The Observer Effect in Internet Measurement: A Case Study of Route Origin Validation

Abbas Mohsenpour
UCLouvain
Louvain-la-Neuve, Belgium

Cristel Pelsser
UCLouvain
Louvain-la-Neuve, Belgium

Abstract

Internet observatories publish reports that shape how networks are assessed and perceived, for example in terms of their security compliance. To ensure transparency and reproducibility, they rely on public, long-lived infrastructure that the networks under evaluation can recognize. We study *measurement awareness*, defined as a network treating recognizable measurement targets differently from functionally equivalent but non-recognizable ones. Measurement awareness describes observable behavior; unlike measurement gaming, it makes no assumption about intent. Using Route Origin Validation (ROV) as a case study, we compare publicly known RPKI-invalid targets against functionally equivalent but undisclosed shadow targets. We identify networks that reject the well-known invalid prefixes while accepting the equivalent shadow prefixes. During responsible disclosure, one operator confirmed that it explicitly filters the prefix used by the *Is BGP Safe Yet?* observatory. Observatories may thus capture not the properties they intend to measure, but the behavior networks expose to them.

CCS Concepts

• **Networks** → **Network measurement**; *Routing protocols*; Network monitoring; • **Security and privacy** → *Security protocols*.

Keywords

Internet measurement, measurement-aware networks, Route Origin Validation, BGP, routing security

ACM Reference Format:
Abbas Mohsenpour and Cristel Pelsser. 2026. The Observer Effect in Internet Measurement: A Case Study of Route Origin Validation. In *The 25th ACM Workshop on Hot Topics in Networks (HotNets '26), November 16–17, 2026, Salt Lake City, UT, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3847290.3859951



This work is licensed under a Creative Commons Attribution 4.0 International License.
HotNets '26, Salt Lake City, UT, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-3029-0/2026/11
https://doi.org/10.1145/3847290.3859951

1 Introduction

Internet measurement is central to understanding the operational state of the network [11]. Researchers, operators, regulators, and policymakers rely on measurement studies to evaluate properties ranging from performance and reliability [17, 30] to routing security [16, 20, 26, 27, 38] and censorship [4, 24, 34, 41].

These studies depend on long-lived observatories and shared infrastructures offering continuous visibility into network behavior. The role of these observatories now extends well beyond improving our understanding of the network. Their outputs increasingly feed into frameworks and decisions with real-world consequences, including security assessments and compliance programs such as MANRS [14], operational datasets such as CAIDA AS Rank [8], and public reporting platforms such as IHR [22] and OONI [15]. As a result, measurement outcomes increasingly affect how networks are assessed and perceived by the broader community.

This evolution echoes phenomena long recognized in economics and the social sciences. Goodhart's Law states that "when a measure becomes a target, it ceases to be a good measure" [18], and Campbell's Law observes that indicators used for decision-making become susceptible to distortion [10]. As the outputs of Internet observatories gain operational and economic significance, their measurements may be subject to similar dynamics, even in the absence of any deliberate intent.

Internet measurement studies implicitly rely on the assumption that networks behave consistently regardless of who is observing them. Measurements are generally interpreted as reflecting the actual policies deployed by the network under observation, rather than policies selectively exposed to the measurement process itself.

However, the very characteristics that make Internet observatories attractive to researchers can also make them distinguishable to the networks they measure. Public documentation, the repeated use of the same infrastructure for longitudinal analysis and reporting, and predictable experimental methodologies may allow operators to recognize measurement traffic and treat it differently. Research ethics further reinforce this exposure, as norms of transparency and reproducibility encourage researchers to publicly document their infrastructure and methods.

This possibility raises a fundamental question for Internet measurement: are we measuring the Internet itself, or the Internet's response to being measured? Answering this question is essential to preserving the validity of Internet measurement. Addressing it requires rethinking how measurements are designed, striking a balance between making measurements less predictable and preserving the transparency and reproducibility on which the community depends.

Our contributions are threefold: (1) we introduce the notion of *measurement-aware networks*, which recognize measurement traffic and treat it differently; (2) we propose a paired-target methodology, based on public and shadow measurement targets, for detecting measurement-aware behavior; and (3) using ROV as a case study, we identify networks exhibiting behavior consistent with measurement awareness and discuss the implications for the design and interpretation of Internet observatories.

Beyond ROV, we identify a set of observatories that rely on similarly recognizable infrastructure and warrant investigation for measurement awareness in future work.

2 Measurement-Aware Networks

Internet measurements implicitly assume that networks behave the same way whether or not they can recognize that they are being measured. We define a *measurement-aware network* as a network that violates this assumption: rather than applying its policies uniformly, it recognizes measurement activity and treats it differently.

In this paper, we use ROV as a case study, in which measurements test the reachability of targets within RPKI-invalid prefixes (Section 3). In this setting, a measurement-aware network treats the RPKI-invalid measurement targets it recognizes differently from functionally equivalent RPKI-invalid targets that are not publicly associated with any measurement (Section 2.3).

Figure 1 illustrates measurement awareness in the context of ROV. Observatories publish their measurement targets and methodologies. A measurement-aware AS (MA-AS) may monitor these documents, identify the public targets, and apply routing or filtering policies to these targets only, leaving traffic to other RPKI-invalid prefixes unaffected. The published scores then reflect the filtering applied to the known targets rather than the network's general treatment of RPKI-invalid routes.

2.1 Working Assumptions

To study measurement-aware behavior, we make the following assumptions about the network. These assumptions are deliberately weak: they require neither insider knowledge of the observatory nor specialized capabilities.

Knowledge. The network has access only to publicly available or observable information, such as documented measurement infrastructure and published methodologies. No insider knowledge of the observatory is required.

Capabilities. The network can distinguish recognizable measurement targets from functionally equivalent targets that are not associated with any measurement activity, and apply different routing or filtering policies to each class. As illustrated in Figure 1, such differentiation can be implemented using standard routing and filtering mechanisms already supported by commodity routers.

Motivation. We make no assumption about why a network differentiates. It may do so unintentionally, or to opt out of measurements altogether, to avoid being tested by a specific observatory, to limit the operational impact of probing, or to obtain more favorable measurement outcomes. Only the last would constitute *measurement gaming*, and distinguishing it from the other cases requires knowledge of intent that our measurements do not provide. Our analysis therefore focuses on the observable difference in behavior, not on its underlying intent.

2.2 Cheap Identification

Measurement awareness requires that measured networks can recognize measurement activity. Such recognition becomes inexpensive when observatories expose information about their infrastructure, which is increasingly common.

Examples include publicly documented probe infrastructures such as RIPE Atlas and CAIDA Ark [7, 36], observation vantage points such as BGP collectors [37, 43], and experimental platforms such as the PEERING testbed [39]. In these systems, resources are intentionally disclosed to support reproducibility and independent validation. The same disclosure, however, allows measured networks to identify measurement activity at little cost, which in turn enables them to adapt their behavior to it.

Prior work shows that measurement systems can be evaded or misled. Researchers have demonstrated attacks that hide malicious routing activity from known BGP collectors [5, 6], manipulate the business relationships inferred by AS Rank to influence how networks are perceived [44], and forge delay measurements to mislead Internet geolocation systems [1]. Conversely, operators can deploy mechanisms that deliberately obfuscate or deceive traceroute-based topology inference, for example to protect their networks [29, 42].

These examples show that reacting to known measurements is technically feasible, for both adversarial and defensive purposes. While they do not show that network operators have business incentives to game observatories, they

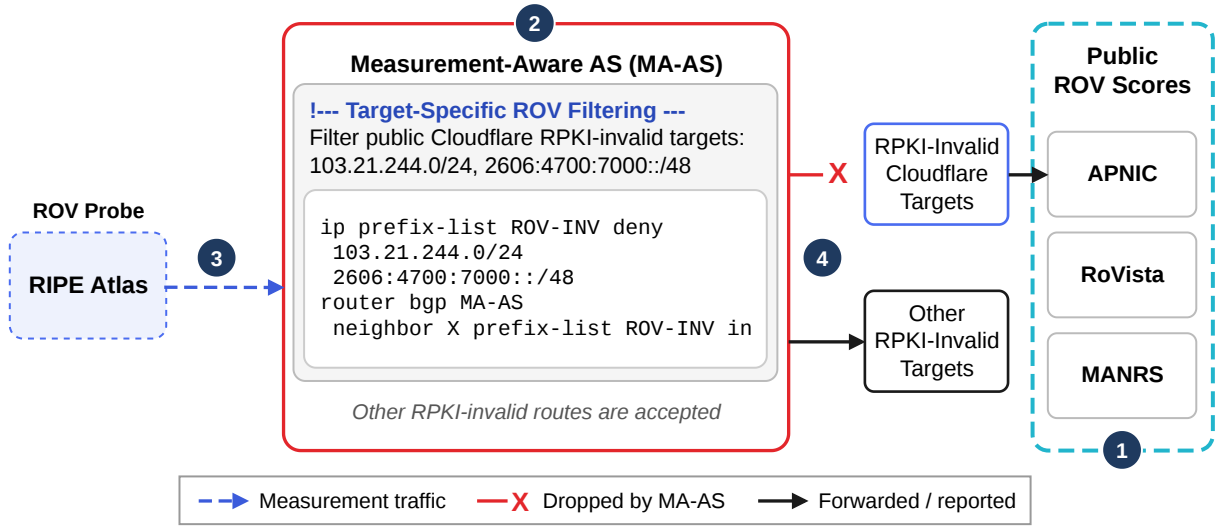


Figure 1: Conceptual overview of a measurement-aware AS in the ROV case study. (1) Observatories use public, long-lived test targets and publish per-network ROV scores. (2) The measurement-aware AS recognizes these public targets and filters them using target-specific rules. (3) Measurement platforms probe the public targets. (4) The measurement-aware AS drops traffic to the public targets but forwards traffic to other RPKI-invalid targets, so the published scores reflect the behavior the measurement-aware AS exposes to observatories rather than its actual ROV deployment status.

do show that recognizable measurement infrastructure lowers the cost of identifying measurements, a prerequisite for reacting to them, whatever the motivation.

2.3 Detection Methodology

Detecting measurement-aware behavior is inherently challenging: researchers lack observations external to the measurements themselves, and therefore cannot directly determine whether a network behaved differently in response to the measurement process.

To address this challenge, we propose a paired-target methodology based on two functionally equivalent measurement targets. The first is a *public target*, which is associated with known measurement infrastructure and is therefore potentially recognizable to the measured network. The second is a *shadow target*, which provides an equivalent measurement function but is not publicly associated with any measurement campaign.

Measurements are performed simultaneously toward both targets under identical experimental conditions. If a network applies its policy uniformly, it should behave equivalently toward the public and shadow targets. Systematic discrepancies between the two observations may therefore indicate that the network adapts its behavior upon recognizing the measurement infrastructure, rather than applying a consistent operational policy.

In the remainder of this paper, we instantiate this methodology using ROV, where publicly known and undisclosed RPKI-invalid prefixes serve as the public and shadow targets, respectively.

3 ROV as a Case Study

ROV [31] uses Route Origin Authorizations (ROAs) [40] published through the Resource Public Key Infrastructure (RPKI) [25] to validate the origin Autonomous System (AS) of Border Gateway Protocol (BGP) announcements. Routes are classified as *Valid*, *Invalid*, or *Unknown*, and operators may reject or de-preference routes classified as invalid.

Several measurement platforms infer ROV deployment by comparing the reachability of valid and invalid prefixes from distributed vantage points. These approaches typically rely on a small number of publicly known invalid prefixes and long-lived measurement infrastructure [3, 13, 32].

For instance, Figure 2 shows the evolution of the publicly documented RPKI-invalid prefixes used by RoVista. Although the platform relies on RPKI-invalid prefixes observed in the wild rather than on dedicated test prefixes, its set of active invalid targets has progressively narrowed: since June 2026, only two prefixes remain active, one of which is the publicly known Cloudflare prefix originally introduced by *Is BGP Safe Yet?* APNIC's ROV measurements rely on this same Cloudflare prefix [21, 26], so all three observatories effectively converge on the same publicly recognizable target.

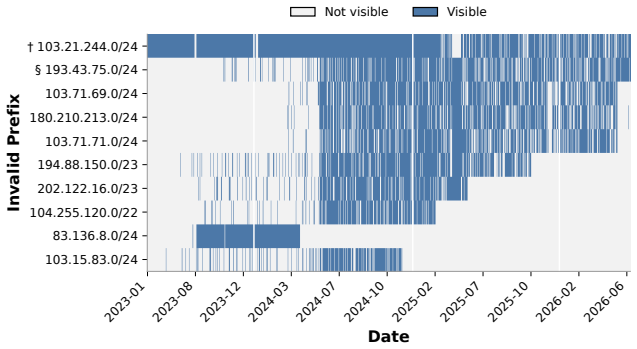


Figure 2: RPKI-invalid test prefixes used by RoVista between January 2023 and July 2026. The number of active invalid targets has steadily decreased, with only two remaining since June 2026, making the measurement targets increasingly predictable. See Table 1 for symbol definitions.

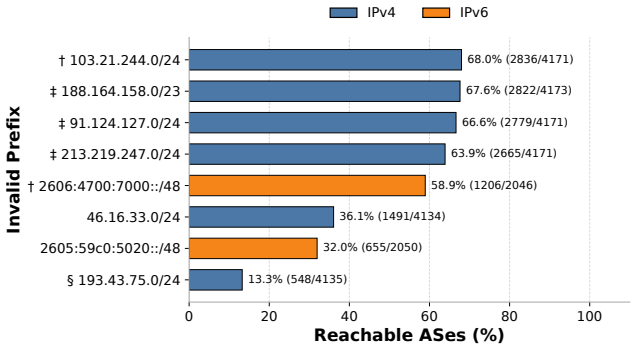


Figure 3: Traceroute reachability of candidate RPKI-invalid prefixes. Values in parentheses give reachable / total source ASes, where the total counts every source AS that was able to launch at least one traceroute toward the prefix, whether or not it reaches the prefix. Totals therefore differ across prefixes, and are lower for IPv6, which fewer probes support. See Table 1 for symbol definitions.

ROV therefore provides a natural case study for measurement awareness. A network that filters only the publicly known invalid prefixes while accepting other invalid prefixes would appear as an ROV adopter, leading existing observatories to overestimate ROV deployment.

Candidate invalid prefixes also vary substantially in reachability. As shown in Figure 3, the fraction of ASes whose traceroutes reach a given candidate prefix ranges from 13.3% to 68%. This variability highlights an additional challenge: target selection itself can substantially influence ROV measurements, which motivates the use of carefully matched public and shadow targets with comparable reachability.

3.1 Selecting Public and Shadow Targets

Applying the paired-target methodology requires public and shadow targets that exhibit comparable reachability and BGP propagation characteristics. We use the publicly known Cloudflare-originated (AS13335) RPKI-invalid prefixes from *Is BGP Safe Yet?*, later adopted by APNIC, RoVista, and the MANRS Observatory metric M9 [28], as public targets.

To identify candidate shadow targets, we collect RPKI-invalid prefixes observed across the BGP data aggregated by bgproutes.io [2], which combines feeds from RIPE RIS, Route Views, PCH, and CGTF collectors together with its own BGP peering sessions. We rank them according to a *Prefix Reachability Score* (PRS):

$$PRS = \text{mean}(V, G, C, P) \tag{1}$$

where V denotes prefix visibility measured as the number of observing vantage points, G the geographic diversity of observing vantage points, C the diversity of observing ASes, and P the diversity of observed AS paths to the prefix.

Table 1: RPKI-invalid prefixes considered in this study, grouped by Prefix Reachability Score (PRS), a composite metric combining the visibility, geographic diversity, observing AS diversity, and AS-path diversity of each prefix. † denotes *Is BGP Safe Yet?* public targets, ‡ BYOIP shadow targets, and § RoVista low-reachability target.

Prefix	Origin AS	PRS (%)
High-reachability targets		
188.164.158.0/23‡	209242	67.58
103.21.244.0/24†	13335	66.52
213.219.247.0/24‡	209242	63.46
91.124.127.0/24‡	209242	62.72
2606:4700:7000::/48†	13335	56.47
Medium-reachability targets		
46.16.33.0/24	208890	38.00
2605:59c0:5020::/48	45700	35.62
Low-reachability targets (RoVista)		
193.43.75.0/24§	202125	22.21

Each component is normalized to the range $[0, 1]$ by dividing by the maximum observed value across all candidate prefixes, and PRS is reported as a percentage. Candidate prefixes are grouped into high-, medium-, and low-reachability categories according to their PRS (Table 1).

To construct suitable shadow targets, we consider RPKI-invalid prefixes advertised through Cloudflare's Bring Your Own IP (BYOIP) service (AS209242) [12]. BYOIP allows organizations to retain ownership of their IP prefixes while

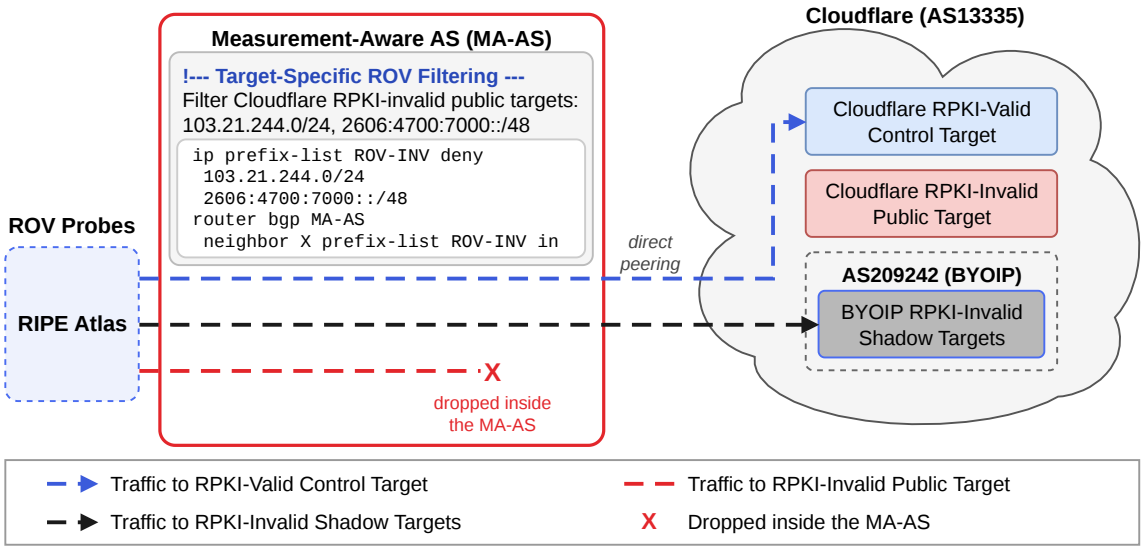


Figure 4: Paired-target methodology for detecting measurement-aware behavior in ROV measurements. To minimize the risk of traffic engineering effects, a RIPE Atlas probe is considered testable only if its traffic toward the RPKI-valid control target and toward all three shadow targets, originated via Cloudflare Bring Your Own IP (AS209242), exits its AS directly to Cloudflare. A probe in a measurement-aware AS cannot reach the public target, yet still reaches both the RPKI-valid control target and the equivalent shadow targets.

Cloudflare advertises them from its network, enabling customers to use Cloudflare's reverse-proxy, CDN, and security services without changing their IP addresses.

In BGP data, Cloudflare is the only neighbor of AS209242, so the shadow targets, like the public targets, reach other networks through Cloudflare. Among the 14 RPKI-invalid prefixes advertised through Cloudflare's BYOIP service, we first exclude those whose reachability could be explained by valid or unknown covering prefixes rather than by the invalid announcement itself. From the remaining six prefixes, we select the highest-reachability prefix from each distinct organization, identified using the Prefix2Org dataset [19], yielding three shadow targets.

Finally, we verify that the Internet Routing Registry (IRR) objects associated with the public targets explicitly authorize Cloudflare to originate them, ruling out IRR-based route filtering as an explanation. We do not perform the same verification for the shadow targets: our methodology only considers cases in which they remain reachable, which makes IRR-based filtering irrelevant to their selection.

3.2 Detecting Measurement-Aware Behavior

We perform traceroute measurements from 12,130 RIPE Atlas probes, distributed across 4,173 ASes, over a period of 14 days. Each probe measures toward the public target, an RPKI-valid Cloudflare control target, and the three shadow targets. We map traceroute interfaces to ASes using Prefix2Org, supplemented with PeeringDB and PCH data [33, 35] to correctly identify Cloudflare's IXP interfaces.

Detection criteria. We classify probes using two conditions:

(C1) Testability. To attribute differences to the measured AS rather than to an intermediate network, we only consider probes whose traffic toward the RPKI-valid control target and the shadow targets exits their AS directly to Cloudflare. Specifically, a probe is *testable* if its traceroutes reach the RPKI-valid control target and all three shadow targets, and each of these four traceroutes leaves the probe's AS directly to Cloudflare. Because our shadow targets are IPv4-only, only IPv4 probes can be testable. In total, 1,208 ASes host at least one testable probe (3,584 probes).

(C2) Filtering of the public target. A testable probe exhibits behavior consistent with measurement awareness if its traceroute toward the public target never leaves its AS, i.e., the public target is filtered inside the MA-AS (Figure 4). The MA-AS learns both the public prefix and the three shadow prefixes from Cloudflare, with the same RPKI status. Uniform ROV deployment within the AS should therefore treat the public and shadow targets alike. We refer to a probe that satisfies both (C1) and (C2) as a *filtering probe*, and call an AS *measurement-aware* if it hosts at least one filtering probe.

IPv6 corroboration. Without IPv6 shadow targets, IPv6 measurements cannot be used to determine whether an AS is measurement-aware, since filtering of the IPv6 public target alone could also reflect ROV deployed only for IPv6.

We therefore use IPv6 only as corroborating evidence: for ASes identified as measurement-aware over IPv4, we check whether the IPv6 public target is also filtered inside the AS, considering only IPv6 probes that reach an RPKI-valid IPv6 target directly via Cloudflare.

Consistent and partial measurement awareness. Because routing policies may be implemented on individual routers rather than uniformly across an AS, different RIPE Atlas probes within the same network may observe different outcomes. We therefore call a measurement-aware AS *consistently* measurement-aware if all of its testable probes are filtering probes, and *partially* measurement-aware if only some of them are, for example, when the filter is deployed on only a subset of the MA-AS's routers.

Results. Among the 1,208 testable ASes, 15 are measurement-aware: 11 consistently and 4 partially. Of the 7 measurement-aware ASes where the IPv6 check applies, 5 also filter the IPv6 public target inside the AS, which suggests a filter on the public targets in both IPv4 and IPv6. To rule out transient routing events and configuration artifacts, we continued monitoring the measurement-aware ASes for two additional weeks. These networks span all five Regional Internet Registry (RIR) regions and have customer cones ranging from 0 to 128 customer ASes, according to CAIDA AS Rank.

Our objective is not to estimate prevalence but to demonstrate existence: even a small number of networks that treat publicly known observatory targets differently from functionally equivalent shadow targets challenges the assumption that observatories provide measurement-independent views of network behavior.

Although measurement-aware networks still reach the shadow targets, they receive high ROV deployment scores from existing observatories, with median scores of 90.3% in APNIC and 100% in RoVista. Measurement-aware networks can thus appear as near-perfect ROV adopters, which exposes a limitation of current Internet-scale ROV measurement methodologies.

Operator feedback. During responsible disclosure, one operator confirmed explicitly filtering the public target used by the *Is BGP Safe Yet?* observatory. The operator explained that they did not wish to participate in ROV measurement campaigns that use Cloudflare-hosted measurement prefixes. This case illustrates the distinction between measurement awareness and gaming: the operator recognized and filtered the public target, but did so to opt out of the measurement rather than to improve their ROV score.

In a separate case, we observed that another network began advertising the *Is BGP Safe Yet?* RPKI-invalid prefix in BGP after responsible disclosure, and continued doing so for the remainder of our observation period. Although we neither attribute this timing to our disclosure nor infer the operator's intent, this case illustrates that publicly known measurement artifacts may become incorporated into operational routing configurations, which in turn affect measurement outcomes.

4 Limits of Attribution

Our detection criteria show that an MA-AS treats the public and shadow targets differently, but not why it does so.

Router-level learning. The main limitation concerns where routes are learned. An AS may peer with Cloudflare on several border routers, and our detection criteria only ensure that all prefixes are learned from the same neighbor, not on the same border router. If ROV is not deployed uniformly, the public RPKI-invalid prefix may be learned on a router that enforces ROV, while all three RPKI-invalid shadow prefixes are learned on routers that do not. Alternatively, measurement-aware filters may not be deployed uniformly across all routers, so that, due to load balancing, some probes observe measurement-aware behavior while others do not. As discussed in Section 3.2, partially measurement-aware ASes are the most susceptible to this explanation.

Selective BGP announcements. Because we study measurement awareness only through the data plane, we cannot verify that Cloudflare announces the public and shadow prefixes identically on every BGP session. However, their similar reachability (Figure 3) and the fact that all three shadow targets are reachable make a systematic difference unlikely.

Other limitations. IP-to-AS mapping errors could affect our results, but only if an intermediate AS remained entirely invisible in the traceroutes, for example by blocking ICMP. In addition, software bugs or misconfigurations in router ROV implementations could cause prefixes with the same RPKI status to be treated differently, and our measurements cannot distinguish such errors from intentional filtering.

5 Discussion

Beyond ROV. Establishing measurement awareness in ROV is demanding, and much of our effort goes into ruling out alternative explanations. The underlying question, however, applies whenever observatories rely on stable, publicly identifiable infrastructure. Below, we outline settings that could be subject to measurement awareness in the same way; whether they actually are remains an open question that we leave for future work.

CAIDA's Spoofer [9] evaluates Source Address Validation (SAV) by sending spoofed packets from Spoofer clients to publicly known CAIDA Ark destinations, and these measurements contribute to the MANRS Observatory metric M5 (Anti-Spoofing) [28]. A network could apply SAV only toward these destinations while continuing to permit spoofed traffic elsewhere, and would thus appear SAV-compliant

without deploying comprehensive filtering. Likewise, censorship observatories such as Censored Planet [41] rely on stable measurement probes; a measurement-aware network could exempt the probes located within its AS from censorship while continuing to censor ordinary users. Similar questions arise for honeypots and network telescopes, which scanners may learn to avoid [23], and more generally for any measurement traffic that may receive preferential treatment.

Measurement awareness may therefore constitute a general source of bias across Internet observatories, rather than a phenomenon unique to ROV.

Implications for observatory design. The design of an Internet observatory plays a critical role in its resilience to measurement awareness. Infrastructures such as RIPE Atlas, which host probes inside individual networks and thereby provide visibility into their internal routing policies, make selective measurement-aware behavior harder to conceal.

More broadly, our results suggest that measurement awareness should become an explicit design consideration for Internet observatories, balancing transparency and reproducibility against resilience to networks that adapt to known measurements.

One promising direction is to make measurements less predictable. Rather than relying on fixed, publicly recognizable targets, observatories could generate measurement artifacts on demand, rotate them over time, or tailor them to the specific networks under evaluation. For example, RScope [27] creates measurement artifacts that are visible only to the networks being evaluated, and LightScope [23] turns production machines into telescopes and honeypots.

However, while network-specific measurement artifacts are feasible in IPv6 thanks to its abundant address space, IPv4 address scarcity makes equivalent methodologies much harder to deploy at scale. This asymmetry may bias future measurement systems toward IPv6 and cause them to overlook behaviors specific to IPv4.

Finally, our disclosure experience suggests that observatories could offer explicit opt-out mechanisms for operators unwilling to participate in measurement campaigns. Such mechanisms would allow operators to opt out openly rather than through filtering, and would improve the interpretation of published results.

Ethical considerations. Our measurements rely exclusively on standard RIPE Atlas traceroutes and publicly observable BGP information. We adhered to RIPE Atlas usage guidelines and rate limits throughout our experiments. We responsibly disclosed our observations to affected operators to validate our interpretation and to seek operational explanations. To support reproducibility, we release our code and artifacts, including the RIPE Atlas measurement IDs and analysis scripts, at https://github.com/AbbasMP17/measurement-awareness-rov.

6 Conclusion

Internet observatories implicitly assume that measured networks behave the same way whether or not they can recognize that they are being measured. We challenged this assumption by introducing the notion of *measurement-aware networks* and a paired-target methodology for detecting them. Using ROV as a case study, we identified 15 networks that filter publicly known measurement targets while still reaching functionally equivalent shadow targets, yet receive near-perfect ROV scores from existing observatories. Publicly identifiable measurement infrastructure can thus lead observatories to capture the behavior that networks expose rather than the property they intend to measure. We hope this work encourages the community to treat measurement awareness as an explicit consideration in observatory design, and to investigate whether other observatories face the same effect.

7 Acknowledgments

We thank our shepherd, Loqman Salamatian, whose feedback greatly improved the framing of this paper, the anonymous reviewers for their valuable comments, and the network operators who responded to our disclosure. This research was partially supported by the Association française pour le nommage Internet en Coopération (Afnic), and by the European Union and Wallonia through the "Wallonia 2021–2027" program.

References

[1] AbdelRahman Abdou, Ashraf Matrawy, and Paul C. van Oorschot. 2017. Accurate Manipulation of Delay-based Internet Geolocation. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security* (Abu Dhabi, United Arab Emirates) *(ASIA CCS '17)*. Association for Computing Machinery, New York, NY, USA, 887–898. doi:10.1145/3052973.3052993

[2] Thomas Alfroy, Thomas Holterbach, Thomas Krenc, K. C. Claffy, and Cristel Pelsser. 2024. The Next Generation of BGP Data Collection Platforms. In *Proceedings of the ACM SIGCOMM 2024 Conference* (Sydney, NSW, Australia) *(ACM SIGCOMM '24)*. Association for Computing Machinery, New York, NY, USA, 794–812. doi:10.1145/3651890.3672251

[3] APNIC Labs. 2026. *APNIC Labs RPKI Statistics*. https://stats.labs.apnic.net/rpki

[4] Abhishek Bhaskar and Paul Pearce. 2024. Understanding Routing-Induced Censorship Changes Globally. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security* (Salt Lake City, UT, USA) *(CCS '24)*. Association for Computing Machinery, New York, NY, USA, 437–451. doi:10.1145/3658644.3670336

[5] Henry Birge-Lee, Maria Apostolaki, and Jennifer Rexford. 2025. Global BGP Attacks that Evade Route Monitoring. In *Passive and Active Measurement: 26th International Conference, PAM 2025, Virtual Event, March 10–12, 2025, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 335–357. doi:10.1007/978-3-031-85960-1_14

[6] Henry Birge-Lee, Liang Wang, Jennifer Rexford, and Prateek Mittal. 2019. SICO: Surgical Interception Attacks by Manipulating BGP Communities. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom)

(CCS '19). Association for Computing Machinery, New York, NY, USA, 431–448. doi:10.1145/3319535.3363197

[7] CAIDA. 2026. Ark Monitor Deployment. [Online]. Available: https://www.caida.org/projects/ark/locations/. Accessed: 2026-06-26.

[8] CAIDA. 2026. AS Rank. https://asrank.caida.org/. Accessed: 2026-06-29.

[9] CAIDA. 2026. The Spoofer Project. Accessed: 2026-06-26. https://www.caida.org/projects/spoofer/

[10] Donald T. Campbell. 1979. Assessing the Impact of Planned Social Change. *Evaluation and Program Planning* 2, 1 (1979), 67–90.

[11] David D. Clark, Matthew Luckie, and K. C. Claffy. 2025. The Current State of the Art in Network Measurement. In *Proceedings of the 53rd Research Conference on Communication, Information and Internet Policy (TPRC)*. doi:10.2139/ssrn.5375475

[12] Cloudflare. 2026. Bring Your Own IPs (BYOIP). https://developers.cloudflare.com/byoip/. Cloudflare Developer Documentation. Accessed: 2026-07-10.

[13] Cloudflare. 2026. *Is BGP Safe Yet?* https://isbgpsafeyet.com/

[14] Ben Du, Cecilia Testart, Romain Fontugne, Gautam Akiwate, Alex C. Snoeren, and kc claffy. 2022. Mind your MANRS: measuring the MANRS ecosystem. In *Proceedings of the 22nd ACM Internet Measurement Conference* (Nice, France) *(IMC '22)*. Association for Computing Machinery, New York, NY, USA, 716–729. doi:10.1145/3517745.3561419

[15] Arturo Filastò and Jacob Appelbaum. 2012. OONI: Open Observatory of Network Interference. In *2nd USENIX Workshop on Free and Open Communications on the Internet (FOCI 12)*. USENIX Association, Bellevue, WA. https://www.usenix.org/conference/foci12/workshop-program/presentation/filast%C3%B2

[16] Justin Furuness, Cameron Morris, Reynaldo Morillo, Arvind Kasiliya, Bing Wang, and Amir Herzberg. 2025. Securing BGP ASAP: ASPA and other Post-ROV Defenses. In *Network and Distributed System Security (NDSS) Symposium*. Internet Society, San Diego, CA, USA, 1–14. doi:10.14722/ndss.2025.240675

[17] Phillipa Gill, Christophe Diot, Lai Yi Ohlsen, Matt Mathis, and Stephen Soltesz. 2022. M-Lab: user initiated internet data for the research community. *SIGCOMM Comput. Commun. Rev.* 52, 1 (March 2022), 34–37. doi:10.1145/3523230.3523236

[18] Charles A. E. Goodhart. 1975. Problems of Monetary Management: The UK Experience. *Papers in Monetary Economics* 1 (1975), 91–121.

[19] Deepak Gouda, Alberto Dainotti, and Cecilia Testart. 2025. Prefix2Org: Mapping BGP Prefixes to Organizations. In *Proceedings of the 2025 ACM Internet Measurement Conference* (USA) *(IMC '25)*. Association for Computing Machinery, New York, NY, USA, 397–414. doi:10.1145/3730567.3764485

[20] Tomas Hlavacek, Haya Shulman, Niklas Vogel, and Michael Waidner. 2023. Keep Your Friends Close, but Your Routeservers Closer: Insights into RPKI Validation in the Internet. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 4841–4858. https://www.usenix.org/conference/usenixsecurity23/presentation/hlavacek

[21] Geoff Huston. 2023. *How we measure RPKI ROA signing and Route Origin Validation.* https://blog.apnic.net/2023/11/09/how-we-measure-rpki-roa-signing-and-route-origination-validation/

[22] IIJ Research Laboratory. 2020. Internet Health Report. https://ihr.iijlab.net.

[23] Eric Kapitanski and Jelena Mirkovic. 2026. LightScope: Turning Production Machines into Telescopes and Honeypots. In *Proceedings of the 2026 ACM Internet Measurement Conference* (Karlsruhe, Germany) *(IMC '26)*. Association for Computing Machinery, New York, NY, USA, 228–245. doi:10.1145/3777912.3839790

[24] Hieu Le, Armin Huremagic, Kevin Wang, Roya Ensafi, and Ram Sundara Raman. 2026. CenRL: A Framework for Performing Intelligent Censorship Measurements . In *2026 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 3816–3834. doi:10.1109/SP63933.2026.00098

[25] Matt Lepinski and Stephen Kent. 2012. An Infrastructure to Support Secure Internet Routing. RFC 6480. doi:10.17487/RFC6480

[26] Weitong Li, Zhexiao Lin, Md. Ishtiaq Ashiq, Emile Aben, Romain Fontugne, Amreesh Phokeer, and Taejoong Chung. 2023. RoVista: Measuring and Analyzing the Route Origin Validation (ROV) in RPKI. In *Proceedings of the 2023 ACM on Internet Measurement Conference* (Montreal QC, Canada) *(IMC '23)*. Association for Computing Machinery, New York, NY, USA, 73–88. doi:10.1145/3618257.3624806

[27] Weitong Li, Yongzhe Xu, Mingwei Zhang, Vasileios Giotsas, and Taejoong Chung. 2026. RScope: Unveiling Global ROV Deployments and Dependencies in the Post-ROV Era. In *Proceedings of the 2026 ACM Internet Measurement Conference* (Karlsruhe, Germany) *(IMC '26)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3777912.3809144

[28] MANRS. 2026. MANRS Observatory Measurement Framework. https://manrs.org/manrs-observatory/measurement-framework/. Accessed: 2026-06-26.

[29] Roland Meier, Petar Tsankov, Vincent Lenders, Laurent Vanbever, and Martin Vechev. 2018. NetHide: Secure and Practical Network Topology Obfuscation. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, Baltimore, MD, 693–709. https://www.usenix.org/conference/usenixsecurity18/presentation/meier

[30] Oliver Michel, Satadal Sengupta, Hyojoon Kim, Ravi Netravali, and Jennifer Rexford. 2022. Enabling passive measurement of zoom performance in production networks. In *Proceedings of the 22nd ACM Internet Measurement Conference* (Nice, France) *(IMC '22)*. Association for Computing Machinery, New York, NY, USA, 244–260. doi:10.1145/3517745.3561414

[31] Prodosh Mohapatra, John Scudder, David Ward, Randy Bush, and Rob Austein. 2013. BGP Prefix Origin Validation. RFC 6811. doi:10.17487/RFC6811

[32] NetSecure Lab. [n. d.]. RoVista: Route Origin Validation Inference and Status Tracking. https://rovista.netsecurelab.org/. Accessed: 2026-07-04.

[33] Packet Clearing House. 2026. Internet Exchange Point Datasets. https://www.pch.net/ixp/data. Accessed: 2026-06-26.

[34] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. 2017. Global measurement of DNS manipulation. In *Proceedings of the 26th USENIX Conference on Security Symposium* (Vancouver, BC, Canada) *(SEC'17)*. USENIX Association, USA, 307–323.

[35] PeeringDB. 2026. PeeringDB: The Interconnection Database. https://www.peeringdb.com/. Accessed: 2026-06-26.

[36] RIPE NCC. 2026. RIPE Atlas. [Online]. Available: https://atlas.ripe.net/. Accessed: 2026-06-26.

[37] RIPE NCC. 2026. Routing Information Service (RIS). [Online]. Available: https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/. Accessed: 2026-06-26.

[38] Nils Rodday, Ítalo S. Cunha, Randy Bush, Ethan Katz-Bassett, Gabi Dreo Rodosek, Thomas C. Schmidt, and Matthias Wählisch. 2021. Revisiting RPKI Route Origin Validation on the Data Plane. In *5th Network Traffic Measurement and Analysis Conference (TMA 2021)*. IFIP.

[39] Brandon Schlinker, Kyriakos Zarifis, Italo Cunha, Nick Feamster, and Ethan Katz-Bassett. 2014. PEERING: An AS for Us. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks* (Los Angeles, CA, USA) *(HotNets-XIII)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/2670518.2673887

[40] Job Snijders, Ben Maddison, Matt Lepinski, Derrick Kong, and Stephen Kent. 2024. A Profile for Route Origin Authorizations (ROAs). RFC

9582. doi:10.17487/RFC9582

[41] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. 2020. Censored Planet: An Internet-wide, Longitudinal Censorship Observatory. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security* (Virtual Event, USA) *(CCS '20)*. Association for Computing Machinery, New York, NY, USA, 49–66. doi:10.1145/3372297.3417883

[42] Samuel T. Trassare, Robert Beverly, and David Alderson. 2013. A Technique for Network Topology Deception. In *MILCOM 2013 - 2013 IEEE Military Communications Conference*. 1795–1800. doi:10.1109/MILCOM.2013.303

[43] University of Oregon Route Views Project. 2026. Route Views Project. [Online]. Available: https://www.routeviews.org/. Accessed: 2026-06-26.

[44] Yi Xu, Yihao Chen, Ke Xu, Qi Li, and Jianping Wu. 2026. ASRogue: Manipulating ASRank-Inferred AS Relationships. In *35th USENIX Security Symposium (USENIX Security 26)*. USENIX Association, Baltimore, MD, 6167–6186. https://www.usenix.org/conference/usenixsecurity26/presentation/xu-yi