

From the Consent of the Routed: Improving the Transparency of the RPKI

Ethan Heilman Danny Cooper Leonid Reyzin Sharon Goldberg
Dept. of Computer Science, Boston University, Boston, MA, USA
{heilman,dannyc}@bu.edu, {reyzin, goldbe}@cs.bu.edu

ABSTRACT

The Resource Public Key Infrastructure (RPKI) is a new infrastructure that prevents some of the most devastating attacks on interdomain routing. However, the security benefits provided by the RPKI are accomplished via an architecture that empowers centralized authorities to *unilaterally* revoke any IP prefixes under their control. We propose mechanisms to improve the transparency of the RPKI, in order to mitigate the risk that it will be used for IP address takedowns. First, we present tools that detect and visualize changes to the RPKI that can potentially take down an IP prefix. We use our tools to identify errors and revocations in the production RPKI. Next, we propose modifications to the RPKI’s architecture to (1) require any revocation of IP address space to receive *consent* from all impacted parties, and (2) detect when misbehaving authorities fail to obtain consent. We present a security analysis of our architecture, and estimate its overhead using data-driven analysis.

1. INTRODUCTION

The RPKI [38] is a new infrastructure for securing interdomain routing with BGP. BGP has traditionally operated as a “default-accept” architecture: any autonomous system (AS) can originate a BGP routing announcement (*i.e.*, claim to be the destination for) for any IP prefix, and other ASes will accept the BGP announcement by default. This has made BGP vulnerable to a number of routing attacks, the most common [19, 21, 41, 50, 51, 56] and devastating [10, 14, 26, 31] of which are prefix- and subprefix hijacks. In a prefix hijack, a hijacking AS originates BGP routes for IP prefixes that were not allocated to it, causing the traffic for those prefixes to be intercepted by the hijacker’s AS.

The RPKI prevents these attacks by providing a trusted mapping from allocated IP prefixes to ASes authorized to originate them in BGP. To do this, the RPKI establishes a top-down hierarchy of *authorities*, rooted at the Regional Internet Registries (RIRs), that allocate and suballocate IP address space, as well as authorize its use in BGP by in-

dividual origin ASes; routers use the RPKI to distinguish between hijacked BGP routes and routes originated by a legitimate AS. The RPKI also turns out to be surprisingly effective against attacks it was not designed to prevent [26]; there is evidence [39] that more advanced secure routing solutions [32, 37] provide limited benefits over what is already provided by the RPKI. The RPKI requires neither changes to BGP nor online cryptographic computations during routing. It is currently being rolled out by RIRs and adopted by individual network operators, and authorizes about 20,000 BGP routes as of January 2014 (Section 2.1).

From default-accept to default-deny. However, the security benefits of the RPKI are accompanied by a drastic shift from BGP’s traditional “default accept” policies, to a new “default deny” mode: to prevent (sub)prefix hijacks, routers should only accept routes authorized by the RPKI, and discard all other routes by default (Section 3). Meanwhile, the RPKI’s hierarchical architecture empowers centralized authorities to *unilaterally* revoke authorization (or *take down*) IP prefixes under their control. This shift has led to concerns [7, 17, 43, 44, 61] that the RPKI creates powerful authorities with the technical means for taking down IP prefixes, and could be exploited by abusive authorities or governments to settle disputes or block undesirable content. This is a stark departure from the status quo, where these authorities (RIRs, National/Local Internet Registries, *etc.*) had the power to allocate IP address space, but not to impact routing to space that has been allocated [27, 54].

Transparency. In light of the risk of takedowns, it would be useful to have mechanisms that can detect when RPKI authorities misbehave; this could create social (and possibly legal) pressure to motivate misbehaving authorities to fall in line. However, the architecture of the RPKI also makes it difficult to distinguish between revocations due to disputes or censorship, and those due to business arrangements that were agreed upon by all impacted parties (Section 3).

Our contributions. As RPKI deployment continues to gain traction, we present an investigation of the risk of RPKI takedowns, and propose technical solutions that mitigate this risk by improving the transparency of the RPKI. Our contributions are:

1. Security audit (Section 3). We untangle the often-unintuitive interactions between the RPKI and BGP, and explain why it is *not* always the case that a revocation in the RPKI can take down an IP prefix in BGP.

2. Tools, measurement & modeling (Section 4). We start by working within the current RPKI specifications, and

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 the author(s) 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.

SIGCOMM’14, August 17–22, 2014, Chicago, IL, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2836-4/14/08 ...\$15.00. <http://dx.doi.org/10.1145/2619239.2626293>.

build two tools that can increase its transparency by detecting and reacting to RPKI problems. Our *detector*, that identifies changes to the RPKI that can takedown IP prefixes, and a *visualizer*, that visualizes the results. We test our tools on the production RPKI and identify real-life errors and revocations (Sections 3,4.1). Since RPKI deployment is still in its infancy, we also develop models of a future full-deployment of the RPKI, based on routing data and RIR information (Section 5.7).

3. Changes to the specifications (Section 5). The current RPKI specifications still place power squarely in the hands of authorities. To remedy this, we propose modest modifications to the architecture of the RPKI. Our design ensures that once a route is authorized by the RPKI, it switches from “default-deny” to “default-accept”; that is, authorities that revoke IP address space must first obtain *consent* from all entities holding allocations (and suballocations) of that space. Our design also allow parties viewing RPKI information from different vantage points to detect when (1) their views are not *consistent*, and (2) authorities fail to properly obtain consent. We prove the security properties of our design. Because our design sometimes requires many entities to provide consent, we also estimate its extra overhead using data-driven analysis (Section 5.7).

2. RPKI PRIMER

We overview the RPKI’s certificate hierarchy, explain how RPKI information determines route validity, and explain how the RPKI can limit threats to BGP.

2.1 The hierarchical structure of the RPKI.

The RPKI arranges authorities in a strict hierarchy that mirrors the IP address allocation hierarchy.¹ An authority may issue cryptographic objects for IP addresses that are covered by its own IP addresses. (An IP prefix P covers prefix π if π is a proper subset of the address space in P (e.g., 63.160.0.0/12 covers 63.160.1.0/24) or if $P = \pi$. Also, prefix 63.160.0.0/12 has *length* 12.) Each authority has a *resource certificate (RC)*, a certificate that contains its cryptographic public key and its set of allocated IP addresses [40]. An authority may issue signed objects for IP addresses covered by its allocation, specifically: (1) an RC that suballocates a subset of its addresses to another authority, or (2) a *route origin authorization (ROA)*, that authorizes a specified AS to originate a set of prefixes, and its subprefixes up to a specified length called *maxLength*, in BGP [38]. ROAs protect BGP from routing attacks (Section 2.2).

Model (Figure 1). We show how an RIR (ARIN) uses its RC to suballocate a prefix to another authority (Sprint), which then issues RCs suballocating this prefix to other authorities (ETB S.A. ESP., Continental Broadband). (This is an excerpt of one of our models of the fully-deployed RPKI; see Section 5.7.) We say Sprint is the *parent* of Continental Broadband, and extend this to child, ancestor, *etc.*, in the obvious way. Sprint issues two ROAs that authorize a specified prefix and its subprefixes up to *maxLength* 24; the remaining ROAs shown authorize only a single prefix.

Status of the RPKI (Table 2). As of January 13, 2014, the production RPKI contains ROAs for about 20K

¹The roots of the RPKI are the five RIRs (Table 2); in the future, IANA could be a single root [38, Section 2.4].

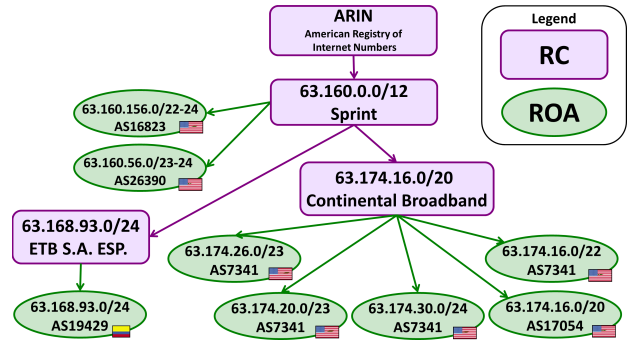


Figure 1: Excerpt of a model RPKI

prefix-to-origin-AS pairs. (About 488K prefixes were announced in BGP that day.) The RPKI’s structure (Table 2) is slightly different than our models. At depth 0, there are trust anchors for each RIR (e.g., ARIN). The trust anchors are long-lived certificates, who issue a handful of shorter-lived intermediate RCs to the RIRs (depth 1); our models and Figure 1 omit intermediate RCs. Intermediate RCs issue leaf RCs to organizations (e.g., Sprint) who then issue ROAs. ARIN has an extra layer of intermediate RCs. RCs for suballocations like Continental Broadband and ETB in Figure 1 are absent; their ROAs could just be issued by Sprint. ROAs usually contain one AS and many prefixes (e.g., all prefixes for AS 7341 could be issued in one ROA).

Publication points. The RPKI was designed to require minimal changes to BGP, and therefore operates entirely out-of-band. RPKI objects are stored in public repositories. Each RC has its own *publication point* (i.e., directory in a file system) where it publishes every object it issued. Each RC also signs a *manifest* that logs the hash of every object present in its publication point. *Relying parties* download RPKI objects from publication points to their *local caches*, validate the objects, push information to their routers, and use it to inform routing decisions in BGP.

2.2 How the RPKI prevents threats to BGP.

BGP is especially vulnerable to subprefix hijacks because of longest-prefix-match routing: when a router learns BGP routes for a prefix and its subprefix, it always prefers the subprefix route. Subprefix hijackers exploit this by originating routes for subprefixes of a victim’s prefix. This leads to a natural desideratum for the RPKI: a subprefix hijacker’s route should be always be invalid when the legitimate route has a matching valid ROA.

Origin authentication. To achieve this desideratum, a relying party uses the RPKI for *origin authentication* as follows. For our purposes, a BGP route is an IP prefix π and an origin AS a . Once a relying party has “access to a local cache of the complete set of valid ROAs” [29, Sec. 2], these valid ROAs are used to classify each route (π, a) learned in BGP into one of three *route validation states* [29, 42]:

- *Valid*: There is a valid *matching ROA*. A matching ROA has (1) a matching origin AS a , and (2) a prefix P that covers prefix π , and (3) the specified *maxLength* no shorter than the length of π .
- *Unknown*: There is no valid *covering ROA*. A covering ROA is any ROA for a prefix that covers π .
- *Invalid*: The route is neither unknown or valid.

The rules above elegantly achieve the desideratum: if a le-

Depth	1	2	3	4
RIPE	4 RC	1909 RC	1512 ROA	
LACNIC	4 RC	282 RC	282 ROA	
ARIN	1 RC	1 RC	99 RC	151 ROA
APNIC from IANA	1 RC	450 RC	58 ROA	
AfriNIC	1 RC	27 RC	48 ROA	

Table 2: Valid ROAs and RCs at each depth of the production RPKI on January 13, 2014.

gitimate route in BGP has a matching valid ROA in the RPKI, a subprefix hijacker’s route will always be “invalid” (because it is “covered” by the valid ROA). To stop subprefix hijacks, relying parties should *drop invalid routes*, *i.e.*, not select routes in BGP that are “invalid” per the RPKI.

3. AUDIT: THE RISK OF RPKI TAKEDOWNS

We explore RPKI threats that take down IP prefixes by causing legitimate BGP routes to be misclassified as “invalid”.

A trace of the production RPKI. To aid our exposition, this section discusses real-life events that our tools detected in the production RPKI (described in Section 4). Events were observed in a trace we collected from 2013/10/23–2014/01/13. Each day, we used *reynic* [2] to pull the state of the production RPKI to an empty local folder, and to cryptographically validate the result. The trace excludes a few days where our collector went down, and a few days where *reynic* failed to properly sync to the repository (reporting `<rsync_transfer_failed kind="bad">`).

3.1 Tradeoff: BGP threats vs RPKI threats.

Now that we have understood how threats to BGP influenced the design of the RPKI, we look at how threats to the RPKI can impact routing with BGP. Specifically, to prevent subprefix hijacks on BGP, every route covered by ROA in the RPKI is classified as “invalid” *by default*, unless it has a matching valid ROA of its own. The RPKI is therefore a “default-deny” architecture,² and this comes with some risks; if a legitimate BGP route is wrongly classified as “invalid” by the RPKI, a relying party that *drops invalid routes* will lose connectivity to that route in BGP. As such, there have been concerns [7, 17, 43, 44, 61] that the RPKI creates a new technical means for IP-prefix takedowns.

Because moving to a default-deny regime is a drastic change for the routing infrastructure, relying parties are explicitly entitled to use their own “local policies” to decide what to do with “invalid” routes [29]. *Drop invalid routes* is the strictest local policy, and comes with the risk that threats to the RPKI can take down legitimate BGP routes. An more lenient policy suggested in RFC6483 [29] is to *depref “invalid” routes*: for a given prefix, a router should prefer “valid” routes over “invalid” routes. This policy implies that a router still selects an “invalid” route when there is no “valid” route for the *exact same* IP prefix. Thus, the router may still be able to reach routes that are wrongly classified as “invalid” as a result of problems with the RPKI. (However, availability of a route at one router can depend strongly on local policy used at other routers. For example, a router that uses the lenient *depref invalid* policy can lose connectivity to an “invalid” route if all its neighboring routers use the

²The idea that the RPKI is a default-deny architecture is un-intuitive, given the “unknown” route validity state. However, a BGP route without a matching valid ROA is “unknown” only when there is *no covering ROA*; see Section 2.2.

policy	routing attack	RPKI manipulation
drop invalid	stops (sub)prefix hijacks	prefix goes offline
depref invalid	subprefix hijacks possible	prefix may stay online

Table 3: Impact of different local policies.

strict *drop invalid* policy.) On the other hand, this policy does *not* prevent subprefix hijacks [12, Section 5].

Table 3 summarizes this difficult tradeoff; the local policy that is best at protecting against attacks on BGP is worst at protecting against problems with RPKI. While it is too early to tell what local policies will be adopted in the long run, better assurances about the trustworthiness of RPKI information are needed before relying parties can start enjoying the RPKI’s full potential for protecting BGP from attack. Sections 4-5 focus on improving these assurances.

A note on granularity. While ROAs can have arbitrary prefix lengths, the longest IPv4 prefix length that is universally accepted by BGP routers is a /24. (Routers usually ignore longer prefixes to avoid bloating their routing tables.) Thus, our discussion on the RPKI’s impact on IP prefix reachability should be thought of as having the granularity of a /24 (or shorter) IPv4 prefix, *i.e.*, no fewer than 256 IPv4 addresses; the RPKI, therefore, can be used for significantly coarser level of blocking than *e.g.*, the DNS [52, 53].

3.2 New threats created by the RPKI.

What threats to the RPKI can cause a legitimate BGP route to be classified as “invalid”? We consider two threat models: threats created by RPKI authorities themselves, and threats to the communication path between relying parties and the public RPKI repositories.

3.2.1 Threats created by RPKI authorities.

Because relying parties download RPKI objects from publication points that are *controlled by their issuer* [38, Section 8], the issuer can manipulate the contents of its publication points in any way it likes. There are two ways an issuer can cause a legitimate BGP route to be classified as “invalid”:

1. Adding a ROA. If a new ROA added to the RPKI covers a legitimate BGP route without its own matching ROA, the legitimate BGP route will be wrongly classified as “invalid”. Consider, for example, this incident:

■ **Case Study 1: ROA misconfiguration.** On December 13, a new ROA was added to the production RPKI rooted at ARIN, authorizing prefix 173.251.0.0/17 with maxlength 24 to AS 6128. This caused a large portion of the address space to downgrade from “unknown” to “invalid”, including several legitimate /24 routes announced in BGP that did not have matching ROAs (AS 53725 originating 173.251.91.0/24, AS 13599 originating 173.251.54.0/24).

2. Whacking a ROA. Alternatively, missing information can cause the RPKI to misclassify a legitimate BGP route as “invalid”. To see how, we first distinguish between route validity and RPKI object validity. An RPKI object (*i.e.*, an RC, ROA, *etc.*) is valid if it has a valid signature, is not malformed or corrupted, *etc.* Meanwhile, the validity of a route depends *exclusively* on the set of valid ROAs in a relying party’s local cache. Thus, any action that prevents the delivery of a ROA object to a relying party has exactly the same effect — the corresponding BGP route becomes (a) “invalid” if the local cache contains some other valid covering ROA, or (b) “unknown” if there is no valid covering ROA. We therefore just say that a ROA is *whacked* whenever a

relying party fails to receive the valid ROA object. There are many ways an authority can whack a descendant ROA:³

a. Revoking. An issuer can always revoke any object it issues using its certificate revocation list (CRL).

b. Deleting or corrupting. An authority can also delete or corrupt any object in its publication point. If the authority logs appropriate change in its manifest, relying parties will accept the change without complaint. For example:

■ **Case Study 2: Deleted ROA.** On December 19, 2013, a ROA for (79.139.96.0/24, AS 51813), for a network in Russia, was deleted from the production RPKI. Meanwhile, since at least November 21, the RPKI also had a covering ROA mapping 79.139.96.0/19-20 to another Russian ISP, AS 43782. (This is a covering ROA because the /19 prefix covers the /24 prefix.) The covering ROA caused the route corresponding to the whacked ROA to downgrade from valid to invalid (per Section 2.2). Both the whacked ROA and covering ROA were issued by the same RC.

c. Overwriting. Each RPKI object is identified by a uniform resource identifier (URI), and an authority may overwrite any RC it issued, so that modified objects can have persistent URIs (thus simplifying operations like certificate renewal and key rollover [30]). An authority can thus overwrite an RPKI object with one for a different set of IP addresses (or a different key, *etc.*); children of the overwritten RC can be whacked as a result, because they are no longer covered by the RC (or signed by the wrong key, *etc.*):

■ **Case Study 3: Overwritten parent RC.** On January 5, 2014, a ROA for (196.6.174.0/23, AS 37688) for a backbone connectivity network in Nigeria was whacked because its parent RC was overwritten. The incident occurred as follows: On January 4, the ROA's parent RC was allocated prefix 196.6.174.0/23. On January 5, the RC was overwritten with an RC with the same key but for an IPv6 prefix 2c0f:f668::/32. The ROA in question (which remained in the publication point) became invalid, because it was no longer covered by its parent. Interestingly, the RC had no valid descendants until January 6, when it issued ROAs covered by the IPv6 prefix to a different AS (AS 37600, in Mauritius). Moreover, if an RC gets whacked, all its descendent objects also get whacked. Our earlier work [18] showed how an authority can whack distant descendant ROAs in a targeted manner that causes no collateral damage to other ROAs.

3.2.2 Threats that disrupt RPKI object delivery.

A third party, who disrupts the communication path from an RPKI repository to a relying party, can whack a ROA just by corrupting a single bit in the ROA. (This is because the ROA would fail cryptographic validation.) Fortunately, however, the RPKI has a mechanism for detecting lost/corrupted information:

Manifests. To provide assurance that no objects have been deleted from a publication point, a collision-resistant hash of the contents of every file issued by an RC is listed in a single file called a *manifest* [28, Section 2.1], which is digitally signed by the RC and stored at its publication point. The manifest must be updated whenever an RC issues, modifies, or revokes an object it issued. To prevent replay at-

³To be sure that the BGP route corresponding to the whacked ROA becomes “invalid”, an authority can just issue a covering ROA, as in Case Study 2.

tacks and the propagation of stale information, manifests are short-lived; they usually expire and are renewed daily.

Of course, a third party can always corrupt or disrupt the delivery of the manifest itself. Consider the following:

■ **Case Study 4: Stale LACNIC manifests.** On December 20, 2013, an error at LACNIC whacked 4217 prefix-to-origin-AS pairs. The day before the incident, the RPKI rooted at LACNIC looked very similar to what is shown in Table 2. On December 20, the manifests and CRLs issued by all four of LACNIC's intermediate RCs (at depth 2) in our local cache all expired. The relying party software [2] raised an alarm and rejected all four of the intermediate RCs as invalid. All objects in subtree rooted at LACNIC became invalid; thus, there were no valid ROAs for any address space allocated by LACNIC, and BGP routes corresponding to LACNIC ROAs downgraded from “valid” to “unknown”.

3.3 Transparency is (sometimes) difficult.

To mitigate risks created by the RPKI, a relying party should be able to detect when a problem with the RPKI has caused BGP route to be wrongly classified as “invalid”. Fortunately, third-party attacks that disrupt object delivery are completely transparent – a relying party can always check that it received all the objects in a manifest, and that the manifest was valid and current. If not, relying parties should raise a *missing-information alarm*; RFC6486 [9, Sect 6.5] states that relying parties should react to such alarms at their own discretion.⁴ Meanwhile, incidents where misconfigured or compromised RPKI authorities wrongly add or whack ROA are much less transparent; if the authority ensures that its manifest is consistent with the objects in its publication point, relying parties will not raise alarms.

Alarms. While the RPKI specifications do not focus on alarms, we think that they are crucial. While alarms alone cannot resolve problems, they do indicate that problems are present, and can trigger mitigation mechanisms. During a missing-information alarm, a relying party might try to reconnect the repository, search for missing objects in its local cache or at other relying parties, or search for unusual routing activity in the IP address space covered in the RC whose manifest triggered the alarm. Today, BGP misconfigurations are resolved by humans picking up the phone; RPKI alarms similarly provide guidance on when picking up the phone may be necessary. Indeed, in Section 5, alarms will be our primary tool for increasing transparency.

There is another important threat to transparency:

Mirror world attacks. In a mirror world attack, an (adversarial) RPKI authority presents one view of the RPKI to some relying parties, and a different view to others. (For example, the subject of an RC is shown a view containing the RC, but other relying parties are shown a view that omits the RC.) The current RPKI standard contains no mechanisms to prevent this attack. Our proposal in Section 5 builds in protections against mirror-world attacks.

4. TOOLS TO IMPROVE TRANSPARENCY

We now develop tools that can improve the trustworthiness of RPKI data, without requiring any modifications to the RPKI itself. Our *detector* tool detects when a change

⁴The relying party software we used [2] in Case Study 4 rejected the stale manifest. Other policies are possible.

to the RPKI causes a route to *downgrade* from “valid” to “invalid/unknown”, (or from “unknown” to “invalid”), and our *visualizer* visualizes these downgrades. We released the source for all our tools,⁵ so they can be incorporated into professionally-maintained systems like [35, 46, 57].

4.1 A tool for detecting downgrades.

One way to detect inconsistencies between BGP and the RPKI is to take a BGP feed [3, 5], run the RPKI validation algorithm in Section 2.2 on each BGP route, and identify BGP routes that the RPKI classifies as “invalid”; indeed, this is done in [35, 46, 57]. Our goal, however, is to detect when a *change* to the RPKI (*i.e.*, an added or whacked ROA) causes a change in the validity status of a route; it can therefore act as alert system for potentially-harmful changes to the RPKI. Moreover, we do not want our tool to be limited to a view of BGP provided by a specific BGP feed; instead, our detector considers the space of *all* possible routes (π, a) , where π is a prefix and a is an origin AS, regardless of whether or not a particular route is visible to a particular BGP collector.

Challenges. The relationship between a single ROA and the validity of BGP routes is complex and dependent on the presence of other ROAs in the system. For example, a ROA giving an AS a prefix π of length 17 up to maxlength 22 actually makes $2^{(23-17)} - 1 = 63$ possible prefixes originated by AS a in BGP become “valid”. If the ROA gets whacked, BGP routes for these prefixes do not necessarily become “invalid”: some may become “unknown” (if there is no covering ROA), while others may remain valid (if there is another ROA for AS a and a super- or subprefix of π). Moreover, when such a ROA appears, it may downgrade “unknown” routes to “invalid” for any subprefix of π of any length, depending on the other ROAs already in the RPKI.

Data structures. Naturally, representing the space of all possible routes (π, a) requires an efficient data structure. Ours is based on the following observations. Consider the complete binary tree of all IP prefixes, with one IP address per leaf, and all possible prefixes as internal nodes. A ROA for prefix π , maxlength m , and origin AS a makes a subtree of this tree “valid” for AS a only; this subtree is rooted at π and goes down to depth m . This ROA also makes “known” (*i.e.*, the complement of “unknown”), for other ASes, the subtree rooted at π and going down to the bottom.

We call these subtrees *triangles* and we represent them using *interval trees*. A triangle has one interval per prefix length. The granularity of each interval depends on the prefix length; intervals at length i have endpoints that are integer multiples of 2^{32-i} . For each prefix length, the total number of intervals kept is bounded by the number of prefixes in ROAs in the RPKI. We can use interval trees to efficiently perform unions, intersections, and complements of sets of triangles. This gives us a “prefix-validity” data structure that stores the validity states for *all* possible routes, based on the ROAs in the RPKI; it can be created in time $O(n \log n)$ given ROAs for n (prefix, AS, maxlength)-tuples.

Our tool uses this data structure to compare a current RPKI state S_{cur} and a previous state S_{prev} . It processes each state once and creates its prefix-validity data structure. It then evaluates the impact of every single change from S_{prev} to S_{cur} in the context of other changes that occurred, by going through each (prefix, AS, maxlength)-tuple




Figure 4: # of invalid IP addresses over time.

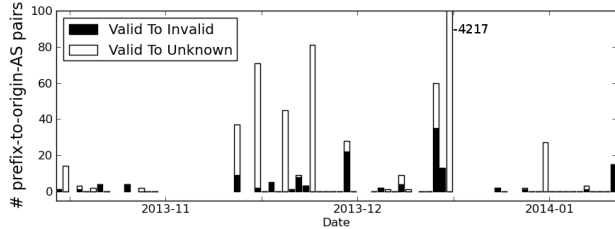


Figure 5: Downgrades due to whacked ROAs.

that appears or disappears from the RPKI (in the transition from S_{prev} to S_{cur}), by extracting relevant information from the prefix-validity data structures.

Utilities for RPKI browsing. Our detector identified all the events in our case studies, and we analyzed each using custom-built python utilities (available upon request) that parse RPKI objects and search RPKI trees based on URI, public key, prefix, and AS number fields.

We also used our detector to check for downgrades between consecutive entries in our trace of the production RPKI:

Downgrades due to added ROAs (Figure 4). Two things happen whenever a ROA is added to the RPKI: (1) routes for the AS authorized in the ROA become “valid”, and (2) the IP address space covered by the ROA transitions to the default-deny state, *i.e.*, all the “unknown” address space it covers downgrades to “invalid”, as in Case Study 1. Figure 4 shows the latter transition over the course of our trace. We show the number of IP addresses that are “invalid” for at least one AS over time, marking each entry in our trace with a diamond. The drop in the number of “invalid” addresses on December 20 was due to the the stale LACNIC manifests described in Case Study 4.

Downgrades due to whacked ROAs (Figure 5). We show the number of (π, a) pairs that downgraded from “valid” to “invalid”, and from “valid” to “unknown”, for each consecutive entry in our trace that contained a successful pull of the RPKI. We show a value of zero if no downgrades occurred at some date, and a gap if our trace was missing an entry. To put these results in perspective, in January 2014, there were about 20,000 prefix-origin AS pairs authorized by ROAs in the RPKI. Most of the incidents in Figure 5 correspond to the whacking of a single ROA containing multiple prefixes; in many incidents, a new ROA appeared authorizing the prefix(es) in the whacked ROA to some other AS. Case Study 4 explains the dramatic event on December 20.

4.2 A tool for visualizing downgrades.

We present a new visualization of the effects of a single ROA on the validity states of multiple BGP routes, based on a modification of the Sierpiński triangle [59]. Our visualization takes in information from our detector and from a specific BGP feed [3, 5], and presents a binary tree that vi-

⁵https://github.com/BUSEC/RPKI_Downgrade_Detector




Figure 6: (l) Visualization of downgrades in Case Study 1. (r) Downgrades when the ROA (63.174.16.0/20, AS 17054) is added to the RPKI in Figure 1.

sualizes the triangles described above, as well as the validity state of specific routes announced in BGP.

Figure 6(l) We visualize Case Study 1. The tool outputs the prefix tree rooted at 173.251.0.0/16, and shows how the triangle rooted at 173.251.0.0/17 transitions from unknown to invalid. The valid route for prefix 173.251.0.0/17 is shown with a grey circle. The added ROA also caused a number of /24 routes in the BGP feed to transition to the “invalid” state; we show these with black circles.

Figure 6(r). Consider the prefixes covered by Continental Broadband in Figure 1, and suppose S_{prev} has the all ROAs issued by Continental Broadband *except* the covering ROA for 63.174.16.0/20, and that the covering ROA is added in S_{cur} . We visualize downgrades from “unknown” to “invalid” in the transition from S_{prev} to S_{cur} . Routes uncovered by the four ROAs in S_{prev} are impacted during the transition; routes covered by the four ROAs in S_{prev} were *already* “invalid” in S_{prev} , and do not appear as downgrades.

5. CHANGES TO RPKI SPECIFICATIONS

While our detector can alarm when changes to the RPKI cause routes in BGP to downgrade to “invalid”, the RPKI still lacks mechanisms that can detect when such changes are legitimate, or when they result from abuse. Because the purpose of the RPKI is to secure routing, not to provide a new technical means for RPKI authorities to take-down IP address space, we now propose modifications to the RPKI specifications to restore the balance of power between RPKI authorities and the subjects of the RCs they issue. We start with our design goals (Section 5.1), overview our design (Section 5.2), and then discuss its details (Sections 5.3-5.4). We then justify our design by presenting its exact security guarantees (Section 5.5), explain why alternative designs would not suffice to provide these guarantees (Section 5.6), and estimate its overhead via data-driven analysis (Section 5.7). This paper overviews our proposal; details and security proofs are in our technical report.

5.1 Design goals.

The following goals have guided our design.

Transparency through alarms. Our overarching goal is to detect RPKI abuse. We want *provable* guarantees that relying parties will raise alarms when something goes wrong; as discussed in Section 3.3, alarms act as a trigger for relying parties to investigate and react to problems with the RPKI. As in the current RPKI specifications, we suppose that relying parties use their own local policies to decide how to resolve alarms [9, Section 6]; alarms could indicate that cer-

tain IP prefixes should be monitored for routing anomalies, or that a relying party should revert to an older (“stale”) set of RPKI objects that did not raise alarms.

The problem of how relying parties should propagate and react to alarms is outside the scope of this paper. First, technical means alone will not be able to resolve a business dispute that results in a unilateral revocation of IP address space in the RPKI; was the issuer of the space at fault, or was it the holder of the address space? Moreover, even requiring different relying parties to agree on a consistent (let alone, correct) course of action to resolve an alarm requires solving the famously-difficult Byzantine agreement problem. We therefore assume that alarms will be resolved by out-of-band mechanisms (phone calls between network operators, social or even legal pressure), just like BGP misconfigurations are resolved today. Nevertheless, to discourage relying parties from raising false alarms, and to facilitate dispute resolution, our secondary goal is *accountability*; allowing relying parties to *provably* demonstrate, in as many situations as possible, (a) why they raised an alarm and (b) who is to blame.

Consent from subjects of RCs. Traditionally, routing has lacked a technical means for revoking IP address space; the RPKI changes this by allowing authorities to *unilaterally* revoke IP address space issued to their descendants. Our goal is to correct this power imbalance, so we require subjects of RCs to *consent* to RPKI modifications that can affect their address space. If consent is refused, this indicates that there is a dispute between the subject of the RC and its issuer, and relying parties should raise an alarm. During emergencies (disputes, lost/stolen keys) when it is impossible to obtain consent from the subject of the RC, the issuer can just unilaterally revoke the RC; these actions will be visible to relying parties, who will raise alarms and investigate the situation out-of-band. Theorem 5.1 formally states how we achieve this goal, and Section 5.7 analyzes the overhead of our consent mechanism.

Consistency for relying parties. We want different relying parties to see consistent views of the RPKI. Our goal is to prevent mirror-world attacks (Section 3.3) by allowing relying parties to compare their local views of the RPKI, and to raise an alarm whenever their views are inconsistent. Theorem 5.2 and 5.3 formally state this goal.

Practicality. Today, the RPKI specifications are still evolving [24, 33], since only a few authorities operate repositories (the five RIRs), and relying parties validate objects using a few evolving software packages like [2]. If our design is to be incorporated in the RPKI specifications, we need to achieve our goals with minimal changes to the current RPKI specifications. Thus, we cannot require synchroniza-

tion between RPKI repositories or relying parties, and must retain most of the structure of the RPKI. Section 5.7 is a data-driven analysis of the practicality of our design.

5.2 Design overview.

To achieve our goals, we propose adding the following new elements to the RPKI. Some seem intuitively necessary; others may be less obvious, but we justify them in Section 5.6.

Consent through .dead objects (Section 5.3.1). Any action that *removes resources* (IP prefixes) or *invalidates* an RC requires the *consent* of the RC and all its impacted descendant RCs. To indicate its consent to having some/all of its resources taken away, an RC will sign a `.dead` object. The authority can then revoke (or remove resources from) the child by publishing the `.dead` object in its publication point and deleting (or overwriting) the child RC.⁶ To simplify the exposition, we will not discuss consent from ROAs; instead we just suppose than any ROA that wishes to be entitled to consent is issued its own covering RC.⁷

Consistency through manifests (Section 5.3.2). In a default-deny architecture like the RPKI, a relying party must know that it has *all* the objects issued by an authority; otherwise, it might misclassify a legitimate BGP route as invalid because its ROA was missing. Fortunately, the RPKI’s manifests contains the hash of every object the authority issues, and thus provide *positive attestation* to the authority’s issued objects [9]. We therefore make manifests central to our design. Relying parties use them to check the consistency of their local caches, because if the manifests are equal, then (by the collision-resistance of the cryptographic hash) so are all the objects the authority issued. Because we cannot require relying parties to synchronize their interactions with the RPKI, we need to make sure they can still check consistency even if they are out of sync. To do this, we enable relying parties to be able to reconstruct all intermediate state of each publication point; to do this, we (1) maintain some extra state at each publication point, and (2) hash-chain manifests, so that each manifest contains the hash of the previous manifest.

Strict checks to raise alarms (Section 5.4). We design validation procedures that, in addition to validating RPKI objects, is also responsible for raising alarms. The procedure has two parts: a local consistency check (done by a single relying party to check consent) and a global consistency check (done by a pair of relying parties to check consistency). In a departure from current RPKI practice, where manifests play a secondary role, our local consistency check treats the manifest as the main cryptographic object, and validates one publication point and one manifest update at a time. Meanwhile, our global consistency check has relying parties compare hashes of the manifests in their local caches.

⁶Per Section 3.2, adding a ROA can also cause routes to downgrade from “unknown” to “invalid”. However, we do not handle this because there is no way to obtain consent from entities originating “unknown” routes, since they do not participate in the RPKI. (If they did, they would already have ROAs, and their routes would not be “unknown”.)

⁷ Actually, each ROA’s parent issues an X.509 end-entity (EE) certificate for an ephemeral one-time-use key, which is used to sign the ROA message [38, Section 2.3]. The ROA and EE cert are stored in a single `.roa` file, so we have treated them as one object. But a ROA could instead consent via its EE cert, instead of asking for its own RC.

5.3 Procedures for RPKI authorities.

We now detail our new `.dead` objects, and our modifications to the RPKI’s specification of manifests [9].

5.3.1 Consent via .dead objects

Suppose an RC A wants to revoke or modify a valid child RC B . (RC B is “modified” if it is overwritten by a new RC B' at the same URL.) A needs the *consent* of B and all the RCs that descend from B that are whacked by the revocation/modification. Consent comes in the form of a `.dead` object, signed each consenting RC.

Revoking B requires a `.dead` object from all the descendants of B . Modifications that remove IP address space from B need `.dead` objects from each descendant of B that overlap with the removed space. No `.dead` objects are required when a modification has no impact on descendants of B (e.g., modifying the parent pointer of B , or increasing the set of IP prefixes or AS numbers it certifies). All other modifications to B must be accomplished by issuing a fresh RC at a different URI. We estimate the number of entities that must be involved in issuing `.dead` objects in Section 5.7.

Constructing a .dead. `.dead` objects are constructed recursively as follows: Let D be a descendant of B that needs to consent. Before D can sign its own `.dead` object, D first collects `.dead` objects from each of its own descendants. D then signs its own `.dead` object, that includes (1) the hash of the `.dead` object issued by each *child* of D , (2) the hash of m_D , the manifest issued by D at the time D signs its own `.dead` object, and (3) the hash of the RC of D . D then provides its own `.dead` object and the `.dead` objects for all its descendants to its issuer C . At this point, D should no longer issue new objects or update its manifest. At the end of this process, A has received the `.dead` objects from B and all relevant descendants of B . (This recursive collection of `.dead` objects protects A , and its descendants, from being falsely accused of revoking a descendant without consent.)

RC A then simultaneously (a) deletes RC B , (b) puts all the `.dead` objects in the publication point of A , and (c) logs all the `.dead` objects in the updated manifest of A .

Key rollover. Any PKI needs a mechanism for refreshing cryptographic keys. Adapting the RPKI’s key rollover mechanism to our design, while still preserving transparency and consistency, requires some care. Details are in our tech report. Importantly, key rollover only requires consent from the RC whose key was rolled.

5.3.2 Manifests as positive attestations

We modify manifests [9] by (a) adding information, and (b) changing the semantics by which they are interpreted.

Normative manifests. We make manifests normative: any objects not logged in the issuer’s current manifest are treated as nonexistent by relying parties. Once manifests become normative, we can simplify other aspects of the RPKI:

1. Only manifests may expire. In the current RPKI, manifest are “updated” when their issuer overwrites them with a fresh manifest with a higher “manifestNumber” [9]; manifests are therefore short-lived objects — most expire and updated with fresh versions within 24 hours. In our design, if a manifest expires before it is updated, then all objects logged in the manifest become “stale”, rather than “invalid”; “stale” indicates that up-to-date information is unavailable, but does *not* indicate that the objects logged in the

expired manifest have explicitly been revoked. A stale manifest (or any object logged in a current manifest but not obtained by a relying party) raises a *missing-information alarm* at the relying party (Section 5.4). Thus, we no longer require expiration dates on ROAs and RCs: any RC/ROA not logged in the current manifest is automatically invalid. This also prevents an authority from issuing short-lived ROAs/RCs so it can circumvent the need to obtain consent.

2. Manifests must log only valid objects. Any issuer that logs an invalid object in its manifest (*e.g.*, an object pointing to the wrong parent, has a prefix that not covered by the issuer’s RC, *etc.*) risks the ire of relying parties, who raise alarms per Section 5.4. (This is necessary for our consistency mechanisms to work properly; see Counterexample 2.) By ensuring that manifests only attest to what is valid, we no longer need CRLs to attest to what is invalid. In fact, RCs/ROAs need not even be signed: signature on a manifest suffices, because the manifest contains the collision-resistant hash of every valid object (as in [25]).

These changes may seem unorthodox, but we note that the RPKI is different from the usual PKI where relying parties obtain and validate objects one-by-one; instead, all objects in a publication point are downloaded *en masse*, so it suffices just to validate the manifest. All these changes can be implemented without any modification to the RPKI object formats by having relying parties ignore CRLs, expiration times/signatures on RCs and ROAs.

Hash chaining. Manifests are hash-chained: each new manifest includes the hash of the contents (excluding the signature) of the manifest it supersedes. A *horizontal chain* is a sequence of consecutively issued manifests, each superseding the next. For each manifest m , hash chaining defines successor and predecessor manifests of manifest m in the obvious way. We also require every manifest to include the hash of the manifest containing its issuer’s RC (“parent manifest”), so that we can definitively determine the resources allocated to an RC at the time it signed its manifest. Thus, a *vertical chain* is a sequence of manifests, each containing the hash of the manifest above it and the hash of an RC that signed the manifest below it.

Reconstructing intermediate states. To prevent mirror world and other attacks (*e.g.*, Counterexample 1 in Section 5.6), relying parties must be able to reconstruct states that could have been seen by other relying parties. Thus, issuers must also provide relying parties with “hints” to enable reconstruction, between two syncs to a publication point, of every intermediate manifest along the horizontal hash chain along with its corresponding publication point state. Our technical report details the small amount of extra information maintained at each publication point for this purpose.

5.4 Validation procedures for relying parties.

The validation procedures for relying parties have three purposes: (1) to determine the set of valid RPKI objects, (2) to raise alarms when something goes wrong, and (3) whenever possible, to hold authorities accountable for causing alarms by proving that they violated procedures in Section 5.3. Since manifests are central to these procedures, we say that a manifest has been *obtained* if the relying party can (a) download a manifest that is validly signed by its issuer, OR (b) reconstruct the manifest as element in the horizontal hash chain that terminates in a manifest that is

alarm	description
missing-information	manifest is stale/missing OR object logged in manifest is missing
bad key rollover	incorrect key rollover procedure
invalid syntax	RC issued malformed object
child too broad	RC issued object it does not cover
unilateral revocation	RC deleted/modified object without <code>.dead</code>
global inconsistency	manifest failed global consistency check

Table 7: Alarms

validly signed by its issuer, per Section 5.3.2.

Initial sync. The initial connection to the repository is as in the current RPKI: objects are downloaded and validated. If a relying party cannot obtain the valid current manifest of an RC R , or any object logged in the manifest, it raises a *missing information alarm*.

Local consistency check. In the current design of the RPKI, updates to a relying party’s local cache are performed all at once. We, instead, require incremental processing of updates: a relying party updates its local cache one publication point and one consecutive manifest along the horizontal hash chain at a time. Updates are only performed for manifests issued by valid RCs. Incremental updates allow us to avoid race conditions that can occur when authorities update their manifests in parallel; see Counterexample 2. New RCs are an exception to this rule: the entire subtree rooted at a new RC should be downloaded and validated immediately to allow new RCs to quickly issue new objects.

Relying parties can sync to publication points in any order and parallelize updates of publication points that are not in an ancestor-descendant relationship. When a relying party Alice obtains an updated state of a publication point, she reconstructs all the intermediate states of the publication point along the horizontal hash-chain. Alice then compares pairs of consecutive states to check that (a) all new or modified objects are valid, and (b) all deletions and overwrites received proper consent. If not, she issues one of the middle alarms in Table 7. In some cases, the alarm is *accountable*, which means that Alice can specifically identify the RPKI authority that misbehaved, and can provide objects that will convince others of the authority’s misbehavior (Section 5.5). Details are in our tech report.

Global consistency check. To defeat mirror-world attacks, relying parties can confirm that other relying parties see the same RPKI objects. A trivial, but unwieldy, solution has relying parties check consistency by synchronizing and exchanging their entire local caches. Our solution dispenses with synchronization and requires the exchange of much less information. We ask only that one party (Alice) is no more than time t_g (“global consistency window”) ahead of the other (Bob).

Bob sends Alice the hash h of the contents (excluding the signature) of latest manifest that he *obtained* for each publication point in his local cache. For every hash value h received from Bob, Alice checks that $h \in H_A$, where H_A is the set that contains the hash of each manifest that Alice obtained, going back for time t_g . The check fails if $h \notin H_A$, and Alice raises a *global-inconsistency alarm* implicating the manifest corresponding to h .

A pairwise interactive protocol is not actually required here; Bob can just post his hash values in a public location for any other relying party to use.

5.5 Security analysis.

We now discuss the security properties of our design. We

only state theorems here; proofs are in our tech report.

Threat model. We suppose that the relying parties named in Theorems 5.1-5.3 are honest; everyone else can be arbitrarily malicious, but cannot break cryptography (e.g., forge digital signatures, find hash collisions, steal keys, etc.). We therefore operate in the standard threat model used in e.g., [20, 36], where trusted auditors (relying parties) audit the behavior of a centralized authority (the RPKI).

We use the following definition: for an RC R and relying party Alice, let the *immediate successor* of R be an RC that (a) overwrites R , or (b) supersedes R after a key rollover, whichever happens first from Alice’s point of view. Define the set of *successors* of R inductively as the set containing R and immediate successors for each of its elements.

Valid remains valid. The following theorem tells us that once an honest relying party Alice sees a valid object, the object will remain valid for Alice until the object consents to revocation (or Alice raises an alarm):

THEOREM 5.1. *Suppose an RC R was valid for a relying party Alice at time t_1 . Consider some time $t_2 > t_1$. Then at time t_2 , at least one of the following is true:*

1. a successor of R is valid in the local cache of Alice and has all the resources of R ;
2. a successor of R is valid in the local cache of Alice, is missing some resources that R had, and Alice observed `.dead` object(s) signed successors of R consenting to the revocation(s) of those resources;
3. at or before time t_2 , Alice saw a `.dead` object signed by a successor of R , consenting to its revocation;
4. at or before time t_2 , Alice raised a unilateral revocation alarm in response to a deleted or overwritten object; the alarm included a successor of R as a victim, and blamed an ancestor (or a successor of the ancestor) of the deleted or overwritten certificate as the perpetrator.

No mirror worlds. We prove robustness to mirror world attacks. Relying parties can be sure that if they see a valid RC in a manifest, and use this manifest in the global consistency check, others who successfully check against them also saw the same valid object:

THEOREM 5.2. *Suppose the local cache of a relying party Bob contains a valid RC R and a manifest m that was issued by the parent of R . Suppose that R is not marked as stale; thus, the version of R in Bob’s local cache is same as the one logged in m . Suppose another relying party Alice performs a global consistency check against Bob; Bob sends the hash of m to Alice, and Alice looks for it in her set H_A of hashed manifests. Suppose the global consistency check does not raise the global-inconsistency alarm for m . Then either*

1. Alice raised an alarm with R as the victim when performing the local consistency check on a manifest in H_A , or
2. at any time t_2 after the global consistency check, at least one condition in Theorem 5.1 holds for R and Alice.

No mirror worlds in the past. Because manifests are hash-chained, a successful global consistency check also implies that parties were consistent in the past. We state the theorem only informally here:

THEOREM 5.3 (INFORMAL). *Suppose neither Alice nor Bob raised alarms for k manifest updates in a row. Then if*

Alice passes a global consistency check against Bob, then for every RC R seen as valid by Bob in any of the past k manifests, at least one of the conditions specified in Theorem 5.1 holds for R and Alice.

Alarms & accountability. Our theorems assume that relying parties Alice and/or Bob are *honest*, and ensure that honest relying parties raise alarms when their local caches are missing/contain objects that violate our specifications. But we cannot, in general, prevent *dishonest* relying parties from raising false alarms, because it is impossible to prevent someone from falsely claiming they are missing information. As such, our alarms come in two categories: *accountable* and *unaccountable*.

If Alice raises an accountable alarm, then she is certain that some RPKI authority misbehaved, she can identify this authority (the “perpetrator”), and she can prove her knowledge to others (even those who do not trust Alice) by publishing the relevant objects from her local cache. Assuming that cryptography is not broken, even a malicious Alice cannot create an accountable alarm that is false.

Alice raises an unaccountable alarm when she is missing information — when, from Alice’s point of view, there is not enough information to distinguish between an attack by an RPKI authority, and a disruption to the delivery of RPKI objects (e.g., by a third-party that tampers with the communication path). While Alice will still raise an alarm in these cases, the alarm will not be provable to others, and even Alice herself cannot tell who perpetrated the alarm. If the missing information is found, then the alarm will either be resolved or become accountable.

To ensure that as many alarms as possible are accountable, we require authorities to commit (by means of hash values) to the context that is necessary to determine the validity of objects. For example, the validity of a `.dead` depends on its signer’s manifest; thus, a `.dead` includes the hash of its signer’s manifest (Section 5.3.1). Similarly, the validity of manifest m depends on the validity of the manifest that logs the RC that issued m ; thus, m includes the hash of its parent’s manifest (Section 5.3.2). An alarm is accountable if these objects are present in Alice’s local cache, and unaccountable if they are missing. A missing information alarm is always unaccountable. Our tech report details when other alarms are accountable.

5.6 On the necessity of our modifications.

We give intuition for our design via two examples.


Counterexample 1: Not checking intermediate state.

Suppose relying parties did not verify every consecutive manifest update. An authority X could then launch a mirror world attack violating Theorems 5.2 and 5.3:

At time t_1 , X issues an RC Y ; at time t_2 , X replaces Y with a new valid RC Y' that contains an additional IP prefix; and continues swapping between Y (at odd-numbered times), and Y' (at even-number times). Alice syncs to X ’s publication point at odd-numbered times, and sees only Y ; Bob syncs at even-numbered times and sees RC Y' . But, since Alice does not check full intermediate states (only manifest chains), she does not notice the transition from Y' to Y , and does not realize that X needs to get a `.dead` object from Y' . Thus, Alice and Bob live in mirror worlds.

Counterexample 2: Race conditions. A major challenge we faced was that relying parties can sync to different

publication points at different times. This creates race conditions that complicate global consistency:



Suppose authority X has small address block. At time t_1 , X issues a child RC Y that is invalid because Y contains more addresses than its parent X . At time t_2 , parent X is overwritten by an RC for those addresses, and its child Y becomes valid. At time t_3 , parent X overwrites Y with an RC that is invalid because it contains even more addresses. This process continues as above. Notice that from time t_1 to t_2 , and from t_3 to t_4 , the manifest of X logs an invalid object. Suppose Alice syncs just after t_1 and just after t_3 ; she decides that Y is invalid. Bob syncs just after t_2 and just after t_4 ; he decides that Y is valid. Hence, mirror worlds!

This problem cannot be caught by checking for `.dead` objects for Y , since in Alice’s view Y is always invalid, and in Bob’s view Y just keeps getting more resources. Nor can it be caught by the global consistency check, because Alice and Bob see the same manifests. Our design eliminates this problem by requiring relying parties to alarm if invalid objects are logged in the manifest; in this example, Alice would raise an alarm. Theorems 5.2 and 5.3 would be false without this requirement.

5.7 Data-driven analysis of our design.

We discuss the impact of our changes on the RPKI.

Less crypto. One immediate improvement is that a single digital signature on the manifest needs to be issued and verified, rather than individual signatures on each RC, ROA, CRL, and manifest (Section 5.3.2). To put this perspective, on January 13, 2014 there were $\approx 10,400$ validly-signed objects in the RPKI; our changes require the validation of only $\approx 2,800$ manifests.

No renewals. RCs/ROAs do not expire in our design (Section 5.3.2); hence recipients of resources are no longer dependent on their issuers for routine renewals.

Mandated interaction for obtaining consent. However, issuers are dependent on recipients to provide consent (via `.dead` objects) for revocations and certain modifications (Section 5.3.1). To find out how often those events happen, we use our trace of the production RPKI for 2013/10/23–2014/01/21 (Section 3). The largest event we observed was in mid-November 2013, when RIPE removed 3,336 RCs/ROAs and issued new ones with new parent/child pointers and public keys, as part of repository restructuring. Our design would require RIPE to obtain `.dead` objects from all of them, which seems to impose a large burden on RIPE. However, unless RIPE holds the secret keys of its descendants (in which case it can just issue `.dead` objects by itself), interaction is needed even if `.dead` objects were not required, because RIPE’s descendant RCs would need to reissue their objects in new publication points.

Besides this event, we saw 4,443 instances of modified/revoked RCs/ROAs. Of these, 3,569 (80%) were renewals, that are not needed in our design. Meanwhile, at most 230 (5%) would need a `.dead` object.

How many parties need to consent? Next, we con-

# ASes	1	2	3	4	5	6-10	10-30	98
RIPE	678	122	51	13	12	30	8	1
LACNIC	123	20	9	2	1	2	0	0
APNIC	26	8	2	0	2	0	0	0
ARIN	30	5	4	4	3	0	0	0
AfriNIC	9	2	1	1	0	0	0	0

Table 8: # of leaf RCs issuing ROAs for X ASes on January 13, 2014; X is in the top row.

# ASes	1-10	11-30	31-100	100-200	200 – 1073
	115,605	594	132	15	11

Table 9: Similar to the distribution in Table 8, except for direct-allocation RCs in our model.

sider how many parties need to be involved in signing a `.dead` when an RC is revoked. Our estimates, made from the production RPKI and a model of the RPKI, suggest that we do not require many `.dead` objects *on average*:

1. **Production RPKI.** Table 2 shows the structure of the production RPKI on January 13, 2014. Suppose that ROAs were able to sign off on `.dead` objects (*e.g.*, because they had requested their own covering RCs, or via the mechanism in footnote 7). How many entities would need to sign a `.dead` object if we wanted to revoke a leaf RC in Table 2? We use ASes as a proxy for entities, and estimate this by counting the number of ASes in ROAs issued by each leaf RC; the results for the production RPKI are in Table 8, which shows that on average, only 1.6 ASes need to consent to revoking a leaf RC, and that 93% of leaf RCs can be revoked with the consent of no more than 3 ASes.

2. **Model.** Because the RPKI is far from fully deployed, we also created a model for a full deployment of the RPKI using routing data for the week starting 2012/05/06. Our model has the RIRs sit at the highest layer of the hierarchy; they issues RCs to “direct allocations,” *i.e.*, IP prefixes directly allocated by the RIRs (*e.g.*, Sprint in Figure 1). We obtained these top two layers of the hierarchy from files retrieved from the FTP site of each RIR. Our model omits intermediate RCs (since they are also held by RIRs). To model ROAs descended from each direct allocation, we extracted (prefix, origin AS)-tuples from BGP feeds [3, 5] for the week starting 2012/05/06. (Any tuple not covered by a directly-allocated prefix was discarded as a bogon.)⁸We grouped tuples by AS and found that, on average, each direct allocation issues ROAs for only 1.5 ASes; the distribution is in Table 9.

With great power comes great responsibility. Tables 8, 9 indicate that there are a small percentage of outlier RCs that issue ROAs for many (even hundreds!) of ASes. (In our model, out of 116,357 total direct-allocation RCs, 26 (0.02%) have more than 100 ASes and 221 (0.18%) have more than 25 ASes.) Revoking these outliers requires a large number of `.dead` objects. However, we consider this to be a feature, not a bug: these RCs can impact routing to a large number of ASes, so revoking them should not be easy. Moreover, outright revocation may not be necessary if the goal is simply to change the resources given to them slightly, because we provide mechanisms for removing resources from RC that only require `.dead` objects from descendants that become invalid as a result (Section 5.3.1).

⁸Figure 1 is derived from this model. We built a subtree of RCs below each direct allocation RC, with one RC for every prefix with (prefix, AS)-tuple in the BGP feeds. The ancestor relationship corresponds to the cover relationship for prefixes, and we collapsed parent-child pairs of RCs generated for the same AS, ASes in the same organization per [16] or when if child was a “stub AS” per [63].

6. RELATED WORK

Routing security. The RPKI is a realization of a series of routing security proposals [6, 32, 48, 64] for origin authentication (Section 2.2). A number of works [10, 26, 39] argue that origin authentication can significantly improve routing security. The RPKI is also the first step towards a comprehensive solution for securing the current routing system with BGPSEC [37] or other proposals surveyed in [14, 31]. See also [65] for clean-slate architectures.

Censorship. Censorship is known to occur at all layers of the Internet’s architecture; see *e.g.*, [45] for an overview. There is already evidence [8, 56] of routing-based censorship with BGP; we explored the risk that the RPKI could also be used for this purpose.

PKI design. Our modified RPKI architecture (Section 5) is related to an long line of work on public key infrastructure (PKI) design, culminating in recent efforts to harden the web PKI [47]. The idea of validating the state of the RPKI over time is related to certificate pinning [23], and the idea of hash-chaining manifests is related to append-only-logs [58] (or see [20] and references therein) and ideas related to certificate transparency [22, 34, 36]. While these works suppose that a single logger tracks a stream of ordered events, we have to deal with the race conditions (Counterexample 2) that result from allowing individual RPKI authorities to maintain their own logs (*i.e.*, manifests). The idea of correcting power imbalances in a hierarchical system also appeared in work on distributing certificate authorities [66] and centralized systems like the DNS [1, 15, 55]; these works distribute the *issuance* of objects, but we only distribute *revocation* (since revocation can harm IP prefix reachability).

The RPKI. Research on the RPKI covers measurement [49, 62] and policy questions [17, 43, 44, 61]. Our earlier HotNets paper [18] discusses the threat of misbehaving RPKI authorities, but does not provide solutions to any the threats it discusses; part of Section 3 overlap with [18], but the remainder of this paper presents new results.

Concurrently to our work, there have been efforts within the IETF to harden the RPKI against authorities that abuse their power [11, 13, 33]. Kent *et al.* [33] considers the threat of whacked ROAs and of “competing ROAs”. A new ROA “competes” with an existing ROA if it contains prefixes covered by the older ROA; a competing ROA is a threat if BGP is attacked, since the AS in the competing ROA can perform a (sub)prefix hijack on the AS in the older ROA. Our architecture (Section 5) ignores this threat because we focus on the risk that the RPKI can take IP prefixes offline in the *absence* of an attack on BGP. Moreover, any authority that issues a competing ROA and then attacks BGP can be held accountable; the competing ROA itself is non-repudiable evidence of the attack. Both [33] and our design defend against whacked ROAs by comparing the state of the RPKI over time, but we also detect mirror world attacks and have a consent mechanism and exact security guarantees.

Systems have also been developed to monitor the RPKI [4, 35, 46, 57, 60]; most use a snapshot of ROAs from the RPKI to determine the validity state of routes in publicly-available BGP route collectors [3, 5]. Our detector (Section 4.1) is complementary because we detect when any *change* to the RPKI alters the validity state of *all* possible routes, not just the ones visible from a particular BGP vantage point at a specific time; it can therefore be used as an alert system

(especially when RPKI deployment reaches steady state), even if a particular vantage point does not obtain a complete view of all routes announced in BGP. Our tools also provide a new way of visualizing downgrade events (Section 4.2).

7. CONCLUSION

We have explored a number of techniques to harden the RPKI against the risk of IP prefix takedowns. We have built tools for detecting takedowns within the existing RPKI specifications. We have also proposed changes to the specifications that (1) entitle parties to *consent* to revocations of their IP address space, and guarantee that relying parties can (2) detect misbehavior by RPKI authorities and (3) obtain a *consistent* view of information in the RPKI. Given the security improvements promised by the RPKI [10, 26, 39], we hope our work will catalyze further efforts to harden the RPKI against abusive authorities.

Acknowledgements. We thank Kyle Brogle for help with early stages of this work, Steve Kent for suggesting the term “ROA whacking”, FCC CSRIC Working Group 6 - Secure BGP Deployment for inspiring this project, Tony Tauber, Ed Felten and various members of the IAB and the IETF SIDR working group for useful discussions, CAIDA for early access to their AS to organization mapping [16], and John Byers, Davide Proserpio, Ben Fuller, Phillipa Gill, Aanchal Malhotra, our shepherd Adrian Perrig, and the anonymous SIGCOMM reviewers for insightful comments on this draft. This work was supported by NSF awards 1017907, 1012798, and 1012910 and a gift from Cisco.

8. REFERENCES

- [1] Github: namecoin repository. <https://github.com/namecoin/namecoin>.
- [2] rcynic software. <http://trac.rpki.net>.
- [3] RIPE RIS raw data. <http://www.ripe.net/data-tools/stats/ris/ris-raw-data>.
- [4] Rpk spider. <http://rpkipspider.verisignlabs.com/>.
- [5] University of oregon route views project. <http://www.routeviews.org/>.
- [6] W. Aiello, J. Ioannidis, and P. McDaniel. Origin authentication in interdomain routing. In *Proceedings of the 10th ACM conference on Computer and communications security*, pages 165–178. ACM, 2003.
- [7] S. Amante. Risks associated with resource certification systems for internet numbers, 2012.
- [8] D. Anderson. Splinternet behind the great firewall of china. *Queue*, 10(11):40, 2012.
- [9] R. Austein, G. Huston, S. Kent, and M. Lepinski. *RFC 6486: Manifests for the Resource Public Key Infrastructure (RPKI)*. Internet Engineering Task Force (IETF), 2012. <http://tools.ietf.org/html/rfc6486>.
- [10] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the Internet. In *SIGCOMM’07*, 2007.
- [11] R. Bush. *Responsible Grandparenting in the RPKI*. Internet Engineering Task Force Network Working Group, 2012. <http://tools.ietf.org/html/draft-ymbk-rpki-grandparenting-02>.
- [12] R. Bush. *RPKI-Based Origin Validation Operation*. Internet Engineering Task Force Network Working Group, 2012. <http://tools.ietf.org/html/draft-ietf-sidr-origin-ops-19>.
- [13] R. Bush. *RPKI Local Trust Anchor Use Cases*. Internet Engineering Task Force (IETF), 2013. <http://www.ietf.org/id/draft-ymbk-lta-use-cases-00.txt>.
- [14] K. Butler, T. Farley, P. McDaniel, and J. Rexford. A survey of BGP security issues and solutions. *Proceedings of the IEEE*, 2010.
- [15] C. Cachin and A. Samar. Secure distributed dns. In *Dependable Systems and Networks, 2004 International Conference on*, pages 423–432. IEEE, 2004.
- [16] CAIDA. AS to organization mapping. <http://as-rank.caida.org/?mode0=as-intro#as-org>.

- [17] Communications Security, Reliability and Interoperability Council III (CSRIC). Secure bgp deployment. *Communications and Strategies*.
- [18] D. Cooper, E. Heilman, K. Brogle, L. Reyzin, and S. Goldberg. On the risk of misbehaving RPKI authorities. *HotNets XII*, 2013.
- [19] J. Cowie. Rensys blog: China's 18-minute mystery. <http://www.renysys.com/blog/2010/11/chinas-18-minute-mystery.shtml>.
- [20] S. A. Crosby and D. S. Wallach. Efficient data structures for tamper-evident logging. In *USENIX Security Symposium*, pages 317–334, 2009.
- [21] A. de Beaupre. ISC Diary: Multiple Banking Addresses Hijacked, 2013. <http://isc.sans.edu/diary/BGP+multiple+banking+addresses+hijacked/16249>.
- [22] P. Eckersley. Sovereign key cryptography for internet domains. Technical report, EFF, 2011.
- [23] C. Evans, C. Palmer, and R. Slevi, editors. *Public Key Pinning Extension for HTTP*. IETF Web Security, Internet-Draft, November 27 2013. <http://tools.ietf.org/html/draft-ietf-websec-key-pinning-09>.
- [24] R. Gagliano, T. Manderson, and C. M. Cagnazzo. *Multiple Repository Publication Points support in the Resource Public Key Infrastructure (RPKI)*. Internet Engineering Task Force (IETF), 2013. <http://tools.ietf.org/html/draft-ietf-sidr-multiple-publication-points-00>.
- [25] I. Gassko, P. Gemmill, and P. D. MacKenzie. Efficient and fresh cerification. In H. Imai and Y. Zheng, editors, *Public Key Cryptography*, volume 1751 of *Lecture Notes in Computer Science*, pages 342–353. Springer, 2000.
- [26] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How secure are secure interdomain routing protocols? In *SIGCOMM'10*, 2010.
- [27] E. Goldman. Sex.com: An update. http://blog.ericgoldman.org/archives/2006/10/sexcom_an_updat.htm, 2006.
- [28] G. Huston, R. Loomans, and G. Michaelson. *RFC 6481: A Profile for Resource Certificate Repository Structure*. Internet Engineering Task Force (IETF), 2012. <http://tools.ietf.org/html/rfc6481>.
- [29] G. Huston and G. Michaelson. *RFC 6483: Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)*. Internet Engineering Task Force (IETF), 2012. <http://tools.ietf.org/html/rfc6483>.
- [30] G. Huston, G. Michaelson, and S. Kent. *RFC 6489: Certification Authority (CA) Key Rollover in the Resource Public Key Infrastructure (RPKI)*. Internet Engineering Task Force (IETF), 2012. <http://tools.ietf.org/html/rfc6489>.
- [31] G. Huston, M. Rossi, and G. Armitage. Securing BGP: A literature survey. *Communications Surveys & Tutorials, IEEE*, 13(2):199–222, 2011.
- [32] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *J. Selected Areas in Communications*, 18(4):582–592, April 2000.
- [33] S. Kent and D. Mandelberg. *Suspenders: A Fail-safe Mechanism for the RPKI*. Internet Engineering Task Force (IETF), 2013. <http://tools.ietf.org/html/draft-kent-sidr-suspenders-00>.
- [34] T. H.-J. Kim, L.-S. Huang, A. Perrig, C. Jackson, and V. Gligor. Accountable Key Infrastructure (AKI): A Proposal for a Public-Key Validation Infrastructure. In *Proceedings of the International World Wide Web Conference (WWW)*, May 2013.
- [35] LACNIC. RPKI looking glass. www.lacnic.net/rpkitools/looking_glass/.
- [36] B. Laurie, A. Langley, and E. Kasper. Certificate transparency. *Network Working Group Internet-Draft, v12, work in progress*. <http://tools.ietf.org/html/draft-laurie-pki-sunlight-12>, 2013.
- [37] M. Lepinski, editor. *BGPSEC Protocol Specification*. IETF Network Working Group, Internet-Draft, July 2012. Available from <http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-04>.
- [38] M. Lepinski and S. Kent. *RFC 6480: An Infrastructure to Support Secure Internet Routing*. Internet Engineering Task Force (IETF), 2012. <http://tools.ietf.org/html/rfc6480>.
- [39] R. Lychev, S. Goldberg, and M. Schapira. Is the juice worth the squeeze? BGP security in partial deployment. In *SIGCOMM'13*, 2013.
- [40] T. Manderson, L. Vegoda, and S. Kent. *RFC 6491: Resource Public Key Infrastructure (RPKI) Objects Issued by IANA*. Internet Engineering Task Force (IETF), 1973. <http://tools.ietf.org/html/rfc6491>.
- [41] S. Misel. "Wow, AS7007!". Merit NANOG Archive, apr 1997. www.merit.edu/mail.archives/nanog/1997-04/msg00340.html.
- [42] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein. *RFC 6811: BGP prefix origin validation*. Internet Engineering Task Force (IETF), 2013. <http://tools.ietf.org/html/rfc6811>.
- [43] M. Mueller and B. Kuerbis. Negotiating a new governance hierarchy: An analysis of the conflicting incentives to secure internet routing. *Communications and Strategies*, (81):125–142, 2011.
- [44] M. Mueller, A. Schmidt, and B. Kuerbis. Internet security and networked governance in international relations. *International Studies Review*, 15(1):86–104, 2013.
- [45] S. J. Murdoch and R. Anderson. *Access Denied: The Practice and Policy of Global Internet Filtering*, chapter Tools and technology of Internet filtering, pages 57–72. MIT Press, 2008.
- [46] NIST. RPKI deployment monitor. <http://www-x.antd.nist.gov/rpki-monitor/>.
- [47] NIST. Workshop on Improving Trust in the Online Marketplace, 2013. <http://www.nist.gov/itl/csd/ct/ca-workshop-agenda2013.cfm>.
- [48] 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*, page 12. ACM, 2011.
- [49] E. Osterweil, T. Manderson, R. White, and D. McPherson. Sizing estimates for a fully deployed rpki. Technical report, Verisign Labs Technical Report, 2012.
- [50] A. Peterson. Researchers say u.s. internet traffic was re-routed through belarus. that's a problem. *Washington Post*, November 20 2013.
- [51] A. Pilosov and T. Kapela. Stealing the internet, 2009.
- [52] D. Piscitello. Guidance for preparing domain name orders, seizures & takedowns. Technical report, ICANN, March 2012.
- [53] D. Piscitello. The value of assessing collateral damage before requesting a domain seizure. Technical report, ICANN, 2013.
- [54] I. G. Project. In important case, RIPE-NCC seeks legal clarity on how it responds to foreign court orders, 2011. <http://www.internetgovernance.org/2011/11/23/in-important-case-ripe-ncc-seeks-legal-clarity-on-how-it-responds-to-foreign-court-orders/>.
- [55] V. Ramasubramanian and E. G. Sirer. The design and implementation of a next generation name service for the internet. *ACM SIGCOMM Computer Communication Review*, 34(4):331–342, 2004.
- [56] Rensys Blog. Pakistan hijacks YouTube. http://www.renysys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml.
- [57] RIPE. RPKI validator. <http://localcert.ripe.net:8088/trust-anchors>.
- [58] B. Schneier and J. Kelsey. Automatic event-stream notarization using digital signatures. In *Security Protocols*, pages 155–169. Springer, 1997.
- [59] W. Sierpiński. Sur une courbe dont tout point est un point de ramification. *Comptes Rendus de l'Académie des Sciences*, 160:302–305, 1915.
- [60] Surfnet. RPKI dashboard. <http://rpki.surfnet.nl/validitytables.html>.
- [61] The President's National Security Telecommunications Advisory Committee. Nstac report to the president on communications resiliency, 2011.
- [62] M. Wählisch, O. Maennel, and T. Schmidt. Towards detecting BGP route hijacking using the RPKI. In *Poster: SIGCOMM'12*, pages 103–104. ACM, 2012.
- [63] L. Wang, J. Park, R. Oliveira, and B. Zhang. Internet topology collection. <http://irl.cs.ucla.edu/topology/>.
- [64] R. White. Deployment considerations for secure origin BGP (soBGP). draft-white-sobgp-bgp-deployment-01.txt, June 2003, expired.
- [65] 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 *IEEE Security and Privacy (SP)*, 2011.
- [66] L. Zhou, F. B. Schneider, and R. Van Renesse. COCA: A secure distributed online certification authority. *ACM Transactions on Computer Systems (TOCS)*, 20(4):329–368, 2002.