Perfect is the Enemy of Good: Setting Realistic Goals for BGP Security

Yossi Gilad
Boston University and MIT

Tomas Hlavacek
Fraunhofer SIT

Amir Herzberg
University of Connecticut and Bar Ilan University

Michael Schapira
Hebrew University of Jerusalem

Haya Shulman
Fraunhofer SIT

"Give them the third best to go on with; the second best comes too late, the best never comes." (Sir Robert Alexander Watson-Watt)

1 INTRODUCTION

Interdomain routing security not in sight. Arguably, the most glaring security vulnerability of today's Internet infrastructure is the insecurity of the Border Gateway Protocol (BGP), which establishes routes between the Autonomous Systems (ASes) that make up the Internet. BGP's susceptibility to configuration errors and attacks is the cause for major Internet outages [28, 32, 33] and traffic hijacking [1, 4, 34]. A crucial step in today's agenda for securing BGP is deploying the Resource Public Key Infrastructure (RPKI), which binds IP address blocks to "owner" ASes via cryptographic signatures. RPKI enables ASes to perform *Route-Origin Validation (ROV)*, i.e., discard BGP advertisements originating in an attacker (or a misconfigured router) that advertises an IP address block that belongs to another ("IP-prefix-hijacks"). RPKI also underlies long-standing [37] and recent [9, 22] proposals for protecting BGP against "path-manipulation attacks" [7].

Despite its significance, RPKI's deployment is discouragingly slow [12, 18, 25, 27] and, worse yet, even the fairly few adopters are not necessarily secure [12]. Prior studies (see [15] and references within) revealed two fundamental reasons for this dismal state of affairs.

(1) Human error. Roughly a third of the records in RPKI repositories are erroneous. A wrong RPKI record might not defend the issuer [15], but the implications can be far worse. Consider an AS that performs ROV. When this AS receives erroneous RPKI records, it may incorrectly drop BGP announcements, disconnecting from thousands of legitimate destinations [12, 19]. This makes RPKI arguably more dangerous than the vulnerability it aims to address. RPKI, in its present state, thus effectively violates the fundamental principle of "do no harm."

(2) The chicken and egg problem. Certifying ownership of IP address blocks in RPKI requires nontrivial effort by the AS that wishes to be certified and potentially also other organizations. This creates a dependency cycle: certification is effective only when ROV is deployed, and ROV is effective only with respect to certified IP address blocks. Consequently, almost no one performs ROV [12, 18, 27], and RPKI certification is very limited [25].

The adoptable imperfect is better than the non-adoptable perfect. We posit that both problems are derived from the same root cause: setting the bar too high in terms of security requirements, at the expense of a nontrivial and error-prone certification process that hinders real-world adoption.

The Pareto Principle (or "80:20 rule") suggests that often most of the desired outcomes are achievable with significantly less effort than needed to realize all one desires, with increasing effort resulting in diminishing returns. We view BGP security as a perfect example. Consider the following two extreme scenarios of IP prefix hijacking: (1) the hijacking AS is remote from the legitimate origin, *vs.* (2) the hijacking AS is the sole ISP of the legitimate origin AS. The 'perfect' defense would prevent both attack scenarios. However, as discussed below, defending against the first (remote hijacker) scenario is significantly *easier and safer to deploy*.

Not only is protecting from the latter attack scenario much more difficult, but the attack itself is of less concern; the sole ISP of an AS can intercept *all* the traffic destined to the AS even without prefix hijacking. Indeed, all high-profile attacks on BGP, to date, belonged to the first category. Today's approaches to routing security, which strive to defend against both attacks, aspire for perfect security and compromise on deployability, as evidenced by the low adoption of RPKI and the non-deployment of BGPsec after many years of effort.

Introducing DISCO. We present DISCO, a Decentralized Infrastructure for Securing and Certifying Origins. DISCO *automatically* certifies ownership of IP addresses and generates filtering rules for discarding bogus BGP route-advertisements (IP prefix hijacks).

HotNets-XVII, November 15–16, 2018, Redmond, WA, USA

ACM ISBN 978-1-4503-6120-0/18/11…$15.00
https://doi.org/10.1145/3286062.3286071

DISCO is carefully engineered to meet the following four fundamental design goals: (1) protect against the most common and alarming attacks against BGP (like [1, 4, 34]), (2) be *fully-automated* to avoid the tedious bureaucracy traditional RPKI certification entails, (3) eliminate the need for coordination with others, and (4) enforce *consistency* between certificates and the BGP control plane, so as to avoid BGP incidents resulting from human neglect in syncing the two.

We argue that meeting goals 1-4 above *concurrently* requires rethinking the standard desideratum of binding IP address blocks to their *legal* owners, which induces a manual, error-prone certification mechanism. Instead, we settle for the lesser goal of certifying *de facto* possession of IP addresses. Specifically, *DISCO verifies that the AS being certified is the destination to which traffic destined for the relevant IP addresses is forwarded de facto over a considerable amount of time*. We show that this is sufficient for achieving strong security in practice (meeting goal 1) while rendering adoption significantly easier and safer (by meeting goals 2-4).

Realizing DISCO requires no changes to BGP routers. We built an open-source prototype implementation of DISCO that can be used to certify IP addresses. We evaluate DISCO and analyze its security guarantees through experiments on the Internet and simulations on empirically-derived datasets.

DISCO is a *first* step. We regard DISCO as a feasible approach for protecting ASes from prefix hijacking. Importantly, however, DISCO can also serve as the basis for recently-introduced countermeasures to BGP path-manipulation attacks, namely, path-end validation [8, 9]. Path-end validation was proposed as an alternative to BGPsec that achieves comparable security benefits while circumventing the potentially insurmountable obstacles facing its adoption. We view the combination of DISCO and path-end validation as a promising and realistic agenda for attaining the long-awaited inter-domain routing security.

2 JUST ENOUGH RPKI BACKGROUND

RPKI [21] certifies the association of network resources (such as IP prefixes) with public keys. Certified owners can advertise records that specify AS numbers authorized to advertise the owner's IP prefixes in BGP. This allows other ASes to filter routes with invalid origins, thus foiling prefix hijacks.

RPKI attests to the ownership of IP prefixes through *resource certificates*, which contain the owner's public key. Under RPKI's *hierarchic* certification process, the owner of an IP prefix can sign a resource certificate that delegates ownership for a subprefix under its control to another organization. A certified owner can use its private key to issue a *route-origin authorization (ROA)*, which lists ASes authorized to originate routes for the owner's IP prefix. The ROA may also record the maximum length (maxLength) of subprefixes the specified AS may advertise. As discussed in [15], erroneous/outdated specifications of maxLength can leave the AS vulnerable to attacks or render its legitimate advertisements invalid.

RPKI's deployment is sluggish. The vast majority of prefixes advertised in BGP is still not protected by ROAs [25] (including most prefixes used by popular web-services [36]), and almost no AS performs ROV [12, 17, 27], leaving the Internet largely exposed to traffic hijacking.

Certification is not easy. RPKI certification is manual and hierarchical. Network operators first need to request their providers (or other entity that allocated the IP address space to them) to issue them a resource certificate. Since many organizations do not yet have resource certificates for their own IP address blocks, in many cases this requires providers to first be certified themselves [12].

Consider the following example, illustrating the deployment challenge arising from hierarchic certification. Level 3 was allocated the prefix 8.0.0.0/8 from ARIN (American Registry for Internet Numbers). Level 3 later allocated subprefixes of its /8 to hundreds of organizations but did not obtain a resource certificate from ARIN. With today's RPKI, hundreds of organizations must wait for Level 3 to certify authority over 8.0.0.0/8 before they can certify their prefixes.

Even if an organization has a resource certificate, it is sometimes difficult for that organization to issue a ROA without generating problems. Consider an organization that held a large prefix P, and transferred a subprefix $p \subseteq P$ to a customer. If the organization issues a ROA covering P, then the customer's announcement for p would be invalid since it would appear similar to a subprefix hijack. This creates a problematic dependency for some of the largest providers [12].

Human error is common. About 10% of the prefixes advertised in BGP that are also covered by ROAs are invalid [25], with the vast majority attributed to mistakes [12, 19, 35]. ASes that use RPKI to filter invalid routes unwittingly discard legitimate BGP-advertisements, and thus disconnect from legitimate destinations.

Another prevalent type of mistakes is issuing ROAs with too flexible maxLength [15]. This can void RPKI's benefits and leave the issuer completely vulnerable to devastating subprefix hijacks. Recent measurements reveal that almost 30% of the prefixes covered by RPKI are in fact unprotected [12, 13, 15].

Human error impacts ROV enforcement. Recent studies explored ROV enforcement, showing that it is very limited at best [12, 18, 27]. According to a recent survey of network operators, the most common reason for not filtering invalid routes is fear of "being disconnected from destinations" due to erroneous RPKI records. Two other popular reasons are "insufficient value" and "no demand from customers", leading to the circular dependency problem in RPKI deployment (described earlier).

3 DESIGN AND ARCHITECTURE

3.1 Goal: Certifying *de facto* Ownership

We posit that to facilitate easy and safe adoption of a certification mechanism, the classical desideratum of binding IP

address blocks to their *legal* owners *must* be relaxed. DISCO certifies *de facto* possession of an IP address block, i.e., it verifies that the AS being certified is, de facto, the destination to which traffic to the relevant IP addresses is forwarded.

We show that verifying de facto ownership enables *automated* and *flat* certification, and thus significantly lowers the bar for adoption and the risk of human error. In addition, DISCO's automated certification process ensures that an AS's records are kept consistent with its BGP advertisements.

We show that this weaker notion of ownership, when coupled with additional precautions, facilitates an easy and safe to deploy security mechanism against prefix hijacking.

DISCO's certification process guarantees the following:

Any AS can certify its IP addresses in DISCO so long as it does not suffer loss of connectivity during the certification process (in which case it can re-try later).

An attacker capable of fooling DISCO's certification either (1) has no point in doing so, i.e., can already *passively* intercept *all* traffic to the victim under *normal conditions* (e.g., is its sole ISP), **or (2) must launch a highly visible, long-running attack**, exposing itself to the victim and leaving public traces in the DISCO system.

3.2 High-Level Overview

In DISCO, a software-implemented *agent* installed within the AS initiates a certification process for address blocks belonging to that AS. Border routers connect to the agent through iBGP. No changes to router software or network infrastructure are needed, only router configuration changes. The agent is responsible for ensuring that the owner AS's public key is attached to its BGP prefix announcements.

DISCO also includes *registrars*, which continuously monitor BGP advertisements from multiple, well-dispersed Internet vantage points. A registrar approves the ownership of an IP address block if during a "*certification (time) interval*" BGP advertisements for that address block reach the registrar with the same public key. An aggregate of more than a threshold of registrar approvals for the same public key and address block is a DISCO-certificate: an association of that prefix with the public key (the DISCO equivalent of traditional RPKI resource certificates).

Once certified, the agent issues a ROA associating its AS number with its IP addresses. Similarly to RPKI, DISCO uses repositories to store and distribute certificates and ROAs.

We next dive into the mechanics of DISCO's certification and origin validation.

3.3 Certification with DISCO

We next describe how DISCO's main software components, the agent and registrars, execute the certification mechanism depicted in Figure 1.

Agent. The agent stores the AS's public/private key pair needed to certify ownership over IP prefixes. BGP routers within the AS connect to the agent using iBGP sessions, which allows the agent to take over local prefix announcements from

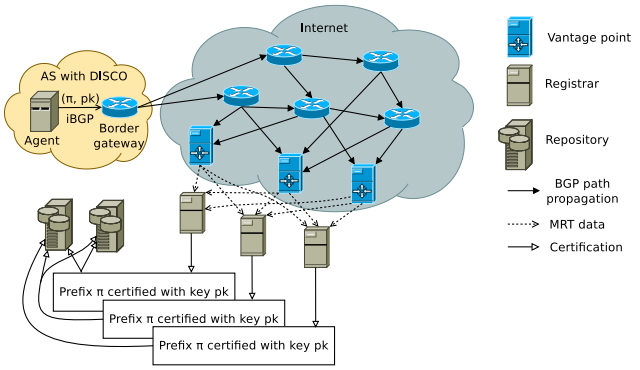

Figure 1. DISCO certification process. The agent associates public key pk with prefix π.

the border routers and become the origin of the announcement within the AS. Using iBGP makes deployment easy and avoids a need for any changes in network infrastructure both inside and outside of the AS that deploys DISCO. Decoupling prefix origin generation from the border routers conforms with common operational practices and BCPs such as [11].

To piggyback a public key over prefix announcements in a BGP-compatible way, DISCO utilizes BGP's optional transitive attributes. A border router must relay optional transitive attributes even if it does not recognize the meaning of the attribute [26, §4.3]. We validate in §5 that this indeed holds (to a large extent) in today's Internet. DISCO defines an optional transitive BGP attribute called *ownership indicator*. To be certified by DISCO, the agent includes the ownership indicator in the BGP advertisements. The ownership indicator specifies the public key for the certifying organization.

Registrars. Registrars continuously monitor BGP announcements from multiple Internet vantage points. If the announced prefix and ownership indicator pair is observed by more than a threshold number of the vantage points and is maintained for the course of a *certification time interval*, the registrar approves the association of an address-block to the specified public key. The associated address-block might be smaller than the actual prefix that the agent announced if some other organization announces a subprefix of it. So, certified address-blocks in DISCO are described by one large prefix, followed by a list of excluded subprefixes (announced by others).

The certification interval must be sufficiently long to ensure that attackers that perform prefix hijacks to certify the hijacked addresses are detected before ownership can be certified. When a prefix-owner detects a hijack, they can announce the hijacked-prefix themselves and prevent attacker-certification (since some registrars will observe the owner's announcement, see §5). In practice, most hijacks last less than a day [34], and so we posit that one week period strikes a good balance between responsiveness and security.

Each registrar has its own private/public key pair, which it uses to cryptographically sign the association of address-blocks with public keys. Registrars send the associations they approve to a repository server (with replicas). If more than a threshold fraction of the registrars approve the same address

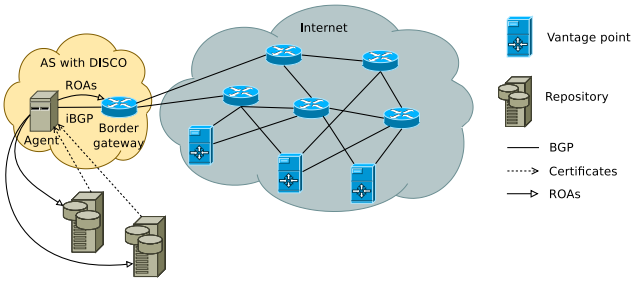

Figure 2. Origin validation with DISCO. The agent issues a ROA, and fetches ROAs issued by others.

block, then the prefix-to-public-key assignment is considered certified. (DISCO's repositories are not trusted; they only store and forward the registrars' signed approvals and cannot fake them.) In §5 we use Internet measurements to inform the selection of the certification threshold.

DISCO's design decouples vantage points, which monitor BGP routes, from the registrars, which make certification decisions. The decision of where to deploy the monitors, and which monitors to trust, is of the organization running a registrar. This decoupling allows DISCO to use existing BGP monitors as readily-available vantage points (see §4).

Certification with holes. In our earlier example with Level 3 (§2), all subprefixes that Level 3 had allocated to other organizations would be excluded from Level 3's DISCO certificate since, de facto, they belong to other organizations. By excluding subprefixes that belong to other organizations, DISCO avoids complicated and potentially risky chains of trust as in RPKI's hierarchy [10].

Continuous certification. Since certification with DISCO is automated, certificates can be relatively short-lived (e.g., twice the certification interval) and automatically renewed. Continuous certification allows DISCO to avoid potentially complicated revocation procedures.

Bootstrapping DISCO. To securely communicate with the registrars (e.g., in the event that the IP addresses of the registrars themselves are hijacked), DISCO agents are configured with the registrars' public keys. Since the number of registrars is fairly small and we do not anticipate frequent changes to the set of registrars, a static list suffices. This resembles TLS, where a static set of root-CAs is responsible for issuing certificates. Unlike TLS, however, trust in DISCO is distributed across several organizations and a minority of corrupt registrars cannot produce a false certificate.

3.4 Origin Authentication with DISCO

The agent authorizes and validates origins, as described in Figure 2.

Issuing DISCO ROAs. Once the AS is certified as the owner of an address block, it can use its private key to issue a ROA for the IP space covered by its certificate. This procedure is automatically initiated by the agent when it observes that a resource certificate was registered through the DISCO repository. Much like RPKI ROAs, a DISCO ROA specifies a list of approved IP prefixes with maxLength and an authorized origin for every prefix. However, like the resource certificate, DISCO ROAs also specify a list of excluded sub-prefixes for every prefix in the ROA (conforming with the "wildcard-ROA" format previously suggested to address some of RPKI's deployment problems [12]). This is important, since in many cases an RPKI ROA issued for one organization may invalidate legitimate announcements by other organizations. Consider the following real-world example. Orange was certified by RIPE for IP prefix 194.2.0.0/15 and issued a ROA for this prefix. However, Orange had also allocated subprefix 194.2.155.0/24 to Ubisoft which does not have a ROA. Hence, Orange's ROA makes Ubisoft's prefix advertisement appear invalid. In contrast, if Orange were certified with DISCO, the certificate and ROA would not cover Ubisoft's subprefix (since in practice messages to those IP addresses reach Ubisoft and not Orange).

For each prefix in the DISCO ROA, the agent configures the maximum length automatically to avoid the pitfalls in manual configurations. More precisely, for each certified address block, the agent gathers the prefixes in that address block that the AS announces. It computes the shortest list of prefixes and maxLength combination that covers precisely the announced prefixes (by running Algorithm 1 from [15]). This prefix list is then specified in the ROA issued by the agent, and each prefix in it is potentially followed by a list of excluded subprefixes. The automated procedure allows mitigating common mistakes in issuing ROAs.

3.5 Some Caveats

We next discuss weaknesses in DISCO's de facto ownership mechanism and explain how DISCO handles them.

Certifying ownership of unannounced IP prefixes. DISCO relies on Internet routing to certify IP address ownership. This limits DISCO to only certifying ownership of IP prefixes that are advertised in BGP. If an attacker announces an otherwise unannounced prefix in BGP, it may certify ownership through DISCO. DISCO deals with this problem by allowing the legal owners *certified by RPKI* for that (unannounced) prefix (i.e., holding a resource certificate for the prefix) to revoke the problematic DISCO certificate. Relying on RPKI in this manner conforms with the do-no-harm principle since, without DISCO, if the owner is not registered through RPKI, the attacker could hijack their prefix (announced or unannounced) anyway.

Certification in the presence of live attacks. Unfortunately, relying on de facto ownership means that DISCO cannot certify an AS's ownership over its IP prefix while someone hijacks that prefix. We believe, however, that in practice this limitation does not undermine DISCO from providing meaningful benefits: most ASes control their IP prefixes most of the time, which allows them to certify ownership and protect against later attacks.

Support for multiple origins. IP prefixes can be advertised from multiple ASes. Our analysis of BGP advertisements observed by RouteViews and RIPE RIS vantage points on

July 15, 2018 reveals that this is true only for a very small fraction, approximately 0.52% and 0.53%, of advertised IPv4 and IPv6 prefixes are announced by multiple ASes (3,741 out of 720,114 and 305 out of 57,141 announced IPv4 and IPv6 prefixes respectively). Importantly, DISCO can support multiple origins by deploying agents with the same public key in all origins.[1]

3.6 DISCO Facilitates Path-End Validation

Path-end validation was recently proposed as a *deployable* alternative to BGPsec that attains a comparable level of security against BGP path-manipulation attacks. Under the original path-end validation proposal [9], an *RPKI-certified* origin of an IP prefix can specify a list of direct neighboring ASes permitted to transit traffic destined for that prefix. This implies that widespread adoption of path-end validation is conditioned on widespread deployment of RPKI. Fortunately, DISCO certification can serve as the basis for path-end validation.[2]

4 IMPLEMENTATION DETAILS

We implemented a prototype of DISCO. Our implementation involves less than 100 and 200 lines of non-library Python code for the agent and registrar, respectively. The code is available at https://github.com/tmshlvck/disco-prototype.

Ownership indicators. In our implementation, ownership indicators consist of a 256-bit public key, expressed as a new BGP optional transitive attribute. Our implementation uses type `0x40` for the attribute to avoid interference with another standardized or squatted attribute types (see [29, 30]) in case of prefix leakages during experiments. We acknowledge the need for standardizing the use of this BGP attribute type before large-scale experiments and adoption of DISCO.

Vantage points. In our implementation, the registrar pulls BGP data from the vantage points afforded by RouteViews and RIPE RIS. Currently, we utilize 21 RouteViews and 23 RIPE RIS vantage points, which provide traces of the full BGP advertisements they observe (optional transitive attributes included) needed for DISCO. Our registrar implementation supports the MRT [6] data input, and so connecting new data vantage points is possible with low effort.

5 EVALUATION

We evaluate DISCO's deployability and security benefits.

5.1 Deploying DISCO

The main obstacle to deploying DISCO on the Internet is that it relies on propagating the ownership indicator with announcements.

To test whether DISCO could be deployed on today's Internet, we used our prototype implementation to attach the ownership indicator to BGP announcements for prefix 188.227.157.0/24 (a prefix under our control) on July 4, 2017. We found which ASes propagated our announcements by observing the AS-PATH as reflected at the public vantage points available to DISCO (see §4). We defer the investigation of how the specific selection of vantage points impacts the results to the future.

We observed that the prefix with the ownership indicator was relayed by 292 distinct ASes. We have also found that 3 distinct ASes appeared on an AS-PATH without the ownership indicator. At least one of these 3 ASes discarded the ownership indicator but relayed the announcement. In the following day, we announced the same prefix, but without the ownership indicator, and observed the AS-PATH as viewed from the public vantage points. We found 11 additional ASes that relayed the announcement without the ownership indicator but did not relay our previous announcement. We, therefore, suspect that these ASes filter announcements with unknown optional transitive attributes. In total, 306 distinct ASes observed either/both announcements.

While clearly a much deeper investigation is needed, these measurements suggest that most of the ASes will accept (and relay) BGP announcements with the ownership indicator, and so DISCO seems to be compatible with today's Internet.

Our experience in practice. We connected an iBGP session from the agent to a router running Cisco IOS 15.2, and announced our test prefix with the ownership indicator attached. We observed the prefix on all of the vantage points used by our implementation (see §4) within less than 5 minutes. We acknowledge that a longer evaluation of more prefixes is required before DISCO can be adopted.

5.2 Security

A prefix protected by a DISCO ROA is guaranteed the same level of security against prefix and subprefix hijacks as with RPKI ROAs. Lychev et al. [23] evaluate these benefits extensively. Our security evaluation is therefore focused on DISCO's certification mechanism. We apply the BGP simulator from [9], which realizes the BGP route-computation framework presented in [16], to the empirically-derived CAIDA AS-level graph from July 2018 [2].

To evaluate DISCO, we build on the measurement results from §5.1. Our simulations presuppose that (1) $\frac{292}{306} = 95.4\%$ of the ASes on the Internet propagate BGP announcements with unknown optional transitive attributes (conforming with the BGP standard), (2) $\frac{3}{306} < 1\%$ discard the unknown optional attribute but forward the rest of the path vectors, and (3) $\frac{11}{306} = 3.6\%$ discard the announcement with unknown optional transitive attributes altogether.

5.2.1 Setting the certification threshold.

We begin by testing how the findings in §5.1 affect DISCO's certification mechanism. In each iteration of the following simulation, we randomly select a different origin AS for the announcement

[1] When there are multiple origins but a single AS is the final destination of all traffic, e.g., one AS is the legal owner and another organization provides it with network services such as anycast, DISCO certification for the legal owner is precisely as in the single origin case.

[2] De facto ownership certification, however, does not facilitate BGPsec adoption, because BGPsec relies on authenticated ownership of AS numbers, which DISCO does not provide (cf. to RPKI).

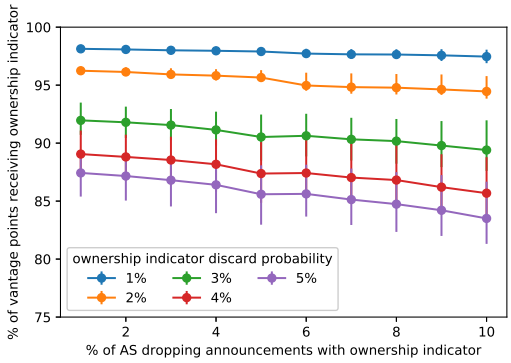

Figure 3. Almost all vantage points observe ownership indicator

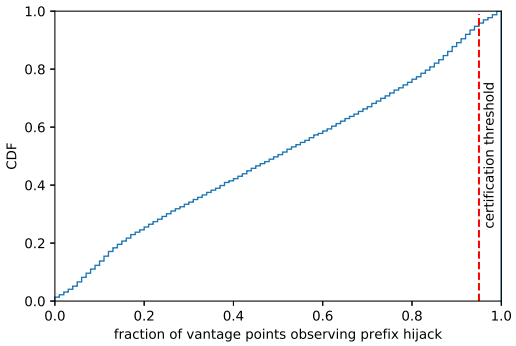

Figure 4. Certification under prefix hijack (real vantage points)

and measure the fraction of our implementation's vantage points (§4) that would observe the ownership indicator. For every data point in the following results, we run 10^5 iterations of the simulator and present the average.

Figure 3 shows the average fraction of vantage points that would observe the ownership indicator. We evaluate different probabilities of discarding announcements with ownership indicators and dropping the ownership indicators by intermediate ASes. Based on the Internet measurements (1% discard the indicator and 3.6% discard the attribute), we find that a threshold of 95% allows certifying ownership (on average).

5.2.2 **Certification is resilient to network attackers.** One concern with the de-facto certification is that the attackers may be able to modify the ownership indicator and certify their public key instead of the legitimate owner. We next show that to do that, attackers would need to hijack traffic to the victim for the entire certification interval, partially or entirely disconnecting the victim for a substantial period, and effectively exposing itself with a highly-visible attack.

Prefix hijack. We first consider a prefix hijacking attacker. To evaluate the attacker's success rate, we select random attacker-victim pairs, where both attacker and victim announce the same prefix. Figure 4 shows the probability that either AS certifies ownership with the existing RouteViews and RIPE RIS vantage points. We find that in less than 4% of cases either the attacker or the victim were able to do so. The reason is that traffic is split among them and so typically neither party achieves the required level of visibility at vantage points to certify ownership. This is a positive result since the attacker cannot hijack a prefix to certify ownership in most cases. (Although it also means that the legitimate owner cannot certify ownership while being hijacked.)

Subprefix hijack. A more aggressive attack is for the attacker to attach their public key to a subprefix of some prefix announced by the victim organization. Since the attacker is the only party announcing the subprefix, their announcement will reach all vantage points, and allow to certify the prefix. However, this certification attack requires the attacker to launch a global subprefix hijack (for an entire certification interval). In this case, the attacker is unlikely to maintain their route to the victim, causing the victim to lose traffic for the whole certification interval. To evaluate this attack strategy in simulation, we consider all possible attacker-victim AS pairs. We have the attacker announce the subprefix to their peers one by one (excluding direct links to the victim) at random order until either the announcement reaches 95% of the vantage points (the certification threshold), or the attacker loses their route to the victim.We found about 281K attacker-victim pairs through simulation where the attacker can carry out this attack yet maintain their route to the victim to avoid detection. However, in over 270K of the cases, the attacker is the single direct upstream provider of the victim, rendering the routing attack useless, because the upstream provider already receives all traffic to the victim.

The above results show that a remote attacker (which is not upstream of the victim) has a rather low-chance of success in certifying ownership without disconnecting the victim. Considering the substantial certification interval, such an attack is very likely to be noticed.

6 RELATED WORK

Some studies explore how to apply anomaly detection to identify routing errors and attacks [20, 31]. Similarly to DISCO, these rely on a global view of Internet routes afforded by vantage points, but only to identify suspicious behavior, as opposed to binding IP prefixes to public keys.

Manual configuration procedures hinder adoption of many cryptographic protocols, and as a result, automatic variants were suggested for setting up IPsec tunnels and TLS [5, 14]. Notably, Let's Encrypt reflects an approach similar to DISCO, issuing X.509 TLS certificates based on de-facto ownership of domain names, and is widely accepted [3, 5, 24].

7 CONCLUSION

We presented DISCO, an automated system for certifying ownership over IP prefixes that addresses RPKI's deployment challenges. Our evaluation shows that DISCO is readily deployable on today's Internet and provides a tangible path to security against prefix and subprefix hijacks. DISCO's deployment is also a first step towards protecting the Internet from more sophisticated attacks on BGP, e.g., via path-end validation [7, 9].

Acknowledgements

We thank the anonymous reviewers and our shepherd Renata Teixeira for valuable comments and suggestions. This work was supported by NSF grants 1414119, 1840041 and 1413920, an ERC starting grant (4th author), award SOW BL 7891, Comcast Corporation, the German Federal Ministry of Education and Research (BMBF), the Hessian Ministry of Science and the Arts and the DFG.

REFERENCES

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

[2] The CAIDA AS Relationships Dataset. http://www.caida.org/data/as-relationships/, Jan. 2016.

[3] M. Aertsen, M. Korczynski, G. C. M. Moura, S. Tajalizadehkhoob, and J. van den Berg. No domain left behind: is Let's Encrypt democratizing encryption? In *ANRW*, pages 48–54. ACM, 2017.

[4] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the Internet. In J. Murai and K. Cho, editors, *SIGCOMM*, pages 265–276. ACM, 2007.

[5] R. Barnesm, J. Hoffman-Andrews, and J. Kasten. Automatic certificate management environment (ACME). https://tools.ietf.org/html/draft-ietf-acme-acme-08, October 2017. Internet-Draft.

[6] L. Blunk, M. Karir, and C. Labovitz. Multi-threaded routing toolkit (mrt) routing information export format, October 2011. RFC6396.

[7] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira. One Hop for RPKI, One Giant Leap for BGP Security. In *HotNets*, pages 10:1–10:7. ACM, 2015.

[8] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira. One hop for rpki, one giant leap for bgp security. In J. de Oliveira, J. Smith, K. J. Argyraki, and P. Levis, editors, *HotNets*, pages 10:1–10:7. ACM, 2015.

[9] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira. Jumpstarting BGP Security with Path-End Validation. In *SIGCOMM*, pages 342–355. ACM, 2016.

[10] D. Cooper, E. Heilman, K. Brogle, L. Reyzin, and S. Goldberg. On the risk of misbehaving RPKI authorities. In *HotNets*, pages 16:1–16:7. ACM, 2013.

[11] J. Durand, I. Pepelnjak, and G. Doering. Bgp operations and security, February 2015. RFC7454.

[12] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman. Are We There Yet? On RPKI's Deployment and Security. In *NDSS*, 2017.

[13] Y. Gilad, S. Goldberg, K. Sriram, and J. Snijders. The Use of Maxlength in the RPKI. https://tools.ietf.org/html/draft-yossigi-rpkimaxlen-01, September 2017. Internet-Draft.

[14] Y. Gilad and A. Herzberg. Plug-and-Play IP Security: Anonymity Infrastructure instead of PKI. In *ESORICS*, pages 255–272. Springer, 2013.

[15] Y. Gilad, O. Sagga, and S. Goldberg. Maxlength considered harmful to the RPKI. In *CoNEXT*, pages 101–107, 2017.

[16] P. Gill, M. Schapira, and S. Goldberg. Modeling on Quicksand: Dealing with the Scarcity of Ground Truth in Interdomain Routing Data. *Computer Communication Review*, 42(1):40–46, 2012.

[17] T. Hlavacek. ROV adoption rate measurement. https://2017.peeringdays.eu/file.php?id=16&n=23_7_cz.nic_hlavacek.pdf, March 2017. CEE Peering Days 2017.

[18] T. Hlavacek, A. Herzberg, H. Shulman, and M. Waidner. Practical experience: Methodologies for measuring route origin validation. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 634–641, June 2018.

[19] D. Iamartino, C. Pelsser, and R. Bush. Measuring BGP Route Origin Registration and Validation. In *PAM*, volume 8995 of *LNCS*, pages 28–40. Springer, 2015.

[20] J. Karlin, S. Forrest, and J. Rexford. Pretty Good BGP: Improving BGP by Cautiously Adopting Routes. In *ICNP*, pages 290–299. IEEE Computer Society, 2006.

[21] M. Lepinski and S. Kent. An infrastructure to support secure internet routing, February 2012. RFC6480.

[22] M. Lepinski and K. Sriram. Bgpsec protocol specification, September 2017. RFC8205.

[23] R. Lychev, S. Goldberg, and M. Schapira. BGP Security in Partial Deployment: Is the Juice Worth the Squeeze? In *SIGCOMM*, pages 171–182. ACM, 2013.

[24] A. Manousis, R. Ragsdale, B. Draffin, A. Agrawal, and V. Sekar. Shedding Light on the Adoption of Let's Encrypt. *CoRR*, abs/1611.00469, 2016.

[25] NIST. RPKI Monitor. http://rpki-monitor.antd.nist.gov/, 2015.

[26] Y. Rekhter, T. Li, and S. Hares. A border gateway protocol 4 (bgp-4), January 2006. RFC4271.

[27] A. Reuter, R. Bush, Í. Cunha, E. Katz-Bassett, T. C. Schmidt, and M. Wählisch. Towards a rigorous methodology for measuring adoption of RPKI route validation and filtering. *ACM Computer Communication Review*, 2018.

[28] RIPE NCC. YouTube Hijacking: A RIPE NCC RIS case study, March 2008.

[29] J. Snijders. BGP Large Communities. SINOG 4, Ljubljana, Slovenia, May 2017.

[30] J. Snijders. Deprecation of bgp path attribute values 30, 31, 129, 241, 242, and 243, February 2017. RFC8093.

[31] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz. Listen and Whisper: Security Mechanisms for BGP. In *NSDI*, pages 127–140. USENIX, 2004.

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

[33] A. Toonk. Turkey Hijacking IP Addresses for Popular Global DNS Providers. BGPMon.

[34] P.-A. Vervier, O. Thonnard, and M. Dacier. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In *NDSS*. The Internet Society, 2015.

[35] M. Wählisch, O. Maennel, and T. C. Schmidt. Towards Detecting BGP Route Hijacking Using the RPKI. *SIGCOMM Comput. Commun. Rev.*, 42(4):103–104, Aug. 2012.

[36] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson. RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem. In *Proc. of Fourteenth ACM Workshop on Hot Topics in Networks (HotNets)*, New York, 2015. ACM.

[37] R. White. Deployment Considerations for Secure Origin BGP (soBGP)., June 2003.