Are We There Yet? On RPKI's Deployment and Security

Yossi Gilad*, Avichai Cohen†, Amir Herzberg‡, Michael Schapira† and Haya Shulman§
*Boston University and MIT CSAIL
†Hebrew University of Jerusalem
‡Bar Ilan University
§Fraunhofer SIT

***Abstract*— The Resource Public Key Infrastructure (RPKI) binds IP address blocks to owners' public keys. RPKI enables routers to perform Route Origin Validation (ROV), thus preventing devastating attacks such as IP prefix hijacking. Yet, despite extensive effort, RPKI's deployment is frustratingly sluggish, leaving the Internet largely insecure. We tackle fundamental questions regarding today's RPKI's deployment and security: What is the adoption status of RPKI and ROV? What are the implications for global security of *partial* adoption? What are the root-causes for slow adoption? How can deployment be pushed forward? We address these questions through a combination of empirical analyses, a survey of over 100 network practitioners, and extensive simulations. Our main contributions include the following. We present the *first* study measuring ROV enforcement, revealing disappointingly low adoption at the core of the Internet. We show, in contrast, that without almost ubiquitous ROV adoption by large ISPs significant security benefits cannot be attained. We next expose a critical security vulnerability: *about a third* of RPKI authorizations issued for IP prefixes do not protect the prefix from hijacking attacks. We examine potential reasons for scarce adoption of RPKI and ROV, including human error in issuing RPKI certificates and inter-organization dependencies, and present recommendations for addressing these challenges.**

I. Introduction

The Border Gateway Protocol (BGP) computes routes between the tens of thousands of smaller networks, called Autonomous Systems (ASes), which make up the Internet. ASes range from large ISPs and content providers to small businesses and universities. BGP is notoriously vulnerable to devastating attacks and configuration errors. Consequently, nation states and corporations are in constant danger from attacks that utilize BGP's insecurity to disconnect ASes from the Internet and to launch highly effective man-in-the-middle attacks. A particularly worrisome and common attack vector is *IP prefix hijacking*, where an AS advertises in BGP an IP prefix not belonging to it. Prefix hijacks are effective and easy to launch, with the extra benefit of a plausible excuse: benign configuration errors [53]. Every year several high-profile incidents resulting from prefix hijacks make the news (e.g., [1], [8], [51], [52]), and many others go under the radar [55].

Permission to freely reproduce all or part of this paper for noncommercial purposes is granted provided that copies bear this notice and the full citation on the first page. Reproduction for commercial purposes is strictly prohibited without the prior written consent of the Internet Society, the first-named author (for reproduction of an entire paper only), and the author's employer if the paper was prepared within the scope of employment.
NDSS '17, 26 February - 1 March 2017, San Diego, CA, USA
Copyright 2017 Internet Society, ISBN 1-1891562-46-0
http://dx.doi.org/10.14722/ndss.2017.23123

The *Resource Public Key Infrastructure (RPKI)* [38] is a hierarchical certification system. RPKI stores *Route Origin Authorizations (ROAs)*, signed records that bind an IP address block to the AS that is allowed to advertise it in BGP. ROAs can be leveraged by BGP routers to perform *Route-Origin Validation (ROV)* [9], [45]: identifying and discarding "invalid" BGP route-advertisements from unauthorized ASes, thus protecting against IP prefix hijacking. Beyond being the leading and, thus far, only standardized solution to prefix hijacking, RPKI is also a prerequisite for prominent routing security mechanisms such as BGPsec [42], and for other proposals for defending against BGP path-manipulation attacks, e.g., [15], [16], [58].

Yet, despite RPKI's crucial role in securing the Internet routing system, RPKI's deployment is frustratingly low [47], leaving the Internet largely exposed to dangerous traffic hijacking attacks. We embark on a systematic study of RPKI's deployment and security. We tackle the following fundamental questions: What is the adoption status of RPKI and ROV? What are the implications for security of *partial* adoption of RPKI/ROV? What are the root-causes for the slow adoption rates? How can deployment be pushed forward? We address these questions through a combination of empirical analyses of multiple datasets, a survey of over 100 network practitioners, and extensive simulations on empirically-derived data. We next describe our main contributions.

A. Deployment Status

Extensive effort is devoted to promoting the adoption of RPKI and ROV. The IETF's SIDR (Secure Inter-Domain Routing) working group standardized the relevant protocols, and router support for ROV is already available from all major vendors. For RPKI to provide meaningful security guarantees, two conditions must be met: (1) ASes should issue ROAs covering their address blocks, and (2) ASes should configure BGP routers to do ROV, i.e., discard BGP advertisements that are incompatible with ROAs. Only about 6.5% of IP prefixes advertised in BGP are covered by ROAs [47], [57]. (We will discuss the root-causes for this scarce deployment later on.) What about adoption of ROV?

Any AS, *anywhere*, can enforce ROV, provided that its routers support ROV (most modern routers do). We aim to understand how many ASes actually do. We present the

first quantitative study of ROV adoption. Our findings are disappointing: only a very small fraction of the large ISPs at the core of the Internet enforce ROV.

B. Security Evaluation

Security in partial deployment. Given our results for ROV adoption, a natural question arises: What is the impact on global routing security of *partial* adoption of ROV? We identify interesting phenomena that can manifest when ROV is not globally adopted. Not surprisingly, ROV enforcement by an AS can yield *collateral benefits*, protecting others routing through that AS from traffic hijacking attacks. Less obviously, however, non-adoption by an AS can cause *collateral damage* resulting in *ROV-enforcing* ASes falling victim to attacks.

To gain insights into how these effects impact security, we present the results of extensive simulations of partial ROV-adoption scenarios. We find that ROV is effective against prefix hijacking *if and only if* almost all of the largest ISPs (e.g., top 100) enforce ROV. Consequently, under today's meagre ROV adoption, an AS that issues a ROA for its IP address blocks is largely unprotected from such attacks.

Insecure ROAs. We next expose a critical security vulnerability of RPKI: about a *third* of the IP prefixes in BGP tables that are covered by ROAs are badly issued and, consequently, the issuing organizations remain vulnerable to traffic hijacking. In fact, an attacker can hijack *all* traffic destined for the issuers of such "insecure ROAs" even if *all* ASes on the Internet do ROV. We show that even large ISPs, e.g., Swisscom and Orange, suffer from this vulnerability. (We present, throughout this paper, real-world examples of security vulnerabilities and other RPKI-related undesirable phenomena. All these problems have been reported to network operators.)

C. Obstacles to Adoption

While the hurdles to deployment of BGPsec were studied in depth (e.g., see [23], [41]), fairly little attention has been given to RPKI's (non)deployment. Clearly, RPKI and ROV suffer from a circular dependency: ASes do not gain sufficient security benefits from ROV so long as so many IP prefixes advertised in BGP are not certified through RPKI, yet there is little incentive to certify ownership over an IP prefix (issue a ROA) when so few ASes do ROV.

This circular dependency is not unique to RPKI and ROV. Similar dependencies hinder adoption of other standards, e.g., DNSSEC and IPv6 (see, for example, DNSSEC adoption rates [18], [22]). RPKI and ROV, however, do not share other major drawbacks of these other standards, e.g., interoperability concerns (i.e., compatibility with legacy protocols), runtime overheads (e.g., resulting from online cryptography), etc. What, then, are the main factors hindering RPKI deployment?

Human error. Beyond insufficient value/incentive for adoption, RPKI and ROV suffer from an even worse problem: (justified) *mistrust*. Past studies [19], [33], [34] show that about 8% of IP prefixes advertised in BGP that are covered by ROAs are "invalid". Consequently, an AS doing ROV would be immediately disconnected from thousands of legitimate IP prefixes owned by hundreds of organizations, an arguably worse outcome than the attack RPKI is meant to prevent. Our survey of network operators confirms that mistrust in the RPKI infrastructure is indeed an important reason for non-deployment. Our empirical measurement results, however, reveal a more nuanced (and somewhat more optimistic) picture: a breakdown of invalid IP prefixes by organizations shows that roughly 20 organizations are responsible for the invalidity of most of these IP prefixes, giving hope that this situation can be greatly improved via highly focused efforts.

Inter-organization dependencies. Another potential reason for RPKI's slow certification rates is *inter-organization dependencies* that force an AS that wishes to issue a ROA to wait for other ASes to issue ROAs first. We show that while such dependencies do not pertain to most ASes, some of the world's largest ISPs cannot easily issue ROAs for this reason.

D. Recommendations

We present several concrete proposals for driving RPKI adoption forward. We advocate concentrating standardization, regulatory, and educational efforts in the following directions.

Eliminating obstacles to deployment via modest changes to RPKI. We introduce a new kind of ROAs, compatible with the current format, which can eliminate troublesome inter-organization dependencies.

Targeting the core of the Internet. Our security evaluation of ROV in partial deployment shows that adoption by the largest ISPs is key to attaining meaningful security guarantees. We thus advocate targeting effort (regulatory, providing proper incentives, etc.) on boosting ROV adoption amongst the top ISPs.

Improving RPKI's reliability and integrity. Our results on RPKI deployment show that human error (when issuing ROAs) is a key factor in RPKI's unreliability and insecurity and that relatively few organizations are responsible for the majority of badly issued ROAs. Thus, concentrating effort on these organizations can lead to a drastic decrease in the number of badly issued ROAs. To alert network operators about "bad ROAs" affecting their connectivity or security, and inform them about their *organization-specific* obstacles to RPKI deployment, we present *ROAlert*, identifying incorrect ROAs and alerting administrators via email and via web. We encourage the reader to enter the ROAlert website at roalert.org to discover whether her/his IP address (or any other IP address of her/his choice) is correctly protected by a ROA in RPKI and, if not, why (the relevant terminology is presented in the sequel). Our experience with ROAlert, including feedback from dozens of network administrators, indicates that ROAlert, if employed and promoted by the relevant entities, could aid in building trust in RPKI and boosting its adoption.

E. Organization

Section II provides the required background on RPKI and ROV. We discuss ROV's adoption status and its impact on global routing security in Sections III and IV, respectively. We present our results on the insecure ROAs in Section V, exposing a vulnerability of RPKI deployment in practice. We discuss and quantify obstacles to adoption in Section VI and introduce, in Section VII, our proposals for facilitating wider adoption. We review related work in Section VIII and conclude in Section IX.

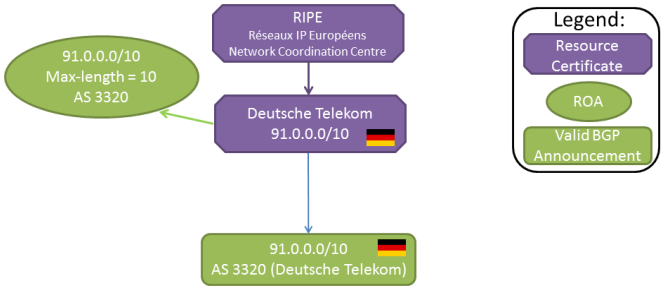

Fig. 1: Deutsche Telekom received an RC from RIPE and issued a ROA to protect its IP-prefix (advertised through BGP)

II. Background: RPKI, RCs, ROAs, and ROV

RPKI [38] maps IP address blocks to organizations that "own" them, thus laying the foundation for combating IP prefix hijacking and for further protection from attacks on inter-domain routing via BGPsec and alternative proposals for defending against BGP path-manipulation attacks [15], [16], [42], [58]. We next elaborate on how ownership of IP prefixes is certified in RPKI and how this information can be used to discard bogus BGP route-advertisements.

A. Certifying Ownership

RPKI assigns an IP prefix to a public key via a *Resource Certificate (RC)*, issued by the authoritative entity for that IP prefix. This allows the owner of the corresponding private key to issue a *Route Origin Authorization (ROA)* specifying the AS numbers of ASes authorized to advertise the IP prefix in BGP. We explain below some of the specifics involved in issuing RCs and ROAs.

RPKI RCs form a certification hierarchy as follows. At the top of the hierarchy are the five Regional Internet Registries (RIRs). Each RIR holds a root (self-signed) RC covering all IP addresses in its geographical region. Organizations that were allocated an IP prefix directly by an RIR can request the RIR to issue them an RC, validating their ownership of the IP prefix. For example, Deutsche Telekom in Figure 1 was certified by RIPE for its address space 91.0.0.0/10.

In case of that ownership was later delagated, that is, if an organization A further allocated a subprefix to organization B, then A is responsible for certifying B as the owner for that subprefix. Of course, to accomplish this, A must itself possess an RC for its assigned IP addresses.

As explained above, once an organization holds an RC, it can issue a ROA to protect its IP prefixes from hijacks.[1] ROAs specify an IP prefix, the number of the AS authorized to advertise that IP prefix in BGP, and the maximum-length of subprefixes the specified AS may advertise. Figure 1 shows that Deutsche Telekom used its RC to issue a ROA so as to protect its IP prefix against hijacking.

[1]Actually, the owner issues an end-entity (EE) certificate for an ephemeral one-time-use key, which is used to sign the ROA [38, Section 2.3]. The ROA and EE cert are stored together, so we have treated them as one object.

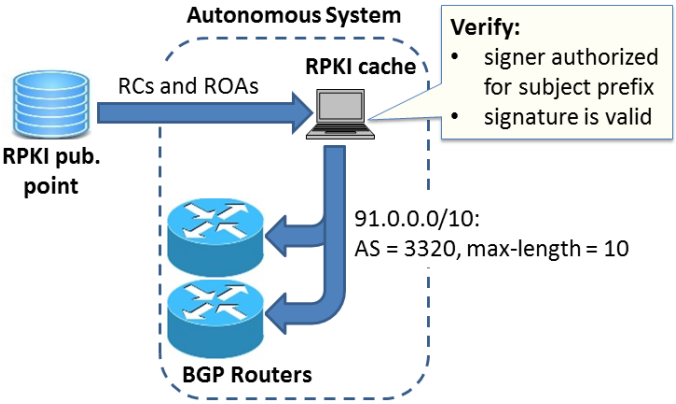

Fig. 2: A local cache syncs with RPKI publication points. BGP routers periodically pull the cache for whitelist updates.

B. Route-Origin Validation (ROV)

Organizations can enforce *route-origin validation (ROV)* [9], [45] to identify and discard BGP route-advertisements that violate ROAs. Namely, advertisements where the destination IP prefix is not mapped to the origin AS specified in the ROA. A local cache machine at the AS periodically syncs with the RPKI database to retrieve RCs and ROAs, validating them from the root of the RPKI hierarchy to its leaves. Valid ROAs mapping IP prefixes to ASes are then used to generate whitelists which BGP routers in that AS pull from the cache periodically using the RPKI-to-Router protocol [10]. See Figure 2. With cryptographic operations outsourced to the cache machine, routers can enforce ROV without changes to their hardware or BGP-message handling architecture (routers supported configuration of filters to BGP messages long before RPKI).

Upon receiving a BGP route-advertisement, the BGP router checks whether the advertised destination IP prefix p is "covered" by a ROA, that is, whether there exists a ROA for a superprefix $P \supseteq p$. The route-advertisement is then assigned one of following three states:

- **Unknown:** p is not covered by any ROA.
- **Valid:** p is covered by a ROA, the origin AS number matches AS number specified in the ROA and the IP prefix is no longer than the maximum length specified in the ROA (e.g., see Deutsche Telekom's advertisement in Figure 1).
- **Invalid:** Otherwise (p is covered, but not "valid").

Routers use this state assignment to realize route-filtering policies. The default action for most routers is to discard invalid routes (e.g., see [14], [35]) and this is also the considered best practice [21] (RFC7454 states that routers SHOULD discard invalid routes). However, an ROV-enforcing AS may instead choose to configure its router to merely de-prefer invalid routes over other routes. Indeed, our survey finds that some operators do this (see Section III). However, as observed in [17], [28], de-preferring invalid routes leaves the AS *completely* vulnerable to subprefix hijacking. We therefore focus our attention henceforth only on ROV adopters that discard invalid routes.

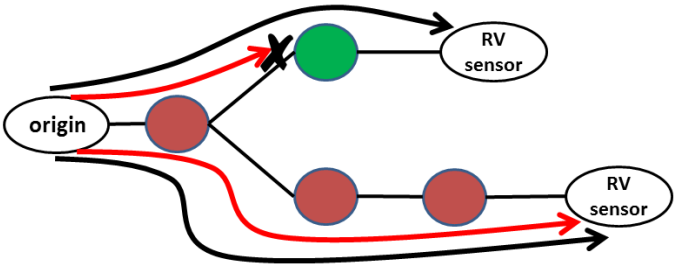

Fig. 3: Using invalid route advertisements (marked by red arrows) and valid/unknown route advertisements (marked by black arrows), to find ASes that do not enforce ROV (marked in red) and ROV-enforcing ASes (marked in green).

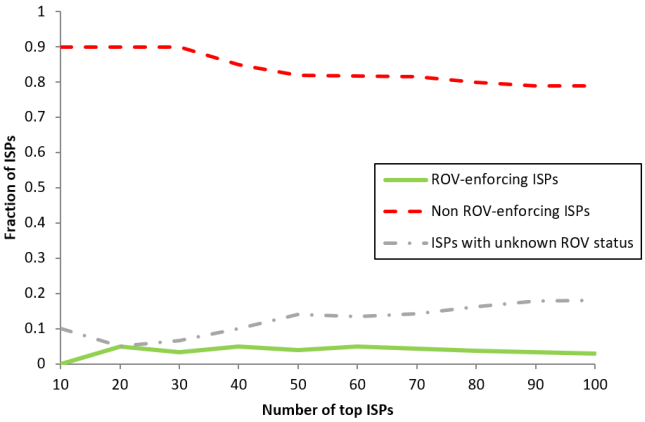

Fig. 4: ROV enforcement among the 100 largest ISPs

III. ROV Adoption Status

Route Origin Validation (ROV), defined in RFC 6483 [32], allows BGP routers to prevent prefix hijacking by detecting that an incoming BGP advertisement is inconsistent with ROAs in RPKI. Major vendors support ROV in their BGP routers with negligible computational overhead (e.g., see [14], [35]). Hence, deploying ROV involves only a modest, one-time installation effort, and no significant operational expenses. We present the first measurements of ROV adoption, showing that it is very limited, in particular at the core of the Internet. To complement and support our measurements, we also surveyed over 100 network practitioners about ROV adoption in their networks.

A. Quantifying ROV Adoption

While measurements regarding the issuing of ROAs exist, e.g., [47], [57], we are unaware of any previous measurement of ROV adoption. Indeed, measuring adoption of ROV by ASes seems challenging: how can we tell if a BGP router performs ROV or not? We identified a way to gain empirical insights regarding the extent to which ROV is adopted by *leveraging invalid BGP route-advertisements*, i.e., advertisements that are incompatible with ROAs (see Section II-B). We first explain our measurement techniques and then discuss our results.

Identifying ASes that do not enforce ROV. We examine BGP paths from the multiple vantage points afforded by 44 Route Views sensors [54] and identify ASes that propagate BGP-path advertisements that are classified as "invalid" according to RPKI (see red nodes in Figure 3).[2] We conclude that these ASes do not enforce ROV (at least not for all invalid IP prefixes). Of course, this measurement identifies *only non-adopting ASes*. When an AS does not appear in any of the received announcements, this might not be due to ROV but to other reasons, e.g., its BGP paths to the relevant IP prefixes did not propagate to the vantage points. Hence, our results set a *lower bound* on ROV *non-adoption*.

Identifying ASes that enforce ROV. To identify ASes that do enforce ROV, we apply the following methodology. We seek an AS that originates both a BGP-route advertisement that is *not* invalid (i.e., is classified as either valid or unknown by RPKI, see the black arrows in Figure 3) and an invalid BGP advertisement, like the origin AS in Figure 3. Intuitively, we then check which ASes discarded the invalid route-advertisement from that AS but relayed the other advertisement from the same AS. Specifically, we check whether there is only one transit AS on the non-invalid advertisement's route that did not relay any invalid route-advertisement (see green node in Figure 3).[3] We categorize an AS as *ROV-enforcing* if it meets the above criterion with respect to three different destination ASes.

Results. Our measurement techniques, as described above, provide a view on the state of ROV enforcement amongst the ASes at the core of the Internet (since these are likely to be on the paths observed by the Route Views sensors). Figure 4 shows the results for the top 100 ISPs (ordered by the number of the customer ASes). We find that almost all of the top 100 ISPs *do not enforce ROV*. Specifically, we found only *one* AS of the top 20 ISPs that enforces ROV, and at least 78 of the top 100 ISPs do not enforce ROV. In fact, we found only 3 of the top 100 ISPs enforce ROV. Note that 19 of the top 100 ISPs could not be classified by our method (as captured by the grey line in Figure 4).

B. Survey Results

To complement and corroborate the above results, we conducted an anonymous survey of over 100 network security practitioners. Our survey was conducted by sending requests over different mailing lists, including 'closed' lists (where membership is limited only to recognized practitioners).[4] Among participants 80% were network operators or managers and most others were security/networking consultants, Appendix A provides additional characterization of the participants.

We asked survey participants whether and how they apply RPKI-based ROV in their BGP routers. Figure 5 illustrates our

[2] We excluded advertisements from the same AS as the Route Views sensor, or from its customers, since these might not be subject to filtering.

[3] Since the BGP decision process at an AS is often identical across different destination IP prefixes in the same AS, this approach is likely to identify ASes that filtered the invalid advertisement and propagated the valid one.

[4] We believe that the results are biased, to some degree, in the direction of stronger security and adoption, since expert and security-aware operators are more likely to participate in such forums and to respond.

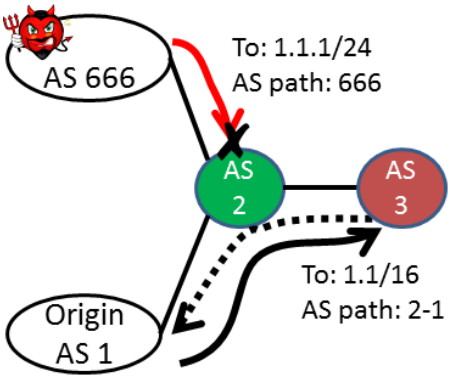

(a) Collateral benefit: AS 3 does not fall victim although it does not perform ROV

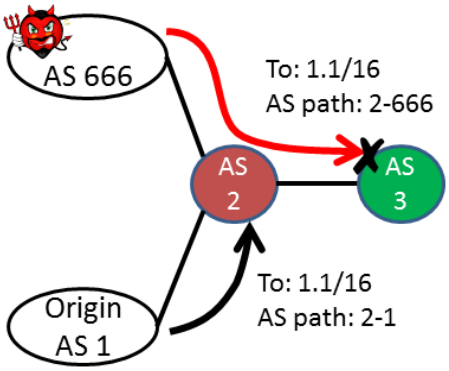

(b) Collateral damage (disconnection): AS 3 receives only an invalid route from (non-enforcing) AS 2, and hence has no valid path to AS 1

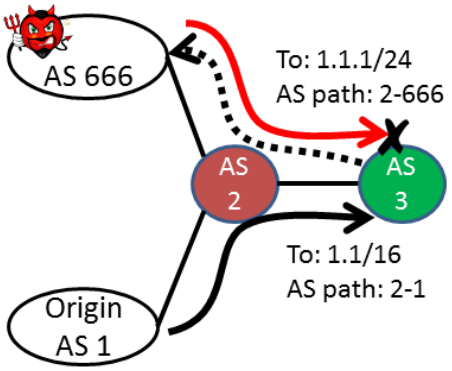

(c) Collateral damage (hijack): AS 3's traffic to AS 1 is forwarded via the attacking AS 666, although AS3 discarded the attacker's advertisement

Fig. 6: Collateral benefit and collateral damage in partial ROV adoption. Solid arrows represent BGP advertisements, dashed arrows represent data packet forwarding.

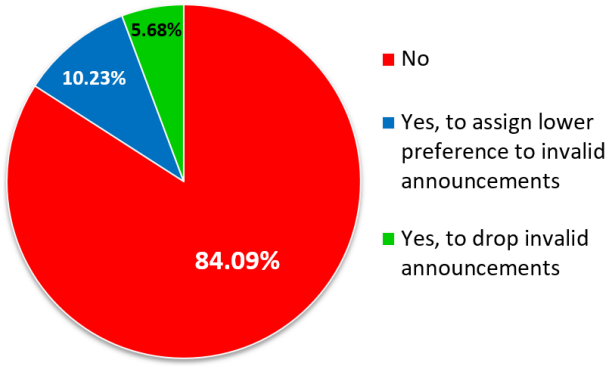

Fig. 5: Survey. Do you apply RPKI-based route origin validation (ROV)?

results. Over 84% of responses indicate not enforcing ROV at all. Less than 6% of the participants reported enforcing ROV. A larger fraction of the participants (over 10%) reported assigning lower preference to BGP route-advertisements whose ROV validation status is "invalid". As described in Section II, this adoption mode leaves adopters completely vulnerable to subprefix hijacks.

IV. Security in Partial Adoption

The measurement results in Section III invite the following question: *What is the impact of partial ROV deployment on security?*

A. Collateral Benefit and Damage

We identify collateral benefit and collateral damage effects that greatly impact the security of ROV in partial deployment. The scenarios in Figure 6 illustrate these effects, which we next describe.

Collateral benefit. ROV allows adopting ASes to protect the ASes "behind" them by blocking malicious route-advertisements even if these ASes do not perform ROV themselves. To see this, consider the simple network depicted in Figure 6a. Suppose that AS 1 is the legitimate owner of prefix 1.1/16, issued a ROA protecting this IP prefix, and advertises this IP prefix in BGP. Suppose also that AS 2 enforces ROV and all other ASes do not. Now, consider the scenario that an attacker, AS 666, announces the subprefix 1.1.1/24. AS 2 will discard the false route-advertisement from 666 and, in doing so, also protect AS 3. We refer to this as a *collateral benefit* from ROV enforcement.

Collateral damage. In contrast, an AS that does *not* apply ROV can cause *ROV-enforcers* to get disconnected from the victim or, surprisingly, have their traffic forwarded to the attacker. We describe such scenarios using the same network topology as before, only now we assume that AS 3 rather than AS 2 performs ROV.

A BGP-speaking router only advertises one route per prefix to neighboring ASes. Assume that the attacker performs prefix hijacking, i.e., advertises the same prefix as the victim. AS 3 that enforces ROV will automatically discard BGP route advertisements that are inconsistent with ROAs. If its provider (AS 2) advertises the attacker's invalid route advertisement, then AS 3 will disconnect from the legitimate origin (AS 1). See Figure 6b.

While AS 3 may be disconnected from legitimate destinations, it may seem that its data traffic to such IP prefixes will never reach illegitimate destinations (i.e., prefix hijackers) since it discards invalid routes. To see why this is incorrect consider the same network but when the attacker performs a subprefix hijack, as described in Figure 6c. AS 2, which does not apply ROV, will fall victim to the attack and select the direct BGP path to AS 666 for IP prefix 1.1.1/24 and the direct BGP path to AS 1 for the larger IP prefix 1.1/16. (Recall that, under BGP, routes to every destination IP prefix are computed independently.) Upon receiving the corresponding two BGP route-advertisements from AS 2, AS 3 will detect the attack and discard the route-advertisement for 1.1.1/24. Consequently,

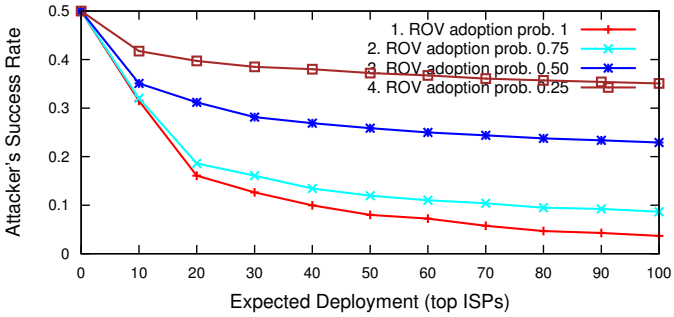

(a) Prefix hijack success rate

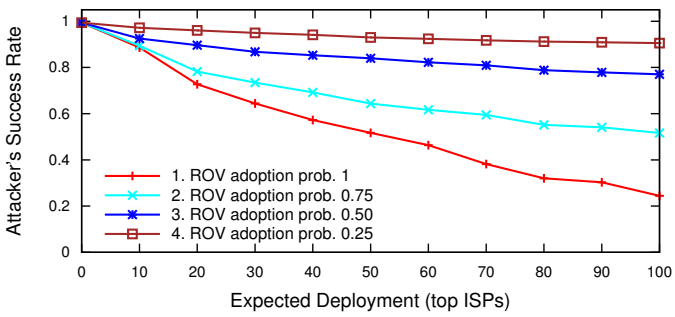

(b) Subprefix hijack success rate

Fig. 7: Collateral benefit: enforcing ROV only at top ISPs dramatically reduces attacker success rates

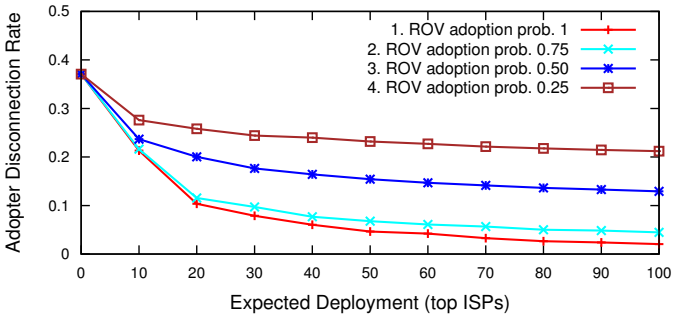

(a) Prefix hijack leads to adopter disconnection

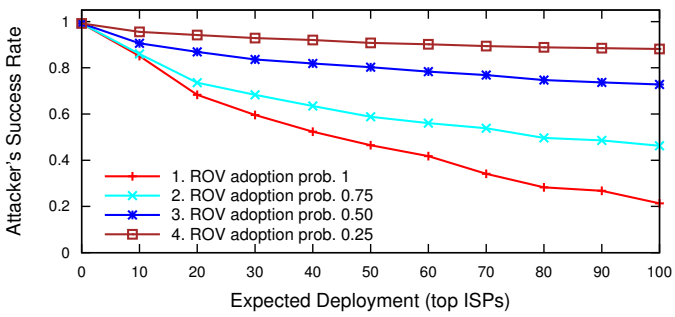

(b) Subprefix hijack steals traffic from adopters

Fig. 8: Collateral damage: without adoption at top ISPs, ASes are "doomed" even if they perform ROV

all packets from AS 1 to addresses in IP prefix 1.1.1/24 will be matched at AS 3 to the BGP path $(2, 1)$. Observe, however, that packets from AS 3 to IP addresses in 1.1.1/24 would be forwarded to AS 2 and then, after being matched to the BGP path for 1.1.1/24 at AS 2, forwarded to the attacker. This phenomenon results from an attack-induced inconsistency between AS 3's perceived path to IP addresses in 1.1/16 (in the control plane) and AS 2's actual forwarding path to IP addresses in 1.1.1/24 (in the data plane).

B. Security Evaluation

We next quantify, through extensive simulations on empirically-derived datasets, the impact of collateral benefit and damage, as described above, on global routing security.

Simulation framework. Our simulations apply the BGP route-computation method presented in [23], [24], [26] to the CAIDA AS-connectivity graph from July 2016. Our results average over 10^6 combinations of attacker and victim (i.e., the legitimate owner of IP prefix) ASes, both selected uniformly at random from the set of all ASes, as in [23], [26].

Impact of collateral benefit and damage. The results in Figure 7 quantify the extent to which partial enforcement of ROV at the core of the Internet can mitigate prefix and subprefix hijacks, for prefixes covered by ROAs. We consider different probabilities of adoption p, and different numbers of expected adopters x (between 0 to 100), chosen randomly from the set of $\frac{x}{p}$ top ISPs (i.e., the ISPs with the highest number of customer ASes). All other ASes do not enforce ROV. The results presented are averaged values after repeating the experiment 20 times for each possible number of adopters $x \in \{0, 10, ..., 100\}$.

We find that, on the one hand, when *all of the top 100 ISPs* enforce ROV, hijack success rates diminish rapidly (as a function of the number of deploying ISPs). Specifically, the success rate of prefix hijacks is below 10% when the 50 largest ISPs enforce ROV and below 5% when the 100 largest ISPs enforce ROV (see Figure 7a). For the subprefix hijack attack, the attacker's success rate falls from almost complete success to about 22% when the 100 top ISPs enforce ROV (see Figure 7b). On the other hand, we find that when adoption rate is low (as it is today), the attacker's success rates diminish much more slowly. For example, the subprefix hijacker's success rate remains as high as 90% even when 25% of the 400 top ISPs perform ROV (i.e., for $x = 100$ adopters). For prefix-hijacking, success rates are lower, but still high, e.g., about 40% even when 25% of the top 400 ISPs adopt.

Benefit from ROV adoption not at the core. We next quantify the benefit that an AS *not* at the core can derive from adopting ROV. Our results in Figure 8 show the attacker's success rate in harming (disconnect or hijack) ASes that *enforce ROV* for different rates of partial ROV adoption at the core of the Internet. Our results show that enforcing ROV does not benefit

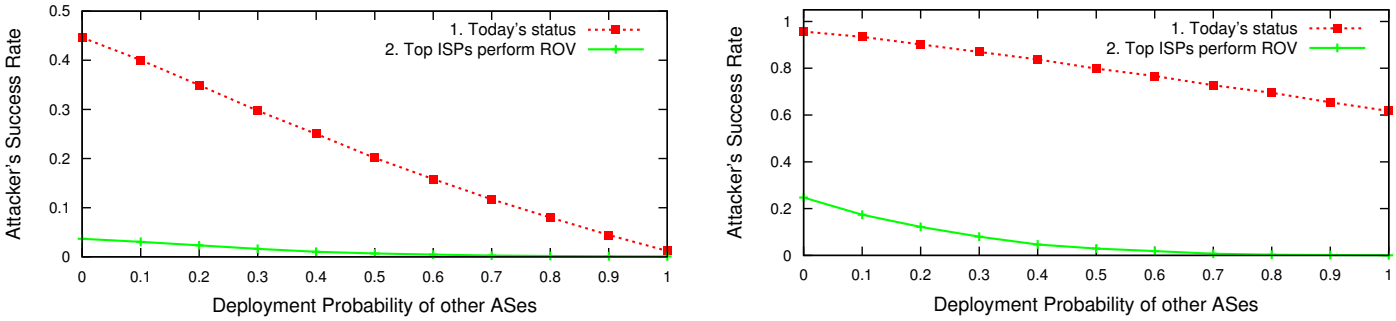

(a) Prefix hijack success rate

(b) Subprefix hijack success rate

Fig. 9: Comparing the benefit of RPKI under (1) today's ROV deployment, and (2) when the top ISPs perform ROV

such an AS much over the benefits gained by only deploying at the core, especially against subprefix-hijacking (observe the similarities between Figure 7a and 8a for prefix hijack, and Figure 7b and 8b for subprefix hijacks). We conclude that the benefit that an AS not at the core derives from ROV adoption at the core of the Internet is essentially *invariant* (on average) to whether or not that AS adopts ROV itself. Put differently, global routing security is primarily affected by the collateral benefit (and damage) from ROV enforcement at the core of the Internet.

Today's status. We explore the security benefits provided by *today's* ROV enforcement at the Internet's core. We compare two cases: (1) ROV's current deployment state, at the top 100 ISPs as presented in Section III-A, and (2) when all of the top 100 ISPs enforce ROV. Figure 9 contrasts these two scenarios (compare the green and red lines). We measure the attacker's success rate when performing prefix and subprefix hijacks for different ROV adoption probabilities of all other ASes on the Internet, i.e., ASes that are not in the 100 largest ISPs, and/or ASes that were classified in "unknown" ROV-enforcement status in Section III-A. We observe that if all top 100 ISPs enforce ROV, the security benefits of RPKI are dramatically improved. In contrast, under today's enforcement of ROV at the Internet's core, prefix hijack attacks are likely to succeed, and subprefix hijacks would remain very effective *even if all other ASes* on the Internet enforce ROV (about 60% success rate, see Figure 9b).

Conclusion. Put together, all our simulation results give rise to the following conclusion. Enforcement of ROV at the core of the Internet, by the top ISPs, appears to be both *necessary and sufficient* for gaining substantial benefits from RPKI. Consequently, until a large fraction of these ISPs enforce ROV, other ASes have limited incentive to issue ROAs for their prefixes (since they will remain largely unprotected).

V. Insecure ROAs

We expose below a new vulnerability of RPKI that renders about a third of IP prefixes covered by ROAs insecure from traffic hijacking attacks even if ROV is ubiquitously adopted.

We say that a ROA is *loose*, when not all subprefixes of the maximum length allowed by the ROA are advertised in BGP, e.g., if the max-length is set to /24, but only a /20 subprefix is advertised in BGP by the *legitimate* origin AS. We found that almost 30% of the IP prefixes covered by ROAs fall in this category. Indeed, our analysis reveals that network administrators often specify very relaxed restrictions on the maximum length. Specifically, over 85% of loose ROAs allowed maximum length of /24 (for IPv4 prefixes) or /48 (for IPv6 prefixes), the most specific prefix lengths typically accepted [21]. Our interactions with network administrators reveal confusion regarding the max-length field and the perception is that specifying high values provides flexibility in future changes to BGP advertisements, namely, enables advertising more specific IP prefixes without revisiting the ROA. We show, however, that IP prefixes covered by loose ROAs are *vulnerable* to the following attack vector, which allows traffic hijacking even if *all* ASes on the Internet perform ROV.

Attacking IP prefixes in loose ROAs. Suppose that a certain ROA specifies that an IP prefix p may be advertised in BGP only by the (legitimate) origin AS a and that AS a advertises prefix p in BGP. Suppose further that this ROA is loose according to the above definition. Then, the attacker may advertise the BGP announcement $666 - a - \bar{p}$, where 666 represents an AS controlled by the attacker and $\bar{p}$ is a permitted subprefix of p. Observe that since $\bar{p}$ is a subprefix of p, *all* traffic destined for IP addresses in $\bar{p}$ will be forwarded to the attacker due to IP's longest-prefix matching. Observe also that this attack cannot be detected by ROV, as the attacker's BGP advertisement specifies the legitimate origin, AS a. We point out that this attack is a hybrid of two well-studied attacks: subprefix hijacking, e.g., the attacker advertising the route $666-\bar{p}$, and the "next-hop attack", i.e., the attacker advertises the route $666 - a - p$. The former is detectable by ROV but guarantees that *all* traffic from victims flow to the attacker (as IP lookup prioritizes more specific prefixes), whereas the latter is undetectable by ROV but typically much less effective [16] since the attacker's success is a function of its location, the BGP routing policies of others, etc. Through the exploitation of loose ROAs *and* BGP path-manipulation, the attacker can enjoy the best of both worlds: attract all traffic to the victim and going unnoticed by ROV.

Real-world example. Swisscom, a large Swiss ISP, issued a ROA for the prefix 81.62.0.0/15, with origin AS 3303. Swiss-

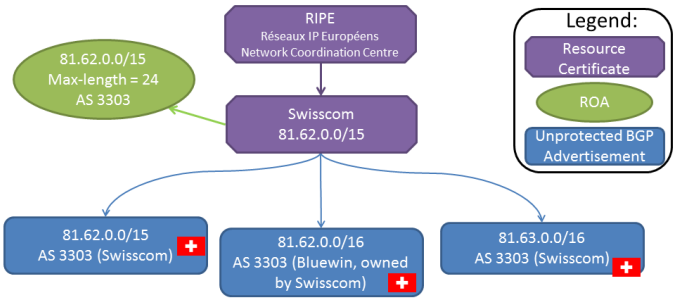

Fig. 10: Swisscom issued a ROA authorizing up to /24 prefixes, yet only announces two /16 subprefixes

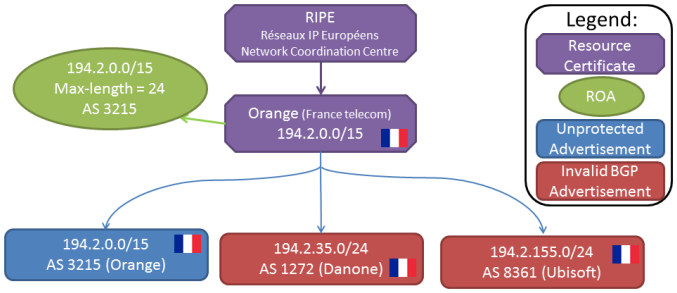

Fig. 11: Orange (formally France Telecom) issued ROA invalidating customer announcements

com also advertises this prefix and its two /16 subprefixes in BGP. See illustration in Figure 10. However, Swisscom's ROA specifies a max-length of 24, exposing it to the above attack. E.g., in order to hijack *all* traffic to the prefix 81.63.0.0/16, the attacker can announce through BGP its two /17 subprefixes 81.63.0.0/17 and 81.63.128.0/17 with AS-path 666-3303. To mitigate the attack, Swisscom should simply change the max-length in their ROA to 16.

VI. Obstacles to Deployment

We investigate the root causes for RPKI and ROV's low adoption. Our measurements in Section III revealed that only few of the top ISPs enforce ROV, i.e., filter routes according to the ROAs in RPKI. Thus, not surprisingly, RPKI suffers from a circular dependency between issuing ROAs and adopting ROV: ASes do not gain sufficient benefits from adopting ROV so long as so many IP prefixes are unprotected by ROAs, yet there is little incentive to issue a ROA so long as so few ASes enforce ROV. We turn our attention to two other main factors contributing to non-deployment: unreliability of RPKI due to badly issued ROAs, and inter-organization dependencies. Unlike past studies of RPKI [19], [33], [34], our analysis relies on a breakdown of ROAs by organizations, thus yielding important new insights into RPKI and ROV's deployability.

A. Mistakes in ROAs

Operators should specify in ROAs the IP prefixes advertised by their organizations through 3 fields: the AS number, IP prefix, and maximum length (see Section II) [39]. Unfortunately, mistakes commonly occur when issuing ROAs [34]. To illustrate the dire implications of such mistakes consider the

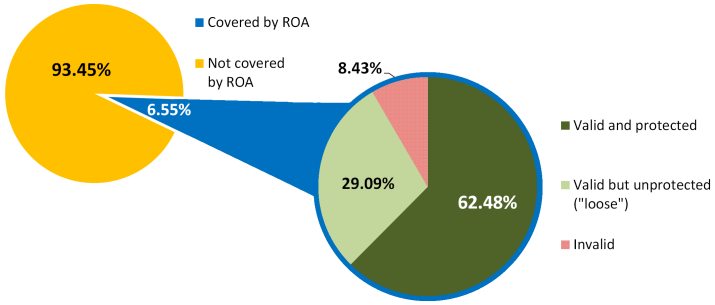

Fig. 12: BGP advertisement status according to RPKI

real-world example in Figure 11, identified by our empirical analyses of ROAs and BGP advertisements. Orange (previously France Telecom), a large French ISP, issued a ROA for IP prefix 194.2.0.0/15 specifying the origin AS number as 3215 (its own AS number). Orange has several customer ASes that advertise in BGP subprefixes of 194.2.0.0/15 under different AS numbers. E.g., Danone advertises the (sub)prefix 194.2.35.0/24 from its own AS, numbered 1272. Most of Orange's customers, including Danone and Ubisoft shown in the figure, did not issue a ROA. Consequently, the BGP advertisement of these customers are viewed as invalid according to RPKI (since they appear like subprefix hijacks, see discussion on the decision process in Section II-B). Thus, any AS enforcing ROV will discard all BGP route-advertisements to these destination IP prefixes. Not only that, Orange's ROA permits /24 advertisements, yet Orange and its customers only advertise 23 of the 512 /24 subprefixes in Orange's address space, rendering the rest vulnerable to our attack from Section V. Other mistakes in ROAs include inadvertently specifying the wrong AS number and advertising in BGP IP prefixes that are more specific (longer) than the maximum length specified in the ROA.

Badly issued ROAs render the BGP route-advertisements of the covered IP prefixes "invalid" to ROV-enforcing ASes. Thus, an ROV-enforcer may not be able to reach the IP prefixes of the issuers of such "bad ROAs" or even the IP prefixes of *other* organizations (as illustrated in the Orange example above)[5]. Figure 12 presents the results of our measurement of the status of ROAs available in the RPKI database (in July 2016), and the corresponding BGP advertisements observed in Route Views (6.5% of IP prefixes advertised in BGP are covered by ROAs). Beyond invalid IP prefixes, as described above, the figure also shows the large fraction (almost 30%) of *valid but unprotected* IP prefixes (due to loose ROAs, as discussed in Section V). As evidenced by Figure 12, a significant number of IP prefixes covered by ROAs are invalid (over 8%), matching the results in [34]. A few of these may be actual prefix or subprefix hijacking attacks, yet the majority is likely a result of configuration errors.

Our survey of network operators validates that being disconnected from legitimate destinations due to badly issued

[5]Note that 38% of the invalid announcements are for a subprefix of some valid announcement by the same origin AS. In this case, traffic from the ROV-enforcer to addresses in the subprefix will reach the origin AS, although the ROV-enforcer discarded the invalid subprefix advertisement. However, ROV-enforcer will not publish the invalid subprefix announcement and hence may 'lose' traffic to this subprefix from multi-home customers.

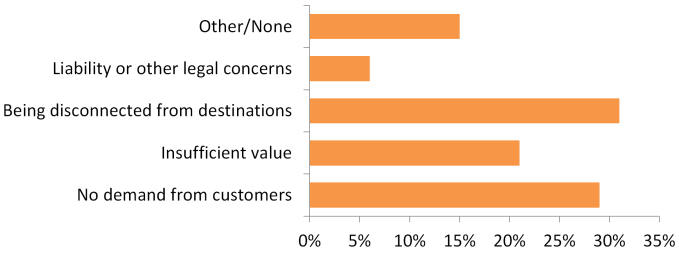

Fig. 13: Survey. What are your main concerns regarding executing RPKI-based origin authentication in your network?

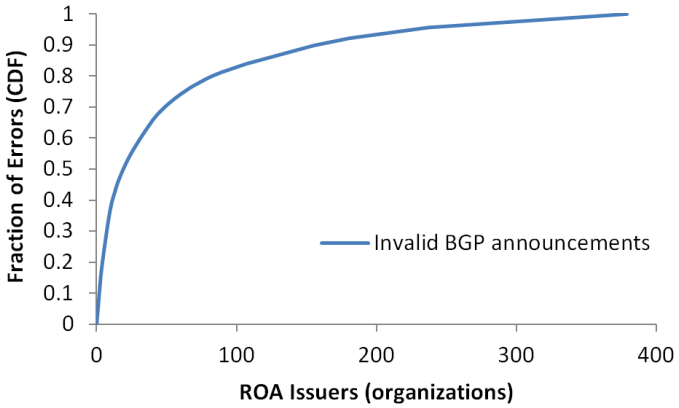

Fig. 14: Organizations responsible for bad ROAs

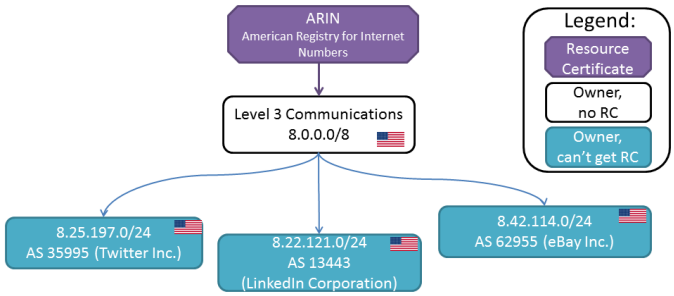

Fig. 15: Level 3 Communications did not obtain an RC. As a result, all its customers that own subprefixes cannot obtain an RC.

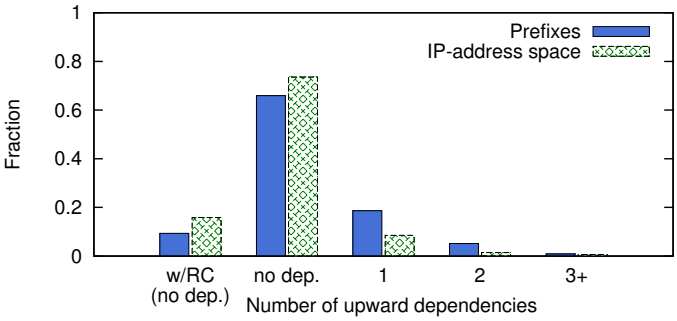

Fig. 16: Upward dependency histogram

ROAs is a major concern. Indeed, *fear from being disconnected from other ASes* (due to invalid ROAs) was the *most common* reason specified for not performing ROV (over 30%), even more common than *no demand from customers* and *insufficient value*; see Figure 13.

We aim to gain deeper insights into the human error factor in RPKI deployment in practice. To this end, we associate badly issued ROAs and the corresponding BGP advertisements with *organizations*. (This turns out to be nontrivial; we explain how it is done in Appendix B.)

We attribute each IP prefix rendered invalid by ROV to the organization responsible for that error. Our analysis reveals some encouraging news. Although hundreds of organizations are responsible for invalid IP prefixes, most of the errors are caused by a small number of organizations, as shown in Figure 14. Specifically, only 20 organizations are responsible for issuing the ROAs causing 50% of the errors. We discuss the significance of this finding, in terms of driving RPKI deployment, in Section VII.

B. Inter-Organization Dependencies

Another potential reason for the disappointingly small number of ROAs in RPKI are *inter-organization dependencies* that arise when to issue an RC and/or ROA for a prefix an organization must wait for other organizations to take action first. Two such dependencies exist: *upward* and *downward*.

The upward dependency is for issuing RCs: the RC for prefix p must be signed by an entity already in possession of a valid RC for a super-prefix containing p. Figure 15 illustrates the hierarchical structure of RPKI and the inter-organizational upward dependencies that it entails. Level 3 Communications is one of the largest ISPs worldwide. It was allocated its IP addresses directly from ARIN, the RIR for North-America, but did not issue an RC. Consequently, over 500 other organizations are unable to obtain RCs and issue ROAs to protect their thousands of prefixes. These organizations include, for example, Twitter, LinkedIn, and eBay, as illustrated in Figure 15.

The downward dependency is for issuing ROAs: if a prefix p originates from some AS a, but p has a subprefix $\bar{p}$ that originates from a *different* AS b (belonging to a different organization), then publishing a ROA specifying a as the origin AS of p, *before* a ROA specifying b as the origin AS of $\bar{p}$, would invalidate BGP advertisements of $\bar{p}$, as in the example of Orange and its customers in Figure 11. (See discussion in Section II on issuing RCs and ROAs.)

We next investigate the extent to which upward and downward dependencies explain the slow adoption of RPKI. Our results show that while there are not many inter-organizational dependencies, downward dependencies pose obstacles to RPKI deployment for some of the largest ISPs worldwide.

Quantifying upward dependencies. We utilize our mapping from RCs and ROAs to quantify upward dependencies. Figure 16 depicts the upward dependencies. See Appendix B for a detailed exposition of our measurement methodology. We find that about 20% of the prefixes are upward-dependent on another organization. In fact, when looking at the portion of the IP address space (the green bars in Figure 16), we find that these prefixes hold less than 10% of the whole IP

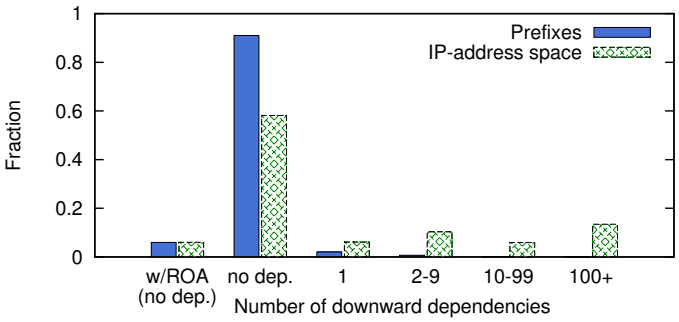

Fig. 17: Downward dependency histogram

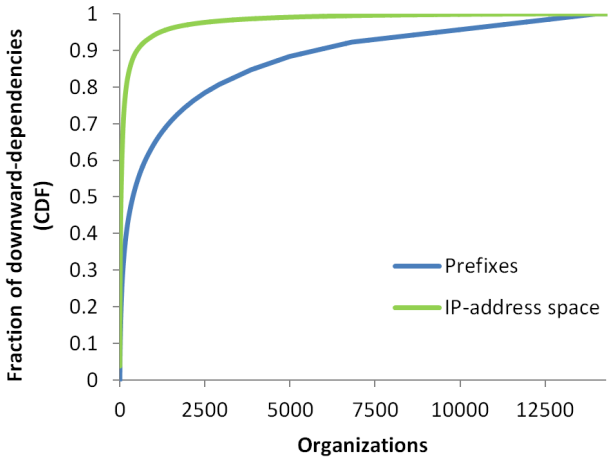

Fig. 18: Organizations with downward dependencies (CDF)

address space, probably since organizations with large prefixes typically received their prefixes directly from the RIRs and so have no upward dependencies. Similar measurements focusing only on the prefixes advertised in BGP exhibit the same trends. We also find that the number of prefixes with no upward dependency, i.e., prefixes that may be covered by a certificate but their owners refrain from doing so, is very high as shown in the blue 'no dep.' bar in Figure 16.

Quantifying downward dependencies. As explained above, downward dependencies arise when to issue a ROA for its IP prefix an organization must wait for other organizations advertising a subprefix of that prefix to issue ROAs first. We next quantify downward-dependencies (see detailed description of our measurement methodology in Appendix B). Figure 17 describes our results: about 90% of the IP prefixes not protected by ROAs are not downward-dependent.

Unfortunately, the relatively few prefixes that are downward-dependent constitute a large portion of the IP-address space (compare blue and green bars for those with more than 100 dependencies in Figure 17) and belong to some of the largest ISPs worldwide (Orange and Level 3 in Figures 11 and 15 are but a few examples). Figure 18 illustrates this problem, observe the sharp curve showing that few organizations have most of the downward dependencies. This is no surprise as the vast majority of ASes are smaller organizations (that do not sell/lease IP address blocks to other organizations), whereas large ISPs own large IP address blocks and suballocate IP addresses to customers. We present a solution to downward dependencies in Section VII.

VII. Driving RPKI Forward

We present below several concrete proposals for addressing the concerns pertaining to RPKI and ROV deployment discussed in the previous sections. We first briefly present three high-level directions and then dive into the details.

Targeting the large ISPs. Our results in Section IV establish that ROV adoption by the largest ISPs can lead to significant improvements in global security. Thus, concentrating efforts on the top ISPs can yield significant incentives for issuing ROAs, hopefully breaking the vicious circular dependency between the non-issuing of ROAs and the non-enforcement of ROV.

Improving RPKI's reliability and integrity. As shown in Section VI-A, relatively few organizations are responsible for

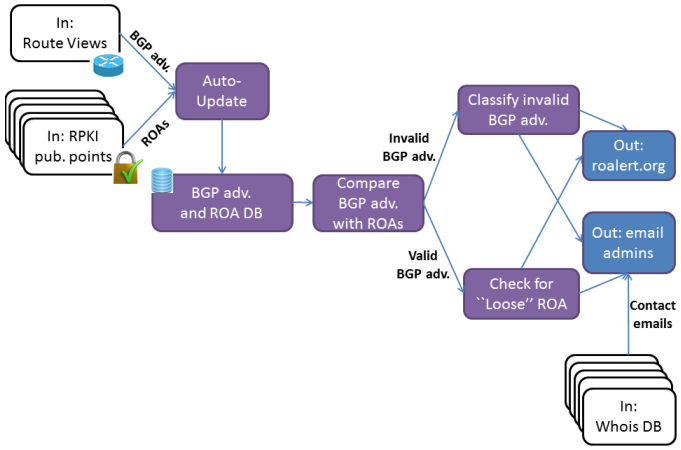

Fig. 19: ROAlert system design

the majority of badly issued ROAs. Thus, the number of BGP advertisements incorrectly categorized as invalid due to human error can be significantly decreased via focused effort. To alert network operators about badly issued ROAs and inform them about their *organization-specific* obstacles to RPKI deployment, we present *ROAlert*. We argue that ROAlert, if adopted by the relevant entities (e.g., the RIRs), could aid in building trust in RPKI and boosting its adoption.

Eliminating downward dependencies with wildcard ROAs. We present a new kind of ROAs, compatible with today's format, which eliminates downward dependencies. As explained in Section VI-B, this form of inter-organization dependencies pertains to a significant portion of IP addresses and to some of the world's largest ISPs.

A. ROAlert

As shown in Figure 12, one of the biggest obstacles to performing ROV is the significant number of bad ROAs. To mitigate this problem we developed *ROAlert*, an automated system that detects bad ROAs and alerts the corresponding network administrators (through emails and web interface).

IP-prefix	Organization	AS number	Problem
91.0.0.0/10	Deutsche Telekom (Figure 1)	3320	None (valid and protected)
81.62.0.0/15	Swisscom (Fig. 10)	3303	Valid but unprotected (has "loose ROA")
194.2.0.0/15	Orange (Fig. 11)	3215	Valid but unprotected (has "loose ROA")
194.2.35.0/24	Danone (Fig. 11)	1272	Invalid (Orange's down dep.)
194.2.155.0/24	Ubisoft (Fig. 11)	8361	Invalid (Orange's down dep.)
8.0.0.0/8	Level 3 (Fig. 15)	3356, 3549	No RC, prevents RC from customers
8.25.197.0/24	Twitter (Fig. 15)	35995	Cannot get RC (dependent on Level 3)
8.22.121.0/24	LinkedIn (Fig. 15)	13443	Cannot get RC (dependent on Level 3)
8.42.114.0/24	eBay (Fig. 15)	62955	Cannot get RC (dependent on Level 3)

TABLE I: Examples of networks and issues raised in the paper

Online, proactive notification. ROAlert periodically retrieves ROAs from RPKI's publication points and then compares them against the BGP advertisements retrieved from Route Views using CAIDA's BGPStream architecture [5] (every 5 minutes) so as to identify and classify bad ROAs (both invalidating BGP advertisements and "loose"). ROAlert uses the RIRs' Whois databases to alert the network administrator (at most once a month for the same violation). Figure 19 illustrates ROAlert's design.

Web interface, ROAlert.org. ROAlert also provides a web interface that allows network operators to check whether their network is (correctly) protected by ROAs. In the event that the network is not properly protected, ROAlert explains the causes for this situation (be it loose ROAs, inter-organization dependencies, etc.). We encourage the reader to enter the ROAlert website to find out whether her/his IP address is protected by RPKI and also to learn more details about the networks mentioned in this paper (Table I) and other networks.

Differences from existing systems. ROAlert is not the first system that compares BGP advertisements and ROAs to detect errors. In particular, RIPE allows network administrators issuing a ROA to check how that ROA would effect the prefixes advertised through BGP [20], [46]. However, in contrast to existing systems, ROAlert is *not an opt-in service*. This is particularly important when an IP prefix of one organization becomes invalid due to a ROA issued by another organization (like in the example in Figure 11). In this scenario, the victim may not even be aware of that ROA issuance or may not fully comprehend the intricacies of RPKI. Furthermore, ROAlert is *constantly running* (whereas [20], [46], for instance, is only active when registering a new ROA). Hence, if BGP announcements change, and "good" ROAs become "bad", alerts are issued in real time. ROAlert also identifies loose ROAs, which are not detected by other systems (since such ROAs are valid according to RPKI). Lastly, ROAlert notifies both offenders and victims with concrete suggestions as to how to fix a given issue.

Importantly, beyond error-detection, ROAlert also serves another important purpose: identifying organization-specific challenges to deployment. To accomplish this, ROAlert relies on careful data analysis of multiple datasets (e.g., to map RCs, ROAs, and IP prefixes to organizations, as explained in Appendix B). Through such data analyses ROAlert can, for instance, inform a network operator interested in issuing a ROA which organizations its network is upward- or downward-dependent on.

Feedback from network operators. We have received feedback and engaged in discussions with 52 of the network administrators who received our alerts. Of these responders, 40 reported fixing the problem. Five administrators who received notification that another organization had caused their prefix to become invalid reported that they had subsequently engaged with the offenders. The rest of the responses mostly explained problematic scenarios that they have difficulty addressing, such as BGP advertisements sent to peering networks, which were supposed to only be used by these internally, but were in fact "leaked" onwards. Two responses indicated that the prefix was, in fact, hijacked (probably due to configuration errors at the offending AS).

Measurements. We have sent notifications to a total of 552 victims and offenders through ROAlert over the course of six months, yet only 168 emails did not "bounce" (see discussion of limitations below). We checked whether one month after notification the problem was fixed by re-examining the ROAs and BGP route-advertisements via ROAlert. The results of these follow-up checks reveal that ROAlert notifications, when reaching the network operator's mailbox, are effective. Over 42% of "bad ROA" alerts resulted in fixing the problematic ROA one month later, this is contrasted with about 15% of ROAs fixed involving operators that ROAlert could not contact. ROAlert also notified administrators of "loose ROAs", i.e., when their prefix is valid but in fact unprotected and had 19% success rate. We conclude that a proactive, non-opt-in alert system can help mitigate bad ROAs.

Limitations. ROAlert relies on information from the Whois databases in order to contact victims and offenders. However, whois entries are often outdated, resulting in unreachable email addresses (see measurements above). Consequently, only a fairly small fraction of victims and violators were alerted. We believe, based on our experience with ROAlert and the above results, that ROAlert can aid in greatly reducing the number of bad ROAs, thus increasing the level of trust in RPKI. We hope that these results will motivate the RIRs, which have more reliable contact points to prefix owners, to adopt ROAlert.

Selective, ROAlert-informed ROV. While ROAlert and similar mechanisms could hopefully contribute to eliminating many badly issued ROAs, surely some bad-ROAs will remain. To address this concern we propose that ASes enforce ROV *selectively*, i.e., enforce ROV for ROAs that have been "well-behaved" for a pre-defined amount of time. ROAlert allows this to be automated (by identifying new ROAs, and marking them "well-behaved" if they do not conflict with BGP advertisements for sufficient period).

Of course, this is not immune to pitfalls. What if an organization changes its set of IP prefixes advertised in BGP in a manner that is incompatible with a previously issued ROA by the same organization? This might still lead to the organization being disconnected by those doing selective ROV. We point

out, however, that using a platform like ROAlert implies that the organization would be quickly notified when its prefix becomes invalid.

We argue that the use of selective ROV could mitigate the justified concerns about discarding BGP advertisements due to mistakes in ROAs, while protecting issuers of "good" ROAs from prefix and subprefix hijacking.

B. Wildcard ROAs

Downward dependencies, discussed in Section VI-B, pose a challenge to deployment of RPKI by some large organizations, e.g., Orange in Figure 11. Consider a large IP prefix p, with origin AS number a, owned by organization Foo. Suppose that several subprefixes $p_1, \ldots p_n$ are contained within p, but belong to other organizations and in particular, have different origin ASes, $a_1, \ldots a_n$. Suppose further that these organizations have not yet issued a ROA for these prefixes.

Foo may go ahead and issue a ROA for p, specifying origin a. That was, for instance, the approach taken by Orange in Figure 11. However, the result is that BGP advertisements for $p_1, \ldots p_n$ (that belong to Danone and Ubisoft in our example), would conflict with the ROA for p and hence be discarded by ROV-enforcing ASes.

Alternatively, Foo may choose to break p into subprefixes that are entirely owned by Foo and only issue ROAs for these prefixes and advertise them in BGP (instead of p in its entirety). However, this may result in significant increase to BGP routing tables. We evaluate this possibility by testing the increase ratio in BGP table size. We used the Route Views dataset [54] to build the Routing Information Base (RIB) for each of the 44 different sensors (BGP routers located in different and geographically dispersed ASes). We then measured, for each sensor, the increase in the number of RIB entries. We find that on average, the increase is over 26% of prefixes in the routing tables, or approximately 160K prefixes (the standard deviation was approximately 2%).

We now present a simple alternative solution, called *wildcard ROA*, which allows organizations to issue ROAs and protect their IP prefixes before their descendants issue ROAs for themselves. Wildcard ROAs are intended to allow an AS holding an RC that covers an address block containing a specific IP prefix p_i not belonging to that AS to specify that *any* AS may be the origin of that prefix (p_i). Thus, any AS with an RC could issue a ROA for itself without having to wait for its descendants to do the same (or giving up on BGP aggregation) simply by issuing another ROA, using a "wildcard AS number", that includes all subprefixes that are no longer in its possession and thus may be announced using other ASNs. Importantly, however, wildcard ROAs must be assigned a lower priority than "regular" ROAs for the same prefix, so that if a specific IP prefix is included in both a regular ROA (issued by the legitimate owner) and (possibly many) wildcard ROAs, the regular ROA prevails, thus protecting the legitimate owner from hijacking and eliminating the downward-dependency problem.

VIII. Related Work

There is a long history of attacks against BGP, which motivated many past studies on analyzing BGP's vulnerabilities and on designing defenses. See an overview of the security issues in [11], [29] and a study of common misconfigurations in [43]. For an example of one of the earliest designs, suggesting many of the basic ideas underlying both RPKI and BGPsec, see [36]. We focus below on the studies most relevant to *deployment*.

Several studies, beginning with [13], explore deployment of path-security mechanisms, mainly BGPsec. Gill et al. [23] study how to encourage adoption of BGPsec by creating appropriate incentives for ISPs to deploy it. Other studies [26] explore the security guarantees of BGPsec, including under partial deployment [41]. Many other studies, e.g., [50], present alternative path-security mechanisms which are easier to deploy. See a survey of the long and winding road to BGP security in [25].

Adoption of RPKI was also studied. In [37], a provider, SURFnet, studied the false-positive rate it would incur if it performed ROV. Several studies [19], [33], [34] explore the issuing of ROAs and measure the number of prefixes that would be discarded with ROV. A study of ROAs covering content provider prefixes was presented in [57]. These studies, however, did not address the issues we focus on, such as quantifying the adoption of ROV and its benefits under partial deployment, loose ROAs, upward and downward inter-organizational dependencies, designing an online alert system, and more.

Wahlisch et al. present preliminary results of a survey of network operators regarding deployment of RPKI and DNSsec [56]. Indeed, many of the challenges of deploying and enforcing DNSsec and RPKI are similar; see [40] for a study comparing deployment of DNSsec with deployment of RPKI, and [30] for discussion on the deployment of DNSsec. Other studies [17], [28] investigate another problematic aspect of RPKI: disproportionate power of centralized authorities to unilaterally revoke any IP prefixes under their control.

Network operators sometimes rely on repositories such as IRR/RADB to filter suspicious BGP route-advertisements. Like RPKI, these repositories bind IP prefixes to origin ASes and can thus be used to realize filtering policies, including additional filtering beyond route-origin validation (see [27]). However, these repositories are not fully reliable/secure and contain numerous errors, as highlighted and measured in [31]. IRR/RADB repositories are typically used, in practice, only to filter incorrect BGP advertisements from customer ASes [21], [44] and thus provide quite limited security benefits.

IX. Conclusion

Extensive efforts have been invested, for many years now, to improve the security and reliability of BGP routing. These efforts range from a patchwork of ad-hoc mechanisms (monitoring services, IRR/RADB registries, etc.) to carefully designed standards such as RPKI and BGPsec. Recent studies have shown that deployment of BGPsec faces significant, possibly insurmountable challenges. In contrast, RPKI deployment, crucial both for combating IP prefix and subprefix hijacking and as a prerequisite to defenses against BGP path-manipulations (incl. BGPsec), is regarded as more feasible. Indeed, monitoring of RPKI certificates and ROAs indicates progress, although at a hardly satisfactory rate.

Our results showed that not only does RPKI suffer from very low deployment, but even Route Origin Validation (ROV), which uses RPKI to defend against prefix hijacking, is rarely enforced. Since ROV is the first and only use of RPKI so far, its deployment and RPKI's deployment suffer from the classic chicken and egg problem. We investigated the implications of ROV's scarce adoption to global routing security, showing that almost ubiquitous deployment of ROV by the largest ISPs (e.g., top 100) is both necessary and sufficient to protect the Internet from prefix and subprefix hijacking. We showed, however, that unless many RPKI records are fixed, even global ROV adoption will not protect many networks from devastating traffic hijacking attacks.

We embarked on a systematic study of the root-causes for RPKI's slow adoption, including *negative* incentives for deployment due to loss of traffic as a result of badly issued ROAs, and inter-organization dependencies. On the positive side, our results show that many of the banes plaguing RPKI and ROV deployment can be vastly improved via highly focused efforts. Specifically, focusing deployment efforts on the large ISPs, modest changes to RPKI, and adoption of a system such as ROAlert by entities such as the RIRs, can go a long way towards mitigating obstacles to adoption. We thus advocate concentrating standardization, regulatory, and operational efforts in these directions.

ACKNOWLEDGEMENTS

This work was supported by ISF grants 420/12 and 1354/11, Israel Ministry of Science grants 1-11807, 3-9772 and 3-10884, the Israeli Center for Research Excellence in Algorithms, NSF grant 1414119, and an ERC Starting Grant. We thank Steve Bellovin, Randy Bush, Sharon Goldberg, Joel Halpern, Ethan Heilman, Tomas Hlavacek, Hezi Moriel, Hank Nussbacher, Alvaro Retana, and Nickolai Zeldovich for their helpful comments and suggestions. Special thanks to Daniel Davidovitch for helping us create ROAlet's web interface, to Matthias Waehlisch and his research group at Freie University Berlin for helping us utilize Miro [48], and to Christian Teuschel from RIPE for helping us utilize the RIPEStat database.

REFERENCES

[1] "The New Threat: Targeted Internet Traffic Misdirection," http://www.renesys.com/2013/11/mitm-internet-hijacking/.

[2] "AFRINIC Whois Database," https://www.afrinic.net/services/whois-query, Jul. 2016.

[3] "APNIC Whois Database," https://wq.apnic.net/apnic-bin/whois.pl, Jul. 2016.

[4] "ARIN Whois Database," https://whois.arin.net/ui, Jul. 2016.

[5] "BGPStream by CAIDA," https://bgpstream.caida.org/, Jul. 2016.

[6] "LACNIC Whois Database," http://lacnic.net/cgi-bin/lacnic/whois, Jul. 2016.

[7] "RIPE NCC Whois Database," https://apps.db.ripe.net/search/query.html, Jul. 2016.

[8] Andree Toonk, "BGP Hijack Incident by Syrian Telecommunications Establishment," BGPMon blog, 2015.

[9] R. Bush, "Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI)," RFC 7115 (Best Current Practice), Internet Engineering Task Force, Jan. 2014. [Online]. Available: http://www.ietf.org/rfc/rfc7115.txt

[10] R. Bush and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol," RFC 6810 (Proposed Standard), Internet Engineering Task Force, Jan. 2013. [Online]. Available: http://www.ietf.org/rfc/rfc6810.txt

[11] K. R. B. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, 2010.

[12] Center for Applied Internet Data Analysis, "Routeviews Prefix to AS mappings Dataset for IPv4 and IPv6," https://www.caida.org/data/routing/routeviews-prefix2as.xml, May 2016.

[13] H. Chan, D. Dash, A. Perrig, and H. Zhang, "Modeling Adoptability of Secure BGP Protocols," in *SIGMETRICS*, R. A. Marie, P. B. Key, and E. Smirni, Eds. ACM, 2006, pp. 389–390. [Online]. Available: http://dblp.uni-trier.de/db/conf/sigmetrics/sigmetrics2006.html#ChanDPZ06

[14] "BGP Origin AS Validation," Cisco IOS guide, Cisco, 2013.

[15] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira, "One Hop for RPKI, One Giant Leap for BGP Security," in *HotNets*, J. de Oliveira, J. Smith, K. J. Argyraki, and P. Levis, Eds. ACM, 2015, pp. 10:1–10:7. [Online]. Available: http://dl.acm.org/citation.cfm?id=2834050

[16] ——, "Jumpstarting BGP Security with Path-End Validation," in *SIGCOMM*, M. P. Barcellos, J. Crowcroft, A. Vahdat, and S. Katti, Eds. ACM, 2016, pp. 342–355.

[17] D. Cooper, E. Heilman, K. Brogle, L. Reyzin, and S. Goldberg, "On the Risk of Misbehaving RPKI Authorities," in *HotNets*. ACM, 2013, p. 16. [Online]. Available: http://dl.acm.org/citation.cfm?id=2535771

[18] S. Crocker, "DNSSEC Deployment Threats - What is Real? What is FUD?" October 2008, DNSSEC and IPv6 Workshop.

[19] R. de Boer and J. de Koning, "BGP Origin Validation (RPKI)," Univeristy of Amsterdam, systems and network engineering group, Tech. Rep., July 2013.

[20] J. de Koning, "RPKI - A Dashboard for BGP Operators," Online in https://labs.ripe.net/Members/javy_de_koning/rpki-dashboard, July 2013.

[21] J. Durand, I. Pepelnjak, and G. Doering, "BGP Operations and Security," RFC 7454 (Best Current Practice), Internet Engineering Task Force, Feb. 2015. [Online]. Available: http://www.ietf.org/rfc/rfc7454.txt

[22] L. Eggert, "DNSSEC Deployment Trends," http://eggert.org/meter/dnssec.

[23] P. Gill, M. Schapira, and S. Goldberg, "Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security," in *SIGCOMM*, S. Keshav, J. Liebeherr, J. W. Byers, and J. C. Mogul, Eds. ACM, 2011, pp. 14–25.

[24] ——, "Modeling on Quicksand: Dealing with the Scarcity of Ground Truth in Interdomain Routing Data," *Computer Communication Review*, vol. 42, no. 1, pp. 40–46, 2012.

[25] S. Goldberg, "Why is it Taking so Long to Secure Internet Routing?" *Commun. ACM*, vol. 57, no. 10, pp. 56–63, 2014. [Online]. Available: http://doi.acm.org/10.1145/2659899

[26] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, "How Secure are Secure Interdomain Routing Protocols?" *Computer Networks*, vol. 70, pp. 260–287, 2014. [Online]. Available: http://dx.doi.org/10.1016/j.comnet.2014.05.007

[27] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. D. McDaniel, and A. D. Rubin, "Working around BGP: An Incremental Approach to Improving Security and Accuracy in Interdomain Routing," in *NDSS*. The Internet Society, 2003. [Online]. Available: http://www.isoc.org/isoc/conferences/ndss/03/proceedings/papers/5.pdf

[28] E. Heilman, D. Cooper, L. Reyzin, and S. Goldberg, "From the Consent of the Routed: Improving the Transparency of the RPKI," in *SIGCOMM*, F. E. Bustamante, Y. C. Hu, A. Krishnamurthy, and S. Ratnasamy, Eds. ACM, 2014, pp. 51–62. [Online]. Available: http://dl.acm.org/citation.cfm?id=2619239

[29] A. Herzberg, M. Hollick, and A. Perrig, "Secure Routing for Future Communication Networks (Dagstuhl Seminar 15102)," *Dagstuhl Reports*, vol. 5, no. 3, pp. 28–40, 2015. [Online]. Available: http://drops.dagstuhl.de/opus/volltexte/2015/5267

[30] A. Herzberg and H. Shulman, "Retrofitting Security into Network Protocols: The Case of DNSSEC," *Internet Computing, IEEE*, vol. 18, no. 1, pp. 66–71, 2014.

[31] T. Hlavacek, "Routing Policies and Real Paths in the Internet," Master's thesis, Charles University, Prague, 8 2016.

[32] G. Huston and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)," RFC 6483 (Informational), Internet Engineering Task Force, Feb. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6483.txt

[33] D. Iamartino, "Study and Measurements of the RPKI Deployment," 2015.

[34] D. Iamartino, C. Pelsser, and R. Bush, "Measuring BGP Route Origin Registration and Validation," in *PAM*, ser. Lecture Notes in Computer Science, J. Mirkovic and Y. Liu, Eds., vol. 8995. Springer, 2015, pp. 28–40. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-15509-8

[35] "Configuring Origin Validation for BGP," Juniper TechLibrary, Juniper Networks, 2013.

[36] S. T. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 4, pp. 582–592, 2000.

[37] J. Kloots, "RPKI Routing Policy Decision-Making: A SURFnet Perspective," https://labs.ripe.net/Members/jac_kloots/, January 2014.

[38] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," RFC 6480 (Informational), Internet Engineering Task Force, Feb. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6480.txt

[39] M. Lepinski, S. Kent, and D. Kong, "A Profile for Route Origin Authorizations (ROAs)," RFC 6482 (Proposed Standard), Internet Engineering Task Force, Feb. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6482.txt

[40] W. Lian, E. Rescorla, H. Shacham, and S. Savage, "Measuring the Practical Impact of DNSSEC Deployment," in *Proceedings of USENIX Security*, 2013.

[41] R. Lychev, S. Goldberg, and M. Schapira, "BGP Security in Partial Deployment: Is the Juice worth the Squeeze?" in *SIGCOMM*. ACM, 2013, pp. 171–182. [Online]. Available: http://dl.acm.org/citation.cfm?id=2486001

[42] E. M. Lepinski, "BGPsec Protocol Specification," RFC, Internet Engineering Task Force, Oct. 2014. [Online]. Available: https://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-10

[43] R. Mahajan, D. Wetherall, and T. E. Anderson, "Understanding BGP Misconfiguration," in *SIGCOMM*. ACM, 2002, pp. 3–16. [Online]. Available: http://doi.acm.org/10.1145/633025.633027

[44] D. McPherson, S. Amante, E. Osterweil, L. Blunk, and D. Mitchell, "Considerations for Internet Routing Registries (IRRs) and Routing Policy Configuration," RFC 7682 (Informational), Internet Engineering Task Force, Dec. 2015. [Online]. Available: http://www.ietf.org/rfc/rfc7682.txt

[45] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, "BGP Prefix Origin Validation," RFC 6811 (Proposed Standard), Internet Engineering Task Force, Jan. 2013. [Online]. Available: http://www.ietf.org/rfc/rfc6811.txt

[46] R. NCC, "RIPE NCC Tools and Resources," Online in https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources.

[47] NIST, "NIST RPKI Monitor," http://rpki-monitor.antd.nist.gov/, 2016.

[48] A. Reuter, M. Wählisch, and T. C. Schmidt, "RPKI MIRO: Monitoring and Inspection of RPKI Objects," in *SIGCOMM*, S. Uhlig, O. Maennel, B. Karp, and J. Padhye, Eds. ACM, 2015, pp. 107–108, http://rpki-miro.realmv6.org/. [Online]. Available: http://dl.acm.org/citation.cfm?id=2785956

[49] RIPE Network Coordination Centre, "RIPE NCC Statistics," https://www.ripe.net/analyse/statistics, 2016.

[50] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz, "Listen and Whisper: Security Mechanisms for BGP," in *NSDI*. USENIX, 2004, pp. 127–140. [Online]. Available: http://www.usenix.org/events/nsdi04/tech/subramanianListen.html

[51] A. Toonk, "Hijack Event Today by Indosat," http://www.bgpmon.net/hijack-event-today-by-indosat/.

[52] ——, "Turkey Hijacking IP Addresses for Popular Global DNS Providers," BGPMon.

[53] ——, "BGP Routing Incidents in 2014, Malicious or Not?" http://www.bgpmon.net/bgp-routing-incidents-in-2014-malicious-or-not/, 2015, BGPMon.

[54] University of Oregon, "Route Views Project," http://bgplay.routeviews.org/, May 2016.

[55] P.-A. Vervier, O. Thonnard, and M. Dacier, "Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks," in *NDSS*. The Internet Society, 2015. [Online]. Available: http://www.internetsociety.org/events/ndss-symposium-2015

[56] M. Wählisch, O. Maennel, C. Perta, T. C. Schmidt, G. Tyson, and S. Uhlig, "Preliminary Results of Survey about RPKI/DNSSEC," in *proceedings of the 92nd IETF*, March 2015.

[57] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, "RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem," in *HotNets*, J. de Oliveira, J. Smith, K. J. Argyraki, and P. Levis, Eds. ACM, 2015, pp. 11:1–11:7. [Online]. Available: http://dl.acm.org/citation.cfm?id=2834050

[58] R. White, "Deployment Considerations for Secure Origin BGP (soBGP)." June 2003. [Online]. Available: http://tools.ietf.org/html/draft-white-sobgp-bgp-deployment-01

Appendix A
Additional Survey Questions

We present the rest of the questions we asked survey participants to characterize their networks in Figures 20 through 24.

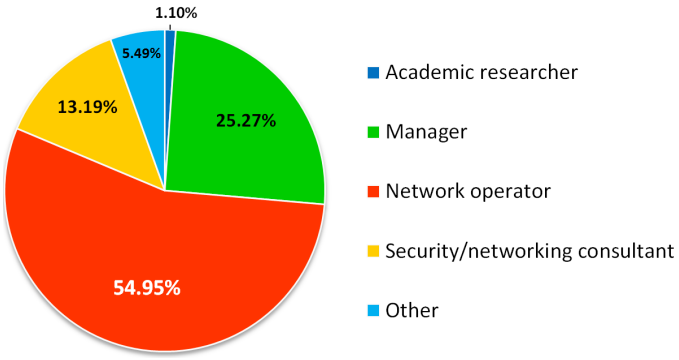

Fig. 20: Survey participant distribution

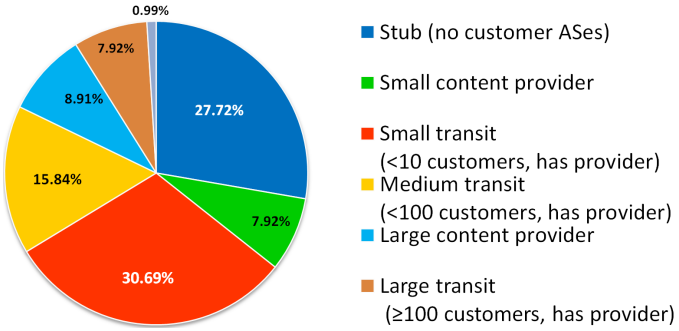

Fig. 21: Survey. Which of the following best describes your network?

Appendix B
Modeling Inter-Organizational Hierarchy

We used the IP prefix allocation database [49] and constructed the tree formed by IP-prefix assignments to organizations, we refer to it as the *prefix allocation tree*. We

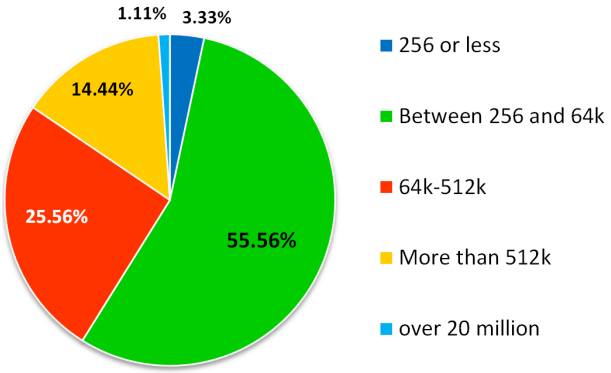

Fig. 22: Survey. How many IP addresses does your network include?

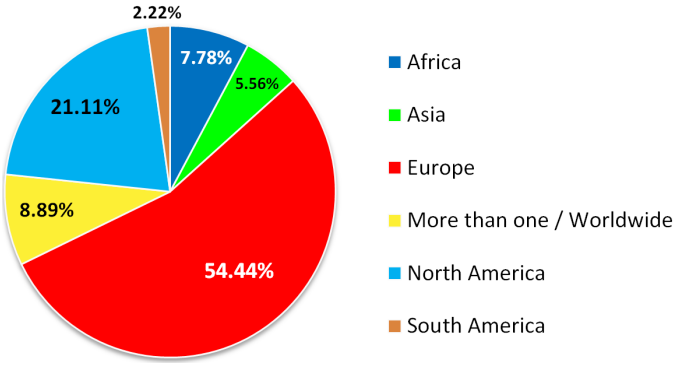

Fig. 23: Survey. Where is your network?

obtain the RCs and ROAs using Miro, a tool developed by Reuter et al. [48]. The datasets are from July 2016. From each RIR we root an *RPKI-object tree* that illustrates the RC and ROA hierarchy (i.e., which RC was used to issue other RCs or ROAs). Each vertex in both trees identifies the owner organization as we next describe.

A. How to map prefixes, RCs, and ROAs to their owners?

To understand RPKI's deployment status we must be able map IP prefixes, ROAs, and RCs, to the organizations that "own" them. This turns out to be a nontrivial feat.

We begin by finding the owner of the prefixes in the allocation tree. We use Route Views [54] to find prefixes advertised through BGP, and identify their owner by the AS number that advertised the prefix (using CAIDA's ASN to organization mapping [12]). For all other prefixes, which are owned by organizations yet unannounced through BGP, we use the Whois databases of the five RIRs [2], [3], [4], [6], [7] to identify the owner. We further use the "contact e-mail address" domain, commonly included in Whois entries, to match IP-prefixes owned by the same organization. In case we could not identify the owner of the prefix, we assume that the parent organization still owns that IP space and remove that prefix from the IP allocation tree. We were able to identify the owners of 92.3% of the prefixes, whereas almost all the prefixes we could not resolve to an owner are leaves in the prefix allocation

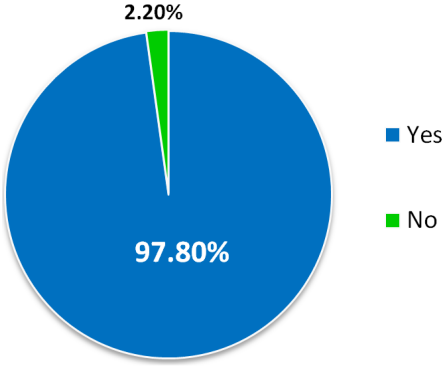

Fig. 24: Survey. Does your organization have an AS number?

tree. In particular, we identified the owners of all prefixes that are announced in BGP (since those are associated with an AS number). We thus believe that our measurements provide a good approximation as to the extent of inter-organizational dependencies.

Next, we identify the owners of ROAs. We use the AS numbers specified in ROAs to indicate which ASes are allowed to announce the protected IP prefixes, and the mapping from ASes to organizations in [12], to identify the owner organizations. Finally, we find the owners of RCs. The RPKI specification [38] does not identify the owner in the RC, to get around this we use a combination of our prefix and ROA mappings. If the RC was used to issue a ROA, we assign the owner of the ROA to that RC. Otherwise, we use the owner of the largest prefix specified in this RC.

B. Measuring upward dependencies

For each prefix in the allocation tree, we find the first "ancestor IP prefix" that is covered by an RC. Importantly, the owner of this prefix is the one who must first issue an RC to its descendant (and the owner of that descendant to its descendant, until reaching the original prefix). We then find the number of intermediate IP prefixes in the allocation tree, i.e., prefixes that must be covered by RCs before the original prefix may be covered by RC. Finally, we count the distinct organizations on the path between the prefix and its ancestor prefix, excluding the organization owning the original prefix and the organization owning the ancestor prefix. We thus quantify the number of organizations that lie between an organization that wishes to issue a ROA and the organization that must initiate the chain of RC issuing.

C. Measuring downward dependencies

For each prefix in the allocation tree, we find the number of prefixes below it (i.e., subprefixes) that are announced in BGP and are not protected by ROAs. These "child prefixes" will be considered invalid by the routers that perform ROV if the parent is issued a ROA. We count the number of distinct organizations that own them.