








The Resource Public Key Infrastructure (RPKI): A Survey on Measurements and Future Prospects

Nils Rodday^{*†} , Ítalo Cunha[‡] , Randy Bush[§] , Ethan Katz-Bassett^{††} 
Gabi Dreo Rodosek^{*} , Thomas C. Schmidt^{||} , Matthias Wählisch^{**} 
^{*}Universität der Bundeswehr München, [†]University of Twente, [§]Arrcus / IJJ,
[‡]Universidade Federal de Minas Gerais, ^{††}USC / Columbia University,
^{||}Hamburg University of Applied Sciences, ^{**}TU Dresden

Abstract—The adoption of the Resource Public Key Infrastructure (RPKI) is increasing. To better understand and improve RPKI deployment, measuring route origin authorization (ROA) objects, RPKI route origin validation (ROV), and RPKI resilience is essential. In this paper, we survey RPKI-related research that aims to understand RPKI deployment. Additionally, we enrich our survey with many industry and IETF-related contributions. Our work provides an in-depth analysis of the many ideas and challenges discussed in studies of the RPKI ecosystem and includes lessons from mistakes made in the past, which we should avoid in the future.

Index Terms—RPKI, ROA, ROV, control plane, data plane, BGP, security, measurement, resilience, survey

I. INTRODUCTION

The Border Gateway Protocol (BGP) does not provide any security guarantees. Any Autonomous System (AS) can announce any route to its peers without giving the receiving peer the option to verify whether the announcement is correct. The absence of such a mechanism allows for intentional attacks and unintentional misconfigurations. An intentional attack could be a targeted prefix hijack that is designed to render services unavailable [1], [2], use allocated address space for spamming purposes [3], [4], or reroute traffic that aims for stealing cryptocurrencies. The latter has happened recently and led to a loss of USD 1.9 million worth of digital money [5], [6]. An unintentional misconfiguration could be a typo that disturbs the routing system by incorrectly forwarding traffic to ASes that should not receive the traffic, *e.g.*, a route leak [7]–[11].

Although a plethora of literature was published during the past two decades, including several surveys that highlight the main challenges of securing BGP [12]–[23], security problems with BGP remain a persistent threat in real deployments.

Many tools have been proposed to detect prefix hijacks, like PHAS [24], Argus [25], BGPAlert [26], ARTEMIS [27], HEAP [28]. Those tools rely on heuristics and not on cryptographically secured material proving whether the announcement is correct. Hence, smart attackers can evade such monitoring [29].

A fundamentally different approach to secure inter-domain routing is the replacement of BGP, *e.g.*, SCION [30]. These

green-field approaches may experience deployment in specific scenarios, but BGP is still the de-facto standard for inter-domain routing. Since BGP is not expected to be replaced anytime soon, the Secure Inter-Domain Routing (SIDR) [31] working group specified BGP extensions to overcome prefix hijacks and path manipulations. These extensions are based on the Resource Public Key Infrastructure (RPKI) [32], a PKI for Internet resources (*i.e.*, IP prefixes and AS numbers). Operators can create Route Origin Authorization (ROA) objects [33] to cryptographically prove legitimate BGP announcements, deploy Route Origin Validation (ROV) to reject or deprefer announcements that violate a ROA [34], or both. Earlier and current RPKI research focuses on three major topics:

ROA measurements: Identifying the address space covered by ROA objects, *i.e.*, which ASes protect their address space.

ROV measurements: Measuring ROV deployment, *i.e.*, which ASes prevent propagation of invalid routes.

RPKI resilience: Measuring weak RPKI components and features, *i.e.*, revisiting design decisions and testing which RPKI software contains vulnerabilities that endanger the availability of RPKI infrastructure components or integrity of RPKI information.

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 May 30, 2023, RouteViews [35] indicated 43.53% of prefixes were valid, 1.12% were invalid, and 55.35% did not have a covering ROA. Up-to-date results can be found at the NIST RPKI deployment monitor [36].

Measuring ROV deployment is more challenging as it requires inferring (private) router configuration changes. The ultimate goal is to infer which ASes are using RPKI data when applying BGP policies (*e.g.*, dropping invalid route announcements). Common measurement setups are based on passively collected data 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 [37]. Instead,




Figure 1. Origin validation binds resources to ASNs. Path validation ensures that an announcement took the path contained in the BGP AS path attribute. Path validation requires origin validation to work properly.

controlled measurements are preferred because they limit the number of independent variables by introducing well-defined ROA and BGP events.

RPKI resilience has already been considered in the standards from the beginning but has become more critical in recent years when the RPKI experienced considerable deployment. While some attacks were discussed on a theoretical level, only recently have researchers attempted to exploit bugs within the RPKI and found vulnerabilities in different software solutions. Many components within the RPKI depend on each other. Moreover, developers make many assumptions about the availability of services and delivery of data, which do not always hold true. As a result, vulnerabilities arise that sometimes threaten the availability of the RPKI ecosystem or the integrity of the data served within.

Contributions. In this work, we classify RPKI measurement research. We survey more than 40 scientific publications as well as many industry and Internet Engineering Task Force (IETF) references. Our goal is to help the reader get a complete picture of the current state of the RPKI ecosystem.

The remainder of this paper is structured as follows: Section II provides details on *origin validation* and *path validation*, where we explain how the RPKI works. The following three sections present our RPKI survey, each analyzing current research chronologically. Section III details earlier work in the ROA measurements domain. Section IV surveys current ROV measurement approaches. Section V provides insights into earlier research dealing with RPKI resilience. Finally, Section VI summarizes our findings.

II. BACKGROUND

Proposals to improve BGP security can be separated into *origin validation* and *path validation*, illustrated in Figure 1. While *origin validation* verifies that an origin AS is allowed to announce a specific prefix, *path validation* verifies that an announcement traversed the path that is stored in the BGP AS path attribute of the announcement.

A. Origin Validation

The lack of proof of address ownership remains a persistent threat in BGP. Many operators use data provided by the Internet Routing Registry (IRR) [38], a globally distributed database, to achieve consistent routing by sharing information (e.g., which AS is allowed to announce a prefix) between

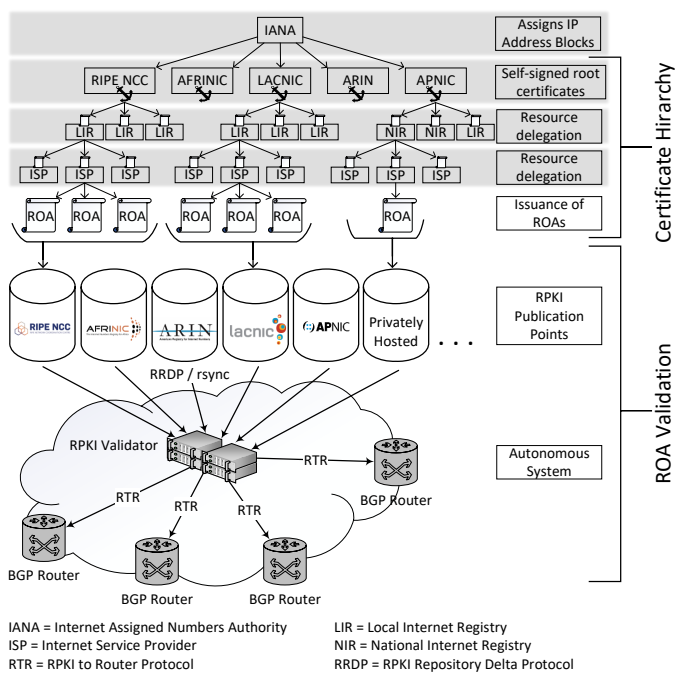


Figure 2. RPKI architecture: The IANA delegates resources to the five RIRs, each of which maintains a trust anchor. They continue to delegate smaller chunks of the resources to their subordinates. An ISP at the end of the chain will create a ROA which is published at a publication point. RPKI validators fetch from the publication points via RRDP or rsync to cryptographically validate the data and send the resulting file via the RTR protocol to the BGP routers for inclusion in the BGP route selection process.

network operators. The IRR serves as source of information to extract peering information [39], infer AS relationships from IRR routing policies [40], or to detect BGP hijacks [25], [28].

Operators create prefix filters based on the IRR to increase inter-domain routing security and prevent hijacks and route leaks. This data, however, is not cryptographically secured and is known to contain incorrect information [41]. To bind ASes to prefixes in a cryptographically secure manner, several approaches have been proposed [42], [43]. Ultimately, the RPKI was designed, proposed, and deployed. The standardization of the RPKI started in April 2008 within the SIDR working group, which led to a set of specifications that define an infrastructure to support secure Internet routing [32]. Between January 2011 and September 2012, all five Regional Internet Registries (RIRs) started deploying RPKI by providing cryptographically signed attestation objects that describe ownerships of IP address spaces. Furthermore, they support the creation of ROAs, objects that define which ASes are allowed to announce IP prefixes. Developing guidelines and providing operational guidance for secure inter-domain networks continues in the SIDR Operations (SIDROPS) [44] working group of the IETF after the start of RPKI deployment in 2012.

An Internet draft that analyses the requirements of the RPKI and the state of the global repository in 2013 is published in [45]. The size of a fully deployed RPKI was estimated in [46] with the time required for synchronization of a fully deployed RPKI rather in the order of days, not hours or minutes.

RPKI architecture overview. The RPKI architecture is shown in Figure 2. The trust chain implements the delegation of IP address resources. The Internet Architecture Board (IAB) initially strongly recommended the use of a single Trust Anchor (TA) [47] instead of one trust anchor run by each RIR but updated its recommendation later in 2018 arguing that it no longer held on to its initial statement [48]. Currently, each RIR maintains its own TA, which is a self-signed root certificate. Therefore, each RIR can theoretically also attest prefix space under the management of other RIRs.

The overarching power of each RPKI root entity that comes with a Public Key Infrastructure (PKI) has been a controversy, with proposals suggesting the use of the Dalskov protocol [49] as a distributed threshold signature model instead, in which only a set of RIRs could jointly sign End-Entity (EE) certificates and delegate address space. These proposals have never been considered for deployment due to additional complexity.

To delegate address spaces to other entities, each RIR signs EE certificates. This process is repeated until an AS gets hold of an EE certificate that proves its authority over the address space. To protect the address space in the global routing system and communicate the authoritative information, a ROA has to be created. A ROA contains, amongst others, the following information:

```
ROA: Prefix/Max-Length, ASN

For example:
ROA: 147.22.0.0/16-24, AS 47065
```

In this example, AS47065 is authorized to announce any BGP route of the prefix 147.22.0.0 of length /16 to /24. To protect unassigned or address space that should not be announced, the RIRs and organizations can issue ROAs setting the origin AS to AS0, which is reserved by the IANA to identify non-routed networks [50].

BGP prefix origin validation. Creating ROAs is one crucial step to protect BGP. It enables other ASes to verify BGP announcements. The verification, however, is a separate process called BGP prefix origin validation [34], commonly also known as Route Origin Validation. Given a BGP announcement and a set of cryptographically correct ROAs, prefix origin validation searches this set for all ROAs that cover the IP prefix of a BGP announcement based on the longest common prefix match. To be *valid*, the announcement must match the ROA in terms of origin AS and prefix length. Announcements that are covered by a ROA but either include another origin AS or are more specific compared to the maximum length are *invalid*. All other announcements (*i.e.*, those that are not covered by any ROA) are *unknown*. In our example above, BGP announcements 147.22.0.0/25, AS 47065 and 147.22.0.0/16, AS 50000 would be invalid. BGP announcement 100.20.0.0/25, AS 47065 would be unknown, given that no other ROA information exists.

Creating resource certificates [51] and ROAs can be perceived as cumbersome, complex, and time-consuming. There has been a proposal to move the RPKI delegation from a certification model to a de-facto ownership model where an AS is considered to be the owner if it used an address space over

a long period consistently [52], [53]. This would temporarily weaken the security guarantees of the RPKI until proper RPKI-signed ROAs are deployed, but it would speed up the process in which ASes would be granted reign over resources. Because of the aforementioned security drawbacks, the proposal has not been considered for adoption.

Prefix filtering. When the validation outcome is available, a router can then implement BGP policies by assigning different preferences (*e.g.*, prefer valid over not found routes and not found over invalid routes) or rejecting routes altogether. It is worth noting that keeping invalid routes is not recommended. If the hijacked prefix is a sub-prefix of the legitimate announcement, no competing valid announcement would be available. As such, the invalid route would be installed and used since a longer common prefix match cannot overrule it. Rejecting invalid routes is recommended and should be adopted by participating ASes.

Hosted vs. delegated model. To manage resource delegation and ROA issuance, the RPKI allows for two separate models, which are supported by most RIRs: hosted model and delegated model. In the hosted model, all certificates and ROAs are stored and managed by a RIR. Users of address space can usually create and configure ROAs using a web portal. In the delegated model, maintenance of the RPKI is delegated to third parties. Owners of Internet resources can run their own Certificate Authority (CA) to manage certificates and ROAs of IP prefixes. Running a publication point is also possible but not required. A dedicated service for ASes who only wish to run the CA but not the publication point is called Publish-In-Parent. This service is already run productively by APNIC [54] and is currently implemented for production by RIPE NCC [55].

The hosted model reduces most of the CA/PKI-related complexity for resource owners. The downside is that private keys remain with the RIR. Moreover, an AS managing resources from different RIRs needs to log into the various portals of each RIR. On the other hand, the delegated model allows for more flexibility and serves as a single point of control. Common RPKI CA software, namely Krill [56] from NLNetLabs and RPKI Toolkit [57] from Dragon Research Labs, can manage resources from different RIRs in a single instance. The private keys remain with the AS running the CA. Larger organizations run their own CA and publication points to implement complete control, while smaller ASes refuse the overhead and mostly use the hosted solution.

Relying party software. Once ROAs are publicly available in the RPKI publication points, an AS might decide to use that information for filtering BGP announcements. There are several Relying Party (RP) software variants, also called RPKI validators, see Table I. The RPKI-Validator 3 [58] and Rcynic [59] have been discontinued. Some of these validators can also be used to manually check single objects in the RPKI via their web Graphical User Interface (GUI). Moreover, RPKI MIRO [60] and RPKIVIZ [61] allow for easy visualization of RPKI objects, while RPKImancer [62] allows creating and dissecting RPKI objects on the command-line.

Friedemann *et al.* [63] compared the validators regarding

Table I
RPKI RELYING PARTY SOFTWARE

Name	Initial Release Year	Language	Maintained
RPKI Validator 1 [67]	2011	Scala	✗
Rcynic [59]	2012	C/Python	✗
RPSTIR [68]	2012	C	✗
OctoRPKI [69]	2019	Go	✓
FORT-Validator [70]	2019	C	✓
rpki-client [71]	2019	C	✓
Routinator 3000 [72]	2019	Rust	✓
RPSTIR2 [73]	2020	Go	✓
rpki-prover [74]	2020	Haskell	✓
RPKI-Validator 3 [58]	2021	Java	✗

their performance and recommended Routinator. An operator usually deploys one or two instances (for redundancy) within the AS. The RP software fetches all available ROAs, either via rsync [64] or the RPKI Repository Delta Protocol (RRDP) [65]. Rsync is known to have some drawbacks: The heavy CPU and memory load makes it an easy target for Denial of Service (DoS) attacks, the lack of library support, and the difficulty of publishing objects atomically. Therefore, the IETF introduced RRDP to replace rsync in the future [66]. RRDP was designed to mitigate the previously mentioned shortcomings and supports HTTPS CDN infrastructure to increase resilience during content provisioning.

RTR implementations. Each validator produces an output called Validated ROA Payload (VRP) in an *out-of-band* fashion. The validation process is performed on dedicated hardware and does not require any BGP router resources. The VRP is a list of all ASNs to prefix combinations that were cryptographically validated during the validation process. There may be more VRP entries compared to the amount of ROAs since a single ROA can be used to authorize multiple prefixes, while a VRP entry only contains a single prefix to ASN combination. The VRP is then delivered to BGP routers within the AS via the RPKI-to-Router (RTR) protocol [75] or alternative out-of-band approaches (*e.g.*, JSON file). Each validator in Table I ships with a RTR-server implementation. On the client side, the router has to run a piece of software that receives the VRP via the RTR protocol from the RPKI validator software. There are closed-source implementations in proprietary router products and open-source implementations, see Table II. Most RTR client implementations provide a function to compare BGP announcements with the received VRPs and assign a validation outcome.

The RPKI is designed in a soft-fail manner. If some RPKI components are unavailable, caches expire, and routes default to unknown. Therefore, all routes should be treated as if RPKI was not deployed. The design decision has been made since an operator's worst-case scenario is dropping traffic due to a recently enabled security feature. This was perceived as a main hindrance to adoption. A short discussion about the soft-fail mechanism and the proposal of an alternative can be found in the ROVER approach [76]–[78]. Further shortcomings have

Table II
OPEN-SOURCE RTR IMPLEMENTATIONS

Name	Language	Client	Server	Maintained
RTRlib [82]	C	✓	✗	✓
rpki-rtr-client [83]	Python	✓	✗	✓
StayRTR [84]	Go	✓	✓	✓
rpki-rtr [85]	Rust	✓	✓	✓

been discussed by Geoff Huston [79].

Comparison of RPKI and IRR deployment. As prefix space covered by ROAs has been constantly increasing during the past years, a recent study showed a comparison between IRR and RPKI and found datasets to be inconsistent in 27.4% and 61.4% of cases for Réseaux IP Européens (RIPE) IRR and Routing Assets Database (RADB), respectively [41]. The findings highlight the need for the RPKI as a cryptographically proven database with accurate information.

Du *et al.* [80] further investigated the level of conformity of ASes participating in the Mutually Agreed Norms for Routing Security (MANRS) project. MANRS is an initiative that attempts to improve routing security. Participating ASes are required to perform certain security-related activities, such as registering their prefixes in the IRR or RPKI. They set out to measure conformity of MANRS ASes and use the Internet Health Report (IHR) of the Internet Initiative Japan (IIJ) [81] as an underlying source. In May 2022, small and medium ASes participating in MANRS were more likely to originate only RPKI-valid announcements (60.1%) compared to non-MANRS participants (24.7%). MANRS participants also originate much less RPKI invalid prefixes (23.6%) compared to non-MANRS participants (68.1%). For larger networks, the gap is much smaller. Large ASes participating in MANRS also propagate RPKI invalid announcements at a lower rate.

B. Path Validation

RPKI combined with route origin validation solves a subset of attacks, such as exact and subprefix hijacks, *e.g.*, when an announcement is sent by an AS that is not legitimate based on a ROA. This prevents many accidental hijacks caused by fat finger incidents. However, other attacks remain, such as path prepending of the legitimate AS, that would still render announcements of the attacker valid. Very recently, an attacker hijacked Amazon's address space by forging the AS path and claiming to be an upstream of an Amazon ASN [86], [87]. To solve these attacks, path validation is required to verify whether the AS path attribute in the BGP announcement matches the actual path that an announcement traversed.

Border Gateway Protocol Security (BGPsec) [88], [89] has been standardized in 2017 in RFC8205 and is designed to satisfy these requirements. It uses forward signing of every AS hop and ensures via strong cryptography that the announcement travels the exact same way on the AS level as described in the AS path attribute. There have been many other proposals made before: soBGP [90], sBGP [91], psBGP [92], [93], pgbgp [94], DPVA [95].

Unfortunately, BGPsec is currently not deployed, and there is little hope for deployment soon. A significant hurdle for adoption is that BGPsec is using *in-band* processing. Every router is required to run cryptographic operations while processing the BGP updates. Since processing capacities directly translate to the amount of BGP updates that a router can handle, BGPsec is too costly in its current state. Developments have made BGPsec faster [96]. But even when the performance problem is solved, another hurdle remains: it is not particularly useful when partially deployed. Therefore, BGPsec and S-BGP have limited security benefits over the RPKI until these solutions are widely adopted [97].

Since partial deployment is a by-product of incremental adoption, some participants within the IETF aim to design lightweight alternative path plausibility solutions that also provide benefits when only partially deployed. In addition to BGP prefix hijacking, they should also be able to mitigate route leaks. Route leaks [98] are identified based on the fact that they violate the Gao-Rexford model [99]. The model describes a set of rules that ASes typically follow to avoid routing valleys. Although routing is not always valley-free, the Gao-Rexford model is used due to a lack of alternatives [100]–[104]. Autonomous System Provider Authorization (ASPA) [105], [106] is a proposal that defines AS relationships from customers to providers. Each AS publishes an ASPA object in which it authorizes its upstream providers to propagate its routes. RP software is able to retrieve these objects, build a tree from them, and decide which paths are valid ones. The processing of cryptographic objects is done in an *out-of-band* fashion within the RP software, similar to RPKI. A BGP router receives a list with validated paths and can compare whether the AS-path attribute in a BGP announcement contains only validated hops. More research is needed to evaluate the introduced performance impact during operation. The outcome is RPKI-like: valid, invalid, or unknown. The proposal is currently moved forward in the SIDROPS working group.

Another idea that aims to achieve the same benefits but has weaker security guarantees is AS-Cones [107]. It takes the opposite direction by allowing upstreams to define AS-Cone objects listing all their customers. The advantage is that large-scale deployment will be much easier to achieve if fewer, large ASes need to participate in issuing cryptographically signed objects. Considering a BGP dump from all RouteViews [35] collectors for 24 hours on October 1, 2022, we obtain a graph with 74,110 ASes and classify each link using the CAIDA AS relationship dataset (as-rel2) [108]. We observe that the vast majority of ASes (62,768) are stubs, without peering or customer relationships. These ASes would not be involved. The disadvantage is that if an AS controls an object in which it can arbitrarily inject other child ASes, and therefore easily fake a non-existing link, a BGP prefix hijack via path prepending will still be possible. A very recent proposal, called ASGroups, modifies the AS-Cones concept slightly by improving upon the ASN.1 formal notations, simplifying the validation concept, and changing the opt-out behaviour [109]. One thing to consider is that operators sometimes hesitate to share too many details about their customer cone since this is business-critical information. This might be a disadvantage in the adoption

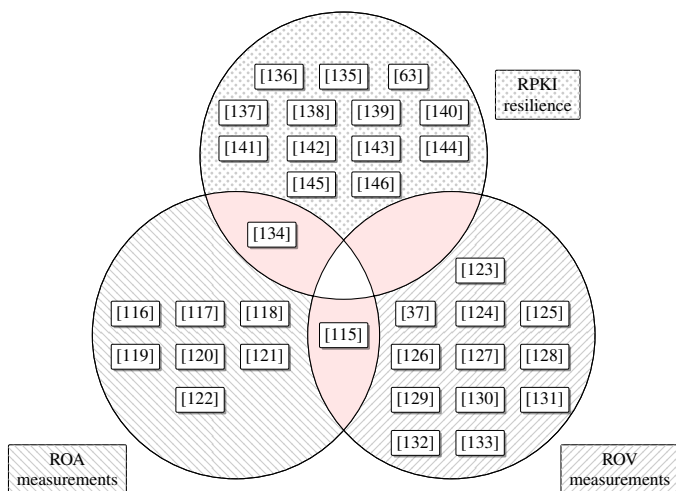


Figure 3. Classification of RPKI research.

process for algorithms relying on such information.

Route Leak Protection (RLP) [110] is one more approach proposed in the Inter-Domain Routing (IDR) working group. It annotates hops in BGP updates with provider-to-customer and peer-to-peer annotations and can therefore detect route leaks. However, it introduces modifications to routers and changes the BGP message format. ASIRIA [111] is an additional proposal to detect route leaks.

Cohen *et al.* [112], [113] proposed in 2015 and 2016, respectively, to use path-end validation. They state that it provides significantly higher security guarantees compared to RPKI in a partial adoption scenario. At the same time, it also requires introducing a new path-end record into the RPKI system. In path-end validation, an AS attests to neighboring ASes via which it can be reached. It does not distinguish between defining upstream relationships (as in ASPA) or downstream relationships (as in AS-Cones/ASGroups). The authors relied on simulations to conduct their research. Instead of relying on simulations, Rodday *et al.* [114] proposed BGPEval, a framework capable of creating large scale testbeds to conduct BGP research.

C. RPKI Research

Many peer-reviewed and non-peer-reviewed publications have been published during the past years. In this work, we aim to categorize existing research directions and highlight the main areas of research within the *origin validation*, specifically RPKI measurement, domain. We analyzed more than 40 papers and found the following predominant categories:

ROA measurements: Identifying the address space covered by ROA objects, *i.e.*, which ASes protect their address space.

ROV measurements: Measuring ROV deployment, *i.e.*, which ASes prevent the propagation of invalid routes.

RPKI resilience: Measuring weak RPKI components and features, *i.e.*, which RPKI deployment scenarios may harm the routing system.

Each category will be explained in detail throughout the following sections. An overview is illustrated in Figure 3.

III. ROA MEASUREMENTS

Within the ROA measurement domain, we identify the following main fields: (i) How to measure ROA coverage, (ii) The problem of the max-length attribute (loose, minimal and hanging ROAs), and (iii) The use of AS0 ROAs. Each publication is listed in Table III, sorted by year, which is linked to the deployment rate at the time of publication. We are able to observe a steady increase in the beginning of RPKI and a steep pick-up of prefix space covered by ROAs during the past couple of years.

A. ROA Coverage

The first scientific publication dealing with measurements of RPKI was published in 2012 by Wählisch *et al.* [116]. They compare the currently available ROAs (which covered roughly 2% of the address space at the time, see Table III) to BGP updates for April 2012 and show that 20% of the verifiable routing table is invalid. A closer analysis attributed most of the invalids to misconfigurations, mentioning that operators were not very familiar with the technology yet. Since the default policy for RPKI was not yet to drop invalid announcements, not much harm was done by these misconfigurations. However, they became a significant hurdle, hindering the adoption of RPKI as the content was considered inaccurate.

In the following years (March 2012 to August 2014), a comprehensive study was performed by Iamartino *et al.* [134]. The authors use every two-hour snapshots of historical BGP data and hourly snapshots of RPKI ROA data over the course of two years to perform validation for each point in time. They find the number of prefix space that is covered by ROAs to increase from 2.05% in the beginning to 5.41% in the end of their measurement window. That finding illustrates that RPKI was slowly being adopted, at least ROAs were gradually created by operators. However, they also found that 80% of prefix space covered by RPKI invalid ROAs was still reachable, as these invalid ranges themselves were covered by RPKI valid or not-found prefixes. Such a finding implies that ROV was not widely deployed yet and filtering of invalid routes did not happen on a large scale. They recommend dropping RPKI invalids as it is safe to do, and their results are consistent with [147], [148]. Kloots [148] studied how many prefixes would be dropped by enabling RPKI ROV and how much traffic is running via these RPKI invalid prefixes that would otherwise be lost.

Wählisch *et al.* [117], [149] continued their analysis in 2015 by analyzing which percentage of the Alexa top 1M domains were protected by RPKI. They found only 6% of webserver prefixes to be covered by RPKI. Out of the covered 6%, 0.09% were RPKI invalid, again most likely due to misconfigurations. An in-depth analysis revealed that more popular websites were less secure than less popular ones since they were hosted in Content Delivery Networks (CDNs), which did not yet deploy RPKI. Out of 199 ASes classified as CDNs, only one CDN with multiple ASNs was found to have four prefixes covered by RPKI.

An important takeaway from these previous studies is the need for monitoring systems. An ongoing RPKI deployment

monitor is provided by National Institute of Standards and Technology (NIST) [36], the MANRS ROA Stats Tool [150], or the Cloudflare RPKI monitoring tool [151] to observe the current deployment status easily.

A longitudinal RPKI analysis was published by Chung *et al.* [119] in 2019. They look at eight years of RPKI data, containing all ROAs ever published, and correlate the data with BGP announcements from public collectors at any point in time. Code and analysis results can be found in [152]. They enrich the public BGP dataset with a private dataset from the Akamai CDN, which they strip of all (private) BGP announcements to avoid a biased view of routing [153]. On February 20, 2019, they report RIPENCC (16.04%) to have the highest percent of ASes that have VRPs published, followed by LACNIC (9.33%), APNIC (8.14%), AFRINIC (3.30%), and ARIN (1.47%). The overall number of prefixes covered by RPKI is increasing throughout all five RIRs, but the uptake differs broadly among RIRs. Overall, they report 12.1% of IP address space to be covered by ROAs in 2019, which has dramatically increased since then.

B. Loose, Minimal, and Hanging ROAs

Iamartino *et al.* [134] not only worked on ROA coverage but also explored the causes for the RPKI invalids they observed. It was assumed that most invalid ROAs were caused by misconfigurations since operators did not fully understand how to use the RPKI yet. Throughout the period from March 2012 to August 2014 the number of RPKI invalids due to invalid max-length dropped from 61% to 54%. Invalids due to an invalid ASN dropped from 24% to 18%, but the invalids due to incorrect max-length and incorrect ASN rose from 15% to 27%. It is evident that RPKI tooling still needed to mature and training of operators was necessary.

Loose ROAs. In 2017, Gilad *et al.* [115] performed a study on ROAs and found roughly 30% of prefixes covered by insecure, which they call *loose*, ROAs. Such ROAs are badly issued and therefore the organizations remain vulnerable to prefix hijacking. A ROA is *loose* if the max-length attribute of the ROA allows for more specific networks than the ones that are announced in BGP. For example, AS47065 makes the following BGP announcement:

```
147.22.0.0/16, AS_PATH: AS47065
```

but signed a ROA that would allow for more specific prefixes to be announced in BGP since the max-length attribute (/16-24) is too coarse-grained, see Section II:

```
ROA: 147.22.0.0/16-24, AS 47065
```

As the RPKI only performs *origin validation*, a subprefix hijack with can be used to attract traffic. The attacker AS61574 could announce the following subprefix:

```
147.22.1.0/24, AS_PATH: AS61574 AS47065
```

The origin at the very right side of the AS path attribute is allowed to issue such an announcement and the attacker AS61574 claims to have received that announcement from AS47065 and forwards it to its peers. The announcement is

Table III
COMPARISON OF ROA MEASUREMENT METHODOLOGIES SORTED BY YEAR

Reference	Measurement Period	Longitudinal	IPv4 Prefix Space Covered by ROAs [%]
Wählisch <i>et al.</i> [116], 2012	April 1–30, 2012	✗	2.00
Iamartino <i>et al.</i> [134], 2015	March 2012–August 2014	✓	5.41
Wählisch <i>et al.</i> [117], 2015	several weeks in 2014/2015	✗	6.00
Gilad <i>et al.</i> [115], 2017	July 2016	✗	6.50
Gilad <i>et al.</i> [118], 2017	June 1, 2017	✗	7.60
Chung <i>et al.</i> [119], 2019	2011–2019	✓	12.10
Hlavacek <i>et al.</i> [120], 2021	February 26, 2019	✗	14.16
Li <i>et al.</i> [121], 2022	January 1, 2022	✗	35.00
Oliver <i>et al.</i> [122], 2022	June 2019–March 2022	✓	35.00

perfectly valid in RPKI terms and would, therefore, propagate via the Internet infrastructure even with Route Origin Validation deployed. Already in 2011, during the early days of RPKI deployment, the RIPE NCC team published a blog post that highlighted challenges because of the max-length attribute [118]. There are, however, legitimate reasons to create *loose* ROAs. One example is traffic engineering mechanisms that must be enabled and disabled when necessary. Another example is Distributed Denial of Service (DDoS) mitigation solutions [117]. These mechanisms need the possibility to issue more specific BGP announcements. Since the RPKI with its *out-of-band* mechanism is relatively slow in propagating updated RPKI information [144], such ROAs have to be issued beforehand. To tackle the problem of misconfigured ROAs and raise awareness among operators, the authors publicly released the tool ROAlert, which is not operational anymore. When it was available, it used WhoIs data to send emails to operators warning them of their *loose* ROAs. One further obstacle to adoption is upward and downward dependencies. Some organizations might wait to issue ROAs for their larger covering prefixes since those ROAs would render subprefix announcements from the same subnet that are lent to customers as RPKI invalid. Therefore, one has to coordinate and wait before issuing ROAs for larger address blocks. They suggest using wildcard ROAs, which basically ‘punch holes’ into a large network, allowing any other AS to announce those sub-prefixes, but the idea was never adopted.

Minimal ROAs. During the same year, 2017, Gilad *et al.* [154] published a follow-up paper studying vulnerabilities caused by the max-length attribute. They suggest only to use minimal ROAs, meaning ROAs that only cover the announced prefix ranges, and point out that RFC7115 already suggests the use of minimal ROAs. Moreover, RFC6907 details use cases and gives explicit examples of how ROAs should be created [155]. Also, RIPE NCC [118] had published advice on how to use the max-length option in 2011. But since reality differs from how things are theoretically supposed to be done, many ROAs are not minimal in practice. The authors developed a Python script that fetches all ROAs and BGP updates at a given point in time and renders all ROAs minimal such that they meet exactly the announced BGP announcements. The tool can be found in [156]. They further recommend only showing the max-length attribute in

GUIs of RIR portals to experienced users and making the creation of minimal ROAs the default policy at RIR interfaces. They observe a 23% increase in VRPs if all ROAs were adjusted to be minimal. All findings and recommendations are summarized in RFC9319 [157].

Chung *et al.* [119] focus on VRPs instead of ROAs, which is the output of a relying party software after the validation of ROAs is performed. The counts of ROAs and VRPs are not directly comparable, as a ROA might contain multiple prefixes, while one VRP is created per prefix. Therefore, one observes a higher count of VRP entries when validating a set of ROAs. Spikes in VRPs were observed two times during the measurement window due to ROA deaggregation. Gilad *et al.* [154] suggested to use minimal ROAs, which would increase the number of VRPs. At APNIC, the effects of such a change were seen when a new management system was implemented. An error led to an uptake of more than 13,000 VRP entries, more than doubling the previously present amount. Moreover, Chung *et al.* found that at the beginning of RPKI, as many as 20.76% of RPKI-covered BGP announcements were invalid, as suggested by earlier research. This situation changed, and due to training and monitoring services provided by RIRs, the share of RPKI invalid prefixes covered by other prefixes dropped to 2.25%–5.39%. A further decrease happened in September 2018, possibly due to Internet Exchange Points (IXPs) who adopted RPKI as a service, forcing participating ASes to fix their ROAs to avoid their routes being dropped. The reasons for invalid announcements found by earlier research were reaffirmed: BGP announcements are too specific, and therefore the issued ROAs do not cover them (48%–51.5%) or the wrong ASN is used to announce the prefix in BGP, compared to the one inserted in the ROA. The authors suggest some plausible causes: First, both ASNs are under management of the same company, but the ROA was not updated. Second, a prefix was suballocated to a customer without updating the ROA. Both are frequent reasons for invalid BGP announcements. Third, but more rarely, DDoS protection services were used by outsourcing scrubbing to external parties without updating the ROAs. Fourth, the invalid announcement was caused by other reasons, possibly also actual hijacks. The duration of prefix hijacks is typically much shorter compared to misconfigurations. The authors report the use of max-length decreasing to 11.2% of prefixes in

VRPs. Most invalid BGP announcements seem to stem from misconfigurations in the BGP/RPKI interplay and not from hijacking attempts. The authors put the suggestion of earlier research [154] into perspective that called for the removal of the max-length attribute by arguing that there needs to be a tradeoff considered. RPKI does still protect against misconfigurations, even with a non-minimal ROA having a longer max-length. It does not protect against intentional hijacks that forge the AS-path. However, since the presented methodology is not performing well enough to identify intentional hijacks, further research is required in this domain to safely differentiate intentional hijacks from misconfigurations.

Hlavacek *et al.* [120] proposed in 2021 and in an extended version in 2022 [158] an approach to further differentiate between misconfigurations of ROAs and actual traffic hijacks. They believe that invalid ROAs are one of the main factors preventing wider ROV adoption on the Internet. By comparing differences between the actual BGP announcements and RPKI ROAs, they find discrepancies and conclude that the majority of conflicts are due to misconfigurations and not actual hijacks. They further analyze BGP hijacks that lay in the past and confirm that BGP hijacks are usually short-lived. For all the conflicts monitored, they repeat the process introduced by Gilad *et al.* in ROAlert [115] where operators were notified via WhoIs data. 760 emails were sent in the first round, followed by another 180 additional emails to newly identified ASes at a later stage. The questionnaire the operators filled out showed that most conflicts were created due to misconfigurations, some promised to fix the errors, while others simply acknowledged their existence. The tool is available online [159].

Hanging ROAs. Li *et al.* [121] pick up the problem of the max-length attribute in 2022 and propose so-called *hanging* ROAs. They apply a bitmap-based encoding scheme that compresses the total size of ROA payloads in RRDP by 26.6% compared to the previously introduced minimal ROA proposal in [154]. Their proposal also reduces the synchronization cost, measured in time using a 10 Mbps link, by 41.3% and 50.4% for the currently used max-length and minimal ROA approach, respectively. Changes in the ROA Protocol Data Units (PDUs) require adoption from RIRs, CA-software developers, and RP software developers. Given the small quantity of data that needs to be transferred during every update interval, it is questionable if the effort of improving compression is worth the introduced complexity required by transitioning from one method to another. Unfortunately, the authors do not state what a transitioning period to introduce the hanging ROAs could look like.

C. AS0 ROAs

Oliver *et al.* [122] investigated the usage patterns of 712 prefixes from Spamhaus' "Don't Route or Peer" (DROP) list [160], which operators frequently use to identify maliciously used address space. They found that for 32% of prefixes' IRR entries were created a month before the prefixes were added to the blacklist, highlighting the fact that the IRR is not a reliable source for filtering information since it can be easily manipulated. Moreover, they find that

attackers do not usually target RPKI-protected prefix space but instead rely on unallocated or unannounced prefix space. They suggest using the AS0 ROA to further increase the security of the inter-domain routing infrastructure. An AS0 ROA allows RIRs to protect unallocated prefix space, which would otherwise be rendered RPKI unknown and could therefore be easily misused. Also, organizations are able to create AS0 ROAs for their address space in case it is currently unused. If such prefix space is not added into an AS0 ROA, getting hijacked with a path prepending attack would still be possible. Three organizations own about 70.1% of this unrouted but covered by non-AS0 ROA prefix space. It would, therefore, be comparably easy to improve security by the effort of a few organizations. They conclude with a call for a re-evaluation of the use of AS0 ROAs by both operators and RIRs.

Overall, we have seen that ROA coverage measurement methodologies were implemented in automated monitoring tools, which provide us nowadays with an up-to-date overview of the current ROA coverage state. Moreover, the initially designed ROA structure leaves room for improvement as *loose* ROAs lead to vulnerabilities that attackers could potentially exploit. *Hanging* ROAs were proposed to improve compression and reduce synchronization costs. Also, AS0 ROAs are not used to the extent necessary to protect unassigned and assigned but announced address space.

IV. ROV MEASUREMENTS

Route Origin Validation, *i.e.*, whether an AS decides on route preferences based on RPKI, can be measured orthogonally to the question whether attestation objects (ROA) are deployed. This section focuses on previous works dealing with ROV. Two main questions arise in this context. First, do ASes deploy ROV? Second, which ASes deploy ROV?

RPKI ROV deployment can be measured on the control plane or data plane, or both. Methods can further be classified into active and passive measurements as well as controlled and uncontrolled experiments. While *active* measurements rely on injecting data into the system under study, *passive* measurements use data that was recorded independently of the actual experiment. *Uncontrolled* experiments refer to a setup that is not under the control of the experimenter. *Controlled* experiments require control of some parameters that may influence the experiment. We summarize work discussed in this section in Table IV.

A. Control Plane Measurements

In 2017, Gilad *et al.* [115] measured the ROV adoption rate using passive control plane measurements. The authors 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 origin 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

Table IV
COMPARISON OF ROV MEASUREMENTS SORTED BY YEAR

Reference	Measurement Period	Plane		Experiment Type		Approach	Longitudinal
		Control	Data	Controlled	Uncontrolled		
Gilad <i>et al.</i> [115], 2017	July 2016	✓	✗	✗	✓	BGP dump analysis	✗
Reuter <i>et al.</i> [37], 2018	February 20–27, 2017	✓	✗	✓	✗	BGP dump analysis with route injection	✗
	May 11–17, 2017	✓	✗	✓	✗		
	August 1–7, 2017	✓	✗	✓	✗		
Hlavacek <i>et al.</i> [124], 2018	February 2017–June 2017	✓	✓	✓	✗	Traceroute & TCP SYN	✗
Cartwright-Cox [125], 2019	N/A	✗	✓	✓	✗	ICMP scans	✗
RPKI WebTest [129], 2019	on-demand, discontinued	✗	✓	✓	✗	HTTP	✗
Rodday <i>et al.</i> [130], 2019	August 22–29, 2019	✓	✗	✓	✗	Control plane extensions	✗
Cloudflare [133], 2020	on demand	✗	✓	✓	✗	HTTP	✗
Testart <i>et al.</i> [131], 2020	April 1, 2017–January, 22 2020	✓	✗	✗	✓	Statistical approach	✓
Huston <i>et al.</i> [127], 2020	June 1–20, 2020	✗	✓	✗	✓	HTTP	✗
Rodday <i>et al.</i> [126], 2021	July 2–19, 2021	✗	✓	✓	✗	HTTP & Traceroute	✗
Chen <i>et al.</i> [128], 2022	May 1–31, 2021	✗	✓	✗	✓	Traceroute	✗
Hlavacek <i>et al.</i> [132], 2023	June 8 and 10, 2022	✓	✓	✓	✗	BGP dump + Traceroute	✗
RoVista [161], 2023	December 24, 2021–September 12, 2023	✗	✓	✗	✓	IP-ID side-channel technique	✓

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. The authors also surveyed operators, which 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. Furthermore, they introduce the concept of collateral benefit and collateral damage. *Collateral benefit* refers to an AS that does not perform ROV sitting behind another AS that implements ROV and drops invalid announcements. Since the upstream performs ROV, the downstream AS benefits from the filtering and does not fall victim to hijacking attacks. *Collateral damage* refers to decisions made by the upstream that negatively impact the downstream peer. The first scenario is disconnection of the downstream peer if the upstream selects the RPKI invalid route in its best-path selection process and the downstream discards it afterward. In this scenario, the downstream could no longer send traffic towards that prefix range. The second scenario is a hijack event, in which the upstream receives both, the more specific hijacked route and the less specific original covering prefix and forwards both to the downstream peer. The downstream peer discards the hijacked more-specific and sends the traffic to the upstream, destined for the less-specific covering prefix. But since the upstream is not performing RPKI filtering, it would continue sending the downstream peer traffic to the attacker. Therefore, although the downstream deploys RPKI filtering, it does not affect how traffic is routed.

Reuter *et al.* [37] reproduced the work by Gilad *et al.* [115] and found that the results heavily relied on the chosen set of BGP collectors. The previously described phenomenon of *collateral benefit* led to measurement errors in the form of wrong attribution of ROV in prior work. 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 the AS operators confirmed. The method was deployed in a live monitoring system [162] and identified 118 ASes as deploying ROV in March 2021, but it is not operational anymore. To be rigorous and exclude false positives, the methodology has two assumptions: The *connected assumption* restricts the analyzed BGP paths to the ones that are directly connected to the announcing AS. That is to make sure that BGP paths only have a length of two, the announcing AS plus the AS under test, and once ROV filtering is identified, it can be attributed to the AS under test. The *visibility assumption* requires the AS under test to be a vantage point and, as such, export routes to a route collector. These rather strict assumptions only allow testing a minimal number of ASes. The PEERING testbed (the AS used to announce routes during experiments) has a limited number of peers and of those only some export routes to route collectors.

Rodday *et al.* [130] extend upon [37] and propose to remove the *connected assumption* for BGP paths that are entirely comprised of vantage points. The extension can only identify the first AS on the path as ROV filtering. Multiple ROV filtering ASes would cover each other, making pinpointing impossible. The authors also propose to extend the previous methodology by relaxing the *visibility assumption* and including isolated non-vantage points. If both ASes, before and behind the AS under test, are exporting routes to route collectors, and the AS behind the isolated non-vantage point is known to be not ROV filtering, a conclusion on filtering for the AS under test can be drawn.

Testart *et al.* [131] introduce a passive approximation methodology of how ROV could be measured on the control plane. The methodology aims to identify 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 fewer RPKI invalid routes compared to the full-feeders. The resulting cluster comprises 21 ASes 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. It was possible to cluster and differentiate the set of RPKI enforcing ASes since there were so few at the time of performing this research, and the vast majority did not perform ROV. A likely problem with this methodology is that the more ASes adopt RPKI, the less statistical difference will be seen in the measurements, and therefore, it will not be possible anymore to flag ROV-enforcing ASes.

Gray *et al.* [163] propose BeCAUSE, an algorithmic framework for inferring network properties based on Bayesian computation for ASes. They apply BeCAUSE to pinpoint ROV-enabled ASes.

Du *et al.* [164] used BGP collector data in 2023 to rank ASes accordingly to their number of propagated RPKI invalids. They emphasize that these ASes should deploy ROV most urgently to shrink the number of RPKI invalids in the wild.

An attempt to compare existing ROV measurement proposals in automated testbeds was presented in [165].

B. Data Plane Measurements

In 2018, Cartwright-Cox [125] presented a *data plane* approach to measure RPKI adoption. He identifies ROV filtering by analysing replies to two types of measurement probes, those that are connected via an RPKI valid IP prefix and those that are connected via an RPKI invalid prefix. In detail, each probe scans the entire IPv4 address space using the Internet Control Message Protocol (ICMP). If all probes receive a reply from a host, the upstream of this host is considered not deploying RPKI. If only the probe connected via an RPKI valid IP prefix receives a reply, it is assumed that the ISP of the replying host is deploying RPKI. This method does not allow identifying the AS that is filtering. Furthermore, to exclude false positives because of ICMP filters on specific paths, a hop-wise AS analysis would be needed. Updates to this study were presented at NLNOG Day in September 2019, and at RIPE 80 in May 2020 [166].

A method similar to the method by Cartwright-Cox [125] has been applied by Huston *et al.* [127]. 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 a 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 [125], Huston *et al.* [127] 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 a few transit providers probably enabled filtering rather than many stub networks.

The RPKI WebTest [129] operated by RIPE is a website to raise awareness of ISPs that do not deploy ROV. It tests whether the local 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 a user's web browser to each IP address.

If both requests succeed, the ISP is considered not rejecting invalids yet. If the IP address from the invalid prefix could not be reached but the IP address from the valid prefix could, the ISP is considered deploying ROV. The RPKI ROV project by Cloudflare [133] uses the same methodology. In addition, it provides the possibility for operators to issue a pull request on Github, updating the list of ASes that enforce RPKI ROV with their own ASN. However, since the process is manual, only a few operators are expected to use this opt-in procedure. Since both HTTP-based tests have the same methodology in common, both suffer from the problem of wrong attribution. Such a wrong attribution happens if an upstream AS performs the filtering, but the website shows that the ISP of the user actually performs RPKI ROV.

Hlavacek *et al.* [124] compare current ROV measurement methodologies. They repeat the controlled *control plane* measurements of [37] 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 Transmission Control Protocol (TCP)-SYN packets sent to the top 1,25M Alexa domains. They find 4 ASes to enforce ROV 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. Shulman *et al.* [167] applied the methodology in June 2022 and found a significantly higher amount of ASes (37.8%) enforcing ROV.

A thorough investigation identifying ROV-enforcing ASes has been published by Rodday *et al.* [126] in 2021. In a nutshell, the authors extend prior work by (i) relaxing the connected assumption of [37] 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 differentiating between partially and fully filtering ASes. Using data from 5,537 vantage points in 3,694 ASes in June 2021, the proposed method detected 207 unique ASes performing ROV: 10 with strong confidence, 12 ASes with weak confidence, and 184 ASes indirectly adopting ROV ASes via filtering by IXP route servers.

A measurement campaign considering both IPv4 and IPv6 prefixes was run by van Hove *et al.* [168], [169] in 2022.¹ They announced RPKI valid prefixes and RPKI invalid prefixes, such that RPKI invalid prefixes were more specific than RPKI valid prefixes. Each group of IPv4 and IPv6 prefix was announced from a different geographic region connected via a different upstream ISP. The assumption of the authors is that traffic is only routed to the covering prefix, in case of full ROV deployment. They found 75% of traffic to be routed to the valid and 25% of traffic to be routed to the invalid prefix ranges.

In 2022, Chen *et al.* [128] also tried to remove the connected-assumption proposed in [37] in order to widen the measurement scope. Instead of announcing prefixes using their own infrastructure, they rely on roughly 6,000 publicly available RPKI invalid BGP prefix-origin pairs, collected via RIPE RIS [170] and Routeviews [35] and accessed via

¹It is worth noting that the description of the setup in the RIPE Labs article [168] differs from the actual presentation [169].

BGPStream [171]. The basic idea is similar to [115]. The prefix-origin pairs are announced by a variety of different ASes. After applying filtering for multi-homed prefixes and prefixes covered by other legitimate announcements, they identify other RPKI-valid prefixes originated by the same origin ASes. Depending on the measurement day, the authors successfully identified 350-500 cases. These prefix pairs are used for further measurements. It should be noted that these prefix pairs are not required to have anything in common other than the origin AS and that one is RPKI invalid, the other RPKI valid. This leads to ambiguity when an origin AS purposefully announces these two different ranges to two different upstreams, creating different AS paths, *e.g.*, due to traffic engineering [37]. For each prefix pair, 200 randomly selected RIPE Atlas [172] and perfSONAR [173], [174] probes are used. ZMap [175] is tasked to identify active hosts within the prefix ranges, and from each of the 200 probes a traceroute is run against the active hosts in the valid and invalid prefix ranges. Each traceroute path is then mapped to an AS path. According to the authors, two possibilities arise: (i) both paths for the valid and invalid range are equal, or (ii) they are not. If they are not equal, the divergence is attributed to ROV, and the probabilistic Stein Variational Gradient Descent (SVGD) model is used to obtain a probability per AS. The basic idea is similar to the work by Grey *et al.* [163]. The results show that 28% ASes deploy ROV ($n=3107$), 43% do not deploy ROV ($n=4716$), 3% partially deploy ROV ($n=357$), and 26% are unknown ($n=2894$), out of 11,074 ASes on BGP paths forwarding RPKI-valid announcements. The methodology is evaluated based on a ground-truth dataset from is-bgp-safe-yet [133] and shows 100% precision and 100% recall.

Such uncontrolled measurements can be heavily skewed by traffic engineering and other factors that are out of control of the experimenter, as Reuter *et al.* [37] have shown. Therefore, it is questionable whether the approach indeed correctly flags ROV-enforcing ASes or simply ASes that perform traffic engineering for some unrelated reason. It is also known that many ASes benefit from the filtering of transit providers and IXPs.

Hlavacek *et al.* [132] perform controlled control plane and data plane measurements focusing on ROV implemented on routeservers at IXPs. Measurements are run on June 8 and 10, 2022. Their control plane data is based on Routeviews [35] BGP dumps, and data plane measurements are obtained by running traceroutes via RIPE Atlas, comparable to [124], [126]. ASes that provide an indication of ROV-filtering are called divergence points. The authors derive seven categories with different confidence levels of ROV-enforcement and report that more than 27% of the ASes filter RPKI invalid routes. Moreover, IXPs are found not to block hijacks, even though IXPs themselves are performing ROV, since many peers use direct peering sessions at the IXP facilities that do not undergo filtering. Validation is performed manually for the 15 tier-1 providers, with 11 correct inferences and 4 non-verifiable. No evidence of false negatives could be found. They also compare to Cloudflare measurements [133], which results in a 75% overlap, and APNIC measurements [127], with 79% overlap.

Another publication dealing with the identification of ROV-enforcing ASes has been published by Li *et al.* [123] in 2023. Measurements were run every four hours from December 24, 2022 to September 12, 2023. The approach is called RoVista and utilizes the IP-ID side-channel technique to infer connectivity between two remote hosts. IP-ID is a field in the IPv4 header. The authors can distinguish between no filtering, inbound filtering, and outbound filtering. The major advantage is that the approach does not require control of vantage points within ASes under test. Therefore, the authors were able to significantly increase the scale of ASes that can be measured (28K ASes). They rely on uncontrolled measurements as they derive a list of RPKI invalid BGP announcements from BGP collector data. First, they use ZMap [175] to find test nodes within the RPKI invalid prefix ranges that reply to TCP SYN packets correctly. According to their observations, nearly 0.7% of global routing table prefixes are RPKI invalid. After applying additional sanitization methods, they derived 31 test nodes residing within invalid prefix space. Next, they attempt to obtain virtual vantage points via ZMap by querying nodes that reply to a TCP SYN/ACK packet with an RST packet. They find 1,396,407 virtual vantage points that cover 28,314 ASes. RST packets allow tracking of the IP-ID counter of the sending host. The idea is to send traffic between pairs of test nodes and virtual vantage points to observe whether the IP-ID counter increases. Since spoofed data-plane packets are used to trigger an increase of the IP-ID counter in the RST packets, they are able to infer whether a successful connection between the test node and the virtual vantage point happened. Based on these results, they calculate an ROV protection score, which is the percentage of test nodes inaccessible from any virtual vantage points within the same AS due to outbound filtering. RoVista is designed to derive an ROV protection score per AS. However, the score might indicate full protection while the AS itself is not filtering based on RPKI since it is sitting behind RPKI filtering ASes. Overall, they find 63.8% of all ASes to have derived benefits from RPKI filtering ASes, and 12.3% to be fully protected by RPKI ROV.

In summary, ROV measurements can be conducted using controlled or uncontrolled experiments on the control or data plane. Attribution of ROV filtering remains a challenging task, though, because of *collateral benefit*, *i.e.*, downstream ASes benefit from filters deployed at upstream ASes. Deploying ROV measurement methods as part of public monitoring projects and making results available becomes popular. Overall, ROV deployment is increasing, following the increase in ROA deployment. The rate at which ROV adoption occurs remain an open research field.

V. RPKI RESILIENCE

Current research focuses on four topics to better understand RPKI resilience. (i) Centrality of the infrastructure, (ii) inconsistencies of Relying Parties, (iii) circular dependencies and usability, and (iv) attacks. In this section, we discuss all four topics. At the end of this section, we also provide an overview of RPKI infrastructure outages.

Table V
COMPARISON OF RPKI RESILIENCE RESEARCH SORTED BY YEAR

Reference	Topic
Cooper <i>et al.</i> [137], 2013	Whacking of ROAs
Heilman <i>et al.</i> [135], 2014	Consent via <i>.dead</i> object
Iamartino <i>et al.</i> [134], 2015	LACNIC/APNIC outage for 9 months
Liu <i>et al.</i> [141], 2015	RPKI risks categorization
Hari <i>et al.</i> [136], 2016	Blockchain proposal
Yan <i>et al.</i> [138], 2018	CA-software suggestions
Kristoff <i>et al.</i> [140], 2020	90% of RPs not falling back to rsync
Shrishak <i>et al.</i> [142], 2021	Threshold-based delegation
Friedemann <i>et al.</i> [63], 2022	Comparison of RP software
Hlavacek <i>et al.</i> [143], 2022	Stalling of RP software
van Hove <i>et al.</i> [139], 2022	Vulnerabilities in RP software
Hlavacek <i>et al.</i> [145], 2022	Attacking DNS to harm RPKI
Fontugne <i>et al.</i> [144], 2023	Delays in the RPKI ecosystem
Hlavacek <i>et al.</i> [146], 2023	RP threshold analysis

Liu *et al.* [141] also summarized research within the RPKI resilience domain and categorized them into technical, economic, and political risks. The summary is also available as an IETF draft [176]. We provide an overview of previous works in Table V.

A. Centrality of the Infrastructure

The RPKI attestation model follows a chain of trust. As such, parents within the trust chain can revoke certificates, thereby invalidating all certificates below them. In 2008, the Internet Governance Project (IGP) warned that those who seek to regulate the Internet would target RIRs [177], and in 2011 the RIPE-NCC sought clarity on how to respond to foreign court orders that force RIRs to withdraw assignments of Internet resources [178].

Following up on questions raised on the centrality of the RPKI infrastructure, Cooper *et al.* [137] worked in 2013 on issues dealing with the trust and power that is given to each RPKI authority by design. It is generally assumed that RPKI authorities behave in alignment with expectations (*e.g.*, they do not revoke child certificates without reason), but what happens if an RPKI authority becomes rogue? Any CA can unilaterally revoke certificates that will impact the business of a descendant. The authors argue that “there is ample evidence of authorities [...] being hacked [179]–[181], misconfigured [182], or compelled by government agencies to delete information (*e.g.*, Domain Name System (DNS) takedowns [183]) or attest to bogus information [184]” [137, reference numbers changed to the reference list of this paper].

Whacking of ROAs. Cooper *et al.* [137] highlight the *whacking* of (great-)grandchildren ROAs and beyond. It is a term coined to describe a process in which an RPKI authority overwrites an existing Resource Certificate (RC) to contain all surrounding address space but the one to be whacked. Consequently, cryptographic validation of the previously valid ROA will fail, and the targeted ROA will become invalid. RP software will therefore not include such ROA into the set of Validated ROA Payloads (VRPs) anymore and the BGP announcement containing that address space in the ROA will

become RPKI unknown, since there is no covering ROA anymore. In case there is a covering ROA of another organization, the BGP announcement might also become RPKI invalid. The attack does not cause collateral damage since it only invalidates the targeted ROA. If the whole RC was revoked instead, possibly more ROAs would be whacked, causing more significant damage to other parties.

A year later, in 2014, Heilman *et al.* [135] follow up on the previous work that the power of RPKI authorities might be too great [18], [21], [137], [185]. They argue that transparency mechanisms would be helpful. This would give the community the option to notice when a certificate authority misbehaves. Social and legal pressure could be created if such mechanisms were in place. However, it is difficult to differentiate between revocations due to disputes, censorship, or business arrangements. Security audits might help to detect malicious behavior and increase transparency. Nowadays, RIRs, such as the RIPE NCC, conduct security audits [186].

Since ROAs can be whacked without the issuing party’s consent, Heilman *et al.* [135] propose introducing a *.dead* object into the RPKI ecosystem. The idea is that from the ROA, which is represented as a leaf in the RPKI tree, the tree is walked towards the root, and every involved entity has to sign a *.dead* object showing their consent that the RC covering the ROA is altered, which results in the revocation of downstream resources. This whole procedure would add transparency since one would know whether everyone agreed before the change, including revocation, happened. However, it would not have any hard security implications as the issuing party still controls the publication point. It could simply ignore the *.dead* objects and go ahead to change the RC immediately. The process is also very cumbersome and would paralyze the RPKI ecosystem. Therefore, it was adopted.

Cross-country certification. Another problem that Cooper *et al.* [137] identified is cross-country certification. Address space is often split into smaller chunks, which are reassigned to other organizations in different countries and, therefore, legal jurisdictions. The centrality of the five RPKI trust anchors comes with the drawback that one country’s court might decide on a case of an ISP located in another country. That makes it difficult, if not impossible, for organizations in another country to push their interests judicially.

Blockchain. To move from a centralized model towards a distributed model, Hari *et al.* [136] proposed in 2016 to use blockchain technology instead of a hierarchical PKI architecture. The chain is designed to create a block of 1MB every 10 minutes. That allows for 3-7 transactions per second. If only information contained in a ROA is captured inside the blockchain the model would theoretically work on average but already fail during peak times (*e.g.*, if an AS changes many ROAs at the same time). If, however, BGPsec is to be deployed, not only ROAs, but also every announcement would have to be tracked in the chain to provide sufficient security guarantees. This is unfortunately not possible with the current technology, as there were already 9,000 changes in BGP per second in 2016. Also, in time, the amount of data that would need to be stored in the chain grows and becomes hard to use.

Using blockchain technology in the inter-domain routing ecosystem is not entirely new. Haeberlen *et al.* [104] proposed the use of blockchain technology as early as 2009 to create a secure log of BGP traces to analyze problems in BGP. Another blockchain approach was proposed in 2018 by Paillisse *et al.* [187]. Here, a proof-of-stake is used in the consensus algorithm. Larger ASes with more space will be more powerful in such a model. Mastilak *et al.* [188] summarize existing blockchain technologies and their applicability to inter-domain routing.

A bigger problem with using blockchain technology is the design principle: BGP is considered an information-hiding protocol. In a blockchain, every transaction is publicly visible. Therefore, many operators might not be willing to move to such a model.

Threshold-based approach. Shrishak *et al.* [142], [189] pick up the risk of legal restrictions for each RIR and suggest in 2020 and 2021, respectively, to use a threshold-based approach for resource delegation instead of giving each of the five Trust Anchor (TA) the power to delegate resources solely. In detail, they propose to use the Dalskov *et al.* [49] protocol as a threshold signature model. In such a model, RIRs would have to issue and sign resource delegations together. Therefore, it would no longer be possible for one RIR to delegate address space allocated to another RIR. It would also no longer be possible for one RIR to take down address allocations if a local court issues an order. Moreover, the risk of a single compromised RIR would be reduced significantly, and single RIRs can be unavailable as resource delegation is only allowed with a majority. In a threshold-based model, three out of five must collaborate to issue a resource delegation. Collaboration, in a malicious sense, is, therefore, improbable.

The RPKI currently provides two models: hosted and delegated. In the hosted model, the RIRs hold the private keys of all descendants and can access the keys to perform signing. In such a scenario, a threshold-based approach can be deployed. In the delegated model, private keys lay within the AS that runs the delegated CA instance. A threshold-based approach would not work in this scenario.

While theoretically possible, such an approach would require much more complex collaboration between the five RIRs. Also, according to their own evaluation, on average 20,000 signatures per day would be required. The suggested protocol would be able to cope with the cryptography processing during such days but would lag behind during peak hours. That would create delays in the RPKI infrastructure, impacting validation.

B. Relying Party Inconsistencies

Cooper *et al.* [137] point out that a ROA that previously led to an RPKI valid BGP announcement might be missing from the repository. Reasons could be delayed ROA renewal, a corrupted file system storing the ROAs, or the unavailability of the RPKI repository. In such a case, RP software would no longer include the address space into the VRRP, rendering related BGP announcements either RPKI unknown (without a covering ROA) or RPKI invalid (with a covering ROA from another organization). Hence, it is paramount that RPKI

RPAs have access to a complete set of ROAs. If only some information is missing, some RPs might have a different view of the world than others, so such an attack is called the *mirror world* attack.

In addition, Heilman *et al.* [135] highlight that manifests can expire. The RPKI is using manifests to track which items in a repository (*e.g.*, ROAs or RCs) have changed via hashes. A RP does not need to fetch all objects again, only the ones that have changed. But since the RIR controls the repository, it is also possible for them to manipulate the manifest itself, also only for certain RPs, mounting the previously described *mirror world* attack.

Another idea presented in [135] is hash-chained manifests, such that an RP can always track back to the previous manifests. Moreover, manifests should not expire but become stale, and once stale, raise a missing-information alarm. They further suggest that if only manifests were signed, but not ROAs, since a collision-resistant hash would be used in the manifests, less cryptography has to be used and Certificate Revocation Lists (CRLs) become useless. ROAs and RCs would not need to expire anymore. This idea remained theoretical; no Internet draft was submitted to the IETF, and it was not considered for standardization. Additional information can be found in Heilman's dissertation [190]. During the same year, in 2014, Kent *et al.* [191] submitted an IETF draft called 'Suspenders: A Fail-safe Mechanism for the RPKI', detailing and tackling the same problem by monitoring if RPKI objects were rightfully changed.

The most important conclusion drawn is that monitoring systems for the RPKI are necessary. This is also emphasized by Iamartino *et al.* [134], who found in their study from 2015 that the LACNIC and APNIC repositories had expired X.509 certificates at the TA for almost nine months, moving all descending ROAs from valid to unknown. Such bad operational practice shows the urgency of monitoring solutions and how they were not adequately deployed back then. Another RIR, ARIN, creates legal barriers before allowing the use of their TA. They require any relying party to sign an agreement. Hence, RP software does not include the ARIN TA by default, which excludes all ROAs under ARIN from validation. An administrator has to manually add the ARIN TA after signing the agreement. The situation regarding monitoring solutions has improved throughout the last couple of years, and monitoring systems are now in place. However, outages of the RPKI infrastructure do occur sometimes [192]–[195].

An extensive study into relying party software was performed by Kristoff *et al.* [140] in 2019. They operate one child, two grandchildren CAs, and three publication points under one RIR CA to record data of connecting RP software, such as timestamp, IP address, originating ASN, reverse DNS records and RP software. One observation is that one publication point encounters up to 20% less traffic than its parent. Therefore, not all RP software has the complete set of RPKI data. During the one-year measurement window, the number of RP software connectivity increased from 25-100 to 75-250, while Facebook's share of RP software increased from 0 to nearly 70 instances in 2020. Most operators use 1-2 distinct RP instances and deploy RP software themselves instead of using instances

hosted in CDNs. Another interesting finding is that some 20% of RPs are slower than 20 syncs per day. Hence, there is a considerable lag when something within the RPKI is changed until every AS has received the updated resources. Most RP software, accounting for nearly 90% of traffic, are not falling back to rsync, which is contrary to the IETF standardized recommendation.

A work by Friedemann *et al.* [63] compares the seven available RPKI RP software solutions in June 2021 regarding their performance during the validation process. They develop a metric (feature-richness, usability, performance, etc.) and run all validators in the same setting to rank them. Routinator 3000 obtains the first rank, followed by the Fort validator. The authors recommend using these validators for production. As an additional insight, they find RPSTIR2 not to fall back to rsync when a publication point does not offer RRDp. Therefore, this validator delivers significantly less VRPs to a BGP router, leading to different routing decisions. All other validators deliver roughly the same output.

In 2023, a study by Fontugne *et al.* [144] examines the issue of delays in the RPKI infrastructure and conducts two experiments. The first experiment involves announcing a /24 prefix pair for each RIR from an AS that is surrounded by ROV filtering ASes. While the control prefix remains static, the test prefix is swapped regularly to make the BGP announcement RPKI valid or invalid. This experiment was conducted over the course of 11 months. In the second experiment, three /24 test prefixes from RIPE NCC are announced by three different networks, with all ROA states changing daily. They record the time taken for user queries in the RIR portal, ROA signing, ROA publication, and RP validation until the information contained in the ROA is deployed productively within an AS. The study finds that there is significant variation in ROA creation times across different RIRs, ranging from a few minutes to over an hour until the ROAs reach publication points. This variability could be due to differences in the underlying processing mechanisms, such as batch processing. Additionally, ROA deletion takes longer to reflect in BGP compared to ROA creation. As expected, most delays in ROA creation are caused by RP software implementations that pull ROAs from publication points at different time intervals.

C. Circular Dependencies and Usability

Cooper *et al.* [137] point out that a general problem of the RPKI is that it is using TCP/IP as its underlying technology, which creates circular dependencies between BGP and RPKI. Reachability information is propagated via BGP, and if certain routes are unavailable due to RPKI filtering, some RPKI repositories might not be reachable, which might lead to a downward spiral.

Yan *et al.* [138] discussed different scenarios CA software needs to be able to support. While their work covers mostly scenarios CA software is designed to support, they suggest to add alerts to the software in case it is supposed to assign resources that are under the control of another CA. That behavior is theoretically possible in RPKI as any CA could issue RCs for any address block, even if not assigned by

IANA. A problem with this approach is that during a key rollover, precisely this alert would be triggered, which is why they suggest exempting such a scenario from the alerting mechanism. Their proposal only protects against misassignment in CA software. If the operator of the CA software would indeed be of malicious intent, the assignment of address space outside of the control of the CA would be on purpose, and a warning would not yield any benefit.

D. Attacks and Threat Models

A very recent attack on the RPKI infrastructure has been presented by Hlavacek *et al.* [143] in 2022. The public announcement that the whole RPKI is broken sparked a lively discussion on the IETF routing mailing list [196]. An accompanying talk was presented at BlackHat USA 2022 [197]. The idea is to inflict packet loss in specific time intervals that are in sync with the refresh intervals of RP software and stall RPs by creating very long delegation chains that prevent them from fetching ROAs and other RPKI data from publication points, forcing the expiry of cached items. This weakness has already been pointed out in RFC7132 [198]: "An attacker could create very deep subtrees with many ROAs per publication point [...]". Upon expiry, cached items will be removed from the VRP list by the RP, and different routing decisions might be taken by the BGP routers. Via a targeted attack on a publication point of an AS that the attacker wants to hijack, the ROAs of that AS will expire and be removed from the local RP cache, allowing the attacker to hijack the target AS prefix ranges. First, the attacker must know the IP addresses the legitimate RP will use to contact the repository. She runs her own publication point or has access to one. All RPs should also connect to the attacker's publication point. Therefore, the attacker can obtain the IP address of the target RP. Second, the attacker sends many spoofed packets to the targeted repository, containing the source IP of the targeted RP instance. Due to rate-limiting, the repository will soon block the legitimate IP address from contacting the repository. As DNS resolvers are anycasted, the attacker's origin for the spoofed packets must be in the same anycast domain as the targeted repository. They find 47% of repositories vulnerable to this attack. The problem is that whenever the RP tries to contact the repository, another attack must be launched to prevent a successful connection. It is hard to know when precisely an RP will initiate a connection and an attack pattern might become visible after a while. The second component of the attack is based on the SlowLoris attack [199], which opens many HTTP connections towards the target and answers very slowly. They find 53.01% of manifests to have a maximum validity of less than 24 hours. Once a manifest becomes stale, it invalidates all ROAs contained within. Therefore, the goal is to keep a RP from refreshing a manifest for this time period. They suggest only allowing delegation chains up to a depth of 32 to avoid this issue. Subsequently, developers of RP software fixed the issues above and introduced thresholds to avoid stalling of RP software.

Hlavacek *et al.* [146] revisited their earlier work in 2023 and found that RP software is still vulnerable via small changes in

methodology. They point out that there will always be a trade-off between permissive-strict thresholds, leading either to too many failures during valid connection attempts or too little security and, therefore, exposure to attacks during malicious use.

Mirdita *et al.* [200] report that in June 2022, 4,344 relying party software instances were deployed on the Internet. They perform black box testing of relying party implementations and find that Routinator and OctoRPKI had exploitable bugs that the developers consecutively fixed.

Another recent work that performed similar research as [143] was done by van Hove *et al.* [139]. Parts of this work were also published as an IETF draft [201]. It is generally assumed that a publication point and certificate authority behave as expected and do not have malicious intent. If these entities are compromised for some reason, what would be the impact on RP software? The IETF is aware of operational problems with rsync and is already pushing for RRDP instead [202], which is why this research focused on RRDP. Since RRDP uses HTTPS and the data is formatted in XML, they apply the OWASP Top10 REST security vulnerabilities [203] to RRDP connections and XML security considerations to the XML formatter. A GZIP bomb, amongst other attacks is also performed. In total, they run 15 different attacks, of which all validators were at least susceptible to one. During the realm of this research a Coordinated Vulnerability Disclosure (CVD) process led by the Dutch NCSC-NL was conducted. Most vulnerabilities have been fixed during the CVD process. The disclosure process led to a lively discussion on the SIDROPS mailing list [204].

Hlavacek *et al.* [145] continued to study RPKI resilience in 2022 but shifted their focus to the DNS infrastructure. Since the relying party instances lie outside the reach of the experimenters, debugging failures during measurements becomes a significant problem. In order to receive multiple queries, they create nested publication points, requiring the relying party software to contact their infrastructure repeatedly until the last object is fetched. Moreover, the authors develop a method to link DNS queries to relying parties by redirecting a request from a relying party to a randomly generated subdomain. The relying party would, therefore, contact the nameserver of the experimenter, and a link between the IP address of the resolver and the IP address of the relying party software is established. The measurements were performed in April 2021 and September 2021. Their results show that 63% of ASes that deploy multiple relying party software instances use DNS resolvers from a single AS, and 42.8% only use a single DNS resolver. This might become a problem when adversaries target the DNS resolver to block lookups of publication points, as the relying party would only receive incomplete information. By attacking the DNS, an adversary can impact RPKI ROV within a specific AS.

E. RPKI Infrastructure Outages

Like any other operational system, the RPKI experiences outages that sometimes might lead to the unavailability of specific components within the ecosystem. Depending on

the importance of the component experiencing the outage, RPKI functionality might be more or less severely impacted. Fortunately, RIRs test their RPKI infrastructure and attempt to discover problems that need mitigation before they are triggered by real-world scenarios [205].

RIPE NCC has a very open policy regarding the remediation of incidents, including their very detailed post-mortem analyses, which are posted on SIDROPS and routing-wg mailing lists soon after incidents are mitigated. Hence, many incidents discussed here were publicly released by RIPE NCC, although they can be expected to happen at the same frequency and with the same impact at other RIRs, which do not necessarily share all information publicly. Such an open policy allows others to learn and avoid the same mistakes and is considered the best way to handle incidents.

One of the early incidents was discovered by Iamartino *et al.* [134], who found that LACNIC and APNIC repositories had expired X.509 certificates for almost nine months. These expired certificates led to ROAs that would not cryptographically validate anymore and hence would not be included in the VRP of RP software. More recently, RIPE NCC published a lessons learned article that highlights some major incidents [195]. In February 2020, a disk problem caused a CRL to expire and ROAs not to get published on the publication server [194]. After fixing the disk problem, only a full CA key roll could mitigate the expired CRL issue. Two months later, 2,669 ROAs were deleted during an update of the internal registry software, causing customers to receive alerts about their deleted ROAs and unprotected IP address space [193]. The incident was mitigated 21 hours later, and ROAs were reinstated. In January 2021, during an outgoing IP transfer from RIPE NCC to another RIR, the update process of the parent CA removed the resource from the parent but left it within the child CA [192]. This led to over-claiming of the child certificate of the RIR member, which in turn triggered older RPKI validators to reject the whole child certificate (as a single resource within was inconsistent) and, therefore, all resources contained within. As a result, RIPE NCC implemented checks that force the renewal of child certificates if parent certificates had resources removed and inconsistencies were found. A similar incident with over-claiming certificates was reported by APNIC in February 2020 [206]. In August 2022, ARIN reported an incident that led to an outage of their RRDP services [207]. Also, ROA publication was delayed during the affected period of 1.5 hours. RP software could not query ARIN repositories, but rsync services were not impacted and used as a fallback for RP software. One month later, another RRDP service degradation was reported by ARIN as RRDP certificates had expired, which prevented updates across all RPKI instances in the ARIN region [208]. Very recently, problems with bandwidth were raised on the SIDROPS mailing list that occur when RIRs serve too many clients but do not provision sufficient bandwidth to sustain recurring queries [209]. While it would be simple to allocate more bandwidth, it is also more costly, and ideas were exchanged about improving compression algorithms used in RRDP.

All these incidents led to a learning and, as a result, additional monitoring systems and checks before specific

actions are performed. It is impossible to think of all possibilities that could potentially go wrong beforehand, which is why the learning phase is crucial until we converge to a stable RPKI infrastructure.

In summary, we have discussed problems related to the centrality of the RPKI infrastructure, *e.g.*, whacking of ROAs and cross-country certification. To move from a centralized toward a decentralized model, blockchain approaches as well as threshold-based approaches have been proposed. None of these have been implemented. Further problems are RP inconsistencies, leading to a *mirror world* attack and circular dependencies between BGP and RPKI. Lately, more research has been performed on attacking the RPKI infrastructure itself. We have discussed the outcome of stalling RP software, performing known REST vulnerabilities against the RRDp protocol, and attacking the DNS to take down specific publication points. Lastly, we discussed recent RPKI outages. It is worth noting that every single outage allows to learn—as long as a post-mortem analysis is published.

VI. CONCLUSIONS AND OUTLOOK

In this work, we surveyed research about the Resource Public Key Infrastructure to provide a comprehensive overview of ROA measurements, ROV measurements, and RPKI resiliency. We presented detailed insights from both the research community and the IETF and network operator community along the historical evolution of RPKI deployment.

Since the beginning of the deployment of RPKI, the method and tools to monitor ongoing progress have been crucial to better understand adoption, identify pitfalls, and justify adjustments of RPKI specifications and operational practice. Several tools are available to assess ROA coverage of IP address prefixes. ROV measurement methods to infer private router configurations and attribute RPKI ROV correctly to single ASes remain challenging, though. Our survey also showed many measurement methods have been introduced, but comparing the effectiveness of the methods is hard because vantage points and coverage differ. This makes it difficult to extend those results to the entire Internet infrastructure.

We identified RPKI resilience research as an emerging trend. The main reason is that increasing deployment leads to additional implementations. This motivates researchers to discover vulnerabilities and look for dependencies that could endanger the availability of RPKI components or the integrity of the data contained within.

In contrast to some prior beliefs, RPKI is becoming an integral part of the larger inter-domain routing infrastructure. To extend the level of security from origin validation to path validation without introducing a significant burden to router hardware, researchers and the IETF discuss options that rely on small extensions to the existing RPKI. In the future, we expect more research about path plausibility and path validation triggered by increasing deployment.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their constructive feedback. We also thank Klement Hagenhoff for a thorough

review of the paper. This work was partly supported by the German Federal Ministry of Education and Research (BMBF) within the projects *PRIME*net and *6G-life*.

REFERENCES

- [1] B. Krebs, “DDoS Mitigation Firm Has History of Hijacks,” 2022. [Online]. Available: <https://krebsonsecurity.com/2016/09/ddos-mitigation-firm-has-history-of-hijacks/>
- [2] OECD, “Routing security,” *OECD Publishing*, no. 330, 2022. [Online]. Available: <https://www.oecd-ilibrary.org/content/paper/40be69c8-en>
- [3] A. Ramachandran and N. Feamster, “Understanding the Network-Level Behavior of Spammers,” in *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*. Pisa, Italy: ACM, 2006, pp. 291–302.
- [4] P.-A. Vervier, O. Thonnard, and M. Dacier, “Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks,” in *Network and Distributed System Security Symposium (NDSS)*. San Diego, California, USA: The Internet Society, 2015.
- [5] J. Stewart, “BGP Hijacking for Cryptocurrency Profit,” 2014. [Online]. Available: <https://www.secureworks.com/research/bgp-hijacking-for-cryptocurrency-profit>
- [6] A. Siddiqui, “KlaySwap – Another BGP Hijack Targeting Crypto Wallets,” 2022. [Online]. Available: <https://www.manrs.org/2022/02/klayswap-another-bgp-hijack-targeting-crypto-wallets/>
- [7] N. Kephart, “Route Leak Causes Amazon and AWS Outage,” 2015. [Online]. Available: <https://www.thousandeyes.com/blog/route-leak-causes-amazon-and-aws-outage>
- [8] A. Toonk, “Hijack Event Today by Indosat,” 2014. [Online]. Available: <http://www.bgpmon.net/hijack-event-today-by-indosat/>
- [9] RIPE NCC, “YouTube Hijacking,” 2009. [Online]. Available: <http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study>
- [10] J. Cowie, “The New Threat: Targeted Internet Traffic Misdirection,” 2013. [Online]. Available: <https://web.archive.org/web/20131121025459/http://www.renexus.com/2013/11/mitm-internet-hijacking/>
- [11] A. Toonk, “Turkey Hijacking IP addresses for popular global DNS providers,” 2014. [Online]. Available: <https://www.bgpmon.net/turkey-hijacking-ip-addresses-for-popular-global-dns-providers/>
- [12] K. Zhang, A. Yen, X. Zhao, D. Massey, S. F. Wu, and L. Zhang, “On Detection of Anomalous Routing Dynamics in BGP,” in *International Conference on Research in Networking*. Athens, Greece: Springer, 2004, pp. 259–270.
- [13] H. Ballani, P. Francis, and X. Zhang, “A Study of Prefix Hijacking and Interception in the Internet,” *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4, pp. 265–276, 2007.
- [14] Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao, “Practical Defenses Against BGP Prefix Hijacking,” in *Proceedings of the 2007 ACM CoNEXT conference*. New York, NY, USA: ACM, 2007, pp. 1–12.
- [15] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, “A Survey of BGP Security Issues and Solutions,” *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, 2009.
- [16] G. Huston, M. Rossi, and G. Armitage, “Securing BGP — A Literature Survey,” *IEEE Communications Surveys & Tutorials*, vol. 13, no. 2, pp. 199–222, 2010.
- [17] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, “How Secure are Secure Interdomain Routing Protocols?” *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 4, pp. 87–98, 2010.
- [18] B. Kuerbis and M. Mueller, “Negotiating a New Governance Hierarchy: An Analysis of the Conflicting Incentives to Secure Internet Routing,” *Communications and Strategies*, no. 81, pp. 125–142, 2011.
- [19] O. Maennel, I. Phillips, D. Perouli, R. Bush, R. Austein, and A. Jaboldinov, “Towards a Framework for Evaluating BGP Security,” in *5th Workshop on Cyber Security Experimentation and Test (CSET 12)*. Bellevue, WA, USA: Usenix Association, 2012.
- [20] J. Li, T. Ehrenkrantz, and P. Elliott, “Buddyguard: A Buddy System for Fast and Reliable Detection of IP Prefix Anomalies,” in *2012 20th IEEE International Conference on Network Protocols (ICNP)*. Austin, TX: IEEE, 2012, pp. 1–10.
- [21] M. Mueller, A. Schmidt, and B. Kuerbis, “Internet Security and Networked Governance in International Relations,” *International Studies Review*, vol. 15, no. 1, pp. 86–104, 2013.

- [22] M. S. Siddiqui, D. Montero, R. Serral-Gracia, X. Masip-Bruin, and M. Yannuzzi, "A Survey on the Recent Efforts of the Internet Standardization Body for Securing Inter-Domain Routing," *Computer Networks*, vol. 80, pp. 1–26, 2015.
- [23] C. Testart, "Reviewing a Historical Internet Vulnerability: Why Isn't BGP More Secure and What Can We Do About it?" in *TPRC 46: The 46th Research Conference on Communication, Information and Internet Policy 2018*. Washington, DC.: TPRC, 2018.
- [24] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A Prefix Hijack Alert System," in *15th USENIX Security Symposium*. Vancouver, B.C., Canada: Usenix Association, 2006.
- [25] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu, "Detecting Prefix Hijackings in the Internet with Argus," in *Proceedings of the 2012 Internet Measurement Conference*. Boston, Massachusetts, USA: ACM, 2012, pp. 15–28.
- [26] M. Candela, "BGP Alerter," 2011. [Online]. Available: <https://github.com/nttgin/BGPalerter>
- [27] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti, "ARTEMIS: Neutralizing BGP Hijacking within a Minute," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2471–2486, 2018.
- [28] J. Schlamp, R. Holz, Q. Jacquemart, G. Carle, and E. W. Biersack, "HEAP: Reliable Assessment of BGP Hijacking Attacks," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 6, pp. 1849–1861, 2016.
- [29] A. Milolidakis, T. Bühler, M. Chiesa, L. Vanbever, and S. Vissicchio, "Poster: Smart BGP Hijacks that Evade Public Route Collectors," in *Presented at ACM Internet Measurement Conference (IMC)*. Virtual: ACM, 2019.
- [30] X. Zhang, H.-C. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. G. Andersen, "SCION: Scalability, Control, and Isolation On Next-Generation Networks," in *2011 IEEE Symposium on Security and Privacy*. Oakland, California, USA: IEEE, 2011, pp. 212–227.
- [31] Internet Engineering Taskforce, "SIDR working group," 2006. [Online]. Available: <https://datatracker.ietf.org/wg/sidr/history/>
- [32] 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>
- [33] M. Lepinski, S. Kent, and D. Kong, "A Profile for Route Origin Authorizations (ROAs)," IETF, RFC 6482, February 2012.
- [34] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, "BGP Prefix Origin Validation," IETF, RFC 6811, January 2013.
- [35] RouteViews Project, "University of oregon routeviews project," 2013. [Online]. Available: <http://www.routeviews.org>
- [36] National Institute of Standards and Technology, "NIST RPKI Monitor," 2020. [Online]. Available: <https://rpki-monitor.antd.nist.gov/>
- [37] 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.
- [38] Merit Networks Inc., "RADb - The Internet Routing Registry." [Online]. Available: <https://www.radb.net/query>
- [39] G. D. Battista, T. Refice, and M. Rimondini, "How to Extract BGP Peering Information from the Internet Routing Registry," in *Proceedings of the 2006 SIGCOMM workshop on Mining network data*. Pisa, Italy: ACM, 2006, pp. 317–322.
- [40] F. Wang and L. Gao, "On Inferring and Characterizing Internet Routing Policies," *Journal of Communications and Networks*, vol. 9, no. 4, pp. 350–355, 2007.
- [41] B. Du, G. Akiwate, T. Krenc, C. Testart, A. Marder, B. Huffaker, A. C. Snoeren, and K. Claffy, "IRR Hygiene in the RPKI Era," in *International Conference on Passive and Active Network Measurement*. Virtual: Springer, 2022, pp. 321–337.
- [42] W. Aiello, J. Ioannidis, and P. McDaniel, "Origin Authentication in Interdomain Routing," in *Proceedings of the 10th ACM conference on Computer and Communications Security*. Washington D.C., USA: ACM, 2003, pp. 165–178.
- [43] E. Osterweil, S. Amante, D. Massey, and D. McPherson, "The Great IPv4 Land Grab: Resource Certification for the IPv4 Grey Market," in *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*. Cambridge, MA, USA: ACM, 2011, pp. 1–6.
- [44] Internet Engineering Taskforce, "SIDROPS working group," 2016. [Online]. Available: <https://datatracker.ietf.org/wg/sidrops/about/>
- [45] T. Bruijnzeels, O. Muravskiy, and B. Weber, "RPKI Repository Analysis and Requirements," Working Draft, IETF Secretariat, Internet-Draft draft-bruijnzeels-sidr-repo-analysis-00, February 2013. [Online]. Available: <https://www.ietf.org/archive/id/draft-bruijnzeels-sidr-repo-analysis-00.txt>
- [46] E. Osterweil, T. Manderson, R. White, and D. McPherson, "Sizing Estimates for a Fully Deployed RPKI," *Verisign Labs, TR*, vol. 1120005, 2012.
- [47] Internet Architecture Board, "IAB statement on the RPKI," 2010. [Online]. Available: <https://www.iab.org/documents/correspondence-reports-documents/docs2010/iab-statement-on-the-rpki/>
- [48] Internet Architecture Board, "IAB statement on the RPKI," 2018. [Online]. Available: <https://www.iab.org/documents/correspondence-reports-documents/2018-2/iab-statement-on-the-rpki/>
- [49] A. Dalskov, C. Orlandi, M. Keller, K. Shrishak, and H. Shulman, "Securing DNSSEC Keys via Threshold ECDSA From Generic MPC," in *European Symposium on Research in Computer Security*. Guildford, United Kingdom: Springer, 2020, pp. 654–673.
- [50] G. Huston and G. Michaelson, "Validation of Route Origination using the Resource Certificate PKI and ROAs," Internet Requests for Comments, RFC Editor, RFC 6483, February 2012.
- [51] G. Huston, G. Michaelson, and R. Loomans, "A Profile for X.509 PKIX Resource Certificates, howpublished = Internet Requests for Comments," RFC Editor, RFC 6487, February 2012.
- [52] Y. Gilad, T. Hlavacek, A. Herzberg, M. Schapira, and H. Shulman, "Perfect is the Enemy of Good: Setting Realistic Goals for BGP Security," in *Proceedings of the 17th ACM Workshop on Hot Topics in Networks*. Redmond, WA, USA: ACM, 2018, pp. 57–63.
- [53] T. Hlavacek, I. Cunha, Y. Gilad, A. Herzberg, E. Katz-Bassett, M. Schapira, and H. Shulman, "DISCO: Sidestepping RPKI's Deployment Barriers," in *Network and Distributed System Security Symposium (NDSS)*. San Diego, CA, USA: The Internet Society, 2020.
- [54] G. Michaelson, "APNIC now supports RFC-aligned 'publish in parent' self-hosted RPKI," 2020. [Online]. Available: <https://blog.apnic.net/2020/11/20/apnic-now-supports-rfc-aligned-publish-in-parent-self-hosted-rpki/>
- [55] F. V. Silveira, "Routing WG Mailinglist - Publish in Parent - input requested," 2022. [Online]. Available: <https://www.ripe.net/ripe/mail/archives/routing-wg/2022-September/004613.html>
- [56] NLNetLabs, "Krill 0.11.0," 2022. [Online]. Available: <https://krill.docs.nlnetlabs.nl/en/stable/>
- [57] Dragon Research Labs, "RPKI Toolkit," 2022. [Online]. Available: <https://github.com/dragonresearch/rpki.net>
- [58] N. Trenaman, "Lifecycle of the RIPE NCC RPKI Validator," RIPE NCC, 20. Oktober 2020. [Online]. Available: https://labs.ripe.net/author/nathalie_nathalie/lifecycle-of-the-ripe-ncc-rpki-validator/
- [59] R. Austein, R. Bush *et al.*, "Dragon Research Labs RPKI Toolkit," 2006. [Online]. Available: <https://github.com/dragonresearch/rpki.net>
- [60] A. Reuter, M. Wählisch, and T. C. Schmidt, "RPKI MIRO: Monitoring and Inspection of RPKI Objects," *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4, pp. 107–108, 2015.
- [61] ZDNS, "RPKIVIZ," 2022. [Online]. Available: <http://rpkiviz.zdns.cn/>
- [62] Ben Maddison, "RPKImancer," 2022. [Online]. Available: <https://github.com/benmaddison/rpkimancer>
- [63] P. H. Friedemann, N. Rodday, and G. D. Rodosek, "Assessing the RPKI Validator Ecosystem," in *2022 Thirteenth International Conference on Ubiquitous and Future Networks (ICUFN)*. Barcelona, Spain: IEEE, 2022, pp. 295–300.
- [64] A. Tridgell, P. Mackerras, and W. Davison, "rsync protocol man page." [Online]. Available: <https://linux.die.net/man/1/rsync>
- [65] T. Bruijnzeels, O. Muravskiy, B. Weber, and R. Austein, "The RPKI Repository Delta Protocol (RRDP)," Internet Requests for Comments, RFC Editor, RFC 8182, July 2017.
- [66] T. Bruijnzeels, R. Bush, and G. G. Michaelson, "Resource Public Key Infrastructure (RPKI) Repository Requirements," Internet Engineering Task Force, Internet-Draft draft-ietf-sidrops-prefer-rrdp-02, Dec. 2022, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-sidrops-prefer-rrdp-02/>
- [67] RIPE NCC, "RPKI Validator 1," 2011. [Online]. Available: <https://github.com/RIPE-NCC/rpki-validator>
- [68] Multiple authors, "RPSTIR v1," 2012. [Online]. Available: <https://github.com/bgpsecurity/rpstir>
- [69] L. Poinssignon, M. Chris, and J. Bampton, "OctoRPKI," GitHub, Cloudflare, 2019. [Online]. Available: <https://github.com/cloudflare/cfrpki>
- [70] LACNIC and NIC.MX, "FORT Validator - Github Repository," 2021. [Online]. Available: <https://nicmx.github.io/FORT-validator/>
- [71] K. Dzonsons, C. Jeker, J. Snijders, T. de Raadt, S. Benoit, and T. Buehler, "rpki-client," OpenBSD, 2021. [Online]. Available: <https://www.rpki-client.org/>

- [72] NLnet Labs, "Routinator Manual," 2021. [Online]. Available: <https://routinator.docs.nlnetlabs.nl/en/stable/>
- [73] S. Qing and D. Ma, "RPSTIR2," GitHub, bgpsecurity, 2020. [Online]. Available: <https://github.com/bgpsecurity/rpstir2>
- [74] M. Puzanov, "rpki-prover," GitHub, 2020. [Online]. Available: <https://github.com/lolepezy/rpki-prover>
- [75] R. Bush and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1," Internet Requests for Comments, RFC Editor, RFC 8210, September 2017.
- [76] A. Malhotra and S. Goldberg, "RPKI vs ROVER: Comparing the Risks of BGP Security Solutions," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 4, pp. 113–114, 2014.
- [77] J. Gersch and D. Massey, "ROVER: Route Origin Verification Using DNS," in *2013 22nd International Conference on Computer Communication and Networks (ICCCN)*. Nassau, Bahamas: IEEE, 2013, pp. 1–9.
- [78] J. Gersch, D. Massey, C. Olschanowsky, and L. Zhang, "DNS Resource Records for Authorized Routing Information," Working Draft, IETF Secretariat, Internet-Draft draft-gersch-grow-revdns-bgp-02, February 2013. [Online]. Available: <http://www.ietf.org/internet-drafts/draft-gersch-grow-revdns-bgp-02.txt>
- [79] G. Huston, "A Reappraisal of Validation in the RPKI," 2014. [Online]. Available: <https://www.dotnxdomain.net/ispcol/2014-04/rpkiv.pdf>
- [80] B. Du, C. Testart, R. Fontugne, G. Akiwate, A. C. Snoeren *et al.*, "Mind Your MANRS: Measuring the MANRS Ecosystem," in *Proceedings of the 22nd ACM Internet Measurement Conference (IMC'22)*. Nice, France: ACM, 2022, pp. 716–729.
- [81] Internet Initiative Japan (IJ) Research Labs, "Internet Health Report," 2022. [Online]. Available: <https://ihr.ijlab.net/ihr/en-us/rov>
- [82] M. Wählisch, F. Holler, T. C. Schmidt, and J. H. Schiller, "{RTRlib}: An {Open-Source} Library in C for {RPKI-based} Prefix Origin Validation," in *6th Workshop on Cyber Security Experimentation and Test (CSET 13)*. Washington, D.C., USA: Usenix Association, 2013.
- [83] Cloudflare, "rpki-rtr-client," 2020. [Online]. Available: <https://github.com/cloudflare/rpki-rtr-client>
- [84] Multiple authors, "StayRTR," 2018. [Online]. Available: <https://github.com/bgp/stayrtr>
- [85] M. Hoffmann, X. Eighteen, and A. Band, "rpki-rtr-rust," 2019. [Online]. Available: https://docs.rs/rpki-rtr/latest/rpki_rtr/
- [86] Siyuan Miao, "Yet another BGP hijacking towards AS16509," 2022. [Online]. Available: <https://mailman.nanog.org/pipermail/nanog/2022-August/220320.html>
- [87] Doug Madory, "Tweet on AS16509 hijack," 2022. [Online]. Available: <https://twitter.com/DougMadory/status/1562089866321698819>
- [88] M. Lepinski and K. Sriram, "BGPsec Protocol Specification," RFC 8205, Sep. 2017. [Online]. Available: <https://rfc-editor.org/rfc/rfc8205.txt>
- [89] K. Sriram, "BGPsec Design Choices and Summary of Supporting Discussions," Internet Requests for Comments, RFC Editor, RFC 8374, April 2018.
- [90] R. White, "Architecture and Deployment Considerations for Secure Origin BGP (soBGP)," Working Draft, IETF Secretariat, Internet-Draft draft-white-sobgp-architecture-02, June 2006. [Online]. Available: <https://www.ietf.org/archive/id/draft-white-sobgp-architecture-02.txt>
- [91] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," *IEEE Journal on Selected areas in Communications*, vol. 18, no. 4, pp. 582–592, 2000.
- [92] T. Wan, E. Kranakis, and P. C. van Oorschot, "Pretty Secure BGP, psBGP," in *Network and Distributed System Security Symposium (NDSS)*. San Diego, California, USA: The Internet Society, 2005.
- [93] P. v. Oorschot, T. Wan, and E. Kranakis, "On Inter-domain Routing Security and Pretty Secure BGP (psBGP)," *ACM Transactions on Information and System Security (TISSEC)*, vol. 10, no. 3, pp. 11–es, 2007.
- [94] J. Karlin, S. Forrest, and J. Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," in *Proceedings of the 2006 IEEE International Conference on Network Protocols*. Santa Barbara, CA, USA: IEEE, 2006, pp. 290–299.
- [95] P. Li, W. Zhou, and K. Li, "An Operational Approach to Validate the Path of BGP," in *International Conference on Algorithms and Architectures for Parallel Processing*. Agia Napa, Cyprus: Springer, 2008, pp. 133–143.
- [96] Ignas Bagdonas, "A Look at BGPsec Performance," 2022. [Online]. Available: <https://ripe84.ripe.net/archives/video/819/>
- [97] R. Lychev, S. Goldberg, and M. Schapira, "BGP Security in Partial Deployment: Is the Juice Worth the Squeeze?" in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. Hong Kong, China: ACM, 2013, pp. 171–182.
- [98] K. Sriram, D. Montgomery, D. McPherson, E. Osterweil, and B. Dickson, "Problem Definition and Classification of BGP Route Leaks," Internet Requests for Comments, RFC Editor, RFC 7908, June 2016.
- [99] L. Gao and J. Rexford, "Stable Internet Routing Without Global Coordination," in *Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*. Santa Clara, California, USA: ACM, 2000, pp. 307–317.
- [100] R. Anwar, H. Niaz, D. Choffnes, Í. Cunha, P. Gill, and E. Katz-Bassett, "Investigating Interdomain Routing Policies in the Wild," in *Proceedings of the 2015 Internet Measurement Conference*. Tokyo, Japan: AMC, 2015, pp. 71–77.
- [101] H. V. Madhyastha, E. Katz-Bassett, T. E. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane Nano: Path Prediction for Peer-to-Peer Applications," in *NSDI*, vol. 9. Boston, MA, USA: Usenix Association, 2009, pp. 137–152.
- [102] R. Mazloum, M.-O. Buob, J. Auge, B. Baynat, D. Rossi, and T. Friedman, "Violation of Interdomain Routing Assumptions," in *International Conference on Passive and Active Network Measurement*. Los Angeles, CA: Springer, 2014, pp. 173–182.
- [103] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, "Building an AS-Topology Model that Captures Route Diversity," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4, pp. 195–206, 2006.
- [104] A. Haerberlen, I. C. Avramopoulos, J. Rexford, and P. Druschel, "NetReview: Detecting When Interdomain Routing Goes Wrong," in *6th USENIX Symposium on Networked Systems Design and Implementation*, vol. 2009. Boston, MA, USA: USENIX Association, 2009, pp. 437–452.
- [105] A. Azimov, E. Uskov, R. Bush, J. Snijders, R. Housley, and B. Madisson, "A Profile for Autonomous System Provider Authorization," Internet Engineering Task Force, Internet-Draft draft-ietf-sidrops-asp-profile-16, Jul. 2023, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-sidrops-asp-profile-16/>
- [106] A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, and K. Sriram, "BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects," Internet Engineering Task Force, Internet-Draft draft-ietf-sidrops-asp-verification-16, Aug. 2023, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-sidrops-asp-verification/16/>
- [107] J. Snijders, M. Stucchi, and M. Aelmans, "RPKI Autonomous Systems Cones: A Profile To Define Sets of Autonomous Systems Numbers To Facilitate BGP Filtering," Internet Engineering Task Force, Internet-Draft draft-ietf-grow-rpki-as-cones-02, Apr. 2020, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/html/draft-ietf-grow-rpki-as-cones-02>
- [108] CAIDA, "The CAIDA AS Relationships Dataset, 20221001," 2022. [Online]. Available: <https://www.caida.org/catalog/datasets/as-relationships/>
- [109] J. Snijders and F. Korsbaeck, "A profile for RPKI Signed Groupings of Autonomous System Numbers (ASGroup)," Working Draft, IETF Secretariat, Internet-Draft draft-spaghetti-sidrops-rpki-asgroup-00, November 2022. [Online]. Available: <https://datatracker.ietf.org/api/v1/doc/document/draft-spaghetti-sidrops-rpki-asgroup/>
- [110] A. Azimov, E. Bogomazov, R. Bush, K. Patel, and K. Sriram, "Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages," Internet Requests for Comments, RFC Editor, RFC 9234, May 2022.
- [111] M. Bagnulo, A. García-Martínez, S. Angieri, A. Lutu, and J. Yang, "Practicable route leak detection and protection with ASIRIA," *Computer Networks*, vol. 211, no. 108966, 2022.
- [112] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira, "One Hop for RPKI, One Giant Leap for BGP Security," in *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*. Philadelphia, PA, USA: ACM, 2015, pp. 1–7.
- [113] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira, "Jumpstarting BGP Security with Path-End Validation," in *Proceedings of the 2016 ACM SIGCOMM Conference*. Florianópolis, Brazil: ACM, 2016, pp. 342–355.
- [114] N. Rodday and G. D. Rodosek, "BGPEval: Automating Large-Scale Testbed Creation," in *Proceedings of the 19th International Conference on Network and Service Management*. Niagara Falls, Canada: IEEE, 2023, accepted for publication.
- [115] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman, "Are We There Yet? On RPKI's Deployment and Security," in *Network*

- and Distributed System Security Symposium (NDSS). San Diego, California: Internet Society, 2017.
- [116] M. Wählisch, O. Maennel, and T. C. Schmidt, "Towards Detecting BGP Route Hijacking Using the RPKI," *SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 103–104, Aug. 2012.
- [117] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, "RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem," in *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*. Philadelphia, Pennsylvania, USA: ACM, 2015, pp. 1–7.
- [118] A. Band, "Using the 'Maximum Length' Option in ROAs," RIPE NCC, 2011. [Online]. Available: <https://labs.ripe.net/author/alexband/using-the-maximum-length-option-in-roas/>
- [119] 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 *Proceedings of ACM Internet Measurement Conference (IMC)*. New York, NY, USA: ACM, 2019, pp. 406–419.
- [120] T. Hlavacek, H. Shulman, and M. Waidner, "Not All Conflicts Are Created Equal: Automated Error Resolution in RPKI Deployments," in *IEEE INFOCOM 2021-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. Virtual: IEEE, 2021, pp. 1–2.
- [121] Y. Li, H. Zou, Y. Chen, Y. Xu, Z. Ma, D. Ma, Y. Hu, and G. Xie, "The Hanging ROA: A Secure and Scalable Encoding Scheme for Route Origin Authorization," in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. Virtual: IEEE, 2022, pp. 21–30.
- [122] L. Oliver, G. Akiwate, M. Luckie, B. Du, and k. claffy, "Stop, DROP, and ROA: Effectiveness of Defenses through the lens of DROP," in *Proceedings of the 22nd ACM Internet Measurement Conference*. Nice, France: ACM, 2022.
- [123] W. Li, Z. Lin, M. I. Ashiq, E. Aben, R. Fontugne, A. Phokeer, and T. Chung, "RoVista: Measuring and Analyzing the Route Origin Validation (ROV) in RPKI," in *Proceedings of ACM Internet Measurement Conference (IMC)*. Montreal: ACM, 2023, accepted for publication.
- [124] 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)*. Luxembourg City, Luxembourg: IEEE, 2018, pp. 634–641.
- [125] B. Cartwright-Cox, "Are BGPs security features working yet?" 2019. [Online]. Available: <https://blog.benjojo.co.uk/post/are-bgps-security-features-working-yet-rpki>
- [126] N. Rodday, I. Cunha, R. Bush, E. Katz-Bassett, G. D. Rodosek, T. C. Schmidt, and M. Wählisch, "Revisiting RPKI Route Origin Validation on the Data Plane," in *Proceedings of Network Traffic Measurement and Analysis Conference (TMA), IFIP*. Virtual: ACM, 2021.
- [127] G. Huston and J. Damas, "Measuring Route Origin Validation," <https://www.potaroo.net/ispcol/2020-06/rov.html>, 2020.
- [128] W. Chen, Z. Wang, D. Han, C. Duan, X. Yin, J. Yang, and X. Shi, "ROV-MI: Large-Scale, Accurate and Efficient Measurement of ROV Deployment," in *Network and Distributed Systems Security (NDSS) Symposium 2022*. San Diego, CA, USA: The Internet Society, 2022. [Online]. Available: <https://www.ndss-symposium.org/wp-content/uploads/2022-214-paper.pdf>
- [129] N. Künneke-Trenaman, E. Aben, J. den Hertog, and J. Snijders, "RPKI Test," 2019. [Online]. Available: <https://www.ripe.net/s/rpki-test>
- [130] N. Rodday, R. Bush, I. Cunha, E. Katz-Bassett, G. D. Rodosek, T. C. Schmidt, and M. Wählisch, "Poster: Extending RPKI ROV Measurement Coverage," in *Presented at ACM Internet Measurement Conference (IMC)*. Amsterdam, The Netherlands: ACM, 2019.
- [131] 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," in *International Conference on Passive and Active Network Measurement*. Springer, 2020, pp. 71–87.
- [132] T. Hlavacek, H. Shulman, N. Vogel, and M. Waidner, "Keep your friends close, but your routerservers closer: Insights into rpki validation in the internet," *arXiv preprint arXiv:2303.11772*, 2023.
- [133] Cloudflare Inc., "Is BGP Safe Yet?" 2022. [Online]. Available: <https://isbgpsafeyet.com/>
- [134] D. Iamartino, C. Pelsser, and R. Bush, "Measuring BGP Route Origin Registration and Validation," in *International Conference on Passive and Active Network Measurement*. New York, NY, USA: Springer, 2015, pp. 28–40.
- [135] E. Heilman, D. Cooper, L. Reyzin, and S. Goldberg, "From the Consent of the Routed: Improving the Transparency of the RPKI," in *Proceedings of the 2014 ACM conference on SIGCOMM*. Chicago, Illinois, USA: ACM, 2014, pp. 51–62.
- [136] A. Hari and T. Lakshman, "The Internet Blockchain: A Distributed, Tamper-Resistant Transaction Framework for the Internet," in *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*. Atlanta, GA, USA: ACM, 2016, pp. 204–210.
- [137] D. Cooper, E. Heilman, K. Brogle, L. Reyzin, and S. Goldberg, "On the Risk of Misbehaving RPKI Authorities," in *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*. College Park Maryland: ACM, 2013, pp. 1–7.
- [138] Z. Yan, G. Geng, H. Nakazato, and Y.-J. Park, "Secure and Scalable Deployment of Resource Public Key Infrastructure (RPKI)," *Journal of Internet Services and Information Security (JISIS)*, vol. 8, no. 1, pp. 31–45, 2018.
- [139] K. van Hove, J. van der Ham, and R. van Rijswijk-Deij, "Rpkiller: Threat Analysis from an RPKI Relying Party Perspective," *arXiv preprint arXiv:2203.00993*, 2022.
- [140] 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.
- [141] X. Liu, Z. Yan, G. Geng, X. Lee, S.-S. Tseng, and C.-H. Ku, "RPKI Deployment: Risks and Alternative Solutions," in *Genetic and Evolutionary Computing*. Fuzhou City, Fujian Province, China: Springer, 2016, pp. 299–310.
- [142] K. Shrishak and H. Shulman, "Privacy Preserving and Resilient RPKI," in *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. Virtual: IEEE, 2021, pp. 1–10.
- [143] T. Hlavacek, P. Jeitner, D. Mirdita, H. Shulman, and M. Waidner, "Stalloris: RPKI Downgrade Attack," in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 4455–4471. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity22/presentation/hlavacek>
- [144] R. Fontugne, A. Phokeer, C. Pelsser, K. Vermeulen, and R. Bush, "Rpki time-of-flight: Tracking delays in the management, control, and data planes," in *Passive and Active Measurement: 24th International Conference, PAM 2023, Virtual Event, March 21–23, 2023, Proceedings*. Springer, 2023, pp. 429–457.
- [145] T. Hlavacek, P. Jeitner, D. Mirdita, H. Shulman, and M. Waidner, "Behind the Scenes of RPKI," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, 2022*, pp. 1413–1426.
- [146] T. Hlavacek, P. Jeitner, D. Mirdita, H. Shulman, and M. Waidner, "Beyond Limits: How to Disable Validators in Secure Networks," in *Proceedings of the ACM SIGCOMM 2023 Conference*, ser. ACM SIGCOMM '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 950–966.
- [147] M. Fincham, "RPKI, NZNOG 2014," 2014. [Online]. Available: <https://hotplate.co.nz/archive/nznog/2014/rpki/>
- [148] J. Kloots, "RPKI Routing Policy Decision-Making - a SURFnet Perspective," 2014. [Online]. Available: https://labs.ripe.net/author/jac_kloots/rpki-routing-policy-decision-making-a-surfnet-perspective/
- [149] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, and S. Uhlig, "When BGP Security Meets Content Deployment: Measuring and Analysing RPKI-Protection of Websites," Technical Report, Tech. Rep., 2014.
- [150] The Internet Society, "MANRS ROA Stats Tool," 2022. [Online]. Available: <https://roa-stats.manrs.org/>
- [151] Cloudflare Inc., "RPKI Monitor," 2022. [Online]. Available: <https://rpki.cloudflare.com/>
- [152] 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, "Coming of Age - Website," 2019. [Online]. Available: <https://trpki-study.github.io/>
- [153] R. V. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "In Search of the Elusive Ground Truth: The Internet's AS-level Connectivity Structure," *ACM SIGMETRICS Performance Evaluation Review*, vol. 36, no. 1, pp. 217–228, 2008.
- [154] Y. Gilad, O. Sagga, and S. Goldberg, "MaxLength Considered Harmful to the RPKI," in *Proceedings of the 13th International Conference on emerging Networking Experiments and Technologies*. Incheon, Republic of Korea: ACM, 2017, pp. 101–107.
- [155] T. Manderson, K. Sriram, and R. White, "Use Cases and Interpretations of Resource Public Key Infrastructure (RPKI) Objects for Issuers and Relying Parties," Internet Requests for Comments, RFC Editor, RFC 6907, March 2013.

- [156] Y. Gilad, “compress roas,” 2017. [Online]. Available: https://github.com/yossigi/compress_roas
- [157] Y. Gilad, S. Goldberg, K. Sriram, Y. Snijders, and B. Maddison, “The Use of maxLength in the Resource Public Key Infrastructure (RPKI),” RFC 9319, Oct. 2022. [Online]. Available: <https://www.rfc-editor.org/info/rfc9319>
- [158] T. Hlavacek, H. Shulman, and M. Waidner, “Smart RPKI Validation: Avoiding Errors and Preventing Hijacks,” in *European Symposium on Research in Computer Security*. Copenhagen, Denmark: Springer, 2022, pp. 509–530.
- [159] F. SIT. Smart Validator. [Online]. Available: <https://rose.smart-validator.net>
- [160] V. G. Li, G. Akiwate, K. Levchenko, G. M. Voelker, and S. Savage, “Clairvoyance: Inferring Blocklist Use on the Internet,” in *International Conference on Passive and Active Network Measurement*. Virtual: Springer, 2021, pp. 57–75.
- [161] W. Li and T. Chung, “RoVista,” 2023. [Online]. Available: <https://rovista.netsecurelab.org>
- [162] A. Reuter, R. Bush, I. Cunha, E. Katz-Bassett, T. C. Schmidt, and M. Wählisch, “Measuring RPKI Route Origin Validation Deployment,” 2016. [Online]. Available: <https://rov.rpki.net>
- [163] C. Gray, C. Mosig, R. Bush, C. Pelsler, M. Roughan, T. C. Schmidt, and M. Wählisch, “BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping,” in *Proceedings of ACM Internet Measurement Conference (IMC)*. New York: ACM, 2020, pp. 492–505.
- [164] B. Du, C. Testart, R. Fontugne, A. C. Snoeren, and K. Claffy, “Poster: Taking the Low Road: How RPKI Invalids Propagate,” in *Proceedings of the ACM SIGCOMM 2023 Conference*, ser. ACM SIGCOMM ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 1144–1146.
- [165] N. Rodday, R. van Baaren, L. Hendriks, R. van Rijswijk-Deij, A. Pras, and G. Dreo, “Evaluating RPKI ROV Identification Methodologies in Automatically Generated Mininet Topologies,” in *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies*. Barcelona, Spain: ACM, 2020, pp. 530–531.
- [166] B. Cartwright-Cox, “The year of RPKI on the control plane,” 2020. [Online]. Available: https://ripe80.ripe.net/presentations/36-Ben_Cox.pdf
- [167] H. Shulman, N. Vogel, and M. Waidner, “Poster: Insights into Global Deployment of RPKI Validation,” in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*. Los Angeles, USA: ACM, 2022, pp. 3467–3469.
- [168] K. van Hove, “Where Did My Packet Go? Measuring the Impact of RPKI ROV,” 2022. [Online]. Available: <https://labs.ripe.net/author/koen-van-hove/where-did-my-packet-go-measuring-the-impact-of-rpki-rov/>
- [169] K. van Hove, “Where Did My Packet Go? Measuring the Impact of RPKI ROV,” 2022. [Online]. Available: <https://www.youtube.com/watch?v=Pz6wluIgaXE>
- [170] RIPE NCC, “RIPE Routing Information Service (RIS),” 2020. [Online]. Available: <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>
- [171] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti, “BG-Stream: A Software Framework for Live and Historical BGP Data Analysis,” in *Proceedings of the 2016 Internet Measurement Conference*. Santa Monica, California, USA: ACM, 2016, pp. 429–444.
- [172] RIPE NCC Staff, “RIPE Atlas: A Global Internet Measurement Network,” *Internet Protocol Journal*, vol. 18, no. 3, 2015.
- [173] B. Tierney, J. Metzger, J. Boote, E. Boyd, A. Brown, R. Carlson, M. Zekauskas, J. Zurawski, M. Swany, and M. Grigoriev, “perSONAR: Instantiating a Global Network Measurement Framework,” *SOSP Wksp. Real Overlays and Distrib. Sys.*, 2009.
- [174] Multiple authors, “The perSONAR Project,” 2022. [Online]. Available: <https://www.perfsonar.net/>
- [175] Z. Durumeric, E. Wustrow, and J. A. Halderman, “ZMap: Fast Internet-wide scanning and its security applications,” in *22nd {USENIX} Security Symposium ({USENIX} Security 13)*, 2013, pp. 605–620.
- [176] X. Lee, X. Liu, Z. Yan, G. Geng, and Y. Fu, “RPKI Deployment Considerations: Problem Analysis and Alternative Solutions,” Working Draft, IETF Secretariat, Internet-Draft draft-lee-sidr-rpki-deployment-02, July 2016. [Online]. Available: <https://www.ietf.org/archive/id/draft-lee-sidr-rpki-deployment-02.txt>
- [177] B. Kuerbis, *Regional Address Registries, Governance and Internet Freedom*. Internet Governance Project, 2008. [Online]. Available: <https://www.internetgovernance.org/wp-content/uploads/RIRs-IGP-hyderabad.pdf>
- [178] M. Mueller, M. van Eeten, and B. Kuerbis, “In Important Case, RIPE-NCC seeks legal clarity on how it responds to foreign court orders,” 2011. [Online]. Available: <https://www.internetgovernance.org/2011/11/23/in-important-case-ripe-ncc-seeks-legal-clarity-on-how-it-responds-to-foreign-court-orders/>
- [179] P. Bright, “How the Comodo certificate fraud calls CA trust into question,” 2011. [Online]. Available: <https://arstechnica.com/information-technology/2011/03/how-the-comodo-certificate-fraud-calls-ca-trust-into-question/>
- [180] E. Galperin, S. Schoen, and P. Eckersley, “A Post Mortem on the Iranian DigiNotar Attack,” 2011. [Online]. Available: <https://www.eff.org/deeplinks/2011/09/post-mortem-iranian-diginotar-attack>
- [181] M. Marquis-Boire, “A Brief History of DNS Hijackings,” 2012. [Online]. Available: <http://archive.icann.org/en/meetings/costarica2012/bitcache/A%20brief%20History%20of%20DNS%20Hijacking%20%E2%80%93%20Morgan%20Marquis-Boire,%20Google-vid=33555&disposition=attachment&op=download.pdf>
- [182] C. Wisniewski, “Turkish Certificate Authority screwup leads to attempted Google impersonation,” 2013. [Online]. Available: <https://nakedsecurity.sophos.com/2013/01/04/turkish-certificate-authority-screwup-leads-to-attempted-google-impersonation/>
- [183] D. Piscitello, “Guidance for Preparing Domain Name Orders, Seizures & Takedowns,” 2012. [Online]. Available: <https://www.icann.org/en/system/files/files/guidance-domain-seizures-07mar12-en.pdf>
- [184] C. Soghoian and S. Stamm, “Certified Lies: Detecting and Defeating Government Interception Attacks against SSL,” in *International Conference on Financial Cryptography and Data Security*. Rodney Bay, St. Lucia: Springer, 2011, pp. 250–259.
- [185] The Communications Security, Reliability and Interoperability Council III, “Secure BGP Deployment - Final Report,” 2013. [Online]. Available: https://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG6_Report_March_%202013.pdf
- [186] RIPE NCC, “RIPE NCC Audit Activity.” [Online]. Available: <https://www.ripe.net/publications/docs/ripe-694>
- [187] J. Paillisse, M. Ferriol, E. Garcia, H. Latif, C. Piris, A. Lopez, B. Kuerbis, A. Rodriguez-Natal, V. Ermagan, F. Maino *et al.*, “IPchain: Securing IP Prefix Allocation and Delegation with Blockchain,” in *2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*. Halifax, NS, Canada: IEEE, 2018, pp. 1236–1243.
- [188] L. Mastilak, P. Helebrandt, M. Galinski, and I. Kotuliak, “Secure Inter-Domain Routing Based on Blockchain: A Comprehensive Survey,” *Sensors*, vol. 22, no. 4, 2022. [Online]. Available: <https://www.mdpi.com/1424-8220/22/4/1437>
- [189] K. Shrishak and H. Shulman, “Limiting the Power of RPKI Authorities,” in *Proceedings of the Applied Networking Research Workshop*. Madrid, Spain: ACM, 2020, pp. 12–18.
- [190] E. Heilman, “Mirror Worlds, Eclipse Attacks and the Security of Bitcoin and the RPKI,” 2022. [Online]. Available: <https://open.bu.edu/handle/2144/44796>
- [191] S. Kent and D. Mandelberg, “Suspenders: A Fail-safe Mechanism for the RPKI,” Working Draft, IETF Secretariat, Internet-Draft draft-kent-sidr-suspenders-04, October 2015. [Online]. Available: <https://www.ietf.org/archive/id/draft-kent-sidr-suspenders-04.txt>
- [192] N. Trenaman, “RPKI Outage Post-Mortem - Inconsistent Certificates,” 2021. [Online]. Available: <https://mailarchive.ietf.org/arch/msg/sidrops/mlFkEcI0DCLvOZXLy3uZmM1x2do/>
- [193] N. Trenaman, “RPKI ROA Deletion: Post-mortem,” 2020. [Online]. Available: <https://www.ripe.net/ripe/mail/archives/routing-wg/2020-April/004072.html>
- [194] N. Trenaman, “RPKI Outage Post-Mortem - Disk Quota,” 2020. [Online]. Available: <https://www.ripe.net/ripe/mail/archives/routing-wg/2020-February/004015.html>
- [195] N. Trenaman, “Lessons Learned on Improving RPKI,” 2020. [Online]. Available: https://labs.ripe.net/author/nathalie_nathalie/lessons-learned-on-improving-rpki/
- [196] Multiple authors, “RIPE NCC Routing Working Group Mail Archive,” 2022. [Online]. Available: <https://www.ripe.net/participate/mail/forum/routing-wg/PGE0ZGM1MmNjLTlyNjgtNTVINS0zMzMyLWl0N2VjNWZjYWQ5Y0BpbmRlcmFsbC5jby5pbD4=>
- [197] P. Jeitner, H. Shulman, M. Waidner, D. Mirdita, and T. Hlavacek, “BlackHat USA 2022: Stalloris,” 2022. [Online]. Available: <https://www.blackhat.com/us-22/briefings/schedule/#stalloris-rpki-downgrade-attack-27348>

[198] S. Kent and A. Chi, "Threat Model for BGP Path Security," Internet Requests for Comments, RFC Editor, RFC 7132, February 2014.

[199] E. Cambiasso, G. Papaleo, G. Chiola, and M. Aiello, "Slow dos attacks: definition and categorisation," *International Journal of Trust Management in Computing and Communications*, vol. 1, no. 3-4, pp. 300-319, 2013.

[200] D. Mirdita, H. Shulman, and M. Waidner, "Poster: RPKI Kill Switch," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 3423-3425.

[201] K. van Hove, "Tree Hints for the Resource Public Key Infrastructure (RPKI)," Working Draft, IETF Secretariat, Internet-Draft draft-kwvanhove-sidrops-rpki-tree-hints-01, December 2021. [Online]. Available: <https://www.ietf.org/archive/id/draft-kwvanhove-sidrops-rpki-tree-hints-01.txt>

[202] T. Bruijnzeels, R. Bush, and G. Michaelson, "Resource Public Key Infrastructure (RPKI) Repository Requirements," Working Draft, IETF Secretariat, Internet-Draft draft-ietf-sidrops-prefer-rrdp-01, October 2021. [Online]. Available: <https://www.ietf.org/archive/id/draft-ietf-sidrops-prefer-rrdp-01.txt>

[203] OWASP, "XML Security Cheat Sheet," 2022. [Online]. Available: https://cheatsheetsseries.owasp.org/cheatsheets/XML_Security_Cheat_Sheet.html

[204] Nanog, "Possible rsync validation dos vuln." [Online]. Available: <https://mailman.nanog.org/pipermail/nanog/2021-October/216309.html>

[205] ARIN, "Notice of Upcoming Test of RPKI Infrastructure," 2022. [Online]. Available: <https://www.arin.net/announcements/20220805/>

[206] APNIC, "Service Announcement: RPKI Outage," 2020. [Online]. Available: <https://www.apnic.net/about-apnic/service-updates/service-announcement-rpki-outage/>

[207] ARIN, "RPKI RRDp Degredation," 2022. [Online]. Available: <https://arin.statuspage.io/incidents/3kzmw8x7rnhb>

[208] ARIN, "RPKI Repository Not updating," 2022. [Online]. Available: <https://arin.statuspage.io/incidents/rq7ldmzxs4b9>

[209] J. Sniijders, "Publication Point - RP synchronization in bandwidth constrained environments," 2023. [Online]. Available: https://mailarchive.ietf.org/arch/msg/sidrops/s70Z3EOJX5tRYKbNA6axbK_Tuo/

VII. BIOGRAPHY SECTION



Nils Rodday is a joint PhD student with the Universität der Bundeswehr München and the University of Twente. His work focuses on the security of the inter-domain routing infrastructure. In particular, he has worked on RPKI measurement methodologies, default route identification, and path validation algorithms. Before starting his PhD he worked as a Senior IT Security Consultant for the IBM Deutschland GmbH where he was responsible for advising large corporations on their IT security strategy. Prior to that he obtained masters degrees from the University

of Trento and the University of Twente.




Ítalo Cunha is an Associate Professor at the Computer Science Department at UFMG, Brazil, since 2012. He developed his Ph.D. under the French CIFRE Program for cooperation between industry and academia at Technicolor Research and Innovation. He graduated from UPMC Sorbonne in 2011. His research focuses on improving network performance, reliability and security. His contributions provide better visibility on Internet topology and routing dynamics; help network operators troubleshoot failures and performance problems; and

empower other researchers. Ítalo has served on the technical committee of flagship networking conferences such as USENIX NSDI and ACM SIGCOMM; he serves as a member of the National Brazilian Research and Education Network (RNP) Monitoring Work Group.



Randy Bush is a Research Fellow and network operator at Internet Initiative Japan, Japan's first commercial ISP. He is also a Member of Technical Staff at the routing platform vendor Arrcus, Inc. He specializes in network measurement and automation especially routing, network security, routing protocols, and IPv6 deployment. He was a lead designer of the BGP security effort. Randy has been in computing for over 50 years, and has a few decades of Internet operations experience. He was a founder of Verio, which is now NTT/Verio. He was among

the inaugural inductees into the Internet Society Internet Hall of Fame in 2012. He has served as a member of the IESG and in various other roles within the IETF. He was also a founder of the Network Startup Resource Center (NSRC), <https://www.nsrc.org/>, an NSF-supported pro bono effort to help develop and deploy networking technology in the developing economies. He was the instigator and founder of a number of NICs and NOGs. In amongst these activities he is still an active researcher, and co-author.



Ethan Katz-Bassett is an Associate Professor in the Electrical Engineering Department and an Affiliate in the Computer Science Department at Columbia University, where he is also a Member of the Computer Engineering Program and the Sense, Collect and Move Data Center at the Data Science Institute. He was previously the Andrew and Erna Viterbi Early Career Chair in the Computer Science Department at the University of Southern California. He completed his Ph.D. in the Department of Computer Science at the University of Washington and worked

for half a year at Google, as part of a great team tasked with making the mobile web fast.



Gabi Dreo Rodosek is a Full Professor of Communication Systems and Network Security at the Bundeswehr University Munich and the Founding Director of the Research Institute CODE. She is the Coordinator of the EU H2020 project CONCORDIA and holds several supervisory and advisory mandates in the industry. Besides, she is a member of the Digital Council of the German Ministry of Defence, a member of the World Economic Forum's Global Future Council on Cybersecurity, a member of the Security Innovation Board at the Munich Security

Conference, etc. Professor Dreo studied computer science at the University of Maribor, Slovenia, and got her PhD and habilitation from the Ludwig-Maximilians-University in Munich. Her research interests include AI-based network security, Software-Defined Networks, 5G/6G, and Moving Target Defence.



Thomas C. Schmidt (Member, IEEE) studied mathematics, physics from Freie Universität Berlin and the German literature from the University of Maryland. He received the Ph.D. degree from FU Berlin in 1993. He is a Professor of Computer Networks and Internet Technologies with the Hamburg University of Applied Sciences (HAW), where he heads the Internet Technologies Research Group. Prior to moving to Hamburg, he was a Director of a scientific computer centre in Berlin. Since then, he has continuously conducted numerous national and

international research projects. He was the principal investigator in a number of EU, nationally funded and industrial projects as well as a Visiting Professor with the University of Reading, U.K. His continued interests lie in the development, measurement, and analysis of large-scale distributed systems like the Internet. He serves as co-editor and a technical expert in many occasions and he is actively involved in the work of IETF and IRTF. Together with his group, he pioneered work on an information-centric Industrial IoT and the emerging data-centric Web of Things. He is a co-founder of several large open source projects and a Coordinator of the community developing the RIOT operating system—the friendly OS for the Internet of Things.



Matthias Wählisch (Member, IEEE) received the Ph.D. degree (Highest Hons.) in computer science from Freie Universität Berlin. He is a Full Professor and holds the Chair of Distributed and Networked Systems at the Faculty of Computer Science at TU Dresden. His efforts are driven by improving Internet communication based on sound research. His research and teaching focus on scalable, reliable, and secure Internet communication. This includes the design and evaluation of networking protocols and architectures, as well as Internet measurements

and analysis. He is the PI of several national and international projects, supported by overall 6.8M EUR grant money. Since 2005, he has been active within IETF/IRTF. He published more than 150 peer-reviewed papers. He co-founded some successful open source projects, such as RIOT and RTRlib, where he is still responsible for the strategic development.