One Hop for RPKI, One Giant Leap for BGP Security

Avichai Cohen
The Hebrew University of Jerusalem
avichai.cohen@mail.huji.ac.il

Yossi Gilad
The Hebrew University of Jerusalem
mail@yossigilad.com

Amir Herzberg
Bar-Ilan University
amir.herzberg@gmail.com

Michael Schapira
The Hebrew University of Jerusalem
schapiram@huji.ac.il

ABSTRACT

Extensive standardization and R&D efforts are dedicated to establishing secure interdomain routing. These efforts focus on two complementary mechanisms: *origin authentication* with RPKI, and *path validation* with BGPsec. However, while RPKI is finally gaining traction, the adoption of BGPsec seems not even on the horizon. This is due to inherent, possibly insurmountable, obstacles, including the need to replace today's routing infrastructure, meagre benefits in partial deployment and online cryptography.

We propose *path-end validation*, a much easier to deploy alternative to BGPsec. Path-end validation is a modest extension to RPKI that does not require modifications to BGP message format nor online cryptography. Yet we show, through extensive simulations on empirically-derived datasets, that path-end validation yields significant security benefits, even with very limited partial deployment. We present an open-source prototype implementation of path-end validation, which does not require changing today's routers, illustrating the deployability advantage over BGPsec.

Categories and Subject Descriptors

C.2.2 [**Network Protocols**]: Routing protocols—*Security*

1. INTRODUCTION

The Internet infrastructure was not designed with security in mind, and is consequently alarmingly vulnerable. We focus on the arguably most acute problem: *securing interdomain routing*, that is, routing between the administrative domains, or "Autonomous Systems" (ASes), which comprise the Internet. As highlighted by many high-profile configuration errors and attacks (e.g., [1, 2, 4]), the Border Gateway Protocol (BGP), today's de facto interdomain routing protocol, is hazardously insecure [7]. However, the adoption of BGP security solutions is difficult and is proceeding slowly [15].

The current prevalent paradigm for securing interdomain routing, as advocated, for instance, by the IETF's Secure Inter-Domain Routing (SIDR) group, consists of two steps: (1) *origin authentication* by deploying the Resource Public Key Infrastructure (RPKI) [19], followed by (2) *path validation* by replacing BGP with BGPsec [8], a secure interdomain routing protocol that extends BGP.

RPKI [19] certifies records binding an IP-prefix with the number and public key of its *origin Autonomous System (AS)*, i.e., the AS that "owns" that prefix. RPKI certificates allow BGP routers to perform *origin authentication* [25]: detect and discard *prefix hijacks*, BGP route advertisements where an IP prefix is announced by an AS that is not its legitimate owner. Prefix hijacks happen frequently (e.g., see [1, 2, 3, 6]) and motivate the adoption of RPKI, which is finally gaining traction [26].

Origin authentication (via RPKI) provides an important first step towards securing interdomain routing, yet it is insufficient to prevent path-manipulation attacks. In particular, even with RPKI fully deployed, the attacker can still perform the *next-AS* attack, i.e., announce a fake link between himself and the victim AS. To address this and other path-manipulation attacks, the IETF is standardizing BGPsec [8], which uses digitally-signed BGP announcements. BGPsec prevents a BGP-speaking router from announcing a path that is not a legitimate extension of a valid path that it received. To ensure this, BGPsec requires each AS to sign every path advertisement that it sends to another AS, and to validate all the signatures of previous ASes along the path. Unlike RPKI, integration of BGPsec necessitates changes to BGP routers and introduces nontrivial run-time computational overhead [15]. Worse yet, recent work on adoption of BGP security [22] shows that in partial deployment, BGPsec is expected to achieve disappointingly meagre security benefits over RPKI, while potentially even leading to *less* security

HotNets '15 November 16–17 2015, Philadelphia, PA USA

DOI: http://dx.doi.org/10.1145/2834050.2834078.

and other undesirable phenomena (e.g., routing instabilities).

The above serious, arguably insurmountable, obstacles facing the adoption of BGPsec, beg the question: are there alternative security measures that are easier to deploy than full-fledged BGPsec, share the advantages of RPKI (no need to replace legacy routers, no online crypto, etc.), yet provide comparable security benefits?

We offer a positive answer by presenting *path-end validation*. Path-end validation has a much more modest goal than (full) path validation with BGPsec: it only attempts to ensure that the last hop along the advertised path is valid, i.e., that the origin AS for the given prefix, authenticated by RPKI, has approved reaching it via the previous AS along the path. What level of security can such a modest security check provide? Our simulations show that even with relatively few adopters, path-end validation suffices to achieve a level of security that is significantly better than RPKI's and, in fact, close to the security guarantees of BGPsec.

In retrospect, these surprisingly good news are easy to comprehend. An attacker cannot fool neighboring ASes regarding the business relationship between him and them, and so should intuitively advertise as short a path to the victim's prefix as it can get away with. However, the attacker is now at a big disadvantage: he cannot pretend to own the prefix or even claim to be directly connected to the victim AS without being detected. Consequently, the path to the victim announced by the attacker to his neighbors must be of length at least 2 (and this path length is increased as the path is further propagated in the Internet). This, combined with the fact that BGP paths are typically short (consistently about 4-hops-long on average [24]), intuitively implies that the vast majority of ASes will not fall victim to the attack.

Importantly, path-end validation constitutes a radical departure from BGPsec's design philosophy, which focuses on achieving "rigorous AS path protection" [27] and does not distinguish between paths that are partially validated and paths that are not validated at all. Our results rely on the insight that while partial validation of paths does not *always* benefit security, there is significant benefit in validation of *path suffixes*, which forces the attacker to announce longer paths. Moreover, the shortness of interdomain routes implies that even validating one-hop suffixes (that is, path-end validation) is enough, as discussed above.

Our design of path-end validation is much easier to realize than BGPsec as it does not require real-time cryptographic operations, nor replacing today's interdomain routing infrastructure. In fact, deploying our path-end validation scheme requires only minor extensions to the mechanisms already established for RPKI.

Contributions

Path-end validation (Section 2). We identify path-end validation as a target security objective that is both achievable without modifying the routing infrastructure and can significantly improve interdomain routing security even in partial deployment.

Evaluation of impact on BGP security (Section 3). We perform extensive simulations to evaluate the security impact of adopting path-end validation, for different adoption rates. The results are encouraging: significant impact is obtained even with very limited partial deployment. We also identify the potential for regional adoption of routing security mechanisms, possibly government sponsored/driven. Specifically, we analyze the impact of adoption of path-end validation in specific geographical/national regions, and its potential to protect local communication within these regions.

Implementation (Section 4). We present an open-source implementation. Our implementation involves grappling with operational issues such as BGP router configuration.

A discussion of related work and our conclusion appear in Sections 5 and 6.

2. PATH-END VALIDATION

To ensure a *deployable* improvement to BGP security, we identify two main goals:

Avoid changes to routers and online cryptography. One of the main concerns with BGPsec and similar proposals is that they require changes to routers. Furthermore, these proposals required validation of (multiple) signatures when processing BGP advertisements, as well as signing upon sending a BGP advertisement. Such requirements make deployment very challenging. We therefore aim to provide a mechanism that can enforce security policies only by configuring today's routers (using their existing capabilities and interfaces).

Provide significant security benefits in partial deployment. Deployment of a new mechanism for securing interdomain routing, involving tens of thousands of independently administered ASes, may take time, even if it does not introduce changes to today's infrastructure. We therefore aim to provide significant security benefits under the realistic partial adoption scenario. This should be contrasted with BGPsec, which provides meagre benefits over RPKI under partial adoption [22].

2.1 Design

Path-end validation builds upon RPKI, i.e., an adopting AS must first authenticate ownership of its IP address block(s) through RPKI. For path-end validation, the AS uses its RPKI-authorized public key to sign the list of approved neighboring ASes through which it can be reached. These lists, received from different ASes, are stored in a database.

Any BGP router anywhere can thus use the information in the database to filter advertisements where the second-to-last AS does not appear in the list specified by the last AS on the advertised path ('path-end forgery'), as well as advertisements where the last AS is not the one specified in the relevant RPKI Route Origin Authorization (ROA). This prevents next-AS, prefix-hijacking and sub-prefix hijacking attacks against the adopting AS.

Example. Consider the network in Figure 1. AS 1 is the "victim", i.e., the AS whose traffic the attacker, AS 2, attempts to hijack. Suppose that AS 1, and also ASes 20, 200, and 300, are adopters. Path-end validation guarantees that they will not fall victim to the following two attacks: (1) AS 2 announces the prefix 1.2.0.0/16 (prefix hijack) or a subprefix of 1.2.0.0/16 (subprefix hijack) and (2) AS 2 announces the bogus route $2-1$ to the prefix 1.2.0.0/16 or its subprefix, i.e., pretends to be directly connected to AS 1 (a next-AS attack). Importantly, as we describe in the following subsection, path-end validation allows even "isolated" adopters (such as AS 20) to protect legacy ASes along the route (cf. to BGPsec).

Path-end validation does not protect against the "2-hop attack", in which AS 2 pretends to be directly connected to a (legacy) neighbor of AS 1 (say, announces the bogus route $2-40-1$). However, such attacks turn out to be quite ineffective, since the path that the attacker can announce to the victim's prefix, must consist of at least two hops, and BGP paths are typically short (about 4-hops-long on average [24]), i.e., most ASes will not fall victim to the attack. We validate this intuition via extensive simulations (Section 3).

Retain BGP protocol and message format. Path-end validation, in contrast to BGPsec, does not use route advertisements to propagate security-related information. Instead, it extends RPKI's *offline* mechanism, which periodically syncs local caches at adopting ASes to global databases, and pushes the resulting whitelists to BGP routers (see details in [9]). This has two advantages. First, deployment is easier as no changes are introduced to routers and no online cryptographic validation is required. Second, it allows validation of BGP advertisements even when there are intermediate legacy routers along the path. Specifically, path-end validation enables protecting *all* adopters (including AS 20 in Figure 1) from next-hop attacks against other adopters (e.g., AS 1). Moreover, even an isolated adopter on the path, such as AS 20, can protect the non-adopters "behind" it by preventing malicious routes from being disseminated and, in particular, a malicious advertisement will not reach AS 30. We show in the following section, that this leads to significant security benefits even in partial deployment, greatly improving over those achievable by BGPsec [22].

Privacy. To accommodate ISPs that are reluctant to disclose the identities of their neighbors (more specifically, customers) for fear of competitors, our design supports a "private mode", where an ISP deploys path-end filtering but does not register its neighbors in the database. This protects privacy-concerned ISPs from falling victim to next-AS attacks against others, without compromising privacy (and increases protection for the other ASes).

We point out, however, that: (1) Over 85% of ASes are not ISPs and have no business interest in keeping the identities of neighbors secret (see, e.g., PeeringDB [28]). In particular, a large fraction of Internet traffic is originated at and destined for ASes with no customers like Google, Netflix, etc. (We present in Section 3 the security benefits of path-end validation for large content providers.) (2) Even if an ISP does not reveal the identity of a customer, that customer can choose to reveal its connection to this ISP so as to protect itself.

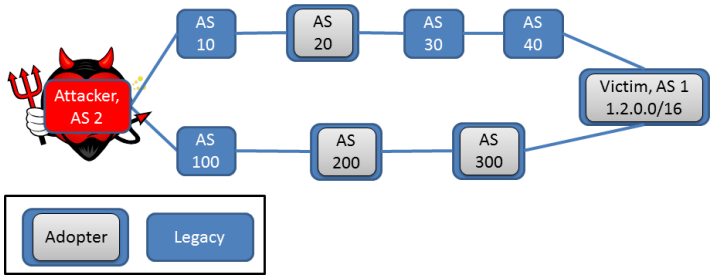

Figure 1: Partial deployment example.

3. IMPLICATIONS FOR BGP SECURITY

We evaluate the level of security provided by path-end validation in *partial* deployment and compare it to RPKI and BGPsec. Our simulation results, presented below, show that even under very partial deployment, path-end validation provides tangible security benefits that come close to those of BGPsec in full deployment (before legacy BGP is deprecated).

Simulation framework. We evaluate different attack strategies. We quantify the attacker's success by the fraction of ASes he is able to attract, as in [12, 22]. Our simulations apply the BGP route-computation framework presented in [12, 13, 16] to the empirically-derived CAIDA AS-level graph [5] from December 2014 (links are annotated with inferred bilateral business relationships and contain previously hidden peering links within IXPs [14]). We averaged our measurements over 10^6 combinations of attacker-victim ASes, where the attacker and victim pairs were selected uniformly at random.

3.1 Results: Protecting Against Attacks

Consider the graphs in Figure 2. The x-axis describes 11 deployment scenarios corresponding to adoption of path-end validation by the set of $0, 10, \ldots, 100$ largest ISPs. The y-axis is the average fraction of ASes on the Internet whose traffic the attacker is able to attract to his network if the victim registers a path-end record. We consider three possible strategies for the attacker: (1) prefix hijacking, (2) next-AS attack, and (3) the 2-hop attack. Path-end validation implies that adopters can detect and ignore the first two attacks, while the third attack goes unnoticed by the defense. The graphs also present three reference lines: the level of security of RPKI in full deployment when the attacker launches the next-AS attack (which cannot be detected by RPKI), the level of security when the attacker launches the next-AS attack under partial deployment of BGPsec (and full deployment of RPKI), and the level of security of BGPsec in full de-

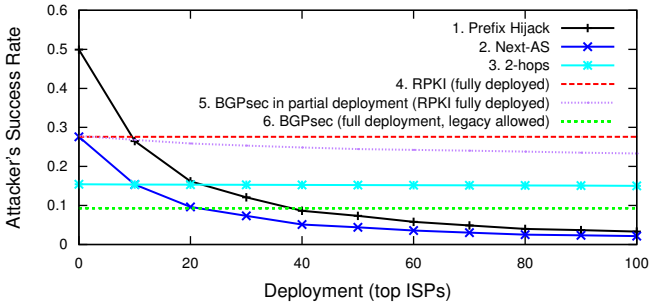

(a) Global security evaluation.

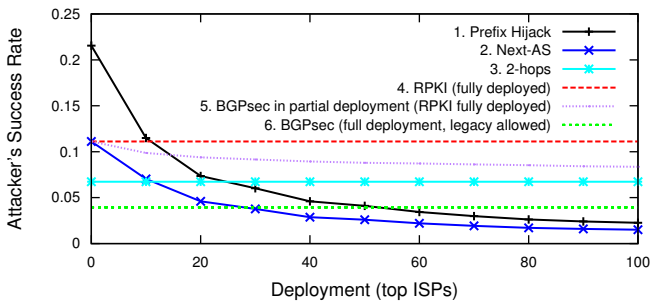

(b) Protection of top content providers.

Figure 2: Attacker success rate for different strategies as function of path-end filters deployment.

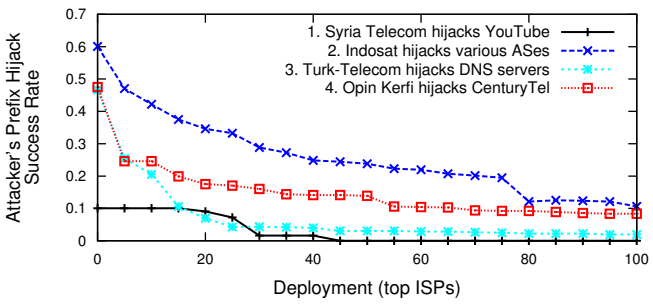

(a) Prefix-hijack attack.

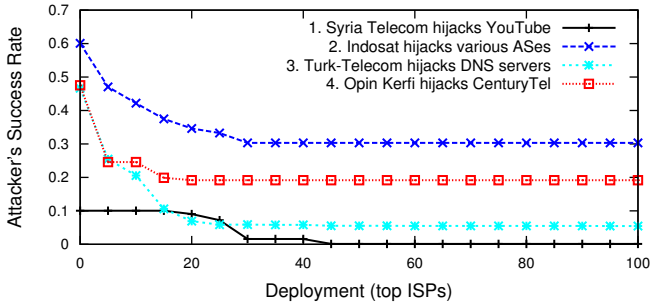

(b) Best attack strategy.

Figure 3: High profile past incidents: attacker's success rate as function of the number of adopters.

ployment so long as legacy BGP is not deprecated, as studied in [22] (our simulations also confirmed the results in [22]).

As shown in Figure 2a, path-end validation is remarkably effective at thwarting prefix hijacks and next-AS attacks, in contrast with the meagre improvement achieved by BGPsec under the same partial deployment scenario (line 5 in the figure). Even with only 20 adopters, the attacker is better off resorting to the 2-hop attack, resulting in success rate of 15%—a big improvement over RPKI in full deployment (28% success rate) and not far from BGPsec (9%). With 100 adopters, the success rates of prefix hijacks and next-AS attacks go down to 5% (from 50% and about 30%, respectively).

Figure 2b describes the protection that path-end validation provides to large content providers (Google, Amazon, Netflix, etc; list taken from [22]), as these generate a significant fraction of traffic on the Internet. Our results here show that path-end validation provides significant security benefits, reducing the attacker's success rate with his best strategy (2-hop attack) to 7% with only 20 adopters, a significant improvement over RPKI, even when RPKI is ubiquitously deployed.

3.2 Revisiting High-Profile Past Incidents

We revisit four recent high-profile prefix-hijack incidents to illustrate the security benefits of path-end validation: (1) Syria-Telecom hijacks YouTube [6] on December 9th, 2014 (2) Indosat hijacks over 400,000 prefixes on April 3rd, 2014 [1]; (3) Turk-Telecom hijacks DNS resolvers in Google, OpenDNS and Level3 on March 29, 2014 [3]; and (4) Opin Kerfi's (an ISP in Iceland) repeated prefix-hijacks [2] in December 2013. All these incidents involved prefix hijacking. Some of these incidents are attributed to benign configuration errors while others are suspected attacks.

Clearly no simulation framework is rich enough to capture all intricacies of interdomain routing (e.g., ASes' actual routing policies). Our aim is therefore not to predict the BGP outcome, but to get a high-level idea of path-end validation's potential influence in these concrete scenarios. We computed, for every attacker-victim pair, the attacker's success rate with X adopters from the largest ISPs, where $X = 0, 5, 10, \ldots, 100$, and for 3 attack strategies: (1) prefix hijacking, (2) the next-AS attack, and (3) the 2-hop attack.

Figure 3a describes the attacker's success rate for prefix-hijack attacks. Observe that path-end validation has a noticeable effect with only 15 adopters, whereas with 80 adopters the attacker's success rate drops significantly. Our results for next-AS attack exhibit the same trends. Hence, even with a modest number of adopters, the attacker's best strategy is to launch the 2-hop attack and avoid detection by the path-end validation mechanism.

Figure 3b plots the attacker's success rate in each deployment scenario ($X = 0, 5, \ldots, 100$ adopters) for his *best* attack strategy among the three (prefix hijacking, next-AS, 2-hop). Consider, for example, the Opin Kerfi versus Century Tel scenario. Before any AS adopts path-end validation, the attacker's best strategy is prefix hijacking, resulting in a success rate of almost 50%, i.e., attracting almost half of the

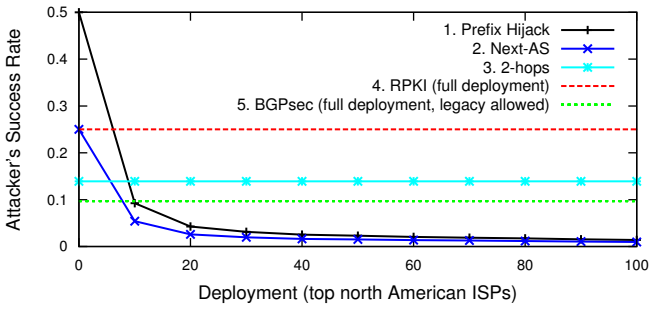

(a) Internal attackers

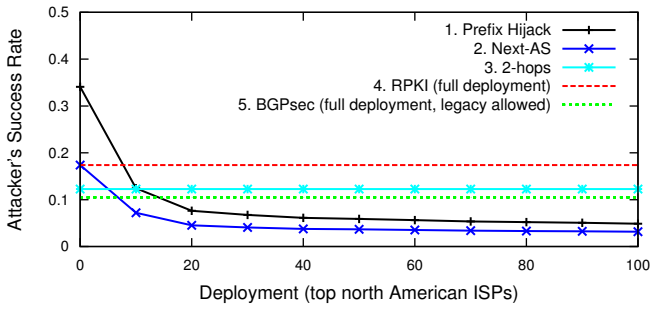

(b) External attackers

Figure 4: Protection for North-American ASes by local adopters

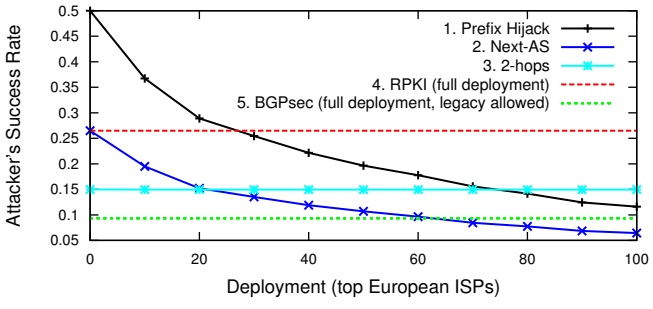

(a) Internal attackers

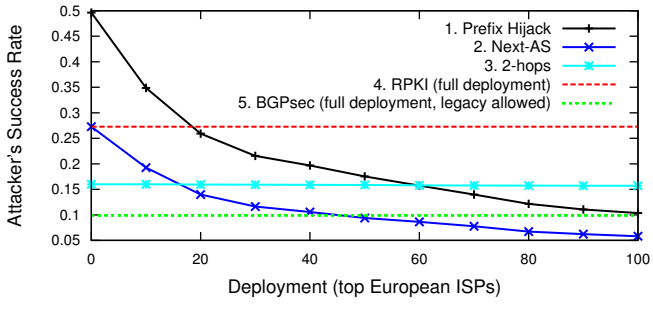

(b) External attackers

Figure 5: Protection for European ASes by local adopters

Internet. As more and more large ISPs adopt the path-end validation defense, the success of prefix-hijacking is significantly decreased. Indeed, even with 15 adopters, the attacker is better off switching to the 2-hop attack, and so the attacker's success rate remains fixed at about 20% henceforth.

3.3 Geography-Based Deployment

One possible strategy for boosting initial deployment of path-end validation is for governments to incentivize large ISPs in their countries to adopt (i.e., install the corresponding filtering rules in their routers, as explained in Section 4). We investigate whether such *local* adoption can protect *l*ocal communication, i.e., protect the ASes in that geographical region. This is important, for instance, since many end-users retrieve content from servers in their geographic region due to the popularity of content delivery networks.

We used the Regional Internet Registries (RIRs) division of the world into five geographic regions and considered adoption only by ISPs in a particular region. We then measured how many benign ASes *in the region* are fooled to take a malicious route to a victim *in the region* advertised by attackers (internal and external to the region).

Figure 4 shows the fraction of North-American ASes that an attacker can attract when trying to capture traffic to a North-American AS. We find that even with only 10 adopters, path-end validation protects the communication between two ASes in North America, even if the attacker is co-located in North America (Figure 4a), reducing the attacker's success rate to just above 15%. The corresponding results for Europe, presented in Figure 5, show that the top 60 European ISPs need to adopt the protocol to achieve a similar effect. While not as good as in the North-American case, these results still provide significant motivation for adoption, considering recent attacks on ASes from Iceland and Belarus [2], and the emphasis on routing security of the European Union.

4. PROTOTYPE

We present a prototype implementation, which complements RPKI and allows to deploy path-end validation with today's routing infrastructure. We envision this prototype as a first step, providing an immediate defense against routing attacks, until the path-end validation mechanism is integrated with RPKI. We design our prototype to be compatible with RPKI in order to simplify future migration. To allow the community to deploy and experiment with path-end validation, our implementation is open-source and available at `https://github.com/routingsec/pathend`.

4.1 Implementation Details

Our implementation defines the path-end record data structure, where an origin AS holding an RPKI certificate specifies a list of approved neighboring ASes. We use the following ASN.1 syntax to describe the record format:

```
PathEndRecord ::= SEQUENCE {
        timestamp Time,
        origin  ASID,
        adjList SEQUENCE (SIZE(1..MAX)) OF ASID}
```

Path-end records are stored in public repositories, similar to RPKI's Route Origin Authorization records (ROAs) [20].

When a repository receives an AS's path-end record to store along with a signature from the origin AS, it retrieves that AS's RPKI certificate and verifies the signature over its path-end record (we utilize RPKI's certificate revocation lists to remove records in case the signing key was revoked), as well as validates that the timestamp is not before an already existing entry for the same origin. An AS can update or delete its path-end records using a signed announcement sent to the repositories, similar to ROAs in RPKI.

Since BGP routers do not yet accept path-end records, we also implement an "agent application" that updates periodically from the repositories and configures BGP routers in the adopter's network with path-end-filtering policies. To avoid trusting the path-end record repository (until the trusted RPKI repositories support distribution of path-end records) the agent retrieves the corresponding record signatures, which path-end repositories store along with the records, and verifies the signature over each record. The agent also retrieves each update from a random repository, so as to ensure that a compromised repository cannot remove a record or provide an obsolete image of the database. The agent application supports an automated mode, where the network administrator provides the credentials needed to configure a BGP router, and the router's IP address, and the agent automatically deploys filtering rules according to the path-end records retrieved. The agent also supports a manual mode, where it only outputs the filtering policies to a router configuration file, which the administrator can later apply.

4.2 Deploying Path-End Filtering Rules

We next describe how the agent configures filtering rules for BGP advertisements on *today's* routers. For each AS the agent deploys only *one* filtering rule. This results in an order of magnitude *less* rules than origin authentication with RPKI, which involves a filtering rule per IP-prefix (there are roughly 50K ASes advertising over 500K IP-prefixes). We therefore believe that path-end validation can scale to support the entire set of ASes. To illustrate the filtering rules installed, we use the network topology in Figure 1 and describe the routing policy for protecting AS1, whose adjacent ASes are $40, 300$. Our description uses the Cisco IOS command-line interface. We verified that routers from other vendors (e.g., Juniper Networks) support the same functionality.

We first use AS1's path-end record to create the following access list (named as1), which blacklists routes containing (invalid) links to AS1 from non-adjacent ASes.

```
// disallow any AS but 40 or 300 to
// advertise a link to AS1
ip as-path access-list as1 deny _[^(40|300)]_1_
```

The agent creates another access list (named allow-all) to allow all other routes. This access list is global, i.e., created once rather than for every adopting AS.

```
ip as-path access-list allow-all permit
```

Finally we apply these policies in order, i.e., first blocking invalid routes, then allowing all others.

```
route-map Path-End-Validation permit 1
    match ip as-path as1
    match ip as-path allow-all
```

5. RELATED WORK

The security vulnerabilities of today's interdomain routing system motivated many proposals for securing BGP routing. Due to length restrictions, we only discuss the main proposals for how to *prevent path manipulation attacks*. We refer the reader to the survey in [10] for an extensive discussion, which includes important complementary directions such as *detection* of attacks (e.g., [17, 21, 31])

Past research analyzed various important aspects of these proposals, including security guarantees [7, 16] and adoptability [11, 12]. Several proposals focused on using cryptography to prevent route manipulation attacks, including S-BGP [18], psBGP [29], and BGPsec [23]. These proposals require changes to the BGP protocol, and extending routers to support online cryptographic computations, which are significant challenges to adoption. In contrast, path-end validation extends RPKI's offline approach, and can be deployed on top of the current Internet infrastructure, with only router configuration changes and offline, off-router cryptography.

Secure-origin BGP (soBGP [30]) was another proposal for securing BGP using offline validation of inter-AS connections. It was proposed to the IETF as a general framework allowing for many possible realizations (from essentially RPKI to variants that require significant changes to routers and BGP message format), but was abandoned in favor of S-BGP and BGPsec, and its properties were not analyzed.

Path-end validation is carefully engineered to be easily deployable and to provide significant security benefits in partial deployment with low overhead. We consider path-end validation as a specific embodiment of soBGP designed to realize these objectives. Our simulation results establish that path-end validation indeed provides an attractive "return on investment" in partial adoption.

6. CONCLUSION

We presented the path-end validation mechanism for improving interdomain routing security while avoiding the hurdles en route to deployment of BGPsec. Our security evaluation shows that path-end validation provides a surprisingly high level of interdomain routing security even with a modest number of adopters, and an open-source implementation shows the feasibility of its deployment on top of today's routing infrastructure. We hence believe that path-end validation provides a tangible path to significant improvements in interdomain routing security in the (seemingly very long) interim period before BGPsec is fully deployed. Our findings motivate the standardization of path-end validation and its integration into RPKI. We propose that governments and industry groups concentrate regulatory efforts and/or financial incentives on convincing large ISPs in their countries to adopt path-end validation (on top of RPKI).

Acknowledgments

This work was supported by ISF grants 420/12 and 1354/11, Israel Ministry of Science grants 3-9772 and 3-10884, the Israeli Center for Research Excellence in Algorithms, and a Marie Curie CIG.

7. REFERENCES

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

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

[3] Turkey Hijacking IP addresses for popular Global DNS providers. http://www.bgpmon.net/turkey-hijacking-ip-addresses-for-popular-global-dns-providers.

[4] Renesys Blog - Pakistan Hijacks YouTube. http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml, Feb. 2008.

[5] CAIDA AS Relationships Dataset. http://www.caida.org/data/as-relationships/, 2014.

[6] Andree Toonk. BGP Hijack Incident by Syrian Telecomunications Establishment. www.bgpmon.net/bgp-hijack-incident-by-syrian-telecommunications-establishment, 2015.

[7] H. Ballani, P. Francis, and X. Zhang. A Study of Prefix Hijacking and Interception in the Internet, 2007.

[8] S. Bellovin, R. Bush, and D. Ward. Security Requirements for BGP Path Validation. RFC 7353, August 2014.

[9] R. Bush and R. Austein. The Resource Public Key Infrastructure (RPKI) to Router Protocol. RFC 6810 (Proposed Standard), Jan. 2013.

[10] K. R. B. Butler, T. R. Farley, P. McDaniel, and J. Rexford. A Survey of BGP Security Issues and Solutions. *proc. of the IEEE*, 98(1):100–122, 2010.

[11] H. Chan, D. Dash, A. Perrig, and H. Zhang. Modeling Adoptability of Secure BGP Protocols, 2006.

[12] P. Gill, M. Schapira, and S. Goldberg. Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security, 2011.

[13] 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.

[14] V. Giotsas, S. Zhou, M. J. Luckie, and kc claffy. Inferring Multilateral Peering. In K. C. Almeroth, L. Mathy, K. Papagiannaki, and V. Misra, editors, *CoNEXT*, pages 247–258. ACM, 2013.

[15] S. Goldberg. Why is it Taking so Long to Secure Internet Routing? *Commun. ACM*, 57(10):56–63, 2014.

[16] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How Secure are Secure Interdomain Routing Protocols? *Computer Networks*, 70:260–287, 2014.

[17] J. Karlin, S. Forrest, and J. Rexford. Pretty Good BGP: Improving BGP by Cautiously Adopting Routes, 2006.

[18] S. T. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, 2000.

[19] M. Lepinski and S. Kent. An Infrastructure to Support Secure Internet Routing. RFC 6480 (Informational), Feb. 2012.

[20] M. Lepinski, S. Kent, and D. Kong. A Profile for Route Origin Authorizations (ROAs). RFC 6482 (Proposed Standard), Feb. 2012.

[21] J. Li, T. Ehrenkranz, and P. Elliott. Buddyguard: A Buddy System for Fast and Reliable Detection of IP Prefix Anomalies, 2012.

[22] 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.

[23] E. M. Lepinski. BGPsec Protocol Specification. Internet draft, https://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-13, July 2015.

[24] Mirjam Kuhne. AS Path Lengths Over Time. https://labs.ripe.net/Members/mirjam/update-on-as-path-lengths-over-time, 2012.

[25] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein. BGP Prefix Origin Validation. RFC 6811 (Proposed Standard), Jan. 2013.

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

[27] K. Sriram. BGPSEC Design Choices and Summary of Supporting Discussions. Internet draft, https://tools.ietf.org/html/draft-sriram-bgpsec-design-choices-08, 2015.

[28] R. Steenbergen. PeeringDB. http://www.peeringdb.com/, 2015.

[29] P. C. van Oorschot, T. Wan, and E. Kranakis. On Interdomain Routing Security and Pretty Secure BGP (psBGP). *ACM Trans. Inf. Syst. Secur*, 10(3), 2007.

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

[31] K. Zhang, A. Yen, X. Zhao, D. Massey, S. F. Wu, and L. Zhang. On Detection of Anomalous Routing Dynamics in BGP, 2004.