reactive IPv6 Network Telescope. In *Proc. of ACM SIGCOMM: Posters and Demos*. ACM, New York. https://doi.org/10.1145/3744969.3748443
[65] Tao Yang and Zhiping Cai. 2024. Efficient IPv6 Router Interface Discovery. In *IEEE INFOCOM 2024*. IEEE, Washington, DC, USA, 1641–1650. https://doi.org/10.1109/INFOCOM52122.2024.10621168
[66] Tao Yang, Zhiping Cai, Bingnan Hou, and Tongqing Zhou. 2022. 6Forest: An Ensemble Learning-Based Approach To Target Generation For Internet-Wide IPv6 Scanning. In *IEEE INFOCOM 2022*. IEEE, IEEE, Piscataway, NJ, USA, 1679–1688.
[67] Tao Yang, Bingnan Hou, Zhiping Cai, Kui Wu, Tongqing Zhou, and Chengyu Wang. 2022. 6Graph: A Graph-Theoretic Approach To Address Pattern Mining For Internet-Wide IPv6 Scanning. *Computer Networks* 203 (2022), 1–12. https://doi.org/10.1016/j.comnet.2021.108666
[68] Li Zhang and Shaowei Fu. 2024. 6MCBLM: Multi-scale CNN and BiLSTM-Attention Hybrid Model for IPv6 Target Generation. In *NNICE 2024*. IEEE, IEEE, Piscataway, NJ, USA, 499–505.
[69] Liang Zhao, Satoru Kobayashi, and Kensuke Fukuda. 2024. Exploring the Discovery Process of Fresh IPv6 Prefixes: An Analysis of Scanning Behavior in Darknet and Honeynet. In *Proc. of PAM (LNCS, Vol. 14537)*. Springer, Berlin Heidelberg, 95–111. https://doi.org/10.1007/978-3-031-56249-5_4
[70] Johannes Zirngibl, Lion Steger, Patrick Sattler, Oliver Gasser, and Georg Carle. 2022. Rusty clusters? Dusting an IPv6 research foundation. In *Proc. of the ACM IMC*. ACM, New York, NY, USA, 395–409. https://doi.org/10.1145/3517745.3561440

## A Ethics

This work does not raise any ethical issues.

## B Artifacts

All artifacts of this paper are publicly available. These include (*i*) all raw measurement data that we captured during our 11 month observation period; (*ii*) all data derived from postprocessing *i.e.,* data that substantiate our arguments and serve as input for our figures; (*iii*) post-processing scripts. All artifacts and details on how to use them are archived here: https://doi.org/10.5281/zenodo.16419096.

## C Testing for randomness: NIST Test Suite

The NIST Test Suite comprises 15 different randomness tests. We exclude tests that either require over 1,000 bits of randomness or additional information. This leaves us with four tests for our analysis.

**Frequency (monobit)** tests the balance between ones and zeros in a sequence to determine uniform randomness. Failure in this initial test make other tests likely to fail as well.

**Runs** examines the number of uninterrupted sequences of identical bits to evaluate the randomness of the oscillation between ones and zeros.

**Discrete Fourier transform (spectral) (FFT)** analyzes peak heights in the Discrete Fourier Transform of a sequence to identify (non-random) periodic features.

**Cumulative sums (cusum)** evaluates deviation of the sum of numbers from the expected values of a random sequence. It can be applied forward (cusum0) or backward (cusum1).

**NIST test input.** We selected sessions with at least 100 packets from our dataset. This filters out all but 2219 sessions, roughly 2.4%, which include 94% of all packets. Since we observed scanners that use different approaches for different parts of the address, we test address sections separately: the first 32 bits after our fixed /32 prefix (subnet) and the last 64 bits (interface identifier (IID)).

**NIST test results.** Figure 18 shows the share of scan sessions in our data set that succeed and fail in the selected NIST tests, separated by IID (left) and subnet (right). Results are further categorized by the temporal behavior of the scanner. Success signifies the selection to be random. For subnets, NIST test mostly fail (right column in Figure 18). In contrast, IID selections pass more frequently the NIST test than selected subnets. Combining both observations suggest that scanners favor a structured approach to select subnets but are more likely to choose random addresses inside the prefixes. These results bolster our observations in Figure 13.

Non-random selection of IIDs occurs most often in *periodic* and *intermittent* scanners. Their repeated scanning cycle might influence their selection as they want to get a comprehensive view of the address space instead leaving it to chance. Interestingly, *one-off* scanners are more likely to randomly select IIDs than the other categories. However, no category shows scans with predominantly random selection.

We do not know how TGAs influence the selection of addresses and how they influence the outcome of this analysis. This will be part of our future work. These results support our assumption. Random scanning cannot be detected by categorizing targeted addresses individually but requires context from the scan session.

Received December 2024; revised June 2025; accepted June 2025

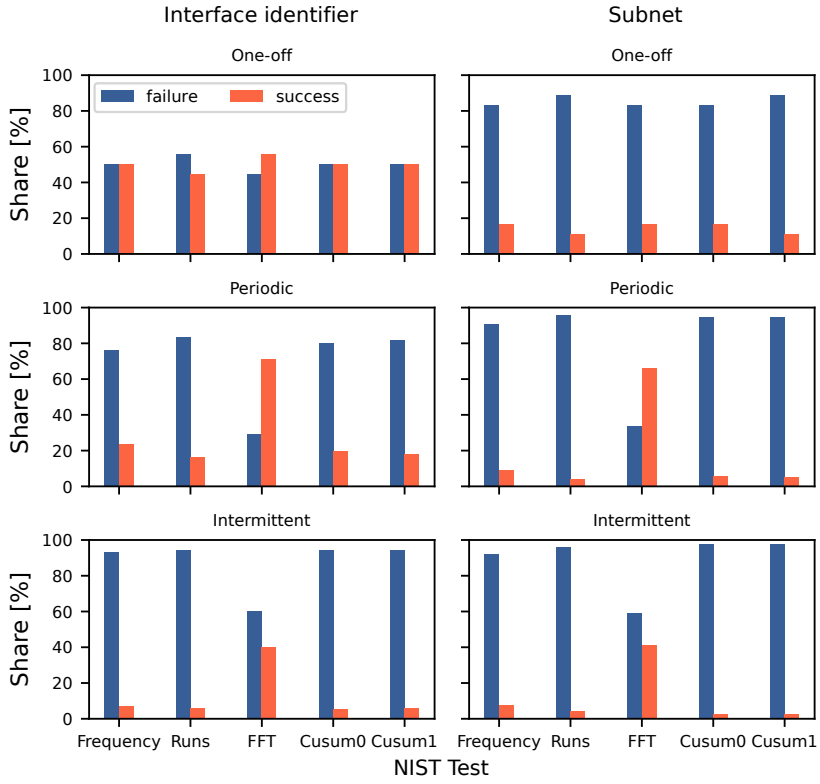


Fig. 18. Results of the NIST tests for T1 sessions with ≥ 100 packets. Scanners tend to iterate IIDs more randomly (*i.e.*, test success) compared to subnets.