

Split Skies: Operational Experiences Extending a Multi-Cloud Network Telescope into the Asia-Pacific

Nils Kempen
University of Münster
Münster, Germany

Benjamin Schumann
University of Twente
Enschede, The Netherlands

Ralph Holz*
University of Münster
Münster, Germany

Abstract

Comprehensive visibility into Internet Background Radiation (IBR) requires diverse measurement infrastructure. Recent work showed that cloud IBR depends on the provider making single-provider telescopes systematically biased. These studies relied entirely on Western providers and left the Asia-Pacific region as a blind spot. In this report we extend a multi-cloud network telescope onto Tencent Cloud inside mainland China. We detail how Internet fragmentation, geopolitics, and provider rules complicate global measurement through strict onboarding barriers and an unreliable management plane. Once deployed, our nodes captured high-volume regional scanning invisible to our Western vantage points. We conclude by connecting our findings to the Responsible Internet framework. This framework relies on independent global observation to build a transparent Network Inspection Plane. We show that building this plane is bottlenecked by sovereign ecosystems and requires a collaborative community effort.

CCS Concepts

• **Networks** → *Network security*; **Network measurement**.

Keywords

Internet Background Radiation, Network Telescope, Cloud Computing, Internet Measurement, Internet Fragmentation

ACM Reference Format:

Nils Kempen, Benjamin Schumann, and Ralph Holz. 2026. Split Skies: Operational Experiences Extending a Multi-Cloud Network Telescope into the Asia-Pacific. In *First International Workshop on a Resilient, Renewable, and Responsible Internet (RE-NET '26)*, December 07–10, 2026, Utrecht, Netherlands. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3843201.3856200>



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

RE-NET '26, Utrecht, Netherlands

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2981-2/2026/12

<https://doi.org/10.1145/3843201.3856200>

1 Introduction

A resilient Internet depends on our ability to observe it. The work by Hesselman et al. established that independent measurement is essential to build a “Responsible Internet” (RI) and achieve network transparency [4]. Network telescopes capturing Internet Background Radiation (IBR) (unsolicited traffic that reaches any routed address) have long been used for observing events on the Internet. They have allowed the detection of scanning campaigns, backscatter from denial-of-service attacks, misconfigurations, as well as Internet outages [8] and measuring these events provides the independent measurement data necessary to support the transparency goals defined in the RI framework.

For two decades, *classical network telescopes* (darknets) that passively capture traffic to large blocks of unused address space have been used for these observations [9]. Unfortunately their sustainability is increasingly threatened. Operators reclaim dark address blocks to meet demand for scarce IPv4 space. The largest publicly available network telescope (UCSD-NT) run by CAIDA, for example, has shrunk to 63.56% of a full /8 [8]. At the same time, scanners are aware of network telescope ranges and shift their scans away from uniformly random targets toward dense, high-value cloud infrastructure [6, 10, 11]. In response, researchers have begun to examine the cloud itself as a vantage point for IBR observation [1, 6, 10]. In our own prior work, we deployed *CLOUD-NT*, a passive multi-cloud network telescope spanning five providers, and found that IBR in the cloud is *provider-dependent rather than geography-dependent* [7]. A telescope hosted on a single provider therefore carries a systematic observational bias. This finding has a practical implication. If the provider determines what a telescope sees, then the set of providers a researcher can actually deploy on determines the limits of visibility. That set is shaped not by science but by geo-politics, provider rules, and verification and payment barriers. These Internet fragmentation problems, already identified in [4], restrict who can observe the network and hinder diverse global observation needed to achieve infrastructure transparency. CLOUD-NT, like most prior cloud telescope work, was built entirely on Western providers (AWS, GCP, Azure, Vultr, and DigitalOcean). For

*Also with University of Twente.

cloud-hosted telescopes the Asia-Pacific region, and mainland China in particular, remained a blind spot. To our knowledge, no prior cloud telescope deployment observed IBR from within mainland China. We set out to extend our cloud-based network telescope to additional providers in these regions.

This paper is an experience report on a second deployment wave of the CLOUD-NT effort that set out to close that gap by extending the telescope eastward onto Tencent Cloud. We describe what this took operationally and what the new vantage point revealed. Our contributions are (i) an account of deploying passive telescope nodes in a politically constrained cloud ecosystem, covering a failed Alibaba Cloud attempt, the pivot to Tencent, per-provider infrastructure divergence, and in-country network interference, distilled into three general lessons (**L1** to **L3**) about Internet fragmentation [2] (§3), (ii) evidence that the Asia-Pacific vantage point captures high-volume activity that our earlier Western-provider deployments missed (§4), and (iii) recommendations for building resilient, representative, and responsibly operated community monitoring infrastructure, which serves the RI as an independent observer.

2 Background on CLOUD-NT

This report builds directly on our CLOUD-NT effort, and we recap only what is needed to follow it, referring to [7] for the full design and analysis. CLOUD-NT is a strictly passive multi-cloud network telescope whose Terraform scripts provision the smallest available virtual machine (VM) in as many regions and availability zones as possible across five providers (AWS, GCP, Azure, Vultr, and DigitalOcean), each listening on a single public IPv4 address. The security groups forward all IPv4 ingress while host firewall rules silently drop every reply, and management runs over IPv6. The first deployment ran 336 VMs across 143 regions for two weeks in April 2025 and was compared against two classical telescopes, the widely known UCSD-NT (a /9 and a /10) and a publicly unknown one, SURF-NT (a /16) [7].

Two of its findings motivate this work. *i*) IBR in the cloud is *provider-dependent rather than geography-dependent*. Traffic collected within a single provider is more self-similar than traffic collected across providers in the same region, so a single-provider telescope is systematically biased. *ii*) the scanner populations seen by cloud and classical telescopes barely overlap. Only 0.60% of the 941k scanners observed were common to all three telescopes, and cloud IPs are disproportionately targeted on cloud-service ports. Neither vantage point subsumes the other. The CLOUD-NT discussion argued, on this basis, for multi-provider deployments and a community-run cloud telescope with a data-governance framework. Yet the providers it used were all Western. Whether provider-dependence extends into the

Asia-Pacific, and how far region matters there, was left open. This is the gap the deployment described here set out to probe.

3 Extending CLOUD-NT Eastwards

Our goal was to extend CLOUD-NT's coverage into the Asia-Pacific, and mainland China specifically, so that we could test whether its provider-dependence finding held beyond the Western cloud market. The second deployment run encompassed 28 VMs in total, 14 on Tencent Cloud (China and Southeast Asia) and 14 on DigitalOcean (North America, Europe, and parts of the Asia-Pacific), across 11 countries from May 18 to June 6, 2025 (20 days). This deployment ran on these two providers only. DigitalOcean served as the multi-region comparison provider that we already knew how to operate, its VMs spanning North America, Europe, and parts of the Asia-Pacific, while Tencent Cloud served as the new Asia-Pacific vantage point. The effort, and the lessons, were almost entirely on the Tencent side. We recount them here because each is a concrete instance of a general obstacle that Internet fragmentation and platform opacity place in the path of global measurement. We number the lessons **L1** to **L3** and return to them in §5.

L1: Who gets to measure is constrained by locally determined business practices. Initially we chose to use Alibaba Cloud as our next cloud provider, because it is the largest provider in the region and could give us a good view into Chinese networks, but we never managed to deploy our network telescope there. Provisioning infrastructure in mainland China required verification procedures we could not easily satisfy from abroad. Even after we managed to successfully create an account on Alibaba Cloud, it was quickly flagged for suspicious activity following our initial API calls. It could only be unblocked with bank statements that the provider judged to show insufficient funds. Provisioning was severely delayed by documentation and management consoles designed exclusively for a domestic market. Translation tooling helped but could not eliminate the substantial operational friction. In the end, we decided to abandon Alibaba and pivoted to Tencent, the region's second-largest provider, where an account could be established and the deployment made to work. Hence, the technical design of the cloud telescope was not the bottleneck, but access to the provider was. For measurement research, this means the vantage points a scientist may occupy are decided by provider policies across different jurisdictions and strict identity requirements. Such onboarding barriers create critical blind spots in global IBR analysis.

L2: The software is portable, the provisioning code is not. The software that runs on the VMs after provisioning, our capture and management scripts, is fully provider-agnostic and moved onto Tencent relatively unchanged. The Terraform provisioning code on the other hand had to be adjusted per-provider, and even though CLOUD-NT already had support for five providers, extending it to Tencent required substantial new code for a different resource model, distinct API and security-group conventions, and an unfamiliar account and billing structure. The general lesson is that provisioning code encodes the inner workings of the ecosystems it was written for, so there is a larger porting cost for distinctly new and different ecosystems. For an RI, this implies that global observation infrastructure cannot be developed unilaterally. Achieving global Internet transparency will require active collaboration across regions to bridge these different technical ecosystems.

L3: The management plane inside China is unreliable. Even once the VMs were running, operating them from outside was unreliable and we faced problems that our other deployments did not have. Public IPv6, on which CLOUD-NT relies for management and data-transfer, was inconsistently available and required manual intervention on a per-VM basis. In-country network conditions interfered with routine tooling. DNS resolution was blocked for some destinations and access to *GitHub* (which hosted some of our binaries) was throttled, so provisioning steps that relied on this had to be reworked around unreliable egress. These problems likely result from Chinese national network controls. Previous research shows these controls are decentralized and highly variable [3]. This unpredictability could explain our inconsistent network access and highlights why establishing reliable observation nodes is so difficult. During this manual setup we also misconfigured the VMs' firewall, allowing only inbound TCP and UDP, blocking external ICMP on the Tencent side, whereas DigitalOcean carried no such rule and received it. We flag the resulting ICMP gap as a limitation of the dataset in §4. The same national filtering that disrupted our management scripts also inherently filters the incoming background radiation we aim to capture. Operating a telescope in these networks requires accounting for local restrictions and data biases.

4 What the New Vantage Point Sees

In this section we analyze what the new vantage point brings us in terms of visibility and how it compares to one of our Western deployments.

Method. We processed the raw packet captures into CAIDA FlowTupleV3¹ records with *Corsaro*² and report packet counts throughout, summed over the observation period from May 18 to June 6, 2025, which is 20 full UTC days. We approximate the three IBR traffic classes from TCP flags. FlowTuple records carry per-packet header information, so we assign a class to each packet on its own. A pure *TCP SYN* or any *UDP* packet counts as scanning, a *TCP ACK*, *TCP RST*, or *TCP FIN* as backscatter, and everything else, including *ICMP*, as other. On top of this per-packet classification we ran CLOUD-NT's sliding-window scan detection, which labels per-source 15-minute windows as port, host, or random scan events [7]. We determined scanner detection thresholds by running the detection algorithm with different thresholds, deciding on the one that lies on the elbow point of the detected scanner counts, as we describe in detail in [7]. We validated this calibration against two blocklists, FireHOL's attacks category³ and the ISC SANS research-scanner list⁴. We run the scanner detection with a common threshold, so that differences between the providers cannot stem from detection settings. We exclude provider-internal traffic from special-use source addresses from every result below. We further exclude two single-node, single-port streams whose volume and shape match the reused-address spillover we described in [7], because their massive volume would otherwise distort the overall baseline.

More pressure from a similar source pool. Averaged over the window, the 14 DigitalOcean VMs together received $289\text{ k} \pm 36.1\text{ k}$ packets per day, of which $283\text{ k} \pm 37\text{ k}$ (98%) were scanning. The 14 Tencent VMs received $540\text{ k} \pm 80.6\text{ k}$ packets per day, 536 k (99.30%) of them scanning. Scanning traffic accounts for almost this entire volume gap. The Tencent VMs saw 1.90 times the packet volume of DigitalOcean, yet from a source pool of similar size. DigitalOcean VMs observed $17.3\text{ k} \pm 1173$ distinct source IPs per day against Tencent's $21.6\text{ k} \pm 4927$, so Tencent drew 1.90 times the packets from only about 1.30 times the sources. Sources send on average more packets toward our Tencent VMs than they do to our DigitalOcean VMs, 25.60 against 16.80 packets per source. The scan detection shows a similar picture. We find 148 k scan events on the Tencent side against 42.2 k on DigitalOcean, a factor of 3.50. Figure 1 shows that scanning dominates the traffic for both providers, that the Tencent volume is steadier over time, and that the event asymmetry holds across the whole window.

Sources concentrate by provider. Among the top 30 source ASes by packet count (Fig. 2 shows the top 15), 12 of them,

¹https://www.caida.org/projects/network_telemetry/docs/data/flowtuple/

²<https://github.com/CAIDA/corsaro3>

³<http://iplists.firehol.org>

⁴https://isc.sans.edu/index_dyn.html




Figure 1: Packets per 6-hour bin by traffic class for each provider (top two panels, shared log scale) and detected scan events per day and provider (bottom panel).




Figure 2: Top 15 source ASes by packet count, split by provider, log scale.

and 5 of the top 10, delivered more than 90% of their packets to a single provider. The few ASes that split evenly are globally deployed infrastructure networks such as Google (AS 396982) at 49.20% and Akamai (AS 63949) at 47.80% Tencent share. This concentration is consistent with CLOUD-NT’s provider-dependence [7] extending into the Asia-Pacific, though in this setup provider and region coincide, so we cannot yet tell whether sources target China as a region or Tencent as a provider. The same pattern holds looking at individ-




Figure 3: Pairwise Jaccard similarity of distinct source IP sets across 28 VMs. Axes are grouped by provider and region. DigitalOcean spans EMEA, NAFTA, and APAC. Tencent spans China and SEA.




Figure 4: Overlap of distinct source IPs between deployments. Rows split the IPs into those unique to DigitalOcean, unique to Tencent, and shared. Values show the union size and Jaccard similarity for all sources, all scanners, and three scan types.

ual source IPs. Figure 3 shows the pairwise Jaccard similarity of the per-VM source sets, mirroring an analysis from [7]. Mean pairwise similarity is 0.38 between DigitalOcean VMs, 0.30 between Tencent VMs, and 0.27 across providers, with off-diagonal values from 0.18 to 0.45. The DigitalOcean VMs form a clear block, more similar to one another (0.38) than to Tencent VMs (0.27). The Tencent VMs are only weakly more similar to each other (0.30) than across providers, so the block structure is asymmetric. Still, which sources a VM shares with another is influenced more by provider than by location, extending the AS-level concentration down to individual sources.

The deployments detect largely disjoint scanner sets. The detected scanner sets overlap only marginally. The scan detection identifies 3181 scanner source IPs at DigitalOcean and 4896 at Tencent, of which only 597 are shared, a Jaccard similarity of 0.08 (Figure 4) against 0.21 across all public sources, so the scanners overlap less than the overall IBR

traffic. The two blocklists used also cover the two populations unevenly. FireHOL’s attacks category lists 57.40% of the Tencent scanners but only 24.60% at DigitalOcean, while the ISC SANS research-scanner list shows the opposite, 6.60% at DigitalOcean against 2.50% at Tencent. Both lists together still miss most DigitalOcean scanners (71.10%) and a large share of Tencent’s (41.30%).

Port traffic volumes vary by vantage point. While a network telescope eventually receives IBR on all ports, the number of packets arriving per port also depends on the vantage point (Figure 5). Widely scanned services reach both providers at similar levels, for example on ports TCP/22 (SSH) and TCP/3389 (RDP) we observe roughly the same number of packets on both providers ($\approx 100k$ and $\approx 40k$ respectively). Other ports are different across providers. SMB port TCP/445 is heavily targeted on Tencent, 180 k against 29.6 k packets, and Telnet port TCP/23 is also more targeted on Tencent. A few ports are single-node spikes rather than provider effects, such as port 3000 with 64.1 k of its 75.2 k packets on one DigitalOcean node (fra1) and port 32580 with 45.1 k of 46.1 k packets on one Tencent node (ap-guangzhou-6).

Region adds a second axis. While providers remain the dominant factor influencing received IBR, geographic region adds a secondary axis of influence. We observe a regional effect within individual providers. Our DigitalOcean VMs in EMEA and NAFTA, for example, have very dissimilar scanner populations. With only three VMs per region, though, we cannot draw any specific conclusions from this observation. Clearer evidence is an event aimed at one region. Between May 22 and 25 we observed a distributed UDP campaign from Psychz Networks (AS 40676), and it targeted mainland China specifically. Over the full observation window this AS sent 502 k packets to almost all UDP ports, 89.20% of them in those four days. 481 k packets (95.70%) landed on the seven Tencent VMs inside mainland China and only 2590 on the seven Tencent VMs outside it in Jakarta, Singapore, and Bangkok. Both sets of VMs ran on Tencent, so this split is regional and not a provider effect. DigitalOcean saw 18.9 k packets in total, its three Asia-Pacific VMs each ≈ 250 , no more than any other DigitalOcean VM. During this burst 447 k packets landed on Tencent against 727 on DigitalOcean. Neither provider diversity nor regional presence alone would have caught it. Only a VM inside the targeted region did. This is a nation-targeted instance of localized scanning, which Richter and Berger showed makes up a substantial share of scan traffic and stays invisible to telescopes outside the targeted region [11].

Traffic the platform injects. We also observed a large amount of internal ICMP traffic on the Tencent VMs, sent by the platform itself. Every Tencent VM received ICMP echo requests

from the provider’s link-local range at a fixed rate of about one every three seconds, steady across the whole observation period (16.03 million packets in total from 14 addresses). DigitalOcean saw no comparable packets. This is the signature of a provider health check probing node liveness, which we could not disable. This is why Tencent’s public other class in Figure 1 stays small and flat. Left in, the health-check stream would account for 60% of Tencent’s total packet volume and drown out the *public other* class in Fig. 1, so we exclude it. As noted in §3, our firewall misconfiguration blocked external ICMP on Tencent, whereas DigitalOcean captured this background Internet traffic. The missing ICMP is a dataset limitation, as it makes the other class incomparable across providers.

Discussion. These results extend CLOUD-NT’s provider-dependence finding eastward and give a first look inside IBR as seen from within Chinese networks. We observe the same provider-dependent traffic structure in the Asia-Pacific, and on top of it we see effects a Western-only deployment would miss, from the campaign aimed at mainland China to the platform’s own injected traffic. This is only a 20 day two-provider snapshot which can serve as a starting point for further work. Because our Tencent nodes sit in China and Southeast Asia and our DigitalOcean nodes mostly in the West, provider and region largely coincide, and only a deployment with several providers per region can say how much of what we see is truly regional. A community telescope, which we argue for in §5, would also show whether Tencent generalizes to other Chinese providers such as Alibaba, and whether Western providers in the region (like our DigitalOcean nodes in Bangalore and Singapore) observe similar traffic patterns. It provides distributed observation an RI needs to enable transparency in cloud infrastructures [4]. Our experience implies that this structure is needed but hard to build for non-technical reasons (§5). One finding already stands on its own. The blocklists fit our two providers very differently. This again hints at targeted scanning and blocklist biases. A single-provider telescope thus inherits not just a skewed traffic mix but a blocklist that matches it unevenly.

5 Lessons and Recommendations

We distill our experience into recommendations for building a diverse cloud-based network telescope, which we think is needed as an addition not only to classical network telescopes, but also to the independent observation a Responsible Internet needs. Our deployment shows that this observation tool is hard to build. To succeed across a fragmented global network, this observation infrastructure must be both resilient and responsibly operated.




Figure 5: Packets per destination IP (rows, DigitalOcean block on top, Tencent below) for the union of each provider’s 30 most-targeted destination ports (37 total), log color scale, empty cells mean zero packets. The empty port-443 cell in the lon1 row is an excluded SYN-flood stream.

For a transparent and resilient Internet. The Responsible Internet relies on a Network Inspection Plane to build Internet transparency. This plane requires independent observers to document Internet security incidents and abuse [4]. Network telescopes provide this kind of data by capturing IBR. Classical telescopes, however, miss modern cloud-specific IBR, and cloud-based telescopes fill this gap. The RI explicitly requires data from multiple vantage points and our findings reinforce this requirement by showing that IBR depends on the provider and the geo-political region. Single-region deployments miss events aimed at other regions entirely and introduce a severe bias. Representative monitoring, therefore, has to be multi-provider and multi-region. This diversity also makes the observation infrastructure itself resilient. A network telescope as diverse as the IBR it captures is harder to evade and survives the loss of any single node. But building this required diversity proves to be difficult. At 2026 list prices our whole 28-VM deployment would only cost about \$120, but deployments must budget not only for VM-hours but for bureaucracy as well. Strict verification rules can rule out a vantage point entirely (L1) and operating inside heavily regulated networks requires managing an unreliable network plane (L3). Furthermore, porting code across divergent ecosystems is not just a change of credentials, but requires extensive knowledge about platforms and ecosystems (L2). These barriers, representing the Internet’s fragmentation, leave us with the conclusion that no single research group can clear these hurdles globally. A community-run cloud telescope therefore seems like the only viable way to pool the onboarding effort and the resulting visibility, necessary to provide a global view of cloud IBR.

For a responsibly operated observer. Cloud-based network telescopes serve as an addition to classical network telescopes, as shown in [7]. Therefore they must use the same passive methods. While reacting to probes is a valid technique [5], it alters scanner behavior and breaks our comparison to classical telescopes. Staying passive also aligns with

the RI mandate to balance transparency against security risks. We minimize data collection and apply equal care across all regions. This duty is even more critical for a community telescope pooling shared data. Responsibility also demands transparency about the biases a deployment introduces. We must therefore empirically verify what each destination receives and actively disclose artifacts like blocked ICMP (§3) or Tencent’s internal probes (§4). Unreported biases invalidate the observation. Ultimately our experience shows how fragmentation impedes global security research.

6 Conclusion

In this report we extended our multi-cloud network telescope CLOUD-NT into the Asia-Pacific and mainland China. Setting up our telescope with local providers was hindered by locally determined business practices and strict onboarding barriers rather than technical design. Our pivot to Tencent and the unreliable management plane inside China show how Internet fragmentation restricts who can observe the network (§3). Once deployed, our new nodes captured high-volume regional scanning invisible to our Western comparison deployment (§4).

Our experience directly informs the design of a Resilient and Responsible Internet. Network resilience requires diverse observation infrastructure to track localized threats. The Responsible Internet framework also depends on independent measurement to build a transparent Network Inspection Plane. Our deployment shows that building this global observation layer is bottlenecked by geo-politics as well as provider restrictions. Overcoming these borders is too complex for any single research group. We urge the community to collaboratively build the distributed and responsibly operated monitoring infrastructure an RI requires.

References

- [1] Fabricio Bortoluzzi et al. 2023. Cloud Telescope: A Distributed Architecture for Capturing Internet Background Radiation. In *Proc. IEEE CloudNet 2023*.

- [2] William J. Drake et al. 2016. *Internet Fragmentation: An Overview*. Technical Report. World Economic Forum.
- [3] Roya Ensafi et al. 2015. Analyzing the Great Firewall of China Over Space and Time. *PoPETs 2015*.
- [4] Cristian Hesselman et al. 2020. A Responsible Internet to Increase Trust in the Digital World. *JNSM* 28, 4 (2020).
- [5] Raphael Hiesgen et al. 2022. Spoki: Unveiling a New Wave of Scanners through a Reactive Network Telescope. In *Proc. USENIX Security 2022*.
- [6] Liz Izhikevich et al. 2023. Cloud Watching: Understanding Attacks Against Cloud-Hosted Services. In *Proc. ACM IMC 2023*.
- [7] Nils Kempen et al. 2026. A Multi-Cloud View of Internet Background Radiation. In *Proc. IFIP TMA 2026*.
- [8] Alexander Männel et al. 2025. Lessons Learned from Operating a Large Network Telescope. In *Proc. ACM SIGCOMM 2025*.
- [9] David Moore et al. 2004. *Network Telescopes: Technical Report*. Technical Report. CAIDA.
- [10] Eric Pauley, Paul Barford, and Patrick McDaniel. 2023. DScope: A Cloud-Native Internet Telescope. In *Proc. USENIX Security 2023*.
- [11] Philipp Richter and Arthur Berger. 2019. Scanning the Scanners: Sensing the Internet from a Massively Distributed Network Telescope. In *Proc. ACM IMC 2019*.