

Demystifying RPKI-Invalid Prefixes: Hidden Causes and Security Risks

Weitong Li[†], Tao Wan[§], Taejoong Chung[†]

[†]Virginia Tech, [§]CableLabs

Abstract—The Resource Public Key Infrastructure (RPKI) enhances Internet routing security by utilizing Route Origin Authorization (ROA) objects to link IP prefixes with their rightful origin ASNs. Despite the rapid deployment of RPKI—over 51.3% of Internet routes now covered by ROAs, there are still 6,802 RPKI-invalid prefixes as of today. This work provides the first comprehensive study to understand and classify the hidden causes of RPKI-invalid prefixes, revealing that ROA misconfigurations often occur during IP leasing and IP transit services. We identify scenarios explaining these misconfigurations and attribute 96.9% of the RPKI-invalid prefixes to such misconfigurations.

We further show their cascading impacts on the data-plane, noting that while most prefixes exhibit negligible effects, 3.1% result in full connectivity loss and 7.1% degrade routing by adding latency and extra hop counts—and, in some cases, also bypassing intended security mechanisms; additionally, we find that such misconfigurations have been triggering false alarms in hijack detection systems. To validate our findings, we build a ground-truth dataset of 294 misconfigured prefixes through direct engagement with 174 network operators. We also interviewed 16 large ISPs and major leasing brokers about their ROA management practices, and we propose suggestions to avert ROA misconfigurations.

Taken together, this study not only fills gaps left by previous research but also offers actionable recommendations to network operators for improving ROA management and minimizing the occurrence of RPKI-invalid announcements.

I. INTRODUCTION

The Border Gateway Protocol (BGP) underpins the Internet’s global routing. However, its lack of intrinsic security features renders it susceptible to a host of exploits, including prefix hijacking, route leaks, and IP address spoofing. To bolster BGP’s security, the Resource Public Key Infrastructure (RPKI) [35] was introduced, creating a cryptographic binding between IP address prefixes and their legitimate originating Autonomous Systems (ASes).

At the core of RPKI lies the Route Origin Authorization (ROA), which explicitly designates which AS is permitted to originate traffic for each IP prefix. By referencing ROAs in the RPKI repository, operators can identify and filter illegitimate route announcements using Route Origin Validation (ROV). Recent data shows that more than 50% of the global IPv4

address space is already validated via ROAs [31], [52], and nearly all top-tier ASes have embraced ROV to discard invalid announcements [37].

Despite this progress, a paradox persists: *over 6,000 RPKI-invalid prefixes propagate daily, with no meaningful decline in recent years*. Prior work attributes many cases to easily detectable misconfigurations, such as customer-provider mismatches or overly broad MaxLength values [9], [26]. These works, however, leave two open questions: *Why do thousands of invalid routes persist, and why do existing heuristics miss them?*

The answer lies in hidden misconfigurations—complex, non-obvious errors entangled with modern network practices like IP leasing or opaque transit services (e.g., BGP tunneling). Unlike simple mistakes, these cases lack visible AS-level relationships in routing data, making them indistinguishable from hijacks to automated tools. For example, a leased prefix announced without ROA updates may appear illegitimate to ROV filters, even if the lessee rightfully operates the IP space.

The consequences of these hidden misconfigurations ripple across the Internet’s security fabric; ROV-enabled networks can inadvertently block legitimate prefixes, triggering outages that cripple time-critical services. Even when connectivity persists, rerouted traffic often traverses suboptimal paths, degrading performance and complicating attack detection. Traffic rerouted through unsecured paths bypasses critical safeguards like DDoS scrubbing services, exposing networks to volumetric attacks or man-in-the-middle surveillance.

This erosion of trust extends to security tools themselves. RPKI’s cryptographic guarantees now underpin critical applications like hijack detection systems, but these tools struggle to distinguish hidden misconfigurations from malicious hijacks. By relying on heuristics to filter “benign” errors—such as same-organization checks—they leave operators inundated with false alarms, masking genuine threats and eroding confidence in automated defenses.

In this paper, we uncover hidden misconfigurations rooted in modern network practices like IP leasing and BGP tunneling, quantify their security risks, and analyze their impact on hijack detection systems. Specifically:

- **Unmasking Hidden Misconfigurations:** We systematically classify 96.9% of RPKI-invalid prefixes—a significant leap over prior work [9], which explained 79% of cases. By dissecting IP leasing dynamics and multi-AS transit arrangements, we expose root causes previously dismissed as “unknown”; notably, 35.5% of those formerly unexplained

RPKI-invalid prefixes arise from leasing scenarios in which the lessor neglects to update the ROA after leasing prefixes to customers. Another 30.1% involve opaque transit setups, where providers announce customer prefixes without appearing in AS paths, rendering traditional detection heuristics ineffective.

- **Quantifying Cascading Impacts:** Our measurements show that hidden misconfigurations reroute traffic onto unsecured paths—18.5% of RPKI-invalid routes diverge because ROV-enabled networks drop them, and 39.2% of the affected prefixes suffer performance degradation, with latency spikes exceeding 100 ms. Worse, misconfigurations involving third-party brokers take 2–5× longer to resolve than internal errors, amplifying exposure to attacks. These misconfigurations also undermine the reliability of security tools: 73.4% of alerts in hijack detection systems like Cloudflare Radar and GRIP falsely flag benign routes as malicious, desensitizing operators to genuine threats; to validate these findings, we contacted 174 network operators whose 349 prefixes were erroneously flagged as hijacked—*all confirmed the alerts stemmed from hidden misconfigurations, not malicious activity.*
- **Operational Insights from ISPs and Leasing Brokers:** Surveys of 8 large network operators and 8 major leasing brokers uncover systemic gaps that perpetuate errors. Organizational silos exacerbate these issues by delaying ROA updates: 3 out of 8 ISPs reported BGP teams propagating routes before IP-management teams updated ROAs, while 5 out of 8 brokers lacked automated ROA monitoring, relying on ad-hoc email coordination.

By bridging the gap between RPKI’s theoretical potential and real-world management challenges, this work equips operators with actionable strategies to close critical ROA management loopholes and reduce false alarms in security systems. Our findings reveal that RPKI’s success depends not merely on adoption rates but on disciplined, end-to-end stewardship of its trust infrastructure—a foundational requirement for safeguarding the Internet’s routing backbone. To foster reproducibility and further research into improving the RPKI ecosystem, we make all code, measurement tools, and datasets at

<https://roa-misconfig.netsecurelab.org>

II. BACKGROUND AND RELATED WORK

A. BGP and RPKI

BGP is the de facto inter-domain routing protocol that connects different Autonomous Systems on the Internet. BGP speakers announce paths towards the origin of IP prefixes through a series of ASes. For example, an AS may receive a BGP route indicating that AS40220 originates the prefix 45.3.0.0/16:

IP Prefix: 45.3.0.0/16
AS_PATH: AS3356 AS174 AS40220

A BGP router builds its routing table based on the BGP messages it receives and applies a route selection process to choose the optimal path for forwarding traffic. Since BGP lacks built-in security mechanisms, an attacker can announce an IP

prefix they are not authorized to announce to intercept traffic, which is known as *prefix hijacking* attacks.

To enhance the security of BGP, the RPKI was introduced. RPKI is a public key infrastructure framework that allows IP address space holders to publish digitally signed certificates binding their IP address space to their ASNs. RPKI secures BGP through two major steps: (1) resource owners register RPKI objects called ROAs, and (2) network operators validate BGP announcements against ROAs to filter out BGP announcements with RPKI-invalid IP prefixes, a process known as ROV.

Registering ROAs: A ROA is a cryptographic document that specifies which AS is authorized to announce a particular IP prefix. To deploy a ROA, network resource owners first create a CA certificate, which binds a set of Internet Number Resources (INRs)—such as ASNs or IP prefixes—to a public key. They then create a ROA that authorizes an AS to announce IP prefixes, which is signed by an End Entity (EE) certificate derived from the CA certificate. These objects must be published in public RPKI repositories operated by either the five Regional Internet Registries (RIRs), which serve as the trust anchors for RPKI, or by National Internet Registries (NIRs) or Local Internet Registries (LIRs) delegated by the RIRs.

Validating BGP Announcements with RPKI: RPKI cannot prevent BGP hijacks unless network operators validate BGP announcements against ROAs. Network operators use RPKI validator software, known as Relying Party software, to fetch and validate ROAs from RPKI repositories. This process produces a list of validated tuples (ASN, ROA prefix, prefix length), called Validated ROA Payloads (VRPs). The set of VRPs is provided to the AS’s routers using the RPKI-to-Router Protocol [7], enabling them to perform ROV, validate incoming BGP announcements based on the standard [39].

With ROV, a router validates incoming BGP announcements using the set of VRPs. First, it determines whether the IP prefix in the BGP announcement is *covered* by any VRP. If so, it then checks if the BGP announcement exactly *matches* a VRP. A BGP announcement matches a VRP when: (1) VRP IP prefix covers the announced IP prefix, (2) the VRP AS matches the announced AS, and (3) the announced prefix length is no longer than the VRP’s *maxLength*—an optional field in a ROA that sets the most-specific (longest) prefix the AS is allowed to advertise for that base prefix—and is at least as long as the VRP’s own prefix length.

A BGP announcement is considered *valid* if it matches a VRP, *invalid* if the IP prefix in the BGP announcement is covered but no VRP matches the announcement, and *unknown* if it is not covered by any VRP.

In this paper, we define the *ROA origin* as the AS authorized to announce an IP prefix in the ROA, and the *BGP origin* as the AS that actually announces the IP prefix.

B. Misconfigured ROAs and Their Impact

We briefly overview previous work on understanding ROA misconfigurations and their impact.

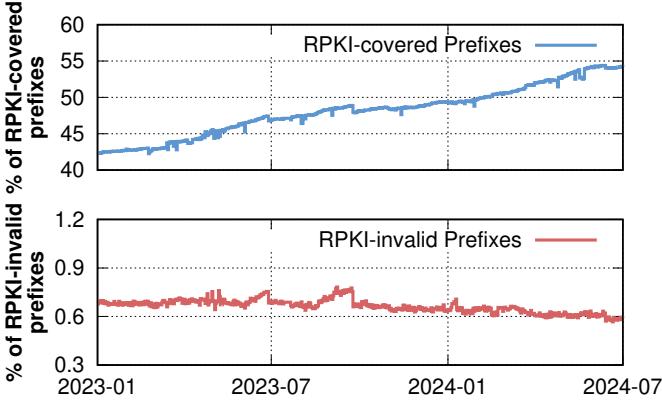


Fig. 1: The percentage of RPKI-covered prefixes and the ratio of RPKI-invalid prefixes to total routes in BGP dumps.

Causes of RPKI-invalid BGP Announcements: Previous studies have identified prevalent ROA misconfigurations among network operators; Chung et al. [9] validated BGP announcements against ROAs and found that nearly 2% of announcements covered by ROAs were invalid, mostly due to misconfigurations. They categorized these into MaxLength errors, same-organization issues, provider-customer relationships, and DDoS protection cases, explaining over 79.0% of RPKI-invalid announcements. Similarly, Xu et al. [59] proposed a case-by-case classification method covering 63.3% of RPKI-invalid routes and Hlavacek et al. [26] used WHOIS and IRR information to infer legitimate misconfigurations when conflicting origins shared records like organization IDs and contact information.

Despite differing methodologies, these studies agree that many RPKI-invalid announcements result from misconfigurations rather than BGP hijacks; however, relying solely on AS relationship datasets and WHOIS/IRR information fails to cover other types of ROA misconfigurations that cannot be explained by BGP data alone (e.g., IP leasing or hidden IP transit), which our work aims to address.

Security risk of RPKI-invalid: When ROA misconfigurations occur, RPKI-invalid prefixes without alternative routes are filtered and become unreachable in ROV networks. This connectivity impact has been discussed in several studies measuring ROV deployment [10], [29], [57], [47], [50]. Recently proposed security extensions of ROV, like ROV++ [40], also aim to limit the propagation of RPKI-invalid prefixes in networks where ROV is only partially deployed. Examining RPKI-invalid prefixes on July 1st, 2024, we find that over 85.6% have alternative routes in current BGP tables. With increasing ROV deployment, this implies that (1) prefixes without alternative routes may experience connectivity problems, and (2) those with alternative routes may face path divergence, leading to suboptimal routing. However, to our knowledge, no prior work has measured the impact of these RPKI-invalid prefixes beyond connectivity issues.

Hijack Detection Using RPKI: RPKI, considered a reliable database for IP resource authority, has applications beyond ROV; it has been used to cross-validate routing databases

like IRR [31] and for source address validation [55]. Another critical application is hijack detection. Recent systems in both industry and academia use ROAs as the authoritative source to detect hijacks. Themis [46] uses RPKI to identify illegitimate BGP announcements. GRIP [1] (based on BGPStream [44]) and Cloudflare Radar [15] also rely on RPKI data for hijack detection.

Crucially, these tools do *not* label every RPKI-invalid route as a hijack, because many such routes stem from benign misconfigurations; instead, they apply heuristics (e.g., same-organization checks or visible provider-customer paths) to suppress obvious false positives. As we show in §VII, hidden misconfigurations that escape these heuristics still generate a significant volume of spurious hijack alerts.

C. RPKI management during IP Leasing and Transit Services

The exponential growth of internet-connected devices and services has intensified the demand for IP address resources and robust network connectivity. Two pivotal mechanisms addressing these needs are IP leasing and IP transit services.

IP Leasing Services: With the exhaustion of IPv4 addresses, organizations often resort to leasing IP addresses as a flexible alternative to purchasing. This approach allows entities to scale their networks without the substantial investment of acquiring IP blocks. Organizations can lease prefixes either *directly* from a resource holder that controls large address blocks such as Cogent [17]—a model known as *direct leasing*, or via a broker such as IPXO [58].

Because a lease is not an official transfer, the ROA stays *under the control of the lessor*. The lessee (or broker) must therefore coordinate with the lessor—typically by email or phone—to add the lessee’s AS to the relevant ROAs. If that coordination fails, prefixes the lessee legitimately originates are still flagged as RPKI-invalid, making them vulnerable to filtering by ROV-enforcing networks.

Previous works have studied the functioning of IP transfer markets [34], [33], methods for identifying leased prefixes using WHOIS logs [20], and the prevalence of malicious activities from leased IP addresses [24], [45], [32]. However, none of these works have focused on the RPKI management during IP leasing processes.

IP Transit Services: In a *conventional* transit relationship, the customer originates its own prefixes and simply hands them to an upstream provider for global reachability. The resulting AS_PATH ends in customer’s ASN, so public AS-relationship datasets (e.g., CAIDA [12]) expose a clear provider-customer link; if the provider fails to keep its ROAs in sync with the customer’s BGP announcements, the mismatch is still easy to detect: the BGP origin (customer) and the ROA origin (customer or provider, depending on who controls the ROA) are both visible in control-plane data [9], [49].

Prior work has shown that this model can suffer from ROA misconfigurations when the transit provider fails to keep its ROAs in sync with its BGP announcements [9], [49].

However, modern transit arrangements can also conceal the true origin of a prefix; a customer may “bring” its own

Previous works	% of Inferred Reported	Misconfigurations Reproduced
Chung et al. [9]	79.0%	81.4%
Xu et al. [59]	63.3%	62.8%
Hlavacek et al. [26]	N/A	59.0%
Combined	-	83.1%
Ours	-	96.9%

TABLE I: Inferred RPKI-Invalid prefixes due to misconfigurations by prior inference methods versus our approach.

address block yet ask the upstream provider to originate that block from the provider’s ASN—typically so the provider can apply DDoS-scrubbing (e.g., Cloudflare’s Magic Transit [42]) or other traffic-shaping policies at the edge. The provider then tunnels the traffic back to the customer, meaning the customer’s ASN never appears in the global BGP table, making the provider’s ASN and the ROA ASN look completely unrelated. If the customer (or the prefix owner) forgets to add the provider’s ASN to the ROA, every such announcement is RPKI-invalid, and the resulting mismatch is indistinguishable from a deliberate hijack—making it much harder for automated tools to diagnose.

III. CURRENT STATUS OF RPKI-INVALID ROUTES

We first examine the current status of RPKI-invalid routes and attempt to reproduce existing works on identifying and classifying misconfigurations in RPKI. To analyze RPKI-invalid routes, we collect the following datasets:

- **BGP Datasets:** We use RouteViews [53] and RIPE RIS [51] BGP table dumps covering the period from January 1st, 2023, to July 1st, 2024. BGP dumps from *all* vantage points of RouteViews and RIPE RIS are collected every four hours during this 18-month measurement period.
- **ROA Datasets:** To identify RPKI-invalid routes, we also collect RPKI ROA objects from all publication points under the five RIRs’ trust anchors. We also deploy relying party software, Routinator [54], to download and validate all ROA certificates on a daily basis from January 1st, 2023, to July 1st, 2024. With the produced VRPs, we perform RPKI validation with the BGP datasets.

Observations: Figure 1 (top) shows the percentage of announced prefixes covered by ROAs. We confirm a rapid growth in coverage, which is encouraging; compared to less than 20% coverage in 2019 [9], we observe an increase from 43.6% to 52.8% during our 18-month measurement period.

However, when we track the number of RPKI-invalid prefixes, we notice that it does not decrease as much as we expected. As Figure 1 (bottom) shows, the number only decreases from 7,989 to 6,802 prefixes, reducing from 0.7% to 0.6% of the total prefixes in the BGP datasets during the 18-month period. This suggests that there must be factors hindering correct ROA setup beyond simple misconfigurations. Combining all snapshots, we observe a total of 42,654 unique RPKI-invalid prefixes during the 18-month period.

A. Reproducing Previous Works

As mentioned in §II-B, previous works [9], [59], [26] studied these RPKI-invalid prefixes and tried to find out the reasons why RPKI-invalid prefixes still exist. It has been shown that most of the RPKI-invalid prefixes are caused by misconfigurations of legitimate origins, rather than malicious hijacking. Table I lists the percentage of RPKI-invalid prefixes each method identified as misconfigurations.

They use the following methods to identify potential legitimate misconfigurations in RPKI-invalid prefixes:

- (a) **MaxLength:** The BGP origin matches the ROA origin AS, but the prefix length is outside the scope specified by the ROA’s `maxLength` field.
- (b) **Same Organization:** The BGP origin and ROA origin are registered under the same organization. [9] uses the `as2org` dataset to identify this kind of misconfiguration, while [26] uses WHOIS and IRR data.
- (c) **Provider-Customer Relationship:** The BGP origin and ROA origin are in a provider-customer relationship, indicating that the issue is caused by misconfigurations during IP transit services.
- (d) **Others:** Additional data sources are also used. For example, [9] uses a list of 36 DDoS providers to identify ROA misconfigurations caused by DDoS protection.

Since these works measured RPKI-invalid prefixes at different times and the datasets they used are not consistent (e.g., partial vantage points in RouteViews, BGP table dumps vs. BGP messages), it is unfair to directly compare the results. Therefore, we reproduce their methodologies and identified potential RPKI-invalid prefixes due to suspected ROA misconfigurations using the BGP datasets we collected, and we are able to attribute 35,445 (83.1%) of them to potential ROA misconfigurations, leaving 16.9% of the RPKI-invalid prefixes as potential unknown or hijacks. This work aims to address this remaining percentage.

IV. HIDDEN TYPES OF ROA MISCONFIGURATION

Our investigation revealed that 42,654 RPKI-invalid prefixes remain present in the global routing table. Using the methodologies from previous works [9], [59], [26], we identified 35,445 (83.1%) of these RPKI-invalid prefixes as misconfigurations, leaving 7,209 (16.9%) of the prefixes classified as *unknown*.

In this section, we examine two modern networking practices that can lead to RPKI-invalid prefixes: transit services and IP-address leasing.

A. Transit Service

In typical IP transit arrangements, a customer announces its own IP prefixes—or, in many cases, obtains sub-prefixes allocated from the provider’s address block—thereby establishing a straightforward provider-customer relationship. In Figure 2(a), for example, AS X announces a /24 derived from AS Y’s larger prefix, so AS Y retains ROA control. If that ROA mistakenly lists AS Y rather than AS X as the origin, the route becomes RPKI-invalid. In most cases, AS relationship datasets

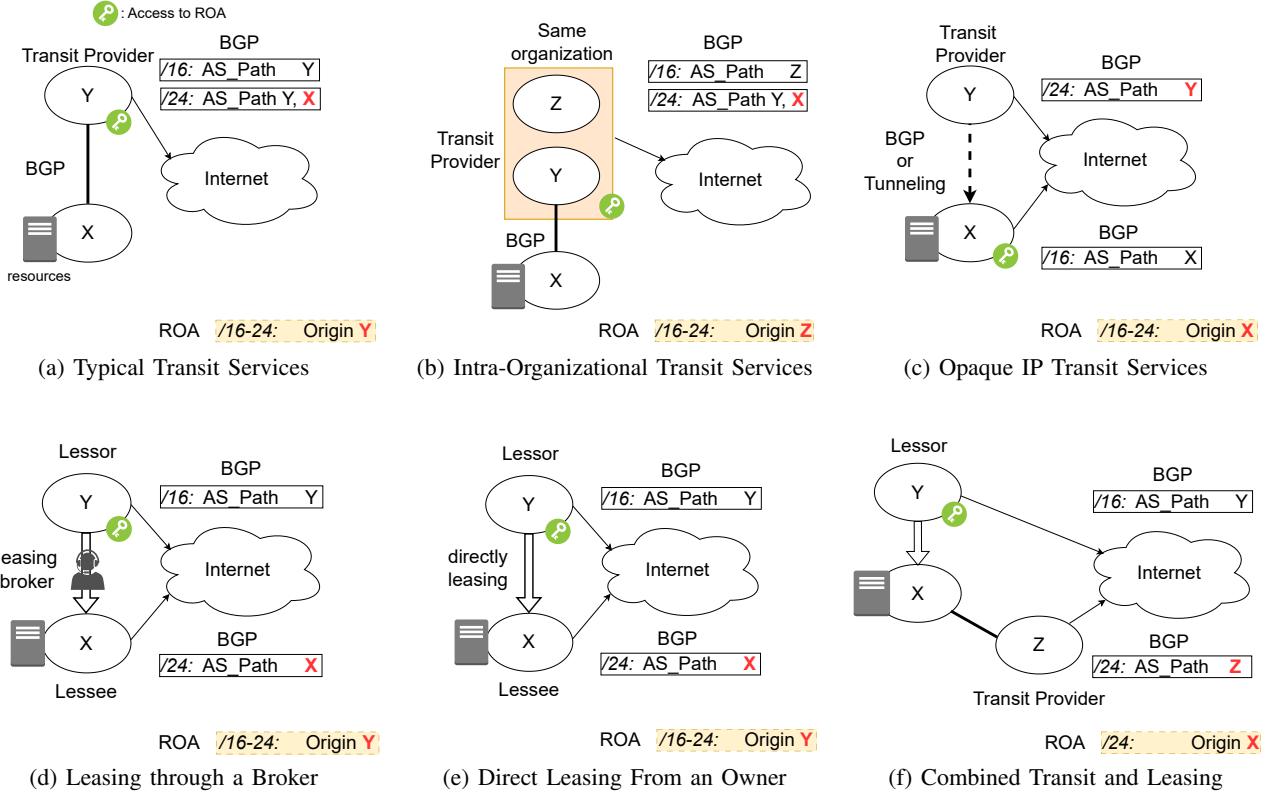


Fig. 2: Misconfigured ROA scenarios due to mismatches between the ROA origin and BGP origin. They can occur when an AS uses a transit provider (a, b, c), leases IP prefixes either through a broker or directly from a lessor (d, e), or combines both leasing and transit services (f). Notably, only scenario (a) can be readily identified through BGP relationships because both the customer and provider appear in the AS_PATH, and prior studies have focused almost exclusively on this case [9], [59], [26]. Please refer Appendix §IV-D for real-world examples of each hidden type of ROA misconfigurations.

detect such misconfigurations by observing the AS_PATH (Y X) and capturing the provider-customer linkage.

Yet, as we discuss below, some subtle or indirect configurations can still evade detection—even though the sub-prefix is technically routed through its provider:

Intra-Organizational Transit Services. IP-transit providers often operate multiple ASNs. As illustrated in Figure 2(b), a provider controlling AS Y and AS Z may allocate a /24 from AS Z to customer AS X while routing upstream through AS Y. If the provider keeps the ROA with AS Z but AS X originates the BGP announcement, no direct link between X and Z appears in BGP data. Detection schemes that rely on inferred provider-customer ties or public AS-relationship datasets therefore miss the misconfiguration.

Opaque IP Transit: A transit provider may announce a customer’s own IP prefixes (i.e., not necessarily allocated from the provider’s block) under the provider’s ASN to apply specific routing policies, such as DDoS mitigation or traffic scrubbing [30], [38]. For example, Chung et al. [9] found that DDoS-scrubbing deployments can create ROA/BGP inconsistencies; however, they did not extend the discussion to the broader class of opaque transit arrangements in which a provider originates a customer’s prefix from its own ASN, nor

did they propose a scalable technique for detecting such cases.

In practice, the provider may employ tunneling protocols (e.g., GRE) to carry the customer’s traffic, as in the case of Cloudflare Magic Transit [42]. We refer to these scenarios as “Opaque IP transit” because the customer’s ASN does not appear in the AS_PATH. In such cases, the customer must configure its ROAs to reflect the provider’s role as the origin ASN, ensuring proper route validation.

For instance, Figure 2(c) illustrates AS X bringing its own /24 prefix to provider AS Y, which then announces the prefix under AS Y’s ASN and tunnels traffic back to AS X. Because this arrangement omits the provider’s ASN from the visible path, AS relationship datasets fail to detect any linkage between AS X and AS Y. Consequently, traditional detection methods—such as those relying on AS_PATH analysis or public relationship datasets like CAIDA’s AS Relationships[12]—often overlook misconfigurations or other anomalies that arise in these opaque transit scenarios.

B. IP Leasing

Since the IP prefixes are leased instead of transferred, the lessor retains ownership of the IP prefixes as well as the

authentication of ROAs.¹

Although brokers may obtain access to ROAs by requiring RIR account access or RPKI CA delegation from the lessor [58], in most cases, the lessor is still responsible for updating the ROAs [48], [14], [28] and brokers need to manually ask the lessor to update the ROAs as lessees request. Unlike prefix sub-allocation during IP transit, the lessor in IP leasing does not need to provide transit services to the lessee; thus, if the lessor has not correctly updated their ROAs to authorize the lessee, BGP announcements made by the lessee could be labeled as RPKI-invalid and could be mistakenly accused of BGP hijacking.

IP Leasing Through Brokers: IP-address brokers do not hold prefixes themselves; they simply serve as intermediaries, matching lessees with prefix owners. As shown in Figure 2(d), the lessor (AS Y), who owns a /16 prefix, can lease a sub-prefix to the lessee AS X through the leasing broker. To facilitate this arrangement, the lessor typically updates their WHOIS and IRR records to reflect the broker's information, allowing the broker to manage the leasing process effectively. After matching the lessee, the lessor is still responsible for updating their ROAs to ensure the lessee's BGP announcements are RPKI-valid.

In principle, a broker could manage a lessor's ROAs in two ways: (1) obtain access to the owner's RIR account, or (2) operate its own RPKI CA and have the lessor delegate ROA control to that CA. In practice this is rare: among 163 registered brokers [20], only one—IPXO—runs its own CA [27]. Consequently, most leased prefixes rely on the lessor to update ROAs, and lapses on the owner's side routinely leave lessee announcements flagged as invalid.

Direct IP Leasing from Owners: An organization can directly lease its prefixes to others without contacting a third-party broker. Figure 2(e) shows the direct leasing scenario where the owner of an IP prefix /16, AS Y, leases a /24 sub-prefix to AS X. AS X will use its own upstream to propagate this /24 to the Internet without using the lessor's networks or ASNs. As with leasing through brokers, usually there is no actual IP ownership transfer during direct leasing. Although the prefix owners can re-allocate the prefix to the lessee in WHOIS records using the Shared WHOIS Project (SWIP [18]), Referral WHOIS (RWhois [19]), or other methods, authority over ROA records will still not be transferred to the lessee. Therefore, the lessee needs to ask the lessor to update the ROA records.

C. Combined IP Transit and Leasing

The aforementioned scenarios are not mutually exclusive; in practice, they can overlap; for example, a customer might lease IP prefixes and announce them through IP transit services, combining both leasing and transit arrangements as illustrated

¹In cases where the lessor reassigns the IP addresses to the lessee through RIRs, the lessee gains control over ROAs as well as WHOIS records. However, since the ownership of the IP prefixes is officially reassigned to another entity (i.e., the lessee), the lessee could refuse to return the addresses to the lessor after the lease ends; thus, in practice, the majority of IP leasing arrangements do not involve reassigning [48], [58], [14], [28].

in Figure 2(f) where an AS X leases an IP prefix from AS Y and uses AS Z for IP transit service. Thus, if AS X decides to use a transit provider AS Z, it also has to communicate with the IP broker or lessor to ask them to update their ROAs to include their transit provider.

Identifying this scenario is more challenging as it involves three ASes—the ROA origin, the BGP origin, and the transit provider's AS—none of which have any relationships in WHOIS or BGP datasets.

D. Case Study

We present real-world examples for hidden types of ROA misconfigurations that we presented in Figure 2.

Intra-Organizational Transit Services: Misconfigurations can arise when one ISPs providing transit services with multiple ASNs, while the origin ASN in ROA can be different from the upstream ASN seen in BGP path. For example, AS 6596 was announcing an RPKI-invalid prefixes 65.50.199.0/24 with its upstream AS 11404 in June 30, 2023, which cannot be explained by existing heuristics [9], [59], [26] since the ROA origin of that prefix is AS 54858. However, AS 11404 and AS 54858 are operated by the same organization, Wave Broadband, that provides IP transit service.

Opaque IP Transit: When misconfiguration happens during opaque IP transit service, the customer ASN may not be visible in the BGP path, and the BGP based AS relationship dataset used in previous works [9], [59] will not capture it. As an example, AS 5413 (Wavenet) was announcing an RPKI-invalid prefix owned by AS 202364 in July 2023 (185.119.109.0/24), and we further see Wavenet provides opaque transit service using VPN tunnel [2] and are providing transit services to multiple customers. On the other hand, AS 202364 and AS 5413 did not have any relationship shown in the CAIDA AS relationship dataset [12], thus previous works will not be able to detect this misconfiguration.

IP Leasing through Brokers: Misconfigurations of ROAs could happen when a prefix is leased through a broker. For example, AS 49870 was announcing an RPKI-invalid prefix 45.67.13.0/24 in July 2023. We found this prefix has a *mnt-by* field in its IRR record pointing to a famous leasing broker, IPXO [58]. The ROA origin of this prefix, AS 51722, was a previous lessee of this prefix, and there is no relationships between AS 49870 and AS 51722 in BGP and registry data.

Direct IP Leasing from Owners: Instead of leasing through a broker, a prefix can be directly leased from its owner. In this scenario, the lessor and lessee also do not have a direct BGP relationship. For example, AS 5650 was announcing an RPKI-invalid prefix, 23.230.45.0/24, in June 2024 when the ROA origin of this prefix was AS 18779. While there was no BGP relationship between AS 5650 and AS 18779, we found that AS 18779 (EGIHosting) is a well-known direct IP leasing provider [20].

V. INFERRING THE HIDDEN TYPES OF ROA MISCONFIGURATIONS

In this section, we focus on 7,209 RPKI-invalid prefixes that remain unexplained by previous works, and present our

methodologies for detecting ROA misconfigurations caused by the four hidden scenarios described in §IV.

A. Intra-Organizational Transit Services

This type of misconfiguration occurs when the BGP announcement originates from an AS that belongs to the same organization as the upstream AS authorized in the ROA record. To identify this type of ROA misconfiguration, for each RPKI-invalid prefix, we extract the *provider* of the BGP origin AS from the AS_PATH (e.g., AS Y in Figure 2(b)) and obtain its organization using the as2org dataset [11]. We then check if this organization matches that of the ROA origin AS (e.g., AS Z in Figure 2(b)). If there is a match, it suggests a strong indication of misconfiguration from the IP transit service provider within the same organization.

This approach attributes 9.5% of the unknown RPKI-invalid prefixes.

B. IP Leasing Through Brokers

We first obtained a ground truth list of leasing brokers from [20], which contains the organization handles of 163 leasing brokers.² We then use Internet Routing Registry (IRR) databases, including all five RIRs' and RADb [6], to identify the RPKI-Invalid prefixes registered under these organization handles, resulting in a total of 1,766 prefixes. However, these leasing brokers may also provide other services, such as cloud hosting, using their own prefixes (not the prefixes leased from lessors). Therefore, the prefixes registered in the IRR databases under the brokers' organization handles do not necessarily represent leased prefixes.

To accurately identify the leased prefixes, we filter out prefixes that are announced by the leasing brokers themselves (i.e., these are likely to be the brokers' own prefixes used for their services). We also exclude prefixes whose covering less-specific prefixes are not announced by ASes other than the brokers. If the less-specific prefix is not announced by another AS, it suggests that the broker owns the entire address block, and the more-specific prefixes are likely not leased.

Applying these criteria, we filter out 317 prefixes, leaving 20.1% (1,449) of the unknown RPKI-invalid prefixes attributed to leasing.

C. Opaque IP Transit and Direct IP Leasing

Identifying RPKI-invalid prefixes caused by opaque IP transit or direct IP leasing is challenging because the ROA-authorized AS (e.g., a lessor or customer) and the BGP-originating AS (e.g., a lessee or provider) lack visible AS-level relationships in BGP data.

To tackle this problem, we focus on identifying *service providers* likely involved in these arrangements. Our reasoning is that an RPKI-invalid prefix advertised by a well-known transit provider (e.g., Cloudflare) or by a block leased from a major lessor (e.g., Cogent) is far more likely to reflect a

²When prefixes are leased through brokers, it is common for the brokers to add their own information to registration databases such as IRRs maintained by RIRs [20].

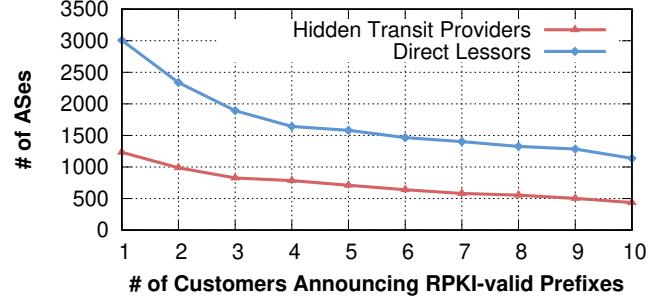


Fig. 3: Number of classified opaque IP transit and direct leasing providers based on their respective customer counts.

hidden misconfiguration than hijacking. Legitimate providers tend to share two observable traits:

- **Operational Scale:** they consistently serve multiple customers over a long horizon (we use 18 months), filtering out one-off or transient networks.
- **RPKI Compliance:** there must be at least one customer that announces RPKI-valid prefixes, as it is unlikely that *all* of their customers are announcing RPKI-invalid prefixes.

1) Classifying Providers: We distinguish between two provider types based on their prefix announcement patterns:

- **Opaque IP Transit Providers:** These providers announce *more-specific* prefixes (e.g., /24) that are subsets of *less-specific* prefixes (e.g., /16) advertised by their customers. For example, in Figure 2(c), the provider (AS Y) announces a customer's /24 prefix under its own ASN, bypassing the customer's ROA. Since the ROA remains tied to the customer's AS, the announcement becomes RPKI-invalid.
- **Direct IP Leasing Providers:** Conversely, these lessors advertise *less-specific* prefixes (e.g., /16), while their lessees announce *more-specific* sub-prefixes (e.g., /24) under their own ASNs (Figure 2(e)). If the lessor's ROA fails to authorize the lessee's AS, these sub-prefixes become RPKI-invalid.

2) Methodology: To identify these providers, we proceed as follows:

- For each RPKI-valid announcement from an AS in our 18-month BGP dataset, we check whether other ASes announce:
 - *Less-specific* RPKI-valid prefixes covering this prefix (to detect potential IP transit providers).
 - *More-specific* RPKI-valid prefixes encompassed by this prefix (to detect potential direct IP leasing providers).
- For each AS identified in step (a), we count how many distinct ASes announce such less-specific or more-specific prefixes. We thereby isolate those providers serving multiple customers. Although identical-length announcements occasionally arise—such as when a provider offers temporary DDoS scrubbing or load balancing—these situations are uncommon; leasing providers seldom delegate their *entire* address space to customers [20], and there is no operational advantage in having both the provider and the

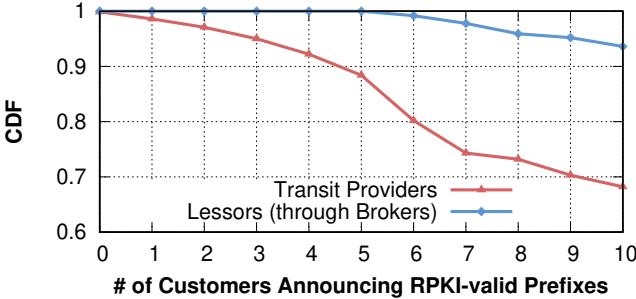


Fig. 4: All IP brokers and 88.4% of known transit providers have more than five ASes announcing RPKI-valid prefixes.

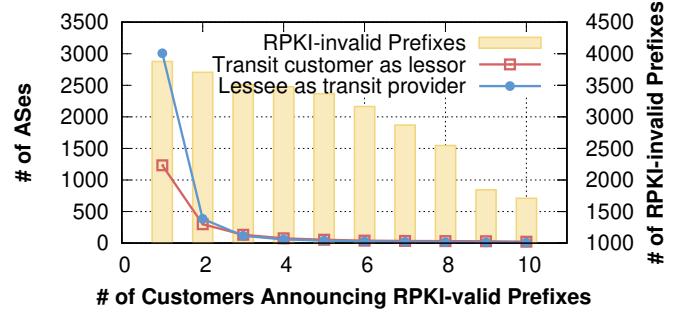


Fig. 5: Misclassifications versus customer threshold, alongside coverage of opaque IP transit providers and direct IP lessors.

customer originate the same prefix. Consequently, we only focus on less-specific and more-specific advertisements; equal-length announcements are examined only when classifying individual prefixes.

Figure 3 illustrates the distribution of potential customers per provider. As shown, only a small number of ASes serve large numbers of customers, and overall, the inferred number of transit providers and direct lessors decreases as the size of their customer base grows. However, selecting an operationally meaningful threshold is challenging: a threshold that is too low (e.g., one customer) may inflate the number of ASes labeled as providers, resulting in excessive false positives.

To refine the threshold, we examine two sets of known providers: (1) Transit providers inferred from standard provider-customer relationships in BGP and (2) identified IP brokers. Although standard transit providers and opaque IP transit providers may differ in their business models, both demonstrate similar customer-prefix announcement patterns. Thus, understanding known providers’ customer distributions informs an appropriate threshold.

Figure 4 illustrates that *all* IP brokers serve more than five RPKI-valid customer ASes, while 88.4% of known transit providers also meet this “at least five customers” criterion. Consequently, we adopt a threshold of five customers as our baseline, striking a balance between coverage and specificity.

3) Evaluation: Evaluating our methodology is challenging because no public ground truth lists which ASes offer opaque transit or lease prefixes. Thus, we adopt a three-pronged validation strategy to gauge potential *misclassification*.

Self-Consistency and Threshold Sanity Checks: Even with our 5-customer threshold, some anomalies can slip through. An AS that merely *borrow*s several prefixes from different lessors could look like a transit provider, and an AS that multi-homes its own space via several upstreams could look like a lessor. Thus, a stronger sign of error is an AS that our method tags simultaneously as a transit provider *and* as a lessor.

Figure 5 shows that when the threshold is only one customer, all 3,007 ASes classified as lessees are also classified as transit providers—a clear overestimate. As we raise the threshold, these misclassifications drop sharply; at five customers, only 55 ASes receive both labels. We consider this threshold to offer a reasonable trade-off between coverage and accuracy:

it is unlikely that a genuine IP transit customer would also use four or more other transit providers.

Validation with Network Operators: We further validate our results by contacting the 55 ASes that our approach labeled both transit providers and lessors. Using the `OrgTechEmail` field in WHOIS, we requested clarification on whether they indeed perform both roles. Of the 18 respondents, 16 *confirmed* our classification was correct. The other two *denied* offering IP leasing services. On closer inspection, these two denials stemmed from mergers and sub-prefix reallocations between newly combined ASes, causing our method to interpret them as a separate leasing relationship. This underscores a limitation in mapping AS numbers to organizations (`as2org`), which we discuss further in §V-D.

Cross-Checking Against Known Hijackers: Our classification could mistakenly label some hijackers, who announce numerous stolen prefixes, as legitimate providers serving multiple “customers”. To check for this, we use 274 confirmed hijacker ASNs from [56], [8]. Of the 208 RPKI-invalid prefixes these hijackers announced in our 18-month dataset, *none* were classified as misconfigurations by our method.

Nevertheless, because these hijacker datasets (primarily collected between 2008 and 2020) could be incomplete or outdated, we further validate our hidden misconfigurations through three additional steps: (1) cross-referencing hijack alerts from Cloudflare Radar [15] and GRIP [1], (2) tracking whether the invalid ROA records are corrected within nine months, and (3) directly confirming with 174 network operators, as described in §VII.

Threshold-based Classification Results: Finally, with 5-customer threshold, we identify 1,582 ASes as transit providers and 710 as lessors, encompassing 2,170 RPKI-invalid prefixes from opaque IP transit (30.1% of the previous unknown prefixes, or 5.1% of the total) and 1,197 from direct IP leasing (16.6% of the previous unknown prefixes, or 2.8% of the total).

Additionally, 367 prefixes are involved in both opaque IP transit and direct IP leasing, while 88 prefixes are involved in both opaque IP transit and broker leasing. In total, 455 prefixes participate in a combined business of IP leasing and IP transit, accounting for 6.3% of the previously unknown prefixes, as illustrated in Figure 6.

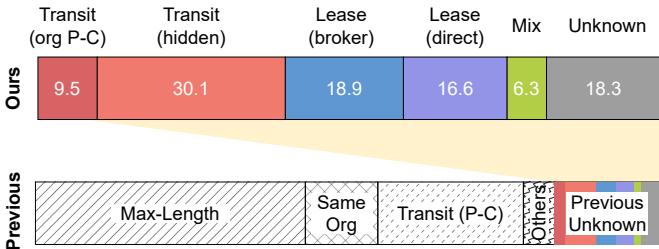


Fig. 6: Percentage of hidden types of ROA misconfigurations covered by our methodologies (13.8%). The gray area represents misconfigurations already reported in the literature.

D. Limitations

Our approach has several limitations; we wish to discuss them below before showing the final classification results.

Potential Hijack Misclassifications: Although no known hijacks appear in our evaluation dataset (§V-C3), certain hijacking incidents could still be incorrectly flagged as misconfigurations if hijackers *target* leased or opaque IP transit prefixes. For instance, a hijacker acquiring control of a leased prefix registered under a broker might appear to be involved in an IP leasing misconfiguration. Moreover, hijackers may exploit legitimate transit services to launch attacks—historical cases involve major ISPs inadvertently (or deliberately) rerouting traffic. While our goal is to identify hidden ROA misconfiguration patterns, rather than precisely detect hijacks, we later show in §VII that our classification results can help reduce hijack false positives.

Mapping ASes to Organizations: To identify intra-organizational transit, we use `as2org`[11], which maps ASes to organizations via WHOIS data. However, this mapping falters when entities operate multiple ASNs under distinct handles or contact information (e.g., in mergers and acquisitions). For example, AS6128 and AS54004 both belong to Optimum[4], [5] but are registered under different names. BGP data alone cannot reveal such corporate relationships; consequently, tools like `as2org`[11] and `as2org+`[3] treat them as separate. As noted in prior research [3], this represents an open challenge in mapping AS numbers to real-world organizations.

E. Classification Results

Figure 6 presents our final classification results. For the 42,654 RPKI-invalid prefixes in our 18-month BGP dataset, three previous works combined successfully identified 83.1% of them as misconfigurations, leaving 7,209 classified as unknown.

Our classifications further explain 5,887 (81.7%) of them: 685 (9.5%) are caused by organization-level provider-customer (P-C) transit, 2,170 (30.1%) by opaque IP transit, 1,363 (18.9%) by leasing business with brokers (excluding 88 prefixes involved in combined scenarios), 1,197 (16.6%) by direct leasing, and 455 (6.3%) are involved in both IP leasing and IP transit businesses. This leaves 1,319 (18.3%) of the previously unknown prefixes, or 3.1% of the total RPKI-invalid prefixes in our dataset, remain unexplained.

VI. IMPACT OF RPKI-INVALID PREFIXES: OPERATIONAL CONSEQUENCES

Misconfigured ROAs can lead to significant operational problems, including connectivity failures and suboptimal routing. While some RPKI-invalid routes remain accessible via alternate paths [10], [50], [47], the rising adoption of ROV—now protecting 23.3% of ASes [37]—increases the likelihood of strict filtering, potentially cutting off access entirely. Such disruptions also pose security risks, as adversaries could exploit unintended gaps or detours to evade defenses; if a resource depends on DDoS protection or firewalls at a specific upstream, diverted traffic bypasses those safeguards and leaves the resource vulnerable. In this section, we investigate how ROA misconfigurations can trigger these operational and security consequences. We focus on two different types of impacts:

- (a) **Connectivity Problems:** RPKI-invalid prefixes may become unreachable due to ROV filtering, leading to connectivity loss. To measure how many RPKI-invalid prefixes are unreachable, we send ICMP echo request to responding hosts under RPKI-invalid prefixes and check whether they can still be reached. Since ROV deployments are different across different networks, we run the measurements from various vantage points on the Internet.
- (b) **Unintended Path Divergence:** Even if RPKI-invalid prefixes remain reachable, traffic may be rerouted through the alternative path (in some cases, unintended paths) due to ROV filtering. Although path changes could happen in any hop of the path, we mainly focus on the upstream AS (the last hop AS before the origin AS), which could lead to the most significant impact on both security and performance: since the upstream AS is usually the one providing transit services as well as security services like DDoS protection, if the traffic is diverted to a different upstream AS, it may bypass these security services and expose the resource to potential attacks. To measure the path divergence, we run traceroutes from various locations on the Internet toward the responding hosts under RPKI-invalid prefixes and compare the penultimate hop ASes with the BGP AS_PATH.

Below, we first describe our experimental design and then present our findings on all RPKI-invalid prefixes captured in one-day BGP dumps.

A. Experimental Design

We use ICMP ping and traceroute to assess connectivity and routing paths to RPKI-invalid prefixes.

- (a) On June 13th, 2024, we used Routinator [54] to validate ROAs and identified 7,043 RPKI-invalid prefixes.
- (b) We then used ZMap to find hosts responding to ICMP echo requests; the scan originated from a non-ROV network to increase the likelihood of reaching more responding hosts, yielding 386,102 hosts covering 6,152 (87.3%) RPKI-invalid prefixes from 4,280 ASes.³ We classify these RPKI-invalid prefixes as 3,627 single-organization

³It is worth mentioning that even when the scan is performed from a non-ROV network, we may still miss responding hosts due to other ASes on the path that have deployed ROV [25].

	Disconnection (%)			Path Divergence (%)		
	0	0-25	25-100	0	0-25	25-100
Total	96.9	2.0	1.1	81.5	12.1	6.4
Max Length	97.1	1.9	1.0	-	-	-
Same ORG	96.7	2.3	1.0	-	-	-
P-C Transit	98.3	1.5	0.2	89.6	7.6	2.8
Org Level Transit	98.0	1.5	0.5	74.2	15.1	10.7
Hidden Transit	97.6	1.8	0.6	70.0	19.4	10.6
Direct Leasing	95.5	3.1	1.3	68.5	21.5	10.0
Broker Leasing	94.3	4.3	2.0	65.1	23.3	11.6
Leasing + Transit	95.2	3.7	1.8	62.3	20.2	17.5
Unknown	91.6	5.0	2.6	63.9	22.8	13.3

TABLE II: Breakdown of how ASes experience disconnection and path divergence due to ROA misconfigurations. Each column group (*Disconnection* and *Path Divergence*) is split into three bins (0%, 0–25%, 25–100%), indicating the fraction of RIPE Atlas probes (per AS) affected by unreachable RPKI-invalid prefixes or altered paths. While 96.9% of ASes have no disconnection at all (0%), 81.5% also have zero path divergence—meaning 18.5% of ASes observe some degree of unintended routing; leasing-related misconfigurations (e.g., “Broker Leasing”) show higher rates of both partial and full disruption compared to transit-only arrangements.

cases, 1,850 from IP transit, 452 from IP leasing, 89 from combined transit and leasing, and 136 as unknown.

- (c) We randomly selected one representative host for each RPKI-invalid prefix to minimize potential negative impact on the hosts.
- (d) We consider only RIPE Atlas probes with at least three probes per AS, resulting in 5,043 vantage points across 1,681 ASes covering 176 countries in 5 continents—1,149 with a 0% ROV score and 532 with a ROV score above 0% based on RoVista [37].
- (e) For each probe, we send ICMP pings and traceroutes to the 6,152 representative hosts under RPKI-invalid prefixes, measuring connectivity, routing paths, and latency information. To ensure reliability, we confirmed that all three probes within each AS exhibited consistent behavior regarding the paths.
- (f) We continuously monitor these RPKI-invalid prefixes and their BGP announcements. Once they are corrected (either by updating ROAs or BGP announcements), we wait two hours to allow the new updates to propagate⁴ before re-conducting the ICMP ping and traceroute experiments.

Finally, throughout the process, we obtained 6,152 responding hosts, each under a unique RPKI-invalid prefix; and 5,043 RIPE Atlas probes from 1,681 ASes.

B. Connectivity Problems

We first examine how different misconfigured ROAs can cause disconnection from the Internet. Table II shows the distribution of ASes that are unable to reach representative hosts under RPKI-invalid prefixes.

First, we find that overall 191 (3.1%, out of 6,152 responding hosts) cannot be reached from at least one RIPE Atlas

⁴According to [23], more than 97% of end-to-end delays for propagating BGP and ROA updates are less than 100 minutes.

probe. While this might be due to client-side artifacts, we further analyze these 159 ASes that announce RPKI-invalid prefixes and find that 184 (96.3%) of them also announce RPKI-valid or unknown prefixes. Using ZMap, we also identify responding hosts and test their connectivity from RIPE Atlas probes. We find that *all of them are able to reach the destination, indicating that the disconnection is indeed due to ROV*.

Second, interestingly, transit misconfigurations exhibit the lowest likelihood of disconnection, ranging from 1.7% for conventional provider-customer transit to 3.4% for opaque transit. This is likely because transit providers typically ensure traffic is forwarded properly, some even whitelist RPKI-invalid prefixes [37]. By contrast, disconnection issues are more prevalent for *leased prefixes*. For instance, 4.5% of RPKI-invalid direct-leasing prefixes and 5.7% of RPKI-invalid broker-leased prefixes result in disconnection.

The higher rate of leasing-related disconnection likely reflects the limited visibility that IP leasing services have into lessee route details, thereby exposing gaps in RPKI management. To verify whether these 191 disconnected prefixes could rely on an alternative RPKI-valid route, we cross-reference the BGP data. Interestingly, 24 (12.6%) of these prefixes *do* have alternative RPKI-valid announcements yet remain disconnected, as highlighted by our BGP routing table analysis. This observation challenges assumptions from earlier studies [10], [50], [47]. A common scenario appears when prefixes are leased through brokers (i.e., Figure 2(d)) where the lessor does not forward traffic to their lessee.

Resolving Connectivity Problems: We now investigate how this connectivity problem is resolved. Figure 7 (top) shows the distribution of the time it takes to fix them, and we compare this with prefixes that do not cause connectivity issues (Figure 7 (middle)).

First, we find that 34.2% of RPKI-invalid prefixes causing disconnection are fixed within one day, and 98.7% within two months. In contrast, non-disruptive prefixes are rarely fixed quickly—only 1.2% within one day and 40.6% within two months. When examining the types of misconfigurations more closely, we observe interesting patterns; if the connectivity issues occur *within an organization*, such as due to `maxLength` issues or misconfigurations within the same organization, the problems tend to be fixed quickly; for example, 35.2% and 39.5% of such connectivity issues are resolved within one day, respectively.

In contrast, cross-organizational issues take longer to fix: only 24.8% (provider-customer transit) and 30.5% (hidden transit) are resolved within one day. Leasing takes even longer, with just 15.2% (direct leasing) and 17.4% (broker leasing) fixed within a day.

This highlights the difficulty of quickly resolving misconfigurations involving multiple organizations or delegated prefixes, as will be discussed in §VIII.

C. Unintended Path Divergence

RPKI-invalid prefixes, even without causing connectivity issues, may reroute traffic through unintended paths. For example, as shown in Figure 2(c), /24 RPKI-invalid routes might

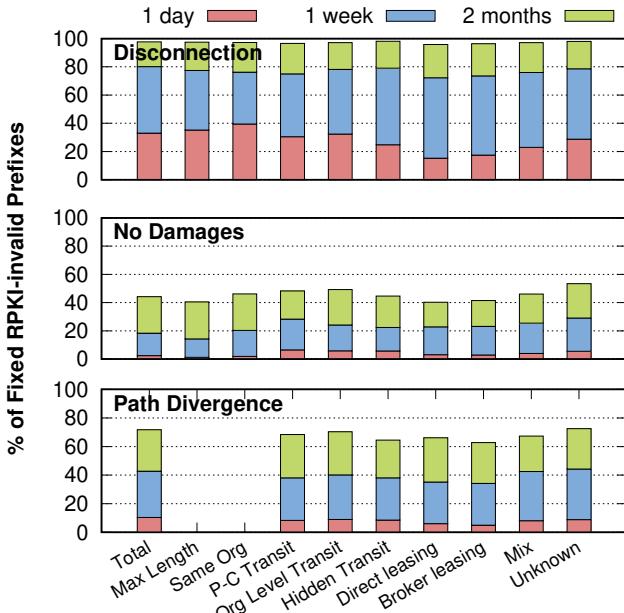


Fig. 7: 98.7% of ROA misconfigurations causing connectivity issues are fixed within two months, compared to 70.2% for those causing path divergence and 40.6% for those causing no damage. Note that there is no path divergence for Max Length and Same Org misconfigurations.

be forced to use valid /16 routes to reach AS X. However, if AS X uses different upstream providers for /16 routes, traffic may bypass AS Y, rendering its transit services non-functional and potentially increasing latency or compromising security—particularly if AS Y provides services like DDoS protection.

Now we aim to measure how many RPKI-invalid misconfigurations result in path divergence.

Identifying Unintended Paths: To determine whether traffic traverses unexpected upstream ASes, we first locate the upstream ASes by analyzing traceroute data. We use the same methodology as [41], [60], examining the penultimate hop AS in the traceroute responses; if the penultimate hop AS is never listed in any AS_PATH of that prefix in our BGP datasets, this may indicate that the penultimate hop is due to an unintended path change caused by RPKI invalidity.

By analyzing the penultimate ASes, we find that 1,268 (18.5%) RPKI-invalid prefixes potentially result in path divergence, which is six times more than the possibility of disconnection (3.1%).

Path divergence occurs only in cross-organizational misconfigurations—it never occurs in the ‘Max Length’ or ‘Same ORG’ scenarios (0%); Table II shows its distribution for IP transit and leasing services. Similar to the disconnection problems, we find that ROA misconfigurations during IP leasing are significantly more likely to cause path divergence (31.5%) compared to IP transit issues (10.4%). A possible reason is that IP transit providers often handle both RPKI-invalid and other prefixes from the same customers, leading to shared upstream ASes and minimizing path divergence.

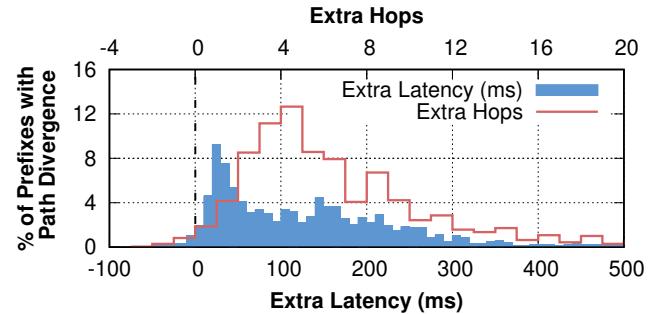


Fig. 8: Extra latency and hop counts caused by path divergence.

Considering that only 31.6% of RIPE Atlas nodes fully protected by ROV, path divergence poses a greater impact than disconnection. Since ROV can occur at *any hop*, 40.1% (509) of the 1,268 RPKI-invalid prefixes causing path divergence actually originate through non-ROV next-hop ASes, showing divergence can happen regardless of the source AS’s ROV status.

Resolving Path Divergence Problems: Now we focus on the 1,268 RPKI-invalid prefixes with potential path divergence issues to understand how they are fixed; Figure 7 (bottom) shows the time taken to get fixed. Interestingly, we find that these issues are *slower to resolve* than disconnections, with only 10.8% fixed within a day and 70.2% within two months, likely because path divergence is harder to detect than disconnection. Furthermore, only 64.5% (574) of these RPKI-invalid prefixes are fixed by correcting the ROA or updating BGP announcements.

Path Changes After Fixes: One limitation of our methodology using the penultimate AS is that path divergence may not be solely caused by RPKI invalidity; route engineering and locally visible routes can also cause the penultimate hops to differ from the BGP AS_PATH. To address this, we focus on the 574 (64.5%) prefixes fixed by updating the ROA; after they are fixed, we find that 94.8% of the potential path divergence issues no longer exist; that is, their penultimate ASes now appear in the BGP dumps. This indicates that most of the path divergence observed during the traceroute measurements was indeed caused by RPKI invalidity.

Performance Issues Due to Path Divergence: When path divergence occurs, it may introduce additional latency due to suboptimal path selection. To measure the latency changes, we re-sent ICMP pings and traceroutes to the RPKI-invalid prefixes 100 minutes after observing that they had been fixed by updating the ROA. Figure 8 shows the average Round Trip Time (RTT) and hop count changes before and after the path divergence was resolved. We make a number of observations: we immediately notice that most (97.1%) RPKI-invalid prefixes with path divergence incurred additional latency, with 39.2% exceeding a 100 ms delay before being fixed.

These additional latencies are mainly caused by the increased number of hops in the path when path divergence occurs. We observe that path divergence also increases hops,

	Detection Platform Radar	Platform GRIP	Combined
Detected Prefixes	17,282	11,450	20,821
RPKI Invalid Prefixes	5,539	4,675	6,114
MaxLength	0	0	0
Same organization	0	0	0
P-C Transit	0	0	0
Org-Level Transit	484	353	521
Opaque IP Transit	1,708	1,479	1,932
Direct Leasing	1,003	980	1,042
Broker Leasing	972	836	1,135
Leasing & transit	324	124	380
Sum.	4,491 (81.1%)	3,772 (80.7%)	5,010 (81.9%)

TABLE III: Number of prefixes flagged by two hijack-detection systems, along with the potential false positives stemming from ROA misconfigurations. Both systems already filter out the misconfigurations documented in [9] such as conventional provider-customer transit (Figure 2(a)).

with 96.9% of affected prefixes showing additional hops and 54.7% exceeding five extra hops.

VII. IMPACT OF RPKI-INVALID PREFIXES: HIJACKING DETECTION

Since ROA records increasingly serve as the “ground truth” for prefix ownership, hijack detection systems such as Cloudflare Radar [15] and GRIP [1] commonly rely on them. These detection systems examine multiple BGP attributes (e.g., AS_PATH, organization mappings) and employ heuristic scoring to determine whether an announcement is a hijack event.

Motivation. Although hijack detectors verify ROA, they suppress most alerts by filtering out familiar misconfigurations—such as announcements made by the same ISP or by a visible provider–customer pair [9], [16]. A prefix may raise a suspicion score at first, but the alarm is cleared once the system links it to that ISP. With the rise of subtler ROA errors, however, we must ask: *Do these “hidden” misconfigurations inflate false-positive rates?* In this section, we measure how often leading detectors misfire on such cases.

A. Datasets

We use two publicly available hijack detection systems from both industry and academia that incorporate ROA: Cloudflare Radar and GRIP, as representative examples. Although these hijack detection systems have their own algorithms that leverage BGP information—such as AS-path attributes and CAIDA’s as2org and AS relationship datasets—to filter legitimate multi-origin announcements, they may not account for hidden types of misconfigurations, which may lead to false positives in hijack detection.

We obtain all inferred hijack events⁵ from these two systems reported from January 1st, 2023, to July 1st, 2024. The combined dataset includes 20,821 unique prefixes labeled

⁵Specifically, we include multi-origin announcements (MOAs) marked as hijacks by Cloudflare Radar (confidence score > 8) and by GRIP (suspicious level > 80).

Category	Detection Platform Radar	Platform GRIP	Combined
Not H-Misconfig.	1,048	903	1,104
ROA Fixed (9 mos)	10 (1.0%)	5 (0.6%)	13 (1.2%)
H-Misconfig.	4,491	3,772	5,010
<i>No Longer Hijack</i>	3,844	2,725	3,916
ROA Fixed (9 mos)	2,338 (60.8%)	1,792 (65.8%)	2,524 (64.5%)
<i>Still Hijack</i>	647	1,047	1,094
ROA Fixed (9 mos)	432 (66.7%)	721 (68.8%)	755 (69.0%)

TABLE IV: Prefixes that remained flagged as hijacks after re-running each detection system; we nullified their scores if classified as hidden misconfigurations. the “Fixed (9 mos)” columns show how many of those flagged prefixes later obtained valid ROAs.

as hijacks, including 17,282 from Cloudflare Radar and 11,450 from GRIP (Table III). We then check the RPKI status of these prefixes and find 6,114 prefixes are RPKI-invalid, 14,597 are not covered by RPKI, and only 110 RPKI-valid prefixes.

B. Results

We apply our classification methodology to these 6,114 RPKI-invalid prefixes and find that 5,010 (82.0%) are likely due to hidden misconfigurations rather than true hijacks; Table IV breaks down the distribution of different misconfiguration types.

We observe that *no* prefixes attributed to well-known ROA misconfigurations (i.e., maxLength, same-organization, or provider–customer transit) appear in the hijack lists—consistent with their built-in mechanisms to exclude common errors [16]. In contrast, “hidden” misconfigurations, including opaque IP transit (1,932 prefixes, 39%) and leasing (2,177 prefixes, 43%), are prominent factors in these potential false positives.

C. Refining Hijack Detection with Hidden Misconfiguration Knowledge

To gauge how recognizing hidden ROA misconfigurations might reduce false positives, we re-run each hijack detection system’s inference algorithms for the misclassified prefixes *without* penalizing them for failing ROA checks. In other words, we treat these hidden misconfigurations the same way both systems already treat other known misconfigurations (such as visible provider–customer relationship):

- (a) *Cloudflare Radar* is not fully open-source, but its scoring logic is public [16]. We obtained each event’s detailed scores and heuristic tags, then re-executed the detection pipeline excluding ROA-based tags for prefixes we identified as misconfigurations.
- (b) *GRIP* also provides detailed scores and rationale via its public API. We similarly re-ran its scripts after omitting ROA-related tags for the relevant misconfigured prefixes.

Table IV summarizes the outcome. Of the 5,010 RPKI-invalid prefixes we classified as misconfigurations in Table III, 3,916 (78.1%) fell below the high-suspicion hijack threshold once their ROA data were excluded.

Confirming at scale that these events were *indeed* non-hijacks is challenging, owing to the covert nature of hijacking

and the limited availability of definitive ground-truth data. While earlier efforts [56], [46] have attempted validation by referring network operators’ mailing list (e.g., NANOG [43]) or using small annotated datasets, they have only covered at most 156 known hijack prefixes, and no fully comprehensive framework for large-scale validation currently exists.

Thus, in the following section, we incorporate operator feedback and additional external validation methods to further corroborate our findings.

Monitoring Fixed ROAs Over Time: As discussed in §VI, genuine prefix owners tend to correct misconfigured ROAs eventually, while malicious hijackers generally lack the authority to acquire valid ROAs. We thus track whether “hijacked” prefixes—especially those classified as hidden ROA misconfigurations—ultimately appear with valid ROAs, indicating potential false positives in the hijack detection systems.

- (a) *Non Hidden-Misconfiguration.* Of the 1,104 RPKI-invalid prefixes not deemed hidden misconfigurations (and thus labeled as hijacks), 13 (1.2%) ever obtained valid ROAs. This low correction rate implies that many of these announcements are genuine hijacks, as the route announcers do not appear to have legitimate authority to fix the ROA status.
- (b) *Hidden-Misconfiguration.* We initially identified 5,010 RPKI-invalid prefixes as hidden misconfigurations. Among these, 3,916 (78.1%) no longer qualified as hijacks after ignoring ROA-based penalties, whereas the remaining 1,094 (21.9%) still triggered hijack flags.
 - (a) *Cleared Hijack Labels.* Of the 3,916 prefixes dropping off the hijack list, 2,524 (64.5%) updated their ROAs within nine months, reinforcing that these were unlikely to be true hijacks.
 - (b) *Persistent Hijack Labels.* Even for the 1,094 still flagged as hijacks, 755 (69.0%) eventually obtained valid ROAs, further underscoring that many presumed hijacks were in fact benign misconfigurations.

These findings reveal two observations: (1) Only about 1.2% of the “non hidden-misconfiguration” prefixes ended up with valid ROAs, confirming that ignoring hidden misconfigurations can indeed reduce false positives. (2) Although the majority of hidden misconfigurations eventually get fixed, *a nontrivial subset lingers*, which needs further investigations.

Resolving Unfixed Cases via Operator Feedback: Even after accounting for prefixes that eventually gained valid ROAs (either by having their legitimate owners fix ROAs or by correcting corresponding BGP announcements), we still find 1,731 cases ($= 5,010 - 2,524 - 755$) where no valid ROAs emerged. However, RPKI-invalid routes can also be corrected through other means (e.g., withdrawal of conflicting announcements), making it difficult to distinguish true misconfigurations from potential hijacks solely by tracking ROA changes.

To address this gap, we conducted a large-scale, direct survey of legitimate prefix owners. Specifically, we retrieved contact information from RIR WHOIS databases for the 1,731 prefixes in question, acquiring valid email addresses. We then contacted these owners, asking whether each RPKI-invalid prefix was a misconfiguration or a genuine hijacking. The

survey methodology and ethical considerations are detailed in Appendix IX-C.

Among the 1,385 email addresses we contacted, messages to 302 were returned as undeliverable. From the remaining set, 174 organizations responded, confirming that *all 349 RPKI-invalid prefixes identified as hidden misconfigurations were indeed misconfigured rather than hijacked*.

VIII. SURVEY AND MITIGATION SUGGESTIONS

Our results indicate that ROA misconfigurations remain prevalent in a way that is not easily noticed. To understand the reasons behind this persistent misconfiguration and to explore potential solutions, we surveyed 8 large ISPs and 8 major leasing brokers involved in the misconfigurations we detected.

- (a) *Surveying ISPs.* We collaborated with CableLabs [13], a research organization that works with broadband operators worldwide. They helped us connect with eight large ISPs—three in the top 100, two in the top 500, and three in the top 1,000 AS rankings [36]—to explore why ROA misconfigurations happen when they provide transit services, including 4 online interviews and 4 email surveys.
- (b) *Surveying leasing brokers.* To survey leasing brokers, we reached out to eight major brokers from top 10 registered leasing brokers with most leased prefixes from a recent study [20] to understand how they manage ROAs for leased prefixes and why ROA misconfigurations happen, including 1 online interview and 7 email surveys.

This outreach consisted of 11 email surveys and 5 online interviews. Although these interactions may not reveal all root causes, they shed light on unanswered questions.

We asked the participants to identify the underlying reasons for ROA misconfigurations and the difficulties of resolving RPKI-invalid announcements (the full set of questions is in §B). From their responses, we highlight several key challenges:

- (a) *Lack of Real-Time Monitoring.* All 8 ISPs reported that they lack automated systems to continuously verify ROA consistency against live BGP announcements. Two ISPs indicated plans to deploy such tools soon. In the interim, ROA misconfigurations often go unaddressed unless external complaints arise or manual checks detect the problem.
- (b) *Organizational Silos in ISPs.* In many ISPs, separate teams manage ROA records and BGP route announcements. Three ISPs acknowledged that their BGP operations teams sometimes propagate transit customer routes *before* the IP-management team updates the corresponding ROA objects. Poor inter-departmental communication can thus introduce persistent misconfigurations.
- (c) *Inconsistent ROV Filtering by ISPs.* Despite all 8 ISPs fully deploying ROV for peer or customer links, 5 were observed announcing RPKI-invalid routes during our measurement window. This suggests that these ISPs are not applying ROV filtering on their own egress routes. In some cases, the rush to onboard customers (e.g., by sales teams) appears to outpace the ROA update processes overseen by IP-management teams.

- (d) *Manual Procedures in IP Leasing.* Most surveyed brokers rely on manual protocols and communication with prefix owners when updating ROAs, since the lessor (prefix owner) retains formal ROA authority. While 5 out of 8 brokers target a 48-hour window to update ROAs after a lessee requests it, the other 3 impose a more stringent 24-hour deadline. Even so, any human-in-the-loop process can introduce delays and errors.
- (e) *Limited RPKI Delegation.* One promising way to avoid manual intervention is to delegate RPKI control of leased prefixes to the brokers themselves. However, RIRs often do not support selective delegation for specific sub-prefixes without transferring ownership entirely, forcing owners to either run their own RPKI publication setups or grant full ROA management of *all* prefixes to the broker. Only one surveyed broker maintains an infrastructure to receive such delegations; uptake is minimal because most owners are unwilling to hand over full ROA control.

Recommendations: Stepping back, our findings indicate that multiple steps can be taken by various network entities to reduce ROA misconfiguration, which will strengthen the RPKI ecosystem.

- *Transit Providers.* ISPs offering transit should implement automated monitoring tools for real-time ROA and BGP consistency checks, ensuring RPKI-invalid announcements are rapidly flagged. Likewise, they should enforce ROV filtering *across all egress routes*, not just for peer or customer links, to prevent the propagation of RPKI-invalid prefixes before ROA updates are in place.
- *Leasing Brokers.* Leasing brokers need to proactively track ROA statuses for sub-leased prefixes and strive to automate the update process. Where feasible, RPKI delegation could offload ROA administration from owners to brokers without a complete ownership transfer.
- *RIRs.* RIRs should consider enabling selective RPKI management delegation for partial address blocks, allowing prefix owners to grant brokers limited control over ROA records. This approach would reduce administrative overhead for both parties and help maintain accurate ROA information.

IX. CONCLUSION

In this paper, we examined the prevalence and impact of RPKI-invalid prefixes on the Internet. We developed a classification method that categorized 96.9% of these invalid prefixes, linking them to underlying ROA misconfigurations; our analysis revealed four different scenarios during IP transit and IP leasing services that can lead to ROA misconfigurations, filling the gap left by previous studies.

We evaluated the broader impact of RPKI-invalid prefixes, including connectivity disruptions, unintended path alterations, and false alarms in hijack detection systems. Notably, 3.1% of RPKI-invalid prefixes caused connectivity issues, and a significant portion of hijacking incidents flagged by public detection systems were actually misconfigurations. Through surveys, we identified operational challenges ISPs face in maintaining accurate ROAs, including lack of monitoring tools and organizational silos.

As RPKI adoption grows, these challenges are likely to intensify. Our findings highlight the critical need for improved ROA management and monitoring practices across all RPKI participants.

ETHICAL CONSIDERATIONS

Our survey is directed at network-operator organizations, not at individual persons. We gather only factual details about deployed systems and collect no personal data. Our Institutional Review Board (IRB) confirmed that, under these conditions, the study does not constitute human-subjects research and therefore did not require a full protocol review. Nonetheless, we address and discuss all relevant ethical considerations in the sections that follow.

A. Data-plane scanning and probing

§VI includes data-plane scanning and probing toward in-the-wild IP addresses. To find responding hosts under RPKI-invalid prefixes, we perform ICMP scanning on 7,043 prefixes using ZMap [22]. Additional ICMP messages are sent to 6,152 individual responding hosts to measure the connectivity, path divergence, and latency impact of RPKI-invalid.

To address potential ethical concerns, we adhere to the ethical scanning guidelines outlined in [22] and follows the Menlo report [21]. Our scanning and probing only involve minimal traffic, we note that (1) we only send ICMP echo packets without any payload throughout our measurement, (2) for each IP addresses under RPKI-invalid prefixes, it will only be scanned once, (3) while 386,102 responding hosts are found, we only pick minimal and necessary number (6,152) of hosts for probing.

For probing with RIPE Atlas, we send two ICMP echo packets towards each hosts from 5,043 vantage points, resulting to a total of 322.8 KB traffic towards each ISP. We spread out our experiments according to a random permutation of each pair of hosts and RIPE Atlas nodes to minimize the traffic bandwidth. To the end, the average bandwidth we generate toward each host will be less than 10 Bps, and the theoretical peak bandwidth during the measurement will be less than 320 Bps, which we believe is minimal and will not result in ethical concerns.

We also inform local network administrators to mitigate risks and handle any inquiries that may arise. Additionally, we ensure that our scans do not overwhelm the upstream provider by limiting the scanning bandwidth to 80 Mbps. Furthermore, we generate only the necessary amount of traffic required for our research objectives, minimizing any excessive network load.

B. Surveys and Interviews

We conduct a survey with eight ISPs on their management of ROA objects in §VIII, whether by email or online interviews. The participants of the survey freely participated. Our survey focus on the company level behavior and does not require any information about individuals, thus does not raise any ethical concern. We also ask for permissions for any information to be included in our paper before publication.

C. Email

To validate hijack detection alerts linked to hidden misconfigurations, we contacted network operators via publicly listed abuse and technical contacts in WHOIS records—channels explicitly designated for reporting routing anomalies and security issues. These contacts are maintained by operators to receive such notifications, ensuring our outreach aligned with intended use cases. We adhered to ethical guidelines for responsible communication:

- **Targeted Messaging:** We limited outreach to operators of prefixes directly implicated in persistent RPKI-invalid announcements, avoiding unsolicited bulk emails.
- **Transparency:** Each email clearly stated the purpose of our inquiry, included evidence of the misconfiguration, and offered opt-out instructions.
- **Zero Complaints:** Despite contacting operators of 294 prefixes, no recipients reported our communications as spam—a testament to the relevance and non-intrusiveness of our approach.

This validation process not only confirmed that 100% of sampled alerts stemmed from misconfigurations but also demonstrated operator engagement in improving routing security. Several operators explicitly thanked us for identifying overlooked configuration errors, highlighting the mutual value of such academic-industry collaboration.

ACKNOWLEDGMENTS

We extend our heartfelt gratitude to the anonymous reviewers and for their invaluable insights. We also thank Tony Tauber, Taylor Harris, Miles McCredie, Mark Goodwin, Leif Sawyer, Ignas Anfalovas, and all participants for their constructive feedback for the interviews and surveys. We are grateful to Brian Scriber for his help and support on this work. This research was supported in part by NSF grant CNS-2323137, CNS-2339378, and Comcast Innovation Grant.

REFERENCES

- [1] GRIP - Global Routing Intelligence Platform. <https://grip.inetintel.cc.gatech.edu>.
- [2] Wavenet Intelligent connectivity. <https://www.wavenet.co.uk/solutions/intelligent-connectivity>.
- [3] A. Arturi, E. Carisimo, and F. E. Bustamante. as2org+: Enriching as-to-Organization Mappings with PeeringDB. *PAM*, 2023.
- [4] Altice USA Announces Closing of Sale of 49.99% of Lightpath Fiber Enterprise Business to Morgan Stanley Infrastructure Partners. <https://www.alticeusa.com/news/articles/press-release/corporate/altice-usa-announces-closing-sale-4999-lightpath-fiber-enterprise-business-morgan-stanley>.
- [5] Altice acquires Cablevision and creates the #4 cable operator in the US market. <https://www.alticeusa.com/news/articles/press-release/corporate/altice-acquires-cablevision-and-creates-4-cable-operator-us-market>.
- [6] Archive of RADB Repository. <ftp://ftp.radb.net/radb/dbase/archive>.
- [7] R. Bush and R. Austein. The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1. RFC 8210, IETF, 2017.
- [8] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill. BGP hijacking classification. *TMA*, 2019.
- [9] T. Chung, E. Aben, T. Bruijnzeels, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, R. van Rijswijk-Deij, J. P. Rula, and N. Sullivan. RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins. *IMC*, 2019.
- [10] B. Cartwright-Cox. Measuring RPKI Adoption via the data-plane. NLNOG Day 2018. https://nlnog.net/static/nlnogday2018/8_Measuring_RPKI_ben_NLNOD_2018.pdf.
- [11] CAIDA ASOrganizations Dataset. <http://www.caida.org/data/as-organizations/>.
- [12] CAIDA ASRelationships Dataset. <http://www.caida.org/data/as-relationships/>.
- [13] CableLabs. <https://www.cablelabs.com/>.
- [14] Can You Lease an IP Address? <https://iptrading.com/blog/can-you-lease-an-ip-address/>.
- [15] Cloudflare Radar. <https://radar.cloudflare.com/routing>.
- [16] Cloudflare Radar’s new BGP origin hijack detection system. <https://blog.cloudflare.com/bgp-hijack-detection/>.
- [17] Cogent Communications Offers 206 Million in Secured Notes Backed by IPv4 Addresses. <https://circleid.com/posts/20240503-cogent-communications-secures-206-million-in-notes-backed-by-ipv4-addresses>.
- [18] Cogent Customer User Guide. https://cogentco.com/files/docs/customer_service/guide/global_cogent_customer_user_guide.pdf.
- [19] Creating a request for additional IP addresses. <https://www.lumen.com/help/en-us/control-center/support/creating-a-request-for-additional-ip-addresses.html>.
- [20] B. Du, R. Fontugne, C. Testart, A. C. Snoeren, and k. claffy. Sublet Your Subnet: Inferring IP Leasing in the Wild. *IMC*, 2024.
- [21] D. Dittrich and E. Kenneally. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. 2012. https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCORE-20120803_1.pdf.
- [22] Z. Durumeric, E. Wustrow, and J. A. Halderman. ZMap: Fast Internet-Wide Scanning and its Security Applications. *USENIX Security*, 2013.
- [23] R. Fontugne, A. Phokeer, C. Pelsser, K. Vermeulen, and R. Bush. RPKI Time-of-Flight: Tracking Delays in the Management, Control, and Data Planes. *PAM*, 2023.
- [24] V. Giotsas, I. Livadariu, and P. Gigis. A First Look at the Misuse and Abuse of the IPv4 Transfer Market. *PAM*, 2020.
- [25] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman. Are We There Yet? On RPKI’s Deployment and Security. *NDSS*, 2017.
- [26] T. Hlavacek, H. Shulman, and M. Waidner. Smart rpki validation: Avoiding errors and preventing hijacks. *European Symposium on Research in Computer Security*, 2022.
- [27] IPXO RPKI Publication Point. <https://magellan.ipxo.com/rrdp/notification.xml>.
- [28] IPv4 Address Leasing: Ownership Rights and RIR Policies. <https://ipv4.global/events/leasing-rights-policies/>.
- [29] Is BGP Safe Yet? <https://isbgpsafeyet.com/>.
- [30] M. Jonker, A. Sperotto, R. van Rijswijk-Deij, R. Sadre, and A. Pras. Measuring the Adoption of DDoS Protection Services. *IMC*, 2016.
- [31] M. Kang, W. Li, R. van Rijswijk-Deij, T. T. Kwon, and T. Chung. IRRedicator: Pruning IRR with RPKI-Valid BGP Insights. *NDSS*, 2024.

[32] T. Krenc and A. Feldmann. BGP Prefix Delegations: A Deep Dive. *IMC*, 2016.

[33] I. Livadariu, A. Elmokashfi, and A. Dhamdhere. On IPv4 transfer markets: Analyzing reported transfers and inferring transfers in the wild. *Computer Communications*, 111, Elsevier, 2017.

[34] I. Livadariu, A. Elmokashfi, A. Dhamdhere, and K. Claffy. A first look at IPv4 transfer markets. *CoNEXT*, 2013.

[35] M. Lepinski and S. Kent. An Infrastructure to Support Secure Internet Routing. RFC 6480, IETF, 2012.

[36] M. Luckie, B. Huffaker, K. Claffy, A. Dhamdhere, and V. Gotsas. AS Relationships, Customer Cones, and Validation. *IMC*, 2013.

[37] W. Li, Z. Lin, M. I. A. Khan, E. Aben, R. Fontugne, A. Phoeker, and T. Chung. RoVista: Measuring and Understanding the Route Origin Validation (ROV) in RPKI. *IMC*, 2023.

[38] G. Moura, C. Hesselman, G. Schapman, N. Boerman, and O. de Weerd. Into the DDoS maelstrom: a longitudinal study of a scrubbing service. *IEEE European Symposium on Security and Privacy Workshops*, 2020.

[39] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein. BGP Prefix Origin Validation. RFC 6811, IETF, 2013.

[40] R. Morillo, J. Furuness, C. Morris, J. Breslin, A. Herzberg, and B. Wang. ROV++: Improved Deployable Defense against BGP Hijacking. *NDSS*, 2021.

[41] S. McQuistin, S. P. Uppu, and M. Flores. Taming Anycast in the Wild Internet. *IMC*, 2019.

[42] Magic Transit. <https://www.cloudflare.com/network-services/products/magic-transit/>.

[43] North American Network Operators' Group. <https://www.nanog.org/>.

[44] C. Orsini, A. King, D. Giordano, V. Gotsas, and A. Dainotti. BGPStream: A Software Framework for Live and Historical BGP Data Analysis. *IMC*, 2016.

[45] L. Prehn, F. Lichtblau, and A. Feldmann. When wells run dry: the 2020 IPv4 address market. *CoNEXT*, 2020.

[46] L. Qin, D. Li, R. Li, and K. Wang. Themis: Accelerating the detection of route origin hijacking by distinguishing legitimate and illegitimate MOAS. *USENIX Security*, 2022.

[47] A. Reuter, R. Bush, I. Cunha, E. Katz-Bassett, T. C. Schmidt, and M. Wählisch. Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering. *CCR*, 48(1), 2018.

[48] A. Rogers. Leasing vs purchasing IPv4: a comparison. <https://www.prefixbroker.com/news/lease-vs-purchase-ipv4>.

[49] N. Rodday, I. Cunha, R. Bush, E. Katz-Bassett, G. D. Rodosek, T. C. Schmidt, and M. Wählisch. The Resource Public Key Infrastructure (RPKI): A Survey on Measurements and Future Prospects. *TMA*, IEEE, 2019.

[50] N. Rodday, I. Cunha, R. Bush, E. Katz-Bassett, G. D. Rodosek, T. C. Schmidt, and M. Wählisch. Revisiting RPKI Route Origin Validation on the Data Plane. *TMA*, 2021.

[51] RIPE Routing Information Service (RIS). <http://www.ripe.net/projects/ris/rawdata.html>.

[52] RPKI Deployment Monitor. <https://rpki-monitor.antd.nist.gov>.

[53] University of Oregon RouteViews project. <http://www.routeviews.org/>.

[54] Routinator. <https://nlnetlabs.nl/projects/rpki/routinator/>.

[55] K. Sriram, I. Lubashev, and D. Montgomery. Source Address Validation Using BGP UPDATEs, ASPA, and ROA (BAR-SAV). IETF, 2024. <https://datatracker.ietf.org/doc/draft-ietf-sidrops-bar-sav/>.

[56] C. Testart, P. Richter, A. King, A. Dainotti, and D. C. Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. *IMC*, 2019.

[57] H. Tomas, H. Amir, S. Haya, and W. Michael. Practical experience: Methodologies for measuring route origin validation. *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2018.

[58] The Ultimate Guide to IPv4 Lease for IP Lessees. <https://www.ipxo.com/blog/ipv4-lease-guide-for-ip-lessees/>.

[59] W. Xu, C. Deliang, and L. Xing. On the classification and false alarm of invalid prefixes in RPKI based BGP route origin validation. *IFIP/IEEE Symposium on Integrated Network and Service Management*, 2019.

[60] M. Zhou, X. Zhang, S. Hao, X. Yang, J. Zheng, G. Chen, and W. Dou. Regional IP Anycast: Deployments, Performance, and Potentials. *SIGCOMM*, 2023.

APPENDIX

A. Reproducibility

We published all code and scripts for measurements and experiments at

<https://roa-misconfig.netsecurelab.org>

for network operators, administrators, and researchers to reproduce our work. Our measurements do not involve any private datasets.

All measurements and analyses can be reproduced using the latest public available datasets, including BGP datasets, RPKI datasets and IRR datasets. Data-plane measurements can be reproduced with the public available platform, RIPE Atlas, using the code we provided.

B. Survey

We conducted a survey with eight ISPs and eight leasing brokers during May 2024 to August 2024. Four of them are interviewed by online meeting, and the remaining four ISPs are communicated through email surveys. Questions to ISPs related with ROA misconfigurations are listed below:

- 1) Do you provide IP transit services to customers using your own prefixes, your customers' onboarded prefixes, or both?
- 2) Do you lease IP prefixes to other customers without providing IP transit services?
- 3) What are your procedures for configuring ROA when offering IP transit versus when leasing prefixes? Do you require transit customers to configure ROA correctly before you enable services for their prefixes?
- 4) Do you monitor for RPKI-invalid prefixes that are announced through your network? If so, what tools or processes do you use, and how frequently do you review this?
- 5) For the following list of RPKI-invalid prefixes observed in your announcements, could you explain the cause (e.g., misconfiguration, outdated ROA, etc.)?
[List of RPKI-Invalid prefixes found originated from each participant in our BGP dataset]
- 6) We observed certain hijack events associated with your network based on public detection systems. Please review

this list and indicate whether each event was due to misconfiguration or a genuine hijack.

[List of hijack events found in §VII related to each participant]

7) Are there scenarios beyond IP transit or prefix leasing that could lead to RPKI-invalid statuses in your network? If so, please describe these scenarios.

Questions for IP leasing brokers related to ROA misconfigurations are listed below:

- 1) When leasing prefixes to a lessee, do you provide ROA configuration services for the leased prefixes?
- 2) How do you configure ROA for leased prefixes? What information do you require from the lessee to set up ROA?
- 3) Who typically has the authority to configure or update ROA for the leased prefixes? (e.g., broker, lessee, upstream provider)
- 4) How long does the ROA configuration process typically take after the prefixes are leased?
- 5) Do you monitor the RPKI status of leased prefixes? If yes, how do you monitor it?