

Revisiting RPKI Route Origin Validation on the Data Plane

Nils Rodday^{*†}, Ítalo Cunha[‡], Randy Bush[§], Ethan Katz-Bassett^{††}
Gabi Dreo Rodosek^{*}, Thomas C. Schmidt^{||}, Matthias Wählisch^{**}

^{*}Research Institute CODE, Universität der Bundeswehr München,

[†]University of Twente, [‡]Universidade Federal de Minas Gerais,

[§]Arrcus / IJ, ^{††}USC / Columbia University,

^{||}Hamburg University of Applied Sciences, ^{**}Freie Universität Berlin

Abstract—The adoption of the Resource Public Key Infrastructure (RPKI) is increasing, as are measurement activities to identify RPKI-based route origin validation (ROV). Several proposals try to identify Autonomous Systems (ASes) that deploy ROV using control plane as well as data plane measurements. We show why simple end-to-end measurements may lead to incorrect identification of ROV. In this paper we evaluate data plane traceroute measurements as a mechanism to extend coverage and provide a reproducible method for ROV identification using RIPE Atlas. Moreover, we extend the current state-of-the-art by identifying ROV performed by route servers at Internet Exchange Point (IXP) and using an include list to differentiate between fully and partially ROV-enforcing ASes. Our measurements from 5537 vantage points in 3694 ASes infer ROV is deployed in 206 unique ASes: 10 with strong confidence, 12 with weak confidence, and 184 indirectly adopting ROV via filtering by IXP route servers.

Index Terms—RPKI, ROV, RIPE Atlas, data plane measurements

I. INTRODUCTION

BGP hijacking is a persisting threat [1], [2] and route leaks are a recurring operational incident that have led to the development of several security add-ons for the Border Gateway Protocol (BGP). Currently, the RPKI [3] is rolled out by operators. Operators can create Route Origin Authorization (ROA) objects [4] to cryptographically prove legitimate BGP announcements, deploy Route Origin Validation (ROV) to reject or deprefer announcements that violate a ROA [5], or both. Two major directions of RPKI research are currently discussed in the community: (i) Identifying the address space covered by ROA objects [6]–[9] and (ii) measuring ROV deployment in the wild [10]–[12].

Measuring ROAs is less challenging as ROA data is publicly provided by RPKI repositories. Combined with public BGP dumps, a fair approximation of the global routing state can be provided. On June 16, 2021, RouteViews indicated 30.06% of prefixes were valid, 0.75% were invalid, and 69.19% did not have a covering ROA. Up-to-date results can be found at the NIST RPKI deployment monitor [13].

Measuring ROV deployment is more challenging as this requires the inference of (private) router configuration changes.

The ultimate goal is to infer which ASes are actually using the RPKI to drop invalid route announcements. Common measurement setups are based on passively collected data (*e.g.*, BGP dumps) or active experiments to observe routing divergence between paths towards valid and invalid prefix announcements of the same origin AS. Active experiments are conducted on the control plane, data plane, or a combination of both. It has been shown that uncontrolled (passive) measurements solely relying on control plane information incorrectly identify ROV enabled ASes [10]. Instead, controlled measurements are preferred because they limit the amount of independent variables by introducing well-defined ROA and BGP events.

Any study of the control plane is limited by the number of vantage points exporting data to public BGP collectors. This becomes even more crucial in the context of ROV measurements because they require path-specific analysis and ROV deployment is still limited. In this paper, we investigate the option to conduct ROV measurements on the data plane to extend control plane coverage. We perform controlled experiments by announcing our own RPKI valid and invalid prefixes in BGP via the PEERING testbed [14] and utilize the RIPE Atlas platform [15] to issue HTTP, ICMP, and traceroute measurements toward these IP prefixes. We show that HTTP and ICMP, which only measure end-to-end connectivity, provide insufficient information to accurately infer ROV deployment. We use additional information provided by traceroute, *i.e.*, the set of ASes traversed on routes toward valid and invalid prefixes, to more accurately infer ROV-enforcing ASes.

We classify our measurements into six cases. The first three cases consider the *connected assumption* of previous work [10], which requires every tested AS to be directly connected to the origin AS. This allows for *strong inferences*. The other three extend the range of our measurements by relaxing the *connected assumption*, which allows inference of ROV at ASes not directly connected to the origin AS provided sufficient route visibility. This allows for *weak inferences*. We also inspect traces for IXP traversals to differentiate between direct ROV filtering at an AS and indirect ROV filtering performed by route servers at an IXP. Additionally, an include list is used in order to reveal fully and partially filtering ASes.

Table I
COMPARISON OF RELATED WORK SORTED BY YEAR

Reference	Year	Plane	Approach
Gilad et al. [11]	2017	Control	Uncontrolled
Reuter et al. [10]	2018	Control	Controlled
Hlavacek et al. [12]	2018	Control+Data	Traceroute & TCP SYN
Cartwright-Cox [16]	2019	Data	ICMP scans
RPKI WebTest [17]	2019	Data	HTTP
Testart et al. [18]	2020	Control	Statistical approach
Huston et al. [19]	2020	Data	HTTP

We make 10 strong inferences, 12 weak inferences, and 184 route server inferences. In total, we identify 206 unique ASes performing ROV of which 146 are fully and 60 are partially filtering.

Contributions. In this work, we focus on controlled data plane measurements to identify ASes that use RPKI Route Origin Validation and make the following contributions:

- 1) We show why simple end-to-end HTTP measurements falsely attribute ROV deployment to ASes under test if transits in between are filtering.
- 2) We develop a new data plane methodology based on controlled measurements using RIPE Atlas that makes strong inferences for ASes directly peering with our announcement sites (1 AS hop) and weak inferences for longer paths (2+ AS hops).
- 3) We take IXP traversals into account and build an include list for ASes seen on invalid paths to differentiate between partially and fully filtering ASes.
- 4) We present up-to-date results and confirm that deployment has increased.

The remainder of this paper is structured as follows: Section II surveys current ROV measurement approaches. Section III discusses fundamental drawbacks of end-to-end measurements to infer ROV. Sections IV and V present our method and results. Section VI reports on lessons learned. Section VII summarizes our findings.

II. CURRENT ROV MEASUREMENT APPROACHES

RPKI measurement methodologies can be coarsely divided in two areas: control plane and data plane measurements. Table I summarizes prior work.

Control plane measurements. In 2017, Gilad et al. [11] measured the ROV adoption rate using passive control plane measurements. They first seek an AS that is originating both an RPKI invalid and a non-invalid (*i.e.*, not found or valid) BGP advertisement. Next, they check whether there is only one transit AS between the propagating AS and the BGP collector. They classify this transit AS as ROV-enforcing if (*i*) the AS is forwarding the RPKI non-invalid route announcements but drops the invalid ones, and (*ii*) this behavior is observed for three different destination ASes. They find that three out of the top 100 ASes on the Internet are performing ROV. An additional survey among operators showed that 84.09% are not

using ROV while 10.23% are assigning a lower preference to invalid announcements and 5.68% are dropping invalids.

Reuter et al. [10] reproduced the work by Gilad et al. [11] and found that the results heavily relied on the chosen set of BGP collectors. They argue that such measurements are *uncontrolled* and propose *controlled* measurements. Instead of merely relying on passive measurements and analyzing existing BGP data, they perform active measurements by announcing their own prefix ranges and controlling the ROA states. This reduces the amount of independent variables present in the setup. The method identified three ASes that were deploying ROV, which were confirmed by the AS operators. The method is deployed in a live monitoring system [20] and identifies 118 ASes as deploying ROV in March 2021.

Testart et al. [18] introduce a passive approximation methodology of how ROV could be measured on the control plane. The methodology aims for identifying statistical anomalies in BGP collector data. First, they extract a set of ASes called full-feeders that report the majority of publicly visible routes to the collectors. Second, they try to find ASes reporting significantly less RPKI invalid routes compared to the full-feeders. The resulting cluster consists of 21 ASes that are identified as filtering. Validation of results is limited to 5 ASes that have publicly been reported to deploy ROV. Overall, the paper identifies the trend of increasing RPKI usage on the Internet.

Gray et al. [21] propose BeCAUSE, an algorithmic framework for inferring network properties based on Bayesian computation for ASes. They apply BeCAUSE to pinpoint ROV-enabled ASes.

Data plane measurements. In 2018, Cartwright-Cox [16] presented a novel approach of how RPKI adoption could be measured on the *data plane*. By performing ICMP scans of the entire IPv4 address space, he identifies responsive hosts. The set of responsive hosts is queried again but from RPKI invalid address space, tracking the amount of replies received. If hosts reply to a control server within RPKI valid address space but not to a control server within RPKI invalid address space, the lack of reachability is attributed to ROV-based filtering. The methodology, however, does not allow for identification of the filtering AS on the path as any transit network may be filtering. Section III elaborates in more detail on this. Updates to this study were presented at NLNOG Day, in September 2019, and at RIPE 80, in May 2020 [22].

A similar method has been applied by Huston et al. [19] to explore protection of end devices by RPKI. They use an experiment prefix that is being swapped after a 36 hours valid period to a 12 hours invalid period and assign an IP address from this address range to an HTTP server. They query this server via HTTP from end-user hosts. Differences in reachability are attributed to ROV. In contrast to the study by Cartwright-Cox [16], Huston et al. [19] do not aim at identifying filtering ASes but instead at determining the share of protected end users. The study reports $\sim 17\%$ of end-users being protected by RPKI filtering and also points out that most

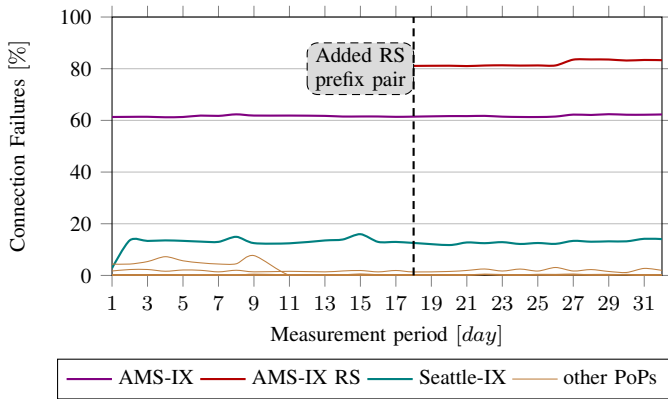


Figure 1. Comparison of the reachability of HTTP measurements initiated from RIPE Atlas probes to RPKI valid and invalid routes, which were announced at different locations in January 15–February 15, 2020. Connectivity is impacted by ROV-filtering at AMS-IX and, to a lesser extent, at Seattle-IX.

probably a few transit providers enabled filtering rather than many stub networks.

The RPKI WebTest [17] operated by RIPE is a website to raise awareness of ISPs that do not deploy ROV. It tests whether the local Internet Service Provider (ISP) of an end host drops RPKI invalid routes or not. IP addresses from two static /24 IP prefixes, one RPKI valid and one invalid, are assigned to a web server. Two HTTP requests are sent from the web browser of a user to each IP address. If both requests succeed, the ISP is considered not rejecting invalids yet. If the IP address from the invalid prefix range could not be reached but the IP address from the valid range, the ISP is considered deploying ROV.

Hlavacek et al. [12] compare current ROV measurement methodologies. They repeat the controlled *control plane* measurements of [10] and argue for data plane measurements because of higher accuracy in their setup. For *data plane* measurements, the authors use both traceroutes via RIPE Atlas and TCP-SYN packets sent to the top 1,25M Alexa domains. They find 4 ASes ROV enforcing based on control plane measurements and 12 ASes based on RIPE Atlas traceroutes. Analyzing the lack of TCP replies they find 201 TCP endpoints protected. This work is most closely related to our work. We extend prior work by (i) relaxing the connected assumption of [10] to increase coverage, (ii) deploy dedicated prefixes announced only to route servers of IXPs and use TraIXroute to identify IXPs, and (iii) build an include list that allows to differentiate between partially and fully filtering ASes.

III. INCORRECT ATTRIBUTION OF END-TO-END MEASUREMENT METHODOLOGIES

Some existing methodologies measuring ROV employ end-to-end measurements. Examples are ICMP scans performed by Cartwright-Cox [16] and HTTP measurements done by the RPKI WebTest [17]. These methods only allow to infer binary connectivity results, *i.e.*, whether a probe could reach its target or not; they do not allow for pinpointing ASes that deploy ROV.

To reproduce the RPKI WebTest study, we announce two distinct /24 prefixes from each of the 11 PEERING Points of Presence (PoPs) [14] to all adjacent ASes (directly and via routeservers), leading to an RPKI valid and an invalid route per PoP. Prefixes of each PoP are not overlapping. We utilize RIPE Atlas to perform HTTP measurements once a day to each prefix over a period of one month. If the prefix of the valid route can be reached but the prefix of the invalid route cannot, ROV is attributed to the tested AS, similar to [17].

Figure 1 shows the ratio of the number of failed HTTP requests to prefixes of invalid routes and the number of successful HTTP requests to prefixes of valid routes. For prefixes that we announce to all peers at AMS-IX, we observe connection failures of about 60%. It is known that AMS-IX deploys RPKI filtering at routeservers based on an opt-out policy since October 20, 2017 [10]. Therefore, many RIPE probes cannot reach the RPKI invalid route. On day 17, we announced an additional prefix pair only via the AMS-IX route server. For this prefix pair we observe that 80% of AMS-IX members do not opt-out of ROV implemented by routeservers and thus cannot reach the invalid route (see red curve in Figure 1).

Narrowing down ROV-inferences for single ASes, we are able to identify five probes that are located in New Zealand. All reached the prefixes of the valid routes but not a single prefix of the invalid routes. Manual investigation revealed that all probes resided within ASes that used AS38022 REANNZ-NZ-A as an upstream. We confirmed our results with the operator of AS38022, who acknowledged that ROV is enabled and invalid routes are being rejected.

We also attempted to repeat the ICMP scans presented in [16]. Similar to the previous reproduction, we announced valid and invalid routes from different PEERING PoPs during different runs. Our results show that filtering transits have significant impact on the outcome of this methodology. When the filtering transit was closer to the PEERING testbed along the AS path (*e.g.*, the upstream) the share of ASes labeled as RPKI ROV enforcing greatly increased. Moreover, results change depending on the PEERING PoP used. Reverse Path Filtering (RPF) also impacts results as filtering upstreams drop data plane packets already on the forward path. The methodology proposed in [16] assumes a reply is dropped on the reverse path, which introduced false positives.

Overall, simple end-to-end methodologies flag ASes as ROV-enforcing, while the upstream might perform the filtering. The obtained attribution does therefore not reliably allow to pinpoint the filtering AS.

IV. MEASUREMENT METHOD AND SETUP

In order to pinpoint ASes that are performing ROV, we propose to consider the path between two hosts hop-wise. To this end, we conduct extensive and reproducible *data plane* measurements using traceroute from RIPE Atlas [15]. To shape the underlying topology, we use controlled, active experiments




Figure 2. Measurement Phases

Table II
POINTS OF PRESENCE AT THE PEERING TESTBED.

Name	Upstream/IXP	# Direct Peers
ams01	AMS-IX*	123
gatech01	Georgia Institute of Technology	1
grnet01	GRNet	1
seattle01	Seattle-IX*	72
uw01	University of Washington	1

* Our AS is also directly connected to the route server of the IXP.

based on the PEERING testbed [14] and a delegated RPKI deployment model. Our study can be divided in four phases, see Figure 2.

Preparation Phase. This phase consists of static BGP announcements and publishing related ROAs at a predefined schedule. We select six PoPs from the PEERING testbed, see Table II. At each PoP, we announce two distinct /24 IP prefixes to all peers, one serves as anchor and the other as experiment prefix, similar to [10]. Additionally, at AMS-IX and Seattle-IX, we announce a distinct set of anchor and experiment prefixes only to the route servers, which will allow for a more fine-grained analysis of RPKI usage at route servers at a later stage. In total, we utilize $14 \times /24$ IP prefixes.

We publish ROAs at our own child Certificate Authority (CA) such that the BGP announcement of the anchor prefix is always RPKI valid while the BGP announcement of the experiment prefix alternates between RPKI valid and invalid every 24 hours. The schedule is illustrated in Figure 3, which was recorded by RIPEstat [23]. It clearly indicates the continuously high reachability of the 147.28.12.0/24 anchor prefix (green) and a reduced reachability of the 147.28.13.0/24 experiment prefix when routes are invalid (red). Our ROAs swap at midnight, but route convergence introduces a small delay before the change becomes visible at route collectors.

Measurement Phase. We randomly choose up to 3 RIPE Atlas probes per AS to perform data plane measurements. We do not select any probes contained in the dataset from [24], since those probes were identified to be located behind middleboxes, which would lead to incorrect results. Overall, we




Figure 3. RIPEstat Routing History for one of our prefix pairs. Visibility of the anchor prefix remains stable while visibility of the experiment prefix drops significantly when ROA configuration leads to invalid routes.

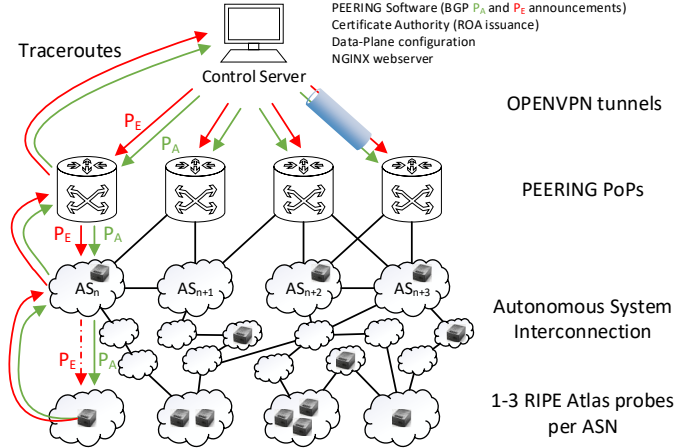


Figure 4. Experiment setup. The control server connects via OpenVPN tunnels to the PEERING PoPs, which announce the prefix ranges. RIPE probes in ASes send traceroutes towards our control server.

choose 5537 probes in 3694 ASes (*i.e.*, on average 1.49 probes per AS). From each probe, we send traceroute measurements to our *anchor* and *experiment* prefixes announced from each PoP. One measurement run needs two days to complete since ROAs are flipped every 24 hours. We conduct 5 experiment runs, from July 2, 2021 to July 19, 2021. All experiments are scheduled using the RIPE Atlas APIv2 in combination with the Cousteau Library [25]. Figure 4 illustrates the experiment setup.

Preprocessing Phase. Once measurement results are downloaded from the RIPE Atlas platform, we perform additional preprocessing in order to make them usable for our ROV identification methodology. First, we translate IP paths obtained from traceroute into Autonomous System Number (ASN) paths. Second, we identify traces that cross IXPs.

We extracted all 38.5K IP addresses from one measurement run and mapped those IP addresses to ASNs as follows: First, we identified 6.5K private IP addresses, which are not mapped further. Second, we identified 22 IP addresses that belong to the PEERING testbed. The remaining 32.4K IP addresses were mapped to ASNs using Team Cymru, which resolved 30.3K addresses.¹ Finally, we utilized PyASN with up-to-date BGP collector dumps from RIPE RIS [26] and Routeviews [27] to further enrich the dataset but could only resolve 0.2K IP addresses additionally.² The remaining 1.8K addresses could not be resolved.

¹<https://team-cymru.com/community-services/ip-asn-mapping/>

²<https://pypi.org/project/pyasn/>

Exposing private IP addresses towards the public, *e.g.*, when replying with a TTL Exceeded message, should be the exception. We observed, however, a fairly high ratio of private IP addresses in our traceroute measurements. In 5174 traces, we find 3380 (65 %) paths that include at least one private IP address.

```

~38.5k IP Addresses
↪ ~6.1K Resolved to private
  ↪ ~23 Resolved to PEERING testbed
    ↪ ~30.3K Resolved with Team Cymru
      ↪ ~0.2K Resolved with pyasn
        ↪ ~1.8k Remain unresolved

```

Listing 1. IP to ASN mapping for a single run.
We observe similar results for other measurement runs.

We sanitize AS-level paths in three steps:

- 1) Remove private IP addresses from the AS path if adjacent (left and right) IP addresses belong to the same ASN.
- 2) Remove unresponsive hops (*i.e.*, “*”) if adjacent (left and right) ASNs belong to the same ASN.
- 3) Remove path prepending (*i.e.*, sequence of duplicate ASNs).

In Listing 2, the first line shows an example of the original AS path derived from a traceroute and IP to AS mapping, and the second line shows the sanitized AS path after applying the method outlined above.

```

48147, private, 48147, 200612, 3257, *, 3257, 209, 2722, 47065
48147, 200612, 3257, 209, 2722, 47065

```

Listing 2. Example of Reduction of an AS Path

It is worth noting that we exclude traces that exhibit the following properties: (*i*) paths that include a private IP address that is not adjacent (left and right) to a public IP address, and (*ii*) paths that include a sequence of unresponsive hops where the first and last unresponsive hops are connected via different ASes.

Each traceroute measurement was executed three times shortly after each other. For each pair of source and destination, we verified that the traceroute results were consistent between multiple runs. Different AS paths across each of three runs were rare (between 0.5–1%), which can often be attributed to load balancers or intra-domain changes.

After mapping IP addresses to ASNs, we identify those paths that cross an IXP. IXPs challenge ROV inference because they may deploy ROV on their route servers but they operate transparently on the network layer and thus are not visible along AS paths. In order to detect IXPs within traceroute paths, we utilize the tool TraIXroute [28]. The detection method is mainly based on IXP membership datasets (IXP peering LAN addresses and AS-to-facility mappings) as well as IXP prefixes. It relies on data from PeeringDB [29], Packet Clearing House (PCH) [30], and Routeviews [27]. Additional work details how TraIXroute can be applied in order to uncover IXP peerings [31].

We evaluate the coverage of TraIXroute by parsing RIPE Atlas measurements for our announcements from AMS-IX

such that the corresponding traceroutes are guaranteed to cross this particular IXP. For paths toward the anchor prefix announced only to AMS-IX route servers, we find that TraIXroute only identifies 1.7% of paths as crossing at least one IXP. We confirmed this surprisingly low detection rate via manual investigation, which showed that most traceroutes do not contain all IP addresses around the IXP crossing required for identification, making it impossible for the tool to correctly flag all paths. This is a common traceroute problem. For paths toward the anchor prefix announced to all AMS-IX peers, TraIXroute identifies at least one IXP in 10% of all traceroutes. We observed that routes via other peers traverse other IXPs more often, which seem to be much easier to identify for TraIXroute compared to AMS-IX deployment. As a result, including TraIXroute into our toolchain provides additional metadata but it does not uncover all IXP crossings present.

ROV Identification Phase. From all measurements we collected during our runs, we only consider results that provide sufficient visibility, *i.e.*, could always reach the *anchor* prefix. Additionally, we require that the experiment prefix traverses the same AS path compared to the anchor prefix when the ROA configuration of the experiment prefix should lead to a valid route. Usable data per PoP varies but generally out of 5362 probes participating, 5143 probes were able to execute traceroutes to all announced prefixes, and roughly 2000–2500 probes satisfied the requirements mentioned above. In this subset, we classified the AS paths as follows. If our probe is not able to reach an IP address within the *experiment* prefix when the route is invalid, we conclude that at least one AS on the path must be dropping invalid routes. If a probe is able to reach the experiment prefix, either no AS on path is performing ROV or, alternatively, default routes at ROV-enabled ASes still provide data plane connectivity (false negatives) [24].

Next, we consider the following six cases of AS paths. Table III illustrates those cases based on a set of examples. The first three cases reflect the strict *connected assumption*, known from previous research [10], which requires every tested AS to be directly connected to the origin network (in our experiments, the PEERING testbed). The remaining three cases aim to expand measurement coverage by thoughtfully relaxing the *connected assumption*. Each row with a ROA state *valid* refers to our setup when ROA configuration and BGP announcement should lead to an RPKI valid route for the experiment prefix, while a row with ROA state *invalid* refers to our ROA configuration that should lead to an RPKI invalid route for the experiment prefix.

I hop—Full reachability without route divergence: The default case in which all four traceroutes could reach the target and exhibit the same route. No ROV is deployed. The single AS on the experiment prefix’s AS path is added to the include list. This AS is underlined in Table III.

I hop—Invalid fail: If traceroutes reach the destination except when the experiment prefix has an RPKI invalid route, we observe filtering due to RPKI and classify the AS as ROV enforcing.

Table III

OVERVIEW OF HEURISTICS TO DETECT ROV USING TRACEROUTE TO THE ANCHOR (P_{anchor}) AND EXPERIMENT ($P_{\text{experiment}}$) PREFIX. BOLD ASNS DEPLOY ROV, UNDERLINED ASNS ARE ADDED TO OUR INCLUDE LIST, A GREEN AS PATH REPRESENTS A PATH WHEN THE ROUTE TO OUR EXPERIMENT PREFIX IS VALID, A RED AS PATH ILLUSTRATES A PATH CHANGE BECAUSE OF AN INVALID ROUTE.

Case	ROA		Traceroute		Example AS paths	
	State	ROV	P_{anchor}	P_{exp}	Probe → Valid Prefix	Probe → Invalid Prefix
<i>1 hop</i>						
Full reachability w/o route divergence	Valid	✗	✓	✓	[111 - 47065]	[111 - 47065]
	Invalid		✓	✓	[111 - 47065]	[<u>111</u> - 47065]
Invalid fail	Valid	✓	✓	✓	[111 - 47065]	[111 - 47065]
	Invalid		✓	✗	[111 - 47065]	[111 - "*" - "*"]
Route divergence	Valid	✓	✓	✓	[111 - 47065]	[111 - 47065]
	Invalid		✓	✓	[111 - 47065]	[111 - 222 - 47065]
<i>2+ hops</i>						
Full reachability w/o route divergence	Valid	✗	✓	✓	[111 - 222 - 333 - 47065]	[111 - 222 - 333 - 47065]
	Invalid		✓	✓	[111 - 222 - 333 - 47065]	[<u>111</u> - <u>222</u> - <u>333</u> - 47065]
Invalid fail	Valid	✓	✓	✓	[111 - 222 - 333 - 47065]	[111 - 222 - 333 - 47065]
	Invalid		✓	✗	[111 - 222 - 333 - 47065]	[111 - "*" - "*" - "*" - "*"]
Route divergence	Valid	✓	✓	✓	[111 - 222 - 333 - 47065]	[111 - 222 - 333 - 47065]
	Invalid		✓	✓	[111 - 222 - 333 - 47065]	[111 - 222 - 666 - 444 - 47065]

1 hop—Route divergence: If traceroutes reach the destination via a consistent and direct AS path but exhibit a route detour for the experiment prefix when the route should be invalid, we infer ROV in the tested AS. This is since the tested AS dropped our invalid route announcement but still forwarded the traffic for other reasons (e.g., a default route) to an upstream provider. This upstream provider then routed the traffic based on its own routing table back towards our PoPs. Alternatively, the tested AS might drop or deprefer RPKI invalid routes announced via our peering link but not routes received from an upstream for policy reasons (partial filtering). In this case, the experiment prefix route must be longer because the tested AS is directly connected to the PEERING testbed.

2+ hops—Full reachability without route divergence: The default case for multiple hops if no AS en route is filtering. The probe can reach both anchor and experiment prefixes along the same AS path, independently of our ROA configurations. ASes on the path to the experiment prefix are added to the include list. They are underlined in Table III.

2+ hops—Invalid fail: For longer paths (relaxing the *connected assumption*) we must be cautious. We require every AS on the anchor path to host a RIPE Atlas node such that we can verify that those ASes are not filtering. Also, we utilize knowledge obtained from the strong 1-hop inference cases. If an AS has already been identified as filtering, we discard any route that traverses it. Having excluded every AS on the path except the probe’s AS, we infer ROV in the probe’s AS.

2+ hops—Route divergence: All four traceroute measurements complete, but we observe a route divergence between anchor and experiment prefix when the route of the experiment prefix is RPKI invalid. By stripping the prefix and the suffix of the AS path toward the experiment prefix from the AS path toward the valid prefix, we are able to isolate the route

divergence. As an additional measure of caution, we only consider cases where exactly one AS is left in the remaining snippet of the path. In the example in Table III, we would flag AS333 as ROV-enforcing.

Include List. Since we would like to be able to differentiate between partially and fully filtering ASes, we maintain an include list. Any AS that has been seen forwarding RPKI invalid announcements will be added to the include list. Once an AS was flagged as ROV enforcing and is also part of the include list we infer that the AS is only partially filtering.

V. RESULTS

The number of ASes that we identify during ROV identification (details see Section IV) are displayed in Table IV. These results include all prefix pairs announced via all five PoPs. In the 1-hop cases, which satisfy the *connected assumption*, we observe ≈ 41 ASes “without route divergence”, while for cases “invalid fail” and “route divergence” we record much more events. This is mostly due to RPKI ROV deployment on route servers at AMS-IX and Seattle-IX. These IXPs cover most of PEERING’s direct peers (Table II).

Relaxing the *connected assumption* (i.e., 2+ hops) we find 731–803 ASes that do not filter based on RPKI. We find ≈ 9 –13 ASes that do deploy RPKI filtering (see “invalid fail” and “route divergence”), though. The overall low number of filtering ASes result from the multiple conditions that we require to be fulfilled to flag an AS as ROV enforcing. We argue, however, that those conditions prevent the introduction of false positives.

Overall, we identify 194–206 ROV-enforcing ASes within our set of 5537 RIPE probes covering 3694 ASes. 48–60 of these ASes are partially filtering ASes, which have been observed filtering in one particular instance while also been observed forwarding RPKI invalid announcements to other

Table IV
RESULTS OF ROV DATA-PLANE MEASUREMENTS FOR ALL POPs [# ASes]

Case	Measurement Run				
	1	2	3	4	5
1 hop					
Full reachability w/o route divergence	43	43	42	41	37
Invalid fail	181	181	182	175	181
Route divergence	15	15	15	13	12
2+ hops					
Full reachability w/o route divergence	803	798	775	731	711
Invalid fail	2	2	2	2	1
Route divergence	11	10	11	7	8
Total unique ROV	206	205	202	194	199
Added to include list	630	628	626	628	587
Partially filtering	60	58	54	55	48
Fully filtering	146	147	148	139	151

peers in other instances. Since we maintain an include list of such ASes, we are able to identify those as partially filtering. The remaining 139-151 ASes have been observed filtering but never forwarded invalid announcements.

ROV at Route Servers. We are able to make more detailed inferences for all 1-hop cases for ASes that peer with PEERING at AMS-IX or Seattle-IX (see Table V). We announced a pair of prefixes to *route servers only* and another pair to *all peers excluding route servers* to accurately determine whether the filtering was done by the AS or by the IXP route server. If the ROV inference happened when we announce to *route servers only* but not when we announce to *all peers excluding route servers*, we conclude that the IXP route server performed the filtering and not the AS (*i.e.*, 160 ASes members of AMS-IX and 33 ASes members of Seattle-IX). In the opposite case, we conclude that the AS is indeed filtering (*i.e.*, 9 AMS-IX members and 2 Seattle-IX members). If both yielded a positive identification then the filtering is deployed at the IXP as well as in the AS (*i.e.*, 4 AMS-IX members and 2 Seattle-IX members).

For all other PoPs without these additional route server prefix pairs, in the 1-hop case and also in all 2+ hop cases, we need to rely on the identification of IXPs in traceroutes via TraIXroute. If an IXP was detected on an AS path that led to the tagging of an AS as ROV-enforcing, we mark the AS with an IXP tag in order to highlight that also the IXP could have been responsible for RPKI filtering instead of the AS. That was the case for two ASes. Both crossing the Digital Realty Internet Exchange.

Sanitization. To make our methodology even more rigorous and minimize the likelihood of false positives we looked at the deviation between the different measurement runs. It turns out that 89% of ASes are present in three or more measurement runs. If traceroutes to the anchor and experiment prefixes complete successfully when both routes are valid, the traceroute to the anchor prefix always completes successfully,

and the traceroute to the experiment prefix does *not* complete successfully when the route is invalid for some reason not related to ROV, a false positive would be introduced. To filter out potential false positives caused by measurement noise, we enforce the restriction that ROV inference in an AS has to occur in three or more runs to be taken into account as ROV-enforcing.

Validation. We validated our findings mainly in a manual process that involved whois records, PeeringDB information, Twitter announcements, and operator mailing lists. Out of 174 ASes that support IPv4 and have direct peering sessions with the PEERING testbed, 73 ASes host at least one RIPE Atlas probe and 10 ASes (9 at Amsterdam and 1 unique additional AS at Seattle) were flagged in the 1-hop case as ROV-enforcing. Given the requirement of positive identifications in three independent measurement runs, we manually vetted 9 of the 10 1-hop inferences to be true positives. For one we could neither confirm nor deny the inference. The remaining 156 (AMS-IX) and 31 (Seattle-IX) ASes in the 1-hop case were flagged via a route server prefix, see Table V. Since 8 ASes were flagged via both IXPs, the number of unique ASes inferred to deploy ROV via announcements to route servers is 172. We expected all such ASes to be members of either Seattle-IX or AMS-IX, which are known to be filtering. Surprisingly, 6 ASes for Seattle and 8 ASes for Amsterdam that were identified via route server prefixes are not members of the respective IXPs, although traceroute data and its mapping suggest such thing. We investigated manually and suspect remote peering to be the underlying mechanism that allows traceroutes to travel the way they do. A company called IX Reach offers precisely such a service that allows remote peering via their infrastructure at multiple IXP facilities. The 2+ hops case yields 6/12 true positives. For the remaining ones no information could be found. Two out of those 12 ROV inferences were also tagged as having crossed an IXP.

All results of this study have been made public: <https://github.com/nrodday/TMA-21>.

VI. DISCUSSION AND LESSONS LEARNED

Measurement infrastructure. Operational maintenance in testbeds such as PEERING may always lead to interruptions of prefix propagation. We also found the dynamic deployment of firewall filters at upstreams. We recommend to monitor the stability of VPN tunnels of the PEERING testbed during measurements when the exchange of BGP traffic between

Table V
IXP FILTERING FOR 1-HOP ‘INVALID FAIL’ AND ‘ROUTE DIVERGENCE’ CASES OF ROV INFERENCES FOR AMS-IX AND SEATTLE-IX MEMBERS [# ASes]

PoP	Total ROV	Route server	Direct	Both
AMS	165	160	9	4
Seattle	33	33	2	2

external hosts and the PEERING testbed PoPs is required, see Figure 4. If prefix space is limited, it seems reasonable to focus on the Amsterdam and Seattle PoPs as they provide the richest connectivity. RIPE Stat [23] helps understanding if a prefix propagates correctly.

On the RIPE Atlas platform we experienced some problems such as none-execution of measurements or unexpected termination after several days. This led to gaps in the dataset and could only be fixed, in most cases, by restarting the whole campaign or parts of it. RIPE Atlas provides the feature to schedule measurements for a specific time. Since we had a strict ROA schedule, measurements needed to be executed roughly around the anticipated time to not reach into a different ROA cycle and therefore measure during the opposite ROA state. We recommend to account for a grace period to be able to use the results of delayed measurements, if possible. Also, one should always check the actual stop date of a measurement to prevent running into an unexpected delay that falsifies measurements.

Middleboxes. Middleboxes that replied instead of the destination [32] in our active data plane measurements initially falsified our results. We used server-side logging to cross-check whether a traceroute has actually reached the PEERING testbed. Depending on server load, roughly 1 % of packets will be dropped during the capture and therefore missing in the server logs, although the traceroute has successfully reached PEERING. In those cases we check whether the traceroute has a PEERING LAN address as the second last hop, implying that it reached the PEERING testbed.

IP address to ASN mapping. Mapping router IP addresses to ASNs is challenging. Adjacent routers that belong to different ASes interconnect via a common peering LAN. The IP prefix of the peering LAN belongs to a single AS, though. When a router replies with the peering LAN IP address (*e.g.*, in traceroute measurements), this leads to incorrect results if the peering LAN is sponsored by the adjacent AS. We observed some traceroute results containing hops which used IP addresses from the peering LAN of the adjacent AS. Based on manual reverse DNS lookups we were able to identify the correct AS of peers. We recommend to use different services and aggregate the obtained data, see Section IV.

ROA schedule. Originally, we planned to adopt the ROA schedule from [10], which changes ROA configurations every eight hours. We observed, however, inconsistent router behavior in our measurement results, which led to further manual investigation. Our initial assumption that all Relying Parties (RPs) fetch ROAs and update their routers within a short period of time does not hold. We observed that some ASes still maintain routes based on stale ROAs. To account for these long update times, we extend the ROA schedule and change ROAs every 24 hours. This gives RPs ample time to update their routers. Our observations triggered a discussion within the IETF SIDROPS working group that led to a draft to narrow down timing parameters for the RPKI supply chain [33].

Further details are discussed in [34].

Limitations. Several limitations remain that we would like to point out. First, the IP to ASN mapping is crucial for all analysis thereafter. If external data sources provide the wrong mapping, inferences made for ROV filtering could be attributed to the wrong ASes. Second, once the AS paths have been obtained, we classify the data into the six cases mentioned in Section IV. Our approach does not consider induced path changes as outlined in [35]. Since the algorithm described in [35] assumes full visibility on all surrounding ASes, it is not usable for the method described in this paper. It further assumes only a single root cause being responsible for a change in routing behavior. Another requirement that cannot be fulfilled here, since a swap of the RPKI ROA to the invalid state will possibly trigger more than just a single AS to drop the invalid prefix and cause a route change. Third, if IXP identification fails although the traceroute traversed an IXP, ROV filtering could be wrongly attributed to the AS instead of the IXP.

VII. CONCLUSION

This work revisited existing ROV identification approaches and presented a novel measurement methodology based on RIPE Atlas vantage points using controlled experiments to identify ROV-enforcing Autonomous Systems. We showed limitations of simple end-to-end methodologies using HTTP and ICMP measurements and argued in favour of traceroute to derive more detailed inferences from the data in order to pinpoint the filtering AS. Our methodology extends the current state-of-the-art by presenting six classification cases and including routeserver identifications via dedicated prefixes as well as TraIXroute. Moreover, we relaxed the *connected assumption* of previous research [10] and extended coverage that allows to make inferences of ROV-enforcing ASes even if the tested AS is more than one hop away from the PEERING testbed. Our measurements from 5537 vantage points in 3694 ASes infer ROV is deployed in 206 unique ASes: 10 with strong confidence, 12 with weak confidence, and 184 indirectly adopting ROV via filtering by IXP route servers of which 146 are fully and 60 are partially filtering.

Future work will look at induced path changes and how to satisfy the assumptions made in [35] such that the methodology can be used to further sanitize our ROV deployment inferences.

ACKNOWLEDGMENT

We thank the RIPE NCC team for lifting our measurement limits to give us the freedom necessary to perform large scale traceroute campaigns and the RIPE Atlas community for supporting our research with credits necessary to schedule the measurements. We also thank the anonymous reviewers and our shepherd Gianni Antichi for their constructive feedback. This work was partly supported by the European Union within the Horizon 2020 project *CONCORDIA* (grant agreement No 830927), and the German Federal Ministry of Education and Research (BMBF) within the project *PRIME*net.

REFERENCES

- [1] RIPE NCC, “YouTube Hijacking,” <http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study>, 2009, [Online; accessed 15-January-2020].
- [2] A. Toonk, “Hijack Event Today by Indosat,” <http://www.bgpmon.net/hijack-event-today-by-indosat/>, 2014, [Online; accessed 15-January-2020].
- [3] M. Lepinski and S. Kent, “An Infrastructure to Support Secure Internet Routing,” RFC 6480, Feb. 2012. [Online]. Available: <https://rfc-editor.org/rfc/rfc6480.txt>
- [4] M. Lepinski, S. Kent, and D. Kong, “A Profile for Route Origin Authorizations (ROAs),” IETF, RFC 6482, February 2012.
- [5] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, “BGP Prefix Origin Validation,” IETF, RFC 6811, January 2013.
- [6] M. Wählisch, O. Maennel, and T. C. Schmidt, “Towards detecting bgp route hijacking using the rpki,” *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 4, pp. 103–104, Aug. 2012.
- [7] D. Iamartino, C. Pelsser, and R. Bush, “Measuring BGP route origin registration validation,” in *Proc. of PAM*, ser. LNCS. Berlin: Springer, 2015, pp. 28–40.
- [8] 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 *Proc. of 14th ACM Workshop on Hot Topics in Networks (HotNets)*. New York: ACM, Nov. 2015, pp. 11:1–11:7.
- [9] T. Chung, E. Aben, T. Bruijnzeels, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, R. van Rijswijk-Deij, J. Rula, and N. Sullivan, “RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins,” in *Proc. of ACM IMC*. New York, NY, USA: ACM, 2019, pp. 406–419.
- [10] 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.
- [11] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman, “Are We There Yet? On RPKI’s Deployment and Security.” in *Proc. of NDSS*, 2017.
- [12] T. Hlavacek, A. Herzberg, H. Shulman, and M. Waidner, “Practical experience: Methodologies for measuring route origin validation,” in *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2018, pp. 634–641.
- [13] National Institute of Standards and Technology, “Nist rpki monitor,” <https://rpki-monitor.antd.nist.gov/>, 2020, [Online; accessed 1-February-2020].
- [14] B. Schlinker, T. Arnold, I. Cunha, and E. Katz-Bassett, “PEERING: Virtualizing BGP at the Edge for Research,” in *Proc. ACM CoNEXT*. New York, NY, USA: ACM, December 2019.
- [15] RIPE NCC Staff, “Ripe atlas: A global internet measurement network,” *Internet Protocol Journal*, vol. 18, no. 3, 2015.
- [16] B. Cartwright-Cox, “Are BGPs security features working yet?” <https://blog.benjojo.co.uk/post/are-bgps-security-features-working-yet-rpki>, 2019, [Online; accessed 10-January-2020].
- [17] N. Künneke-Trenaman, E. Aben, J. den Hertog, and J. Snijders, “RPKI Test,” <https://www.ripe.net/s/rpki-test>, 2019, [Online; accessed 31-December-2019].
- [18] C. Testart, P. Richter, A. King, A. Dainotti, and D. Clark, “To Filter or not to Filter: Measuring the Benefits of Registering in the RPKI Today,” *PAM*, 2020.
- [19] G. Huston and J. Damas, “Measuring Route Origin Validation,” <https://www.potaroo.net/ispcol/2020-06/rov.html>, 2020, [Online; accessed 16-October-2020].
- [20] A. Reuter, R. Bush, I. Cunha, E. Katz-Bassett, T. C. Schmidt, and M. Wählisch, “Measuring RPKI Route Origin Validation Deployment,” <https://rov.rpki.net>, [Online; accessed 14-June-2020].
- [21] C. Gray, C. Mosig, R. Bush, C. Pelsser, M. Roughan, T. C. Schmidt, and M. Wählisch, “BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping,” in *Proc. of ACM Internet Measurement Conference (IMC)*. New York: ACM, 2020, pp. 492–505.
- [22] B. Cartwright-Cox, “The year of RPKI on the control plane,” https://ripe80.ripe.net/presentations/36-Ben_Cox.pdf, 2020, [Online; accessed 14-May-2020].
- [23] RIPE NCC, “RIPE Stat,” <https://stat.ripe.net/>, 2020.
- [24] N. Rodday, L. Kaltenbach, I. Cunha, R. Bush, E. Katz-Bassett, G. D. Rodosek, T. C. Schmidt, and M. Wählisch, “On the Deployment of Default Routes in Inter-domain Routing,” in *ACM SIGCOMM 2021 Workshop on Technologies, Applications, and Uses of a Responsible Internet (TAURIN 2021)*. ACM, 2021, accepted for publication.
- [25] RIPE NCC, “RIPE Atlas Cousteau,” <https://github.com/RIPE-NCC/ripe-atlas-cousteau>, 2019, [Online; accessed 10-January-2020].
- [26] —, “RIPE Routing Information Service (RIS),” <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>, 2020, [Online; accessed 16-October-2020].
- [27] RouteViews Project, “University of oregon routeviews project,” <http://www.routeviews.org>, 2013, [Online; accessed 16-Juli-2020].
- [28] G. Nomikos and X. Dimitropoulos, “traixroute: Detecting ixps in traceroute paths,” in *International Conference on Passive and Active Network Measurement*. publisher, 2016, pp. 346–358.
- [29] PeeringDB, “The Interconnection Database,” <https://www.peeringdb.com/>, 2020, [Online; accessed 10-October-2020].
- [30] Packet Clearing House, “Internet Exchange Directory,” <https://www.pch.net/ixp/dir>, 2020, [Online; accessed 10-October-2020].
- [31] G. Nomikos, V. Kotronis, P. Sermpezis, P. Gigis, L. Manassakis, C. Dietzel, S. Konstantaras, X. Dimitropoulos, and V. Giotsas, “O peer, where art thou? uncovering remote peering interconnections at ixps,” in *Proceedings of the Internet Measurement Conference 2018*. New York, NY, USA: ACM, 2018, pp. 265–278.
- [32] K. Edeline and B. Donnet, “A first look at the prevalence and persistence of middleboxes in the wild,” in *2017 29th International Teletraffic Congress (ITC 29)*, vol. 1. IEEE, 2017, pp. 161–168.
- [33] R. Bush, J. Borkenhagen, T. Bruijnzeels, and J. Snijders, “Timing Parameters in the RPKI based Route Origin Validation Supply Chain,” Internet Engineering Task Force, Internet-Draft draft-ietf-sidrops-rpki-rov-timing-00, May 2020, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-rpki-rov-timing-00>
- [34] J. Kristoff, R. Bush, C. Kanich, G. Michaelson, A. Phokeer, T. C. Schmidt, and M. Wählisch, “On Measuring RPKI Relying Parties,” in *Proceedings of the ACM Internet Measurement Conference*. New York, NY, USA: ACM, 2020, pp. 484–491.
- [35] U. Javed, I. Cunha, D. Choffnes, E. Katz-Bassett, T. Anderson, and A. Krishnamurthy, “Poiroot: Investigating the root cause of interdomain path changes,” *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, pp. 183–194, 2013.