

To Filter or not to Filter: Measuring the Benefits of Registering in the RPKI Today

Cecilia Testart¹, Philipp Richter¹, Alistair King², Alberto Dainotti², and David Clark¹

¹ MIT {ctestart, richterp, ddc}@csail.mit.edu

² CAIDA, UC San Diego {alistair, alberto}@caida.org

Abstract. Securing the Internet’s inter-domain routing system against illicit prefix advertisements by third-party networks remains a great concern for the research, standardization, and operator communities. After many unsuccessful attempts to deploy additional security mechanisms for BGP, we now witness increasing adoption of the RPKI (Resource Public Key Infrastructure). Backed by strong cryptography, the RPKI allows network operators to register their BGP prefixes together with the legitimate Autonomous System (AS) number that may originate them via BGP. Recent research shows an encouraging trend: an increasing number of networks around the globe start to register their prefixes in the RPKI. While encouraging, the actual *benefit* of registering prefixes in the RPKI eventually depends on whether transit providers in the Internet enforce the RPKI’s content, *i.e.*, configure their routers to validate prefix announcements and *filter* invalid BGP announcements. In this work, we present a broad empirical study tackling the question: To what degree does registration in the RPKI protect a network from illicit announcements of their prefixes, such as prefix hijacks? To this end, we first present a longitudinal study of filtering behavior of transit providers in the Internet, and second we carry out a detailed study of the visibility of legitimate and illegitimate prefix announcements in the global routing table, contrasting prefixes registered in the RPKI with those not registered. We find that an increasing number of transit and access providers indeed do enforce RPKI filtering, which translates to a direct benefit for the networks using the RPKI in the case of illicit announcements of their address space. Our findings bode well for further RPKI adoption and for increasing routing security in the Internet.

Keywords: Internet security · Routing · RPKI · BGP.

1 Introduction

The inter-domain routing system of the Internet continues to suffer from major routing incidents, including accidental route leaks causing widespread disruptions [28], and intentional prefix hijacks for malicious purposes [8, 14, 29]. At the heart of the problem lies BGP’s lack of mechanisms for route authentication: a network that receives a route advertisement from a neighbor has no easy means

to validate its correctness. The RPKI [20] represents one of the most recent attempts to increase BGP security, providing networks in the Internet with a trustworthy database that maps BGP prefixes to the Autonomous System (AS) number that is authorized to originate them. The RPKI is backed by strong cryptography, with the Regional Internet Registries (RIRs) serving as trust anchors. Networks can leverage this data to validate that incoming BGP announcements point to the correct origin AS. Recent research shows an encouraging trend of both increasing global registration of prefixes in the RPKI (17% of routed prefixes are registered in the RPKI as of September 2019), as well as increasing data quality of actual RPKI records [11]. The RPKI has thus the potential to finally provide a universally trusted route origins database, a major building block to greatly improve routing security.

While encouraging, we point out that increasing registration of prefixes in the RPKI only represents a first step towards securing BGP. The eventual benefit of RPKI registration depends on whether the networks of the Internet enforce the RPKI's contents, *i.e.*, drop invalid announcements and hence do not propagate them to their neighbors. Recently, AT&T, a major transit ISP, publicly announced that they started filtering BGP announcements that are invalid as per the RPKI [2], suggesting increasing acceptance and trust by major transit providers in the RPKI. However, besides such anecdotal evidence, we know little about current levels of RPKI *enforcement* in the Internet and, as of today, have no way to assess the resulting benefits of RPKI registration.

To tackle these questions, we empirically study to what degree networks in the Internet filter BGP announcements based on RPKI validation and show to what extent registration in the RPKI benefits networks in situations in which RPKI is needed the most: instances of conflicting prefix announcements in the global routing table, such as those caused by misconfiguration and prefix hijacking. Our key contributions are as follows:

- Leveraging historical snapshots of the global routing table and validated RPKI records, we develop a passive method to detect filtering of RPKI invalid prefixes for IPv4 and IPv6 and study filtering deployment over time. While RPKI filtering was virtually nonexistent just two years ago, RPKI enforcement has increased substantially: we found that—as of January 2020—approximately 10% of the networks we considered, including major transit providers, filter invalid announcements.
- We study the effect of RPKI filtering on global prefix reachability in the case of conflicting announcements: Multiple-Origin AS (MOAS) conflicts, and subprefix announcements, contrasting our findings with a baseline of non-RPKI-registered prefixes. We find that, already as of today, RPKI filtering starts to show effect in real-world cases: in all considered scenarios, registration of prefixes in the RPKI results in limited reachability of conflicting and invalid (potentially illicit) prefix announcements in the global routing table.

Our findings are encouraging for the research, standardization, and operator communities. Increasing RPKI *enforcement* starts to translate to a direct benefit

for a network registering its prefixes. Our results bode well for increasing routing security in the Internet, and our metrics allow for easy assessment of current levels of filtering and the resulting benefit in conflicting-announcement scenarios. Our study is entirely based on publicly available datasets, allowing both for reproducibility, and for continuous monitoring.

2 Background and Datasets

2.1 Related Work

The IETF has devoted substantial efforts to develop, and document in detail, the RPKI over the last years [9, 15–17, 19–21, 24]. Recently, the research community started to measure RPKI deployment in the Internet. Chung *et al.* provide both an accessible overview of today’s RPKI deployment and an extensive study of RPKI registration and usage patterns. They find increasing registration of prefixes and networks in the RPKI, and overall higher data quality of RPKI records, resulting in lower numbers of RPKI-invalid prefixes caused by misconfiguration by the respective operators [11]. Iamartino *et al.* had previously measured problems with RPKI registered ROAs and the potential impact that validation and filtering of RPKI-invalid announcements could have in production [18].

To the best of our knowledge, only two previous academic studies, using two different methods, touched upon the adoption of RPKI-invalid filtering, finding only negligible RPKI filtering in 2016 and 2017. Gilad *et al.* analyze a month of BGP RIB dumps from 44 ASes [13]. Their passive approach uses all the ASes but the last hop in the AS path of RPKI-valid and -invalid announcements to identify ASes filtering invalid announcements. They find that, in July 2016, only 3 of the top 100 ASes (by customer cone size) were enforcing RPKI-invalid filtering. Reuter *et al.* instead, actively advertise RPKI-valid and -invalid prefixes of address space under their control [26]. They infer which ASes filter RPKI-invalid announcements based on the propagation path of their announcements, finding only 3 ASes filtering in 2017. Measuring RPKI filtering also caught attention from the operator community: Cartwright-Cox uses active measurements to infer filtering based on presence or absence of ICMP responses from probed IP addresses in RPKI-valid and -invalid prefixes [10].

Our study complements and extends prior work: our passive method to detect filtering of RPKI-invalid announcement focuses on networks that provide a direct and full feed to BGP collectors, which allows for definitive and detailed assessment of RPKI filtering of these networks. Our study is longitudinal, revealing a strong uptake in RPKI filtering deployment in recent years. Most importantly, however, we present a first-of-its-kind assessment of RPKI enforcement and its actual impact and benefit in situations in which the RPKI is needed the most: instances of conflicting prefix announcements in the global routing table.

2.2 RPKI and BGP Datasets

To study the visibility of RPKI-valid and RPKI-invalid announcements in the global routing table, we leverage the following datasets.

Longitudinal BGP dataset: To study long-term trends of RPKI filtering behavior, we download and process—using CAIDA BGPStream [25]—snapshots of the routing tables (RIB dumps) of all RouteViews and RIPE RIS collectors on the first day of each month³ from April 1, 2017 until January 22, 2020.

Fine-grained BGP dataset: To assess the visibility of RPKI-invalid announcements in detail, we process all the BGP updates generated over the month of September 2019 by RouteViews and RIPE RIS collector peers’ and we compute 5-minute snapshots of their routing tables using CAIDA BGPStream [25].

RPKI data: We take daily snapshots of validated Route Origin Authorizations (ROAs) for every day in September 2019, made available through the RIPE NCC RPKI validator [5]. For longitudinal analysis, we instead leverage the historical dataset of validated ROAs made publicly available by Chung *et al.* [11], selecting snapshots that align with our BGP dataset. A validated ROA consists of a prefix and the AS number authorized to originate that prefix in BGP according to cryptographically signed records in the RPKI. ROAs may include a *maxLen* attribute specifying up to which prefix length the de-aggregation of the ROA prefix is to be considered valid.

2.3 Preprocessing

From BGP snapshots to prefix-origin pairs: As a first step, we remove *bogon* prefixes from our BGP dataset, these include IETF reserved address space, and portions of address space not allocated by IANA to RIRs [3]. We further remove any IPv4 prefixes more specific than /24 or less specific than /8 (more specific than /64 or less specific than /8 for IPv6). Then we extract, for each BGP snapshot (both RIB dumps and those we derive from updates), all visible prefixes together with the advertised origin AS, obtaining *prefix-origin pairs*.⁴ For each prefix-origin pair, we save the set of *feeders*—that is, ASes that directly peer with any of the RouteViews and RIPE RIS route collectors—that have a route to the given prefix-origin in their routing table. In the following, we will leverage the set of feeders to assess filtering and to estimate visibility of prefix-origin pairs in the global routing table.

Tagging prefix-origin pairs: We next tag each individual prefix-origin pair in our dataset with its corresponding RPKI state. For each prefix-origin pair, we find the closest snapshot available of validated ROAs and tag the prefix-origin pair with one of the following states: (i) *unknown*: the prefix is not covered by any prefix of validated ROAs in the RPKI; (ii) *valid*: the prefix is covered by a validated ROA, the AS number in BGP matches the one in the ROA, and the prefix length in BGP is at most the *maxLen* attribute of the ROA; (iii) *invalid ASN*: the prefix is covered by a validated ROA, but the origin AS in BGP does not match the origin AS in any ROA covering the prefix; (iv) *invalid length*: the prefix is covered by a validated ROA, the origin AS in BGP matches the

³ Or the closest day for which validated historical RPKI data is available.

⁴ Note that a prefix can have multiple origins in the global routing table, in this case we extract multiple prefix-origin pairs.

origin AS in the ROA, but the prefix length in BGP is longer than the maxLen attribute, *i.e.*, the prefix is more specific than what is allowed as per the ROA.

3 To Filter or not to Filter: Longitudinal Study

In this section, we provide a macroscopic perspective on RPKI filtering deployment in today’s Internet. In particular, we study to which extent some of the transit networks in the Internet do filter BGP announcements with invalid RPKI state and how this filtering behavior evolved over time.

3.1 Detecting Filtering

While there is no practical way to comprehensively study filtering behavior of all networks, we introduce a method to infer RPKI filtering with high confidence for a small but relevant set of ASes. At a high-level, our method is made of two steps: *(i)* we select *full-feeder* ASes, *i.e.*, ASes that share with BGP collectors a number of routes (and thus prefix-origin pairs) comparable to what is globally visible in BGP—in other words, they tend to share the vast majority of, if not all, their preferred routes; *(ii)* we leverage our set of RPKI-invalid prefix-origin pairs to look for significant presence/absence of them in what full-feeders share.

The essence of this approach is to look for statistically significant absence of RPKI-invalid prefix-origin pairs: *e.g.*, the absence of a single invalid pair in the routes shared by a full-feeder is not a strong indication of RPKI-based filtering; similarly, the absence of a large number of invalid pairs in a shared routing table that is already missing many other valid routes (*i.e.*, from a *partial-feeder*) is not a strong indication of RPKI-based filtering either. The combination of the two factors instead, provides a high degree of confidence. In § 3.3, we validate our method for a few ASes that have publicly stated when they started applying RPKI-based filtering. In detail, we operate as follows.

(i) Selecting full-feeders: We consider a collector’s peer a *full-feeder* if the number of prefix-origin pairs shared by that AS is at least 75% of the maximum prefix-origin pair count sent by all feeders. We perform our analysis for IPv4 and IPv6 independently. In Figure 1a, the orange line shows this threshold for IPv4 in September 2019: out of 578 ASes peering with the collectors, we consider 276 to be full-feeders for IPv4 (232 for IPv6, see the Appendix). We chose 75%, since it separates recent and historical snapshots well.

(ii) Detecting filtering of RPKI-invalid announcements: From the set of full-feeder ASes, we infer an AS to be filtering RPKI-invalid announcements if the number of RPKI-invalid prefix-origin pairs received from that AS is less than 20% of the maximum number of invalid records sent by all full-feeders. Here, we leave some leeway, since previous research [26] has shown that, even if ASes are filtering *most* RPKI-invalid announcements, they usually never filter *all* invalid announcements due to churn in RPKI records and selective filtering (*cf.* § 3.3). The green dashed line in Figure 1a, shows this threshold for IPv4, we infer 21 ASes were filtering RPKI-invalids announcements in September 2019.

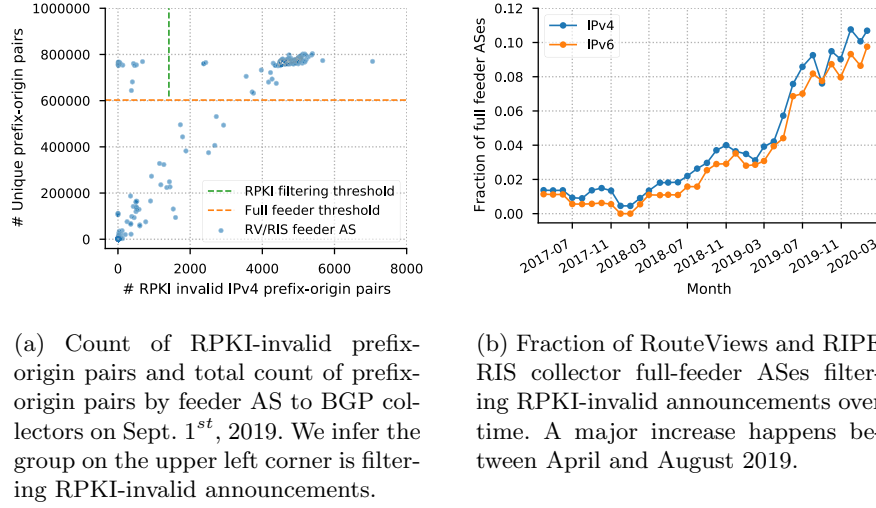


Fig. 1: Full-feeder ASes filtering of RPKI-invalid announcements.

The representativeness of our approach is limited by the comparably small number of full-feeder ASes: 290 ASes for IPv4 and 246 ASes for IPv6 in January 2020. However, we find that these networks include many global transit providers and mid-sized networks: 187 transit and access ASes (of which 12 are Tier-1 ASes), 36 content providers, and 47 educational/non-profit networks, according to PeeringDB [4]. In total there are 36 ASes in the top 100 CAIDA AS rank and 93 in the top 1,000. This set of ASes thus provides a reasonable approximation to study macroscopic filtering trends of major networks in the Internet.

3.2 Filtering Networks: Longitudinal Trends and Current Status

With our method in hand, we now present a longitudinal analysis of RPKI-invalid filtering behavior. Figure 1b shows the evolution of the fraction of full-feeder ASes that filter RPKI-invalid announcements for IPv4 and IPv6. Both protocols follow a similar trend, with slightly fewer ASes filtering RPKI-invalid IPv6 announcements compared to IPv4. We detect that in April 2017, less than 2% full-feeders were filtering RPKI invalid announcements: 3 out of 219 full-feeder ASes for IPv4 and 2 out of 176 for IPv6. We witness overall low levels of RPKI filtering until April 2018, when a few full-feeder ASes start to filter each month, reaching about 3% one year later in March 2019. From April until August 2019, we see a 3-fold increase in the rate of RPKI filtering adoption. In late January 2020, 11% of full-feeder ASes filter RPKI-invalid announcements in IPv4 and 10% in IPv6, 30 out of 290 and 23 out of 246 respectively.

The bulk of the networks filtering RPKI-invalid announcements are either transit or access network providers (17 ASes, 9% of such networks) or educational-research/non-profit networks (9 ASes, 19% of such networks). We find lower lev-




Fig. 2: RPKI-invalid IPv4 prefix-origin pairs from networks that publicly announced RPKI filtering deployment, vertical lines show the announcement date of deployment completion (dashed) or beginning of deployment (dotted).

els of filtering deployment in larger networks: only 2 of the 36 full-feeder ASes in the top 100 CAIDA AS Rank do filter invalid prefix-origins and 10 out of the 93 ASes in the top 1,000 CAIDA AS Rank filter. We only find one out of 36 content providers filtering invalid prefix-origins. RIPE, ARIN and APNIC are the regions with most full-feeder ASes, and we find 22, 5, 1 filtering ASes, representing 13%, 8% and 3% of full-feeders ASes from these regions respectively.

3.3 A closer Look at Filtering Networks

Comparison with public announcements of RPKI filtering: Five transit ISPs that provide direct and full BGP feeds to one of our considered collectors have publicly stated that they have deployed or are currently deploying RPKI-invalid filtering: AT&T (AS7018), KPN (AS286), Seacom (AS37100), Workonline Communications (AS37271) and Telia (AS1299) [1, 2, 6, 7]. Figure 2 shows the count of invalid prefix-origin pairs propagated by these five ASes during 2019, annotated with their public announcement date of filtering implementation.

In our data, we see over 4,000 invalid prefix-origins from all networks in early 2019. In mid February 2019, AT&T publicly stated that they started filtering RPKI-invalid route announcements and afterwards we only detect a few hundred invalid prefix-origins sent to collectors by AT&T. In early April 2019, two major African ISPs, Workonline Communication and Seacom, announced completion of deployment of RPKI filtering, after which we observe only several hundred invalid prefix-origins from these two networks. However, these ASes have encountered operational issues when deploying RPKI filtering and have (partially) stopped filtering for some periods of time, see intermittent upticks [22].

In late June 2019, KPN announced completion of deployment of RPKI-filtering and has only propagated a few dozen invalid prefix-origins to collectors since. Finally, in mid September 2019, Telia announced that it began to deploy RPKI filtering. Shortly after their announcement, we detect a continual decline

in the number of invalid prefix-origins forwarded by Telia. However, since RPKI enforcement deployment is not finished at the time of this writing, we still see over 2,000 invalid prefix-origins from Telia in early 2020, hence not meeting our detection thresholds yet. Our method detected RPKI-invalid filtering after all announcements of completion of full deployment of RPKI filtering.

Partial RPKI filtering: In our longitudinal study, no full-feeder network ever filters *all* RPKI-invalid announcements. Besides some expected short-term churn, *e.g.*, caused by delays when updating filtering rules, we identified 3 main reasons for persistent partial RPKI filtering: (i) selective RPKI Trust Anchor (TA) filtering: we find 6 networks not validating ROAs from the ARIN TA, resulting in a higher share of propagated invalid prefix-origins. Indeed, legal barriers limiting availability of ARIN ROAs have been reported [30]. (ii) Selective filtering depending on AS relationships: several network operators announced to implement filtering only for routes received from peers, but not customer networks [2]. (iii) Operational deployment issues: network operators reported compatibility issues with RPKI validator implementations and router software, prompting them to deploy RPKI-filtering in a subset of their border routers [22].

4 RPKI to the Rescue: Conflicting Announcements

Our findings of increasing deployment of RPKI filtering in the recent years motivate us to study the effect of filtering in more detail. We first introduce how we process our dataset to allow for analysis of visibility of individual routing events and study the overall visibility of valid/invalid prefixes. Next, we showcase several relevant real-world case studies of conflicting, and hence potentially malicious, prefix announcements. Visibility of a prefix in the global routing table translates directly into its *reachability*, and thus serves as a proxy to study the benefit of RPKI filtering in the wild. In this section, we present our findings for IPv4. Our findings for IPv6 are similar and can be found in the Appendix.

4.1 Tracking Visibility in the Global Routing Table

Aggregating prefix-origin snapshots into *timelines*: To study the visibility of RPKI-registered prefixes, we leverage our fine-grained BGP dataset, consisting of per-feeder snapshots of all prefix-origin pairs every 5 minutes in September 2019 (*cf.* § 2.2). As a first step, we aggregate adjacent prefix-origin pairs into continuous *timelines*. We require (i) that the maximum deviation in visibility within each timeline is less than 10%, otherwise we terminate the timeline and start a new one. We express visibility of a prefix-origin pair timeline as the fraction of active feeder ASes that propagate a route to given prefix and origin AS. Secondly, (ii) we require consistent RPKI state (valid/invalid ASN/invalid length/unknown) for each prefix-origin timeline.⁵ The resulting timelines consist of a tuple of a prefix, an origin AS, a visibility level, its RPKI state, and

⁵ For 0.37% IPv4 prefix-origin timelines, the RPKI state changed due to churn in the RPKI database caused by changes of RPKI entries during our measurement window. We remove these instances.

Prefix-origin timelines	count	%
IPv4 total	883,400	100%
RPKI covered	147,870	16.7%
RPKI-valid	139,537	15.8%
RPKI-invalid ASN	4,203	0.47%
RPKI-invalid length	4,130	0.46%




Table 1: Properties of our IPv4 prefix-origin timelines and their respective RPKI validity states (September 2019).

Fig. 3: CDF of IPv4 prefix-origin pairs by visibility during September 2019 for different RPKI states.

timestamps. We filter prefix-origin timelines with a private AS number or AS-Set as origin, and prefix-origin timelines with very low visibility, *i.e.*, seen by 3 or fewer peers, since such very low visibility prefixes are unlikely to represent actual events in the global routing table. Table 1 shows the properties of our resulting dataset.

Overall prefix-origin visibility by RPKI state: Figure 3 shows CDFs of the visibility of prefix-origin timelines, expressed as percentage of active feeder ASes seeing a prefix-origin. Overall, we find that RPKI-valid as well as RPKI-unknown prefix-origins (*i.e.*, prefixes not covered by validated ROAs) show similar visibility levels, with 80% of all prefix-origins seen by 80% or more of feeder ASes (see green and blue lines). RPKI invalid prefix-origins, however, show vastly different visibility: some 20% of these prefix-origins are very localized announcements (seen by less than 5% of feeder ASes, see orange and red lines), and we speculate that these cases are instances of misconfigurations, whether in BGP or RPKI records, which happen to also show up as RPKI-invalid artifacts. More importantly, we find that even invalid prefix-origins with higher visibility show distinctively lower visibility when compared to valid prefix-origins (see concentration of RPKI-invalid at around 70%, compared to over 80% for RPKI-valid). This difference in prefix-origin propagation is the direct result of filtering of RPKI-invalid announcements.

4.2 Conflicting Prefix Announcement Scenarios

Next, we study *RPKI in action*, *i.e.*, we want to understand if registration in the RPKI benefits networks in cases of conflicting announcements. In particular, we cover 3 scenarios: *(i)* Multiple Origin AS (MOAS) announcements: instances where two equal prefixes are announced with two different origins, often caused by intentional or unintentional prefix hijacks; *(ii)* subMOAS announcements: instances where an announcement of a more specific prefix points to a different origin AS, also a potential prefix hijack scenario; *(iii)* same-origin subprefixes,

instances where a more specific prefix is visible, points to the same origin AS as its parent, but fails RPKI validation due to max length restrictions. This scenario is what we would expect to see in the case of a path hijack, the most advanced form of prefix hijacks [27]. We note that in this work, we do not attempt to classify instances of conflicting prefix announcements into malicious activity vs. misconfigurations. Instead, we base our notion of illicit announcements on the RPKI state of the involved prefixes: if two prefix announcements are in conflict, and only one of them passes RPKI validation, in our analysis we treat the invalid one as if it is an illicit announcement (while it might also be due to incorrect/unissued ROAs). Our argument here is that, irrespective of the root cause of these conflicts, we can study the effectiveness of RPKI filtering under the same conditions that would also hold when a malicious actor injects BGP prefixes to hijack address space.

4.3 Visibility of Multiple Origin AS (MOAS) Prefixes

To study the visibility of prefixes that are concurrently originated by multiple origin ASes, we first isolate our prefix-origin timelines that show *(i)* two origin ASes for the same prefix and *(ii)* one of these prefix-origins is registered in the RPKI and valid. In total, we find about 90,000 instances of MOAS prefix-origin pairs in September 2019 for IPv4, of which some 10% are cases in which at least one prefix-origin is RPKI-valid, while others are not. Of these cases, about 20% ($N=1898$) are cases of exactly 2 MOAS prefix-origin pairs one valid and the other invalid according to RPKI records.

Figure 4 shows the distribution of the maximum visibility of prefix-origin timelines during MOAS conflicts of two prefix-origin pairs, where we partition RPKI-valid and -invalid state, see positive y -dimension in Figure 4. We see a stark difference: RPKI-valid prefixes clearly dominate visibility, with more than 70% of valid prefixes having visibility greater than 70%, and we only see few instances of RPKI-valid prefixes with low visibility (only 12% of instances with less than 30% visibility). Their invalid counterparts, on the other hand, show distinctively lower visibility: some 60% have a visibility level lower than 30%. Some invalid prefixes do reach substantial visibility levels, but we do point out that even those higher-visibility invalid prefixes cluster at around $\approx 65\%$, that is, significantly lower when compared to valid prefixes, which cluster at around $\approx 80\%$. These results are consistent with our expectations: the RPKI benefit should be significant in instances of exact MOAS conflicts, since two prefixes compete for reachability in the global routing table, and even when RPKI filtering is not enforced, some routers still give preference to RPKI-valid announcements over RPKI-invalid announcements as part of the route selection process (discarding an invalid route only if a valid one is available) [12].

To assess the potential benefit of registering a prefix in the RPKI vs. not registering it, we next compare the above studied instances of MOAS conflicts in which the concerning prefix is registered in the RPKI against vanilla cases of MOAS, in which the concerning prefix is not registered, and hence both prefix-origins are of type RPKI-unknown. Here, in the absence of RPKI information, we




Fig. 4: Visibility of prefix-origin pairs during MOAS conflicts: RPKI-valid and invalid ASN MOAS prefix pairs in the positive y -dimension, RPKI-unknown MOAS prefix pairs in the negative y -dimension, partitioned as stable/unstable according to total advertisement time during September 2019.

face the difficult problem of determining which of the conflicting announcements represents the legitimate announcement vs. the illicit one. Taking a pragmatic approach, we leverage stability of announcements as a proxy: In the case of a MOAS conflict where neither prefix-origin is registered in the RPKI, we tag the prefix-origin that was visible for a longer period of time as *stable*, and the conflicting prefix-origin that was visible for a shorter period of time as *unstable*. We pick only MOAS cases where the stable prefix-origin is announced for a period at least 3 times longer⁶ than the unstable prefix-origin counterpart (N=6,374 MOAS events for IPv4). We acknowledge that this heuristic is not a hard-and-fast rule, since there are many potential root causes for unstable announcements (*e.g.*, rewiring, address space transfers, etc.) but it allows us to present a one-to-one comparison of RPKI vs. non-RPKI scenarios.

We plot the distribution of prefix-origin visibility of RPKI-unknown prefixes in the negative y -dimension in Figure 4. We find that, overall, stable prefixes show much greater visibility when a MOAS conflict happens, when compared to their conflicting unstable counterparts. However, contrasting the vanilla case (no RPKI registration, negative y -dimension) against the case in which the prefix is registered in the RPKI (positive y -dimension), we see a difference: unstable RPKI-unknown prefixes generally reach higher levels of visibility when compared to RPKI-invalid prefixes. This difference manifests both for very low visibility cases, where RPKI-unknown cluster at around $\approx 15\%$ visibility, higher than their RPKI-invalid counterparts which cluster at $\approx 8\%$, as well as for cases of higher

⁶ We tested different thresholds, finding that the modes of the distribution do not change much.

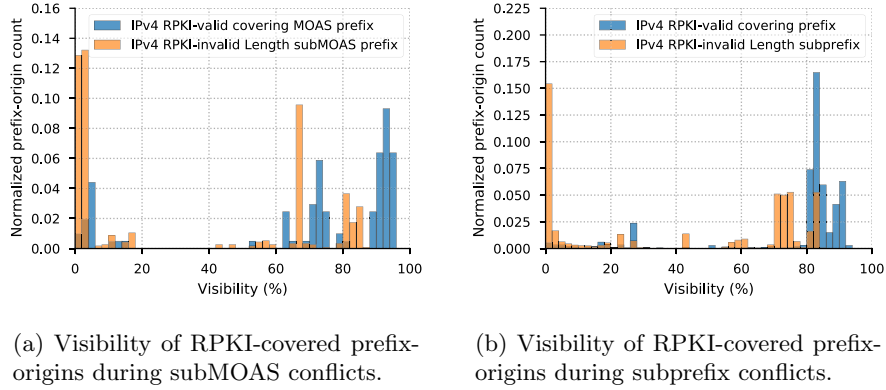


Fig. 5: Impact of RPKI registration in subMOAS and subprefix conflicts.

visibility: unstable RPKI-unknown prefixes reach visibility levels of some 70%, while RPKI-invalid cluster below 60%. Indeed, less than 14% of RPKI-invalid MOAS instances reach a visibility over 60% compared to 37% for unstable RPKI-unknown MOAS instances. RPKI registration shows a clear effect on prefix visibility when MOAS conflicts happen.

4.4 Visibility of Subprefix Announcements

We next study instances of subprefix announcements, which instead do not compete with the covering prefix for visibility, since routers use longest-prefix matching, preferring more-specific routes for packet forwarding. For this reason, subprefix announcements can be a powerful and dangerous tool to, *e.g.*, hijack address space and redirect traffic, and their effect has been also evidenced in large-scale routing incidents, including route leaks [23, 28].

To study the impact of RPKI registration on subprefix announcements, we first isolate all incidents of subprefix announcements in our dataset, *i.e.*, we observe a covering (that is, less specific) prefix, covered by a validated ROA in the RPKI, and concurrently a more specific prefix announcement that does not pass RPKI validation—either because of an invalid ASN (subMOAS) or invalid prefix length (subprefix). In total, we find 10,450 instances of RPKI-invalid subprefix and subMOAS announcements in IPv4, conflicting with 2,291 RPKI-valid covering prefixes. Figure 5a shows the distribution of prefix visibility in the case of subMOAS: if a more-specific prefix announcement fails RPKI validation because it has a different origin AS ($N=5,401$ subMOAS prefixes, $N=966$ covering prefixes). While the RPKI-valid covering prefixes show high visibility, their invalid counterpart, subMOAS prefixes, show two modes of visibility: some 35% of invalid subMOAS show very low visibility, *i.e.*, lower than 10%. More importantly, however, is the finding that none of the subMOAS prefixes reach the same visibility level as their valid parent: while subMOAS prefixes barely exceed

75% visibility, their valid covering prefixes typically reach some 85% - 90% visibility and 75% reach at least 80% visibility. These observations are consistent with our earlier findings of increasing RPKI filtering, and highlight that RPKI registration also benefits registrants in the case of difficult-to-combat subMOAS situations.

Figure 5b shows the visibility for invalid-length subprefix announcements having the same origin AS as their covering RPKI-valid counterpart (N=5,049 subprefix, N=1,325 covering prefixes). Recall that the RPKI permits to specify a `maxLength` attribute, limiting the prefix length of any prefix matching the RPKI record, irrespective of the origin AS. Besides cases of misconfiguration, this scenario also applies in the case of *path hijacks*: instances where an attacker injects a subprefix that allegedly points to the same origin AS as its valid covering prefix, but in fact the attacker redirects traffic to its network. Such attacks can, *e.g.*, be carried out by prepending the valid origin AS at the end of the path after the hijacker’s AS number. Such path hijacks present advanced forms of prefix hijacks and are difficult to detect using today’s methods [27]. In Figure 5b, we see similarly lowered levels of visibility for RPKI-invalid subprefix announcements, even if they point to the registered origin AS. Invalid announcements reach some 70% of visibility, substantially lower when compared to their valid covering prefix. These results show that RPKI registration can benefit networks even in this most advanced case of illicit announcements: subprefix path hijacks.

5 Discussion and Conclusion

Recent research has shown increasing registration in the RPKI by networks around the globe. Our work complements these observations, adding an important dimension: RPKI enforcement. We find that a substantial, and growing, number of ISPs in the Internet begin to filter invalid RPKI announcements, including major players such as AT&T. Increasing RPKI enforcement starts to bring direct value to networks, since registration in the RPKI benefits them in real-world scenarios, such as prefix hijacks. Our findings show that already as of today, registration in the RPKI limits the propagation of illicit announcements, in MOAS conflicts as well as in instances of subMOAS and subprefix announcements. Evidence of direct value for networks could incentivize even more transit providers to deploy RPKI filtering to benefit their customers. While the RPKI protects its registrants in the case of such illicit announcements, we can also expect that increasing RPKI enforcement provides further incentives for networks to keep their RPKI records up-to-date, since stale records and other misconfigurations will have a direct impact on reachability of the respective address blocks. Our method provides a simple way to track current levels of RPKI filtering and to study its impact on illicit prefix announcements. Continuous monitoring of deployment of filtering allows for more transparency in the process, and empirical evidence of benefits of registration provides further incentives for network operators to join the growing group of networks that protect their prefixes by registering them in the RPKI.

Acknowledgments

We thank the anonymous reviewers for their thoughtful feedback. This work was partially supported by the MIT Internet Policy Research Initiative, William and Flora Hewlett Foundation grant 2014-1601. We acknowledge funding support from the NSF Grants CNS 1705024 and OAC 1724853. This material is based on research sponsored by Air Force Research Laboratory under agreement number FA8750-18-2-0049. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions in this paper are those of the authors and do not necessarily reflect the opinions of a sponsor, Air Force Research Laboratory or the U.S. Government.

References

1. AS286 Routing Policy. <https://as286.net/AS286-routing-policy.html>
2. AT&T/as7018 now drops invalid prefixes from peers. <https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html>
3. Cymru BGP Bogon Refence. <http://www.team-cymru.org/bogon-reference-bgp.html>
4. PeeringDB. <https://www.peeringdb.com>
5. RIPE NCC RPKI Validator. <https://rpki-validator.ripe.net/>
6. RPKI Route Origin Validation - Africa. <https://mailman.nanog.org/pipermail/nanog/2019-April/100445.html>
7. Telia Carrier Takes Major Step to Improve the Integrity of the Internet Core. <https://www.teliacarrier.com/Press-room/Press-releases/Telia-Carrier-Takes-Major-Step-to-Improve-the-Integrity-of-the-Internet-Core-.html>
8. The Hunt for 3ve: Taking down a major ad fraud operation through industry collaboration. Tech. rep. (Nov 2018), https://services.google.com/fh/files/blogs/3ve_google_whiteops_whitepaper_final_nov_2018.pdf?_hstc=&_hssc=&hsCtaTracking=c7b87c5c-1676-4d53-99fb-927a07720b17%7C9d63bf77-0926-4d08-b5ec-46b1a06846bc
9. Bush, R., Austein, R.: The Resource Public Key Infrastructure (RPKI) to Router Protocol. RFC 6810 (Proposed Standard) (Jan 2013), <https://www.rfc-editor.org/rfc/rfc6810.txt>, updated by RFC 8210
10. Cartwright-Cox, B.: The year of RPKI on the control plane. <https://blog.benjojo.co.uk/post/the-year-of-rpki-on-the-control-plane> (September 2019)
11. Chung, T., Aben, E., Bruijnzeels, T., Chandrasekaran, B., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Rijswijk-Deij, R.v., Rula, J., Sullivan, N.: RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins. In: Proceedings of the Internet Measurement Conference. pp. 406–419. IMC '19, Association for Computing Machinery, Amsterdam, Netherlands (Oct 2019), <https://doi.org/10.1145/3355369.3355596>
12. Cisco: IP Routing: BGP Configuration Guide, Cisco IOS XE Release 3S. https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xs-3s/irg-xe-3s-book/bgp-origin-as-validation.html

13. Gilad, Y., Cohen, A., Herzberg, A., Schapira, M., Shulman, H.: Are We There Yet? On RPKI's Deployment and Security. In: Proceedings 2017 Network and Distributed System Security Symposium. Internet Society, San Diego, CA (2017)
14. Goodin, D.: Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency (Apr 2018), <https://arstechnica.com/information-technology/2018/04/suspicious-event-hijacks-amazon-traffic-for-2-hours-steals-cryptocurrency/>
15. Huston, G., Michaelson, G., Loomans, R.: A Profile for X.509 PKIX Resource Certificates. RFC 6487 (Proposed Standard) (Feb 2012), <https://www.rfc-editor.org/rfc/rfc6487.txt>, updated by RFCs 7318, 8209
16. Huston, G., Michaelson, G., Martinez, C., Bruijnzeels, T., Newton, A., Shaw, D.: Resource Public Key Infrastructure (RPKI) Validation Reconsidered. RFC 8360 (Proposed Standard) (Apr 2018), <https://www.rfc-editor.org/rfc/rfc8360.txt>
17. Huston, G., Michaelson, G.: RFC 6483: Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs) (Feb 2012), <https://tools.ietf.org/html/rfc6483>
18. Iamartino, D., Pelsser, C., Bush, R.: Measuring BGP Route Origin Registration and Validation. In: Mirkovic, J., Liu, Y. (eds.) Passive and Active Measurement, vol. 8995, pp. 28–40. Springer International Publishing, Cham (2015), http://link.springer.com/10.1007/978-3-319-15509-8_3
19. Kent, S., Kong, D., Seo, K., Watro, R.: Certificate Policy (CP) for the Resource Public Key Infrastructure (RPKI). RFC 6484 (Best Current Practice) (Feb 2012), <https://www.rfc-editor.org/rfc/rfc6484.txt>
20. Lepinski, M., Kent, S.: An Infrastructure to Support Secure Internet Routing. RFC 6480 (Informational) (Feb 2012), <https://www.rfc-editor.org/rfc/rfc6480.txt>
21. Lepinski, M., Kent, S., Kong, D.: A Profile for Route Origin Authorizations (ROAs). RFC 6482 (Proposed Standard) (Feb 2012), <https://www.rfc-editor.org/rfc/rfc6482.txt>
22. Maddison, B.: RIPE Forum - Routing Working Group - RPKI Route Origin Validation - Africa (Apr 2019), <https://www.ripe.net/participate/mail/forum/routing-wg/PDZ1MzAzMzhhlWVhOTAtnZIxOC11MzIOLTBjZjMyOGI1Y2NkMOBzZWVjb20ubXU+>
23. Newman, L.H.: Why Google Internet Traffic Rerouted Through China and Russia. Wired (Nov 2018), <https://www.wired.com/story/google-internet-traffic-china-russia-rerouted/>
24. Newton, A., Huston, G.: Policy Qualifiers in Resource Public Key Infrastructure (RPKI) Certificates. RFC 7318 (Proposed Standard) (Jul 2014), <https://www.rfc-editor.org/rfc/rfc7318.txt>
25. Orsini, C., King, A., Giordano, D., Giotsas, V., Dainotti, A.: BGPStream: A Software Framework for Live and Historical BGP Data Analysis. In: Proceedings of the 2016 Internet Measurement Conference. pp. 429–444. IMC '16, Association for Computing Machinery, Santa Monica, California, USA (Nov 2016), <https://doi.org/10.1145/2987443.2987482>
26. Reuter, A., Bush, R., Cunha, I., Katz-Bassett, E., Schmidt, T.C., Waehlich, M.: Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering. ACM SIGCOMM Computer Communication Review **48**(1), 9 (2018)

27. Sermpezis, P., Kotronis, V., Gigis, P., Dimitropoulos, X., Cicalese, D., King, A., Dainotti, A.: ARTEMIS: Neutralizing BGP Hijacking within a Minute. arXiv:1801.01085 [cs] (Jan 2018), <http://arxiv.org/abs/1801.01085>
28. Strickx, T.: How Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline Today (Jun 2019), <https://blog.cloudflare.com/how-verizon-and-a-bgp-optimizer-knocked-large-parts-of-the-internet-offline-today/>
29. Testart, C., Richter, P., King, A., Dainotti, A., Clark, D.: Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. In: Proceedings of the Internet Measurement Conference on - IMC '19. pp. 420–434. ACM Press, Amsterdam, Netherlands (2019), <https://dl.acm.org/doi/10.1145/3355369.3355581>
30. Yoo, C., Wishnick, D.: Lowering Legal Barriers to RPKI Adoption. Faculty Scholarship at Penn Law (Jan 2019), https://scholarship.law.upenn.edu/faculty_scholarship/2035

Appendix: IPv6 results




Fig. 6: Count of RPKI-invalid IPv6 prefix-origin pairs and total count of prefix-origin pairs by feeder AS to BGP collectors on Sept. 1st, 2019.




Fig. 7: CDF of IPv6 prefix-origin pairs by visibility during September 2019 for different RPKI states.

Detecting RPKI-filtering in IPv6: We apply the method described in 3.1, setting equivalent thresholds to those used for IPv4. In September 2019, out of 402 ASes peering with collectors for IPv6, we consider 232 to be full-feeders, and of those 232 we infer 18 are filtering RPKI-invalid announcements.

Tracking visibility in the global IPv6 routing table: Using the methodology described in 4.1, we build prefix-origin timelines for IPv6 prefixes⁷. Table 2 shows the properties of our resulting dataset.

Overall IPv6 prefix-origin visibility by RPKI state: Figure 7 shows CDFs of the visibility of prefix-origin timelines, which show very similar behavior to the

⁷ 0.13% of IPv6 prefix-origin timelines whose RPKI state changed during our measurement window were removed.

Prefix-origin timelines	count	%
IPv6 Total	91,313	100%
RPKI covered	19,173	20.1%
RPKI-valid	17,656	19.3%
RPKI-invalid ASN	362	0.40%
RPKI-invalid length	1155	1.26%

Table 2: Properties of our IPv6 prefix-origin timelines and their respective RPKI validity state.





Fig. 8: Visibility of RPKI covered IPv6 prefix-origin pairs during MOAS conflicts.

ones described in 4.1 for IPv4. In IPv6, there are even fewer RPKI-valid prefix-origins with low visibility compared to IPv4: less than 10% IPv6 prefix-origins have less than 80% visibility compared to 20% for IPv4.


Visibility of Multiple Origin AS (MOAS) IPv6 prefixes: In total, we find about 41,000 instances of MOAS prefix-origin pairs in September 2019 for IPv6, of which some 133 are cases in which at least one prefix-origin is RPKI-valid while others are not. Figure 8 shows the distribution of the maximum visibility of prefix-origin timelines during MOAS conflicts.

Visibility of IPv6 subprefix announcements: We find 575 subMOAS prefix conflicting with 102 covering prefixes (figure 9a) and 1,903 subprefixes conflicting with 235 covering prefixes (figure 9b).

Issuing RPKI records for IPv6 prefixes also benefit networks in the case of conflicting (and potentially malicious) announcements.



(a) Visibility of RPKI-covered IPv6 prefix-origins during subMOAS conflicts.



(b) Visibility of RPKI-covered IPv6 prefix-origins during subprefix conflicts.

Fig. 9: Impact of RPKI registration in subMOAS and subprefix conflicts.