



PDF Download
3730567.3764502.pdf
20 January 2026
Total Citations: 0
Total Downloads: 121

 Latest updates: <https://dl.acm.org/doi/10.1145/3730567.3764502>

SHORT-PAPER

Time To Scan: Digging into NTP-based IPv6 Scanning

MICHAEL KLOPSCH, RWTH Aachen University, Aachen, Nordrhein-Westfalen, Germany

CONSTANTIN SANDER, RWTH Aachen University, Aachen, Nordrhein-Westfalen, Germany

KLAUS WEHRLE, RWTH Aachen University, Aachen, Nordrhein-Westfalen, Germany

MARKUS DAHLMANN, RWTH Aachen University, Aachen, Nordrhein-Westfalen, Germany

Open Access Support provided by:

RWTH Aachen University

Published: 28 October 2025

[Citation in BibTeX format](#)

IMC '25: ACM Internet Measurement Conference

October 31, 2025

WI, Madison, USA

Conference Sponsors:

SIGMETRICS

SIGCOMM

Time To Scan: Digging into NTP-based IPv6 Scanning

Michael Klopsch
michael.klopsch@rwth-aachen.de
Communication and Distributed Systems
RWTH Aachen University
Germany

Klaus Wehrle
wehrle@comsys.rwth-aachen.de
Communication and Distributed Systems
RWTH Aachen University
Germany

Constantin Sander
sander@comsys.rwth-aachen.de
Communication and Distributed Systems
RWTH Aachen University
Germany

Markus Dahlmanns
dahlmanns@comsys.rwth-aachen.de
Communication and Distributed Systems
RWTH Aachen University
Germany

Abstract

Due to its large address space, IPv6 remains a challenge for Internet measurements. Thus, IPv6 scans often resort to hitlists that, however, mainly cover core Internet infrastructure and servers. Contrarily, a recent approach to source addresses leveraging NTP servers promises to discover more user-related hosts. Yet, an in-depth analysis of hosts found by this approach is missing and its impact remains unclear.

In this paper, we close this gap by sourcing client IPv6 addresses from our NTP Pool servers and scanning related hosts. We get 3 040 325 302 IPv6 addresses, unveiling 283 867 deployments of consumer products underrepresented in a state-of-the-art hitlist, only leading to 37 858 finds. Security-wise, we find that only 28.4 % of 73 975 NTP-sourced SSH and IoT-related hosts appear to be securely configured, compared to 43.5 % of 854 704 hosts in the hitlist, revealing previously underestimated security issues. Last, we switch sides and identify first (covert) actors adopting NTP-based address sourcing in their scanning.

CCS Concepts

• **Networks** → **Network measurement; Security protocols.**

Keywords

Internet measurements; IPv6; NTP; address sourcing

ACM Reference Format:

Michael Klopsch, Constantin Sander, Klaus Wehrle, and Markus Dahlmanns. 2025. Time To Scan: Digging into NTP-based IPv6 Scanning. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25)*, October 28–31, 2025, Madison, WI, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3730567.3764502>

1 Introduction

Over a decade ago, the introduction of ZMap [21] established large-scale IPv4 address scanning and enabled various research [20] that, e.g., unveiled numerous security issues: Missing encryption [6],

private keys shared among different hosts [17, 25], the use of outdated and broken cryptography [5], and missing access control [17] across a variety of protocols. However, all ZMap-based works leave out IPv6 as a substantial part of the Internet as it would take eons to scan the whole IPv6 address space [16, 21, 34, 37, 43, 44, 52, 53, 57].

Instead, to gain insights into the configuration of IPv6 deployments, common practice is to only scan a promising subset of the IPv6 address space [16, 26, 44, 57]. State-of-the-art methods compile such subsets through the use of traceroute-like probes [10, 12], extraction of data from the DNS [11, 24, 53, 55], and machine learning to extend already existing subsets [14, 23, 49]. A widely used [16, 30, 52] and regularly updated subset is the TUM IPv6 Hitlist [26, 52, 57] that is built by a combination of these methods.

Yet, analysis results based on such hitlists heavily depend on the included addresses. For instance, the aforementioned methods and the TUM IPv6 Hitlist are known to overrepresent Internet infrastructure and servers [15, 26, 44, 49], while they tend to include less end-user deployments. This can lead to overseen end-user security issues. To overcome this limitation, Rye and Levin [44] recently introduced a novel approach leveraging NTP servers and the NTP Pool to collect further IPv6 addresses.

While the authors proved that their concept collects IPv6 addresses with less focus on infrastructure and servers, its impact on application layer scans is unknown. Hence, it is unclear whether the additionally found deployments really complement current state-of-the-art IPv6 measurements.

In this paper, we thus address the research gap of understanding how NTP servers can help to complement the current view on IPv6 reachable deployments. To this end, we (i) replicate Rye and Levin’s setup to collect IPv6 addresses via NTP servers, (ii) scan the sourced IPv6 addresses for running HTTP, SSH, AMQP, MQTT, and CoAP services, (iii) analyze their security configuration, as well as (iv) compare our results against the TUM IPv6 Hitlist. Additionally, we shed light on scanners that apply Rye and Levin’s approach.

Contributions: Our main contributions are as follows.

- We collect 3 040 325 302 IPv6 addresses over four weeks and validate our implementation against Rye and Levin.
- We find 283 867 consumer deployments that are overlooked or belittled by hitlist-based analyses.
- We compare 854 704 hitlist-found SSH and IoT hosts against 73 975 hosts unveiled via NTP and find a significant decrease in secure deployments from 43.5 % to 28.4 %.



This work is licensed under a Creative Commons Attribution 4.0 International License. *IMC '25, Madison, WI, USA*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1860-1/2025/10
<https://doi.org/10.1145/3730567.3764502>

- We identify 2 actors relying on NTP-sourcing, with one attempting to conceal the scanning activity.

2 Background and Related Work

Our IPv6 application layer scans on previously undiscovered deployments are driven by (i) state-of-the-art research in building subsets of IPv6 addresses and subsequent application layer scans, and (ii) Rye and Levin’s recent approach of sourcing IPv6 addresses from the NTP Pool, which we discuss in the following.

2.1 State-of-the-Art IPv6 Scanning

Various methods have been established to source or generate IPv6 addresses interesting to scan. Based on these methods researchers performed different host and protocol analyses.

2.1.1 Sourcing IPv6 Addresses. To initially discover IPv6 addresses, research mainly relies on the DNS. Hostnames from, e.g., certificate transparency logs [24, 39], zone walks [27, 55], and IPv4 reverse DNS lookups [53] as well as the IPv6 reverse zone [11, 22] are used to compile interesting-to-scan addresses. Yet, DNS entries for hosts of interest need to exist [15], thus nudging scans towards content-providing hosts. Hence, traceroute-like probes [12, 34, 43, 45] allow to add core Internet infrastructure [10, 12] and last-hop routers [45].

Additionally, target generation algorithms extrapolate already found addresses through statistical analysis and machine learning. For instance, Entropy/IP [23] and its extensions [10, 37] employ entropy clustering, while other approaches [14, 29, 49] use machine learning to predict new addresses. Yet, the algorithms still tend to remain biased toward their input addresses [49, 52, 56], although research also tries to generate addresses outside the seed space [49].

As such, the existing approaches for sourcing IPv6 addresses are focused away from end-user “eyeball” networks.

2.1.2 Analyzing IPv6 Hosts. Besides sourcing IPv6 addresses and analyzing addressing properties [30], scanning the hosts is vital for extracting new findings on Internet deployments. Various IPv6 studies look at the adoption of MPTCP [8, 48] or NAT64 [31], analyze IPv6 off-nets [28] or the IoT [16, 33, 46], and fingerprint routers [4].

Scans with security-focus unveil widespread problems with DNSSEC support [38], vulnerabilities of VPNs [35], open ports on routers [15], as well as issues of IoT and HTTPS deployments: IPv6-enabled IoT deployments have a low TLS adoption [16, 33], miss access control [16], often use expired or self-signed certificates [33], and seldom support TLS 1.3 [16, 33]. Similarly, HTTPS measurements found numerous openly accessible configuration pages and a high proportion of not-publicly-trusted certificates [56], as well as high shares of already revoked certificates [50].

While these measurements provide important insights into the IPv6 Internet, due to their reliance on hitlists, they focus more on servers or Internet infrastructure and less on end-user deployments.

2.2 NTP-based IPv6 Address Sourcing

Focusing on these previously neglected end-user devices, Rye and Levin (R&L) [44] recently proposed to source IPv6 addresses using the NTP Pool. As the NTP Pool is a widely-used global collection of timeservers, billions of clients, including end-user devices, regularly synchronize their time and thus connect to the servers and transmit

their (IPv6) addresses. By running 27 NTP servers announced in the NTP Pool for seven months, R&L were able to collect nearly 8 billion distinct IPv6 addresses. Multiple factors, such as differences in the address structure compared to the TUM IPv6 Hitlist, ASes the addresses belong to, and information from embedded MAC addresses, show that the addresses indeed include many end-user devices. However, so far, no work exists that performs active scans with the NTP-sourced addresses.

Takeaway: *So far, IPv6 measurements largely overlook end-user devices due to their focus toward servers or Internet infrastructure. To compensate, NTP-based sourcing allows focusing on end-user devices.*

3 Collecting IPv6 Addresses

To understand potential influences of missing end-user devices in active IPv6 scans, we first need to source addresses via NTP servers, reproducing Rye and Levin’s (R&L’s) setup.

3.1 Methodology

Under comprehensive ethical measures (cf. Appendix A), we deploy 11 NTP servers modified to capture client addresses and add them to the NTP Pool. Since we deploy less servers than R&L (to diminish financial expenses and influences on the NTP Pool), we cannot match their original geographic distribution. Additionally, we could not determine the criteria used to select the server locations. Still, to ensure an efficient address collection, we target countries having a small proportion of already existing NTP servers [40] in comparison to their routed IPv6 addresses [1, 42]. This estimation leads us to deploy a server in each of Australia, Brazil, Germany, India, Japan, Poland, South Africa, Spain, the Netherlands, the United Kingdom, and the United States.

After adding the servers, we monitor the number of requests and increase our servers’ operator-configurable weight in the NTP Pool until reaching, at peak times, a request rate close to our maximum scanning rate (cf. Appendix A). This approach allows us to perform continuous real-time scans to analyze the host configuration (cf. Section 4.1). For comparison with results from state-of-the-art address sources, we rely on the frequently [4, 8, 16, 30, 31, 33, 46, 48] used TUM IPv6 Hitlist which we scan after reaching a similar number of IP addresses collected by our NTP servers. Focusing on the timespan between finalizing our pool configuration and the conclusion of the hitlist scans, we consider NTP-sourced data from between July 20th 2024 and August 16th 2024.

3.2 Comparison of Collected Addresses

Table 1 reports the number of distinct addresses, containing /48 networks and ASes to gain a first understanding how our set of collected IPv6 addresses compares to the addresses that R&L sourced and the TUM IPv6 Hitlist (both the full and public variant with responsive addresses only) includes.

Using 11 instead of 27 NTP servers over one month instead of seven, we collected over five times as many distinct addresses as a proportional scale-down of R&L’s results. While we see a constant rate of new addresses over the complete collection period, it is possible that a longer collection time results in diminishing returns. Looking closer at the collection rates of our NTP servers, we find multiple orders of magnitude of differences in the number of distinct

	Our Data	Rye and Levin (R&L)	TUM IPv6 Hitlist	
	Jul./Aug. '24	Jan.–Aug. '22	public	full
	Jul. '24 (publ. Aug. '24)			
IP addresses	3 040 325 302	7 914 066 999	21 994 977	2 104 362 964
... overlap	—	(unknown)	1 072 543	7 008 330
/48 networks	5 089 323	7 205 127	1 024 223	16 026 551
... overlap	—	2 648 083	514 285	3 310 338
ASes	10 515	10 886 ¹	21 438	27 488
... overlap	—	3131	9555	10 311
median IPs in /48s	5	(unknown)	1	2
median IPs in ASes	708.5	(unknown)	10	86

Table 1: Number of distinct IPs/networks per dataset.

IP addresses per NTP server (India: 2.6 billion vs. the Netherlands: 9 million, full list in Appendix D). Thus, we see the geographic server location as one important factor next to, e.g., the operator-configurable server weight parameter netspeed in determining how many clients a server sees. We attribute our overall scale-adjusted higher efficiency to a mix of the various factors, as well as possible shifts in the underlying data since R&L’s research in 2022.

Looking further into the distribution of addresses, we can also see that we achieve a broader coverage of different networks. For the ratio of included /48 networks and ASes to total addresses, we find similarly many ASes and 70% of the /48 networks with less than half as many distinct IPs.

Also, the overlap of /48 networks and ASes between our and R&L’s sets shows that while some networks and ASes emit a sufficient amount of NTP requests to reach R&L’s and our deployments, we still find many new networks and ASes providing a broader perspective. Nevertheless, both versions of the TUM IPv6 Hitlist contain most of the ASes we found and show high overlap with our /48 networks—a result matching R&L’s findings in their paper and confirming that the methods used by the TUM IPv6 Hitlist lead to a wide range of covered networks.

Like R&L, we find that the average density of /48 networks (i.e., average IPs per network) is higher for our NTP-sourced data compared to the TUM IPv6 Hitlist, suggesting client-side networks. As this difference also holds for the median density of networks and ASes (cf. Table 1, bottom), it is not caused by single outliers, but widespread across the data sets.

3.2.1 Properties of Collected Addresses. To get a first intuition on the hosts behind the IPv6 addresses, in alignment with R&L [44], we group the addresses based on their interface identifiers (IIDs): whether these are zeroes, have only the last (two) byte(s) set (“structured” addresses), and, for others, by their entropy. The distribution in Figure 1 shows that the TUM IPv6 Hitlist contains a higher share of structured addresses than R&L’s and our data set, indicating manually configured server or router addresses. For the public version of the Hitlist, the difference is more pronounced, matching the expectation that servers are more responsive to scans.

In direct relation to R&L’s results, our dataset shows less entropy. Yet, more of R&L’s addresses also only have the last two bytes set compensating this shift. Hence, we argue that the shape of included addresses in our dataset is still similar to R&L’s including more

¹Mapping of /48 subnets from R&L’s public dataset based on RIPE RIS data from 2022. While different from the 9006 ASes R&L originally reported, we rely on our count to uphold comparability.

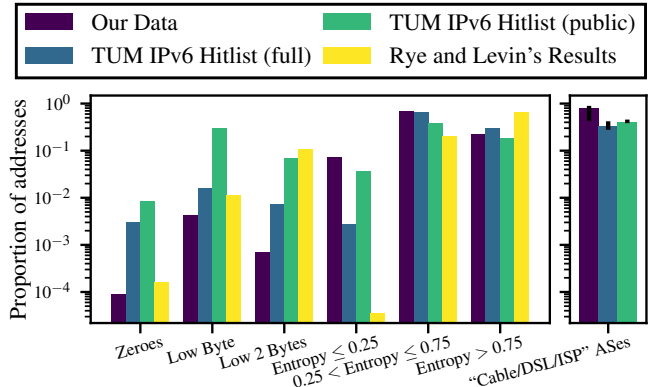


Figure 1: Prop. of addresses grouped by IID and AS.

addresses from end-user “eyeball” devices than the state-of-the-art TUM IPv6 Hitlist. Underpinning this result, we also find the proportion of addresses whose AS is labeled as “Cable/DSL/ISP” in the PeeringDB [2] to be higher than in the TUM IPv6 Hitlist. To get an even better intuition on the devices behind found addresses, we include an analysis of embedded EUI-64 identifiers in Appendix B.

Takeaway: In comparison to R&L, our setup reaches largely similar results, with differences largely caused by a mix of various configuration options and other factors. The address structure indicates that our NTP-sourced data focuses more on end-user devices compared to state-of-the-art hitlists.

4 Usefulness for Active Scans

To analyze whether the structural differences of addresses in our NTP-sourced set indeed imply finding new hosts with a different configuration than in the TUM IPv6 Hitlist, we now perform active scans.

4.1 Methodology

During the design and execution of our scans, we strictly follow our ethical considerations (cf. Appendix A). To gain a broad view on the hosts behind the IPv6 addresses while keeping their measurement-incurred load low, we focus our scans on widely-used classical protocols and, as we expect such deployments in end-user networks, modern IoT protocols. Specifically, we scan HTTP and SSH as protocols that yield the most responsive hosts in the IPv4 space [13] as well as AMQP, MQTT, and CoAP as modern IoT protocols that were subject to scans in the past [16, 33] on their respective IANA port. We also scan the TLS variants of HTTP, AMQP, and MQTT.

For scanning, we rely on version of zgrab2 [54] extended to support the protocols we scan and fed in real time with every address our NTP servers found during our collection period. In the last week of our collection period (August 9th to August 16th), we also scan the full TUM IPv6 Hitlist (not the public list filtered to responsive addresses) for comparison.

4.2 Data Set Overview

We summarize our scan results in Table 2. For all protocols except CoAP, we find more active endpoints using the TUM IPv6 Hitlist than using our NTP source (#Addr-columns in Table 2). This result is in line with the assumption that the TUM IPv6 Hitlist focuses on

Protocol (Ports)	Our Data			TUM IPv6 Hitlist			Overlap
	#Addrs	#Addrs w/ TLS	#Certs/Keys	#Addrs	#Addrs w/ TLS	#Certs/Keys	#Certs/Keys
HTTP (80, 443)	508 799	396 351 (77.9%)	285 615 (72.3%)	379 136 782	16 234 079 (4.28%)	1 135 845 (7.00%)	22 504
SSH (22)	293 229	— —	73 923 (25.2%)	2 218 005	— —	852 760 (38.4%)	4553
MQTT (1883, 8883)	4316	334 (7.73%)	43 (12.9%)	48 987	1062 (2.17%)	843 (79.4%)	5
AMQP (5672, 5671)	1152	14 (1.22%)	9 (64%)	3083	111 (3.60%)	101 (91.0%)	2
CoAP (5683 (UDP))	5093	— —	— —	1511	— —	— —	—

Table 2: Successful scans by protocol. Number of unique certificates/keys provides lower bound of actual hosts.

servers. Only with CoAP—specifically targeted at IoT end-devices instead of servers—we find over 3 times more endpoints via NTP, which again hints at seeing more eyeball devices via NTP.

The TLS adoption among scanned HTTP servers is comparatively high for the NTP-sourced addresses (#Addrs w/ TLS-columns in Table 2), likely due to the increased focus this area had in the last years, e.g., as driven by Let’s Encrypt [3]. In contrast, the very low TLS adoption rate of HTTP servers behind the TUM IPv6 Hitlist addresses is due to numerous (356 million) Cloudfront addresses with failed TLS handshakes (probably due to our requests missing a hostname), showing hyperscalers’ impact on focused IPv6 Internet measurements. For AMQP and MQTT, the TLS adoption across both address sources is low as well, indicating security problems with these protocol deployments.

Looking at the number of unique certificates of TLS-enabled HTTP, MQTT, and AMQP hosts as well as SSH host keys we found via NTP-sourcing (#Certs/Keys-columns in Table 2), we see that MQTT and SSH present a much lower proportion of unique keys than HTTP and AMQP. For CoAP, we filter by the embedded MAC addresses (not visible in Table 2) and find a proportion of 70%—similarly to what we see for HTTP and AMQP. As such, for HTTP, AMQP and CoAP, we are confident that we did not excessively find the same hosts over and over again. For MQTT and SSH, we suspect widespread key-reuse [17, 25] (cf. Section 6) to occur, also causing some unique keys to end up in multiple categories in our later analysis.

To account for finding the same host at different addresses multiple times, e.g., due to dynamic IP addresses, we filter for unique TLS certificates and SSH host keys as proxy for unique hosts. While this approach yields a hard lower bound for the actual number of unique hosts, considering results based on grouping by network provides numbers that are potentially more realistic, but less reliable. Hence, we report those results in Appendix C while focusing on the more reliable estimate based on unique TLS certificates and SSH host keys in the following.

4.3 New Types of Devices

To identify which type of deployments were, until now, underrepresented in IPv6-based scans and estimate the impact of NTP-based address sourcing on active Internet measurements, we next analyze protocol-specific indicators from the found hosts. Specifically, we focus on the TLS-enabled HTTP deployments, SSH hosts, as well as CoAP devices and extract the site title, OS description, and CoAP resources respectively. While Table 3 summarizes our results, we now traverse through the protocols individually.

4.3.1 HTTP. To approximate the deployment type for HTTP servers, we extract the HTML page titles (to exclude CDN error pages,

status code 200 only). Furthermore, to disregard minor variations in, e.g., version numbers, we group titles if their Levenshtein distance normalized to 0–1 is at most 0.25.

The resulting groups show that endpoints from the TUM IPv6 Hitlist mostly respond with default or error pages as well as pages indicating D-LINK network infrastructure and 3CX telecommunication servers. In contrast, our NTP-sourced hosts comprise mostly consumer devices from AVM (FRITZ!), as well as a consumer Wi-Fi device from Cisco. The much higher number of FRITZ! devices and the existence of other devices not found by the hitlist-based scans show that the NTP Pool yields completely new device types that were missed before. We attribute the comparatively high prevalence of AVM (FRITZ!) products to the fact that AVM makes it extremely easy to make devices available from the Internet, while other similar devices might not offer a web interface accessible from the Internet. Concretely, our scans were able to find 16 852 devices of types missed by the TUM IPv6 Hitlist next to 257 195 FRITZ!Box devices massively underrepresented in the TUM IPv6 Hitlist.

4.3.2 SSH. For investigating device types via SSH, we leverage that SSH server IDs often contain the OS’s name. Despite an overall lower hit rate, the NTP-sourced addresses account for the vast majority of Raspbian systems which most likely reside in end-user networks. In contrast, with the TUM IPv6 Hitlist we found more FreeBSD systems likely used in core Internet infrastructure, even considering the hitlists’ overall higher response rate. Hence, a more complete view, e.g., including Raspbian devices, requires address sources complementing state-of-the-art hitlists, such as the NTP-based approach unveiling 4765 Raspbian devices.

4.3.3 CoAP. In order to identify CoAP device types, we analyze their advertised resource prefixes. One very popular type of devices offers the `/castDeviceSearch` resource, but cannot be found relying on the TUM IPv6 Hitlist only. While the name suggests a relation to end-user media devices, we could not find any meaningful documentation of this resource. Contrarily, the devices offering `/qlink/*` resources relate to a cryptocurrency-based Wi-Fi service [32]. In total, we found significantly more devices based on NTP-sourced addresses with significantly different advertised resources. Combined, our NTP servers found 5055 devices missed or underrepresented by the TUM IPv6 Hitlist. Thus, again, a more complete view on IPv6-reachable CoAP devices requires address sources beyond the TUM IPv6 Hitlist.

Takeaway: Across different protocols, hitlist-based IPv6 scans miss whole classes of end-user devices. Using NTP-based sourcing, we find 283 867 new or underrepresented devices of at least six distinct types the TUM IPv6 Hitlist neglects.

HTTP			SSH		
HTML Title Group	#Certificates		OS	#Host Keys	
	Our Data	TUM IPv6 Hitlist		Our Data	TUM IPv6 Hitlist
(no title present)	3435 (1.21 %)	309 729 (34.2 %)	Ubuntu	28 522 (38.6 %)	392 207 (46.0 %)
FRITZ!Box	257 195 (90.8 %)	35 841 (3.96 %)	Debian	13 830 (18.7 %)	180 748 (21.2 %)
D-LINK	0 (0 %)	46 548 (5.14 %)	Raspbian	4765 (6.4 %)	658 (0.1 %)
FRITZ!Repeater	14 751 (5.20 %)	7 (0.00 %)	FreeBSD	140 (0.2 %)	14 014 (1.6 %)
(IP) was not found	0 (0 %)	41 384 (4.57 %)	other/unknown	26 677 (36.1 %)	265 219 (31.1 %)
FRITZ!Powerline	1480 (0.52 %)	0 (0 %)	CoAP		
Host Europe GmbH – (IP)	0 (0 %)	38 270 (4.22 %)	resource group	Our Data	TUM IPv6 Hitlist
Common UI	748 (0.26 %)	486 (0.05 %)	castdevice	2967 (58.2 %)	0 (0.0 %)
3CX Webclient	164 (0.06 %)	16 729 (1.85 %)	qlink	2088 (41.0 %)	1352 (75.1 %)
Webinterface	651 (0.23 %)	20 (0.00 %)	efento	4 (0.1 %)	55 (3.1 %)
3CX Phone System Mgmt.	322 (0.11 %)	14 575 (1.61 %)	nanoleaf	1 (0.0 %)	49 (2.7 %)
(Cisco Wi-Fi AP)	621 (0.22 %)	0 (0 %)	empty	21 (0.4 %)	311 (17.3 %)
(other)	3971 (1.40 %)	402 187 (44.4 %)	other	15 (0.3 %)	34 (1.9 %)

Table 3: The NTP-sourcing finds deployments the TUM IPv6 Hitlist does not or only minimally contains.

4.4 Different Security Configurations

To further estimate the influence of including NTP-sourced addresses in active scans, we next analyze the security configurations of the deployments per protocol.

4.4.1 SSH. Installing the latest security patches on SSH servers is of utmost importance to prevent attackers from hijacking sessions or complete systems [9, 41]. Hence, we analyze the version number of found servers. To avoid false positives due to backported fixes [25], we only assess servers unveiling their patch-level, which restricts our analysis to Debian-derived distributions. As updates to the stable distributions only include security or other important bug fixes [47], we consider every non-latest version outdated.

Figure 2 shows a worryingly high proportion of outdated servers that are thus potentially vulnerable to security issues. However, the proportion of outdated servers is far higher when found via NTP. We conjecture that the NTP Pool allows finding less carefully administrated end-user devices, while the TUM IPv6 Hitlist tends to contain better, potentially professionally managed devices. Hence, NTP-based address sourcing adds a different view on security configurations in the wild.

4.4.2 AMQP&MQTT. Since AMQP and MQTT brokers often transmit sensitive (IoT) data, they are targeted by attackers [7, 51]. Yet, such brokers are still often not sufficiently secured and, e.g., lack access control [16, 17, 51]. Figure 3 shows that more than half of the MQTT brokers found via NTP do not implement access control, indicating a severe problem underestimated by our TUM IPv6 Hitlist-based scan where 80 % of MQTT brokers enabled access control. Contrarily, access control is widely deployed for AMQP brokers, which we deduce to it being more heavyweight and potentially

deployed in professional settings. Still, the security configuration differences for MQTT show that NTP-sourced servers open up a diverging view on IPv6 hosts.

Takeaway: The much lower proportion of end-user deployments leads hitlist-based scans to overestimate the security state of devices: The proportion of secure deployments drops from 43.5 % to 28.4 % when instead scanning the NTP-sourced addresses.

5 NTP-Sourcing by Others

Finally, we evaluate the use of R&L’s approach in the wild to see who is using the approach for which use-cases.

5.1 Methodology

In alignment with our ethical considerations (cf. Appendix A), we continuously query NTP servers from the NTP Pool during our scanning period. Based on the NTP Pool info pages, these servers served, on average, 86% of responses. For each query, we use a distinct source IPv6 address and capture the incoming traffic on that address to decide which NTP server triggered a scan. To identify NTP-unrelated scans that found our addresses by chance, we also monitor the surrounding address space for potential scattering.

5.2 NTP-Sourcing on the Rise?

We were able to match all captured scan packets to an NTP query, leading us to identify, next to our own scans, scan campaigns of two additional actors. First, the Georgia Institute of Technology (GT) sources addresses from 15 NTP servers and scans 1011 different ports, indicating an interest in diverse running services such as FTP, BGP, or Postgres. The scans started less than an hour after receiving the NTP response and lasted about 10 minutes, suggesting




Figure 2: NTP-sourcing unveils more outdated hosts.




Figure 3: NTP-sourced servers show worse security.

no attempt to disguise. Hence, NTP address sourcing currently leads to upcoming measurement research that includes a different viewing angle beyond addresses included in today's hitlists.

The second actor does not provide any identifying information and also operates their NTP servers and scanning hosts in ASes of two different cloud providers (Amazon and Linode). Hence, we were not able to identify this actor who neglects established measurement best-practices [21]. Additionally, the scanned ports 443, 8443 (both HTTPS), 3388, 3389, 5900, 5901, 6000, 6001 (all for remote graphical access), 9200 (Elasticsearch), and 27017 (MongoDB) are assigned to services ordinarily protected by access control, suggesting a security focus. Given that the scans span over multiple days with long delays between each attempt and not every address receives connection attempts on all ports, the actor likely aims to avoid detection. Overall, the behavior indicates a covert actor, using NTP-based address sourcing to detect (previously unseen) end-user deployments, where our previous results show a deficient security configuration to be more likely (cf. Section 4.4).

Takeaway: Besides research, NTP-sourcing for IPv6 scans is also used by covert actors. In combination with our findings that end-users are more strongly represented in NTP-sourcing, their security might be affected, too.

6 Discussion, Limitations & Recommendations

We have shown that the focus of NTP-sourced IPv6 addresses away from Internet infrastructure and servers influences the results of analyses. In this section, we discuss additional aspects of our procedure, limitations, and recommendations for future research.

Dynamic IP Addresses: Due to its reactive nature, NTP-sourcing is prone to double-counting hosts with dynamic IP addresses. To still give a lower bound, we relied on certificates and host keys as fingerprints to identify hosts.

Alternatively, globally unique MAC addresses embedded in the IPv6 addresses can help to deduplicate deployments. However, not all deployments use EUI-64 addresses, MAC reuse [44] makes it impossible to estimate definitive bounds, and we detected less distinct MAC addresses than certificates and keys (cf. Appendix B). A more comprehensive fingerprinting method, e.g., based on more application-level data or meta-data from the NTP requests such as request times, for tighter bounds remains for future work.

A further consequence of the comparatively higher proportion of dynamic addresses inherent in the end-user focus is that aggregating NTP-sourced addresses into a list is not useful, as such a list would be outdated almost immediately. Hence, using NTP-sourced addresses for measurements requires an appropriate setup for gathering addresses.

Certificate and Key Reuse: Due to our reliance on host keys and certificates for deduplication, we did not evaluate secret reuse, a widespread and serious problem [17, 25]. To nevertheless get an intuition of this problem, we analyze the number of addresses that reuse secrets which appear in more than two ASes to account for double-homed hosts. We again only consider HTTP status code 200 responses. 45 377 NTP-sourced hosts spanning over 27 ASes relied on the most-used key while the most widespread key was seen in 315 ASes. Overall, 91 773 IP addresses relied on 304 reused keys. In contrast, 23 303 hosts based in 108 ASes found using the TUM

IPv6 Hitlist sent the most-used and -widespread key, with 3846 keys accounting for 143 460 IPs overall. Hence, the much higher proportion of addresses per key for the NTP-based data indicates a more severe reuse problem, possibly due to a higher reliance on pre-built system or container images containing secrets [18].

Hit Rate: While IPv6 scans in general have a low scan-success-ratio [16, 52], our NTP-sourced scans inherently show an even lower overall hit rate (0.42%) as user deployments are typically less reachable from the Internet. Hence, scans require a higher throughput and/or duration to achieve similar data set sizes. However, given the broader view offered by NTP-sourced target addresses the advantages for measurements and analyses outweigh any disadvantages.

Recommendations for IPv6 Measurements: Our results show that NTP-based address sourcing uncovers end-user devices missed by current hitlists. Thus, measurements which want to include end-user devices should consider integrating address sources with focus on end-user devices. Due to dynamic addresses, static hitlists are less useful when targeting end-user devices, so such measurements should make sure to (also) make use of address sources yielding up-to-date addresses. While the NTP-based address sourcing approach we evaluate in this paper is one possible method of attaining live end-user addresses, finding and evaluating the usefulness of other potential sources, including address generators trained on such addresses, remains for future work.

Takeaway: While NTP-based address sourcing introduces challenges, e.g., deduplicating hosts with dynamic IP addresses, the broader view clearly outweighs any disadvantages. Thus, the NTP-based approach is a useful additional address source for researchers wanting to include end-user devices in IPv6 measurements.

7 Conclusion

Rye and Levin [44] recently introduced NTP-based address sourcing to provide insights into end-user devices current IPv6 measurements tend to disregard. However, an analysis of hosts behind the sourced addresses and the impact of the approach was missing.

Our results show that NTP-sourced IPv6 addresses indeed provide a more diverse view, outweighing drawbacks such as a low hit rate or scanning a host multiple times: We identified 283 867 instances of six different consumer products mainly overlooked by hitlist-based scans. Even for well-represented deployment types, we unveil underestimations in analysis results: While 43.5% of 854 704 hitlist-found SSH and IoT hosts appear to be secure, only 28.4% of 73 975 NTP-sourced hosts are, showing a larger problem than initially anticipated. In this light, our finding that a covert actor already performs NTP-sourcing for scans is troublesome and underlines security risks for end-users.

Summarizing, NTP-based IPv6 address sourcing is an essential building block that future IPv6 research should not ignore. The method, which is already used in the wild, finds devices hidden from previous research and exposes underestimated security issues.

Acknowledgments

This work has been funded by the German Research Foundation (DFG) under Grant No. WE 2935/20-1 (LEGATO). We thank the anonymous reviewers and our shepherd for their valuable feedback.

A Ethics

Although our research does not involve human subjects, it still requires ethical considerations as Internet-wide measurements could have unintended implications.

From the deployment of our NTP servers to our active scans, we consistently follow basic ethical guidelines [19] and adhere to rigorous ethical principles through a number of measures. Most importantly, we carefully deploy our NTP servers without disrupting the NTP Pool (cf. Appendix A.1), follow established guidelines during our scans [21] (cf. Appendix A.2) to reduce their impact, and handle collected data with care (cf. Appendix A.3).

A.1 NTP Pool Handling

The collection of IPv6 addresses from the NTP Pool warrants considerations for server deployment, the communication of research goals, and for data handling. We align these considerations in our work with those of R&L [44] and the NTP Pool’s policies, including the NTP Pool’s Terms of Service we carefully reviewed.

A.1.1 Server Deployment. Generally, adding new servers to the NTP Pool improves its performance and decreases the load of already existing NTP servers. However, we address three situations that still can cause issues: *First*, only stable servers that reliably answer NTP requests are a valuable addition for the NTP Pool. Consequently, we thoroughly tested our NTP server software adjustments before their deployment and selected cloud servers guaranteeing nearly 100 % host and network uptime. *Second*, due to the NTP Pool’s client mapping process, adding a server to a currently empty zone (typically a country) *decreases* the service quality for this country’s clients [36]. Thus, we refrained from adding servers to empty zones, which affected a single country in our study. *Third*, the removal of our servers after our study may disrupt clients currently using it. We thus stop advertising our servers in the NTP Pool four weeks before their shutdown, doubling the recommended value of two weeks.

A.1.2 Communication. To detect any unforeseeable impact of our research, during our whole study, we monitored the community forum of the NTP Pool and looked for the description of issues potentially introduced by our study. Additionally, we further communicate the use of our servers for research transparently and prominently by including a note and contact information on our server info pages on the NTP Pool’s website.

Still, we did not see any inquiry related to our study in the forum and did not receive any email.

A.2 Active Internet Measurements

For our active scans, we follow the Good Internet Citizenship guidelines by Durumeric et al. [21] and procedures imposed by our institution. Hereby, we defer to the ethical trade offs already made for existing IPv4-wide scans, where, like in our case, clients are scanned without explicit consent.

A.2.1 Measurement Load. We coordinate our scans closely with our institution’s network operations center and further limit our scans (Section 4) to 100 000 outgoing packets per second to not overwhelm any network. Additionally, we rely on the randomly




Figure 4: The distribution of NTP server location for IP addresses based on whether and what MAC address is embedded

received NTP requests to spread our scans across time and target addresses thus not overloading single ASes. As we expect to also scan low-powered devices, we moreover introduce delays of 10 seconds to 10 minutes between the scans of the different protocols and refrain from scanning the same IP addresses for three days after each scan.

A.2.2 Communication. Similar to our communication within the NTP Pool (cf. Appendix A.1.2), we also communicate the intent of our scans. We clearly identify as performing a research scan in the reverse DNS entries and host web pages on the addresses used for scanning. Moreover, we identify ourselves in protocol-specific fields where possible. The web pages explain purpose and scope of the scans we perform, and also include contact information and instructions for how to opt out. We honor all opt-out requests our institution received during previous scans and promptly reply to communications related to our scans.

A.3 Responsible Data Handling

Given the potential privacy and security impact of the collected data, i.e., either IP addresses that might allow tracking [44] or received protocol-specific information, such as version numbers, we store all data on secured servers. We do not track users nor do we exploit security bugs and, additionally, refrain from publishing any of the collected data to also avoid others from doing so.

B EUI-64 Analysis

Of the IPv6 addresses we gathered, 903 million use one of 675 million EUI-64 IIDs. Only a fraction, 20 million IP addresses based on 9.2 million EUI-64s, have the “unique” bit set. Thus, our proportion of MAC addresses to total IPv6 addresses deviate strongly from the 3% observed by Rye and Levin.²

Taking a closer look at the (claimed) globally unique MAC addresses, we find that 9.1 million (used by 19 million IPs) have an OUI listed in the IEEE’s database [?]. The top vendors as determined from OUIs, shown in Table 4, largely match R&L’s results

²Unfortunately, we could not determine whether R&L count all EUI-64 addresses or only those with the unique bit set. The deviation compared to our data is significant in any case.

Manufacturer	#MACs	#IPs	R&L Rank
AVM Audiovisuelles Marketing und Computersysteme GmbH	6 008 344	14 751 238	↓
Amazon Technologies Inc.	1 121 310	1 474 996	2
AVM GmbH	320 867	819 647	↓
Samsung Electronics Co.,Ltd	186 090	335 621	3
Sonos, Inc.	144 489	192 247	4
vivo Mobile Communication Co., Ltd.	110 794	127 243	5
Shenzhen Ogemray Technology Co.,Ltd	91 961	107 369	↓
(Unlisted)	73 129	96 763	1
China Dragon Technology Limited	69 791	135 935	↓
GUANGDONG OPPO MOBILE TELECOMMUNICATIONS CORP.,LTD	52 405	94 786	↓
Shenzhen iComm Semiconductor CO.,LTD	49 322	68 189	↓
Qingdao Haier Multimedia Limited.	48 100	60 304	↓
QING DAO HAIER TELECOM CO.,LTD.	45 007	54 455	↓
Hui Zhou Gaoshengda Technology Co.,LTD	31 376	45 123	7
Fiberhome Telecommunication Technologies Co.,LTD	29 177	34 029	↓
Tenda Technology Co.,Ltd.Dongguan branch	28 206	32 593	↓
Beijing Xiaomi Electronics Co.,Ltd	27 290	46 554	↓
Earda Technologies co Ltd	26 482	26 984	↓
Guangzhou Shiyuan Electronics Co., Ltd.	26 408	45 632	↓
Shenzhen Cultraview Digital Technology Co., Ltd	25 229	32 925	↓

Table 4: Number of MAC and IP addresses observed by manufacturer indicated in the OUI, with comparison to R&L’s ranking (↓ corresponds to manufacturers not under R&L’s top ten)

with two exceptions: First, we observe significantly less unlisted OUIs, which may be due to methodological differences to R&L, e.g., server locations. Second, our top vendor, AVM GmbH, which is responsible for nearly two thirds of the assigned MAC addresses, is not part of R&L’s top ten vendors. This results correlates well with our HTTP scans showing many AVM devices. AVM’s high prevalence is especially concerning given that R&L found many of these devices to be geolocatable within meters, posing a significant privacy risk. We assume that, similar to the differences in address count, various factors including drift in the underlying data and geographic NTP server location—AVM has a large European market share, where our proportion of servers was higher—contribute to the different results.

Focusing on the potential effect of the geographic server distribution, we plot the distribution of collecting NTP servers for addresses based on the MAC embedding in Figure 4. We see that the majority of addresses whose embedded MAC address was listed in the IEEE’s database were collected by our European NTP servers, confirming our assumption that the manufacturer distribution is influenced by geographic factors. The different distributions when instead looking at unlisted addresses with the “unique” bit and at locally-assigned MAC addresses show structural differences in the address sets collected by our different servers, underlining that the server location greatly influences the types of gathered addresses.

C Scan Results Grouped by Different Criteria

While we focused on unique devices as counted by SSH host keys or TLS certificate fingerprints in our main analysis, we present results considering other metrics here.

Table 5 shows an extended summary of our scan results, showing the number of responsive networks, ASes, and countries (as

determined with MaxMind’s GeoLite2 database [?]). While the TUM IPv6 Hitlist’s higher success rate is still evident across all protocols (except CoAP), the gap lowers when considering more aggregated groups: for example, for SSH, the gap drops from nearly one order of magnitude to less than half an order of magnitude when considering /56 networks instead of addresses. The at least double-digit numbers for ASes and countries our NTP-based scans achieve (except for AMQPS) show that our results are based on devices from a wide range of different networks and not unduly influenced by single operators.

Nevertheless, we still look into our specific results from a network-based angle for additional insights. Table 6 shows the HTML title groups, SSH OSes, and CoAP resource groups when counting based on networks, respectively. The top HTML title groups do change noticeably, indicating that many deployments share keys or are only reachable over insecure plain HTTP. Still, the results continue to show that some device types remain underrepresented or not present at all in the TUM IPv6 Hitlist. For SSH, we observe an even greater difference in numbers due to key reuse, however the changes to the OS distribution are minor, with the results still showing that NTP-sourcing finding many Raspbian devices missed by the TUM IPv6 Hitlist while not being able to find many FreeBSD devices.

Considering our security analysis next, we plot the up-to-date-ness of SSH servers by network in Figure 5. The overall much higher proportion of outdated servers visible here is probably due to outdated servers also reusing keys, as these would be counted only once in our main analysis, but multiple times here. The large gap in up-to-date-ness between NTP- and Hitlist-sourced addresses widens.

Protocol	HTTP	HTTPS	SSH	MQTT	MQTTS	AMQP	AMQPS	CoAP	
Port	80	443	22	1883	8883	5672	5671	5683 (UDP)	
Our Data	IPv6 Addrs	508 472	396 141	292 686	4308	334	1148	14	5080
	/32 nets	3350	2593	3081	228	43	121	6	258
	/48 nets	64 142	34 022	50 694	2339	113	684	9	3717
	/56 nets	283 784	200 291	191 278	3203	231	900	9	4184
	/64 nets	437 934	349 264	224 196	3219	232	903	9	4189
	ASes	2212	1661	2288	145	37	80	6	79
Countries	133	124	131	46	22	26	4	14	
TUM IPv6 Hitlist	IPv6 Addrs	379 136 782	16 188 460	2 218 005	48 987	1062	3083	111	1511
	/32 nets	14 255	12 290	14 463	606	238	349	53	101
	/48 nets	305 527	86 588	599 940	1678	482	710	79	478
	/56 nets	641 264	228 559	759 299	10 918	582	1310	83	1165
	/64 nets	1 274 538	723 913	1 218 980	47 469	632	1527	86	1222
	ASes	12 323	10 641	12 439	482	198	301	43	73
Countries	194	187	208	67	43	51	17	27	

Table 5: Successful scans by protocol per network, AS, and country

HTTP HTML Title Group	Our Data				TUM IPv6 Hitlist			
	IPs	/48	/56	/64	IPs	/48	/56	/64
Welcome to nginx!	6972	2708	3109	3166	116 316	59 044	67 296	75 540
FRITZ!Box	354 934	21 319	174 852	320 204	35 872	6538	19 669	25 718
Apache2 Ubuntu Default Page: It works	10 933	3915	8481	8700	133 892	46 082	97 599	105 795
Nothing Page (empty)	25 150	7711	19 647	20 624	68 682	4578	11 277	37 752
FRITZ!Repeater	17 292	7108	10 959	11 912	5 758 555	41 173	83 018	146 025
FRITZ!Repeater	11 696	2082	9268	10 486	6	1	1	1
Index of /pub/	971	556	702	863	23 960	2726	4739	6071
Login - Join	6093	2817	4022	4030	6421	1657	2074	2365
D-LINK	0	0	0	0	47 287	543	9521	46 898
Home	5937	1826	5571	5614	4240	635	900	1105
Unknown Domain	0	0	0	0	35 993	50	1433	35 993
UFI配置管理-ZHXL_V2.0.0	2503	1839	2498	2503	0	0	0	0
Plesk Obsidian 18.0.34	433	278	334	371	17 172	2560	5670	8188
My Modem	1975	1786	1972	1975	8	1	1	1
GPON Home Gateway	0	0	0	0	34 338	589	4651	31 006
Ms Portal	1812	1420	1758	1758	185	123	130	151
Hier entsteht eine neue Webseite.	0	0	0	0	27 332	38	448	27 283
UFI-JZ_V3.0.0	1218	995	1213	1217	0	0	0	0
Hello! Welcome to Synology Web Station!	4	4	4	4	2265	1953	2141	2152
GAID - WIFI NG BAYAN	1248	1007	1197	1197	0	0	0	0

SSH OS	Our Data				TUM IPv6 Hitlist				CoAP resource group	Our Data				TUM IPv6 Hitlist			
	IPs	/48	/56	/64	IPs	/48	/56	/64		IPs	/48	/56	/64	IPs	/48	/56	/64
Ubuntu	92 886	21 213	57 971	64 892	862 453	340 381	394 750	515 135	castdevice	2967	2863	2925	2925	0	0	0	0
Debian	115 536	18 270	93 481	106 778	647 420	211 972	278 599	433 432	qlink	2088	846	1239	1244	1352	356	1034	1090
Raspbian	42 249	12 327	23 346	27 277	822	576	595	682	efento	4	4	4	4	55	40	45	45
FreeBSD	233	131	138	145	29 309	1564	2094	2767	nanoleaf	1	1	1	1	49	48	49	49
other/unknown	42 374	14 558	23 255	26 388	678 001	60 848	109 157	280 593	empty	18	16	18	18	21	17	17	17
									other	15	5	5	5	34	29	30	31

Table 6: New devices found when counting by networks. With this view, NTP-sourcing still finds new device types.

For AMQP and MQTT, we plot the networks with open and access controlled servers in Figure 6. While for AMQP, the difference between our NTP- and Hitlist-based scans is marginal, suggesting that our previous result was due to small sample size, we see a more interesting result for MQTT. Here, the overall proportion of access control increases, reaching nearly 100% for the TUM IPv6 Hitlist

and individual IPs or /64 networks. Thus, TLS-secured MQTT brokers seem more likely to disable access control, possibly indicating a misunderstanding of what security TLS provides on the side of operators. The absolute gap between NTP- and Hitlist-sourced devices however remains very similar to our previous result of about 40 percentage points (with a corresponding much higher relative

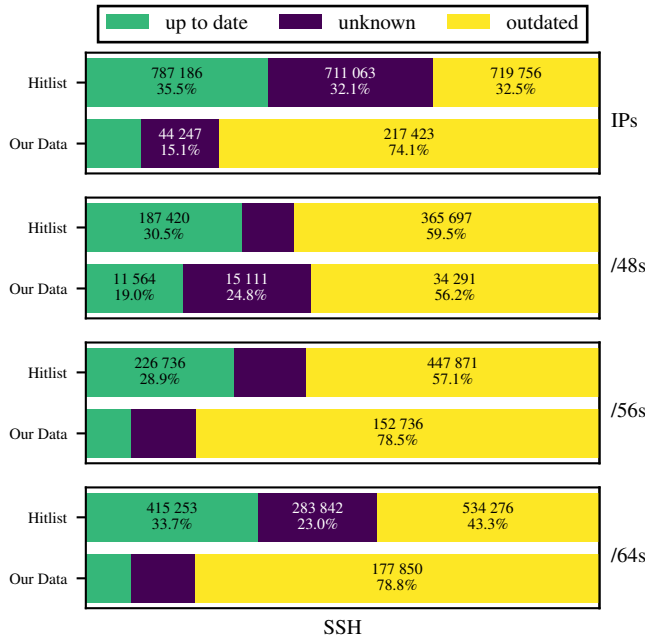


Figure 5: Counting networks instead of unique keys yields much more outdated SSH hosts.

gap due to the overall improvement), confirming our finding that the TUM IPv6 Hitlist misses this large security problem.

D Additional Data

In order to make it easier for follow-up research to compare results with ours, we provide an extended version of our main data tables.

The number of addresses collected by each of our servers is contained in Table 7, clearly showing the large differences between different servers.

Additionally, we list the top 100 OSeS extracted from SSH server IDs in Table 9 and the top 100 HTML titles in Table 8.

Location	#Addresses
India	2 569 110 445
Brazil	224 407 144
Japan	68 729 590
South Africa	36 634 220
Spain	32 921 871
United Kingdom	31 334 399
Germany	25 694 654
United States	24 316 424
Poland	19 103 584
Australia	10 120 272
the Netherlands	9 093 946

Table 7: Number of collected addresses per server.

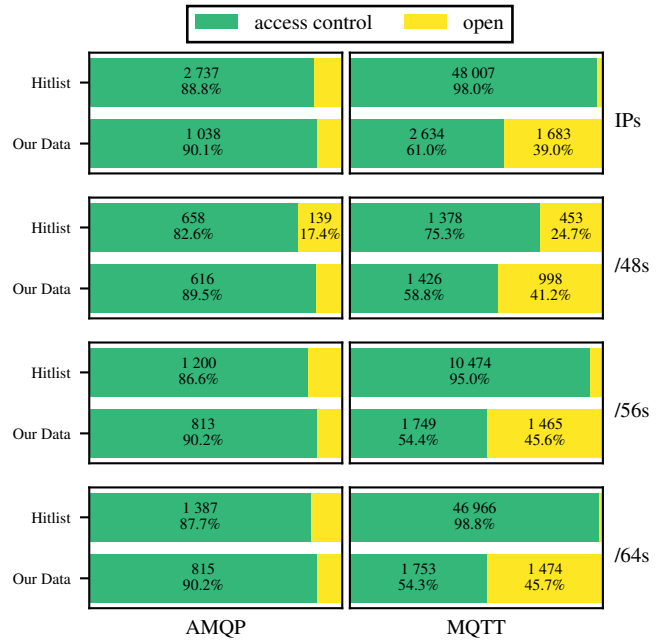


Figure 6: While counting MQTT brokers by networks shows a higher rate of access control overall, the gap between NTP- and Hitlist-sourced brokers remains.

The CoAP resources classified as “other” in the main body are, for the devices found via NTP:

- /maha, /.well-known/core (13 IPs)
- /.well-known/core, /cit, /cit/s (1 IP)
- /.well-known/core, /window, /maha, /loginid, /phonename, /internet_status (1 IP)

and for the devices found via the TUM IPv6 Hitlist:

- /api, /api/v1, /.well-known/core (14 IPs)
- /dp, /rd (5 IPs)
- /gnssPosition, /gpsTime, /.well-known/core (2 IPs)
- /c, /t, /i, /m, /.well-known/core (2 IPs)
- /rd, /virtual_notify, /.well-known/core (1 IP)
- /bs (1 IP)
- /bs, /fw, /rd (1 IP)
- /rd, /dp (1 IP)
- /test, /validate, /hello, /bl%C3%A5b%C3%A6rsyltet%C3%B8y (1 IP)
- /create1 (1 IP)
- /, /time, /async, /example_data (1 IP)
- /.well-known/core, /cit, /cit/s (1 IP)
- /.well-known/eris/blocks (1 IP)
- /bs, /rd, /software, /dp, /firmware, /.well-known/core (1 IP)
- /rd, /.well-known/core (1 IP)

HTML Title Group	Our Data	TUM IPv6 Hitlist
(empty)	3435 (1.20%)	309729 (29.71%)
FRITZ!Box	257195 (90.16%)	35841 (3.44%)
D-LINK	0 (0.00%)	46548 (4.46%)
FRITZ!Repeater 6000	14751 (5.17%)	7 (0.00%)
(IP) was not found	0 (0.00%)	41384 (3.97%)
FRITZ!Powerline 1260	1480 (0.52%)	0 (0.00%)
Host Europe GmbH – (IP)	0 (0.00%)	38270 (3.67%)
Common UI	748 (0.26%)	486 (0.05%)
3CX Webclient	164 (0.06%)	16729 (1.60%)
WebInterface	651 (0.23%)	20 (0.00%)
3CX Phone System Management Console	332 (0.12%)	14575 (1.40%)
WAP150 Wireless-AC/N Dual Radio Access Point with PoE	621 (0.22%)	0 (0.00%)
Plesk Obsidian 18.0.34	447 (0.16%)	13398 (1.29%)
Nothing Page	226 (0.08%)	9519 (0.91%)
Index of /pub/	77 (0.03%)	9451 (0.91%)
Welcome to nono!	121 (0.04%)	7713 (0.74%)
Apache2 Ubuntu Default Page: It works	106 (0.04%)	6301 (0.60%)
Home	117 (0.04%)	2566 (0.25%)
FASTPANEL2	14 (0.00%)	5696 (0.55%)
Remote Console on LAN	88 (0.03%)	0 (0.00%)
Login - Join	39 (0.01%)	4483 (0.43%)
pfSense-nat - Login	73 (0.03%)	314 (0.03%)
Selamat, website (IP) telah aktif!	0 (0.00%)	4024 (0.39%)
UniFi OS	66 (0.02%)	250 (0.02%)
Domain Default page	34 (0.01%)	3098 (0.30%)
Login Absensi	59 (0.02%)	173 (0.02%)
Hier entsteht eine neue Webseite.	0 (0.00%)	2352 (0.23%)
OctoPrint Login	42 (0.01%)	12 (0.00%)
Freebox OS :: Identification	0 (0.00%)	2309 (0.22%)
awdial - Login page	39 (0.01%)	162 (0.02%)
Account suspended	6 (0.00%)	2229 (0.21%)
MS Console	38 (0.01%)	1 (0.00%)
Hello! Welcome to Synology Web Station!	2 (0.00%)	1901 (0.18%)
Cloudea	35 (0.01%)	16 (0.00%)
YunoHost admin	18 (0.01%)	1840 (0.18%)
Red Hat OpenShift	28 (0.01%)	82 (0.01%)
Welcome !	26 (0.01%)	1810 (0.17%)
One moment, please...	27 (0.01%)	874 (0.08%)
IC Hosting	25 (0.01%)	1747 (0.17%)
Invisible Internet Protocol Daemon	25 (0.01%)	2 (0.00%)
mail UI	10 (0.00%)	1606 (0.15%)
FreePBX Administration	25 (0.01%)	189 (0.02%)
NAS1 - Synology DiskStation	0 (0.00%)	1538 (0.15%)
Index	24 (0.01%)	212 (0.02%)
Nová doména u Váš Hosting	0 (0.00%)	1355 (0.13%)
this is a mail-in-a-box	23 (0.01%)	1327 (0.13%)
Outlook	1 (0.00%)	1234 (0.12%)
Router	22 (0.01%)	145 (0.01%)
Ceet Webmail :: Welcome to Ceet Webmail	14 (0.00%)	1202 (0.12%)
Issabel - Página de Ingreso	20 (0.01%)	73 (0.01%)
vaiton – Just another WordPress site	3 (0.00%)	1184 (0.11%)
Avaya J139 Phone	19 (0.01%)	0 (0.00%)
Sign in - GitLab	6 (0.00%)	1094 (0.10%)
Opening...	18 (0.01%)	24 (0.00%)
Web Hosted by Hostico	0 (0.00%)	1079 (0.10%)
Projects	18 (0.01%)	83 (0.01%)
Jitsi Meet	12 (0.00%)	1026 (0.10%)
Ubiquiti EdgeSwitch	17 (0.01%)	11 (0.00%)
Mineral – My WordPress Blog	3 (0.00%)	939 (0.09%)
IISA Windows	15 (0.01%)	27 (0.00%)
- Laravel	1 (0.00%)	934 (0.09%)
Homebridge	14 (0.00%)	3 (0.00%)
Shared IP	13 (0.00%)	910 (0.09%)
ZeroShell	14 (0.00%)	0 (0.00%)
Default Parallels Plesk Page	7 (0.00%)	852 (0.08%)
EdgeOS	14 (0.00%)	372 (0.04%)
Home - Mine	4 (0.00%)	839 (0.08%)
alivaris.com – Domain default page	13 (0.00%)	808 (0.08%)
IIS Windows Server	2 (0.00%)	838 (0.08%)
Admin Panel	13 (0.00%)	39 (0.00%)
Site is under construction	3 (0.00%)	814 (0.08%)
Elastix - Login page	12 (0.00%)	4 (0.00%)
Site in Maintenance	6 (0.00%)	811 (0.08%)
C Cloud	11 (0.00%)	303 (0.03%)
Website (IPv4).cloudvps.regruhosting.ru is ready. The content is to be added	3 (0.00%)	807 (0.08%)
Poweradmin	10 (0.00%)	47 (0.00%)
Login – NextCloud	7 (0.00%)	781 (0.07%)
恭喜, 站点创建成功!	9 (0.00%)	322 (0.03%)
Grafana	5 (0.00%)	710 (0.07%)
Welcome to AutoSMTP.com Ultimate Email Marketing Solution	9 (0.00%)	28 (0.00%)
Sign In - gatos	4 (0.00%)	704 (0.07%)
EPE Journals	9 (0.00%)	3 (0.00%)
phpMyAdmin	4 (0.00%)	698 (0.07%)
(IP)	9 (0.00%)	15 (0.00%)
404 Not Found	2 (0.00%)	620 (0.06%)
Pritunl	8 (0.00%)	511 (0.05%)
Synology Router - Router	0 (0.00%)	617 (0.06%)
Web management Home	8 (0.00%)	263 (0.03%)
Cloudron Not Found	5 (0.00%)	593 (0.06%)
SmokePing Latency Page for Network Latency Grapher	8 (0.00%)	25 (0.00%)
ACASA	6 (0.00%)	591 (0.06%)
Error 404	7 (0.00%)	149 (0.01%)
BigBlueButton@FAU	4 (0.00%)	565 (0.05%)
ISL Conference Proxy	7 (0.00%)	8 (0.00%)
Unknown address	2 (0.00%)	545 (0.05%)
Aria - Business HTML Landing Page Template	7 (0.00%)	3 (0.00%)
CGNV2	0 (0.00%)	534 (0.05%)
Core Management	7 (0.00%)	71 (0.01%)
Request rejected :(0 (0.00%)	504 (0.05%)
PHP 8.3.1 - phpinfo()	7 (0.00%)	157 (0.02%)

Table 8: Top 100 extracted HTML title groups by unique certificate fingerprint

OS	Our Data	TUM IPv6 Hitlist
Ubuntu	28 522 (38.58%)	392 207 (45.99%)
(empty)	26 622 (36.01%)	260 804 (30.58%)
Debian	13 830 (18.71%)	180 748 (21.19%)
Raspbian	4 765 (6.44%)	658 (0.08%)
FreeBSD	140 (0.19%)	14 014 (1.64%)
Endless	25 (0.03%)	0 (0.00%)
PKIX	1 (0.00%)	1 828 (0.21%)
NetBSD	22 (0.03%)	245 (0.03%)
FIPS	0 (0.00%)	1 442 (0.17%)
Trisquel	3 (0.00%)	57 (0.01%)
OVH	1 (0.00%)	528 (0.06%)
Deepin	1 (0.00%)	3 (0.00%)
FlowSsh	1 (0.00%)	35 (0.00%)
ConfigManager	1 (0.00%)	2 (0.00%)
SSH	0 (0.00%)	31 (0.00%)
SimpleRezo	0 (0.00%)	30 (0.00%)
uio	0 (0.00%)	22 (0.00%)
Unknown	0 (0.00%)	20 (0.00%)
server	0 (0.00%)	15 (0.00%)
5e	0 (0.00%)	12 (0.00%)
2	0 (0.00%)	10 (0.00%)
3	0 (0.00%)	8 (0.00%)
Vyatta	0 (0.00%)	7 (0.00%)
DjilOS	0 (0.00%)	6 (0.00%)
SENTINEL	0 (0.00%)	6 (0.00%)
Wait	0 (0.00%)	6 (0.00%)
compatible	0 (0.00%)	5 (0.00%)
FTP	0 (0.00%)	4 (0.00%)
Celeonet	0 (0.00%)	3 (0.00%)
FlokiNET	0 (0.00%)	3 (0.00%)
Linux	0 (0.00%)	3 (0.00%)
All	0 (0.00%)	3 (0.00%)
Windows	0 (0.00%)	3 (0.00%)
Globalscape	0 (0.00%)	3 (0.00%)
Clebian	0 (0.00%)	2 (0.00%)
8	0 (0.00%)	2 (0.00%)
MidnightBSD	0 (0.00%)	2 (0.00%)
MirBSD	0 (0.00%)	2 (0.00%)
ssh	0 (0.00%)	2 (0.00%)
null	0 (0.00%)	2 (0.00%)
Transfer	0 (0.00%)	2 (0.00%)
CM4all	0 (0.00%)	2 (0.00%)
Server	0 (0.00%)	2 (0.00%)
virtuel	0 (0.00%)	2 (0.00%)
sftp	0 (0.00%)	1 (0.00%)
ootanuki	0 (0.00%)	1 (0.00%)
kitty	0 (0.00%)	1 (0.00%)
DEARDHSTHISISPATCHEDDAMMIT	0 (0.00%)	1 (0.00%)
CentOS	0 (0.00%)	1 (0.00%)
Microsoft	0 (0.00%)	1 (0.00%)
Suppegi	0 (0.00%)	1 (0.00%)
NFX	0 (0.00%)	1 (0.00%)
Zoo	0 (0.00%)	1 (0.00%)
Sentinel	0 (0.00%)	1 (0.00%)
Never	0 (0.00%)	1 (0.00%)
1989	0 (0.00%)	1 (0.00%)
WeirdSSH2	0 (0.00%)	1 (0.00%)
1	0 (0.00%)	1 (0.00%)
Camxos	0 (0.00%)	1 (0.00%)
bshax	0 (0.00%)	1 (0.00%)
UNIX	0 (0.00%)	1 (0.00%)
rollback	0 (0.00%)	1 (0.00%)
fabSD	0 (0.00%)	1 (0.00%)
Devuan	0 (0.00%)	1 (0.00%)
IBM11	0 (0.00%)	1 (0.00%)
alcatraz	0 (0.00%)	1 (0.00%)
Udfxos	0 (0.00%)	1 (0.00%)
cuatro	0 (0.00%)	1 (0.00%)
b726	0 (0.00%)	1 (0.00%)
mombe	0 (0.00%)	1 (0.00%)
TUNNP	0 (0.00%)	1 (0.00%)
secretroad	0 (0.00%)	1 (0.00%)
Epopen	0 (0.00%)	1 (0.00%)
GET	0 (0.00%)	1 (0.00%)
fyag	0 (0.00%)	1 (0.00%)
HelloWorld	0 (0.00%)	1 (0.00%)
YouWillNotSeeMyDistro	0 (0.00%)	1 (0.00%)
SunOS	0 (0.00%)	1 (0.00%)
evil	0 (0.00%)	1 (0.00%)
hackerhorse	0 (0.00%)	1 (0.00%)
OpenSSH	0 (0.00%)	1 (0.00%)
NULL	0 (0.00%)	1 (0.00%)
in	0 (0.00%)	1 (0.00%)
RF	0 (0.00%)	1 (0.00%)
NSA	0 (0.00%)	1 (0.00%)
Google	0 (0.00%)	1 (0.00%)
Nosey	0 (0.00%)	1 (0.00%)
SFTP	0 (0.00%)	1 (0.00%)
c579553f35fd	0 (0.00%)	1 (0.00%)
usage	0 (0.00%)	1 (0.00%)
3be07f2e542d	0 (0.00%)	1 (0.00%)
be	0 (0.00%)	1 (0.00%)
da069adcfdd9	0 (0.00%)	1 (0.00%)
IIROC	0 (0.00%)	1 (0.00%)
GlobalSCAPE	0 (0.00%)	1 (0.00%)
VShell	0 (0.00%)	1 (0.00%)
Arts	0 (0.00%)	1 (0.00%)
Greenbone	0 (0.00%)	1 (0.00%)
3a31031c2a76	0 (0.00%)	1 (0.00%)
0bb7d2bed2f2	0 (0.00%)	1 (0.00%)

Table 9: Top 100 extracted OSes from SSH server IDs by unique host key

References

- [1] 2024. RIR Statistics. <https://ftp.lacnic.net/pub/stats/>.
- [2] 2025. PeeringDB. <https://peeringdb.com>
- [3] Josh Aas, Richard Barnes, Benton Case, Zakir Durumeric, Peter Eckersley, Alan Flores-López, J. Alex Halderman, Jacob Hoffman-Andrews, James Kasten, Eric Rescorla, Seth D. Schoen, and Brad Warren. 2019. Let’s Encrypt: An Automated Certificate Authority to Encrypt the Entire Web. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (2019).
- [4] Taha Albakour, Oliver Gasser, Robert Beverly, and Georgios Smaragdakis. 2021. Third time’s not a charm: exploiting SNMPv3 for router fingerprinting. In *Proceedings of the 21st ACM internet measurement conference*. 150–164.
- [5] Martin R Albrecht, Jean Paul Degabriele, Torben Brandt Hansen, and Kenneth G Paterson. 2016. A surfeit of SSH cipher suites. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 1480–1491.
- [6] Johanna Amann, Oliver Gasser, Quirin Scheitle, Lexi Brent, Georg Carle, and Ralph Holz. 2017. Mission accomplished?: HTTPS security after dignotar. In *Proceedings of the 2017 Internet Measurement Conference, IMC 2017, London, United Kingdom, November 1-3, 2017*, Steve Uhlig and Olaf Maennel (Eds.). ACM, 325–340. doi:10.1145/3131365.3131401
- [7] Syaiful Andy, Budi Rahardjo, and Bagus Hanindhito. 2017. Attack scenarios and security analysis of MQTT communication protocol in IoT system. In *2017 4th International Conference on Electrical Engineering, Computer Science and Informatics (EECSI)*. 1–6. doi:10.1109/EECSI.2017.8239179
- [8] Florian Aschenbrenner, Tanya Shreedhar, Oliver Gasser, Nitinder Mohan, and Jörg Ott. 2021. From single lane to highways: Analyzing the adoption of multipath TCP in the internet. In *2021 IFIP Networking Conference (IFIP Networking)*. IEEE, 1–9.
- [9] Fabian Bäumer, Marcus Brinkmann, and Jörg Schwenk. 2024. Terrapin Attack: Breaking SSH Channel Integrity By Sequence Number Manipulation. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Philadelphia, PA, 7463–7480.
- [10] Robert Beverly, Ramakrishnan Durairajan, David Plonka, and Justin P Rohrer. 2018. In the IP of the beholder: Strategies for active IPv6 topology discovery. In *Proceedings of the Internet Measurement Conference 2018*. 308–321.
- [11] Kevin Borgolte, Shuang Hao, Tobias Fiebig, and Giovanni Vigna. 2018. Enumerating active IPv6 hosts for large-scale security scans via DNSSEC-signed reverse zones. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 770–784.
- [12] CAIDA. 2024. The CAIDA Macroscopic Internet Topology Data Kit - 2024-02. <https://www.caida.org/catalog/datasets/internet-topology-data-kit>.
- [13] Censys. 2024. https://search.censys.io/search?resource=hosts&sort=RELEVANCE&per_page=25&virtual_hosts=EXCLUDE&q=
- [14] 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 Conference on Computer Communications*. 1–10. doi:10.1109/INFOCOM42981.2021.9488912
- [15] Jakub Czyz, Matthew J. Luckie, Mark Allman, and Michael Bailey. 2016. Don’t Forget to Lock the Back Door! A Characterization of IPv6 Network Security Policy. In *Network and Distributed System Security Symposium*.
- [16] Markus Dahlmanns, Felix Heidenreich, Johannes Lohmöller, Jan Pennekamp, Klaus Wehrle, and Martin Henze. 2024. Unconsidered Installations: Discovering IoT Deployments in the IPv6 Internet. In *Proceedings of the 2024 IEEE/IFIP Network Operations and Management Symposium (NOMS ’24), May 6-10, 2024, Seoul, Korea*. IEEE.
- [17] Markus Dahlmanns, Johannes Lohmöller, Jan Pennekamp, Jörn Bodenhausen, Klaus Wehrle, and Martin Henze. 2022. Missed opportunities: measuring the untapped TLS support in the industrial internet of things. In *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*. 252–266.
- [18] Markus Dahlmanns, Constantin Sander, Robin Decker, and Klaus Wehrle. 2023. Secrets Revealed in Container Images: An Internet-wide Study on Occurrence and Impact. In *Proceedings of the 2023 ACM Asia Conference on Computer and Communications Security*. ACM, 797–811. doi:10.1145/3579856.3590329
- [19] David Dittrich and Erin Kenneally. 2012. *The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research*. Technical Report. U.S. Department of Homeland Security.
- [20] Zakir Durumeric, David Adrian, Phillip Stephens, Eric Wustrow, and J. Alex Halderman. 2024. Ten Years of ZMap. In *Proceedings of the 2024 ACM on Internet Measurement Conference (Madrid, Spain) (IMC ’24)*. Association for Computing Machinery, New York, NY, USA, 139–148. doi:10.1145/3646547.3689012
- [21] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and Its Security Applications. In *22nd USENIX Security Symposium (USENIX Security 13)*. USENIX Association, Washington, D.C., 605–620. <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/durumeric>
- [22] Tobias Fiebig, Kevin Borgolte, Shuang Hao, Christopher Kruegel, and Giovanni Vigna. 2017. Something from nothing (There): collecting global IPv6 datasets from DNS. In *Passive and Active Measurement: 18th International Conference, PAM*

- 2017, Sydney, NSW, Australia, March 30-31, 2017, *Proceedings 18*. Springer, 30–43.
- [23] Pawel Foremski, David Plonka, and Arthur Berger. 2016. Entropy/IP: Uncovering structure in IPv6 addresses. In *Proceedings of the 2016 Internet Measurement Conference*. 167–181.
- [24] Oliver Gasser, Benjamin Hof, Max Helm, Maciej Korczyński, Ralph Holz, and Georg Carle. 2018. In Log We Trust: Revealing Poor Security Practices with Certificate Transparency Logs and Internet Measurements. In *Passive and Active Measurement Conference 2018*.
- [25] Oliver Gasser, Ralph Holz, and Georg Carle. 2014. A deeper understanding of SSH: Results from Internet-wide scans. In *2014 IEEE Network Operations and Management Symposium (NOMS)*. IEEE, 1–9.
- [26] Oliver Gasser, Quirin Scheitle, Pawel Foremski, Qasim Lone, Maciej Korczyński, Stephen D Strowes, Luuk Hendriks, and Georg Carle. 2018. Clusters in the expanse: Understanding and unbiasing IPv6 hitlists. In *Proceedings of the Internet Measurement Conference 2018*. 364–378.
- [27] Sharon Goldberg, Moni Naor, Dimitrios Papadopoulos, Leonid Reyzin, Sachin Vasantand, and Asaf Ziv. 2015. NSEC5: Provably Preventing DNSSEC Zone Enumeration.
- [28] Fahad Hilal, Patrick Sattler, Kevin Vermeulen, and Oliver Gasser. 2024. A First Look At IPv6 Hypergiant Infrastructure. *Proceedings of the ACM on Networking 2*, CoNEXT2 (2024), 1–25.
- [29] 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. *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications* (2021), 1–10.
- [30] Amanda Hsu, Frank Li, and Paul Pearce. 2023. Fiat lux: Illuminating ipv6 apportionment with different datasets. *Proceedings of the ACM on Measurement and Analysis of Computing Systems 7*, 1 (2023), 1–24.
- [31] Amanda Hsu, Frank Li, Paul Pearce, and Oliver Gasser. 2024. A first look at NAT64 deployment in-the-wild. In *International Conference on Passive and Active Network Measurement*. Springer, 112–129.
- [32] Wayne Jones. 2022. Altcoin Explorer: QLC Chain, the Next Generation Public Chain for Network-as-a-Service (NAAS). <https://crypto.news/altcoin-explorer-qlc-chain-generation-public-chain-network-service-naas/>
- [33] Peter Jose, Said Jawad Saidi, and Oliver Gasser. 2023. Analyzing iot hosts in the ipv6 internet. (2023). arXiv:2307.09918 <https://arxiv.org/abs/2307.09918>
- [34] Xiang Li, Baojun Liu, Xiaofeng Zheng, Haixin Duan, Qi Li, and Youjun Huang. 2021. Fast IPv6 network periphery discovery and security implications. In *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 88–100.
- [35] Aniss Maghsoudlou, Lukas Vermeulen, Ingmar Poese, and Oliver Gasser. 2023. Characterizing the VPN Ecosystem in the Wild. In *International Conference on Passive and Active Network Measurement*. Springer, 18–45.
- [36] Giovane C. M. Moura, Marco Davids, Caspar Schutjser, Cristian Hesselman, John Heidemann, and Georgios Smaragdakis. 2024. Deep Dive into NTP Pool's Popularity and Mapping. *Proc. ACM Meas. Anal. Comput. Syst.* 8, 1, Article 15 (feb 2024), 30 pages. doi:10.1145/3639041
- [37] Austin Murdock, Frank Li, Paul Bramsen, Zakir Durumeric, and Vern Paxson. 2017. Target generation for internet-wide IPv6 scanning. In *Proceedings of the 2017 Internet Measurement Conference*. 242–253.
- [38] Yevheniya Nosyk, Maciej Korczyński, and Andrzej Duda. 2023. Guardians of DNS Integrity: A Remote Method for Identifying DNSSEC Validators Across the Internet. In *2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*. IEEE, 1470–1479.
- [39] Stijn Pletinckx, Thanh-Dat Nguyen, Tobias Fiebig, Christopher Kruegel, and Giovanni Vigna. 2023. Certifiably Vulnerable: Using Certificate Transparency Logs for Target Reconnaissance. In *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*. 817–831. doi:10.1109/EuroSP57164.2023.00053
- [40] NTP Pool project. 2024. pool.ntp.org: NTP Servers in Global, pool.ntp.org. <https://www.ntppool.org/zone/@>
- [41] Qualys. 2024. regreSSHion: RCE in OpenSSH's server, on glibc-based Linux systems (CVE-2024-6387). <https://www.qualys.com/2024/07/01/cve-2024-6387/regression.txt>
- [42] RIPE. 2024. RIPE Routing Information Service. <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/>.
- [43] Justin P Rohrer, Blake LaFever, and Robert Beverly. 2016. Empirical study of router IPv6 interface address distributions. *IEEE Internet Computing 20*, 4 (2016), 36–45.
- [44] Erik Rye and Dave Levin. 2023. IPv6 Hitlists at Scale: Be Careful What You Wish For. In *Proceedings of the ACM SIGCOMM 2023 Conference*. 904–916.
- [45] Erik C Rye and Robert Beverly. 2020. Discovering the IPv6 network periphery. In *Passive and Active Measurement: 21st International Conference, PAM 2020, Eugene, Oregon, USA, March 30–31, 2020, Proceedings 21*. Springer, 3–18.
- [46] Said Jawad Saidi, Srdjan Matic, Oliver Gasser, Georgios Smaragdakis, and Anja Feldmann. 2022. Deep dive into the IoT backend ecosystem. In *Proceedings of the 22nd ACM Internet Measurement Conference (Nice, France) (IMC '22)*. Association for Computing Machinery, New York, NY, USA, 488–503. doi:10.1145/3517745.3561431
- [47] Christian Schwarz, Adam Di Carlo, Raphaël Hertzog, Andreas Barth, Lucas Nussbaum, Hideki Yamane, and Holger Levsen. 2025. Debian Developer's Reference: Managing Packages. <https://www.debian.org/doc/manuals/developers-reference/pkgsg.en.html>
- [48] Tanya Shreedhar, Danesh Zeynali, Oliver Gasser, Nitinder Mohan, and Jörg Ott. 2022. A longitudinal view at the adoption of multipath TCP. (2022). arXiv:2205.12138 <https://arxiv.org/abs/2205.12138>
- [49] Guanglei Song, Jiahai Yang, Lin He, Zhiliang Wang, Guo Li, Chenxin Duan, Yaoyong Liu, and Zhongxiang Sun. 2022. AddrMiner: A Comprehensive Global Active IPv6 Address Discovery System. In *2022 USENIX Annual Technical Conference (USENIX ATC 22)*. USENIX Association, Carlsbad, CA, 309–326. <https://www.usenix.org/conference/atc22/presentation/song>
- [50] Markus Sosnowski, Johannes Zirngibl, Patrick Sattler, Juliane Aulbach, Jonas Lang, and Georg Carle. 2024. An Internet-wide View on HTTPS Certificate Revocations: Observing the Revival of CRLs via Active TLS Scans. In *2024 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 297–306.
- [51] Shreyas Srinivasa, Jens Myrup Pedersen, and Emmanouil Vasilomanolakis. 2021. Open for hire: attack trends and misconfiguration pitfalls of IoT devices. In *Proceedings of the 21st ACM Internet Measurement Conference (IMC'21)*. Association for Computing Machinery, New York, NY, USA, 195–215.
- [52] Lion Steger, Liming Kuang, Johannes Zirngibl, Georg Carle, and Oliver Gasser. 2023. Target Acquired? Evaluating Target Generation Algorithms for IPv6. In *Proceedings of the Network Traffic Measurement and Analysis Conference (TMA)* (Naples, Italy).
- [53] Stephen D. Strowes. 2017. Bootstrapping Active IPv6 Measurement with IPv4 and Public DNS. (2017). arXiv:1710.08536 <https://arxiv.org/abs/1710.08536>
- [54] The ZMap Project. 2024. ZGrab 2.0. <https://github.com/zmap/zgrab2>.
- [55] Matthäus Wander, Lorenz Schwittmann, Christopher Boelmann, and Torben Weis. 2014. GPU-Based NSEC3 Hash Breaking. In *2014 IEEE 13th International Symposium on Network Computing and Applications*. 137–144. doi:10.1109/NCA.2014.27
- [56] Grant Williams and Paul Pearce. 2024. 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*. 295–313.
- [57] Johannes Zirngibl, Lion Steger, Patrick Sattler, Oliver Gasser, and Georg Carle. 2022. Rusty Clusters? Dusting an IPv6 Research Foundation. In *Proceedings of the 2022 Internet Measurement Conference (Nice, France)*. ACM, New York, NY, USA, 15 pages. doi:10.1145/3517745.3561440