Analyzing adoptability of secure BGP routing proposals

Max Crone
mccrone@kth.se
KTH Royal Institute of Technology
Stockholm, Sweden

ABSTRACT

With the internet ever-increasing in importance for this world, the issue of insecure BGP routing has still not been solved. Despite the existence of many proposed solutions to secure BGP routing and prevent any of the attacks that are prevalent on the internet today, none have seen significant adoption amongst network operators. This paper identifies concrete metrics that aim to predict adoptability of secure BGP routing protocols. The main finding is that the most significant predictor is the factor of adoption itself, paradoxically. More practical metrics include interoperability, implementation diversity, hardware cost, and incremental deployment value. These metrics are further explained and connected to data on actual deployment status of the proposals. The concluding recommendation is to include adoptability as a major part of future protocol design processes.

CCS CONCEPTS

• **Networks** → **Routing protocols**; **Security protocols**;
• **Security and privacy** → **Security protocols**.

KEYWORDS

secure routing, bgp, bgpsec, trust, pki, rpki, disco

1 INTRODUCTION

Security of internet routing is an important and fundamental topic in the internet. Today's internet routing is made possible by the Border Gateway Protocol (BGP). The internet consists of autonomous systems (AS), which are groups of routing prefixes under control by a single organization (e.g., an internet service provider). These ASes use BGP to exchange routing and reachability information amongst each other. No single AS has a complete view of the internet, but all strive to maintain a list of neighboring ASes, or BGP peers, that are one hop closer to the intended recipient on the internet. Ultimately, a packet is routed to the AS that has authority over an IP prefix space that includes the recipient's IP address. However, routers running BGP will accept advertised routes they receive from peers by default. There is no authentication of these announcements. This insecurity of BGP has lead to many types of attacks on internet routing. Examples include BGP hijacking, interception, and AS path forgery attacks [15].

Researchers have proposed solutions to increase security or replace BGP altogether [17], many of which have never seen widespread adoption. There are many reasons why adoption has proven to often be difficult. Furthermore, all solutions differ in their security guarantees, properties, and trust models. These factors lead to a diverse landscape of security proposals, and adoption procedures and failures. This paper investigates metrics that predict the success of adoption, because clear results will contribute to more effective protocol designs and deployment strategies. Ultimately, this will improve the security of routing on the internet.

2 RELATED WORK

Research into routing security is generally classified into two different categories. Firstly, there are the works that evaluate security properties of approaches to secure BGP. These are useful for understanding the technical details on which approaches differ. Secondly, there are the works that evaluate the deployment of approaches. These provide a better overview of the challenges faced during adoption and how different approaches differ on that aspect. However, there is very little work on *adoptability* as a property of security proposals. This section briefly summarizes recent work from both categories, and discusses a single paper that proposes adoptability as a protocol property.

A survey from 2018 compares the properties of sixteen different BGP security proposals [15]. It concludes that many of these approaches only resolve a relatively small part of the problems, and additionally cause high computational overhead on hardware from Autonomous Systems (ASes). The survey also notes that the approach with the current highest rate of adoption is Resource Public Key Infrastructure (RPKI), which contrasts with the very low rate of adoption for almost all other methods.

For research into deployment, there are generally three categories of work. Firstly, there are surveys on the deployment status of concrete solutions. A longitudinal study from 2019 strengthens the notion that RPKI is seeing significant adoption [4]. More recent work by Cloudflare confirms this trend [18].

Secondly, there are proposals for increasing deployment. Researchers recently published a method to increase adoption of RPKI by automating certification of de facto ownership of IP address blocks [10]. Earlier work suggested the creation of economic incentives by industry consortiums or government organizations to motivate individual network operators to adopt secure BGP solutions [7].

The third category concerns analyses of partial deployments. These works quantify the practical security benefits of certain protocols at multiple points throughout their deployment process. A study from 2013 found that many BGP security solutions offer little practical security as long as they are not ubiquitously adopted [14].

The main shortcoming in all of these works, is that they do not focus on the intrinsic adoptability properties of these security proposals, while those might lead to the most effective changes for increasing adoption.

A paper from 2006 argues that it is important to consider the dimension of adoptability in protocol design [3]. However, this work has been published before RPKI existed. Since then, the secure internet routing landscape has changed as a result of RPKI coming onto the stage.

3 PROBLEM DEFINITION

With the abundance of proposed solutions to secure BGP routing, it has only become more clear that none of the approaches have succeeded in reaching widespread adoption yet. Internet routing remains largely insecure to this day. Over time, there have been many issues with deployment processes for these secure solutions. Existing research evaluates the status of deployment and its consequences on security properties. However, little work has been done to reflect on the inherent adoptability of these protocols. Therefore, this paper looks into the metrics that might predict or influence adoptability. Additionally, it will quantify and compare these metrics across multiple of the solutions currently being proposed to secure BGP routing. The intention is to achieve a better understanding of what specific features influence deployment and to cause a shift of perspective for future protocol design, as to make adoptability an integral part of the process.

4 APPROACH

In order to identify and compare adoptability metrics, this paper investigates several sources. The paper will first survey the most relevant proposed approaches to secure BGP routing. The discussion on every one of these respective solutions shall mainly focus on three points. Firstly, it will briefly discuss the historical development and workings of the proposal. Secondly, it overviews the current state of deployment. Lastly, the discussion will identify limiting factors in deployment, and related these to the protocol design.

Based on these individual investigations, this paper continues by consolidating the identified limiting issues for deployment. These are used to distill abstract commonalities that are candidates for adoptability metrics. For each of these, the paper will quantify the extent to which every protocol adheres to it. The results will be presented in table that allows for direct comparison between all protocols. Combining these results with the data on deployment status gives rise to predictive metrics.

The expected result is a better framework for reasoning about properties of security and deployment for current and future secure routing protocols. As part of this result, the paper will have defined clear and unambigious metrics for practicality of approaches.

5 SURVEY OF BGP SECURITY PROPOSALS

This section introduces and describes several BGP security proposals from literature and open standards organizations such as the Internet Engineering Task Force (IETF). The focus will be on the current status of deployment or adoption for each respective approach.

5.1 Route filtering

ASes often use route filtering as a way of mitigating attacks or the effects of misconfigurations [19]. This method is different from other proposed solutions, because it does not include any additional cryptographic procedures or systems that would achieve a theoretical level of security. Instead, route filtering is a truly practical approach to increasing security.

Filtering decisions are frequently made using information published in the Internet Routing Registry (IRR). The IRR consists of multiple databases where network operators submit routing data in an effort to coordinate avoidance of routing problems. None of the records are strictly validated, and many entries might therefore contain errors.

Additionally, network operators consider AS business relations, special-use IP addresses, and length of AS paths when designing filter rules. They often also restrict announcements for networks smaller than /24.

The reason that route filtering is so widespread in use, is because of its simplicity. ASes are not required to modify any of their legacy protocols, nor do they need to replace their hardware or routers to upgrade them with more computational power. This makes route filtering a relatively cheap, albeit incomplete, solution.

However, the sharing of routing information might also be against commercial interests of some ASes. The information might reveal routing or peering agreements they have made with other parties, but which are not meant to be publicly announced for corporate reasons. The sharing of information might additionaly violate privacy requirements and laws. It has also been found that maintaining these routing databases and filter lists, ends up creating significant overhead and thus additional costs. These form disincentives to more effective leverage of the route filtering solution, thus negatively impacting deployment and adoption.

5.2 Resource Public Key Infrastructure (RPKI)

RPKI (RFC 6480) describes an infrastructure to support improved security for internet routing [13]. The RFC was first published in 2012, and many other documents expanding on the system have been published since by the Secure Inter-Domain Routing working group [1].

RPKI's aim is to provide cryptographically verifiable associations between a route and an originating Autonomous System Number (ASN). This would ensure that all advertised routes on the internet are verifiably originating from

[1] https://datatracker.ietf.org/wg/sidr/documents/

the entity that has ownership over the IP prefix, which helps to prevent major BGP hijacking attacks. The foundation is a distributed repository system for storing signed records that bind a route origin to an ASN. These so called Route Origin Authorizations (ROA) authorize a certain AS to originate routes to prefixes within that space. The existing hierarchy of IP address allocation lends itself well for RPKI. At the root resides the Internet Assigned Numbers Authority (IANA), directly below which there are five Regional Internet Registries (RIRs) that span the geopolitical regions of earth; RIPE NCC, LACNIC, ARIN, APNIC, and AFRINIC. Each of the these RIRs acts as a Certificate Authority (CA) or Trust Anchor (TA) in RPKI. This means that they can issue and sign certificates to anyone whom they allocate address space to. All five RIRs provide a method for their members to sign ROA records for IP/ASN pairs. Figure 1 gives an overview.

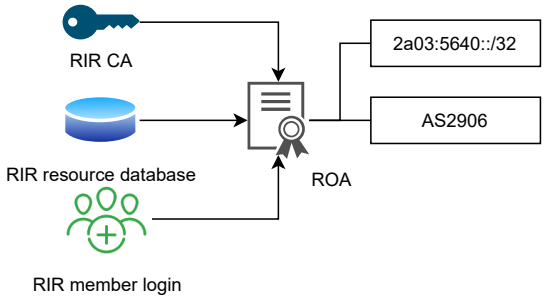

Figure 1: Overview of a ROA in RPKI.

With the routes signed and the ROAs stored with the RIRs, this information can be propagated to any network operator who wants to use it to filter routing. The benefit thus comes from network operators that now have a verifiable way to detect BGP hijacking attacks that originate from an unauthorized AS, and can thus filter out those announcements. This practically limits the spread and therefore the impact of BGP hijacking attacks.

The RPKI trust model requires actors to place trust in the Trust Anchors. So far, these are limited to the five RIRs. Because every network operator already has a relationship with their RIR, the additional trust placed with the RIRs in RPKI seems reasonable. However, this now does provide RIRs with the power to take down prefixes, by removing ROAs or even revoking certificates. Researchers recently proposed a system to limit this power by requiring joint coordination by multiple RIRs before such action could be taken [20].

RPKI does not require any changes to BGP message formats. Nor does it require any cryptographic operations to be performed on all routers on the network. The IETF has standardized a protocol to deliver RPKI prefix origin data to routers as RFC 8210 called the RPKI to Router (RTR) protocol [2]. All the cryptographic validation is done on separate servers from the network operator.

Deployment of RPKI, which started in 2011, has been slow, but saw a significant increase in recent years [4][8][18]. From data by Cloudflare, 30% of ASes in their *Is BGP safe yet?* project deployed RPKI, by either signing their routes, filtering based on ROAs, or both [5].

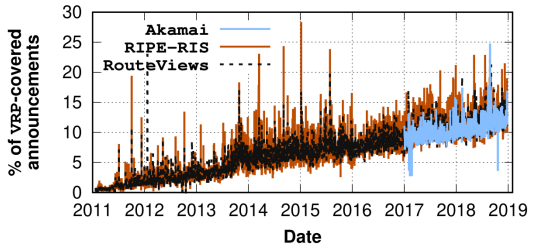

Figure 2: The percentage of BGP announcements covered by Validated ROA Payloads (VRPs). From [4].

Figure 2 shows a consistent and steady increase in the number of verifiable BGP announcements being made on the internet. This proves that current RPKI adoption has become significant.

However, this development was slow, and there still remain issues that hinder further deployment. Deployment was found to be hindered by human error in configuration of RPKI, which lead to mistrust in the RPKI infrastructure among network operators [6]. Additionally, RPKI too suffered from the chicken-and-egg problem. The implementation of RPKI certification and route-origin validation (ROV) based filtering require non-trivial amounts of work from network operators. However, there is little incentive for network operators to start this process as long as RPKI is not widely deployed yet. Solutions such as DISCO try to sidestep these issues by automating certification, the population of public repositories, and the generation of route filtering rules for ROV [10]. The automation of this process also resolves parts of the human error prone configuration. DISCO has only recentely been published, so there has not been large scale use of it yet.

5.3 BGPsec

The aim of BGPsec is to provide path validation; security for the path of ASes through which BGP announcements pass. That is, BGPsec provides cryptographic assurance that every AS on the path has explicitly authorized the advertisement of the route to the subsequent AS in the path. This is in contrast with RPKI, which only provides path origin validation.

BGPsec is being standardized by the IETF as RFC 8205 [21]. BGPsec leverages RPKI and extends it by including signatures to BGP messages. It requires each AS on the path to sign its BGP messages using their certificate as distributed under the RPKI. Upon receiving a BGPsec announcment, an AS validates all signatures included by all ASes on the path, and filters the route if any of them are invalid. This prevents AS path forgery and interception attacks [15].

However, the approach is such that the receiving AS only learns a path via BGPsec if every AS on the path has adopted

BGPsec. This is clearly necessary to guarantee path validation, but it makes adoption much harder. In order to guarantee backwards compatability while deployment would be ongoing, ASes want to continue supporting nodes in the network that have not adopted BGPsec yet. This effectively undermines the most important security guarantee of BGPsec, because an attacker can still always pretend they do not know BGPsec, thereby resetting the security situation to regular BGP again. Effective deployment of BGPsec thus becomes an "all-or-nothing" case, combined with the chicken-and-egg problem. Researchers additionally showed that partial adoption of BGPsec has marginal benefits, and in some cases might even create new vulnerabilities [14].

Furthermore, routers in networks using BGPsec must now cryptographically sign and verify all BGP messages they send. These are computationally expensive operations for which many routers are not sufficiently equipped. To deploy BGPsec, network operators would thus also have to replace large parts of their hardware infrastructure.

5.4 S*BGP

Researchers have previously proposed other methods for secure path validation for BGP besides BGPsec, including S-BGP [12], soBGP [16], and psBGP[22]. This set of protocols is generally referred to as S*BGP.

All three introduce a PKI and use multiple different certificates, provided in-band or out-of-band. S-BGP has large costs of operation, similar to BGPsec, which makes it unattractive for adoption. psBGP, similar to S-BGP, uses signatures to perform AS path validation, but incorporates a method based on expressed trustworthiness between ASes to decide whether it should validate all signatures. This reduces the computational overhead slightly, but also decreases the achieved security. soBGP does not mandate clear choices for implementation, thus interoperability can not be ensured. For example, it allows distribution of signed data either via repositories, or in-band using BGP messages. Additionally, it allows for computation of authorized routes by routers itself, or by a NOC that distributes results to routers at unspecified intervals.

Even though S-BGP, soBGP, and psBGP have been proposed twenty years ago, all three have seen practically no adoption, because of high cost of deployment and ambiguity in specifications. Neither have there been any production-ready software implementations for these protocols.

5.5 Secure Multi-Party Computation (SMPC)

SMPC introduces a separate set of computational servers to which route computation is outsourced [9]. This additionally preserves the privacy of ASes' routing policies. However, others have questioned the scalability and computational overhead of this approach [15]. The predefined set of servers in particular require high processing time, as they simulate BGP by using the input data from the ASes.

SMPC's trust model now includes a set of computational servers.

5.6 Symmetric Key approach

Because of the inherently high cost of full path validation using asymmetric cryptography, researchers have also proposed solutions based on an approach with symmetric cryptography. An example is Secure Path Vector (SPV) routing, which uses a sequence of one-time offline signatures [11], and which is improved upon by work using keychain-based signatures [23].

Another solution based on symmetric keys was proposed by Bruhadeshwar et al. in 2011 [1]. The threat model in this solution explicitly trusts one BGP router in any path of a certain length. It performs more efficiently than earlier methods like SPV, because after the trusted router has verified an UPDATE message, the signing material from the previous BGP routers are no longer required to be forwarded to successors in the path.

The main reason that holds back adoption of any of these schemes is the lack of long-term security of the symmetric keys, because of their inherent vulnerability to brute-force attacks [15].

6 DEPLOYMENT ISSUES

This section briefly discusses two recurring issues with deployment for all respective protocols.

Individual entities in the internet routing domain act according to local business objectives, and those do not provide incentives to transition their deployments to support any of the secure BGP proposals [7].

Furthermore, earlier research has shown that partial deployment of secure path validation solutions for BGP (i.e., S-BGP, soBGP, BGPsec, etc.) provides marginal security value [14]. This enforces the meagre set of options for deployment to just "system-wide change overnight" before any network operator would even consider putting in the effort.

7 METRICS FOR ADOPTION

This section presents metrics for adoptability of proposed solutions to secure BGP routing, distilled from the literature on deployment status and deployment limitations. It additionally gives a quantification of the extent to which the discussed BGP security proposals fulfill those metrics. These results combined with the deployment data, give an indication of adoptability metrics that are most effective predictors.

The main finding is paradoxical and non-deterministic. The metric with the highest predictive value for adoptability, is the factor of adoption itself. This paradoxical metric has also been referred to as the circular dependency and the chicken-and-egg problem to internet routing by other researchers. There does not seem to be a protocol property that objectively determines whether a particular methods will be adopted or not. Instead, a method requires industry-wide backing of router vendors, network operators, and standard bodies. Solely evaluating security properties or security guarantees does not seem to predict adoption. For example, the S-BGP protocol generally provides the same path validation as BGPsec, but there is widespread consensus now that BGPsec should eventually be adopted, and not S-BGP, as

Metric	Route filtering	RPKI	BGPsec	S*BGP	SMPC	SPV
Hardware cost	■	■	-	-	□	□
Interoperability	■	■	-	-	-	-
Incremental deployment	■	□	-	-	-	□
Implementation diversity	■	□	-	-	-	-

Table 1: Quantification of adoptability metrics for secure BGP routing proposals. Evaluation legend: ■ good, □ okay, - bad.

proven by the effort put in to the IETF working groups that standardize the BGPsec protocol and ecosystem.

There are however several derivative metrics that stem from the paradoxical metric. The following two metrics are proxies to the paradox, but more practical to take into account when setting out to increase deployment.

Interoperability refers to how well understood the protocol specification is by every entity involved and whether the specification is unambigious enough to ensure that different implementations will still be able to communicate with each other. Network operators said to be hesitant of adopting any new procedure if it was not absolutely clear that it would not negatively impact routing performance. The status of standardization is a good indicator for quality of interoperability. The fact that some protocols — RPKI and BGPsec — are being formally standardized means that representatives from industry, network operators, RIRs, and others are writing the specification together, thereby agreeing on how it should operate. Naturally, interoperability only happens once a large number of different parties agree on the specific operation of a protocol. This will thus only significantly occur when the protocol has already seen non-trivial adoption.

Implementation diversity is considered to be important before a network operator or AS decides to invest their resources in the process of deploying a certain protocol. Only a single implementation does not signal strength of the protocol ecosystem. This metric also derives from the adoption paradox, because as adoption increases, an increasing number of parties will start to develop their own implementations, thereby contributing to the implementation diversity.

The following paragraphs identify two additional metrics that are less intertwined with the adoption paradox metric.

Hardware cost is an important factor in the decision of network operators. A network operator has invested in many hardware routers to build up their AS. In general, these routers are not well-equipped for computationally intensive cryptographic operations. Thus, protocols such as BGPsec and protocols from the S*BGP set become less likely to be adopted by network operators, because they require also the routers in the network to calculate and verify signatures. Before network operators could deploy these respective protocols, they would have to make significant investments to upgrade large parts of their network infrastructure such that it can perform these cryptographic operations adequately fast as to not impact routing performance too much.

Incremental deployment is required for the process of internet-wide deployment to be practical. Because of the distributed nature of the routing infrastructure, enforcing changes across the entire internet is infeasible, as it would involve coordinating action between thousands of individually operating parties. Thus, for a protocol to be adopted, it needs to be able to communicate with legacy parts of the network. An extra dimension to this metric comes in the form of incremental value with deployment. As observed by earlier work [14], protocols such as BGPsec and S*BGP provide no marginal value to any network operator that deploys them. Thus, there is no incentive for any individually operating entity to commence the deployment. If a protocol were able to actually provide marginal value for an individual operator, it would make deployment more attractive.

These four metrics form a basis for a perspective switch to adoptability as an inherent design property of protocols. In order to now assess how relevant these are, Table 1 lists the four metrics discussed, and quantifies the extent to which the listed proposals to secure BGP routing fulfill them.

Clearly, route filtering is predicted to be most practical for adoption. This corresponds to the real-world situation, in which practically all network operators perform some form of filtering on their routers as a way of securing their networks and cleaning up their databases from bogus advertisements.

RPKI follows in similar vain. It performs well on the hardware cost metric, because RPKI does not require any additional computations to be performed on the routers themselves. Thus, network operators do not have to replace large parts of their networks. Incremental deployment of RPKI is favourable, because the effort of signing ROAs that a network operator puts in, translate to improved security when all other ASes that perform filtering on ROAs are now better suited to recognize hijacking attacks aimed at the original network operator's IP prefixes.

Most other protocols are observed to perform badly on all metrics. This corresponds to their level of global deployment, which is often practically zero.

The quantification of metrics showed that, in general, cheaper and easier to deploy protocols have a higher chance of wide adoption. Despite many parties stating that the security of full path validation is needed on the internet, none of them have taken considerable steps to actually deploy e.g. BGPsec yet. The world remains stuck with partial deployments of protocols that provide limited protection against the many attacks that plague the internet routing system.

8 CONCLUSION

This paper surveyed proposals to secure BGP routing, and identified the deployment issues that often accompany them. Despite the amount of research done into deployment status and solutions for increasing deployment, there appeared to be a lack of work on adoptability as a protocol design property itself. This paper presented a list of metrics that predict adoptabilty of secure BGP routing protocols. The main finding was that the most significant predictor of adoption, is the factor of adoption itself. Metrics that derived from this, are interoperability and implementation diversity. Two other metrics that were found to predict adoption are hardware cost and incremental deployment value. It is recommended that future protocol design processes consider adoptability as an important factor. Future work should investigate additional metrics that predict adoptability, and research how they can be more effectively leveraged to ultimately provide secure internet routing.

REFERENCES

[1] B. Bruhadeshwar, S. S. Kulkarni, and A. X. Liu. 2011. Symmetric Key Approaches to Securing BGP—A Little Bit Trust Is Enough. *IEEE Transactions on Parallel and Distributed Systems* 22, 9 (Sept. 2011), 1536–1549. https://doi.org/10.1109/TPDS.2011.19

[2] Randy Bush and Rob Austein. [n.d.]. The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1. https://tools.ietf.org/html/rfc8210

[3] Haowen Chan, Debabrata Dash, Adrian Perrig, and Hui Zhang. 2006. Modeling Adoptability of Secure BGP Protocols. *ACM SIGCOMM Computer Communication Review* 36, 4 (2006), 279–290.

[4] Taejoong Chung, Emile Aben, Tim Bruijnzeels, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, Roland van Rijswijk-Deij, John Rula, and Nick Sullivan. 2019. RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins. In *Proceedings of the Internet Measurement Conference*. ACM, Amsterdam Netherlands, 406–419. https://doi.org/10.1145/3355369.3355596

[5] Cloudflare. [n.d.]. Is BGP safe yet? · Cloudflare. https://isbgpsafeyet.com

[6] Yossi Gilad, Avichai Cohen, Amir Herzberg, Michael Schapira, and Haya Shulman. 2017. Are We There Yet? On RPKI's Deployment and Security. In *Proceedings 2017 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA. https://doi.org/10.14722/ndss.2017.23123

[7] Phillipa Gill, Michael Schapira, and Sharon Goldberg. 2011. Let the market drive deployment: a strategy for transitioning to BGP security. In *Proceedings of the ACM SIGCOMM 2011 conference*. ACM, 14–25. https://doi.org/10.1145/2018436.2018439

[8] Sharon Goldberg. 2014. Why is it taking so long to secure internet routing? *Commun. ACM* 57, 10 (Sept. 2014), 56–63. https://doi.org/10.1145/2659899

[9] Debayan Gupta, Aaron Segal, Aurojit Panda, Gil Segev, Michael Schapira, Joan Feigenbaum, Jenifer Rexford, and Scott Shenker. 2012. A new approach to interdomain routing based on secure multi-party computation. In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks - HotNets-XI*. ACM Press, Redmond, Washington, 37–42. https://doi.org/10.1145/2390231.2390238

[10] Tomas Hlavacek, Italo Cunha, Yossi Gilad, Amir Herzberg, Ethan Katz-Bassett, Michael Schapira, and Haya Shulman. 2020. DISCO: Sidestepping RPKI's Deployment Barriers. In *Proceedings 2020 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA. https://doi.org/10.14722/ndss.2020.24355

[11] Yih-Chun Hu, Adrian Perrig, and Marvin Sirbu. 2004. SPV: Secure Path Vector Routing for Securing BGP. In *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 179–192. https://doi.org/10.1145/1015467.1015488

[12] S. Kent, C. Lynn, and K. Seo. 2000. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected Areas in Communications* 18, 4 (April 2000), 582–592. https://doi.org/10.1109/49.839934

[13] Matt Lepinski and Stephen Kent. [n.d.]. An Infrastructure to Support Secure Internet Routing. https://tools.ietf.org/html/rfc6480

[14] Robert Lychev, Sharon Goldberg, and Michael Schapira. 2013. BGP security in partial deployment: is the juice worth the squeeze?. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. ACM, 171–182. https://doi.org/10.1145/2486001.2486010

[15] Asya Mitseva, Andriy Panchenko, and Thomas Engel. 2018. The state of affairs in BGP security: A survey of attacks and defenses. *Computer Communications* 124 (June 2018), 45–60. https://doi.org/10.1016/j.comcom.2018.04.013

[16] James Ng. [n.d.]. Extensions to BGP to Support Secure Origin BGP (soBGP). https://tools.ietf.org/html/draft-ng-sobgp-bgp-extensions-01

[17] Adrian Perrig, Pawel Szalachowski, Raphael M. Reischuk, and Laurent Chuat. 2017. *SCION: A Secure Internet Architecture*. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-67080-5

[18] Louis Poinsignon. 2020. The Internet is Getting Safer: Fall 2020 RPKI Update. https://blog.cloudflare.com/rpki-2020-fall-update/

[19] Pavlos Sermpezis, Vasileios Kotronis, Alberto Dainotti, and Xenofontas Dimitropoulos. 2018. A Survey among Network Operators on BGP Prefix Hijacking. *ACM SIGCOMM Computer Communication Review* 48, 1 (April 2018), 64–69. https://doi.org/10.1145/3211852.3211862

[20] Kris Shrishak and Haya Shulman. 2020. Limiting the Power of RPKI Authorities. In *Proceedings of the Applied Networking Research Workshop on ZZZ*. ACM, Virtual Event Spain, 12–18. https://doi.org/10.1145/3404868.3406674

[21] Kotikalapudi Sriram and Matthew Lepinski. [n.d.]. BGPsec Protocol Specification. https://tools.ietf.org/html/rfc8205.html

[22] Tao Wan, Evangelos Kranakis, and P C van Oorschot. 2005. Pretty Secure BGP (psBGP). 16.

[23] H. Yin, B. Sheng, H. Wang, and J. Pan. 2010. Keychain-Based Signatures for Securing BGP. *IEEE Journal on Selected Areas in Communications* 28, 8 (Oct. 2010), 1308–1318. https://doi.org/10.1109/JSAC.2010.101008