A Multi-Cloud View of Internet Background Radiation

Nils Kempen[§], Ricky K. P. Mok[†], Bernhard Degen[‡], Syed Mujtaba Jafri[†], Ralph Holz[§¶]
[§]University of Münster, Münster, Germany
[†]CAIDA/UC San Diego, La Jolla, CA, USA
[‡]University of Twente, Enschede, The Netherlands
[¶]University of Sydney, Sydney, Australia

***Abstract*—As services are increasingly centralized in public clouds, understanding the nature of Internet Background Radiation (IBR) hitting these particular environments is an important part of understanding their overall security posture. Classical network telescopes, long the cornerstone of IBR research, face hurdles here: their surface area is shrinking, and their well-known address ranges are easily avoided. We present a multi-cloud view of IBR in this paper. We deploy a passive, distributed network telescope with 336 IPs across five major cloud providers. We compare traffic from our cloud telescope with data from two classical telescopes, a large well-known (/9 + /10) and a small unknown /16, to analyze observational biases. To enable a fair comparison across very different telescope sizes, we tune a scan detection algorithm to determine appropriate thresholds. Our findings reveal that IBR in the cloud is strongly provider-dependent rather than geography-dependent, highlighting the necessity of multi-cloud deployments for comprehensive visibility of IBR in the cloud. Our cloud telescope also captures a distinct set of scanners, confirming that scanning activity is not uniform across the IPv4 space, and we confirm that small, distributed telescopes are ill-suited for observing random events like DDoS backscatter. Our work underscores that monitoring must evolve beyond classical telescopes to include diverse, multi-cloud vantage points to accurately capture IBR.**

I. Introduction

Classical network telescopes, also known as darknets, leverage unused routed IP address space to passively capture incoming Internet traffic. As no active users/services are present in the darknets, all received traffic is by definition *unsolicited* and often referred to as Internet Background Radiation (IBR). A large body of research over the last two decades has analyzed IBR to detect scanning traffic, backscatter from DoS attacks, and traffic stemming from misconfigurations.

Unfortunately, the efficacy and sustainability of *classical network telescopes* are increasingly threatened by *i) changes in scanning strategies:* classical telescopes rely heavily on the assumption that attackers select targets uniformly at random [1]. This makes their well-known address ranges susceptible to evasion by modern scanners and limits their ability to observe targeted attacks. Furthermore, attackers shift their focus away from sparsely populated IPv4 spaces toward dense, high-value cloud infrastructure, which sometimes manifests as Highly Responsive Prefixes (HRPs) [2] due to Internet consolidation [3, 4] and dynamic IP mapping [5], leading scanners to preferentially target specific prefixes, and move away from uniformly random scanning [6, 7, 8]; *ii) exhaustion of the IPv4 address space:* since network operators can no longer obtain new address space from Regional Internet Registries, they are increasingly reclaiming telescope address blocks to accommodate a growing number of users and devices. For example, CAIDA's network telescope has been reduced to only 63.56% of its original size (a full /8) due to internal demand and resale of IP ranges [9].

Accordingly, researchers have begun to examine the use of public cloud platforms as a vantage point to collect IBR [6, 10, 11]. However, we observe two gaps in the literature. Firstly, there is a lack of studies that would compare classical network telescopes against a purely *passive* (listen-only) configuration of vantage points in the cloud. Secondly, it is not yet understood how IBR *differs between cloud providers*.

Our research aims to bridge these major gaps in prior studies. First, we characterize the IBR *passively* captured from virtual machines (VMs) deployed across *multiple* cloud providers, providing comprehensive data that are also better comparable to those from classical network telescopes. Second, we compare the data obtained from the cloud with two classical network telescopes (UCSD-NT and SURF-NT) with complementary characteristics. SURF-NT, operated by a research-and-education network in the Netherlands, has a smaller aperture and is not yet publicly known, reducing the likelihood that attackers will avoid scanning its address space.

We designed and implemented CLOUD-NT, an Infrastructure-as-code multi-cloud network telescope that supports IBR collection across five major cloud providers. We deployed CLOUD-NT for two weeks, using 336 VMs spanning 143 cloud regions. Our key contributions are: 1) a method for adjusting the detection-threshold parameters of Zeek's scanner-identification algorithm to accommodate telescopes of varying sizes, validated against known IP blocklists, 2) an extensive analysis of IBR across different cloud providers, 3) an in-depth comparison and analysis of scanners targeting cloud versus classical telescopes.

We find that: 1) The IBR collected within the same

cloud provider shows higher similarity in source IPs than IBR collected across different providers in the same region (average Jaccard index of 0.35 vs. 0.29). 2) We identify a total of 941k scanners across all telescopes. Only 0.60% of these are commonly detected by all three network telescopes (766k/96.60k/6.80k of scanners are uniquely observed by UCSD-NT/SURF-NT/CLOUD-NT, respectively). 3) We examine the destination ports targeted by the scanners and find that different telescopes observe different types of scanning traffic, with cloud deployments being heavily targeted on cloud-service associated ports and classical telescopes on ports associated with remote-control services.

II. Related work

The composition of IBR has shifted substantially over the past two decades. Pang et al. [12] characterized IBR at four unused networks and found it dominated by self-propagating worms and autorooters targeting Windows services such as CIFS, RPC and NetBIOS, with backscatter from RSDoS attacks as a secondary component. Almost a decade later, Dainotti et al. [13] documented a horizontal scan of the entire IPv4 address space conducted by a botnet from approximately 3 million distinct IP addresses, using a coordinated and unusually covert scanning strategy targeting *SIP* server infrastructure. More recent studies have reported a continued shift away from indiscriminate worm propagation toward targeted, service-specific scanning campaigns and away from uniformly random target selection [6, 7, 8].

Deploying and operating large-scale network telescopes is increasingly costly and challenging, motivating researchers to investigate how reducing telescope size affects threat visibility [14] and explore alternative IBR-collection strategies. These include analyzing firewall logs from Content Delivery Networks (CDNs) [7], analyzing network flows traversing Internet Exchange Points [15], recovering IBR from ICMP errors [16], and leveraging unused IPs in active networks (greynets) [17]. All these approaches require access to traffic data from production networks, which is often unavailable publicly and may raise privacy concerns.

Cloud platforms offer researchers resources to rent (e.g., IPv4 addresses, compute, and storage) to build network telescopes and honeypots for IBR collection and threat intelligence. Izhikevich et al. [6] deployed honeypots across several networks, including the cloud, and combined them with *Greynoise* data. They show that scanners avoid networks without live services and differentiate by geography. *Cloud Telescope* [10] used 26 AWS regions (1 IP per region) to gather IBR. Such a small deployment is vulnerable to distortion by a few heavy hitters. *DScope* [11] gathered unsolicited traffic on low-cost AWS spot instances, which could be frequently terminated and reassigned, causing IP churn. The captured traffic may include spill-over from former users rather than genuine malicious packets. The VMs deployed in all these prior studies reacted to ingress traffic to different degrees. For example, the VMs used in *Cloud Telescope* [10] had their SSH port changed to 65535, meaning they remained reachable from the Internet. Such interactivity can alter scanner behavior [18, 19], making the resulting data hard or impossible to compare to that obtained by classical passive network telescopes. Both *DScope* [11] and *Cloud Telescope* [10] deployments used only a *single* cloud provider, limiting the generalizability of their findings across the broader cloud ecosystem.

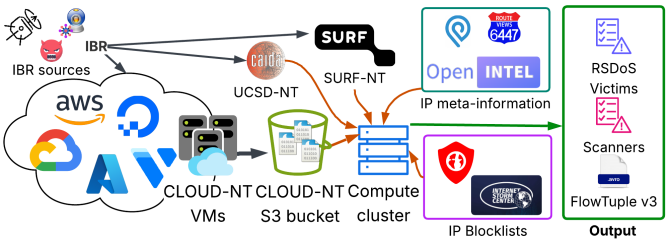

Fig. 1: Our data collection and analysis pipeline integrates multiple sources of data to gather insights into IBR.

In this work, we carefully design our experiments to eliminate potential confounding factors, allowing us to reveal the true visibility afforded by building network telescopes in the cloud. Our comparative analysis uses data collected from two classical network telescopes, allowing us to provide a comprehensive perspective on IBR.

III. Methodology

We develop CLOUD-NT, a distributed network telescope composed of virtual machines (VMs) across various regions of major cloud providers, to collect IBR from April 5-23, 2025 (§III-A). For comparison, we obtain IBR data from two classical network telescopes, during the same period (§III-B). We process the data and annotate them with additional meta-information and datasets (§III-C). Because the telescopes differ vastly in size, we develop methods to calibrate the thresholds for detecting scanning events (§III-D) and describe how we detect RSDoS events (§III-E). Fig. 1 shows an overview of our data collection and analysis pipeline.

A. Deploying a Distributed Multi-Platform Cloud Telescope

We design and implement CLOUD-NT, a multi-platform cloud-based network telescope to collect IBR from diverse networks and geographic locations. Our deployment strategy aims for provider independence, although some platform-specific adjustments are necessary. To ensure scalability and reproducibility, we develop a set of Terraform scripts to deploy the smallest available VMs across five major cloud providers: AWS, GCP, Azure, Vultr, and DigitalOcean. We select these providers to provide broad geographic coverage and diverse network topologies, leveraging their respective strengths in region availability and pricing. To maximize geographic and topological coverage, we deploy one VM per availability zone to target the smallest and most granular deployment unit available, resulting in a total deployment of 336 instances.

Table I summarizes CLOUD-NT's per-provider deployment, including captured data volume over the measurement period

TABLE I: CLOUD-NT Deployment Overview: Per-Provider Region/Instance Counts, Capture Sizes, and Cost.

Provider	Region	Instances	Capture Size	Prices[1] per day	Prices[1] Total
AWS	23	73	1.04 GiB	$19.37	$348.66
GCP	36	106	1.40 GiB	$25.70	$462.60
Azure	40	113	1.30 GiB	$26.50	$477.00
Vultr	30	30	0.62 GiB	$5.37	$96.58
Dig.Oc.	14	14	0.21 GiB	$2.00	$35.99
Total	143	336	4.57 GiB	$78.93	$1420.83

and incurred costs. The total deployment cost was approximately 1420.83 USD over the measurement period ≈78.93 USD per day.

Launching many VMs in a public cloud is costly. As cloud providers not only charge for the compute resources of the VMs but also for the use of public IPv4 addresses, we need to be selective in where we deploy VMs. Despite the economical advantage they offer, we chose not to use spot instances as in [11], as these can be terminated at any time, which also means their IP addresses are released. In contrast to [11], our method allows longitudinal observation from a stable set of IPs, so we can also distinguish between IBR and residual (spillover) traffic caused by prior users of the IPs.

We leverage IPv6 for management and data transfer to ensure the IPv4 interfaces of the VMs are completely non-responsive to ingress traffic. Also, we configure the cloud providers' default firewalls/security groups to allow all ingress traffic to be forwarded to the VMs and apply `iptables` rules to prevent the VM's OS from responding to any IPv4 traffic. We disable all VM "health check" services offered by the cloud providers to reduce internal probing traffic to the VMs.

B. Obtaining Data From Classical Network Telescopes

For comparison, we obtain the traffic data collected during the same time period from two classical network telescopes:

(a) **UCSD-NT:** This telescope is the largest network telescope available to researchers. It is operated by CAIDA at UC San Diego, capturing traffic from adjacent /9 and /10 network prefixes. Because the telescope has been operational for a long time, its comprising prefixes have been exposed and are widely known within the research community. We use it to represent a well-established, widely used vantage point in our study. (b) **SURF-NT:** Operated by the Dutch national research and education network, this telescope consists of one /16 network. Because SURF-NT was set up very recently, its comprising prefixes have not been exposed to the same degree. We include SURF-NT because its operational setup closely resembles UCSD-NT, but its relatively unexposed address space provides a valuable contrast for evaluating visibility biases.

C. Data Processing and Annotation

To facilitate analysis, we convert raw packet captures collected from the cloud, UCSD-NT, and SURF-NT into the *FlowTuple* [20] format using CAIDA's *Corsaro* software [21]. This process aggregates packets into flows at 5-minute intervals and maps source IP addresses to Autonomous System Numbers (ASNs) using CAIDA's prefix-to-AS mappings [22].

To be able to later analyze the geographic locations and network types of IBR sources, we annotate the flow data with geolocation and ASN type classifications using IPinfo [23], which prior evaluations have shown to provide accurate results [24]. Additionally, we resolve the hostnames of source IPs using historical reverse DNS (PTR) records collected by OpenINTEL [25]. This allows us to identify the specific infrastructure and hostnames associated with the sources of IBR and extract geo-information included in PTR records using CAIDA's *Hoiho* [26]. We provide an overview of the used enrichment datasets in Table VI; it additionally includes the FireHOL blocklist [27] and the ISC SANS research scanner list [28], which we use as proxy ground truth for the scanner-detection calibration in §III-D.

D. Detecting Internet Scanners and RSDoS Attack Events

To detect scanning and Randomly Spoofed Denial of Service (RSDoS) attacks in IBR, allowing for comparison between the different network telescopes, we need an adaptable method to detect scans in different telescope setups.

To identify scanning activities, we adopt a prevalent heuristic that is derived from the *Zeek* Intrusion Detection System [29]. Zeek's default policy identifies a source as a scanner if it attempts to contact 25 or more unique destination IP addresses on a single port within a 5-minute moving window ($N = 25, T = 5m$). Prior studies have adopted these parameters for scan detection without adjusting them for the specific telescope context (e.g., darknets and greynets) or size (e.g., /9, /14, and /16) [29, 30].

To maintain the rigor of threshold-based detection while ensuring accurate temporal boundaries, we apply an overlapping sliding window workflow [31, 32]. Unlike fixed-window approaches that may artificially fracture scan events across temporal boundaries, we employ a 15-minute window with 5-minute increments, to align with our *FlowTuple* time intervals. This creates a 10-minute overlap between consecutive windows, guaranteeing that any isolated scan event with a duration $D \leq 10$ minutes is encapsulated within at least one continuous processing frame, mitigating potential boundary loss effects.

We aggregate the traffic within these windows by source IPs and examine them using a strict, top-down priority sequence. We evaluate the number of unique destination hosts (N_h), ports (N_p), and IP-port pairs (N_r) contacted against their respective thresholds (T_h, T_p, and T_r). Applying this hierarchy classifies scans into three mutually exclusive categories: (a) *Block/Random Scan* identifies high-entropy sources targeting a high volume of unique destination IP and port combinations. This triggers if a source exceeds the unique pairs threshold ($N_r \geq T_r$) OR reaches both base thresholds in the examined window ($N_h \geq T_h \wedge N_p \geq T_p$). This captures aggressive Internet-wide sweeps that would otherwise count towards the more targeted scan metrics. (b) *Horizontal Scan*

[1]Based on 2026 list prices; see Appendix A for more details.

identifies sources targeting multiple hosts strictly below the port threshold ($N_h \geq T_h \wedge N_p < T_p$). This isolates horizontal scanning. (c) *Vertical Scan* identifies sources targeting a limited number of hosts on many ports ($N_p \geq T_p \wedge N_h < T_h$). Following the original Zeek implementation, we adopt $T_p = 25$ across all datasets to define this boundary.

This approach prevents a single massive Internet-wide scan, one that probes many IP-and-port combinations, from independently exceeding both the horizontal and vertical thresholds, which could split the event into multiple smaller scan records.

Threshold Calibration and Validation: Applying the same set of threshold across all the telescope data is inappropriate due to the vast difference in telescope sizes, ranging from 336 IPs (CLOUD-NT) to over 10M destination IPs (UCSD-NT).

Because modern asynchronous scanners distribute probes globally, smaller telescopes experience a reduced statistical hit-rate per time window compared to large subnets. For example, at a scanning rate of 100k packets per second, about 226.43k and 7.04 packets are expected in a 15-minute window for UCSD-NT and CLOUD-NT, respectively. Thus, a static default threshold (e.g., $T_h = 25$) cannot adapt the sensitivity of the detection; it sets an unrealistically high bar for small telescopes, causing them to miss global scanners [33].

To set the thresholds for each telescope, we sweep a range of threshold values for horizontal scans ($T_h \in \{2, 5, 10, 15, 25, 50, 100\}$) and random scans ($T_r \in \{25, 50, 100, 200, 500, 1000\}$). We identify the optimal trade-off point between scanning and other IBR with the Kneedle algorithm to locate the point of maximum curvature (the "elbow") in the resulting scanner volume distributions with per-telescope results presented in §IV-E.

Furthermore, to empirically validate these derived elbows without labeled network telescope scan data, we compare all unique source IPs extracted from the varying threshold bounds against two IP blocklists—FireHOL blocklists [27] and the research scanner list published by the Internet Storm Center [28]. This validation assumes that maintaining overlap with known abuse lists indicates higher confidence in our thresholds. Both feeds are by construction skewed toward IPs already noisy enough to have been reported; an aggregated blocklist, like FireHOL, mitigates individual-provider bias and establishes a practical external baseline.

E. Identifying RSDoS Events and Victims

Our approach to detect RSDoS attacks relies on CAIDA's *Corsaro* software and the method proposed in [34]. Specifically, the `rsdos` plugin filters telescope traffic to isolate unsolicited response packets, such as TCP SYN/ACKs or ICMP errors, that attack victims inadvertently send to spoofed addresses. The pipeline then aggregates these backscatter packets into discrete attack events by grouping them by protocol and victim's IP addresses. To effectively eliminate background noise and isolate genuine attacks, the system uses thresholds regarding flow duration, packet rates, and timeout intervals. For a comprehensive breakdown of the exact *Corsaro* configuration and the specific mathematical thresholds used in this analysis, we refer to [35, Appendix J].

IV. Results

We provide an overview of IBR captured by all the telescopes during our two-week measurement period in §IV-A. We conduct in-depth analysis to study the similarity of IBR in different cloud providers in §IV-B and the overall cloud IBR in §IV-C. Next, we compare the cloud-based network telescope approach to classical network telescopes in two different use cases: RSDoS detection in §IV-D, scanner detection in §IV-E and scanner overlap in §IV-F.

A. Data Overview

Over the measurement period, UCSD-NT captured 51.18 TiB of traffic while SURF-NT captured 354.20 GiB, compared to 4.57 GiB for CLOUD-NT (Table I).

We summarize the temporal changes in volume and the source statistics of received traffic across telescopes in Fig. 2 and Table II, respectively. CLOUD-NT received more unique sources per IP than the classical telescopes (46.16 vs. 1.14 and 0.09), while the latter observed broader source-network diversity due to their larger apertures (e.g., 75.34k ASNs for UCSD-NT). Within CLOUD-NT, the combined number of average unique source IPs observed is lower than any single provider, since some sources target only a subset of providers (cf. §IV-B, Fig. 4). Averaging over all destinations dilutes their contribution. UCSD-NT and SURF-NT maintained more stable baselines than CLOUD-NT throughout the measurement period (Fig. 2), suggesting that the variability observed in CLOUD-NT is inherent to cloud address space rather than measurement artifacts.

Deploying VMs in cloud environments introduces unique observational challenges. While we cannot definitively exclude the presence of edge filtering or proprietary anti-DDoS scrubbing at the cloud provider level, CLOUD-NT captured a sustained volume of unsolicited traffic. Our VMs logged an average of ≈5.70M packets per day, originating from ≈16K unique ASNs. We identified ≈5.90M connection attempts from well-known Internet-wide scanners, such as Shodan, Censys, and Shadowserver, via DNS `PTR` records. This diverse traffic profile indicates that our deployment environment was

TABLE II: Overview of Network Telescope Datasets During the Measurement Period (April 5-23, 2025).

Telescope	avg. uniq. src IPs / dst IP / h	uniq. ASNs overall	avg. uniq. ASNs seen / h
UCSD-NT	0.09	75,335	17,911.39
SURF-NT	1.14	16,858	5,059.86
CLOUD-NT (comb.)	46.16	12,791	1,545.98
GCP	77.02	10,077	803.67
Azure	79.86	9,432	819.94
AWS	95.83	8,011	632.20
Vultr	156.58	8,293	590.81
DigitalOcean	212.68	6,865	394.80

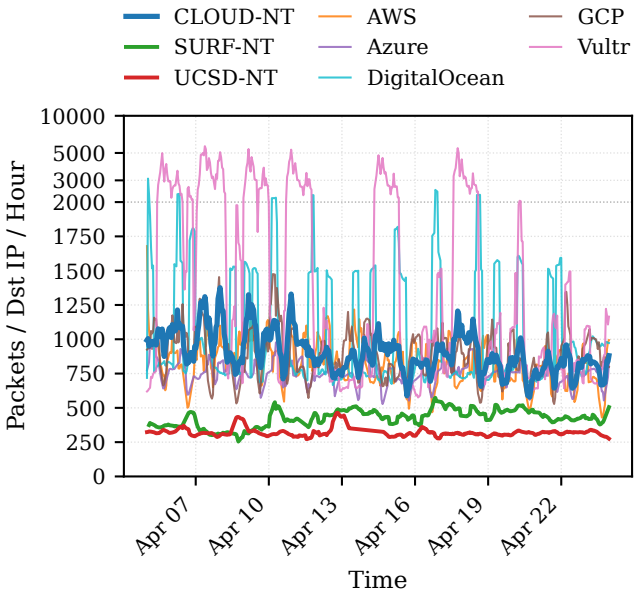

Fig. 2: Traffic volume across our three network telescopes in terms of packets received per hour per destination IP address. CLOUD-NT is additionally split up into its individual components.

sufficiently permissive to capture representative IBR, without noticeable filtering by the provider.

As cloud providers routinely recycle dynamic IPv4 addresses, our VMs inevitably captured some traffic intended for previous tenants of those IPs: "traffic spillover". To quantify this, we fitted an ordinary least-squares linear regression to each provider's daily mean packet rate over the measurement period and tested whether the slope differed from zero using a two-sided t-test at $\alpha = 0.05$. Overall, CLOUD-NT exhibited a statistically significant downward trend of -14.2 packets/IP/day ($p < 0.001$, $R^2 = 0.57$), declining from $\approx$1k to under 800 packets/IP. AWS showed a similarly significant decline (-9.3 packets/IP/day, $p < 0.01$), while the regression results for other providers were not statistically significant.

While we expected higher variability in Vultr and DigitalOcean due to smaller deployments, Vultr exhibited a steep overall decline (-105.60 packets/IP/day, $p < 0.01$). We identified AS 48090 (TECHOFF SRV) as the main cause. Instead of steady traffic, AS 48090 showed a strict on-off periodicity, alternating between inactivity and traffic bursts (peaking at 41–74% of total packets on active days) which ceased late in our observation window. The cause of this on-off periodicity remains an unresolved provider-specific artifact. We therefore interpret Vultr's regression slope as descriptive of this measurement period rather than evidence of a stable trend.

Together, these density, stability, and per-provider-artifact differences mark cloud IBR as qualitatively distinct from IBR in classical telescopes.

B. Cloud-Provider Similarity

To quantify the overlap of IBR collected across different cloud providers, we computed the pairwise Jaccard similarity of unique source IPs across all cloud region pairs. The resulting similarity matrix (see the full heatmap in Appx. A) is exceptionally dense. Visualizing the data grouped by cloud provider revealed a structural pattern. The main diagonal blocks of the heatmap, which represent comparisons within the same provider, showed distinctly high similarity. In contrast, the blocks outside this main diagonal, which represent comparisons between different cloud providers, showed markedly lower similarity regardless of their geographic locations. On average, the Jaccard similarity for intra-provider/inter-region pairs was 0.346, compared to 0.293 for inter-provider/intra-region pairs. The absolute difference of 0.053 represents an 18% relative increase in shared source IPs within a single provider's infrastructure compared to the overlap between different providers in the same geographic area.

As a representative example of the consistency of this phenomenon, we show the Jaccard similarities of selected region pairs in Fig. 3.

Frankfurt, Germany: AWS, GCP, and Azure each have set up a region in the city. Analyzing GCP traffic in the `europe-west3` (Frankfurt) region revealed a Jaccard similarity of 0.49 with GCP's `me-west1` region (Tel Aviv) (Fig. 3a). In contrast, comparing `europe-west3` to other providers physically located in the same city yielded notably lower similarity scores: 0.30 with AWS `eu-central-1` (Frankfurt) and 0.25 with Azure `germanywestcentral` (Frankfurt). This indicate that the GCP Frankfurt region has nearly twice as much overlap in IBR sources with a geographically distant GCP data center than with an Azure data center located in the same city.

North America vs. Europe: We observed a comparable trend in AWS. The `ca-central-1` region (Canada) had a Jaccard similarity of 0.48 with the geographically distant `eu-south-1` region (Milan, Italy). In contrast, when compared to a competitor in the same country, AWS `ca-central-1` only showed a 0.32 similarity with Azure's `canadacentral` region (Fig. 3b). This constituted a 50% increase in overlapping source IPs for the intra-provider pair compared to the geographically proximate inter-provider pair.

We examined all regions where multiple providers are active and found a consistent pattern expect for DigitalOcean's `sfo1` region. Instead of targeting specific geographic regions across different networks, IBR sources primarily interacted with the global IP space of individual cloud providers. This indicates that an instance's IBR traffic profile depended more on its hosting provider than on its physical geographic placement.

C. Destination IP Coverage vs. Packets

Characterizing IBR captured by CLOUD-NT prior to applying any specific behavioral thresholds, Fig. 4 presents a heatmap showing the behavior of each unique source IP observed by CLOUD-NT. Adapted from Richter and Berger [7], source IPs are binned across two dimensions: the fraction

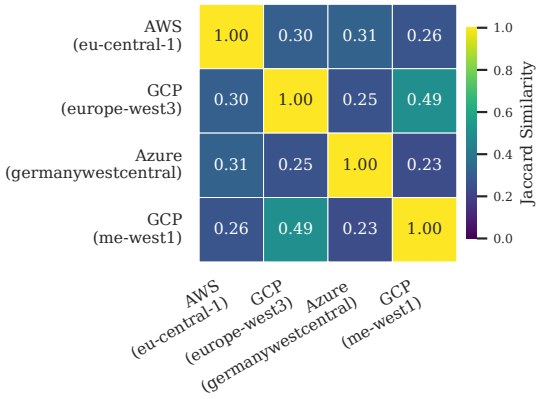

(a) GCP regions and co-located competitors. High intra-provider similarity contrasts with low cross-provider overlap in Frankfurt.

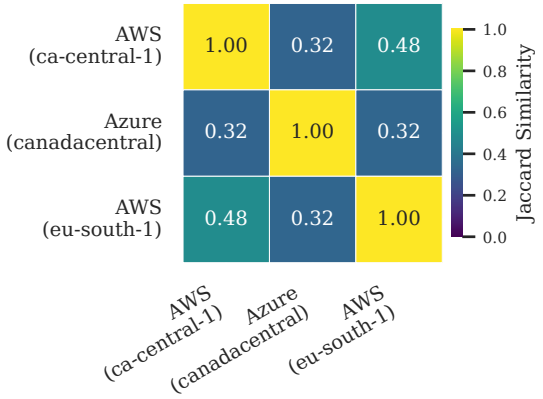

(b) AWS regions and co-located competitors. Geographically distant AWS regions share more source IPs than co-located cross-provider pairs.

Fig. 3: Jaccard similarity of observed source IPs for selected cloud regions. Diagonal entries (1.00) are self-comparisons. Off-diagonal intra-provider pairs consistently exhibit higher similarity than inter-provider pairs in the same geographic location.

of monitored destination IPs contacted, or the fan-out (x-axis), and the total number of packets sent (y-axis). The aggregated data shows a dense concentration of source IPs clustered near the y-axis. These hosts contacted only a small fraction of the available destination IPs but generated a high volume of traffic, a pattern that typically indicates traffic spillover or highly localized backscatter. We also observed several distinct vertical artifacts at specific destination fractions (between 0.8-1.0). Our investigation reveals that these vertical lines arise because some source IPs interacted with only a subset of cloud providers. Differences in the number of VMs deployed across providers produce these artifacts in the figure. For example, a source scanning our entire deployed infrastructure except AWS would contact 263 of our 336 total instances, producing the prominent vertical line observed at $x = 0.78$. Some sources only contact a subset of providers, matching the provider-dependency pattern from §IV-B.

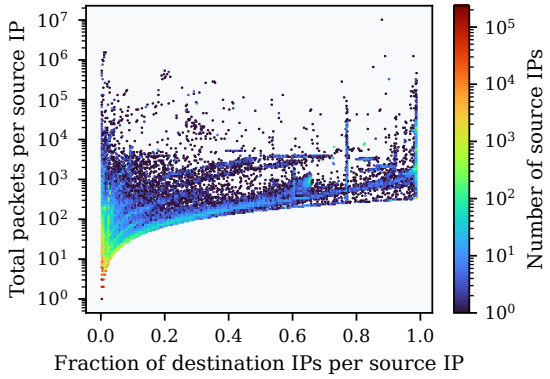

Fig. 4: CLOUD-NT: heatmap of source IPs and packet counts across all deployed IPs. Most sources only contact a small fraction of IPs; vertical artifacts arise from sources targeting a subset of all providers (§IV-C).

D. RSDoS Visibility

To evaluate cloud-based RSDoS monitoring, we compared the CLOUD-NT with SURF-NT. We found that CLOUD-NT's address space is too small for reliable detection. SURF-NT vastly outperformed CLOUD-NT, capturing ≈136K events (1.33B packets) across ≈19.50K unique target IPs, compared to CLOUD-NT's ≈15.10K events (61.10M packets) and 109 IPs. Temporally, CLOUD-NT experienced prolonged periods of zero activity with brief spikes, in contrast to SURF-NT's continuous traffic. In line with [14], the visibility of RSDoS scales with the size of the telescope and 336 IPs across five providers does not substitute for a large classical telescope.

E. Scanner Visibility

As outlined in §III-D, we sweep T_h and T_r per telescope and select the elbow point in the retained scanner-IP curve, and cross-reference against the FireHOL and ISC SANS blocklists (§III-C). Because traffic volumes varied significantly across telescopes, we performed this calibration independently for each telescope. To empirically validate these selections, we cross-referenced the retained source IPs against the *FireHOL* and *ISC SANS* blocklists ("attacks" and "research" category lists, respectively). Fig. 5 shows the intersection and disjoint sets of IPs across varying host scan thresholds, demonstrating how aggressive filtering impacts the retention of known scanners versus other IBR. As network telescopes only capture a limited subset of Internet attacks/events, only a fraction of the detected scanners overlap with the blocklisted IPs.

Evaluating scanner detection across thresholds for UCSD-NT, we observed a steep initial decline in source IPs classified as scanners, followed by a distinct plateau. At a baseline threshold of $T_h = 2$, the telescope captured over 1.20M horizontal host scanners. Increasing the threshold to $T_h = 25$ (the computed elbow) reduced this count by 96% to 47.01k. Importantly, this steep drop was not purely elimination. While low-volume sources were discarded entirely, some IPs were simply reclassified as port or random scanners since they no

longer met the elevated host threshold relative to their port count. Despite this overall reduction, the intersection with the blocklists remained highly concentrated, retaining 1339 IPs from FireHOL and 670 from ISC. Increasing the threshold further to $T_h = 50$ negatively impacted detection, dropping the total host scanner count to 1098 and the FireHOL intersection to just 103. For random scans, the volume stabilized entirely at $T_r = 100$. Scaling from 100 up to 1000 resulted in less than a 1% change in the FireHOL overlap (7716 vs. 7637). Thus, we select thresholds of $T_h = 25$ and $T_r = 100$ for UCSD-NT.

SURF-NT data demonstrated a scaled correlation between the size of the telescope's address space and the used thresholds. Increasing the host threshold from 2 to 15 reduced the horizontal host scanner count by over 950k IPs (dropping from ≈1.10M down to ≈153K). This reduction in overall IP count resulted in a negligible drop in blocklist overlap; the intersection with the ISC list only shifted from 3322 at $T_h = 2$ to 3263 at $T_h = 15$ (a 1.70% loss). Because random scans hit this mid-sized telescope more sparsely, a higher volume bound was required to prevent targeted scans from being misclassified as random scans, establishing calculated elbows of $T_h = 15$ and $T_r = 200$.

For the much smaller CLOUD-NT, employing a lower threshold of $T_h = 10$ provided the best trade-off. Compared to a baseline of 2, setting $T_h = 10$ reduced the number of captured host scanners from 141.17K to 17.59K. However, the FireHOL intersection remained stable at 3626 overlapping IPs. If we apply UCSD-NT's threshold ($T_h = 25$) to this smaller telescope, the FireHOL overlap would drop to 2670. Therefore, we use the derived thresholds of $T_h = 10$ and $T_r = 100$ for CLOUD-NT. In short, fixed default thresholds do not transfer across telescope sizes and operators should tune them to their own telescope and check the result against an external feed.

F. Scanner Overlaps and Telescope Scale

With the chosen thresholds, we evaluated the overlap of identified scanners. Fig. 6 shows that each telescope captured a large exclusive set: 765.67k IPs in UCSD-NT, 96.62k in SURF-NT, and 6.85k in CLOUD-NT. Showing high marginal visibility, CLOUD-NT observed 54.88 scanners per IP, compared to 2.40 for SURF-NT and 0.07 for UCSD-NT.

Table III summarizes the behavior of these exclusive scanners. While median port counts were comparable across telescopes (10 to 32), the tails diverged sharply. The 99$^{\text{th}}$ percentile reached 2438 ports in CLOUD-NT versus 273 in UCSD-NT and 37 in SURF-NT, revealing a population of scanners that exclusively targeted cloud prefixes with deep vertical enumeration while bypassing classical telescopes entirely. We computed the Pearson correlation between ports and hosts scanned. We found that UCSD-NT exhibited a strong positive correlation ($r = 0.68$), indicating indiscriminate sweeps where scanners that probe more ports also target more hosts. In contrast, CLOUD-NT showed a slight negative correlation ($r = -0.10$, $p < 0.001$), meaning scanners probing many ports concentrate on fewer hosts, consistent with targeted scans of specific cloud IPs.

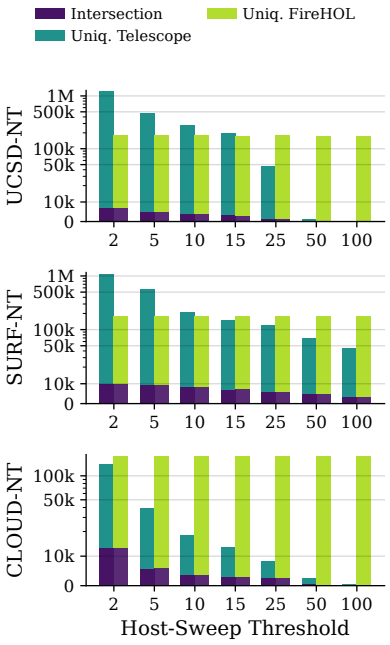

Fig. 5: Intersection of identified host scanner IPs with FireHOL blocklists across varying host scan thresholds for all three telescopes.

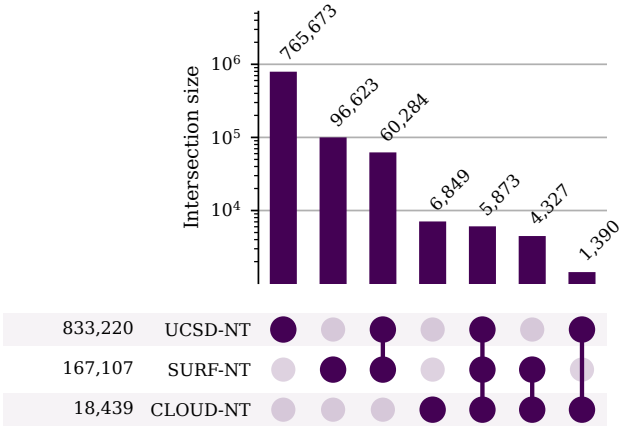

Fig. 6: Scanner-IP overlap across telescopes. Only 5873 scanners are seen by all three. The telescopes are complementary, not redundant.

Targeted Ports: To study the services targeted by scanners in cloud environments, we examined the top destination ports of the telescopes' exclusive scanners by calculating the total number of scan events in which each port was present (Table IV). In both UCSD-NT and SURF-NT, port 23 (Telnet) was the top port, generating nearly 8M hits in UCSD-NT alone. This predominance of Telnet traffic may be caused by IoT botnets (e.g., Mirai) [36]. Additionally, UCSD-NT observed substantial traffic destined to port 8728 (MikroTik

TABLE III: Exclusive Scanner Behavior by Scan Type (§III-D)

	Type	Count	Med. Ports	Med. Hosts
UCSD-NT	Host	980k	23	27
	Port	3.10M	28	22
	Random	4.30M	57	53
SURF-NT	Host	381k	4	31
	Port	414	27	3
	Random	95k	31	132
CLOUD-NT	Host	30k	8	14
	Port	1.40k	32	9
	Random	5.40k	36	12

TABLE IV: Top 5 Targeted Ports by Telescope Exclusive Scanners (Percentage of Scan Events Including the Port)

Rank	UCSD-NT		SURF-NT		CLOUD-NT	
	Port	%	Port	%	Port	%
1	23	0.60	23	3.93	443	0.04
2	80	0.46	80	2.94	8443	0.04
3	8728	0.44	8080	2.67	990	0.03
4	0	0.41	8000	2.54	993	0.03
5	22	0.39	8081	2.45	995	0.03

RouterOS), highlighting the continuous automated exploitation of networking devices. Additionally, we observed frequent scans targeting port 0. This phenomenon has already been extensively discussed in existing literature [37, 38, 39] and we did not investigate it further.

In contrast, Telnet and HTTP were absent from the list of top ports in CLOUD-NT (coming in on place 20 and 144 respectively). Cloud-exclusive scanners focused on higher-level enterprise and web infrastructure. The most targeted ports in CLOUD-NT were 443 (HTTPS) and 8443 (often used for Web Admin panels), followed by mail servers (990, 993, 995, 110) and remote access or database ports like 3389 (RDP) and 6379 (Redis). The presence of TLS-heavy ports implies that cloud directed scanners would perform complete handshakes to gather certificate data, hunt for misconfigured enterprise applications, or seek out exposed administrative interfaces.

Scanner Origins: Table V depicts the geographic and AS origins of detected scanners. UCSD-NT traffic was highly distributed and Asia-centric, dominated by China (35.3%) and India (9.0%), with the US accounting for only 6.4%. Furthermore, its ASN distribution lacked concentration; each top ASNs only contributed ≈0.5% of traffic. This lack of centralization within major hosting providers suggests the scanning originated from widely distributed individual sources, aligning with prior observations of decentralized IoT botnets operating heavily from residential ISP address spaces [36].

We observed a cloud-on-cloud scanning phenomenon in CLOUD-NT. To ensure this observation was not an artifact of cloud-internal traffic, we compiled IPv4 prefixes for each cloud provider and removed all intra-provider traffic (*i.e.*, traffic where the source IP belongs to the same provider hosting the receiving telescope VM). Regardless of the specific provider evaluated, incoming scans were predominantly of U.S. origin (≈73–80%). The ASN distribution was highly

TABLE V: Top 3 Source Countries and ASNs of Scanners per Telescope

Telescope	Top Countries (%)		Top ASNs (%)	
UCSD-NT	CN	35.36	Microsoft (AS8075)	0.53
	IN	9.01	Censys (AS398324)	0.52
	US	6.40	SS-Net (AS204428)	0.52
SURF-NT	US	20.32	DigitalOcean (AS14061)	42.28
	IN	11.76	HiNet (AS3462)	10.87
	TW	11.62	NIB India (AS9829)	5.96
CLOUD-NT	US	73.34	Linode (AS63949)	45.58
	CN	5.27	Microsoft (AS8075)	14.31
	BR	4.79	Amazon (AS16509)	7.10

concentrated among competing cloud providers: ASN63949 (Linode/Akamai) alone accounted for roughly 45% of all scanners hitting CLOUD-NT, followed closely by Microsoft (AS8075) and AWS (AS16509). Among the top 1000 scanner ASes, 82.60% of all CLOUD-NT scanners originated from hosting ASes, compared to only 39.50% for UCSD-NT.

The three telescopes thus see largely disjoint scanner populations. CLOUD-NT misses the sparse, randomized scans characteristic of botnets, which are visible to UCSD-NT, but it exposes targeted vertical scanning from commercial datacenters aimed at enterprise and web infrastructure. CLOUD-NT demonstrates that location in the address space can be as important as network size for observing modern scanning campaigns.

V. Limitations

Our study has several limitations: 1) Our measurement only spans a bit over two weeks, which is insufficient to capture seasonal patterns or longer-term shifts in scanning behavior. Our similarity and scanner overlap analyses likewise reflect a single observation window. How these evolve over longer periods, remains future work. 2) CLOUD-NT's aperture of 336 IPs inherently limits visibility into rare, randomly-distributed events such as RSDoS backscatter (§IV-D) and sparse scans. As we chose to only deploy one VM per region / availability zone for each provider, results for individual regions should be interpreted with caution, while our cross-provider findings remain robust. 3) Our spillover analysis (in §IV-A) reports a single regression slope per provider and thus captures only the aggregate decay trend. An address that previously hosted a popular public-facing service, with PTR records and client-side DNS caches still pointing to it, would plausibly receive residual traffic far longer than one previously used only for internal or outbound workloads. A per-IP characterization would refine this estimate. 4) As discussed in §IV-A, we cannot rule out provider-side DDoS scrubbing or edge filtering, but the consistent presence of well-known scanners in our captures provide a lower-bound for actual traffic. The traffic captured still accurately represents what actually is observed by cloud tenants.

VI. Discussion

Our results show that cloud and classical telescopes observe different scanning activity. The scanner overlap between the telescopes is small, indicating that neither approach subsumes the other. Classical telescopes capture broad sweeps from distributed residential networks, while cloud telescopes show some targeted vertical enumeration from commercial hosting infrastructure. A researcher using only one type would miss the specific scans visible to the other. For studying threats to cloud-hosted services, a cloud telescope is necessary. Cloud-exclusive scanners probe enterprise services that are barely visible in classical telescope traffic (Table IV). The slightly negative port-host correlation in CLOUD-NT further hints at deliberate scanning campaigns rather than other IBR. Classical telescopes do not show this activity because some of these scanners focus on cloud providers' address ranges. Conversely, cloud telescopes are inadequate for RSDoS detection and for tracking globally distributed botnets.

The cloud provider similarity analysis has a direct operational implication. Since IBR profiles depend more on the hosting provider than on geographic location, deploying a cloud-based network telescope in only one cloud provider introduces a systematic observation bias. Comprehensive cloud threat monitoring therefore requires capturing IBR across multiple providers, not just regions.

Methodologically, we use external blocklists as a proxy ground truth to determine scanner detection thresholds. Albeit imperfect, this empirical baseline offers a better alternative to arbitrary parameter tuning, and demonstrates that telescopes capture many active scanners missing from public feeds.

Cloud telescopes also offer a favorable cost-visibility trade-off. The scanners observed per monitored IP in CLOUD-NT exceed those in UCSD-NT by three orders of magnitude. While scanners-per-IP is an unconventional metric for classical telescopes, it serves as a practical measure of cost-efficiency in cloud environments where operators pay per allocated address. This efficiency arises because cloud IP space is actively targeted, unlike the IPv4 prefixes that host classical telescopes. A subsampling analysis (Appendix A) further shows sub-linear scaling. Over our measurement period, each additional USD of VM rental yielded $\approx$1000 new source IPs at moderate deployment sizes, falling to $\approx$350 at full scale, while 200 VMs already recover $\approx$74% of full-scale visibility.

Our results suggest that a community-run cloud telescope would provide benefit to the research community. The funding seems to be an entirely solvable issue if several groups partner or win small-sum support of various cloud providers. Our setup already demonstrated meaningful visibility at modest cost. We are aware that providers generally hesitate to share data about attacks on their networks, but releasing such data from a community telescope to the academic community, under a framework governing use of the data, disclosure and academic publications, seems unlikely to be a major concern.

VII. Summary

This paper compares Internet Background Radiation (IBR) captured by a multi-provider cloud network telescope (CLOUD-NT) with two classical network telescopes (UCSD-NT and SURF-NT). We found that IBR is highly provider-dependent, with scanners frequently targeting specific cloud environments rather than geographic regions. While classical telescopes remain necessary for observing widespread background radiation and RSDoS attacks, they could miss targeted enterprise scans. Some scanners targeting the cloud exhibited deep vertical enumeration, primarily originating from other hosting providers rather than ISPs, and were often absent from IP blocklists. Our findings demonstrate that multi-cloud deployments provide a highly efficient complement to classical darknets for monitoring modern, service-specific threats.

Acknowledgements

We thank our shepherd and the anonymous reviewers for their constructive feedback. We also extend gratitude to SURF for providing the telescope data and to the CATRIN project for providing additional data. This material is based on research sponsored by the National Science Foundation (NSF) grants CNS-2212241, CNS-2450552, OAC-2319959, and OAC-2531134. Parts of the calculations for this publication were performed on the HPC cluster PALMA II of the University of Münster, subsidised by the DFG (INST 211/667-1). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the funding agencies.

References

[1] D. Moore *et al.*, "Network Telescopes: Technical Report," CAIDA, Tech. Rep., 2004.

[2] P. Sattler *et al.*, "Packed to the brim: Investigating the impact of highly responsive prefixes on Internet-wide measurement campaigns," in *Proc. ACM CoNEXT*, 2023.

[3] Internet Society, "2019 Internet Society Global Internet Report: Consolidation in the Internet Economy," Internet Society, Report, 2019.

[4] J. Arkko *et al.*, "Considerations on Internet consolidation and the Internet architecture," IETF, Internet Draft draft-arkko-iab-internet-consolidation-02, 2019.

[5] M. Fayed *et al.*, "The ties that un-bind: Decoupling IP from web services and sockets for robust addressing agility at CDN-scale," in *Proc. ACM SIGCOMM*, 2021.

[6] L. Izhikevich *et al.*, "Cloud watching: Understanding attacks against cloud-hosted services," in *Proc. ACM IMC*, 2023.

[7] P. Richter and A. Berger, "Scanning the scanners: Sensing the Internet from a massively distributed network telescope," in *Proc. ACM IMC*, 2019.

[8] K. Benson *et al.*, "Leveraging Internet background radiation for opportunistic network analysis," in *Proc. ACM IMC*, 2015.

[9] A. Männel *et al.*, "Lessons learned from operating a large network telescope," in *Proc. ACM SIGCOMM*, 2025.
[10] F. Bortoluzzi *et al.*, "Cloud Telescope: A distributed architecture for capturing Internet background radiation," in *Proc. IEEE CloudNet*, 2023.
[11] E. Pauley, P. Barford, and P. McDaniel, "DScope: A cloud-native Internet telescope," in *Proc. USENIX Security*, 2023.
[12] R. Pang *et al.*, "Characteristics of internet background radiation," in *Proc. ACM IMC*, 2004, pp. 27–40.
[13] A. Dainotti *et al.*, "Analysis of a "/0" stealth scan from a botnet," in *Proc. ACM IMC*, 2012, pp. 1–14.
[14] B. Degen *et al.*, "Through a smaller lens: Revisiting opportunistic analysis using network telescopes," in *Proc. PAM*, 2026.
[15] D. Wagner *et al.*, "How to operate a meta-telescope in your spare time," in *Proc. ACM IMC*, 2023.
[16] E. Chan *et al.*, "Analyzing Internet background radiation with reflective network telescopes," in *Proc. ACM/IRTF ANRW*, 2025.
[17] L. Miao *et al.*, "Extracting Internet Background Radiation from raw traffic using greynet," in *Proc. IEEE ICON*. Singapore: IEEE, Dec. 2012, pp. 370–375.
[18] R. Hiesgen *et al.*, "Spoki: Unveiling a New Wave of Scanners through a Reactive Network Telescope," in *Proc. USENIX Security*, 2022.
[19] D. Ferrero *et al.*, "Revealing informed scanners by colocating reactive and passive telescopes," in *Proc. IEEE RAID*, 2025.
[20] CAIDA, "FlowTuple," https://www.caida.org/projects/network_telescope/docs/data/flowtuple/, Mar. 2021.
[21] CAIDA, "Corsaro3," https://github.com/CAIDA/corsaro3, Mar. 2024.
[22] Network Startup Resource Center, "Routeviews," https://archive.routeviews.org/, 2025.
[23] "IPinfo | The Trusted IP Data Provider for Developers & Enterprises," https://ipinfo.io/.
[24] O. Darwich *et al.*, "Replication: Towards a Publicly Available Internet Scale IP Geolocation Dataset," in *Proc. ACM IMC*, Oct. 2023.
[25] OpenIntel, "Active dns measurement project," https://www.openintel.nl/data/reverse-dns/.
[26] M. Luckie *et al.*, "Learning to extract geographic information from internet router hostnames," in *Proc. ACM CoNEXT*, Dec. 2021.
[27] C. Tsaousis, "FireHOL IP Lists," http://iplists.firehol.org.
[28] SANS Internet Storm Center, "SANS.edu Internet Storm Center," https://isc.sans.edu/index_dyn.html.
[29] V. G. Li *et al.*, "Reading the tea leaves: A comparative analysis of threat intelligence," in *Proc. USENIX Security*, 2019.
[30] J. Mazel and R. Strullu, "Identifying and characterizing ZMap scans: A cryptanalytic approach," Aug. 2019.
[31] S. Lagraa *et al.*, "Deep mining port scans from darknet," *Int. J. Netw. Manag.*, vol. 29, no. 3, May 2019.
[32] E. Bou-Harb *et al.*, "Cyber Scanning: A Comprehensive Survey," *IEEE Commun. Surv. Tutor.*, vol. 16, no. 3, 2014.
[33] A. V. C. Camargo *et al.*, "Less is More? Exploring the Impact of Scaled-Down Network Telescopes on Security and Research," in *Proc. SBRC*, May 2024.
[34] D. Moore *et al.*, "Inferring Internet denial-of-service activity," *ACM Trans. Comput. Syst.*, vol. 24, no. 2, 2006.
[35] R. Hiesgen *et al.*, "The age of DDoScovery: An empirical comparison of industry and academic DDoS assessments," in *Proc. ACM IMC*, 2024.
[36] H. Heo and S. Shin, "Who is knocking on the Telnet Port: A Large-Scale Empirical Study of Network Scanning," in *Proc. ACM ASIACCS*, May 2018.
[37] E. Bou-Harb *et al.*, "Multidimensional investigation of source port 0 probing," *Digital Investigation*, 2014.
[38] M. Luchs and C. Doerr, "The curious case of port 0," in *Proc. IFIP Networking*. IEEE, 2019.
[39] A. Maghsoudlou *et al.*, "Zeroing in on port 0 traffic in the wild," in *Proc. PAM*, 2021.
[40] M. Bailey, D. Dittrich, E. Kenneally, and D. Maughan, "The menlo report," *IEEE Security & Privacy*, 2012.

APPENDIX

COMPLETE JACCARD SIMILARITY HEATMAP

Fig. 7 shows the complete Jaccard similarity matrix across all monitored cloud regions. The diagonal structure indicates that intra-provider source IP overlap consistently exceeds inter-provider overlap, regardless of geographic proximity. Regions are grouped by provider (e.g. AWS, Azure, GCP) to highlight this effect. We also identified outliers in the data. The dark line in *Azure* represents the `koreacentral` region. This region is paired with the `koreasouth` region, which did not display the same anomaly. Even though our VMs in `koreacentral` did not share the same announced prefix as other Azure nodes, they were located within the same AS (8075) and used IP addresses from shared prefixes. An analysis of the source countries of the IP addresses contacting these VMs revealed a distinct difference; the `koreacentral` region received vastly more traffic from China than any other *Azure* region. In fact, we observed more than 30 times the number of source IPv4 addresses located in China in `koreacentral` than we observed on average across all other *Azure* regions. Similarly, we observed outliers in *DigitalOcean's* `ams3` and `lon1` regions, where we found substantially more source IPs originating from Pakistan, Ireland, and Nigeria (`ams3`) or from Argentina and France (`lon1`).

CLOUD-NT

Pricing: Although we deployed CLOUD-NT in 2025, we evaluate costs using 2026 regional list prices, without any discounts or free-tiers. The GCP *e2-micro* instances (2 shared vCPUs, 1 GB RAM) total approximately 25.70 USD per day. Azure's equivalent *Standard_B2ats_v2* yields a daily cost of roughly 26.50 USD, while AWS's comparable *t3a.micro* costs 19.37 USD. Furthermore, our deployment utilizing DigitalOcean's smallest droplets costs 1.99 USD per day, and the Vultr *vc2-1c-1gb* instances amount to a total daily cost of 5.37 USD.

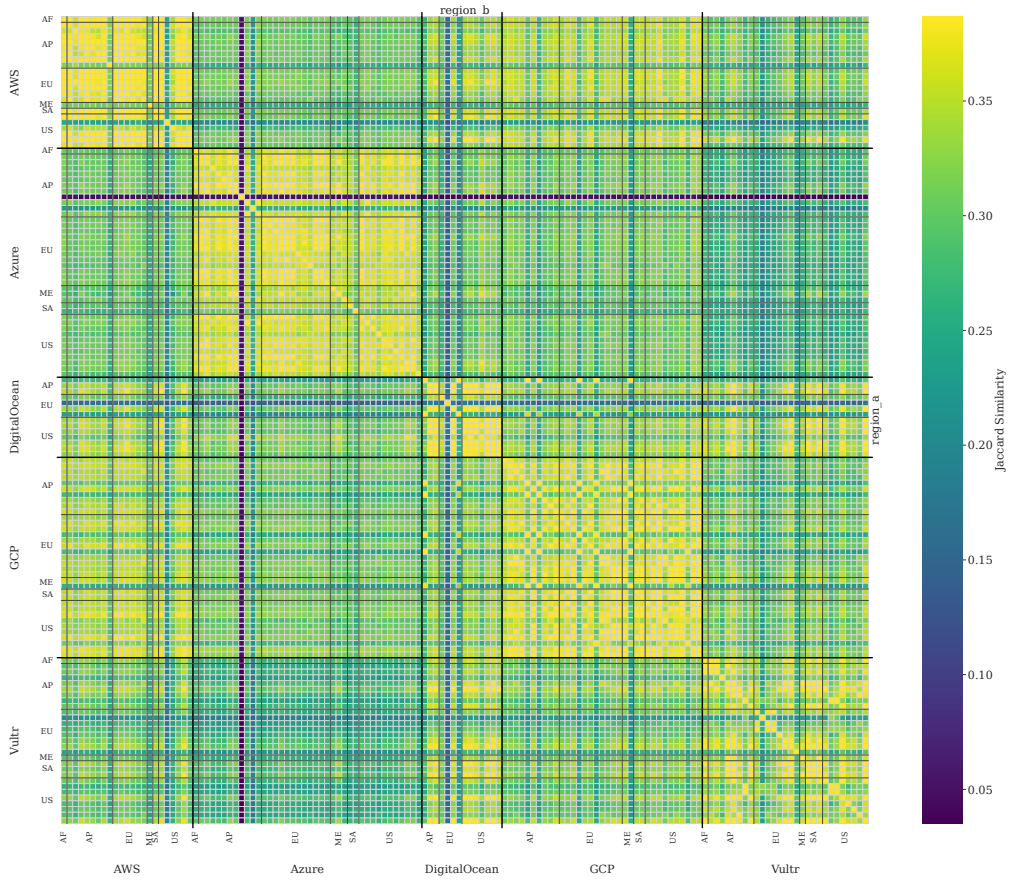

Fig. 7: Complete Jaccard similarity matrix of source IPs across all cloud regions. Diagonal blocks (intra-provider) show markedly higher similarity than off-diagonal blocks, even for co-located data centers.

TABLE VI: Overview of Reference Datasets

Dataset Name	Information Extracted	Dataset Date
IPInfo	Geolocation, ASN type	2025-04-28
IP2Asn	Autonomous System Number (ASN)	2025-04-18
OpenIntel	Reverse DNS records (PTRs)	2025-04-24
Hoiho	Reverse DNS context (PTRs)	-
FireHOL	Blocklisted IPs	2025-04-18
ISC SANS	Research scanner IPs	2026-02-09

ISC does not offer historical lists. We filter to only include IPs already present during our observation period.

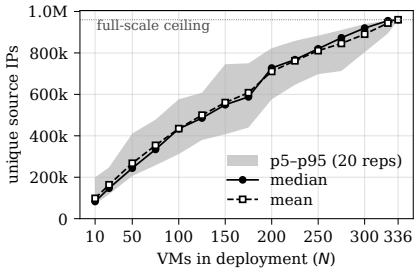

Fig. 8: Unique src. IPs observed when subsampling CLOUD-NT to N VMs (20 random draws per N; band: p5–p95).

Across all providers, operating the full 336-VM deployment for our measurement period would cost approximately 1421 USD in 2026. To translate these per-VM costs into a cost-per-visibility figure, we subsampled the CLOUD-NT VM set to $N \in \{10, 25, \ldots, 325, 336\}$, drawing 20 uniform-random samples per size and counting distinct source IPs over the full observation window. Fig. 8 shows that visibility scales sub-linearly. Each USD of VM rental yields ≈2300 unique source IPs at small deployments, falling to ≈680 at full scale, and 200 randomly drawn VMs already recover ≈74% of full-scale visibility. Note that uniform IP-level subsampling treats all VMs equally and therefore underestimates the value of adding additional providers.

Deployment: The necessary *Terraform* scripts to recreate our cloud deployment are available at: `https://github.com/CAIDA/cloudtelescope-public`.

Ethical Considerations

In this study, we adhered to the ethical guidelines established by the Menlo Report [40]. Our data collection strictly involved unsolicited Internet Background Radiation (IBR) using well-established network telescope methodologies prevalent over the past two decades [1]. Because these IPs hosted no legitimate services, the captured traffic primarily consists of automated scanning, backscatter, and misconfigurations. In the event that misguided legitimate traffic reached

our monitors, privacy was preserved by strictly analyzing flow data rather than inspecting packet payloads. We minimized potential harm to the network ecosystem by ensuring our network telescopes were strictly passive; unlike interactive honeypots, our deployments did not respond to, provoke, or amplify malicious traffic. All deployments were conducted in compliance with the Terms of Service and Acceptable Use Policies of the respective cloud providers, and the resulting insights are shared to improve the understanding of IBR.