RPKI is Coming of Age

A Longitudinal Study of RPKI Deployment and Invalid Route Origins

Taejoong Chung
Rochester Institute of Technology

Emile Aben
RIPE NCC

Tim Bruijnzeels
NLNetLabs

Balakrishnan Chandrasekaran
Max Planck Institute for Informatics

David Choffnes
Northeastern University

Dave Levin
University of Maryland

Bruce M. Maggs
Duke University and
Akamai Technologies

Alan Mislove
Northeastern University

Roland van Rijswijk-Deij
University of Twente and
NLNetLabs

John Rula
Akamai Technologies

Nick Sullivan
Cloudflare

ABSTRACT

Despite its critical role in Internet connectivity, the Border Gateway Protocol (BGP) remains highly vulnerable to attacks such as prefix hijacking, where an Autonomous System (AS) announces routes for IP space it does not control. To address this issue, the Resource Public Key Infrastructure (RPKI) was developed starting in 2008, with deployment beginning in 2011. This paper performs the first comprehensive, longitudinal study of the deployment, coverage, and quality of RPKI.

We use a unique dataset containing *all* RPKI Route Origin Authorizations (ROAs) from the moment RPKI was first deployed, more than 8 years ago. We combine this dataset with BGP announcements from more than 3,300 BGP collectors worldwide. Our analysis shows the after a gradual start, RPKI has seen a rapid increase in adoption over the past two years. We also show that although misconfigurations were rampant when RPKI was first deployed (causing many announcements to appear as invalid) they are quite rare today. We develop a taxonomy of invalid RPKI announcements, then quantify their prevalence. We further identify suspicious announcements indicative of prefix hijacking and present case studies of likely hijacks.

Overall, we conclude that while misconfigurations still do occur, RPKI is “ready for the big screen,” and routing security can be increased by dropping invalid announcements. To foster reproducibility and further studies, we release all RPKI data and the tools we used to analyze it into the public domain.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '19, October 21–23, 2019, Amsterdam, Netherlands

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6948-0/19/10...\$15.00

<https://doi.org/10.1145/3355369.3355596>

1 INTRODUCTION

The Border Gateway Protocol (BGP) is *the* mechanism that allows routers to construct routing tables across the Internet. Unfortunately, the original BGP protocol lacked many security features (e.g., authorization of IP prefix announcements), making BGP vulnerable to attacks such as prefix hijacking [3, 5, 7, 14] and route leaks [5]. To defend against these threats, the Resource Public Key Infrastructure (RPKI) was developed in April 2008 as part of the IETF in the SIDR Working Group [54]. Beta deployments followed in the years after, until all Regional Internet Registries (RIRs) started production deployment of RPKI in January 2011.

At its core, RPKI is a hierarchical Public Key Infrastructure (PKI) that binds Internet Number Resources (INRs) such as Autonomous System Numbers (ASNs) and IP addresses to public keys via certificates. The corresponding private keys can be used by certificate holders to make attestations about these INRs—most importantly, Route Origin Authorization (ROA) objects. ROAs allow a certificate holder to authorize an ASN to announce certain IP prefixes, and are signed using the private key of a certificate covering the IP space.

Each of the five RIRs operate their own RPKI *trust anchor* (equivalent to a root certificate in other PKIs), the private key of which is used to sign such certificates. The RIRs also offer hosted services to their members, enabling them to obtain RPKI certificates and generate ROAs.

RPKI objects including certificates, ROAs, and supporting structures such as manifests and certificate revocation lists (CRLs) are published in so-called RPKI *repositories*. RPKI validation software—called Relying Party (RP) software—retrieves objects from these repositories and performs cryptographic validation of the content, ultimately producing a set of valid ROAs. A *validating router* can then use this set to verify incoming BGP announcements. If the router finds a BGP announcement to be in conflict with the set of valid ROAs, it should reject the announcement as (by definition) the origin AS is not authorized to announce the IP prefix(es).

While RPKI sounds straightforward, in practice it can be complex, creating many opportunities for mistakes. For example, an AS may sub-allocate an IP prefix to a customer AS without updating its ROAs or mistakenly include the wrong range of IP prefixes in a

ROA, thus accidentally making its announcements invalid. If such mistakes are pervasive, a validating router cannot rely on RPKI validation to drop invalid routes, as doing so might have too great of an impact on valid traffic. Such mistakes would weaken the basis of RPKI and may hamper adoption of RPKI.

While past studies [25, 43, 58] have looked at snapshots in time of various aspects of the RPKI ecosystem, little is known about how RPKI has developed since its inception in 2011. This situation makes it hard to draw conclusions about the quality of RPKI data and the viability of actually relying on RPKI to filter invalid announcements. In this study, we aim to change this situation by taking a longer view. Using a unique dataset covering *all* RPKI data published by the RIRs on a daily basis since its early origins in 2011 (even before its full standardization), we study over eight years of RPKI data, combined with publicly available data on BGP announcements covering the same period. We augment this dataset with more detailed BGP data from a large CDN from 2017 until the present day to understand the impact of RPKI validation on an operator. With this dataset, we present a comprehensive study of how the RPKI ecosystem has evolved and what fraction of BGP announcements today are actually verifiable using RPKI. We look at common misconfigurations of RPKI, and how these affect the validity of BGP announcements. We also examine whether RPKI meets its goal of preventing the acceptance of intentionally malicious announcements. Our main findings and contributions are as follows:

- We perform the first, detailed day-to-day longitudinal study of RPKI in the context of real-world routing data;
- We study the pervasiveness of common misconfigurations and how these develop over time;
- We attempt to isolate intentional malicious announcements by filtering out common misconfigurations;
- We show that, today, RPKI is ready for “the big screen” and can safely be used to filter invalid announcements.

To foster reproducibility and further research into the RPKI ecosystem, we publicly release all of our analysis code and data (where possible¹) to the research community at

<https://rpki-study.github.io>

The remainder of this paper is organized as follows. Section 2 provides background on RPKI objects and the Route Origin Validation process as well as the related work. Section 3 describes our dataset for this study, and Section 4 shows how RPKI has been deployed since its launch. Sections 5 and 6 examine why and how some BGP announcements are RPKI invalid. Sections 7 and 8 provide a concluding discussion and future work.

2 BACKGROUND AND RELATED WORK

In this section, we provide background information on RPKI and give an overview of related work.

¹Our Akamai dataset is provided under agreement with Akamai; we are not permitted to release this data. However, we provide links for the other two sources and also provide a link for RPKI objects where other researchers can obtain access themselves.

Figure 1: Organization of the RPKI repositories

2.1 BGP

Interdomain routers use the Border Gateway Protocol (BGP) to establish routing tables. In brief, BGP speakers announce *paths* to destination IP prefixes through a series of ASes. Simplifying somewhat, an example BGP announcement looks like the following:

129.21.0.0/16, AS_PATH: AS3549 AS3356 AS4385

This announcement advertises a path to the IP prefix 129.21.0.0/16 through the three ASes listed. Routers process announcements and withdrawals from peers in order to maintain routing tables, and generally pick the most specific prefix in their table when choosing how to forward a packet.

The original BGP protocol lacked many security features, opening the door to a number of attacks. *First*, a malicious AS could make an announcement for an IP prefix that it did not own, which would cause some of the traffic for that IP prefix to be sent to it; this attack is called *prefix hijacking*. *Second*, a malicious AS could make a *more specific* announcement than the originator’s prefix for an IP prefix it did not own (e.g., announcing 129.21.128.0/17 in the example above). Because routers will pick the most specific prefix when forwarding traffic, *all* traffic for that IP prefix will be forwarded to the malicious AS; this attack is called *sub-prefix hijacking*. These attacks have occurred with frequency in practice, with significant effects for the attacked IP prefix holders [3, 5, 7, 14].

2.2 RPKI Objects

RPKI [33] is a public key infrastructure designed as an out-of-band system to help prevent BGP address prefix (and sub-prefix) hijacking attacks. Briefly, RPKI employs cryptographic signatures to limit the set of entities who can announce IP prefixes. There are multiple types of supporting objects in the core RPKI system; the two we use in this paper are:

- (1) a CA certificate, which binds a set of Internet Number Resources (INRs) such as Autonomous System Numbers (ASNs) and IP prefixes to a public key;
- (2) a Route Origin Authorization (ROA), which authorizes an AS to announce certain IP prefixes and is signed by a CA certificate.

These objects are all published in public RPKI repositories operated by the Regional Internet Registries (RIRs). Figure 1 shows how these repositories are organised. Each RIR has a separate hierarchy

starting at its *trust anchor* and certificate [28, 33]. These trust anchors are each owned by an RIR and are akin to root certificates in other PKIs. The trust anchors are used to sign CA certificates for each RIR's members so that the members can make different kinds of *assertions*. In some regions (currently the LACNIC and APNIC regions) there is sometimes an intermediate level at a National Internet Registry (NIR). Equally, RIR or NIR members can also delegate resources to their customers (as shown in Figure 1).

The most important assertion a resource holder can make is a *Route Origin Authorization* (ROA), which authorizes an AS to advertise IP prefixes through BGP. An ROA is a signed attestation that the holder of a set of IP prefixes has authorized a *single* AS to originate routes for those prefixes.² Thus, it contains a single ASN and (multiple) IP prefixes with their prefix lengths, which the AS is authorized to announce.³

2.3 Max-Length

An AS may wish to de-aggregate an allocated IP prefix into multiple so-called sub-prefixes. For example, AS 4385 may wish to de-aggregate its prefix 129.21.0.0/16 into multiple /20 blocks (e.g., 129.21.0.0/20) for their own purposes (fine-grained traffic control, assignment to customers, etc). In the protocol described thus far, the AS would have to create and sign a ROA containing each of the sub-prefixes:

```
129.21.0.0/20, AS 4385
...
129.21.240.0/20, AS 4385
```

Alternatively, the AS can use the *MaxLength* attribute [34] in the ROA, which specifies the longest prefix length for the authorized IP prefix that the AS may announce. Continuing with the example, the AS could instead sign a single ROA:

```
ROA: 129.21.0.0/16-20, AS 4385
```

that would authorize AS 4385 to announce any of the sub-prefixes of 129.21.0.0/16 in CIDR blocks of length between /16 and /20.

The *MaxLength* attribute is therefore efficient, as it acts as a macro that allows a single ROA to match many sub-prefixes. Those sub-prefixes that are not actually advertised, however, but matched by the ROA can be vulnerable to forged-origin sub-prefix hijacks [24, 26]. Thus, it is often recommended to use *MaxLength* *only* if all sub-prefixes are actually advertised in BGP.

2.4 Route Origin Validation

ASes in the RPKI use so-called Relying Party (RP) software in order to download and validate RPKI objects. From all of the ROAs, RP software constructs a set of tuples (ASN, ROA prefix, prefix length, max length), which are called Validated ROA Payloads (VRPs). The set of VRPs can then be made available to the AS's routers using the RP protocol [6].

²In fact, an ROA can contain any AS number (and not just the AS number of the AS signing the ROA). This enables ASes to outsource the BGP operations to another party or include another AS in a multi-homing relationship.

³If a prefix holder wishes to authorize multiple ASes they can simply create multiple ROA objects.

When such a router receives a BGP announcement, it attempts to validate the announcement using the set of VRPs [37]. In order to do so, it determines first if the announced IP prefix is *covered* by *any* VRP; if so, it then determines whether the announcement *matches* the VRP. In more detail, an IP prefix in a BGP announcement is said to be *covered* by a VRP when the IP prefix address and the VRP IP prefix address are identical for all bits specified by the VRP IP prefix length. A BGP announcement is considered to match a VRP when (1) the VRP IP prefix covers the announcement's IP prefix, (2) the VRP AS matches the announcement's AS, and (3) the length of the announcement's prefix is no greater than the *MaxLength* in the VRP.

Hence, a BGP announcement received by a validating router is in one of the three possible RPKI validity states:

- **Valid:** the BGP announcement is matched by a VRP,
- **Invalid:** the IP prefix in the BGP announcement is covered by a VRP, but no VRP matches the announcement,⁴
- **Unknown:** the IP prefix in the BGP announcement is not covered by any VRP.

For example, consider VRPs published by Rochester Institute of Technology (AS4385). AS 4385 can announce one of its IP prefixes by sending a BGP announcement to its neighbors:

```
129.21.0.0/16, AS_PATH: AS4385
```

The neighbors can verify the origin of the BGP announcement by looking up VRPs, and will find that there is a VRP which matches. Thus, this announcement is considered valid.⁵ However the following BGP announcement is considered invalid as it is covered by at least one VRP (for 129.21.0.0/16), but not matched by any VRP:

```
129.21.240.0/24, AS_PATH: AS4385
```

Finally, the following BGP announcement is considered unknown as there is no VRP that covers the announced prefix:

```
129.22.128.0/17, AS_PATH: AS4385
```

Also, routers do not need to do any cryptographic verification to perform this analysis, as it done purely on the basis of VRPs obtained from RP software. Because of this, BGP origin prefix validation is supported by many routers, and does not incur a significant cost in terms of memory or CPU usage on routers.

2.5 Related Work

In this section, we discuss work related to understanding the RPKI ecosystem and other approaches for securing BGP.

RPKI ecosystem There have been a number of studies [15, 25, 29, 43, 58] that focused on the deployment status of RPKI; Gilad et al. [25] studied RPKI adoption from the perspective of network operators; they tried to understand the challenges and incentives to deploy RPKI by performing a survey among network practitioners. Cohen et al. [11] showed that a partial deployment can also yield

⁴Hence, the origin AS can authorize only IP prefixes that match the VRPs; this is to enforce aggregation and prevent (sub-)prefix hijacking, in which a more specific prefix is announced than specified at the origin.

⁵Note that RPKI does not protect against "AS-in-the-middle attacks" where an attacker prepends its AS to the origin AS on the AS_PATH; the neighbors will only attempt to validate the origin.

Trust Anchor	Measurement Period	VRPs	
		Number	Percent of ASes
APNIC	2011-01-21 – 2019-02-20	14,025	8.14%
LACNIC	2011-01-21 – 2019-02-20	4,510	9.33%
RIPENCC	2011-01-21 – 2019-02-20	40,830	16.04%
ARIN	2012-09-24 – 2019-02-20	4,575	1.47%
AFRINIC	2011-01-21 – 2019-02-20	176	3.30%

Table 1: Overview of the RPKI datasets across five RIRs. The number of VRPs and percentage of ASes that have VRPs published is as of February 20, 2019.

significant security benefits through simulations. Reuter et al. [43] proposed active measurement techniques using BGP announcements under their control to study the uptake of RPKI validation among network operators. Cartwright [15] proposes a data plane approach to achieve the same goal through ICMP messages.

Wählisch et al. [58] focused on the deployment of RPKI in the Web ecosystem by checking announcements for prefixes hosting Alexa 1M websites. They found, surprisingly, that less popular websites are more likely to be secured than prominent sites.

Other studies focused on the security of RPKI; Gilad et al. [26] pointed out that the MaxLength attribute of ROAs could weaken BGP security unless all sub-prefixes matched on an ROA with the MaxLength attribute were actually announced. Cooper et al. [12] showed that sophisticated attacks on RPKI repositories could cause transient failures of RPKI, thus taking some IP prefixes offline.

Researchers have also studied and developed RPKI looking glasses and software to inspect the current state of deployment [44, 48, 49] and help operators verify correct RPKI deployment [18, 51, 52].

Our study extends these prior works in three ways. *First*, we examine **all** ROAs from all RIRs since the beginning of RPKI over 8 years ago. *Second*, we examine the current RPKI deployment status using both RPKI objects as well as actual BGP announcements secured by RPKI. *Third*, we examine more types of misconfigurations and potentially suspicious BGP announcements, which requires longitudinal data.

Deploying BGP security protocols There is a large body of work that studies security issues in BGP [5, 27], investigates common misconfigurations [38], proposes security extensions to BGP such as soBGP [59], S-BGP[32], BGPsec [23], or identifies overall challenges to securing interdomain routing [21]. Due to the massively distributed nature of the network, however, it has been challenging to estimate the deployment status of these security protocols or even compare the pros and cons across different security protocols. Gill et al. pointed out that security concerns alone do not provide sufficient motivation for network operators to deploy new security protocols, and thus proposed a strategy to encourage adoption of BGP security protocols (e.g., BGPsec) by providing appropriate incentives to ISPs [20]. On the other hand, Subramanian et al. sought an alternative, easier to deploy, way to ensure path security using cryptographic functions to check bogus route advertisements in the control plane [53]. Several studies compared the effectiveness of BGP security protocols by quantifying the impact of attacks

Dataset	Measurement Period	VPs	Number of Prefixes	
			Uniq.	Orig.
RIPE-RIS	2011-01-21 – 2018-12-27	24	905K	938K
RouteViews	2011-01-21 – 2018-12-27	23	958K	1.00M
Akamai	2017-01-01 – 2018-12-31	3,300	1.94M	1.98M

Table 2: Overview of BGP announcement datasets: The number of (1) vantage points (VPs, collectors), (2) IP prefixes, and (3) IP prefixes with its origin AS observed in the datasets during December 2018.

(e.g., fraction of ASes for which an attacker could intercept traffic) through simulations, assuming that the security protocols were either fully [22] or partially [13] deployed. Lychev et al. also showed that BGP security protocols that aim for path validation such as BGPsec actually provide only modest benefits over origin authentication [35] protocols such as RPKI.

3 DATASETS AND APPROACH

We start out by briefly discussing the datasets we use and our general approach for analyzing this data.

3.1 RPKI data

Each of the five RIRs maintains an rsync repository with RPKI data that relying parties can query in order to perform RPKI validation. The RIPE NCC has maintained a daily archive of the repositories for *all five* RIRs since the beginning of 2011;⁶ we are grateful to RIPE NCC for making this data available for analysis. Table 1 provides an overview of this dataset for each of the RIRs.

3.2 BGP data

In order to understand how ROAs affect routing table construction, we need BGP announcements as well. Thus, we leverage three datasets with BGP announcement data as shown in Table 2. The first two datasets are publicly available and cover the entirety of the period for which we also have RPKI data. However, these public datasets rely on a relatively limited number of vantage points, which can lead to a biased view of routing [41].

To mitigate this, we have also obtained a much larger dataset from a large CDN that contains BGP announcement data from thousands of vantage points globally. While the large CDN dataset provides us a much greater coverage of BGP announcements, it comes with two caveats. *First*, the dataset was only available beginning in 2017, and thus we only have data for the final two years of the study. *Second*, the dataset comes from direct peering between the CDN and various ASes, and it contains many private BGP announcements (i.e., those announced only to the CDN). Thus, from the large CDN dataset, we remove the private BGP announcements by only keeping announcements for IP prefixes where we observed a corresponding BGP announcement for that prefix on that day in one of the public data sets.

⁶There are a few days over the eight year time period during which data was not recorded correctly: out of 2,952 days during the measurement period, data was unavailable on 45 (1.52%) of them.

Across all of the data sets, we use over 46 billion BGP announcements for analysis in the remainder of the paper. Finally, a note about terminology: whenever we refer to authorizations published in RPKI repositories, we will use the term ROAs. In most cases, however, we are discussing validation, in which case we will use the term VRPs, as that is how such data is typically processed by routers.

3.3 IPv4 vs. IPv6

This paper focuses exclusively on IPv4, and does not analyze or compare our findings with IPv6. We do so for a number of reasons. *First*, in a preliminary analysis we performed of the IPv6 data, we did not observe apparent differences between IPv6 and IPv4 in terms of trends in growth of the number of VRPs [46]⁷. Hence, analyzing IPv6 did not provide much additional information about the development of the RPKI as an ecosystem. *Second*, it is difficult to conduct apples-to-apples comparisons between IPv4 and IPv6 deployments. For example, consider our analysis of the fraction of the address space covered by the RPKI. The IPv4 address space is much more densely allocated and announced relative to IPv6, and thus the fraction of address space covered by the RPKI for IPv4 would be much larger. However, this says nothing of the disparity in terms of how much traffic such prefixes cover, making it difficult to understand the impact of such differences between IPv4 and IPv6. To avoid confusion, we limit ourselves to IPv4 and leave analysis of IPv6 (along with corresponding traffic volumes) to future work. We note that the datasets we released include data for IPv6 and most of the tools we ship to use with the datasets (most notably the validation tool called “Ziggy”) support processing of IPv6 data.

4 RPKI DEPLOYMENT

We begin our analysis of RPKI by focusing on the deployment in terms of the number of ROAs we see in the RPKI repositories, and the fraction of ASes that are using RPKI, and the fraction of all IPv4 space that it covers. To do so, we perform a longitudinal analysis of all RPKI objects along with over 46 billion BGP announcements collected from more than 3,300 different vantage points to answer two questions: 1) how have network operators published ROAs to protect their resources? and 2) how many BGP announcements are actually covered by VRPs?

4.1 Deployment of VRPs

Our goal in this section is to conduct a large-scale, longitudinal, and detailed study of RPKI adoption. To observe how network operators have deployed RPKI since the early days of its launch,⁸ we use the RPKI repository data we received from RIPE. Table 1 shows the number of VRPs, derived from ROAs, in each of the trust anchors (i.e., RIRs) as well as the percentage of ASes that have at least one VRP.⁹ Figure 2 plots the number of VRPs, the percentage of ASes that have at least one VRP, and the IP space covered by VRPs in

⁷As of May 2019.

⁸APNIC, LACNIC, RIPENCC, AFRINIC launched their RPKI service in January 2011 and ARIN did so in September 2012 [39].

⁹To do this, we calculate the number of ASes and the IP space allocated to each of the RIRs by analyzing all NRO statistics (Number Resource Organization) [47].

Figure 2: The growth of RPKI in terms of the # of VRP IP prefixes, the % of ASes where some of their IPv4 addresses are covered by VRPs to all ASes managed by the RIR, the % of IPv4 addresses covered by VRPs to all assigned IPv4 addresses for the RIR.

each of the RIRs. From the table and figures, we make a number of observations.

First, we observe a general increasing trend in all three graphs in Figure 2, indicating a significant and increasing adoption of RPKI both in terms of the number of ASes that have VRPs and the fraction of IP space covered by a VRP. This is encouraging as previous work [25] showed that 84% of network practitioners were not interested in deploying RPKI through a survey in 2016.

Second, we observe that overall RPKI deployment varies significantly between RIRs: between 1.38% (ARIN) and 15.11% (RIPENCC) of ASes are included in one or more VRPs in our latest snapshot, and between 2.7% (AFRINIC) and 30.6% (RIPENCC) of the total IPv4 address space administered by RIRs is covered by VRPs. Interestingly, a few registries have a rapidly growing RPKI coverage. For example, the fraction of the total IP space covered by VRPs rose from 19.2% to 30.6% between January 1, 2017 and February 27, 2019.

Third, we observe a few upward “spikes” in the data set. For example, the sharp spike between June 6th and 19th, 2017 for the number of VRPs in APNIC was primarily due to ROAs for three ASNs.¹⁰ In fact, they used to have ROAs that included the MaxLength attribute to cover a large number of IP prefixes with a single ROA. During the aforementioned period, however, more than 13,000 VRPs were introduced separately by disabling *all of the* MaxLength attributes. This explains why the number of VRPs spiked, but the fraction of

¹⁰AS4775, AS10091, and AS9299

Figure 3: The percentage of BGP announcements covered by VRPs. We observe that the fraction of verifiable announcements is consistently increasing across all datasets.

IP space covered by VRPs remains unchanged. This change was in fact not initiated by the resource holders, but instead was caused by the introduction of a new management system at APNIC that unified management of RPKI and IRR. This system mistakenly started disaggregating ROAs when existing data was imported into the system upon launch. When APNIC noticed the sharp increase in ROAs, they rolled back this change and reintroduced ROAs with MaxLength attributes on June 19th, 2017, which makes the number of VRPs drop back to 7,200 [57]. Similarly, a large jump in VRPs for RIPENCC on July 31st, 2017 was due to AS8551 (Bezeq International, Israel). In this case, the resource holder themselves disaggregated ROAs that used the MaxLength attribute into separate ROAs. This led to an increase by 3,486 VRPs, whereas again, the fraction of IP space covered by VRPs remained unchanged.

4.2 BGP Announcements with RPKI

With the knowledge of the number of VRPs that exist, we now examine how many actual BGP announcements are covered by VRPs over time. Specifically, we focus on the percentage of BGP announcements of which IP prefixes are covered by at least one VRP. Note that having an IP prefix covered by a VRP does not by itself imply that the BGP announcement is valid (to be valid, it must be exactly matched with the range of IP prefixes specified in the VRP); we examine the prevalence of *invalid* announcements in the next section.

First, we observe that the number of unique pairs of IP prefixes with origin ASes is 2.0% ~ 4.4% higher than the number of unique IP prefixes as shown in Table 2. This implies that some IP prefixes are announced by *multiple origins*. This could be due to intended purposes such as multi-homing or unintended purposes such as route leaks or BGP hijacking. We discuss this later in the paper.

Second, Figure 3 plots the fraction of BGP announcements that are covered by VRPs. A key observation is that the number of BGP announcements that are verifiable using RPKI is *consistently* increasing across all datasets: between 9.98% and 11.28% of BGP announcements are covered by VRPs in our latest snapshot.

In summary, we observe a surprising level of deployment for RPKI, both in terms of the number of ASes and the fraction of IP space covered by VRPs. Next, we explore whether the BGP announcements covered by VRPs are actually valid or not.

Figure 4: The percentage of invalid BGP announcements from Akamai, RIPE-RIS, and RouteViews datasets: for the first two years of its deployment, about 20.76% of the RPKI-covered BGP announcements are invalid.

5 ROUTE ORIGIN VALIDATION

We now turn to examining the central question in this paper: *what would happen if all ASes deployed RPKI validating routers?* Thus, we focus only on BGP announcements that are verifiable using RPKI by finding at least one VRP covering the announced IP prefixes; consistent with prior work [8, 43], we refer to these prefixes as *covered prefixes* and these announcements as *covered announcements*. We validate *all* BGP announcements in our datasets by comparing them with covering VRPs. We do so over the entire history of our dataset to not only understand the fraction of valid/invalid announcements today, but to also understand the overall trends.

5.1 Invalid announcements

Recall that a BGP announcement is considered *invalid* when the IP prefix is *covered* by at least one VRP but *no VRP matches it*; an announcement is considered *unknown* when the IP prefix is not covered by any VRPs.

During the entire measurement period, we observed a total of 46 billion BGP announcements. Of these, 43 billion (91.9%) were unknown and 3.8 billion (8.1%) were covered; of the covered announcements, we find that 3.45 billion (90.4%) were valid and that 344 million (9.6%) were invalid.

Figure 4 plots the fraction of all covered BGP announcements that were invalid that we observed during our measurement period. However, according to the recommended best practices for network operators, BGP routes for prefixes more specific than /24 are *not* usually accepted to prevent routing table deaggregation [19, 31]. Thus, to obtain the effective BGP announcements that will end up in BGP tables, we also plot the same graph and filter out the BGP announcements more specific than /24 in the bottom plot. We make a number of observations.

Figure 5: The percentage of invalid BGP announcements begins to drop on September 2018.

First, we observe that a large fraction of covered BGP announcements were actually *invalid* at the very early stages of RPKI deployment across all datasets regardless of /24 filtering; for example, during 2011, 8,005,538 out of 16,363,056 (48.92%) covered announcements were invalid (in the RouteViews dataset). Interestingly, we find that 2,199,715 (27.47%) of invalid announcements were due to announced IP prefixes being covered, but *not matched* with VRPs even though their ASNs matched with the VRP ASNs, which implies potential misconfigurations of VRPs.

Starting from 2012, however, the situation became significantly better: in our final snapshot, only between 2.25% and 5.39% of the covered IP prefixes were invalid (depending on the dataset). We believe the sharp decrease in the fraction of RPKI invalid announcements are due to the RIRs who improved their hosted services and RPKI training from 2012. For example, the RIPE NCC-hosted interface started to show BGP announcement validity to prefix owners and offered the option to operators to receive alerts about invalid announcements. The interfaces of LACNIC and APNIC were similarly modified to show invalid announcements to their users more clearly.

Second, when we focus on the last 12 months of our measurement period, we also notice that the overall percentage of invalid prefixes has further been decreasing since September 2018. A zoomed-in version of the graph is presented in Figure 5. We believe this is due to recent efforts from IXPps who adopted RPKI as a service. Some networks started to drop RPKI invalids either by using this service or by deploying RPKI validation themselves; for example, DE-CIX deployed RPKI and started to *drop* RPKI invalid prefixes in 2019 [50]. Thus, the prefix owners who published invalid RPKI prefixes had two choices to prevent their announcements from being filtered by either (a) removing their invalid ROAs or (b) fixing them to match their announcements. Since we did not observe a drop in RPKI coverage in Figure 2, we believe that the owners preferred to fix their ROAs.

5.2 Why BGP announcements are invalid

We just observed that while the fraction of invalid BGP announcements was originally over half of the covered announcements, it is now much smaller and stable. However, it appears that a non-trivial fraction of the covered BGP announcements—between 2% and 5%—are still invalid. We now examine the underlying “reasons” that make these BGP announcements invalid.

Figure 6: The number of invalid BGP announcements due to *too specific* or *wrong ASN*. Note that the *y* axis extends up to 21,028 (wrong ASN) on the RouteViews and up to 11,193 (Too Specific) on the RIPE-RIS dataset.

As we mentioned in Section 2, BGP announcements can be marked as invalid primarily due to two reasons:

- **Too-specific:** An announcement would be considered invalid if the announcement is otherwise valid but the announced IP prefix is *too specific*. In other words, the IP prefix is covered by at least one VRP, and the origin in the announcement is identical to the VRP ASN, but the announced IP prefix does not exactly match with the VRP IP prefix. In such a case, it is likely due to misconfigurations rather than malicious attempts such as prefix hijacking as the origin AS and VRP ASN are identical.¹¹
- **Wrong ASN:** The announced prefixes are covered by at least one VRP, but none of the VRPs match the ASN in the announcement. These announcements could be malicious as the announcing AS is not supposed to announce such a prefix. However, it may also be a configuration error (e.g., ROAs that were not updated when the IP prefixes were moved to another ASN, or AS multi-homing where the ROAs were created for one ASN only).

Figure 6 shows the distribution of the reasons during our measurement period. We now dig deeper into each of these two reasons.

5.3 Too-specific announcements

During our measurement period, we observe that on average 48.0% ~ 51.5% of the invalid announcements are too-specific. For example,

¹¹In theory, an attacker can announce a more specific prefix and prepend the victim’s AS to the path. However this is unlikely because the announcement would still appear as RPKI invalid.

we observe a spike on January 21, 2012 in the RouteViews dataset; this was due to AS 12322 (Free SAS, France), who announced 7,671 invalid IP prefixes, about 96.0% of a total of 7,988 invalid BGP announcements on that date. When we investigated this, we found that they published 6 ROAs, which contain 8 IP prefixes; however, *none of them* specified the MaxLength attribute to include more specific IP prefixes. They immediately fixed the issue by adding the MaxLength field to include more specific IP prefixes on January 22, 2012. Interestingly, we observe that they introduced 8 more ROAs, six of which with a MaxLength attribute on October 23, 2018. However, they again failed to specify a proper value for MaxLength, which caused nearly 8,800 IP prefixes to go invalid. Similarly, we also observe another spike on April 16, 2018 on both the RouteViews and RIPE-RIS datasets; this was due to AS 5089 (Virgin Media Limited), which announced more than 3,200 IP prefixes without setting the MaxLength in the ROAs. Those invalid BGP announcements would not have occurred if their covering ROAs had been specified correctly, either by setting a more specific prefix length in the MaxLength attribute or publishing separate ROAs that cover the IP prefixes that they announced.

Interestingly, a recent survey study [25] showed that some network administrators were confused regarding the MaxLength attribute; for example, network administrators would incorrectly assume that the prefix length specified in the ROA would validate all IP prefixes that are more specific or would not know how to properly set the MaxLength attribute in ROAs.¹²

5.3.1 Usage of MaxLength: Now we examine how the MaxLength attribute is currently being used in ROAs, and why there have been so many too-specific BGP announcements that end up being marked as invalid.

We first obtain the IP prefixes from VRPs that *do not* contain the MaxLength attribute from all VRPs; Figure 7(a) plots the fraction of IP prefixes in VRPs without the MaxLength attribute. Interestingly, we find that the use of MaxLength has been decreasing and only 11.2% of IP prefixes in VRPs use it in our latest snapshot. This aligns with a previous report [26].

However, when we focus on the actual prefixes that have been announced through BGP, we see a different behavior. 7(b) compares the fraction of valid (top) and too-specific IP prefixes (bottom) where their matched and covered VRPs do not have the MaxLength attribute. We make a number of observations; we first note that the MaxLength attribute is widely used to validate BGP announcements; even though the usage is decreasing, we still see that 52.3% of the valid IP prefixes are validated through VRPs with the MaxLength attribute.

We then focus on too-specific announcements; the bottom graph in 7(b) shows that more than 90% of the too-specific announcements (and more than 92% in our latest snapshot) are due to VRPs that do not have the MaxLength attribute. In other words, the majority of the too-specific announcements could have been validated if their VRPs were to correctly set the MaxLength attribute or create explicit VRPs for all the more specifics. This result empirically shows

¹²For this reason, a recent Internet Draft [24] recommends to avoid using the MaxLength attribute in ROAs, and publish separate ROAs covering IP prefixes that authorized ASes may announce.

(a) % of all IPv4 prefixes in VRPs without MaxLength

(b) % of valid (top) and too-specific (bottom) BGP announcements that do not have the MaxLength attributes.

Figure 7: More than 92% of the too-specific announcements do not set the MaxLength at the time of writing.

that some network operators are likely confused about setting the MaxLength attribute correctly.

As a final example, we look at a steep increase in BGP announcements that are invalid because they are too specific on May 6, 2014; in fact, the number of too-specific announcements without the MaxLength attribute in their covered VRPs did not change much (only less than 40 VRP IP prefixes were added). However, AS 6147 (Telefónica del Perú S.A.A, Peru) who previously announced 609 IP prefixes too-specific (in the RouteViews dataset) due to their narrow MaxLength attributes (/19) in their VRPs updated *all* of their VRPs to cover more specific prefixes by increasing the MaxLength attribute to /24. In the plot, this then leaves mostly prefixes that are too-specific without a MaxLength attribute, as evidenced by the line rising from around 50% to nearly 90%. This also indicates that the MaxLength attribute can effectively fix misconfigurations if used correctly.

As we have seen, it seems to be highly likely that invalid BGP announcements caused from too-specific IP prefixes are due to misconfigurations rather than suspicious attempts such as hijacking. Next, we turn and focus on the BGP announcements originated from the wrong ASN.

5.4 Wrong ASN announcements

We now examine the fraction of BGP announcements that are covered, but that are invalid because the origin AS in the announcement does not match the one in the VRPs. As these announcements are originated from different (unauthorized) ASes, they could be an attempt to hijack the IP (sub-)prefixes. However, it does *not always* mean that all invalid BGP announcements with wrong ASNs are hijacking attempts; there could be a number of causes, including many representing misconfigurations:

- **Two different ASNs managed by the same operator:** An operator that owns and manages multiple ASNs may update the IP prefixes without updating the ROAs, thus making the originating AS in the BGP announcement mismatch with the ASN in the VRP. To identify this case, we use CAIDA’s AS-organizations datasets [16] to map ASNs to ISPs, looking for invalid announcements where the conflicting ASes are owned by the same ISP.
- **Provider–Customer Relationship:** An AS can sub-allocate part of its IP prefixes to its customers. In such a case, if the AS publishes ROAs containing the sub-allocated IP prefixes with its ASN instead of the customer’s ASN, the BGP announcements originated from the customer will be invalid. We use CAIDA’s AS relationship dataset [17] to identify relationships between ASes, looking for invalid announcements where the conflicting ASes are known to have a provider–customer relationship.
- **DDoS Protection:** Origin ASes may outsource “scrubbing” of their traffic by using traffic diversion to a DDoS protection service (DPS) [30]. These services usually announce IP prefixes on behalf of owners, which may cause their BGP announcements to be invalid if the prefix owner has not updated the ROAs. For this analysis, we obtain a list of DDoS protection ASes from a recent report by Forrester [56], and look for invalid announcements where their announcing AS is a known DDoS protection AS.
- **Other:** If none of the prior cases hold, we do not know the exact cause of the invalid announcement. This case would include attempted (sub-)prefix hijacking, as such announcements would not likely fall under the three categories above. We therefore label these as “other”.¹³

Based on this classification, we plot the number of the announcements falling into each of the categories in Figure 8. We immediately observe non-trivial fractions of the same ISP, provider–customer, and other invalid announcements; only rarely do we observe invalid announcements due to DDoS protection services. We make a number of observations below.

First, we can confirm multiple cases where an ISP swaps and announces their IP prefixes between two ASes that it manages; for example, Telmex Colombia S.A that manages two ASes, AS 10620 and AS 14080, announced 1,518 IP prefixes (RouteViews dataset) and 1,118 (RIPE–RIS dataset) from AS 10620 between June 8, 2011 and March 2, 2012; however, these IP prefixes were supposed to be announced from AS 14080 as the ASN of their matched VRPs

¹³We acknowledge that the cases we listed above are not exhaustive; however, even this limited list allows us to understand general behavior of the BGP announcements that do not fall into these categories.

Figure 8: The number of BGP announcements with wrong ASNs in four categories. Note that y axis on the RouteViews dataset extends to 16,498!

was actually AS 14080, thus causing all of them to be considered invalid.

Unfortunately, it took almost 9 months for the problem to be corrected. They first *deleted all ROAs* on March 2, 2012, changing the status of all of the announcements to unknown.¹⁴ After that, they reverted to the same ROAs on March 6, 2012 and stopped announcing the prefixes from the invalid origin.

We also observe a similar misconfiguration in the Akamai dataset where Altice Dominicana—which manages both AS 12066 and AS 28118—announced 545 IP prefixes from AS 28118, which were supposed to be announced from AS 12066. Similar to the misconfiguration of Telmex Colombia S.A., these invalid BGP announcements lasted for more than 17 months from July 26, 2017 until the latest snapshot of the Akamai dataset.

Second, surprisingly, we rarely see announcements from DDoS protection ASes. We found only 15 IP prefixes (in the RouteViews dataset) in our latest snapshots: AS 26415 (Verisign Global) announcing 6 IP prefixes owned by AS 13285 (TalkTalk Communications), AS 19905 (Neustar) announcing an IP prefix owned by AS 21599 (Cable Onda) and 3 ASes from Level3 announcing 8 IP prefixes owned by 3 different ISPs.

Third, we observe that mismatches between ASNs who are in a provider–customer relationship happen frequently. For example, AS 6128 (Cablevision Systems Corp.) who sub-allocated its IP prefixes to 9 different ASes has ROAs that cover all of the sub-prefixes

¹⁴This means validating routers will accept the announcements, but all the security benefits of RPKI have been effectively stripped.

Figure 9: The CDF of the days that an AS has sent unauthorized BGP announcements with wrong ASNs. Because DDoS protection ASes have announced less than 15 prefixes, we have left this category out. Note that the x axis extends to over 2,234 days (C-P).

but set the ASN as their own ASN, thus making the IP prefixes announced from all of their customer ASes invalid from October 27, 2013 to our latest snapshot.

We also observe that the case where a provider AS publishes the ROAs but its customer announces the IP prefixes has happened more frequently than the opposite case; during our measurement period, we find that 87.95% (Akamai), 89.46% (RIPE-RIS), and 84.40% (RouteViews) of those announcements are due to providers that have not updated the ASN in their ROAs to be the customers' full prefix. We believe that the main cause of this would be the cases where a provider announces *covered*-prefixes, but it (a) does not have information on more specific announcements (and how they change) made by their customers, and (b) it simply *cannot* delegate ROA management for this space to their customers when using RIR-hosted services.

Four, we observe a number of invalid announcements in the "other" category with different behaviors; we present a few notable examples where a single AS announced more than 1,000 prefixes that are actually owned by other ASes or more than 1,000 prefixes of an owner AS are announced by unauthorized ASes *at a given date* across our datasets;

- From January 12, 2017 to March 9, 2017 (in the Akamai dataset), AS 395561 (Absolute Connections) announced more than 28,322 IP prefixes owned by 694 other ASes, which suggests that it had attempted to hijack many IP prefixes.

Figure 10: The percentage of unauthorized BGP announcements that are in the "other" category. Note that y axis ends at 30%.

- We also find a hijack attempt concentrated in a certain country: AS 55649 (a private ISP in Hong Kong) announced 1,091 IP prefixes originally owned by 12 ASes on February 28, 2018, 10 of which are located in China.¹⁵
- We also observe cases where one AS issued large numbers of BGP announcements with prefixes owned by multiple ASes; for example, AS 37468 (Angola Cables) announced more than 3,500 IP prefixes originally owned by 82 ASes on May 11, 2018. They did so again on July 19, 2018 by announcing more than 15,000 IP prefixes owned by 1,554 ASes.¹⁶
- Interestingly, we also often observe a case where one AS becomes a target from many ASes; we observe, e.g., 401 IP prefixes owned by AS 27738 (Ecuadortelecom S.A.) being announced by 743 ASes on January 7, 2018, but we could not find corroborating evidence for why this happened.

From these examples, we observe that ASes who *misconfigured* their ROAs (e.g., an ISP swapping the IP prefixes between two ASes it manages, a provider AS not updating the ROAs) have generally announced unauthorized prefixes a bit longer. Based on this observation, we plot Figure 9, which shows the cumulative distribution of the number of dates on which we observed a same pair of origin AS and IP prefix of invalid BGP announcements during our measurement period in each of the categories. We make a number of observations; *First*, we observe that invalid announcements in the "other" category are generally announced shorter than the other categories across the datasets; for example, 34.6% (Akamai), 55% (RIPE-RIS), 86.9% (RouteViews) of "other" announcements are observed only for a *single day*.

Second, we also find invalid BGP announcements from the ASes in a customer-provider relationship last longer; for example, 10.3% (Akamai), 9.8% (RIPE-RIS), 9.6% (RouteViews) of these are observed more than 365 days.

We have shown that many unauthorized BGP announcements are not necessarily suspicious attempts. Rather, they are likely to be due to misconfigurations such as setting ROA IP prefixes

¹⁵AS 4837, AS 17785, AS 17799, AS 17897, AS 4809, AS 23650, AS 132719, AS 17896, AS 4134, AS 17923.

¹⁶This incident was reported at <https://twitter.com/bgpstream/status/1020007234082889728>. However, the report shows only one IP prefix from AS 12343.

Figure 11: CDF of the # of victim ASes that an attacking AS has announced. Note that the x axis extends to 701 ASes (Akamai)

too wide making BGP announcements too-specific or announcing prefixes from a different ASN managed by the same ISP. Our simple classification methods allow us to (loosely) estimate what fraction of invalid BGP announcements is possibly due to misconfigurations. This leaves us with a rest category of “other” invalid announcements. As Figure 10 shows, the number of “other” invalids is low, but not negligible (1.39% for Akamai, 1.07% for RIPE-RIS, and 1.13% for RouteViews at the end of our measurement period). In the next section, we attempt to find explanations for this last category.

6 OTHER INVALID ANNOUNCEMENTS

We now turn our attention to examining the remaining invalid announcements in the “other” category, for which we have not yet found an explanation. As these announcements originate from an origin which has likely nothing to do with the authorized origin, we first compare them with well-known hijacking incidents. For the remainder of this section, we will refer to the AS sending out unauthorized and unknown announcements as the “attacking AS” and the owner as the “victim AS” regardless of the actual intent behind the announcements, which is consistent with prior BGP work (such as [36]).¹⁷

6.1 Case study: BGPStream

BGPStream [9] monitors the real time BGP announcements from multiple datasets such as RouteViews and RIPE-RIS. Among other features, it attempts to detect hijacking attempts [40]. The project also announces suspected hijacking incidents via their Twitter account [10]; for cross-validation purposes, we crawled all reports from this account on suspected BGP hijacking attempts.

Out of 2,361 IPv4 hijack reports collected from the account, 2,082 IP prefixes are unknown to RPKI, but 279 were *covered* by at least one of the VRPs we have. The 279 hijack reports contain (1) the time when they detect the hijack, (2) IP prefix address, (3) prefix length, (4) the authorized AS to announce the IP prefix, and (5) the attacking AS. From this information, we are able to find each of the ROAs and their VRPs in our dataset. We next validate and classify them based on the same classification method we have introduced;

¹⁷We point out, again, that RPKI cannot be used to distinguish those that prepend their ASN to the origin AS or impersonate the origin AS because RPKI only attempts to verify the origin.

Figure 12: CDF of the # of attacking ASes that a victim AS has been targeted by. Note that the x axis extends to 1,947 ASes (RouteViews)

- For 6 (2.15%) hijack reports, we observe that victim AS and attacking AS were actually from the same ISP.
- For 10 (3.58%) hijack reports, we find that they are customer and provider relationships and 9 of them are the case where a provider AS registers ROAs, but the IP prefixes are announced by its customer.
- We are not able to find any DDoS protection ASes that announce prefixes on behalf of the origin.
- For 263 (94.27%) hijack reports (the remainder), we find that they are in our “other” category.

These results show a potential impact of RPKI when it is deployed correctly without misconfigurations; the few misconfigurations can be easily patched by publishing correct ROAs. The remaining hijack reports are at least suspicious and can be effectively filtered by relying on RPKI validation.

6.2 Attacking vs. Victim AS

We now attempt to further understand unexplained invalid announcements by looking who pretended to be whom (attacker vs. victim) in these announcements. First, we examine for how many ASes an attacking AS has tried to steal prefixes. Figure 11 shows the cumulative distribution of the # of victim ASes that an attacking AS has tried to hijack. We observe that the majority of suspected attacking ASes focus on a single victim. 75.4% (Akamai), 78.7% (RIPE-RIS) and 71.8% (RouteViews) respectively are cases where a suspected attacker targets just a single victim AS. However, we also observe a few cases where some ASes attacked many victims; The 99th percentile of the # of victim ASes are over 12 ASes (Akamai), 11 ASes (RIPE-RIS), and 19 ASes (RouteViews). In fact, AS 395561 (Absolute Connection) announced 11,512 prefixes owned by 701 ASes (in the Akamai dataset), AS 200759 (Innofield) announced 362 prefixes owned by 76 ASes (in the RIPE-RIS dataset), and AS 37468 (Angola Cable) announced 15,364 prefixes owned by 841 ASes.

Then, we shift focus to victim ASes to understand what are likely popular targets because we observe invalid announcements from many attacking ASes. Figure 12 shows the cumulative distribution of the number of attacking ASes that have attempted to steal the prefixes of a single AS.

We observe that the average # of attacks that a victim AS receives, overall, is more than the average # of ASes targeted by attacking

ASNs (Figure 11). This implies that there are popular ASes that are targeted by many different ASes *preferentially*. For example, AS 60458 (Xtudio Networks S.L.U) was attacked from 138 ASes (from the Akamai dataset) and AS 8048 (CANTV Servicios Venezuela) was attacked from 173 (RIPE-RIS) and 1,947 ASes (RouteViews dataset) during our measurement period. Considering that the prefixes of the victim ASes are announced from more than 100 ASes that are not in any provider-customer relationship, nor in the same ISP, nor from the DDoS protection AS, we argue that these are highly likely to be targeted attacks.

We find that we lack sufficient evidence to attribute intent to all of the unexplained invalid announcements. Nevertheless, circumstantial evidence suggests that at least part of these invalid announcements are likely hijack attempts. If routers apply RPKI validation, then such suspicious announcements will be filtered out, effectively protecting against hijacks. If, on the other hand, these invalid announcements are due to some other, unknown configuration error, then at least such errors are detected and can be resolved in collaboration with the legitimate prefix holder.

6.3 Traffic from the “other” category

We now turn to examining the amount of traffic from the “other” category, which might be helpful for network operators to estimate the impact of dropping potentially suspicious prefixes using the RPKI. In collaboration with Akamai Technologies, we calculated the portion of all HTTP/HTTPs traffic that came from the IP prefixes in “other” category between December 1st and December 28th, 2018. Figure 13 shows these daily percentages of the HTTP/HTTPs traffic from the “other” category. We find that a very small fraction of traffic (less than 0.3%) was exchanged with the “other” prefixes, indicating that Akamai would have lost at most 0.3% of traffic if they had dropped only the invalid prefixes that are, in all likelihood, not announced the authorized origin. However, we also note that this is a non-negligible amount of traffic and that we cannot prove that the invalid announcements are hijack attempts. As such, there remains a need for reliable techniques for detecting hijacking attempts, so that this information can be used in concert with RPKI validation to safely drop unauthorized prefix announcements and their traffic.

7 DISCUSSION

7.1 MaxLength

It has been argued that the use of MaxLength is harmful [24, 26]. The reasoning is, that by allowing more specific announcements the prefix holder makes itself vulnerable to *malicious* hijacks where the origin ASN is spoofed by prepending it to the BGP path.

However, on the other hand, operators who announce prefixes limit their flexibility by not allowing the more specifics. For example, there may be needs for traffic engineering, or rerouting traffic through DDoS mitigation services that warrant making more specific announcements at unpredictable moments. Authorizing such announcements just in time may not work, as it will take time for new ROAs to be published to the repositories and validated, and routers may not use the VRPs immediately (e.g., some operators build static filter lists every 24 hours).

Figure 13: The percentage of Akamai’s HTTP/S traffic coming from the IP prefixes in the “other” category between December 1st and December 28th, 2018. Note that y axis ends at 0.3%.

Hence, there is a trade-off to consider for these operators. Strictly speaking, their networks will be more secure against malicious hijacks by not allowing more specifics, but at the same time traffic engineering will be indistinguishable from such hijacks. If one does not ever use such traffic engineering, then the decision is simple. However, for networks that do use traffic engineering techniques regularly, there is something to be said for allowing the more specifics. The networks will be more vulnerable to *malicious* attacks, but they will still be protected against accidental hijacks, which is considered to be the majority.

This balance will continue to be a discussion point between operators until practical path verification can be done. Currently operators may have some information about the plausibility and validity of BGP paths based on the information through out-of-band mechanisms such as Peering DB.

7.2 RPKI and Path Validity

The RPKI standards also include a specification for BGPsec, which makes entire BGP paths verifiable. Sadly, however, BGPsec is not getting deployed [20, 35]. There are several reasons for this; for example, BGPsec only works when every ASN on a path participates. Also, it requires that cryptographic signing and validation is done by BGP speakers, which is currently supported by only a handful of BGP software such as Quagga [42] and Bird [55]. As a result the SIDROPS working group [54] in the IETF is now considering a pragmatic approach and they recently introduced two new drafts [1, 2], which seem promising. In a nutshell, the proposal is that any ASN may publish an exclusive list of ASNs that they can be seen to advertise announcements to as an RPKI signed object. These objects are called Autonomous System Provider Authorizations (ASPAs). Just like for ROAs, cryptographic validation is done by Relying Party software, and the routers get a simple list of tuples of verified ASN adjacencies. This allows routers to evaluate adjacencies in a path and reject any path where an adjacency is present where the origin ASN authorized some ASN(s), but not the ASN that is found. Note that this would mitigate the concerns around using the MaxLength attribute in VRPs, as spoofing the origin ASN (by prepending an ASN to the origin on the BGP path) would become much more detectable if ASPA objects are published.

Furthermore, the ASPA approach also allows for an incremental opt-in adoption path: if no ASPA object is published for an ASN then any adjacency is simply *unknown*, rather than invalid. Thus, there is a benefit for individual ASNs to protect themselves against spoofing by publishing objects. Obviously, the more ASes participate the better BGP as a whole is protected, but there is no requirement that *all* ASes have to participate before individual ASes start to benefit.

8 CONCLUSIONS

In this paper we studied the Resource Public Key Infrastructure from its inception over 8 years ago to the present day. Over this period RPKI saw very significant deployment to the extent that globally 12.1% of the IPv4 address space is now already covered by Route Origin Authorizations (ROAs). As our analysis showed, in the early days of RPKI deployment the number of misconfigurations that could lead to BGP announcements being marked as invalid was significant. Data quality, however, improved dramatically over the years to the point where nowadays over 94.3% of announcements (in the RouteViews dataset) covered by ROAs are valid.

This does not mean there are no more misconfigurations. Our analysis identified four common types of misconfigurations and classified how pervasive these misconfigurations are over time. Identifying misconfigurations also allowed us to filter these out to leave what we dubbed “potentially malicious announcements”. At present, this filtering is not accurate enough to be able to use RPKI to actually *detect* malicious announcements.

Yet detection was never the goal of RPKI; the goal was to be able to *filter out* BGP updates with unauthorized announcements. Our analysis shows that RPKI is highly successful at this, especially as data quality improved dramatically over the years. With this in mind, we believe RPKI is “ready for the big screen” and operators can start relying on RPKI to drop invalid announcements. And we are not alone in this; a number of prominent operators have already started dropping invalids (e.g., AT&T [4]). Furthermore, common practices are emerging in the operator community such as guidelines for reaching out to owners of prefixes with broken ROAs and temporarily making exceptions to prevent these prefixes from being marked as invalid [45].

Future Work At present, there is no way to detect if an announcement is actually malicious with a high degree of confidence. Existing systems such as BGPStream [40] rely on heuristics to do this. As RPKI coverage expands and data quality keeps improving, invalid announcements detected by RPKI may become a valuable source of evidence of malicious intent. Observations in this paper may help find a way to do this; as the CDFs in Figure 9 show, for example, there appears to be evidence that malicious announcements have a much shorter lifetime than actual misconfigurations. This could help separate the wheat from the chaff when identifying hijacks.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and our shepherd, Olaf Maenel, for their helpful comments. This research was supported in part by NSF grants CNS-1850465, CNS-1564143, CNS-1901325, CNS-1900879, CNS-1563320, CNS-1901090, CNS-1901047 and EC H2020 Project CONCORDIA GA 830927, and made possible by Akamai Technologies and Cloudflare.

REFERENCES

- [1] A. Azimov, E. Bogomazov, R. Bush, K. Patel, and J. Snijders. Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization. IETF, 2018. <https://tools.ietf.org/html/draft-azimov-sidrops-aspa-verification-01>.
- [2] A. Azimov, E. Uskov, R. Bush, K. Patel, J. Snijders, and R. Housley. A Profile for Autonomous System Provider Authorization. IETF, 2018. <https://tools.ietf.org/html/draft-azimov-sidrops-aspa-profile-00>.
- [3] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the internet. *SIGCOMM*, 2007.
- [4] J. Borkenhagen. AT&T/AS 7018 Now Drops Invalid Prefixes from Peers. <https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html>, 2019.
- [5] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford. A survey of BGP security issues and solutions. *Proceedings of the IEEE*, 98(1), IEEE, 2010.
- [6] R. Bush and R. Austein. <https://tools.ietf.org/html/rfc8210>. RFC 8210, IETF, 2017.
- [7] M. A. Brown. Pakistan hijacks YouTube. 2008. <https://dyn.com/blog/pakistan-hijacks-youtube-1/>.
- [8] R. D. Boer and J. D. Koning. BGP Origin Validation (RPKI). University of Amsterdam, 2013. https://www.os3.nl/_media/2012-2013/courses/rp2/p59-report.pdf.
- [9] BGPStream. <https://bgpstream.com/>.
- [10] BGPStream. <https://twitter.com/bgpstream/>.
- [11] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira. One Hop for RPKI, One Giant Leap for BGP Security. *HotNets*, 2015.
- [12] D. Cooper, E. Heilman, K. Brogle, L. Reyzin, and S. Goldberg. On the risk of misbehaving RPKI authorities. *HotNets*, 2013.
- [13] H. Chan, D. Dash, A. Perrig, and H. Zhang. Modeling adoptability of secure BGP protocol. *SIGCOMM*, 2006.
- [14] J. Cowie. China’s 18-Minute Mystery. 2010. <https://dyn.com/blog/chinas-18-minute-mystery/>.
- [15] B. Cartwright-Cox. Measuring RPKI Adoption via the data-plane. NLNOG Day 2018. https://nlnog.net/static/nlnogday2018/8-Measuring_RPKI_ben_NLNOG_2018.pdf.
- [16] CAIDA ASOrganizations Dataset. <http://www.caida.org/data/as-organizations/>.
- [17] CAIDA ASRelationships Dataset. <http://www.caida.org/data/as-relationships/>.
- [18] Cloudflare RPKI Validator Tools and Libraries. <https://github.com/cloudflare/cfrpki>.
- [19] C. Dietzel, A. Feldmann, and T. King. Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. *PAM*, 2016.
- [20] P. Gill, M. Schapira, and S. Goldberg. Let the market drive deployment: a strategy for transitioning to BGP security. *SIGCOMM*, 2011.
- [21] S. Goldberg. Why is It Taking So Long to Secure Internet Routing? *ACM Queue*, 12(8), 2014.
- [22] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How secure are secure interdomain routing protocols? *SIGCOMM*, 2010.
- [23] W. George and S. Murphy. <https://tools.ietf.org/html/rfc8206>. RFC 8206, IETF, 2017.
- [24] Y. Gilad, S. Goldberg, K. Sriram, J. Snijders, and B. Maddison. The Use of Maxlength in the RPKI draft-ietf-sidrops-rpkimaxlen-02. IETF, 2019.
- [25] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman. Are We There Yet? On RPKI’s Deployment and Security. *NDSS*, 2017.

- [26] Y. Gilad, O. Sagga, and S. Goldberg. MaxLength Considered Harmful to the RPKI. *CoNEXT*, 2017.
- [27] A. Herzberg, M. Hollick, and A. Perrig. Secure Routing for Future Communication Networks (Dagstuhl Seminar 15102). 2015. <http://drops.dagstuhl.de/opus/volltexte/2015/5267/>.
- [28] G. Huston, R. Loomans, and G. Michaelson. A Profile for Resource Certificate Repository Structure. RFC 6481, IETF, 2012.
- [29] D. Iamartino. Study and measurements of the RPKI deployment. Master’s Thesis, Politecnico di Milano, 2015.
- [30] M. Jonker, A. Sperotto, R. van Rijswijk-Deij, R. Sadre, and A. Pras. Measuring the Adoption of DDoS Protection Services. *IMC*, 2016.
- [31] A. Lutu, M. Bagnulo, and O. Maennel. The BGP Visibility Scanner. *INFOCOM*, 2013.
- [32] C. Lynn, J. Mikkelsen, and K. Seo. Secure BGP (S-BGP). IETF, 2003.
- [33] M. Lepinski and S. Kent. An Infrastructure to Support Secure Internet Routing. RFC 6480, IETF, 2012.
- [34] M. Lepinski, S. Kent, and D. Kong. A Profile for Route Origin Authorizations (ROAs). RFC 6482, IETF, 2012.
- [35] R. Lychev, S. Goldberg, and M. Schapira. BGP security in partial deployment. Is the juice worth the squeeze? *SIGCOMM*, 2013.
- [36] C. McArthur and M. S. Guirguis. Stealthy IP Prefix Hijacking: Don’t Bite Off More Than You Can Chew (Poster). *SIGCOMM*, 2008.
- [37] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein. BGP Prefix Origin Validation. RFC 6811, IETF, 2013.
- [38] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP misconfiguration. *SIGCOMM*, 2002.
- [39] Meeting of the NRO Executive Council – 101221. <https://www.nro.net/meeting-of-the-nro-executive-council-101221/>.
- [40] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti. BGPStream: A Software Framework for Live and Historical BGP Data Analysis. *IMC*, 2016.
- [41] 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. *SIGMETRICS*, 2008.
- [42] Quagga Routing Suite. <https://www.quagga.net/>.
- [43] A. Reuter, R. Bush, I. Cunha, E. Katz-Bassett, T. C. Schmidt, and M. Wählisch. Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering. *CCR*, 48(1), 2018.
- [44] A. Reuter, M. Wählisch, and T. C. Schmidt. RPKI MIRO: Monitoring and Inspection of RPKI Objects (Poster Paper). *SIGCOMM*, 2015.
- [45] N. Raijer. RPKI for Managers. *5th NLNOG Day*, 2018.
- [46] RIPE RPKI Certification Stats. <https://certification-stats.ripe.net>.
- [47] RIR Regional Statistics. <https://www.nro.net/about/rirs/statistics/>.
- [48] RPKI - Validator. <https://github.com/RIPE-NCC/rpki-validator>.
- [49] RPKI Deployment Monitor. <https://rpki-monitor.antd.nist.gov>.
- [50] RPKI at the DE-CIX route servers. <https://www.de-cix.net/en/resources/route-server-guides/rpki>.
- [51] RTRlib: The RPKI RTR Client C Library. <https://rpki.realmv6.org/>.
- [52] Routinator. <https://nlnetlabs.nl/projects/rpki/routinator/>.
- [53] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz. Listen and whisper: security mechanisms for BGP. *NSDI*, 2003.
- [54] Secure Inter-Domain Routing (sidr). <https://datatracker.ietf.org/wg/sidr/about/>.
- [55] The BGPsec enabled Bird Routing Daemon. <http://www.securerouting.net/tools/bird/>.
- [56] The Forrester Wave: DDoS Mitigation Solutions, Q4 2017. <https://www.cloudflare.com/media/pdf/forrester-wave-ddos-mitigation-solutions-q4-2017.pdf>.
- [57] Tim Harrington, APNIC. Personal Communication.
- [58] 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. *HotNets*, 2015.
- [59] R. White. Architecture and Deployment Considerations for Secure Origin BGP (soBGP). IETF, 2006.