TU Delft

Delft University of Technology

Document Version
Final published version

Licence
CC BY

Citation (APA)
Ferrero, D., Sordello, A., Griffioen, H., Drago, I., Smaragdakis, G., & Mellia, M. (2026). Understanding Different Perspectives: Analyzing UDP Scanning Through Reactive Telescopes. In *IMC 2026 - Proceedings of the 2026 ACM Internet Measurement Conference* (pp. 246-255). (Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC). Association for Computing Machinery (ACM). https://doi.org/10.1145/3777912.3839810

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Understanding Different Perspectives: Analyzing UDP Scanning Through Reactive Telescopes

Dario Ferrero
d.ferrero@tudelft.nl
Delft University of Technology
Delft, The Netherlands

Andrea Sordello
andrea.sordello@polito.it
Politecnico di Torino
Torino, Italy

Harm Griffioen
h.j.griffioen@tudelft.nl
Delft University of Technology
Delft, The Netherlands

Idilio Drago
idilio.drago@unito.it
Università di Torino
Torino, Italy

Georgios Smaragdakis
g.smaragdakis@tudelft.nl
Delft University of Technology
Delft, The Netherlands

Marco Mellia
marco.mellia@polito.it
Politecnico di Torino
Torino, Italy

Abstract

Reactive telescopes reveal scanning behavior that passive telescopes cannot observe, but prior deployments have focused exclusively on TCP. We present RT-UDP, the first transport-layer reactive telescope for UDP, which returns empty UDP packets across all ports while applying rate limits to limit its contribution to reflection attacks. We deploy RT-UDP alongside passive telescopes and TCP responders at four university networks in three countries and compare their observations over 30 days. Activating RT-UDP increases observed UDP source addresses by 13.0–43.8% and flows by 11.7–99.6%, although the additional traffic is concentrated among a small subset of sources. It also shifts the application-layer landscape: approximately half of the traffic targets the top 10 destination ports, and services rarely visible in passive telescopes become prominent. Despite carrying no application-layer data, an empty UDP response alters sender behavior: retransmissions increase by 87–602%, different-payload follow-ups by more than 200%, and TCP fallback by over an order of magnitude. Payload analysis further reveals protocol switching and multi-stage probing triggered by interaction. These findings demonstrate that minimal transport-layer responses expose UDP scanning strategies and service ecosystems that remain largely invisible to passive measurement.

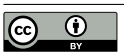
CCS Concepts

• **Security and privacy → Network security**.

Keywords

Network Telescope, Internet Scanning

ACM Reference Format:
Dario Ferrero, Andrea Sordello, Harm Griffioen, Idilio Drago, Georgios Smaragdakis, and Marco Mellia. 2026. Understanding Different Perspectives: Analyzing UDP Scanning Through Reactive Telescopes. In *Proceedings of the 2026 ACM Internet Measurement Conference (IMC '26), October 12–16, 2026, Karlsruhe, Germany*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3777912.3839810



This work is licensed under a Creative Commons Attribution 4.0 International License.
IMC '26, Karlsruhe, Germany
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2327-8/2026/10
https://doi.org/10.1145/3777912.3839810

1 Introduction

Internet scanning is at an all-time high [4]. Such scanning activity is the first step of a chain of actions to enumerate and potentially exploit vulnerable hosts and services on the Internet [22]. Over the years, network telescopes [18] and honeypots [25] have been collecting data to study scanners' behavior and shed light on this threat landscape [5]. Interactive honeypots and, more recently, reactive telescopes have been shown to offer deception, i.e., engage scanners and attackers to deliberately gain insights about their techniques, tactics, and procedures (TTPs).

Despite significant progress in collecting and analyzing scanning and attack traffic in passive [2], reactive [9, 23], and interactive [10] network monitoring infrastructures, the focus has been on Transmission Control Protocol (TCP) traffic. This is to be expected, as studies have shown that the vast majority (80%-90%) of scanning traffic is TCP [7], the rest being almost exclusively User Datagram Protocol (UDP). Yet UDP scan traffic has remained understudied until today. Unlike TCP, UDP is stateless, raising the question of whether reactiveness can elicit further scanning behavior at all.

Being able to leverage and extend existing monitoring infrastructure to better characterize UDP scanning at scale would be an important step towards minimizing the harm of UDP-based attacks [26]. A more fundamental question is whether the stateless nature of UDP makes reactive telescopes irrelevant for the study of UDP scanning. In this paper, we show that even minimal UDP interaction fundamentally changes the observable scanning landscape, revealing services and behaviors that remain largely unobservable in passive telescopes.

Our contributions can be summarized as follows: **(1)** To the best of our knowledge, we implement the first controlled Reactive Telescope based on UDP (RT-UDP), with the capability of responding to both UDP and TCP traffic; **(2)** We deploy RT-UDP at four vantage points across two continents and three countries, comparing transport-protocol responsiveness and response behavior across locations; **(3)** We characterize UDP Internet scanning, analyzing how destination ports, source countries, DDoS reflection attempts, and payloads evolve when responding to incoming UDP traffic.

Monitoring tools, post-processing scripts and instructions to download the anonymized data are available online.[1] The remainder of the paper is organized as follows: Section 2 reviews prior work

[1] https://github.com/DelftThreatIntelligenceLab/rt-udp

on internet measurement methodologies and studies on scanning; Section 3 describes the design and deployment of our UDP Reactive Telescope; Section 4 presents the analysis of UDP traffic observed through our system; and Section 5 concludes the paper and outlines directions for future work.

2 Related work

Network telescopes[2] – advertised but inactive subnets – have been used to measure Internet Background Radiation for two decades [1, 2, 7, 16, 18, 32]. Their strength lies in scale and generalizability across vantage points, but by design they cannot engage scanners. This leaves more complex activity invisible whenever a payload arrives only after a response.

To address this limitation, L4-Responders [29] and Spoki [9] complete the TCP handshake at the transport layer. The design has since been extended to elastic IP space with DScope [23], and used to mine application-layer payloads from active subnets [29]. All transport-layer reactive telescopes deployed so far have responded only to TCP traffic. Whether such interaction is meaningful for UDP, given its stateless nature, remains an open question.

The closest UDP analogues are the amplification honeypots surveyed in [19]. AmpPot [13], in particular, emulates a fixed set of services (including DNS, NTP, SSDP, SIP) to study reflection attacks [8, 14, 15, 26], and supports "agnostic" mode, which allows the deployment of sensors on all UDP ports with randomized payloads. AmpPot is designed to study amplification, thus even in "agnostic" mode, it provides a large service amplification response attractive to certain types of attackers [13]. In contrast, RT-UDP is designed to avoid amplification, operates purely at the transport layer and responds across all UDP ports without emulating specific services. In this way, RT-UDP aims to let the scanner – not the honeypot – guide the observed traffic mix.

Beyond honeypot studies, we are aware of only one prior work analyzing UDP scanning [27]. The authors cluster four months of payloads collected at a single passive sensor. Our reactive setup, deployed across four sites in three countries, complements that view by exposing behaviors that only emerge once a sensor responds, such as protocol switches between probes and multi-stage scans.

3 Methodology

In this section, we present the design of our reactive UDP telescope, describe the countermeasures to prevent involvement in DRDoS attacks, and outline the networks in which our RT-UDP operates.

3.1 RT-UDP Responder Design

Similar to previous works [9, 23], we design a *Reactive Telescope* (RT), i.e., a transport-layer responder that captures all traffic while it can reply on a programmable subset of IPv4 addresses and ports. For TCP traffic, it completes the 3-way handshake and continues to acknowledge data segments it receives. For UDP traffic, it sends back an empty UDP packet.

Sending UDP responses requires caution. In fact, service emulation attracts amplification traffic, which is well covered in prior work [13, 26]. To avoid this, we never send any application data. For UDP, an empty reply is the minimal payload that may activate a previously dark IP address at the transport layer without the risk of amplifying incoming packets. We recognize that, even without a payload, this type of reply might still carry protocol-specific semantic value (e.g., generating application errors), but this approach provides a responsive endpoint to scanners without an explicit protocol fingerprint in the reply.

To prevent contribution to reflection attempts, we additionally implement a rate-limiting mechanism discussed in Section 3.2. We develop the RT-UDP in the Go programming language, using separate native channels for packet processing of each transport protocol. The tool can be configured to react to none, either, or both protocols, i.e., acting as a *Darknet*, *RT-UDP*, *RT-TCP*, or *RT-Both* platform. It is available as open source on GitHub (link in Section 1).

3.2 Rate Limiting and DRDoS Safety

Although we never amplify a single request, a reactive UDP sensor can still be misused for *Distributed Reflective Denial of Service (DRDoS)* attacks. In this scenario, an attacker spoofs the victim's address as the source of a UDP request, so that our reply would land on the victim. To limit our potential contribution to the attack, we apply rate limits with the same rationale as AmpPot [13]. When the incoming packet rate associated with a source exceeds a threshold, we stop responding to it for a blocking period $T = 1$ hour while continuing to log incoming packets. We do not implement a rate-limiting approach for the TCP responder, as the three-way handshake protects reflection attempts [3, 15].

Specifically for UDP, we set three thresholds for a /24 RT: (i) per source IP: 300 ppm, to allow the RT to reply to a single source performing a full block scan of a /24 on one port. (ii) per source /24: 900 ppm, to limit DRDoS attempts that rotate spoofed source addresses across the victim's /24 to bypass per-IP limits. This /24 threshold catches this rotation while still permitting legitimate distributed scans whose sources belong to the same /24; (iii) per destination (IP, port): 1,000 ppm, to cap direct flooding of any single endpoint of the RT-UDP, preventing one heavily-targeted endpoint from dominating the captured scan dataset.

These thresholds are looser than AmpPot's 10 ppm per source IP and /24, and stay well within the budget defined in [8], which caps a reactive sensor's worst-case output at 0.5 Mbps, a fraction of average Internet uplink [21]. With an empty IPv4/UDP reply, that budget corresponds to roughly 90k ppm. In our case, a DRDoS attempt that fully saturates our per-/24 limit would lead us to reply at 900 ppm, i.e. two orders of magnitude below the suggested limit.

As discussed in [19], looser thresholds make the boundary between aggressive scanning and low-rate attack traffic harder to draw. We discuss the effect of this trade-off, and the population of sources we end up blocking, in Section 4.2.

3.3 Deployment and Activation

We deploy RT across four university networks (Table 1): NL_1 at Delft University of Technology (the Netherlands), IT_1 at Politecnico di Torino and IT_2 at Politecnico di Milano (Italy), and BR_1 across the Federal Institute of Espírito Santo (IFES) and the Federal University of Espírito Santo (UFES) (Brazil). The IT_1, IT_2, and BR_1 sensors operate through the Holoscope platform [28]; at each sensor, we partition a /24 network into four consecutive /26 blocks assigned

[2]We use the terms "Network Telescope" and "Darknet" interchangeably.

Table 1: Daily traffic statistics across passive and reactive telescope deployments. Reference values are the 10-day mean ± SD; deployment values are changes from this mean, reported as percentages and Reference SD units (in parentheses).

Sensor	Deployment	Size	UDP IPs	UDP Flows	TCP IPs	TCP Flows
NL_1	*Reference*	/24	10.9k ± 1.8k	162.5k ± 37.3k	28.5k ± 2.2k	2.4M ± 573.9k
	Darknet	/24	+18.2% ± (+1.1)	+6.0% ± (+0.3)	+1.7% ± (+0.2)	-21.2% ± (-0.9)
	RT-UDP	/24	+13.0% ± (+0.8)	+75.4% ± (+3.3)	+20.6% ± (+2.6)	-42.1% ± (-1.8)
	RT-Both	/24	+41.2% ± (+2.6)	+49.9% ± (+2.2)	+100.9% ± (+12.9)	+151.1% ± (+6.3)
IT_1	*Reference*	/26	5.5k ± 760	30k ± 3.9k	20.4k ± 901	749k ± 255.4k
	Darknet	/26	+11.1% (+0.8)	-5.4% (-0.4)	-4.7% (-1.1)	-38.6% (-1.1)
	RT-UDP	/26	+21.5% (+1.5)	+11.7% (+0.9)	-1.0% (-0.2)	-36.7% (-1.1)
	RT-TCP	/26	+20.2% (+1.5)	+7.6% (+0.6)	+50.6% (+11.5)	+280.0% (+8.2)
	RT-Both	/26	+29.1% (+2.1)	+39.6% (+3.1)	+49.7% (+11.2)	+282.8% (+8.3)
IT_2	*Reference*	/26	5k ± 844	23.7k ± 2.7k	18.5k ± 881	496.5k ± 249.6k
	Darknet	/26	+22.8% (+1.4)	+14.1% (+1.2)	+2.0% (+0.4)	-32.0% (-0.6)
	RT-UDP	/26	+33.2% (+2.0)	+40.2% (+3.5)	+3.8% (+0.8)	-29.6% (-0.6)
	RT-TCP	/26	+33.2% (+2.0)	+30.0% (+2.6)	+48.8% (+10.3)	+153.6% (+3.1)
	RT-Both	/26	+33.8% (+2.0)	+41.3% (+3.6)	+59.7% (+12.5)	+154.4% (+3.1)
BR_1	*Reference*	/26	4.1k ± 521	24.3k ± 4.1k	17.7k ± 702	744k ± 340.8k
	Darknet	/26	+20.0% (+1.6)	-0.8% (0.0)	-2.8% (-0.7)	-9.5% (-0.2)
	RT-UDP	/26	+43.8% (+3.5)	+99.6% (+5.8)	+6.8% (+1.7)	+18.1% (+0.4)
	RT-TCP	/26	+39.3% (+3.1)	+93.0% (+5.5)	+65.2% (+16.5)	+281.4% (+6.1)
	RT-Both	/26	+51.3% (+4.1)	+93.0% (+5.5)	+57.1% (+14.4)	+248.5% (+5.4)

*We exclude traffic from IPs involved in DRDoS attempts (see Section 4.2)

to RT-UDP, RT-TCP, RT-Both, and the Passive Telescope (Darknet). At NL_1, instead, each deployment occupies a non-contiguous /24 subnet, and the Darknet is part of a larger passive /16 telescope. Moreover, we used NL_1 for preliminary RT-UDP tests before the stable data-collection period. In Section 4.1, we denote the complete /24 subnets as NL_1 and sampled /26 subnet as NL_1^* to enable fair comparisons with the other /26 deployments. Overall, this configuration provides a controlled setting for assessing the relative effects of "activating" dark IP addresses with different protocol combinations.

We started our collection on March 19^{th}, 2026. For the first 10 days, all deployments but the NL_1 RTs were darknets (*Reference* period). On March 29^{th}, we activated the IT_1, IT_2, BR_1 RTs and consider data collected until April 19^{th}, for a total of 10+20 days.

Table 1 reports daily traffic flow and source-IP statistics for each deployment relative to the Reference period. Reference values are the daily mean and standard deviation (± SD) over the first 10 days; for NL_1, they are computed from its /24 Darknet over the same interval. For subsequent deployments, we report the percentage change in the mean over the 20-day capture and, in parentheses, the shift in units of the Reference SD. Values within ±1 SD are comparable to the Reference period day-to-day variability, whereas larger shifts indicate changes exceeding this baseline variability.

Activating RT-UDP increases UDP senders by 13.0%–43.8% and UDP flows by 11.7%–99.6% across all sensors. Relative to Reference variability, IT_1 shows modest shifts (+1.5 SD for UDP IPs and +0.9 SD for UDP flows), whereas BR_1 shows substantially stronger responses (+3.5 and +5.8 SD, respectively). As expected, TCP statistics are only marginally affected, with shifts between −1.1 and +1.7 SD.

With RT-TCP, senders increase by 48.8%–65.2% across all sensors, corresponding to pronounced shifts of +10.3 to +16.5 SD. TCP flows increase by 153.6%–281.4%, but their greater day-to-day variability limits the corresponding shifts between +3.1 and +8.2 SD.

RT-Both preserves and sometimes amplifies the effects of activating each protocol individually: UDP and TCP senders increase by

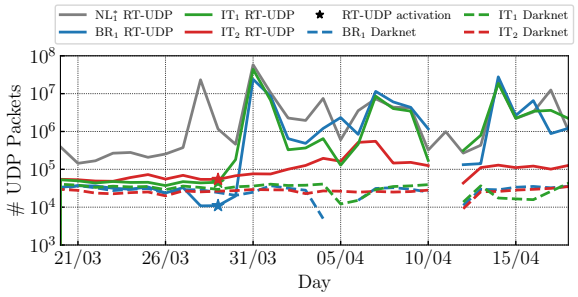


Figure 1: Daily packet counts before and after replying to UDP traffic. Each ★ marks an RT-UDP activation. A short outage is visible in some sensors.

29.1%–51.3% and 49.7%–100.9%, respectively, while UDP and TCP flows increase by 39.6%–93.0% and 151.1%–282.8%. Thus, both reactive mechanisms remain visible when operating simultaneously.

For comparison, the 20-day Darknet statistics characterize variation without reactive behaviour. Relative to the Reference period, sender counts range from −1.1 to +1.6 SD and flow counts from −1.1 to +1.2 SD. These results show that shifts on the order of 1 SD can occur even in the absence of reactive behaviour, providing useful context for the substantially larger shifts observed under the reactive deployments. Finally, because the IT_1, IT_2, and BR_1 deployments occupy contiguous address blocks, horizontal scanners attracted by a reactive subnet may continue into neighboring subnets and affect their traffic; measurements from these deployments are therefore not entirely independent.

Takeaway. Activating RT-UDP increases UDP senders by 13.0–43.8% and UDP flows by 11.7–99.6%, while having negligible impact on TCP traffic. RT-TCP produces a more pronounced increase in TCP senders and flows, consistent with previous work [9, 23, 29]. To the best of our knowledge, this is the first measurement of the traffic impact induced by an RT-UDP deployment.

4 RT-UDP Data Characterization

We report the results of deploying our UDP Reactive Telescope across multiple sensors over a month. Section 4.1 quantifies the activation effect of UDP replying. Section 4.2 characterizes rate-limited sources, and Section 4.3 analyses UDP scanning behavior, including targeting, origin, and payload.

4.1 Traffic Changes Induced by Activation

Figure 1 shows how UDP traffic suddenly increases sharply immediately after activating the RT-UDP (solid lines), and stays persistently above the darknet baseline (dotted lines). Note the y-log scale, with peaks reaching up to 1000× the baseline volume for IT_1 and BR_1. Not all of this volume is scanning: days after activation, IT_1 and BR_1 are picked up by the same DRDoS campaign that NL_1 had been receiving, with peaks of 10^6–10^7 packets per day. These source addresses – which may also be spoofed victims – trigger the rate-limiting, and we investigate them in Section 4.2. They are excluded from the following analysis to avoid bias.

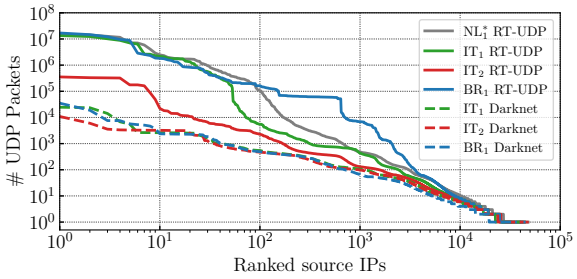
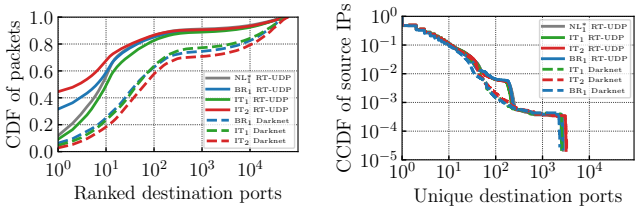


Figure 2: UDP packets received per-source for 20 days from 29^{th} March. Rate-limited IPs are excluded (see Section 4.2).



(a) Packets across ranked destination ports.

(b) Unique destination ports per source IP address.

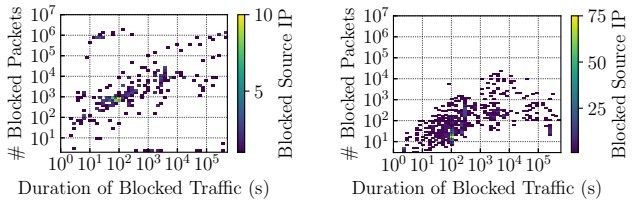
Figure 3: Distribution of UDP traffic across destination ports. Rate-limited source addresses are excluded.

Figure 2 compares the packet volumes generated by individual source IPs across the Darknet and RT-UDP deployments. In the Darknet, the most active sources generate approximately 10^4 packets, whereas under RT-UDP their activity reaches up to 10^7 packets, an increase of three orders of magnitude. However, this increase is concentrated among no more than 4k of the 40k observed source IPs, while the remaining ones exhibit comparable traffic volumes across the two deployments. These results indicate that RT-UDP does not uniformly trigger additional traffic across all sources; rather, the increase is driven by a relatively small subset of highly active addresses.

Activation affects both the distribution of traffic across ports and the destination-port fan-out of individual sources, as observable in Figures 3a and 3b.[3] In the Darknet baseline, the top 100 ports account for more than 50% of UDP traffic, consistent with prior TCP observations [7]. SIP (5060), DNS (53), SNMP (161), and IKE (500) are among the most frequently contacted services, in line with previous studies [26, 27]. This pattern is consistent across the IT_1, IT_2, and BR_1 darknets, with only minor regional variation.

After activation, traffic becomes substantially more concentrated: the top 10 ports alone account for roughly 50% of the observed traffic. The composition of these ports also changes markedly, with ports 8080–8083 (alternative HTTP), 11211 (Memcached), and 1434 (SQL Server) emerging among the most targeted. Activation therefore not only concentrates scanner traffic on a smaller set of ports, but also exposes probing toward services that receive little traffic in the Darknet. We further investigate these effects in Section 4.4.

[3]We omit the baseline for NL_1, since Darknet and RT-UDP are in separate /16 network, thus its deployment differs from the other sensors.



(a) Blocked traffic due to *Source-IP* rate limiting.

(b) Blocked traffic due to *Endpoint* rate limiting.

Figure 4: Distribution of blocked traffic at NL_1.

Despite this shift in aggregate traffic, the destination-port fan-out of individual source IP addresses changes more moderately. Figure 3b shows that the RT-UDP and Darknet curves remain close across most of the distribution, indicating broadly a similar number of destination ports contacted. A separation nevertheless emerges around 100 contacted destination ports: while only about 0.1% of sources reach this fan-out in the Darknet, the fraction increases to nearly 1% under RT-UDP. Taken together, these results suggest that RT-UDP primarily redistributes and concentrates traffic toward specific services, while only a small fraction of sources expands the range of ports they probe.

Takeaway. Activating UDP responders increases traffic by up to 1000×. After excluding attempted DRDoS sources, however, only one in ten source IP addresses sends more packets; the top 10 ports account for about 50% of traffic; rarely observed services emerge, such as Memcached and SQL Server; and most scanners retain similar port fan-out.

4.2 Blocked Source IP Addresses

Section 4.1 showed that the post-activation traffic is contaminated by DRDoS, motivating a closer look at the sources we rate-limit. Recall that these source IP addresses, besides aggressive scanners, may represent the spoofed addresses of the victims of DRDoS. Across the RT-UDP and RT-Both deployments, since March 29th, the rate limiter has blocked 4k source IP addresses, representing 1.6% of all observed sources responsible for 1557M packets – about 96% of the incoming UDP packets.

Figure 4 shows, for NL_1, the average number of blocked packets and the duration of each blackhole period. We split the analysis by the cause of rate limiting: Figure 4a focuses on sources blocked for exceeding the per-source-IP rate. Two clusters emerge: (i) About 30 IPs in the upper-left of the plot each send more than 10^6 packets within a few tens of seconds. These predominantly contact ports 389 (LDAP) and 11211 (Memcached) and originate from different ASes. The pattern – short, intense bursts – is consistent with DRDoS activity; (ii) A larger cluster sits in the central and lower regions: moderate packet counts that cease after a few hundred seconds. These appear to be aggressive scanners that briefly trigger our rate limiter but do not sustain DRDoS-level volume.

Figure 4b reports the sources blocked due to exceeding the per-endpoint rate limit in NL_1, which is triggered when a specific (IP, port) endpoint is contacted too aggressively. The activity is dominated by a single destination IP on port 37810 (DVR DHCPDiscover),

with most sources originating from ASes in Indonesia. This pattern is consistent with known DVR DHCPDiscover reflection/amplification activity, where large populations of IoT devices (e.g., CCTV/DVR systems) respond to unsolicited UDP probes and are leveraged in coordinated scanning or DDoS campaigns [20]. We also report, for the sake of completeness, that in NL_1 the per-source-network rate limiting never gets triggered.

Takeaway. Our rate limiter identifies DRDoS-scale bursts and aggressive scanners through per-source thresholds, and coordinated endpoint targeting through per-endpoint thresholds.

4.3 Origin of UDP Scanners

We now examine the countries associated with UDP source IP addresses. To reduce bias from high-volume attacks, we exclude sources associated with DRDoS and known institutional scanners identified through the MISP warning lists [17, 31]. As institutional scanners generate approximately 30–40% of daily UDP traffic, this filtering focuses our analysis on less-explored traffic.

Table 2 reports the top countries associated with UDP source IPs, geolocated using IPInfo [11] during our 20-day collection campaign. For each sensor, we report absolute counts and the RT-UDP change from its darknet, $(N_{\text{RT-UDP}} - N_{\text{Darknet}})/N_{\text{Darknet}} \times 100$, highlighting increases of at least 10% in bold to emphasize countries for which RT-UDP exposes a substantially larger source population.

IP geolocation alone cannot account for source-address spoofing, and our measurements do not provide sufficient information to reliably distinguish legitimate senders from spoofed victims. This limitation is relevant for RT-UDP. While spoofing provides little benefit when targeting a non-responsive darknet for reflection, a responsive sensor can be discovered and exploited as a reflector, causing spoofed victim addresses to appear among the sources. Conversely, RT-UDP may also expose legitimate scanners that become visible only after the sensor replies to their probes. Therefore, an increase in source IPs under RT-UDP may reflect both mechanisms, and we interpret it as additional *observed source addresses*, rather than directly as additional scanners or spoofed victims.

The darknet deployments show a consistent distribution across IT_1, IT_2, and BR_1, dominated by Brazil (6.9k–8.0k unique source IPs), China (5.3k–6.2k), and the United States (3.2k–4.1k). The prominence of China and the United States aligns with their established role as major sources of scanning activity [9], while our measurements also identify Brazil as a major origin of UDP source addresses. RT-UDP preserves this ranking but affects countries differently. China shows the largest consistent increase, reaching 7.3k–9.0k source IPs (+36.9%–44.4%), followed by the United States at 4.4k–4.8k (+15.9%–36.1%). Hong Kong also consistently increases by 24.8%–34.1%, despite lower absolute counts. In contrast, Brazil remains nearly unchanged, while most other countries vary by less than our 10% highlighting threshold.

Thus, RT-UDP does not produce a uniform increase across countries: it exposes substantially more addresses from China, the United States, and Hong Kong, while large populations such as Brazil remain stable. This uneven increase is consistent with a mixture of response-induced scanners and addresses associated with spoofed reflection traffic, although our measurements cannot distinguish the two phenomena.

Table 2: Geographic distribution of source IPs.

Country	Source IPs targeting Darknet			Source IPs targeting RT-UDP			
	IT_1 (36.7k IPs)	IT_2 (39.4k IPs)	BR_1 (33.6k IPs)	NL_1 (84.1k IPs)	IT_1 (39.5k IPs)	IT_2 (43.0k IPs)	BR_1 (37.7k IPs)
China	6.1k	6.2k	5.3k	21.7k	**8.3k (+36.9%)**	**9.0k (+44.4%)**	**7.3k (+38.6%)**
Brazil	6.9k	8.0k	6.9k	14.0k	6.9k (+0.0%)	8.0k (-0.6%)	6.9k (+0.8%)
USA	4.0k	4.1k	3.2k	6.9k	**4.8k (+18.6%)**	**4.8k (+15.9%)**	**4.4k (+36.1%)**
India	2.0k	2.2k	2.1k	6.6k	1.9k (-5.4%)	2.2k (+2.4%)	2.0k (-5.6%)
UK	1.4k	1.4k	0.8k	2.1k	1.3k (-2.8%)	1.4k (-2.1%)	0.9k (+4.3%)
Mexico	1.0k	1.2k	1.1k	2.3k	1.0k (-0.6%)	1.2k (+1.3%)	1.1k (-2.5%)
South Africa	0.7k	0.8k	0.7k	1.4k	0.7k (-0.9%)	0.8k (+1.2%)	0.7k (+4.1%)
Argentina	0.7k	0.8k	0.6k	1.5k	0.7k (+1.5%)	0.8k (-1.4%)	0.7k (+3.0%)
Malaysia	0.6k	0.6k	0.5k	1.3k	0.6k (+7.7%)	0.6k (+0.2%)	0.6k (+7.3%)
Indonesia	0.6k	0.6k	0.5k	1.2k	0.5k (-3.8%)	0.6k (-0.3%)	**0.6k (+11.4%)**
Singapore	0.7k	0.4k	0.6k	1.0k	0.7k (+3.3%)	0.4k (-0.2%)	**0.7k (+15.2%)**
Hong Kong	0.3k	0.3k	0.5k	0.5k	**0.4k (+34.1%)**	**0.4k (+25.9%)**	**0.7k (+24.8%)**

*We exclude traffic from NL_1 Darknet as it is not co-located in the same network as RT-UDP

Takeaway. RT-UDP preserves the geographic ranking but reveals disproportionately more source IPs from China, the United States, and Hong Kong, potentially including spoofed victims rather than only additional scanners.

4.4 UDP Application-Layer Scans

Hiesgen et al. [9] showed that TCP-based Reactive Telescopes primarily attract HTTP, SSH, TLS, and ADB requests, whereas UDP scanning has mainly been studied in the context of DRDoS [13, 26] and amplification-vulnerable services.

Table 3 compares the top 10 protocols and services detected by the NL_1 Darknet and RT-UDP; results for the other sensors appear in Appendix Tables 7 - 9, while RT-Both is omitted as it yields similar results. We manually label 46 protocol and notable-payload classes using regular expressions derived from recurring byte patterns in payload prefixes and bodies (Table 10). This conservative approach reduces false positives but leaves 50% of Darknet traffic unlabeled. Classification is particularly difficult for encrypted, binary, proprietary, poorly documented, TLV-encoded, malformed, or handcrafted messages; our results therefore represent a lower bound on confidently classified traffic.

RT-UDP reveals a substantially different application-layer landscape: SIP traffic grows from 329k to 1.3M packets, while most other Darknet protocols remain stable. Moreover, several protocols absent from the Darknet top 10 enter the RT-UDP ranking (underlined in Table 3). Most notably, Mozi Botnet [30] traffic generates 5.8M packets from only 2,010 sources using six distinct payloads, all containing the Mozi XOR encrypted configuration. Dahua DHIP enters the top 10 alongside Memcached, a protocol known to enable UDP DRDoS attacks [6, 26].

Prior work documented TCP scanning for protocols on non-standard ports [29] and services deployed on unexpected transport ports [12]. We similarly observe UDP sources probing services across multiple ports, although activation affects their port coverage differently: SIP expands from 373 distinct ports in the Darknet to 1.1k in RT-UDP, whereas NTP contracts from 16k to 4.4k. This may indicate that an empty reply either triggers deeper probing or is interpreted as a valid response, reducing broader port scanning.

4.5 Follow-ups to UDP replies

We examine how UDP senders react after RT-UDP replies, including retransmissions, TCP fallback, and different-payload follow-ups that may indicate further reconnaissance or adaptive scanning.

Table 3: Top 10 Application protocols across the /24 within NL_1. We underline protocols that appear after replying.

Sensor	Label	# Packets	# Payloads	# Flows	# Sources	# Ports
Darknet	*Unknown*	1,892,404 (50.61%)	1,497,566	1,737,368	72,179	65,536
	SIP	329,057 (8.8%)	306,626	284,711	5,063	373
	DNS	228,855 (6.12%)	34,515	193,020	5,999	268
	DTLS	156,836 (4.19%)	9,742	155,905	1,337	286
	Valve SEQ	144,171 (3.86%)	3	142,740	2,892	53
	NTP	129,620 (3.47%)	29,445	125,285	10,418	16,334
	RakNet	110,320 (2.95%)	58,233	109,897	18,749	52,136
	SSDP	76,520 (2.05%)	22	75,986	2,606	251
	CoAP	62,882 (1.68%)	1,664	62,119	2,673	1,355
	IKE	51,832 (1.39%)	18,332	45,717	7,501	8,889
	BitTorrent	48,342 (1.29%)	10,585	42,968	6,045	234
RT-UDP	Mozi Bot Config.	5,829,627 (57.06%)	6	2,230	2,010	15
	SIP	1,322,309 (12.94%)	1,301,265	1,183,578	6,808	1,166
	Unknown	1,080,514 (10.58%)	628,060	963,341	60,908	64,669
	Dahua DHIP	396,013 (3.88%)	4	163,763	1,914	228
	DNS	229,100 (2.24%)	61,602	202,203	7,212	230
	Memcached	149,275 (1.46%)	2,976	17,166	235	2
	DTLS	141,479 (1.38%)	9,095	141,472	1,136	260
	NTP	125,637 (1.23%)	16,857	119,198	7,104	4,410
	Valve SEQ	123,409 (1.21%)	3	123,326	2,813	47
	RakNet	92,763 (0.91%)	50,593	92,737	18,296	48,150
	SSDP	71,749 (0.7%)	31	71,744	3,215	250

* We exclude sources associated with DRDoS (Section 4.2) and known scanners from MISP [17, 31]

Table 4: Follow-up behavior of UDP senders observed during the 20-day RT-UDP activation period.

Sensor	Type	# UDP Pseudo-flow*	Single Packet	2nd Packet		
				Retransmission	TCP Fallback	Different Payload
IT_1	Darknet	558*k*	552*k* (98.86%)	1.1*k* (0.20%)	20 (0.00%)	5.3*k* (0.94%)
	RT-UDP	597*k*	566*k* (94.81%)	7.8*k* (1.31%)	2.4*k* (0.41%)	21*k* (3.47%)
	$\Delta_{\text{RT-UDP}}$		-4.1%	+567.6%	+11,302.9%	+268.0%
IT_2	Darknet	542*k*	536*k* (99.06%)	941 (0.17%)	13 (0.00%)	4.1*k* (0.76%)
	RT-UDP	574*k*	549*k* (95.61%)	7.0*k* (1.22%)	1.7*k* (0.30%)	16*k* (2.87%)
	$\Delta_{\text{RT-UDP}}$		-3.5%	+602.3%	+12,348.2%	+277.7%
BR_1	Darknet	452*k*	446*k* (98.63%)	1.4*k* (0.31%)	151 (0.03%)	4.7*k* (1.03%)
	RT-UDP	595*k*	569*k* (95.58%)	3.4*k* (0.57%)	3.0*k* (0.51%)	20*k* (3.33%)
	$\Delta_{\text{RT-UDP}}$		-3.1%	+87.6%	+1,430.4%	+222.2%

* We exclude sources associated with DRDoS (Section 4.2) and known scanners from MISP [17, 31]

Table 4 compares sender behavior in the Darknet and RT-UDP. In the Darknet, 98% of pseudo-flows (`src_ip`, `dst_ip`, `dst_port`)[4] contain a single packet. Retransmissions of the same payload account for only 0.20%, approximately 1% continue with a different payload, and TCP fallback to the same destination is nearly absent. RT-UDP markedly changes this behavior. We compute the relative change from the Darknet as $\Delta_{\text{RT-UDP}} = (\text{RT-UDP}_{\%} - \text{Darknet}_{\%})/\text{Darknet}_{\%} \times 100$. Although single-packet pseudo-flows remain dominant, their fraction decreases by 3–4%. Retransmissions increase by 87–602%, different-payload follow-ups by approximately 250%, and TCP fallbacks by 1,400–12,000% from their near-zero Darknet baseline. Thus, even an empty reply alters a small fraction of pseudo-flows, inducing retransmissions, new payloads, and TCP connection attempts.

Reply-induced traffic. To isolate traffic likely induced by our replies, we select sources that send exactly one packet to a destination port in the Darknet and match their pseudo-flows (`src_ip`, `dst_port`, 1^{st}-Payload) with RT-UDP traffic within one hour. We then check for a 2^{nd} packet and inspect its payload. Because these sources generate no Darknet follow-up, an RT-UDP follow-up can be attributed to our reply with greater confidence.

We analyze IT_1, IT_2, and BR_1, whose deployments share the same /24 space, allowing us to compare a scanning campaign contacting the Darknet (1^{st} payload only) and RT-UDP (1^{st} and 2^{nd} payloads). The one-hour window balances shorter intervals, which favor aggressive senders, against longer ones, which may combine distinct campaigns. Appendix Figure 5 evaluates alternative windows.

[4]A pseudo-flow terminates after five minutes without additional packets.

Table 5: Notable examples of different follow up UDP packets with respect to the initial one.

Destination Port	# IP Source	# IP Darknet	# IP RT-UDP	1st Packet	2nd Packet
Different Payload Version					
53 (DNS)	9	127	81	DNS (EDNS) Query	EDNS (DNS) Query
5632 (pcAnywhere)	36	43	29	"ST"(or "NQ") String	"NQ"(or "ST") String
123 (NTP)	4	111	76	48B GET monlist	8B GET monlist
69 (TFTP)	111	59	56	Read file "/a" as "netascii"	Read file "file" as "octet"
Protocol Switch					
5060 (SIP)	1	8	1	SIP INVITE	QUIC Handshake
1701 (L2TP)	3	17	8	MikroTik discovery	BitTorrent DHT ping
5683 (CoAP)	2	26	10	GET /.well-known/core	BitTorrent DHT Ping
Two-stage probing					
53(DNS)	1	8	1	HEX encoded DNS Query	Valid DNS Query
443 (QUIC)	1	7	7	QUIC with "SCANNING" String	QUIC Handshake
9034	1	38	23	Minimal HTTP GET	Injection shell downloader

Across 20 days, we match 140k pseudo-flows between the Darknet and RT-UDP and identify three behaviors: no follow-up (92%), an identical second payload (3.5%), and a different second payload (2.5%), the latter suggesting more advanced scanning. Manual inspection of representative sequences in Table 5 reveals increasing sophistication: within-protocol payload variation (e.g., DNS or NTP), protocol switching (e.g., BitTorrent messages), and two-stage probing in which benign initial packets trigger targeted or malicious follow-ups (e.g., shell injection).

Takeaway. RT-UDP enables the observation of previously unexplored service scans, as well as reply-triggered follow-up traffic, revealing scanner strategies that only emerge upon interaction.

5 Conclusion and Future Work

Reactive telescopes have been used to study TCP scanning, but their applicability to UDP remained unclear. We show that even minimal interaction extends to UDP and fundamentally changes the observable scanning landscape. Across four networks in three countries, RT-UDP increases observed UDP traffic by up to three orders of magnitude during peak bursts. After removing the senders performing DRDoS amplification attempts, RT-UDP influences the top services requested while exposing adaptive scanner behavior – payload variation, protocol switching, and multi-stage probing – that is invisible without interaction. Several directions remain open. The application-layer payloads we observe point to under-explored scanning ecosystems (e.g. DTLS, IKE), each warranting deeper characterization. The sequences of payloads observed after interaction also enable fingerprinting of scanning campaigns beyond individual source addresses. To support this line of work, we release RT-UDP and anonymized data to the community.

Acknowledgments

This research was supported by the Dutch Research Council (NWO) under the ADAPTIve project, and by the European Commission under the Horizon Europe Programme as part of the projects Safe-Horizon (Grant Agreement #101168562) and RECITALS (Grant Agreement #101168490). The authors would like to thank Francesco Musumeci, Massimo Tornatore, and Sebastian Troia from Politecnico di Milano for their assistance in operating the IT_2 sensor, and Willen Borges Coelho for operating the BR_1 sensor.

References

[1] Michael D Bailey, Evan Cooke, Farnam Jahanian, Jose Nazario, and David Watson. 2005. The Internet Motion Sensor - A Distributed Blackhole Monitoring System. In *Proceedings 2005 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA.

[2] 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*. ACM, Tokyo Japan, 423–436. doi:10.1145/2815675.2815702

[3] Kevin Bock, Abdulrahman Alaraj, Yair Fax, Kyle Hurley, Eric Wustrow, and Dave Levin. 2021. Weaponizing Middleboxes for TCP Reflected Amplification. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 3345–3361. https://www.usenix.org/conference/usenixsecurity21/presentation/bock

[4] 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

[5] Zakir Durumeric, Michael Bailey, and J. Alex Halderman. 2014. An Internet-Wide View of Internet-Wide Scanning. In *23rd USENIX Security Symposium (USENIX Security 14)*. USENIX Association, San Diego, CA, 65–78. https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/durumeric

[6] Vincent Ghiëtte and Christian Doerr. 2018. How Media Reports Trigger Copycats: An Analysis of the Brewing of the Largest Packet Storm to Date. In *Proceedings of the 2018 Workshop on Traffic Measurements for Cybersecurity* (Budapest, Hungary) *(WTMC '18)*. Association for Computing Machinery, New York, NY, USA, 8–13. doi:10.1145/3229598.3229606

[7] Harm Griffioen, Georgios Koursiounis, Georgios Smaragdakis, and Christian Doerr. 2024. Have you SYN me? Characterizing Ten Years of Internet Scanning. In *Proceedings of the 2024 ACM on Internet Measurement Conference* (Madrid, Spain) *(IMC '24)*. Association for Computing Machinery, New York, NY, USA, 149–164. doi:10.1145/3646547.3688409

[8] Harm Griffioen, Kris Oosthoek, Paul van der Knaap, and Christian Doerr. 2021. Scan, Test, Execute: Adversarial Tactics in Amplification DDoS Attacks. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security* (Virtual Event, Republic of Korea) *(CCS '21)*. Association for Computing Machinery, New York, NY, USA, 940–954. doi:10.1145/3460120.3484747

[9] 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 *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 431–448. https://www.usenix.org/conference/usenixsecurity22/presentation/hiesgen

[10] Niclas Ilg, Paul Duplys, Dominik Sisejkovic, and Michael Menth. 2023. A survey of contemporary open-source honeypots, frameworks, and tools. *Journal of Network and Computer Applications* 220 (2023), 103737. doi:10.1016/j.jnca.2023.103737

[11] IPInfo. 2026. *IPInfo Lite Dataset*. https://ipinfo.io/lite Accessed at 2026-03-01.

[12] Liz Izhikevich, Renata Teixeira, and Zakir Durumeric. 2021. LZR: Identifying Unexpected Internet Services. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 3111–3128. https://www.usenix.org/conference/usenixsecurity21/presentation/izhikevich

[13] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. 2015. AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. In *Research in Attacks, Intrusions, and Defenses*, Herbert Bos, Fabian Monrose, and Gregory Blanc (Eds.). Springer International Publishing, Cham, 615–636. doi:10.1007/978-3-319-26362-5_28

[14] Johannes Krupp, Michael Backes, and Christian Rossow. 2016. Identifying the Scan and Attack Infrastructures Behind Amplification DDoS Attacks. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, Vienna Austria, 1426–1437. doi:10.1145/2976749.2978293

[15] Marc Kührer, Thomas Hupperich, Christian Rossow, and Thorsten Holz. 2014. Hell of a Handshake: Abusing TCP for Reflective Amplification DDoS Attacks. In *8th USENIX Workshop on Offensive Technologies (WOOT 14)*. USENIX Association, San Diego, CA. https://www.usenix.org/conference/woot14/workshop-program/presentation/kuhrer

[16] Alexander Männel, Jonas Mücke, K. C. 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 *Proceedings of the ACM SIGCOMM 2025 Conference* (São Francisco Convent, Coimbra, Portugal) *(SIGCOMM '25)*. Association for Computing Machinery, New York, NY, USA, 826–841. doi:10.1145/3718958.3754347

[17] MISP. 2026. *misp-warninglists*. https://misp.github.io/misp-warninglists/ Accessed at 2026-03-01.

[18] David Moore, Colleen Shannon, Geoffrey M Voelker, and Stefan Savage. 2004. Network Telescopes: Technical Report. (2004). https://escholarship.org/uc/item/1405b1bz

[19] Marcin Nawrocki, John Kristoff, Raphael Hiesgen, Chris Kanich, Thomas C. Schmidt, and Matthias Wählisch. 2023. SoK: A Data-driven View on Methods to Detect Reflective Amplification DDoS Attacks Using Honeypots. In *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*. 576–591. doi:10.1109/EuroSP57164.2023.00041

[20] NETSCOUT ASERT. 2021. DHCPDiscover Reflection/Amplification DDoS Attack Mitigation Recommendations. https://www.netscout.com/blog/asert/dhcpdiscover-reflectionamplification-ddos-attack-mitigation.

[21] Ookla. 2026. *Speedtest Global Index*. https://www.speedtest.net/global-index Accessed at 2026-03-01.

[22] Eric Pauley, Paul Barford, and Patrick McDaniel. 2023. The CVE Wayback Machine: Measuring Coordinated Disclosure from Exploits against Two Years of Zero-Days. In *Proceedings of the 2023 ACM on Internet Measurement Conference* (Montreal QC, Canada) *(IMC '23)*. Association for Computing Machinery, New York, NY, USA, 236–252. doi:10.1145/3618257.3624810

[23] Eric Pauley, Paul Barford, and Patrick McDaniel. 2023. DScope: A Cloud-Native Internet Telescope. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 5989–6006. https://www.usenix.org/conference/usenixsecurity23/presentation/pauley

[24] Eric Pauley and Patrick McDaniel. 2023. Understanding the Ethical Frameworks of Internet Measurement Studies. In *The 2nd International Workshop on Ethics in Computer Security (EthiCS 2023)*. San Diego, CA. doi:10.14722/ethics.2023.239547

[25] Niels Provos et al. 2004. A Virtual Honeypot Framework. In *USENIX Security Symposium*, Vol. 173. 1–14.

[26] Christian Rossow. 2014. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Proceedings 2014 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA. doi:10.14722/ndss.2014.23233

[27] Sina Rostami, Taha Albakour, and Tiago Heinrich. 2025. State of UDP Scanners on the Internet. In *2025 9th Network Traffic Measurement and Analysis Conference (TMA)*. 1–4. doi:10.23919/TMA66427.2025.11097012

[28] Andrea Sordello, Marco Mellia, Idilio Drago, Rodolfo Valentim, Francesco Musumeci, Massimo Tornatore, Federico Cerutti, Martino Trevisan, Alessio Botta, and Willen B. Coelho. 2026. Holoscope: Open and Lightweight Telescope & Honeypot Platform. *IEEE Communications Magazine* (2026), 1–6. doi:10.1109/MCOM.001.2500784

[29] Francesca Soro, Thomas Favale, Danilo Giordano, Idilio Drago, Tommaso Rescio, Marco Mellia, Zied Ben Houidi, and Dario Rossi. 2023. Enlightening the Darknets: Augmenting Darknet Visibility With Active Probes. *IEEE Transactions on Network and Service Management* 20, 4 (2023), 5012–5025. doi:10.1109/TNSM.2023.3267671

[30] Teng-Fei Tu, Jia-Wei Qin, Hua Zhang, Miao Chen, Tong Xu, and Yue Huang. 2022. A comprehensive study of Mozi botnet. *International Journal of Intelligent Systems* 37, 10 (2022), 6877–6908. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/int.22866 doi:10.1002/int.22866

[31] Cynthia Wagner, Alexandre Dulaunoy, Gérard Wagener, and Andras Iklody. 2016. MISP: The Design and Implementation of a Collaborative Threat Intelligence Sharing Platform. In *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*. ACM, 49–56.

[32] Eric Wustrow, Manish Karir, Michael Bailey, Farnam Jahanian, and Geoff Huston. 2010. Internet Background Radiation Revisited. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*. ACM, Melbourne Australia, 62–74. doi:10.1145/1879141.1879149

[33] Jun Xu, Jinliang Fan, M.H. Ammar, and S.B. Moon. 2002. Prefix-preserving IP Address Anonymization: Measurement-based Security Evaluation and a New Cryptography-based Scheme. In *10th IEEE International Conference on Network Protocols, 2002. Proceedings.* 280–289.

Appendix

Ethics

The primary ethical concern in our work is the potential for our infrastructure to be exploited in reflective DDoS attacks. To mitigate this risk, we implement strict rate-limiting measures in our UDP responder, detailed in Section 3.2.

Our rate-limiting policy is inspired by established reactive deployments [13, 14], from which we derive temporal thresholds governing when and for how long our responder ceases to emit UDP packets. The volumetric thresholds – governing response magnitude – are conservatively calibrated to ensure our contribution to aggregate Internet traffic remains negligible compared to production services. Throughout our measurement campaign, we observe no successful evasion of our rate limiter, with empirical traffic volumes remaining within designed bounds as reported in Section 4.

Table 7: Top 10 application-layer protocols for BR_1. We underline protocols that appear after replying.

Sensor	Label	# Packets	# Flows	# Payloads	# Sources	# Ports
Darknet	*Unknown*	153,123 (51.32%)	148,902	114,836	21,405	37,404
	RakNet	25,550 (8.56%)	25,548	17,395	9,302	20,682
	SIP	21,255 (7.12%)	21,242	19,134	718	69
	Valve SEQ	14,402 (4.83%)	14,384	3	923	29
	DNS	14,268 (4.78%)	13,835	5,015	1,479	50
	NTP	9,729 (3.26%)	9,577	2,517	1,624	935
	Minimal 1-Byte Payload	6,304 (2.11%)	6,304	173	417	111
	Large ASCII 'X' Scanner marker	4,695 (1.57%)	388	2	1	1
	SSDP	4,146 (1.39%)	4,127	21	329	30
	BitTorrent	3,733 (1.25%)	2,815	1,891	901	236
	Status probe	3,558 (1.19%)	3,558	1	59	12
RT-UDP	*Unknown*	621,474 (47.39%)	615,475	572,340	23,984	39,882
	Mozi Bot Config.	498,942 (38.05%)	256	3	249	14
	SIP	29,741 (2.27%)	29,720	26,860	1,222	107
	RakNet	26,874 (2.05%)	26,874	17,948	9,362	21,533
	DNS	24,798 (1.89%)	22,708	9,025	2,222	36
	Valve SEQ	16,960 (1.29%)	16,949	3	981	26
	NTP	15,337 (1.17%)	15,196	3,111	1,948	1,059
	Minimal 1-Byte Payload	7,230 (0.55%)	7,230	200	631	116
	QUIC	6,466 (0.49%)	3,957	2,666	710	20
	SSDP	6,114 (0.47%)	6,097	25	490	40
	IKE	5,087 (0.39%)	4,427	2,277	1,680	544

Table 8: Top 10 application-layer protocols for IT_1. We underline protocols that appear after replying.

Sensor	Label	# Packets	# Flows	# Payloads	# Sources	# Ports
Darknet	*Unknown*	169,689 (45.11%)	165,573	122,582	23,438	39,278
	RakNet	23,732 (6.31%)	23,727	16,403	9,289	19,503
	SIP	21,775 (5.79%)	21,775	19,672	1,798	65
	DTLS	21,165 (5.63%)	21,165	1,339	891	60
	DNS	20,935 (5.57%)	20,577	5,496	3,055	59
	Valve SEQ	17,297 (4.6%)	17,285	3	1,012	33
	NTP	12,257 (3.26%)	12,221	2,685	2,622	1,020
	SSDP	7,649 (2.03%)	7,625	23	1,780	37
	IKE	6,718 (1.79%)	6,022	1,944	2,410	561
	Minimal 1-Byte Payload	6,711 (1.78%)	6,710	209	639	119
	BitTorrent	5,736 (1.52%)	5,381	946	1,708	58
RT-UDP	Mozi Bot Config.	706,938 (62.88%)	236	4	225	13
	Unknown	176,872 (15.73%)	170,942	125,853	24,954	39,264
	DNS	27,318 (2.43%)	24,873	8,352	3,335	58
	RakNet	23,292 (2.07%)	23,291	16,152	9,116	19,050
	SIP	21,589 (1.92%)	21,589	19,572	1,859	53
	DTLS	20,937 (1.86%)	20,937	1,322	895	51
	Valve SEQ	16,852 (1.5%)	16,835	3	1,015	29
	NTP	16,596 (1.48%)	15,910	3,601	3,424	977
	IKE	10,091 (0.9%)	9,420	4,236	2,697	554
	SSDP	8,681 (0.77%)	8,663	25	1,826	36
	BitTorrent	6,903 (0.61%)	6,167	1,433	1,746	64

Table 9: Top 10 application-layer protocols for IT_2. We underline protocols that appear after replying.

Sensor	Label	# Packets	# Flows	# Payloads	# Sources	# Ports
Darknet	*Unknown*	175,852 (51.06%)	171,896	140,300	25,580	44,469
	RakNet	31,530 (9.15%)	31,528	20,589	10,760	24,426
	DTLS	26,311 (7.64%)	26,311	1,673	881	66
	Valve SEQ	16,624 (4.83%)	16,606	3	1,097	24
	DNS	12,867 (3.74%)	12,527	2,624	2,344	51
	IKE	7,019 (2.04%)	6,168	2,154	2,673	603
	BitTorrent	5,481 (1.59%)	5,122	959	1,472	47
	OpenVPN	4,349 (1.26%)	4,349	463	1,014	40
	SIP	4,344 (1.26%)	4,344	2,732	1,955	47
	SQL Server	4,134 (1.2%)	4,134	1	1,606	15
	pcAnywhere	3,944 (1.15%)	3,944	2	1,375	39
RT-UDP	*Unknown*	1,095,001 (61.13%)	177,762	422,107	27,985	44,692
	Mozi Bot Config.	488,344 (27.26%)	200	4	193	12
	RakNet	31,367 (1.75%)	31,363	20,272	10,580	24,332
	IKE	28,051 (1.57%)	9,092	10,564	2,954	636
	DTLS	26,129 (1.46%)	26,129	1,688	876	67
	Valve SEQ	16,476 (0.92%)	16,455	3	1,100	32
	DNS	14,696 (0.82%)	14,152	3,272	2,414	58
	BitTorrent	6,746 (0.38%)	5,906	1,713	1,642	65
	BACnet	4,967 (0.28%)	4,967	15	1,270	49
	Large ASCII 'X' Scanner marker	4,815 (0.27%)	389	2	1	3
	SIP	4,690 (0.26%)	4,690	2,924	2,210	53

Our study adheres to ethical guidelines for Internet measurement research [24]. All experiments are conducted within controlled network address spaces that were previously dark, ensuring no interference with end-user traffic or exposure of privacy-sensitive data. To facilitate reproducibility and community validation, we release anonymized network traces following established best practices. Specifically, destination subnets are replaced with placeholder prefixes that preserve distinctions between experimental configurations, while source IPv4 addresses are anonymized using the prefix-preserving Crypto-PAN algorithm [33] to maintain topological properties. Application-layer payloads are anonymized as much as possible to prevent disclosure of infrastructure-specific information that may be embedded in probe content, while keeping the data useful for research purposes.

Well-known Institutional Scanner List

We report the classes of well-known institutional scanners considered in this work, provided by [17]: `Alphastrike`, `Bufferover`, `Census`, `Censys`, `Coalition`, `Cypex`, `Cyberresilience`, `F6`, `Intrinsec`, `Ipinfo`, `Ipip`, `Modat`, `Netsecscan`, `Probethenet`, `Rapid7`, `Research scanner` (provided by CIRCL Network Telescope), `Shadowforce`, `Shadowserver`,`Shodan`, `Skipa`, `Stretchoid`.

Payload Label References

Table 6 describes the protocol labels referenced in Section 4.4 and in the per-sensor tables below. Labels are assigned manually by inspecting recurring byte patterns in the payload prefix and body, and turned into regular expressions matched across the whole dataset. Payloads that do not match any known pattern are aggregated under the *Unknown* category.

Table 6: Descriptions of payload labels.

Label	Description
BACnet	BACnet/IP
BitTorrent	BitTorrent P2P protocol
Dahua DHIP	Proprietary protocol used by Dahua devices
DNS	Domain Name System
DTLS	Datagram Transport Layer Security. Client Hello messages, v1.0 - 1.2
Mozi Bot Config	Mozi XOR encrypted configuration
IKE	Internet Key Exchange (IKE) v1 and v2
pcAnywhere	discontinued remote-access software program
Memcached	Distributed memory caching system. Generally it uses port 11211/UDP
NTP	Network Time Protocol control queries, notably MONLIST requests
QUIC	Quick UDP Internet Connections
RakNet	RakNet Open Connection Request, commonly used by Minecraft servers
SIP	Session Initiation Protocol (mostly OPTIONS, REGISTER, INVITE)
SQL Server	SQL Server Browser discovery probe
SSDP	Simple Service Discovery Protocol M-SEARCH
Valve SEQ	Valve Source Engine Query
OpenVPN	OpenVPN protocol
CoAP	Constrained Application Protocol
Large ASCII 'X' Scanner marker	Payload containing repeated 'X' character
Minimal 1-Byte Payload	1-Byte payload
Status probe	Payload containing the string 'status'

UDP Payloads Across Deployments

Tables 7–9 extend the NL_1 analysis to other sensors.

Additional Plots

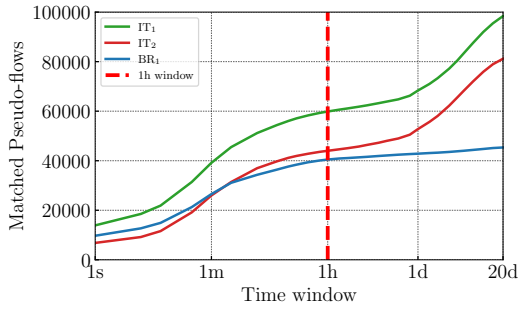


Figure 5: Follow-ups to UDP Replies: Sensitivity analysis of the time windows. The plot reports the number of matched pseudo-flows between the darknet and the RT-UDP subnet across the three deployments as the observation window expands from seconds to the full 20-day period.

Payload Regexes

The following Table 10 shows the regexes adopted to match on the application-layer payloads collected. For brevity, we omit the links to references for the protocol message formats and scanners datasets, but we will provide them along with the shared data.

Table 10: Regular expressions used to categorize payloads.

Label	Regular Expression
NTP	^\x17\x00\x03\x2a
NTP	^\x17\x00\x02\x2a
NTP	^\xe3\x00\x04\xfa\x00\x01\x00\x00\x00\x01\x00\x00
NTP	^\xe3(?:\x00){47}$
NTP	^(?:\x1b|\x5b|\x9b|\xdb|\x23|\x63|\xa3|\xe3).{47}$
NTP	^(?:\x16|\x26)\x02\x00\x01\x00\x00\x00\x00\x00\x00\x00\x00$
BitTorrent	^\x13BitTorrent protocol
BitTorrent	^\x13BitTorrent protocol.{48}$
BitTorrent	^\x00\x00\x00\x00$
BitTorrent	^\x64.*1:q(?:4:ping|9:find_node|9:get_peers|13:announce_peer).*1:y1:q.*\x65$
BitTorrent	^d1:d2:id.*1:q9:get_peers.*1:y1:q(?:\x65)?$
BitTorrent	^d1:ad2:
BitTorrent	^.{4}\x14\x00\x64.*1:md.*\x65$
BitTorrent/LSD	^BT-SEARCH * HTTP/1.1
BitTorrent	^\x00\x00\x04\x17\x27\x10\x19\x80\x00\x00\x00\x00.{4}$
BitTorrent	^\x00\x00\x04\x17\x27\x10\x19\x80\x00\x00\x00\x00
BitTorrent	^\x41\x00.{18}$
BitTorrent	^\x41\x00.{6}\x00\x00\x00\x00\x00\x10\x00\x00.{2}\x00\x00$
BitTorrent	.*ut_metadata.*
BitTorrent	.*ut_holepunch.*
BitTorrent/PEX	.*ut_pex.*
Memcached	.*stats\x0d\x0a$
Memcached	.*stats items\x0d\x0a$
Memcached	.*flush_all(?:\x0d\x0a)?$
BACnet	^\x81(\x0a|\x0b|\x0c)
RakNet	\x00\xff\xff\x00\xfe\xfe\xfe\xfe\xfd\xfd\xfd\xfd\x124Vx
SQL Server	^\x02$
SLPv2	\x02.*service:
SLPv2	^\x02\x09\x00\x00\x1d\x00\x00\x00\x00\x00\x9d\xf6\x00\x02en\x00\x00\xff\xff\x00\x07DEFAULT$
SSDP	^M-SEARCH
SSDP	^NOTIFY * HTTP/1.1
SSDP	^HTTP/1.1 200 OK
DTLS	^(?:\x14|\x15|\x16|\x17)\xfe(?:\xff|\xfd)
Mikrotik MNDP	^\xc8[\x02\x03].{2}\x00\x00\x00\x00\x00\x00\x00\x00(?:.{4,})+$
Valve SEQ	^\xff\xff\xff\xff\x54(?:\x54)?Source Engine Query(?:\x00)?$
SNMP	^\x30.*\x02\x01(?:\x00|\x01).*\x04public.*\xa0
HTTP	^(?:[\x41-\x4f]|[\x50-\x5a])+\x20(?:[\x21-\x2f]|[\x30-\x6f]|[\x70-\x7e])+ HTTP/1.(?:\x30|\x31)(?:\x0d)?\x0a
Dahua DHIP	.*DHIP
Quake III	^\xff\xff\xff\xff(?:getstatus|getinfo)
SIP	^(OPTIONS|INVITE|REGISTER|ACK|BYE|CANCEL|PRACK|SUBSCRIBE|NOTIFY|PUBLISH|REFER|INFO|MESSAGE|UPDATE) sip
RIPE Atlas	^MGLNDD_
pcAnywhere	^NQ$
pcAnywhere	^ST$
IBM DB2	^DB2GETADDR
DIGI ADDP	^DIGI
Kerberos	^(?:\x6a|\x6b|\x6c|\x6d|\x6e|\x6f|\x7e)(?:[\x00-\x7f]|\x81.|\x82.{2}|\x83.{3}|\x84.{4})\x30.+$
SCADA Moxa	^\x01\x00\x00\x08\x00\x00\x00\x00$
Latronix	^\x01\x00\x00\x00$
Gardasoft	^Gardasoft
HTML	^(?:\x0a)*<!DOCTYPE html
Apple Remote Desktop	^\x00\x14\x00(?:\x00|\x01\x03)
LDAP	^\x30(?:[\x00-\x7f]|[\x81-\x84].{1,4})\x02\x01.\x63
TFTP	^\x00[\x01\x02](?:.)+\x00(?:netascii|octet|mail)\x00
MGCP	^(?:RQNT|RSIP|CRCX|DLCX|NTFY|AUEP|AICK|EPCF|MDFX)\x20.* MGCP (?:0.1|1.0)
Omron FINS	^FINS
RIP	^\x01\x01\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x10$
RIP	^\x01\x01\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x10\x00$
tha_net_loop_scan! Scanner marker	^tha_net_loop_scan$!
Large ASCII 'X' Scanner marker	^(?:\x58){50,}$
Xpanse Scanner marker	.*Xpanse.*
Minimal 1-Byte Payload	^.$
status probe	^\status\\x00$
NBNS	^.{2}\x00\x00\x00\x01\x00\x00\x00\x00\x00\x00\x00\x20.{32}\x00\x00\x21\x00\x01$
EtherNet/IP	^\x63\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00.{8}\x00\x00\x00\x00$
CoAP	^(?:[\x40-\x48]|[\x50-\x58])(?:\x01|\x02|\x03|\x04).{2,}$
XDMCP	^\x00\x01(?:\x00[\x01-\x09]|\x00\x0a|\x00\x0b).{2,}$
OpenVPN	^\x38.{8}\x00\x00\x00\x00\x00(?:\x38\xb1\x26\xde)?$
OpenVPN	^\x38\x12+\x00+$
WS-Discovery	.*http://schemas.xmlsoap.org/ws/2005/04/discovery.*
WS-Discovery	^\x3c.*Probe.*
XML block	^<?xml
Canon BJNP	^BJNP
STUN	^[\x00-\x3f].{3}\x21\x12\xa4\x42.{12,}$
STUN Probe	^\x00\x01.{2}\x21\x12\x0a\x44.{12,}$
NTP	^(?:\x1b|\x23).{47}$
NTP	^(?:\x16|\x26).{11,}$
NTP	^(?:\x17|\x27).{7,}$
NTP	^\x17\x00\x03\x2a
QUIC	^\x0d.{8}\x51(?:[\x30-\x39]){3}\x00$
QUIC	^\x0d.{19}\x51(?:[\x30-\x39]){3}\x00\x00\x00\x00$
QUIC	.*CHLO.*PAD.*SNI.*VER
QUIC	^[\x80-\xff]\x00\x00\x00\x00
QUIC	^[\xc0-\xff]\x00\x00\x00\x01
QUIC	^[\xc0-\xff]k3C\xcf
QUIC	^[\xc0-\xff](?:[\x0a\x1a\x2a\x3a\x4a\x5a\x6a\x7a\x8a\x9a\xaa\xba\xca\xda\xea\xfa]){4}
DNS	^.{2}[\x00-\x07].\x00\x01.{4}\x00\x01.*\x00\x00\x01\x00\x01.*\x00\x00\x29.*$
DNS	^.{2}[\x00-\x07].\x00\x01.{4}\x00\x01.*\x00\x00\x1c\x00\x01.*\x00\x00\x29.*$
DNS	^.{2}[\x00-\x07].\x00\x01.{4}\x00\x01.*\x00\x00\x10\x00\x01.*\x00\x00\x29.*$
DNS	^.{2}[\x00-\x07].\x00\x01.{4}\x00\x01.*\x00\x00\xff\x00\x01.*\x00\x00\x29.*$
DNS	^.{2}[\x00-\x07].\x00\x01.{4}\x00\x01.*\x00\x00\x30\x00\x01.*\x00\x00\x29.*$
DNS	^.{2}[\x00-\x07].\x00\x01.{4}\x00\x01.*\x00\x00\x29.*$
DNS	^.{2}[\x00-\x07].\x00\x01.{6}.+$
DNS	^.{2}[\x00-\x07].\x00\x01.{6}.*\x00\x00\x01\x00\x01(?:.*)$
DNS	^.{2}[\x00-\x07].\x00\x01.{6}.*\x00\x00\x1c\x00\x01(?:.*)$
DNS	^.{2}[\x00-\x07].\x00\x01.{6}.*\x00\x00\x10\x00\x01(?:.*)$
DNS	^.{2}[\x00-\x07].\x00\x01.{6}.*\x00\x00\xff\x00\x01(?:.*)$
DNS	^.{2}[\x00-\x07].\x00\x01.{4}\x00\x01.*\x00\x00\xff\x00\x01.*\x00\x00\x29.*$
DNS	^.{2}\x10\x00\x00\x00\x00\x00\x00\x00\x00\x00$
DNS	.*\x09_services\x07_dns-sd.*
DNS	^\x00\x06\x01\x00\x00\x01\x00\x00\x00\x00\x00\x00\x07version\x04bind\x00\x00\x10\x00\x03$
IKE	^.{17}\x20.{10}.*$
IKE	^.{17}\x10.{10}.*$
Mozi Encrypted Configuration	(?i)^\x15\x15\x29\xd2.{320,400}?(?:NfZ\x8f\x80\xc8\xac\x23\x8d\xac\x47\x06\xd5Oo~){2,}.*$