# Crack in the Armor: Underlying Infrastructure Threats to RPKI Publication Point Reachability

Yunhao Liu*†, Jessie Hui Wang*†✉, Yuedong Xu‡, Zongpeng Li*, Yangyang Wang*†, Jilong Wang*†
*INSC, BNRist, Tsinghua University †Zhongguancun Laboratory ‡Fudan University
lyh22@mails.tsinghua.edu.cn, {jessiewang, zongpeng}@tsinghua.edu.cn, ydxu@fudan.edu.cn, {wangyy, wjl}@cernet.edu.cn

***Abstract*—The effectiveness of the RPKI in preventing BGP prefix hijacking relies not only on the presence of valid ROAs but also on the successful retrieval of ROAs from publication points (PPs) by relying parties (RPs). Guaranteeing the integrity of data and uninterrupted connectivity during this retrieval process necessitates the proper implementation of security measures in the underlying infrastructure, *i.e.*, the DNS and routing infrastructures.**

**In this paper, we collect information on the specific DNS and routing infrastructures used during the information retrieval process and analyze the infrastructure threats to the reachability of RPKI PPs. Regarding the DNS infrastructure, we report that 31 PPs (48.4%) are susceptible to DNS spoofing attacks and pinpoint the reasons for the appearance of DNSSEC-unprotected zones, such as CNAME redirections to unprotected zones and NS delegations to third-party insecure DNS servers. Regarding the routing infrastructure for communicating with nameservers, our analysis shows that a significant 55 PPs (85.9%) have at least one ROA-unprotected nameserver on their resolution paths, and highlights that the absence of ROA registration for gTLD nameservers accounts for vulnerabilities in 44 of these 55 PPs. Regarding the routing infrastructure for RP-PP communications, we report that 5 PPs fail to register ROAs for the IP addresses of their PP servers. Simulations of routing hijack attacks show that, in the case of the most vulnerable PP, up to 65% to 83% of ASes may experience a loss of connectivity to the PP.**

**Furthermore, we investigate the deterministic and probabilistic dependencies among publication points and uncover a critical issue: some RIR-operated PPs rely on less secure lower-level PPs, which can significantly amplify the impact of vulnerabilities within insecure PPs, potentially leading to cascading failures.**

## I. Introduction

The Border Gateway Protocol (BGP) is the de facto inter-domain routing protocol that connects Autonomous Systems (ASes) across the global Internet. Despite its critical role in propagating IP prefix reachability, BGP lacks built-in authentication and authorization, making it vulnerable to prefix hijack attacks [1], [2], [3], [4]. In a prefix hijack, a malicious AS illegitimately announces ownership of IP prefixes it does not control. Other ASes that accept this forged announcement may then misroute traffic, redirecting data intended for the legitimate prefix holder to the attacker instead.

To address the lack of origin authentication in BGP, the Internet Engineering Task Force (IETF) standardized the *Resource Public Key Infrastructure* (RPKI) [5]. RPKI allows IP address holders to create Route Origin Authorizations (ROAs), which are cryptographically signed objects that bind an IP prefix to the ASes authorized to originate it. These signed objects are published in a distributed repository which includes tens of *publication points* (PPs). To utilize these objects, an AS must run *relying party* (RP) software that periodically retrieves and validates ROAs from these PPs. The validated information is then supplied to routers, which perform Route Origin Validation (ROV) by checking incoming BGP announcements against the ROA data. Using ROV, ASes can filter out routes with invalid origins, effectively blocking prefix hijack attacks. Beyond its role in preventing hijacks, RPKI also serves as the foundational trust infrastructure for other routing security mechanisms such as BGPsec [6] and ASPA [7].

The deployment of RPKI, encompassing both ROA creation and ROV enforcement, has grown steadily in recent years. As of July 2025, ROAs cover approximately 58% of BGP-announced IPv4 prefixes and 63% of IPv6 prefixes [8], and 28.9% of ASes have adopted ROV to filter out routes with invalid origins [9].

Existing studies on RPKI effectiveness typically focus on the deployment ratios of ROAs and ROV. However, even if a prefix owner (e.g., for prefix $p$) has registered an ROA and a BGP route receiver (e.g., $AS_a$) has implemented ROV, protection against prefix hijack attacks on traffic flows from $AS_a$ to $p$ is not fully assured. This is because the process for the receiver's RP software to retrieve the ROA is a complex, multi-step process. For the ROA to be usable, *data integrity* and *end-to-end connectivity* of all communication data flows involved in this *retrieval process* must be ensured.

Specifically speaking, the process involves two key sub-processes. The first sub-process is that the RP resolves the domain name of each individual PP through the DNS infrastructure to obtain the PP's IP address. This *resolution sub-process* involves iterative interactions with authoritative name servers across varying levels of the DNS hierarchy. The second sub-process involves the RP establishing communication with the PP via its IP address to retrieve ROAs using either RRDP or rsync, which we refer to as the *download sub-process*. To ensure data integrity during the resolution sub-process,

✉ Jessie Hui Wang is the corresponding author.

Network and Distributed System Security (NDSS) Symposium 2026
23-27 February 2026, San Diego, CA, USA
ISBN 979-8-9919276-8-0
https://dx.doi.org/10.14722/ndss.2026.241141
www.ndss-symposium.org

all involved entities must adopt DNSSEC, which provides cryptographic validation of DNS responses. In contrast, the download sub-process ensures data integrity through RPKI's built-in security mechanisms. Notably, the download sub-process benefits from strong cryptographic protections inherent in RPKI, whereas the resolution sub-process remains vulnerable due to partial DNSSEC deployment on the Internet. To ensure end-to-end connectivity, routing security for all data flows within both sub-processes must be meticulously upheld. Regrettably, despite several routing security mechanisms proposed in academia and industry, RPKI stands out as the sole well-developed and practically viable solution, albeit limited in its capability to solely address prefix hijacks. Similar to DNSSEC, RPKI faces partial adoption within the current Internet landscape.

In summary, *the reliable delivery of RPKI data from PPs to RPs is reliant on the security of two underlying infrastructures: the DNS and the routing infrastructures*. This requires the adoption of DNSSEC for DNS security and the deployment of RPKI itself to prevent prefix hijacks on the routing infrastructure. Therefore, in this paper, we conduct a measurement study to collect information on the specific DNS and routing infrastructures involved in RP-PP communication. We then analyze the deployment status of their respective security solutions, *i.e.*, DNSSEC for DNS infrastructure and RPKI for routing security, to assess the infrastructure threats to the reachability of RPKI PPs. We also pinpoint that the structural inter-PP dependencies can significantly amplify the risks of these infrastructure threats, as well as other potential threats to PP reachability.

Our contributions can be summarized as follows.

- We measure DNSSEC deployment status across the DNS resolution paths of RPKI PPs. Among the 64 PPs discovered in this study, 31 PPs (48.4%) are vulnerable to DNS spoofing attacks in resolving their PP domain names due to incomplete DNSSEC protection. A further analysis of these 31 vulnerable PPs shows that 7 PPs suffer from PP-controlled DNSSec-unprotected zones, usually resolvable with minimal coordination, 22 PPs suffer from third-party-hosted unprotected zones on their resolution paths, highlighting operational risks of NS delegations to third-party DNS providers without DNSSEC, and 2 PPs suffer from CNAME redirections to unprotected zones outside their control. Additionally, RIPE NCC's RRDP domain and rsync domain names experience inconsistent DNSSEC protection.
- We evaluate the ROA coverage and ROV deployment for networks that accommodate the authoritative nameservers used in resolving PP names. A significant 55 PPs (85.9%) feature at least one nameserver on their resolution paths lacking ROA coverage. This discovery prompts a deeper investigation, uncovering that out of the 13 authoritative gTLD nameservers, only three (`e`, `i`, and `m`) have registered ROAs. The absence of ROA registration for this crucial zone contributes to 44 out of the 55 vulnerable PPs. Furthermore, local resolvers capable of exploring alternative resolution paths in the event of failures can markedly diminish the risks of disconnections stemming from routing hijack attacks on a ROA-unprotected nameserver.
- We evaluate the ROA coverage and ROV deployment for networks that accommodate RPKI PPs. Out of the 60 PPs resolved to a single IP address, 4 PPs have their IP addresses unprotected by ROA. Among the 4 PPs resolved to varying IP addresses depending on the user's geographic location, 1 PP has some of its IP addresses unprotected by ROA.
- We evaluate the practical risk of the identified routing hijack threats. Simulations show that the effectiveness of attacks varies significantly among these 10 vulnerable PPs. In the case of the most vulnerable PP, up to 65% to 83% of ASes may experience a loss of connectivity to the PP.
- We investigate the deterministic and probabilistic dependencies among publication points and uncover a critical issue: some RIR-operated PPs rely on less secure lower-level PPs. This situation poses a significant risk as vulnerabilities in these insecure PPs can be greatly magnified through inter-PP dependencies, potentially leading to cascading failures. Specifically, DNS spoofing attacks and low-rate attacks could compromise non-RIR PPs on which the RIPE PP probabilistically relies, triggering cascading disruptions that affect 39 and 42 PPs. This could widen the impact, expanding the reduction of the size of ROA-protected IP addresses from 14.58% and 6.62% to 50.89% and 42.92%, respectively.

In addition to the aforementioned findings, our analysis highlights the risks associated with deploying publication points in CDNs and leveraging NS delegations to external domains. To foster reproducibility, we publish our measurement and analysis code as well as the anonymized results at [10].

## II. Related Work

**Deployment of RPKI.** The deployment of RPKI has steadily increased as more ASes create ROAs and adopt ROV to filter invalid BGP announcements. Longitudinal measurements over the past decade have shown a consistent rise in both ROA registration [11], [12], [13], [14] and ROV enforcement [15], [16], [17], [9]. Despite this growing adoption, subprefix hijacks are still possible under partial ROV deployment [18], and subsequent works have proposed enhancements to ROV to mitigate such attacks [19], [20].

**RPKI security.** Researchers have identified various attacks targeting the RPKI system, which are summarized in Table I.

The first category targets RP software vulnerabilities that may cause crashes. An early analysis [21] showed that many RP implementations were susceptible to both generic attacks, such as gzip bombs, XML billion laughs, and path traversals, and RPKI-specific bugs that could crash the software. Later, [22] reported two zero-day vulnerabilities in the most popular RP implementations, which could trigger crashes even with slightly malformed objects or oversized CA subtrees, affecting

84.9% of global RPs. To systematically analyze these bugs, researchers developed CURE [23], a general-purpose, language-agnostic RP fuzzer, which discovered 18 new vulnerabilities.

The second category exploits malicious RPKI authorities to launch attacks. RPKI is vulnerable to resource subversion and manipulation [24], because any high-level resource owner can unilaterally delegate or revoke resources without verification, affecting arbitrary child or grandchild PPs.

The third category targets RP-PP communication, preventing RPs from fetching RPKI data from PPs. [25] first used a delegated PP and an RP to map RPKI deployments, discovering RP addresses and query frequencies. They showed that an attacker could trigger rate limiting on PP servers and nameservers by sending spoofed packets, causing legitimate RP queries to be dropped, and could further prolong validation rounds using connection slowdown techniques or unlimited PP delegations until manifests expire. [26] identified several DNS-related vulnerabilities in RPKI, including reduced redundancy due to multiple RPs sharing the same resolver, recursive resolvers lacking proper DNSSEC validation, and DNS components hosted on prefixes without ROA coverage that are exposed to BGP hijacks. [27] showed that despite community patches, stalling attacks that delay or block validation remain viable in multiple forms. In addition to these attacks, a recent study [28] provided a comprehensive overview of RPKI security, showing that 56% of global validators suffer from at least one documented vulnerability. It also reported persistent issues in repository management, including frequent RRDP/rsync failures and malformed content, highlighting the fragility of current RPKI operations.

While [26] examined vulnerabilities at RP-side local resolvers, our study systematically evaluated the security of PPs across the entire RP–PP communication process. In the resolution sub-process, we analyzed DNSSEC deployment along resolution paths to identify PPs that remain vulnerable even when RPs use secure local resolvers, and investigated the causes of these vulnerabilities. We also evaluated end-to-end connectivity under potential prefix hijacks targeting authoritative nameservers, showing that the likelihood of such attacks is lower than previously reported in [26]. In the download sub-process, we analyzed ROA coverage using data from multiple resolvers worldwide and assessed ROV deployment for PPs, which allowed us to understand how network configurations and CDN deployment impact RPKI data availability. Finally, to the best of our knowledge, we are the first to analyze inter-PP dependencies and demonstrate how they can amplify attack impact, potentially leading to cascading failures.

## III. Background

### A. Overview of RPKI

RPKI provides a cryptographically verifiable binding between IP prefixes and their authorized origin ASes, enabling route origin validation in interdomain routing. It consists of hierarchically organized repositories, known as publication points, and relying party software that fetches, parses, and validates RPKI objects.

| Study | Year | Attack Target | Attack Vector |
|---|---|---|---|
| [22] | 2022 | RP software | Exploiting software bugs to crash or stall RPs |
| [21] | 2023 | | |
| [23] | 2024 | | |
| [24] | 2013 | Child PPs | Abuse of PP delegation and revocation mechanisms |
| [25] | 2022 | RP-PP communication | DoS attacks on PPs and nameservers<br>Stalling attacks via unlimited PP delegations |
| [26] | 2022 | | DNS hijacking targeting DNS resolvers |
| [27] | 2023 | | Stalling and delta-snapshot attacks on RP validation |
| Our work | 2025 | | DNS spoofing and prefix hijacking targeting RP–PP communication |

TABLE I: A summary of attacks targeting the RPKI system.

**RPKI Publication Points.** The RPKI hierarchy is anchored by five Regional Internet Registries (RIRs), each serving as a trust anchor that certifies its allocated IP resources. These RIRs, as well as other entities that hold resource certificates, such as National Internet Registries (NIRs) and ISPs, can act as Certificate Authorities (CAs). A CA issues subordinate certificates to delegate portions of its resource holdings and allows prefix holders to create ROAs, which authorize specific ASes to originate the associated prefixes. Each CA publishes its signed objects at a PP, whose URL is specified in the Subject Information Access (SIA) extension of its certificate. Depending on how the publication infrastructure is managed, a PP may operate in either a hosted mode, where the RIR provides publication services on behalf of the certificate holder, or a delegated mode, where the certificate holder maintains their own repository infrastructure. Each PP stores all signed objects issued by its CA, including ROAs, subordinate certificates, Certificate Revocation Lists (CRLs), and a manifest that enumerates all valid files.

**Relying Party Software.** Networks that enforce ROV typically operate an RP that periodically performs validation sessions starting from Trust Anchor Locators (TALs) [29]. Each TAL specifies the URL and the public key used to authenticate the trust anchor certificate maintained by an RIR. The RP authenticates each trust anchor and extracts its SIA to locate the associated repository. It then recursively traverses the certificate delegation chain by visiting all reachable PPs and retrieves their published objects using either the RRDP [30] or the rsync protocol. For each retrieved object, the RP verifies the signatures and checks certificate validity using the issuer's public key. Once the entire validation process is complete, the RP compiles a set of Validated ROA Payloads (VRPs), each representing an authorized IP prefix and its corresponding origin AS. These VRPs are then provided to BGP routers via the RPKI-to-Router (RTR) protocol [31], where they are used to validate route announcements received via BGP.

**Interaction between Components.** Figure 1 illustrates the interactions among components involved in RPKI validation. In Step (1), the RP software uses a DNS resolver to look up the addresses of hardcoded TALs to locate the PPs of the five RIRs. The DNS resolver then sends a lookup request to the authoritative nameserver responsible for the corresponding RIR domain to obtain the IP address of the PP (Step 2). Once the IP address is received, the RP downloads the repository data using RRDP or rsync (Step 3). This repository includes

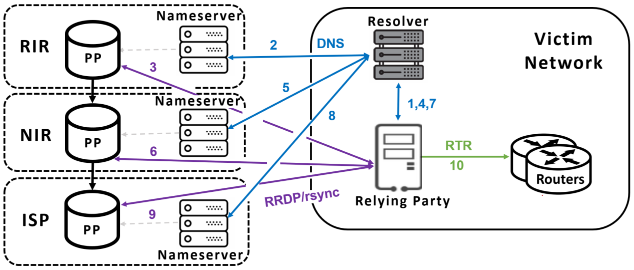

Fig. 1: Interaction between components in RPKI process.

signed metadata delegating to subordinate PPs, such as those operated by NIRs. Following this delegation, the RP extracts the NIR PP's domain name from the RIR's metadata, resolves it similarly (Steps 4-5), and downloads the corresponding repository data (Step 6). The same procedure applies recursively to the ISP-operated PP (Steps 7-9), traversing the hierarchical RPKI certificate chain. After all reachable PPs have been contacted and their RPKI objects validated, the RP compiles a complete set of VRPs and delivers them to BGP routers via the RTR protocol (Step 10). The correctness and timeliness of this process depend critically on the availability and integrity of the DNS resolution and download sub-process at each stage.

### B. Infrastructure Threats to RPKI PP Reachability

This paper focuses on the critical process by which an RP software retrieves ROAs from PPs. This process is composed of two sub-processes, *i.e.*, DNS resolution and data download, and its success depends on the data integrity and end-to-end connectivity of all involved communication flows. Given that the data integrity of the download sub-process is already protected by X.509 certificates and manifest files, our analysis zeroes in on three types of threats to the remaining parts of this process.

**Data integrity of the DNS resolution sub-process**. DNS responses from authoritative nameservers are vulnerable to spoofing unless DNSSEC [32] is implemented properly. Illustrated in Figure 2, Attacker A injects a fake DNS response to deceive the local DNS resolver, redirecting the RP to a malicious server instead of the PP. These DNS spoofing maneuvers are typically carried out via man-in-the-middle interception or DNS cache poisoning. Protection from this type of attack is only possible if all authoritative nameservers participating in the iterative resolution of the PP's domain name have adopted DNSSEC.

**End-to-end connectivity of the DNS resolution sub-process**. A disruption in end-to-end connectivity during the DNS resolution of a PP domain name will cause the process to fail, impeding the RP from obtaining the PP's IP address and then undermining PP reachability. Therefore, robust routing security between the local resolver and each nameserver is essential. One necessary (though not sufficient) condition for this is that the participating nameservers must register ROAs for their IP addresses. An ROV-enabled local resolver can use these ROAs to ensure connectivity to the nameservers. Furthermore, the nameservers must be hosted in ROV-enabled networks. This allows them to filter out counterfeit BGP announcements and protect the return path to the local resolver (assuming the resolver's IP is covered by a registered ROA). Our analysis assumes that the local resolver is properly secured, as its security should be addressed locally. This assumption is also reasonable given that the best current operational practices [33] recommend positioning RPs near BGP routers to reduce hijacking risks, which implies that the RP's AS deploys ROV to enforce routing security. Additionally, prior studies [25] show that RPs commonly use DNS resolvers within their own AS or from major providers such as Google and Cloudflare, both of which have deployed ROV.

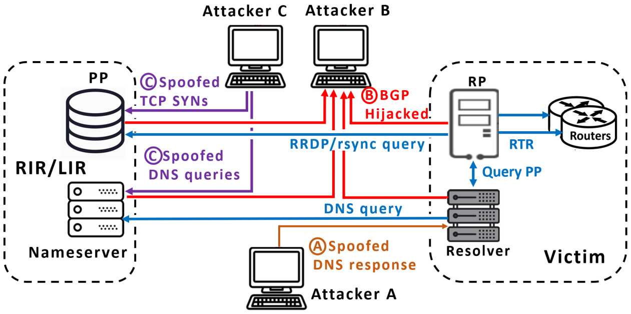

Fig. 2: Attacks exploiting security weaknesses of underlying RPKI infrastructure.

**End-to-end connectivity of the download sub-process**. Similar to the above end-to-end connectivity concern, the end-to-end connectivity between RPs and PPs can be disrupted if the routing between them is hijacked. As shown in Figure 2, the attacker $B$ launches a prefix hijacking attack to redirect traffic flows to or from a PP, undermining PP reachability. A necessary yet insufficient measure to safeguard the routing security between PPs and RPs is for a PP to register ROAs for its IP addresses and to operate within ROV-enabled networks. We do not study the deployment status of security mechanisms on RPs, as RP security issues should be addressed locally, which are beyond the scope of this paper.

Besides the above threats, additional threats exist within the RPKI system. An adversary could disrupt the functionality of nameservers or PPs directly through *low-rate attacks*. As depicted in Figure 2, attacker $C$ initiates low-rate attacks by sending a large number of spoofed queries to trigger rate limiting on nameservers or PPs. This form of attack has been extensively studied in previous research [25]. While this paper does not delve into this specific threat, it is considered within our evaluation of the cascading risks amplified by inter-PP dependencies in Section V. Inter-PP dependencies emerge when PPs utilize nameservers or IP prefixes protected by ROAs issued by other PPs, thereby introducing the potential for cascading failures.

## IV. Analyzing Infrastructure Threats to RPKI PP Reachability

In this section, we present a measurement-driven analysis of the underlying infrastructure employed by RPs to retrieve RPKI data from PPs. We first measure the deployment status of DNSSEC and RPKI on PP servers and all relevant authoritative nameservers. Based on it, we report the diverse threats to the reachability of PPs throughout the DNS resolution and data download subprocesses. Finally, we evaluate the effectiveness of attacks that exploit these threats.

### *A. Measurement of Security Mechanism Deployment*

We measure the deployment of two key security mechanisms, DNSSEC for DNS infrastructure and RPKI-based ROA and ROV for routing security, in the underlying infrastructure supporting RPKI data retrieval. The following describes our measurement methodology.

(1) **Extract PP domain names.** Using a snapshot of the global RPKI repository provided by RIPE NCC [34], dated June 1, 2025, we extract the repository URLs from the SIA fields of all resource certificates and derive their corresponding domain names. This yields 64 distinct PP domain names, with duplicates removed across certificates. Among them, 17 PPs use different domain names for rsync and RRDP. We focus on the RRDP domain names in the following analysis, while the rsync domain names are discussed separately in Section IV-B2.

(2) **Construct DNS resolution paths.** For each PP domain name, we construct its DNS resolution path by performing iterative DNS queries starting from the DNS root zone. We define a *resolution path* as the ordered sequence of DNS zones traversed during the resolution process, where each parent zone delegates authority to a child zone via NS records.

A zone along this path may have multiple authoritative nameservers that can respond to queries. While the resolution path itself refers to the zone-level hierarchy, the presence of multiple nameservers in each zone gives rise to different *nameserver-level resolution paths*, where queries may follow different routes depending on which nameservers are selected at each step. We extract and analyze all nameserver-level resolution paths for each PP.

An important operational detail involves glue records, which are address (A/AAAA) records provided in a parent zone to facilitate the resolution of nameservers whose hostnames lie within the delegated child zone. If glue records are absent, resolvers need to perform additional queries starting from the root zone to resolve these nameserver hostnames, increasing both the complexity and length of the resolution process.

(3) **Measure DNSSEC deployment.** A zone is considered *DNSSec-protected* only when both of the following conditions are satisfied: (1) its parent zone publishes a Delegation Signer (DS) record that matches a DNSKEY record in the zone, ensuring continuity of the DNSSEC trust chain; and (2) its authoritative nameservers correctly serve DNSSEC-related records, including valid RRSIG signatures and NSEC or NSEC3 records for authenticated denial of existence.

Based on the DNSSEC status of individual zones, we determine whether a PP domain name is DNSSEC-protected. A domain name is considered as *DNSSEC-protected* when a complete and valid signature chain exists from the root zone to its authoritative zone, *i.e.*, all zones on the resolution path are DNSSEC-protected.

(4) **Measure ROA coverage.** We first identify the IP addresses of PPs and nameservers on their resolution paths. Considering potential geographic variation in DNS resolution, we perform queries using a globally distributed set of public resolvers obtained from [35].

Although the resolver list spanned 193 countries and regions, we received responses from only approximately 15,000 resolvers across 153 of them. Using these responsive resolvers, we queried the domain names of all 660 identified nameservers and 64 PPs to get their corresponding IP addresses.

Only 42 (5.67%) of these domain names are resolved to IP addresses that vary by geographic location, indicating the majority of domain names yield consistent resolutions globally. This uniform resolution pattern mirrors conventional authoritative DNS setups, where anycast servers are commonly employed to furnish consistent responses on a global scale. Conversely, domain names showing geographic disparities in resolution are probably supported by Content Delivery Networks (CDNs) or other geo-distributed infrastructures. These systems optimize performance by steering clients to different addresses according to their locations.

Using BGP routing table snapshots from RouteViews [36], we match each resolved IP address with the longest matching prefix from the BGP tables and also determine the originating AS of this prefix. Subsequently, we verify whether this prefix is covered by a valid ROA that authorizes the observed origin AS. Upon successful verification, the nameserver or PP is considered ROA-protected.

### *B. Data Integrity Threats in DNS Resolution Subprocess*

In this subsection, we first summarize the characteristics of DNS resolution paths of PPs. Then we introduce the DNSSEC deployment status on the nameservers along the resolution paths of each PP and identify PPs vulnerable to DNS spoofing attacks.

*1) Characteristics of PPs' Resolution Paths:* We analyze PPs' resolution paths and report the following observations.

**Top-Level Domain Distribution.** We examine the top-level domains (TLDs) of the PP domain names. As shown in Figure 3, 29 out of 64 PPs use either the `.com` or `.net` TLD, accounting for approximately 45% of the total. Both TLDs are operated by Verisign, which manages a unified set of gTLD root servers, i.e., `[a-m].gtld-servers.net`. Consequently, nearly half of the PPs depend on Verisign's infrastructure for their name resolution. In addition, we also observe that 28% of PPs use country-code TLDs (ccTLDs). We further categorize these ccTLDs by continent and find that their distribution is geographically balanced across regions. This suggests that operators and ISPs in diverse parts of the world actively adopt the delegated mode to operate their own

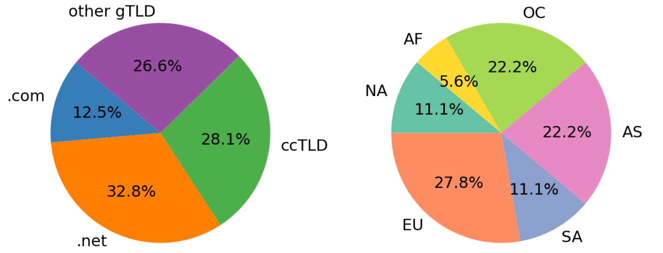

Fig. 3: Distribution of TLD categories and ccTLD continents for PPs.

PPs, contributing to the global decentralization of the RPKI infrastructure.

**CDN Deployment.** Among the 42 domain names exhibiting geographic variation in DNS resolution, 4 are PP domain names. Furthermore, we identify 10 PPs hosted on CDNs using the CDN identification methodology from dRR [37], which utilizes the presence of CDN provider identifiers in CNAME records and HTTPS response headers.

Among the 4 PPs exhibiting geographic variation, 3 are confirmed to use CDN services. The remaining PP, operated by LACNIC, shows geographic variation without association to any known CDN provider, suggesting it employs custom geo-distribution strategies. Table II summarizes our findings. Cloudflare is the most commonly used CDN, supporting 8 PPs, including those operated by AFRINIC and APNIC. RIPE NCC relies on Akamai, while Amazon utilizes its proprietary CloudFront service. These providers implement distinct infrastructure models: Cloudflare employs globally consistent anycast IP addresses, resulting in uniform DNS resolution across locations. In contrast, Akamai and CloudFront use localized edge servers that produce region-specific resolution behaviors. This can potentially lead to inconsistencies in availability and increase vulnerability to routing-based attacks in certain areas.

| RRDP Domain (CNAME) | Geo-IP | ASN | Header |
|---|---|---|---|
| magellan.ipxo.com | Yes | 13335 | Cloudflare |
| krill.stonham.uk | No | 13335 | Cloudflare |
| rrdp.apnic.net<br>(rrdp.apnic.net.cdn.cloudflare.net) | No | 13335 | Cloudflare |
| rrdp.sub.apnic.net<br>(rrdp.sub.apnic.net.cdn.cloudflare.net) | No | 13335 | Cloudflare |
| rrdp.afrinic.net<br>(rrdp.afrinic.net.cdn.cloudflare.net) | No | 13335 | Cloudflare |
| rpki-rrdp.mnihyc.com<br>(rpki-rrdp.mnihyc.com.cdn.cloudflare.net) | No | 13335 | Cloudflare |
| rpki-publication.haruue.net<br>(rpki-publication.haruue.net.cdn.cloudflare.net) | No | 13335 | Cloudflare |
| rrdp.paas.rpki.ripe.net<br>(rrdp.paas.rpki.ripe.net.cdn.cloudflare.net) | No | 13335 | Cloudflare |
| rrdp.ripe.net<br>(rrdp.ripe.net.akamaized.net)<br>(a1513.dscd.akamai.net) | Yes | 20940 | Akamai |
| rpki-rrdp.us-east-2.amazonaws.com<br>(d27hng7mrok6ld.cloudfront.net) | Yes | 16509 | Amazon |

TABLE II: CDN-hosted PPs. *Geo-IP* indicates whether the domain resolves to different IP addresses based on the resolver's geographic location. *Header* presents the CDN provider derived from HTTPS response headers.

**Resolution Path Lengths.** The reliability of resolving PP domain names depends significantly on the security of the nameservers along the resolution path. Longer resolution paths involve more intermediate nameservers, each increasing the chance of interception or interference during the resolution process. We report the distribution of resolution path lengths in Figure 4, excluding zero-occurrence lengths for clarity. For non-CDN PPs, resolution paths are generally shorter, typically ranging from 3 to 6. In contrast, CDN-backed PPs tend to have longer resolution paths, often between 6 and 7, with some extending beyond 10. This additional length mainly results from CNAME-based redirections, which introduce more intermediate nodes and create cross-organizational dependencies. Longer resolution paths also appear in non-CDN cases due to out-of-zone NS delegations.

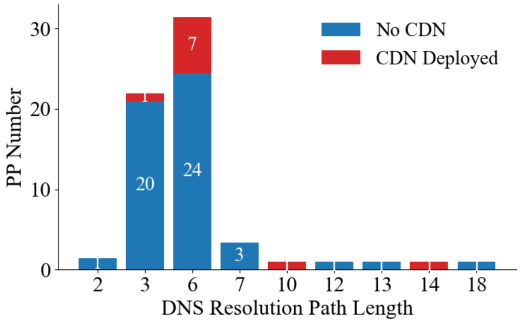

Fig. 4: Distribution of DNS resolution path lengths for PPs with and without CDN deployment.

*2) DNS Spoofing Threats Analysis:* Spoofed DNS responses can redirect RPs to malicious servers rather than PP servers, disrupting the reachability of PPs from RPs. Based on our analysis of DNSSEC deployment across resolution paths, we pinpoint unprotected zones and assess the associated threat landscape.

**DNSSEC Deployment Along Resolution Paths.** Figure 5 shows the distribution of PPs categorized by the number of non-DNSSEC-protected zones along their resolution paths. We find that 33 PPs (51.6%) have resolution paths with all zones protected by DNSSEC. The remaining *31 PPs each have at least one non-DNSSEC-protected zone on the resolution path*. Among these 31 PPs, 18 (58.1%) PPs have exactly one unprotected zone, which may be comparatively easier to address.

Despite the presence of non-DNSSEC-protected zones, DNSSEC adoption by nameservers used by RPKI PPs is substantially higher than that observed in the general web ecosystem. According to ICANN's June 2025 statistics [38], fewer than 10.5% of domain names in the Majestic Million list, which is commonly used as a replacement for the retired Alexa ranking, are protected by DNSSEC.

**Classification of Unprotected Zones.** We further analyze the operators of the unprotected zones to understand the operation complexity involved in securing them.

We consider an unprotected zone as *PP-controlled* if it meets any of the following conditions: (1) the zone corresponds to the PP domain or its subdomain; (2) the zone's Start

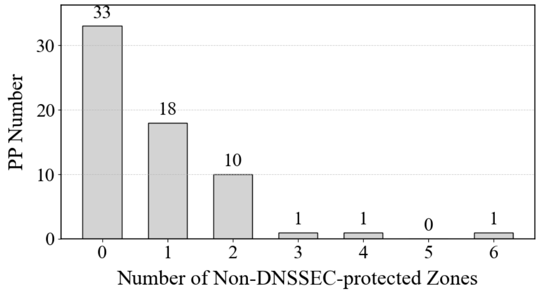

Fig. 5: Distribution of PPs by the number of non-DNSSEC-protected zones in their resolution paths.

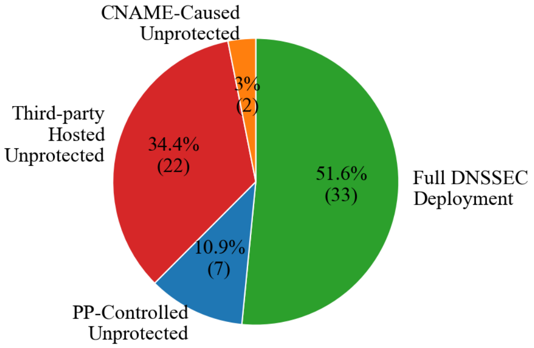

Fig. 6: Classification of PPs by type of non-DNSSEC-protected zones.

of Authority (SOA) record includes an administrative contact email address associated with the PP domain; (3) WHOIS registration information matches the PP operator or affiliated entity. This classification represents a best-effort inference, with the outlined conditions serving as indicators of a close relationship between the PP and the zone's operator. This closeness suggests that the PP operator likely possesses the influence to advocate for DNSSEC deployment within the zone.

Some of the other unprotected zones are clearly hosted on recognized third-party DNS platforms such as Cloudflare or AWS, and they are classified as *third-party hosted*. The third class is *CNAME-caused unprotected* zones. These unprotected zones appear on the resolution path due to CNAME redirections to domains beyond the authoritative DNS hierarchy of the PP domain name.

This threefold classification covers all unprotected zones on all resolution paths of PPs. Accordingly, we classify PPs with unprotected zones into three groups, *i.e.*, *PP-Controlled Unprotected PPs*, *Third-Party Hosted Unprotected PPs*, and *CNAME-Caused Unprotected*.

Figure 6 shows the classification of PPs. *Overall, 31 PPs (48.4%) have unprotected zones in their resolution paths.* Among them, 7 PPs involve PP-controlled unprotected zones. These cases are often due to configuration oversights and can typically be resolved with minimal coordination. 22 PPs have third-party-hosted unprotected zones on their resolution paths, which reflects the operational risks of relying on third-party DNS providers without DNSSEC. Nearly half (45.5%) of these 22 PPs are hosted on Cloudflare, with the remainder spread across providers such as AWS, Freerange Cloud, and DNSPod.

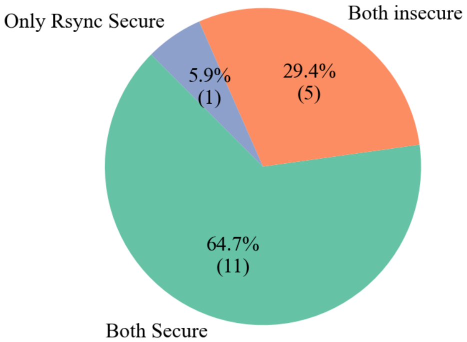

Fig. 7: Comparison of DNSSEC protection status between RRDP and rsync domain names.

In addition, 2 PPs involve CNAME redirections that lead to unprotected zones outside the administrative control of the PP, breaking the DNSSEC validation chain and introducing DNS spoofing vulnerabilities. To mitigate such risks, operators should audit all cross-domain CNAME references and ensure that DNSSEC is properly deployed across the entire resolution path.

**DNSSEC Deployment for rsync Domains.** As noted in the previous subsection, 17 PPs employ different domain names for rsync and RRDP access, typically differing only in the lowest-level subdomain, *e.g.*, `rrdp.arin.net` and `rpki.arin.net`. While these domain names are generally operated by the same organization, they may resolve through different DNS paths and follow separate operational configurations, which can result in divergent security deployments, including DNSSEC.

Given that RP implementations typically fall back to rsync when RRDP retrieval fails, we separately evaluate the DNSSEC deployment of these rsync domains to assess their reliability as fallback channels. As shown in Figure 7, most PPs exhibit consistent DNSSEC protection across both domain types. An exception to this pattern is observed with RIPE NCC: while the resolution path of its RRDP domain name includes CNAME-caused unprotected zones, its rsync domain name enjoys complete protection. This inconsistency hints that RP implementations without fallback support might be vulnerable to DNS spoofing attacks on the RRDP domain name.

### C. Connectivity Threats in DNS Resolution Subprocess

Routing security between the local resolver and each nameserver is essential for the end-to-end connectivity of traffic flows for the DNS resolution process. We assume that the local resolver is properly secured, as its security should be addressed locally by the ASes who use RPKI. We focus on the routing security issues related to the authoritative nameservers, *i.e.*, nameservers must register ROAs for their IP addresses, and they must be hosted in ROV-enabled networks.

*1) ROA Registration by Nameservers used by the RPKI System:* Our analysis shows that 55 PPs (85.9%) have at least one nameserver uncovered by ROAs on their resolution paths. We use the same analysis method as the prior work [26], which reported that 50 PPs among the 53 PPs in 2022 have resolution paths involving at least one nameserver without ROA protection.

The widespread lack of ROA coverage indicates that the majority of PPs are vulnerable to prefix hijacking in the resolution process. Shocked by this result, we further investigate the underlying cause of incomplete ROA coverage along the resolution paths of PPs. As reported in the previous section, 45% of PPs use names whose top-level domains are `.com` or `.net`. These two top-level domains are resolved by a unified set of 13 authoritative gTLD nameservers `[a-m].gtld-servers.net`. We examine the ROA deployment status of these nameservers. Our measurements reveal that only three of these servers, `e`, `i`, and `m`, have registered ROAs. Other nameservers are not protected by ROAs, which is the reason for 44 PPs (68.8%) with ROA-unprotected nameservers on their resolution paths. Excluding the influence of gTLD nameservers leads to a sharp reduction in the fraction of PPs with incomplete ROA protection, dropping from 85.9% (55 PPs) to 17.19% (11 PPs).

Given the critical role of gTLD nameservers, we conduct a further investigation, which reveals that the three gTLD nameservers with ROAs still have problems in preventing prefix hijacks. Verisign registers them only for the /32 IP addresses. However, BGP tables [36] show that only the corresponding /24 prefixes are advertised globally. To verify this, we query 27 Looking Glass servers from the dataset in [39] by running the `show bgp route` command for each of the three nameserver IPs. The results show that only 3 ASes observe the exact /32 prefixes, while the other 24 ASes see only the aggregated /24 prefixes. This indicates that the /32 prefixes are commonly filtered or aggregated during route propagation, rendering the /32 ROAs ineffective in most ASes and leaving these nameservers vulnerable to prefix hijacking. We reported this issue to Verisign in June 2025, after which ROAs are registered for the corresponding /24 prefixes of the `e`, `i`, and `m` servers.

*2) ROV Enforcement of Networks Hosting Nameservers:* We examine whether the networks hosting the authoritative nameservers have deployed ROV.

We utilize the RoVista dataset [9], which reports ROV deployment for about 32,000 ASes, offering the most comprehensive coverage to date. For each nameserver, we identify its hosting AS and check its ROV status in RoVista, which classifies ASes as either *protected*, *unprotected*, or *unknown* if not covered by the dataset. If a hosting AS is not protected, we further check whether any of its upstream providers have deployed ROV, based on CAIDA's AS relationship data [40]. A nameserver's hosting AS is classified as *protected* if it or all of its providers enforce ROV; otherwise, it is considered *unprotected*. To account for the uncertainty introduced by unknown ASes, we evaluate two extreme scenarios: an upper bound assuming all unknown ASes enforce ROV and a lower bound assuming none do.

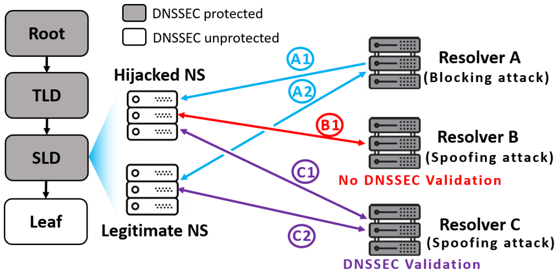

Fig. 8: Prefix hijacking Attacks on DNS resolution Process.

Based on this analysis, under the upper bound all PPs have all their authoritative nameservers on their resolution paths hosted in protected ASes. Under the lower bound, only 43.9% of PPs have all their authoritative nameservers in protected ASes, with 56.1% of PPs having at least one nameserver in a potentially unprotected AS. The actual protection level is expected to fall between these two bounds and is likely closer to the lower bound, as many unknown ASes are smaller networks with limited deployment of routing security mechanisms.

*3) Impact of Resilient Local Resolvers:* The analysis in the preceding subsections of Section IV-C is predicated on the assumption that each PP possesses solely a single nameserver-level resolution path. Consequently, if any nameserver along this path falls victim to prefix hijacking or accepts a spurious route to the local resolver, the resolution of the PP domain name will fail. However, in instances where a PP has multiple nameserver-level resolution paths, the success of attacks on resolving the PP domain name hinges on additional factors, including the capabilities of the local resolver.

During resolution, a resolver queries one of the authoritative nameservers for each zone along the DNS resolution path. If no response is received (due to routing hijack attacks), it typically retries with another nameserver for the same zone to improve the resilience. As shown in Figure 8, the resolver $A$ first queries a hijacked nameserver that silently drops the query, then retries another legitimate nameserver and successfully obtains a valid response. Therefore, to mount a successful routing attack to block the resolution process of one PP domain name, the attacker must hijack the prefixes of all authoritative nameservers for a given zone; otherwise, the resolver can still complete resolution.

According to our analysis, the resolution process of only 2 PPs are vulnerable to prefix hijack attacks, given the presence of resilient local resolvers. In each of its nameserver-level resolution paths, there is at least one nameserver lacking ROA protection.

*Remarks.* Once a prefix hijacking attacker attracts data flows, they can launch a man-in-the-middle attack by manipulating the data packets before forwarding them to the legitimate receiver. Therefore, routing hijacking is not only a means to disrupt connectivity but also a sophisticated method for data manipulation and eavesdropping.

We examine the effectiveness of prefix hijacking attacks for DNS spoofing. If the domain name of one PP is DNSSEC-protected, any spoofing attacks can be detected. Previously, we reported that 31 PPs each have at least one non-DNSSEC-protected zone on the resolution path. Among them, we observe that 6 PPs have resolution paths containing at least one unprotected zone with at least one nameserver missing ROA protection. These 6 PPs can be attacked by hijacking the insecure nameserver and injecting falsified DNS responses.

*4) Vulnerability of Nameservers to Subprefix Hijacking:* Even if all nameservers involved in the DNS resolution subprocess have registered ROAs, they can still be vulnerable to subprefix hijacking under partial ROV deployment. This phenomenon, referred to as collateral damage [18], [19], [20], occurs when an attacker hijacks a more-specific subprefix of an ROA-protected prefix whose length is shorter than /24 for IPv4 or /48 for IPv6.

A necessary but not sufficient condition for subprefix hijacking of an ROA-protected prefix is that the prefix is shorter than /24 for IPv4 or /48 for IPv6. Our analysis shows that the nameservers of 17 PPs are susceptible to subprefix hijacking. Among them, 5 PPs have a zone whose nameservers are all announced from such shorter prefixes, making the zone directly vulnerable to subprefix hijacking. The other 12 PPs have a zone that lacks DNSSEC and includes at least one nameserver announced under a shorter prefix, enabling an attacker to hijack that prefix and inject forged DNS responses.

### D. Connectivity Threats in Download Subprocess

Retrieving RPKI data by RPs from PPs depends on the security of the underlying routing infrastructure. We assume that the RP is properly secured, as its security should be addressed locally by the ASes who need to retrieve RPKI data. We focus on the routing security issues related to the PPs, and evaluate the exposure to prefix hijacking attacks by analyzing whether the PPs are protected by ROAs and whether the networks hosting them enforce ROV.

*1) ROA Registration by PPs:* We evaluate the susceptibility of PPs to prefix hijacking by examining whether the IP prefixes covering their server addresses are protected by valid ROAs.

Our measurements show that among 60 PPs with globally identical IP addresses, 4 lack ROA protection. Two of these are operated by Chinese network providers, consistent with the generally lower ROA adoption rate observed in this region.

Among the 4 PPs with geographically distributed IP addresses, 3 have full ROA coverage for all of their IP addresses across all regions.

In contrast, RIPE NCC's `rrdp.ripe.net` is served through Akamai's CDN and resolves to different IP addresses based on the user's geographic location. We observe that in 12 countries, including Azerbaijan, Brazil, and Morocco, these IP addresses lack ROA coverage. By identifying the ASes responsible for these addresses, we find that in Azerbaijan the affected IPs belong to AS29049 (Delta Telecom Ltd), in Brazil to AS26599 (TELEFÔNICA BRASIL S.A), and in Morocco to AS36925 (MEDITELECOM). This indicates that in these regions, CDN nodes rely on IP addresses provided by local network operators rather than Akamai itself, which further complicates coordinated ROA registration. This region-specific deployment increases the operational complexity of maintaining complete ROA coverage, as each geographic node may be hosted by a different network. Without coordinated ROA registration across all hosting ASes, parts of the infrastructure remain vulnerable to prefix hijacking. Compared to anycast-based approaches like those used by Cloudflare, which serve all regions using a consistent set of IP prefixes, region-specific CDN deployment strategies introduce additional challenges for securing global routing.

*2) ROV Enforcement of Networks Hosting PPs:* Deploying ROV on the networks hosting PPs is essential to prevent attackers from disrupting the connectivity between PPs and RPs by hijacking the IP prefixes of RPs, which would otherwise allow attackers to intercept or redirect RPKI objects sent from PPs to RPs. Based on the ROV deployment classification, we evaluate the ROV protection status of the ASes hosting each PP's servers. In the lower bound, 42.4% of PPs are protected, while in the upper bound all PPs are considered protected.

*3) Vulnerability of PPs to Subprefix Hijacking:* Similarly, we examine the routing table entries covering PP server IPs and find that 37 PPs are potentially vulnerable to subprefix hijacking, far more than in the DNS resolution subprocess. This is largely because many PPs are hosted on cloud platforms that announce aggregated prefixes to serve a large number of tenants. Major hosting networks include Cloudflare (8 PPs), Hetzner (5), Akamai (4), Vultr (3), OVH (3), and Alibaba Cloud (2), indicating that such hosting environments are common in operating PPs and contribute to their vulnerability.

### E. Evaluation of Attack Effectiveness and Impact

This section assesses the effectiveness of prefix hijacking attacks through simulations on the Internet topology. We further analyze the potential impact of both prefix hijacking and DNS spoofing attacks by examining the scale of ROA objects hosted by affected PPs, reflecting the extent of ROA data potentially impacted by these attacks.

*1) Effectiveness of Prefix Hijacking Attacks:* The analysis above, focusing on the ROA registration of the IP addresses of nameservers and PPs, reveals that certain prefixes are susceptible to prefix hijacking attacks owing to the absence of corresponding ROAs. We summarize the security property of the 64 PPs being evaluated as follows. Due to the presence of ROA-unprotected nameservers, the resolution process of 2 PPs is vulnerable to prefix hijack attacks in the presence of resilient local resolvers, and 6 PPs can be attacked by hijacking the insecure nameserver and injecting falsified DNS responses. Due to the presence of ROA-unprotected IP addresses of PP servers, 4 PPs with globally consistent IP addresses are vulnerable to prefix hijack attacks, and 1 PP (the one operated by RIPE that resolves to different addresses across the world) is vulnerable to prefix hijack attacks in some regions. 2 PPs are insecure due to both unprotected nameservers and unprotected PP servers. Therefore, in total, our analysis identifies 10 PPs

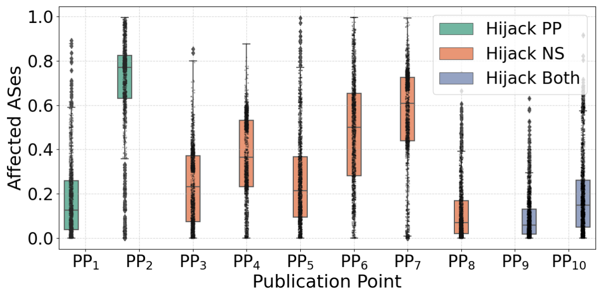

Fig. 9: Fraction of ASes affected by prefix hijacking attacks on PPs, nameservers and both.

with globally exploitable vulnerabilities and 1 with regionally exploitable vulnerability to prefix hijack attacks, all resulting from insufficient ROA protection.

Although being unprotected by ROAs presents an opportunity for attacks, the success of routing hijacks is contingent on additional factors, such as the location (AS number) of the attackers. Therefore, to evaluate the practical risk of such vulnerabilities, we simulate prefix hijacking attacks using the BGPy simulator [41]. For each vulnerable prefix, we identify the victim AS using CAIDA's prefix-to-AS mapping [36], which corresponds to the AS legitimately originating the prefix used by a PP or nameserver. We then model an attacker located in a randomly selected AS that announces the same prefix to perform prefix hijacking. A hijack is considered successful if traffic from relying parties or resolvers is diverted away from the legitimate infrastructure to the attacker. In addition, we focus on globally vulnerable PPs in our simulation, as regionally vulnerable ones may resolve to different origin ASes depending on client location, making their attack surfaces less consistent and difficult to model.

The simulation results shown in Figure 9 illustrate the fraction of ASes whose traffic would be redirected to an attacker during prefix hijacking of PPs or their nameservers. To avoid disclosing information about potentially vulnerable PPs, we represent each PP using an anonymized identifier (e.g., $PP_i$) instead of its real domain name in the figure. We can see that the effectiveness of attacks varies significantly among them. Notably, $PP_2$ exhibits the most severe impact, with approximately 65% to 83% of ASes losing reachability to the PP in most attacks. For the other PPs, prefix hijacking generally results in more than 20% of ASes experiencing disrupted connectivity to the affected PP.

It is important to note that the above analysis primarily focuses on equal-length prefix hijacking. Our further analysis indicates that 7 PPs or their nameservers are hosted on prefixes shorter than /24, making them susceptible to subprefix hijacking. This vulnerability allows an attacker to attract traffic from nearly all ASes, substantially broadening the potential impact.

Additionally, the previous analysis assumes that forged routes propagate without restrictions, making attacks relatively easier to detect. However, attackers could adopt more targeted hijacking strategies by leveraging mechanisms such as BGP communities to limit route propagation [42]. This approach enables disruption of RPKI data retrieval for specific ASes while reducing the risk of detection.

TABLE III: Impact of different attack types on impacted ROA coverage and required attack duration. The *Impacted ROA (%)* column reports IPv4 coverage first, with IPv6 coverage shown in parentheses.

| Subprocess | Attacks | PPs | Impacted ROA (%) | Median Duration (h) |
|---|---|---|---|---|
| DNS Resolution | DNS spoofing (fallback) | 30 | 14.6% (4.2%) | 117.4 |
| | DNS spoofing (no fallback) | 31 | 49.3% (41.3%) | 24.0 |
| | Prefix hijack | 8 | 0.27% (0.11%) | 6.0 |
| | Subprefix hijack | 17 | 9.8% (0.73%) | 27.9 |
| Download | Prefix hijack (Global) | 4 | 0.27% (0.10%) | 6.0 |
| | Prefix hijack (Regional) | 5 | 35.0% (37.2%) | 24.0 |
| | Subprefix hijack | 37 | 29.7% (14.9%) | 168.0 |

*2) Impact of Unreachable PPs Resulting from Attacks:* This subsection evaluates the impact of DNS spoofing, prefix hijacking, and subprefix hijacking attacks using two key metrics. The first metric, *impacted ROA coverage*, quantifies the reduction ratio of the size of ROA-protected IP addresses. The second metric, *required attack duration*, represents the worst-case time needed for an attack to invalidate an RP's cached RPKI data. It is typically determined by the validity of the manifests, as the RP treats all repository objects as "suspicious" and discards them once the manifest becomes stale. The reported duration in Table III represents the median of the worst-case attack durations among all affected PPs.

**Impacted ROA coverage.** For DNS spoofing attacks, the extent of impact depends on whether the relying party is able to fall back to rsync when RRDP endpoints are unavailable. When fallback is supported, 14.6% of the initially protected IPv4 ROA-protected space is left vulnerable post-attack, as the corresponding ROAs for these IP addresses are stored within PPs rendered unreachable due to the attacks, and 4.2% of IPv6 ROA-protected space is left vulnerable post-attack.

However, without the fallback option, the impacted ROA coverage (i.e., reduction of ROA-protected IP address space) increases significantly to 49.3% for IPv4 and 41.3% for IPv6. This amplification is primarily attributed to the vulnerability of the RRDP domain operated by RIPE, which lacks full protection by DNSSEC, rendering it susceptible to spoofing. In contrast, RIPE's rsync domain is properly signed and generally resilient against such attacks. Yet, fallback to rsync is only triggered when the relying party software is configured to do so. If DNS spoofing effectively prevents access to the RRDP domain and no fallback mechanism is in place, relying parties may be unable to retrieve timely validation data.

Prefix hijacking generally affects fewer PPs, which correspond to a relatively small portion of the ROA-covered address space. However, an exception arises in regional hijacking scenarios related to RIPE's CDN deployment strategy. RIPE delivers RRDP data through geographically distributed CDN nodes, each using different IP prefixes depending on the user's location. Some of these IP prefixes are not covered by ROAs, leaving them vulnerable to potential prefix hijacking attacks within specific countries. As a result, prefix hijacking in these regions reduces the size of ROA-protected address space by as much as 35.0% for IPv4 and 37.2% for IPv6.

These findings highlight two security challenges caused by CDN deployment in the RPKI ecosystem. First, the use of CNAME records by CDNs complicates DNSSEC deployment, resulting in incomplete protection of some PP domains. Second, the geographic distribution of CDN nodes leads to varying IP prefixes across regions, some of which lack ROA coverage and are therefore vulnerable to prefix hijacking. To address these issues, operators must ensure comprehensive DNSSEC protection for all domain names involved in CDN redirection, including CNAME targets. Additionally, all IP prefixes used by CDN nodes to serve RPKI data should be covered by valid ROAs.

**Required attack duration.** The required duration varies significantly across attack types. Attacks targeting small non-RIR PPs, such as prefix hijacks, typically need to persist for only about 6 hours before more than half of the cached manifests of these PPs expire. In contrast, attacks affecting a larger number of ROAs usually need to last around 24 hours before a majority of manifests become invalid. This difference is mainly caused by the validity periods configured by different PPs. RIR-hosted PPs generally assign longer manifest lifetimes, around seven days for APNIC and two to three days for ARIN, while most non-RIR PPs set validity periods shorter than 24 hours. Consequently, even short-lived disruptions can cause cached manifests of small PPs to expire, temporarily disabling ROV for the prefixes covered by the affected VRPs.

## V. PP Dependencies and Cascading Risks

In the previous section, we analyzed threats to the reachability of individual PPs arising from their underlying infrastructure. However, a PP's reachability can also depend on data from other PPs, creating dependency relationships. In this section, we identify these dependencies by examining certificate chains and ROAs issued by other PPs that cover the IP prefixes of PPs or their nameservers. Using this, we build a dependency graph to evaluate how individual failures can cascade among PPs.

### A. PP Dependency Analysis

PPs exhibit two distinct forms of interdependencies within the RPKI system.

- **Deterministic dependency:** Each certificate's SIA field specifies the URL of the PP storing the RPKI objects it signs, including ROAs and child certificates. To retrieve these objects, the RP must first obtain the parent certificate from its PP. Thus, downstream RPKI data availability depends directly on the availability of the parent certificate's PP, forming a deterministic dependency.
- **Probabilistic dependency:** $PP_{\mathrm{A}}$ probabilistically depends on $PP_{\mathrm{B}}$ if $PP_{\mathrm{B}}$ issues ROAs covering the IP prefixes of $PP_{\mathrm{A}}$ or its authoritative nameservers. This dependency is termed probabilistic because the unavailability of $PP_{\mathrm{B}}$ results in loss of ROA protection for $PP_{\mathrm{A}}$, exposing it to prefix hijacking threats. However, whether the hijacking succeeds depends on the attacker's capabilities and network conditions.

Using the collected RPKI snapshot and the measured IP addresses of PPs and their authoritative nameservers, we identify both types of inter-PP dependencies. For deterministic dependencies, we recursively parse the SIA field of each certificate to build the complete chain of PPs from the current PP up to the trust anchor. All PPs appearing in this chain are recorded as direct dependencies.

For probabilistic dependencies, we examine whether a PP or any authoritative nameserver associated with it resides within a prefix covered by an ROA issued by another PP. For the PP itself, a dependency is recorded if its IP prefix is protected by an ROA issued by another PP. For nameservers, we take into account the DNSSEC status of the zone they belong to. A successful hijack attack on a DNSSEC-protected zone requires compromising all its authoritative nameservers. In contrast, zones without DNSSEC protection are vulnerable even if only a single nameserver can be hijacked. Accordingly, we establish a dependency only if all nameservers in a DNSSEC-protected zone have prefixes covered by ROAs from the same PP. For unprotected zones, a single such nameserver suffices to introduce a dependency.

For each PP, we compute how many other PPs exhibit deterministic or probabilistic dependencies on it. The results are shown in Figure 10, where orange bars represent deterministic dependencies and blue bars represent probabilistic ones. We observe that a small number of PPs act as critical hubs, with numerous other PPs depending on them via both dependency types. In particular, the PPs operated by RIPE, APNIC, and ARIN serve as key hubs, each referenced by over 15 other PPs for parent certificate retrieval. Among them, ARIN PP stands out by providing ROA protection for the prefixes of 36 other PPs or their associated authoritative nameservers, representing the largest share of probabilistic dependencies. This concentration creates significant threats, as disruptions to these PPs can cascade across the ecosystem by blocking certificate retrieval or increasing the prefix hijacking threats for PPs that depend on them.

Additionally, the other two RIR-operated PPs, LACNIC and AFRINIC, have few other PPs depending on them. This reflects a preference for the hosted mode, in which the RIR centrally manages and publishes RPKI data on behalf of its members. In contrast, RIPE, ARIN, and APNIC PPs predominantly adopt delegated mode, enabling their members to operate independent RPKI repositories. This difference in operational modes contributes to the variation in dependency patterns observed among RIR-operated PPs.

However, we also observe that some RIR-operated PPs themselves depend on non-RIR PPs, meaning that critical

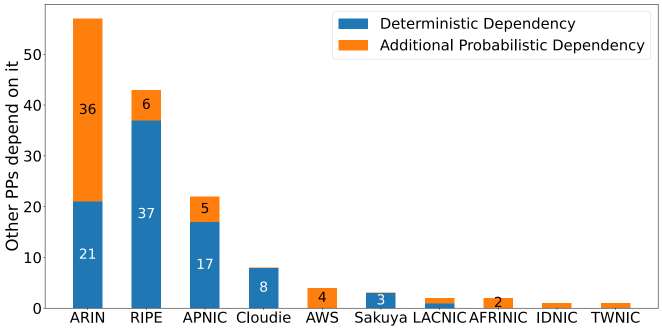

Fig. 10: Number of other PPs that deterministically or probabilistically depend on each PP.

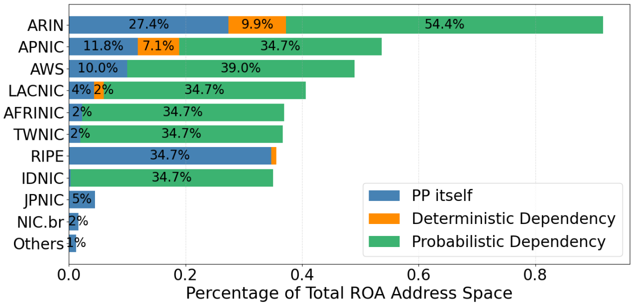

Fig. 11: Fraction of ROA-Covered address space published by each PP and by PPs that depend on it.

infrastructure components may rely on third-party PPs for reachability. Such dependencies introduce potential cascading risks because failures or attacks on these non-RIR PPs may cause RIR-operated PPs to lose routing protection, increasing their exposure to prefix hijacking. We illustrate this dependency phenomenon with the following cases.

- **RIPE PP.** The RIPE PP exhibits probabilistic dependencies on TWNIC, IDNIC, and AWS-hosted PPs due to its reliance on Akamai CDN infrastructure. Its domain name is mapped to the CNAME `a1513.dscd.akamai.net`, whose IP addresses and authoritative nameservers (e.g., `n4dscd.akamai.net`, `n6dscd.akamai.net`) vary across regions. The prefixes of these servers are covered by ROAs issued by different PPs: IDNIC PP in Indonesia, TWNIC PP in Taiwan, and AWS-hosted PP in Ireland. If any of these PPs become unreachable, the corresponding ROAs may be lost, weakening routing protection in affected regions. This undermines the routing security of both the RIPE PP and its resolution infrastructure, increasing the risk of prefix hijacking that could block access even in networks enforcing ROV.
- **LACNIC PP.** The LACNIC PP exhibits probabilistic dependencies on AWS-hosted PP due to the routing protection of its authoritative nameservers. Specifically, the domain `rrdp.lacnic.net` is delegated to authoritative nameservers hosted on the AWS DNS infrastructure, such as `ns-1094.awsdns-08.org` and `ns-394.awsdns-49.com`. The IP prefixes of these nameservers are protected by ROAs published via AWS-hosted PP. If the AWS-hosted PP becomes unreachable, the associated IP prefixes will lose ROA protection, exposing the LACNIC PP's resolution process to prefix hijacking threats. This weakens the routing security of the LACNIC PP globally, potentially allowing attackers to disrupt RPs' access to the LACNIC PP by hijacking the prefixes of its authoritative nameservers.

To further understand the potential impact of these dependencies, we quantify the address space covered by the ROAs published on each PP, as well as the cumulative address space of those PPs that rely on it. Specifically, for each PP, we compute the fraction of the ROA-covered address space that is (i) directly published by the PP, and (ii) associated with other PPs that exhibit deterministic or probabilistic dependencies on it.

Figure 11 shows the distribution of ROA-covered address space based on the two categories defined above. The results reveal a clear imbalance across PPs. Those operated by RIPE, ARIN, and APNIC together provide the ROA hosting service for the majority of ROA-secured space, with RIPE contributing about 34.7%, ARIN 27.4%, and APNIC 11.8%. In contrast, LACNIC and AFRINIC contribute only 4.3% and 2.2%. Although PPs operated by LACNIC, AFRINIC, TWNIC, IDNIC, and AWS cover smaller portions of the ROA-protected space, their unavailability can still weaken routing security. Some high-coverage PPs like RIPE depend on these smaller PPs for ROA protection of the prefixes associated with the PP server or the authoritative nameservers. Losing such support in certain regions reduces their reachability and exposes RPs to localized prefix hijacking.

This concentration of ROA coverage and cross-PP reliance creates potential failure points that can cascade across the ecosystem, amplifying risks to RPKI validation. We analyze these cascading risks in the next subsection.

### *B. Cascading Risks Amplified by Inter-PP Dependencies*

Sections IV analyzed vulnerabilities in the underlying infrastructure supporting individual PPs. Here, we explore how disruptions to one PP can cascade through dependencies, amplifying risks by causing wider ROA unavailability and weakening RPKI validation across the ecosystem.

We first identify PPs that are especially vulnerable to attacks targeting infrastructure components beyond the control of ROV-enforcing ASes, as described in Section III. The susceptible PPs are summarized below.

- **DNS Spoofing Attack:** As identified in Section IV-B2, 30 PPs have neither their RRDP nor rsync domain names fully protected by DNSSEC. This exposes them to risks of DNS spoofing attacks, which can lead to forged DNS responses that misdirect RPs to malicious PP servers.
- **Prefix Hijacking Attack:** As detailed in Section IV-E1, our analysis identifies 10 PPs and their authoritative nameservers with IP prefixes that lack adequate ROA

protection globally, making them vulnerable to prefix hijacking regardless of DNSSEC validation. In addition, 1 PP operated by RIPE exhibits regional vulnerability due to incomplete ROA coverage of CDN-backed IP addresses resolved in specific areas. In both cases, attackers can inject forged BGP routes to hijack these prefixes, disrupting access to the affected PPs.

- **Low-Rate Attack:** We replicated the experiment proposed by [25] and identified 34 PPs whose authoritative nameservers are vulnerable to low-rate attacks. Additionally, 14 PPs themselves are also susceptible to such attacks. In this attack, an adversary can forge the source IP address of a victim resolver or RP to send a large volume of queries to the targeted nameserver or PP. This induces rate-limiting on the target, causing it to drop or ignore subsequent legitimate requests from the victim. Accounting for the overlap between vulnerable nameservers and PPs, we find that 37 distinct PPs in total are exposed to this attack.

We find that RIR-operated PPs are generally well protected against the three classes of attacks discussed earlier. Most deploy DNSSEC for their domains, host repositories in ROV-enforcing networks, and avoid aggressive rate-limiting configurations. However, some PPs, such as RIPE and LACNIC, remain vulnerable to region-specific prefix hijacking and DNS spoofing attacks.

Moreover, even well-secured PPs can become indirectly vulnerable due to dependencies on other PPs. For instance, the RIPE PP depends on the LACNIC PP in South America and an AWS-hosted PP in Ireland, both of which lack full DNSSEC protection. If these upstream PPs are compromised, attackers can hijack prefixes used by the RIPE PP in those regions, disrupting RPs' access to the RIPE repository. This demonstrates how strong protections at one PP can be weakened by weaker security in its upstream infrastructure.

Table IV summarizes the impact of three types of attacks on PPs, distinguishing between their direct effects and cascading effects amplified through inter-PP dependencies. The scope column distinguishes regional impacts, limited to specific areas, from global impacts affecting locations worldwide. We observe that DNS spoofing attacks and low-rate attacks can compromise non-RIR PPs on which the RIPE PP probabilistically depends, triggering cascading disruptions that affect 39 and 42 PPs, and expanding the impacted ROA coverage from 14.58% and 6.62% to 50.89% and 42.92%, respectively. These results demonstrate how inter-PP dependencies can amplify infrastructure weaknesses and escalate localized failures into broader validation disruptions.

Additionally, the AWS-hosted PP has relatively low rate-limiting thresholds, making it more vulnerable to low-rate attacks. Because it provides ROA protection for the prefixes of LACNIC PP's authoritative nameservers, an outage or attack on this PP could increase the global impacted ROA coverage to 8.24%, potentially preventing RPs from resolving and accessing the LACNIC PP.

TABLE IV: The direct and cascading impact of three types of attacks on RPKI PPs, evaluated by *impacted ROA coverage*, *i.e.*, the reduction of the size of ROA-protected IP addresses.

| Attack | Scope | Direct | | Cascading | |
|---|---|---|---|---|---|
| | | PPs | ROA Cov. | PPs | ROA Cov. |
| DNS Spoofing | Regional | 30 | 14.58% | 39 | 50.89% |
| | Global | 30 | 14.58% | 38 | 16.20% |
| Prefix Hijacking | Regional | 11 | 34.96% | 48 | 35.88% |
| | Global | 10 | 0.27% | 12 | 0.27% |
| Low-Rate Attack | Regional | 37 | 6.62% | 42 | 42.92% |
| | Global | 37 | 6.62% | 41 | 8.24% |

## VI. Discussion

This section presents an analysis of possible causes of the vulnerabilities identified in our measurements and provides corresponding mitigation recommendations.

### A. Possible Causes of PP Vulnerabilities

Identifying the root causes of PP vulnerabilities is inherently challenging, as security deployment decisions made by PP operators are internal and rarely disclosed. To better understand the possible causes, we integrate our measurements with insights from prior work and operator feedback, although only 8 operators responded. Our analysis suggests two possible causes discussed below.

First, deploying ROAs and DNSSEC can be operationally challenging. Registering ROAs for globally distributed infrastructures is complex, as such systems often rely on anycast or DNS-based deployments. In anycast setups, multiple legitimate origin ASes may announce the same prefix, and omitting any of them from the ROA can cause valid routes to be rejected. In DNS-based deployments, a single domain or nameserver may map to multiple IP addresses across regions, and missing ROA coverage for any address can leave them exposed to prefix hijacking. In addition, regional and policy-related factors pose further challenges. In some countries, regulatory and economic factors result in low ROA adoption, making it difficult for operators using local infrastructure to achieve full protection. Another factor is that certain TLDs, such as `.ng` and `.im`, have not deployed DNSSEC and therefore lack DS records in the root zone. As a result, domains under these TLDs cannot be validated through the DNSSEC chain of trust and remain effectively unprotected.

Second, some PP operators are reluctant to deploy security mechanisms. As noted in [12], privacy concerns may lead operators to avoid registering ROAs, since this can expose sensitive inter-organizational relationships. In addition, some PPs are still in an experimental phase and have not yet considered security issues. However, although these PPs may not serve RPKI data, their presence in the ecosystem introduces potential risks, as attackers could exploit them to delay RP synchronization, potentially causing cached RPKI data to expire.

### B. Recommendations for Mitigating Risks

Our analysis indicates that the DNS and routing infrastructures engaged in the retrieval process between RPs and

PPs lack comprehensive protection through DNSSEC and RPKI implementation. This deficiency suggests that networks employing ROV could encounter difficulties in reaching PPs. To mitigate the risks highlighted in our study and fortify the resilience of the RPKI system, we put forth the following recommendations.

**On the deployment of DNSSEC.** 1) PP operators should select domain names whose authoritative nameservers have adopted DNSSEC. Our findings show that 10.9% of all PPs currently lack DNSSEC protection solely because they use domain names with non-DNSSEC-enabled authoritative nameservers. 2) PP operators should carefully audit the entire DNS resolution path for their PPs to avoid vulnerabilities such as insecure CNAME redirections or NS delegations to unsecured zones. 3) RP operators should implement DNSSEC-enabled local resolvers. This measure can significantly reduce the exposure to DNS spoofing attacks, dropping the percentage of affected PPs from 85.9% to just 12.5%.

**On the deployment of RPKI.** 1) PP operators should register ROAs for all of their IP addresses - note that some PPs are resolved to different IP addresses in diverse geographical regions. Our findings show that a total of 5 PPs, 4 with single IP addresses and 1 with multiple IP addresses, have not properly registered their ROAs. 2) PP operators should ensure that their PPs are hosted in networks that have enabled ROV. While no direct evidence was found of ROV-unprotected PPs, the ROV status of 57.8% of PPs remains unknown due to the inherent difficulties in measuring ROV status. 3) PP operators should carefully audit the entire DNS resolution path to ensure that all authoritative nameservers involved have registered ROA for their IP addresses and are deployed in ROV-enabled networks.

**PP redundancy and distribution.** To improve resilience, it is desirable for each PP to operate multiple servers with distinct IP addresses distributed across different regions. Similarly, authoritative nameservers are preferably geographically dispersed to reduce the impact of local outages. This diversity helps RPs retrieve RPKI data even if some servers or nameservers become temporarily unreachable.

**Be careful if deployed in CDNs.** While deploying PPs in CDNs can enhance availability, scalability, and latency, this approach must be carefully evaluated due to the security challenges it introduces. First, CDNs often employ CNAME redirection, which lengthens the DNS resolution path and increases the risk of involving insecure DNS zones. Second, because CDNs always provide service from multiple IP addresses across different geographical regions, it is essential that all associated IP prefixes have corresponding ROA entries to ensure reliable PP reachability worldwide.

**Monitoring and alerting mechanisms for PP availability.** A distributed monitoring system can be deployed globally to continuously track the availability and security of all PPs. When repeated polling attempts to a specific PP fail, the system can issue alerts to operators, enabling timely investigation and coordinated recovery before such failures lead to the expiration of cached RPKI data.

**Be aware of the risks of inter-PP dependencies.** Some PPs create dependencies by using nameservers or prefixes protected by ROAs issued by other PPs, introducing potential cascading risks. To enhance resilience, critical infrastructures such as DNS and PP servers should use IP addresses protected by ROAs from trusted and secure PPs. When multiple PP servers or nameservers are deployed for redundancy, they should be assigned to prefixes protected by ROAs from different PPs to avoid dependence on a single one.

## VII. Conclusion

In this paper, we show that while RPKI is essential for securing Internet routing, its reliance on underlying infrastructure introduces vulnerabilities that have not been thoroughly examined. Our measurements reveal that incomplete DNSSEC deployment and insufficient ROA coverage leave many PPs exposed to attacks, including DNS spoofing and prefix hijacking. Beyond these direct risks, we also uncover dependencies among PPs that could lead to cascading failures within the ecosystem.

These findings highlight the need for improved operational practices, such as broader DNSSEC deployment, more consistent ROA coverage, and reducing risky dependencies among PPs. Future work should focus on monitoring these dependencies and evaluating mitigation strategies to enhance the resilience and trustworthiness of RPKI.

## Ethical Considerations

Ethical issues have been carefully considered throughout this study. First, in our reproduction of the methodology from [25] to identify nodes prone to rate limiting, we adopted the same cautious, stepwise probing strategy. We gradually increased the query rate and immediately halted upon observing signs of rate limiting. The thresholds used to detect such behavior were consistent with those reported in the original study. Second, we made our best efforts to notify the operators of all vulnerable publication points. In total, there are 43 PPs in our dataset that are susceptible to attacks, including prefix hijacking, subprefix hijacking, and DNS spoofing. To obtain contact information for the operators, we employed four approaches: (1) searching the official websites of the PPs via search engines, (2) querying PEERINGDB for operator contact emails, (3) examining the ROAs stored at each PP to determine which ASes are authorized to originate the prefixes, thereby inferring the likely operator of the PP, and (4) performing WHOIS queries on the relevant domains to locate email addresses. Using these methods, we successfully identified emails for 37 PPs and sent a total of 35 notifications, with two operators managing two PPs each.

## Acknowledgement

We sincerely thank our shepherd and all anonymous reviewers for their valuable comments. We also thank the participants who responded to our emails for their helpful feedback. This work was supported in part by the National Natural Science Foundation of China under Grant 62472246 and Grant 62072269, and in part by Zhongguancun Laboratory.

# REFERENCES

[1] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, 2009.

[2] N. Ripe, "YouTube Hijacking: A RIPE NCC RIS case study," *http://www.ripe.net/news/study-youtube-hijacking.html*, 2008.

[3] H. Ballani, P. Francis, and X. Zhang, "A study of prefix hijacking and interception in the Internet," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4, pp. 265–276, 2007.

[4] P.-A. Vervier, O. Thonnard, and M. Dacier, "Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks," in *Proceedings of the 2015 Network and Distributed System Security (NDSS) Symposium*, 2015.

[5] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," RFC 6480, Feb. 2012. [Online]. Available: https://www.rfc-editor.org/info/rfc6480.

[6] M. Lepinski and K. Sriram, "BGPsec Protocol Specification," RFC 8205, Sep. 2017. [Online]. Available: https://www.rfc-editor.org/info/rfc8205.

[7] A. Azimov, E. Bogomazov, R. Bush, K. Patel, and J. Snijders, "Verification of AS_PATH using the resource certificate public key infrastructure and autonomous system provider authorization," *Network Working Group Internet Draft. November*, 2020.

[8] NIST, "NIST RPKI Monitor," 2025. [Online]. Available: https://rpki-monitor.antd.nist.gov/.

[9] W. Li, Z. Lin, M. I. A. Khan, E. Aben, R. Fontugne, A. Phokeer, and T. Chung, "RoVista: Measuring and Understanding the Route Origin Validation (ROV) in RPKI," in *Proceedings of the ACM Internet Measurement Conference (IMC'23)*, Montreal, Canada, October 2023.

[10] "RPKI PP Infrastructure Security Measurement," 2025. [Online]. Available: https://github.com/internetsys/RPKI-PP-Infrastructure-Security-Measurement.

[11] M. Wählisch, O. Maennel, and T. C. Schmidt, "Towards Detecting BGP Route Hijacking using the RPKI," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 103–104, 2012.

[12] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, "RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem," in *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, 2015, pp. 1–7.

[13] D. Iamartino, C. Pelsser, and R. Bush, "Measuring BGP route origin registration and validation," in *International Conference on Passive and Active Network Measurement*. Springer, 2015, pp. 28–40.

[14] T. Chung, E. Aben, T. Bruijnzeels, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, R. v. Rijswijk-Deij, J. Rula *et al.*, "RPKI is coming of age: A longitudinal study of RPKI deployment and invalid route origins," in *Proceedings of the Internet Measurement Conference*, 2019, pp. 406–419.

[15] 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," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 1, pp. 19–27, 2018.

[16] W. Chen, Z. Wang, D. Han, C. Duan, X. Yin, J. Yang, and X. Shi, "ROV-MI: Large-Scale, Accurate and Efficient Measurement of ROV Deployment," in *Proceedings of the 2022 Network and Distributed System Security (NDSS) Symposium*, 2022.

[17] T. Hlavacek, H. Shulman, N. Vogel, and M. Waidner, "Keep Your Friends Close, but Your Routeservers Closer: Insights into RPKI Validation in the Internet," in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 4841–4858.

[18] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman, "Are we there yet? On RPKI's deployment and security," in *Proceedings of the 2017 Network and Distributed System Security (NDSS) Symposium*, 2017.

[19] R. Morillo, J. Furuness, C. Morris, J. Breslin, A. Herzberg, and B. Wang, "ROV++: Improved Deployable Defense against BGP Hijacking," in *Proceedings of the 2021 Network and Distributed System Security (NDSS) Symposium*, 2021.

[20] W. Li, Y. Li, and T. Chung, "ImpROV: Measurement and Practical Mitigation of Collateral Damage in RPKI Route Origin Validation," in *34th USENIX Security Symposium (USENIX Security 25)*, 2025, pp. 3631–3647.

[21] van Hove, Koen and van der Ham-de Vos, Jeroen and van Rijswijk-Deij, Roland, "rpkiller: Threat Analysis of the BGP Resource Public Key Infrastructure," *Digital Threats*, vol. 4, no. 4, Oct. 2023.

[22] D. Mirdita, H. Shulman, and M. Waidner, "Poster: RPKI kill switch," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 3423–3425.

[23] D. Mirdita, H. Schulmann, N. Vogel, and M. Waidner, "The CURE to vulnerabilities in RPKI validation," in *Proceedings of the 2024 Network and Distributed System Security (NDSS) Symposium*, 2024.

[24] Cooper, Danny and Heilman, Ethan and Brogle, Kyle and Reyzin, Leonid and Goldberg, Sharon, "On the risk of misbehaving RPKI authorities," in *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*, 2013, pp. 1–7.

[25] T. Hlavacek, P. Jeitner, D. Mirdita, H. Shulman, and M. Waidner, "Stalloris: RPKI downgrade attack," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 4455–4471.

[26] Hlavacek, Tomas and Jeitner, Philipp and Mirdita, Donika and Shulman, Haya and Waidner, Michael, "Behind the scenes of RPKI," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 1413–1426.

[27] T. Hlavacek, P. Jeitner, D. Mirdita, H. Shulman, and M. Waidner, "Beyond Limits: How to Disable Validators in Secure Networks," in *Proceedings of the ACM SIGCOMM 2023 Conference*, ser. ACM SIGCOMM '23, 2023, p. 950–966.

[28] D. Mirdita, H. Schulmann, and M. Waidner, "SoK: An Introspective Analysis of RPKI Security," in *34th USENIX Security Symposium (USENIX Security 25)*, 2025, pp. 3649–3665.

[29] S. Weiler, S. Kent, G. Huston, and G. Michaelson, "Resource Public Key Infrastructure (RPKI) Trust Anchor Locator," RFC 6490, Feb. 2012. [Online]. Available: https://www.rfc-editor.org/info/rfc6490

[30] T. Bruijnzeels, O. Muravskiy, B. Weber, and R. Austein, "The RPKI Repository Delta Protocol (RRDP)," RFC 8182, Jul. 2017. [Online]. Available: https://www.rfc-editor.org/info/rfc8182.

[31] R. Bush and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1," RFC 8210, Sep. 2017. [Online]. Available: https://www.rfc-editor.org/info/rfc8210.

[32] S. Rose, M. Larson, D. Massey, R. Austein, and R. Arends, "DNS Security Introduction and Requirements," RFC 4033, Mar. 2005. [Online]. Available: https://www.rfc-editor.org/info/rfc4033

[33] Mutually Agreed Norms for Routing Security (MANRS), "MANRS ISP Guide," https://www.manrs.org/isps/guide/.

[34] RIPE NCC, "RIPE RPKI Repository," https://ftp.ripe.net/rpki/, 2025.

[35] Public DNS, "Public DNS Server List," https://public-dns.info.

[36] CAIDA, "Routeviews prefix to as mappings dataset for ipv4 and ipv6," Jul. 2025. [Online]. Available: https://www.caida.org/catalog/datasets/routeviews-prefix2as/

[37] Y. Su, D. Li, L. Chen, Q. Li, and S. Ling, "dRR: A Decentralized, Scalable, and Auditable Architecture for RPKI Repository," in *Proceedings of the 2024 Network and Distributed System Security (NDSS) Symposium*, 2024.

[38] ICANN, "Majestic Million DNSSEC Deployment Statistics," https://ithi.research.icann.org/graph-m11.html, 2025.

[39] S. Zhuang, J. H. Wang, J. Wang, Z. Pan, T. Wu, F. Li, and Z. Zhang, "Discovering Obscure Looking Glass Sites on the Web to Facilitate Internet Measurement Research," in *Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '21, 2021, p. 426–439.

[40] CAIDA, "CAIDA Serial 1 Data Set," http://www.caida.org/data/as-relationships/serial-1/, Jul. 2025.

[41] J. Furuness, C. Morris, R. Morillo, A. Herzberg, and B. Wang, "Bgpy: The bgp python security simulator," in *Proceedings of the 16th Cyber Security Experimentation and Test Workshop*, 2023, pp. 41–56.

[42] H. Birge-Lee, L. Wang, J. Rexford, and P. Mittal, "SICO: Surgical Interception Attacks by Manipulating BGP Communities," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 431–448.

# APPENDIX

## A. Multi-Week Measurement

We conducted measurements over multiple periods to observe the evolution of PP security. First, prior to our initial submission in July, we performed continuous measurements for three days. In the week leading up to our disclosure in October, we conducted continuous measurements to establish

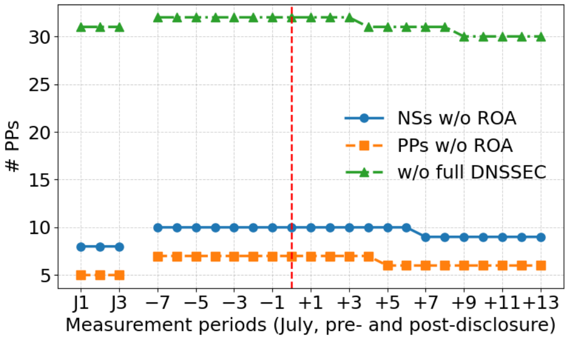

Fig. 12: Number of PPs whose reachability was not fully secured by DNSSEC or ROAs. The red dashed line indicates the disclosure date.

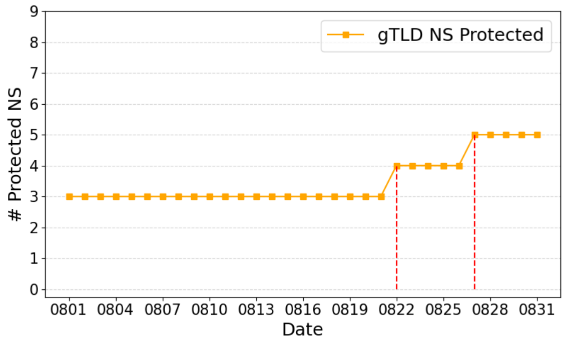

Fig. 13: Daily number of gTLD authoritative nameservers with ROA protection in August 2025.

a precise pre-disclosure baseline. Finally, after the disclosure, we continued measurements for two weeks to track potential improvements. To reduce the impact of short-term fluctuations, each measurement was repeated three times, and the most frequently observed result was used to provide a more reliable view of PP security. Figure 12 shows the number of PPs whose reachability was not fully secured by DNSSEC or ROAs across these periods.

1) **DNSSEC deployment.** In the week prior to disclosure, 32 PPs had at least one zone in their resolution path not secured by DNSSEC, compared with 31 during the three-day measurement in July 2025. This increase was mainly due to the total number of PPs growing from 64 to 67, while only one of the 31 previously unprotected PPs from July deployed DNSSEC, indicating limited improvement within the original dataset. In the two weeks following disclosure, the number of PPs with unprotected zones gradually decreased to 30, including one case in which an operator confirmed that they enabled DNSSEC on their SLD on November 6 and published the corresponding DS record in the `.net` zone.
2) **ROA coverage.** In July, 5 PPs were hosted on IP prefixes that did not have valid ROAs, and 8 PPs were served by authoritative nameservers whose IP prefixes were not covered. By the week prior to disclosure, these numbers had increased slightly to 7 and 10, respectively, mainly because several newly added PPs were hosted on prefixes without ROAs. Following disclosure, the number of PPs hosted on unprotected prefixes decreased to 6, and the number of PPs with nameservers lacking ROA protection decreased to 9. This improvement was primarily driven by one PP whose operator registered an ROA for the PP's hosting prefix on November 7 and an ROA for the prefix of its last-hop authoritative nameservers on November 9. We also measured ROA coverage for all CDN nodes used by RIPE PP and found no changes compared with July.

Since the IP addresses of authoritative nameservers for gTLD domains are generally stable, we tracked changes in their ROA coverage from August to the present. Although these prefixes may be announced by dozens of ASes, making comprehensive ROA protection challenging, the number of gTLD nameservers protected by ROAs increased from three to five during August, demonstrating that Verisign is actively working to secure these nameservers. The number of protected nameservers has remained at five since then. Figure 13 shows the number of gTLD nameservers protected by ROAs during August 2025. We plan to repeat these measurements one year from now.