Securing BGP - Mechanisms to Prevent Routing Leaks

Leon Spörl, Michael Oberrauch*
*Chair of Network Architectures and Services
School of Computation, Information and Technology, Technical University of Munich, Germany
Email: leon.spoerl@tum.de, oberrauc@net.in.tum.de

***Abstract*—The Border Gateway Protocol (BGP) is today's prevailing interdomain routing protocol, and due to its widespread adoption, it is likely to remain so for the foreseeable future. BGP was designed with scalability and efficiency in mind but lacks fundamental security features. Routing leaks caused by misconfiguration or malicious actions are among the main vulnerabilities. We evaluate the leading security mechanisms IRR, RPKI, and BGPsec with respect to performance, scalability, adoption, and their ability to fulfill specific security objectives. Often, inconsistent adoption is the limiting factor of the described mechanisms as security features only take full effect if rules are uniformly applied and strictly enforced. We conclude by reviewing promising new advancements regarding the implementation of ASPA and outline potential future developments.**

***Index Terms*—Autonomous System Provider Authorization (ASPA), Border Gateway Protocol (BGP), Border Gateway Protocol Security (BGPsec), Internet Routing Registry (IRR), Resource Public Key Infrastructure (RPKI)**

1. Introduction

Since its introduction over 30 years ago, the Border Gateway Protocol has become the global de facto standard of inter-AS routing. Despite its excellent routing capabilities, it is insecure by design.

One of the key problems of BGP is that an AS has no control over its resources, i.e., how the AS number and IP prefixes that have been uniquely assigned by the Regional Internet Registry are announced and propagated. This lack of control allows malicious or misconfigured networks to distribute incorrect routing statements. Routing leaks are a common type of security violation caused by erroneous route propagation. Minor routing errors can affect large regions, with even simple attacks posing a major threat to the Internet's backbone. There have been numerous incidents, one of them being an accidental route leak induced by Google in August 2017, which caused massive Internet disruptions in Japan [1].

Incidents like this stress the need for globally deployed BGP security mechanisms. In this paper, we analyze the most relevant innovations that allow us to prevent routing leaks and other BGP-related attacks. We begin by reviewing related work in this research area, followed by some background information and definitions. We then examine popular security approaches and end with an outlook on future developments.

2. Related work

In the early days of BGP, most research papers only proposed new solutions or evaluated a single security proposal [2], [3]. A few years later, the first systematic reviews were published. As BGP security is a quickly evolving field of research, the findings of literature surveys from around 2010 are outdated and do not address most of today's security mechanisms [4]–[6]. A reasonable recent survey by Mitseva et al. [1] provides an extensive overview of BGP security properties and solution approaches.

Regarding the individual security mechanisms, recent research includes a quantitative analysis of RPKI deployment by Chung et al. [7] and a survey by Rodday et al. [8] summarizing RPKI-related studies, challenges and solution proposals. Research by Du et al. [9] highlights vulnerabilities of the Internet Routing Registry, by quantifying and analyzing irregular IRR records. The challenges of BGPsec deployment and optimization approaches are summarized by Abdelhafez and Fadlalla [10].

3. Background

Before discussing the attributes of BGP and the need for specific security measures, we have to take a look at the general structure of the Internet and its historical background.

3.1. Autonomous Systems

The Internet consists of many interconnected, independently administered networks, so-called autonomous systems (ASes). From a technical perspective, an autonomous system is a set of routers that follow a uniform routing plan to allow for intra- and inter-AS network communication. [11] An AS is managed by a single organization and is identified by a unique 16 or 32-bit number. The AS numbers (ASNs) and IP address spaces are assigned by the Regional Internet Registries (RIRs) to the respective autonomous systems. [1], [12] Depending on the geographical location, a different RIR applies [13]:

- **AfriNIC**: Africa
- **APNIC**: parts of Asia, Pacific Region
- **ARIN**: North America, parts of the Caribbean
- **LACNIC**: Latin America, parts of the Caribbean
- **RIPE NCC**: Europe, Middle East, parts of Asia

A variety of interior gateway protocols are used for routing within an AS, i.e., intradomain routing [14]. Some

 doi: 10.2313/NET-2025-11-3_19

organizations may even run multiple routing protocols in parallel or use proprietary standards. While intradomain routing is transparent to other ASes and does not have to follow a common standard, interdomain routing requires an operational standard so all autonomous systems can span a global network.

3.2. The Border Gateway Protocol

The Border Gateway Protocol (BGP) was first introduced in 1989 [15], specified in RFC 1105, and replaced the Exterior Gateway Protocol (EGP) [16]. The most recent version, BGP-4 [11], is today's de facto standard for global inter-AS routing [1]. In the early days of the Internet, there were also other interdomain routing protocols proposed, e.g. the OSI Inter-Domain Routing Protocol (IDRP) [17], which was expected to replace BGP but has no relevance today [18], [19].

The BGP specification describes how autonomous systems can exchange reachability information and announce IP prefixes to each other [11]. BGP is a path-vector protocol, meaning that not only the distance but the whole path to a destination is announced. Therefore, each AS adds itself to the path before propagating the route to others. [5] Using that information, every AS can maintain its own reachability graph, make informed routing decisions, and derive a routing table. Topology changes are propagated within minutes, making BGP a dynamic and resilient protocol. An AS can take different metrics into account when making routing decisions, the most relevant ones being path and prefix length. The shortest path method is used to minimize the number of hops between source and destination. If announced IP address ranges overlap, the announcement with the more specific IP prefix is usually prioritized. [12]

Business agreements and financial considerations also play a role when selecting paths and propagating routes to other ASes. Two connected autonomous systems can have a customer-provider relationship, i.e., the customer pays the provider for the traffic routed over the path or a peer-to-peer relationship where traffic is routed over the link free of charge. [20] If each AS tries to maximize its revenue and only propagates routes that incur financial gain, we speak of valley-free routing. After a packet has traversed a provider-to-customer or peering link, it is only allowed to take additional provider-to-customer links and must not again move upwards in the hierarchical BGP model. Valley routes may lead to increased path lengths and ASes unintentionally transiting traffic [21].

3.3. BGP Route Leaks

Although the term route leak already appears in literature from the early 2000s, it lacked a uniform definition for a long time [22]. RFC 7908, published in 2016, defines a route leak as "the propagation of routing announcement(s) beyond their intended scope" [23]. A route leak occurs when an AS propagates a learned route to another AS and thereby violates the policies of an AS alongside the resulting path. More specifically, policies typically refer to the business relations described in Section 3.2. A policy violation can be the infringement of the valley-free property [1].

RFC 7908 provides a taxonomy of route leaks based on observed incidents and differentiates between six types. The first five depict distinct scenarios where learned BGP routes from non-customer ASes are announced to provider or peering ASes. Figure 1 illustrates a type 1 route leak where an AS receives a route announcement from a provider AS and propagates it to another provider AS. The leaked route is probably preferred by AS 3 because customer ASes are more highly prioritized than peering ASes, resulting in a hairpin turn at AS 2.

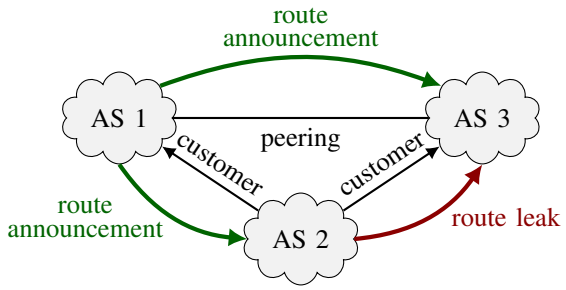

Figure 1: BGP route leak type 1

The remaining route leak category, type 6, refers to an issue that can occur if BGP is used as a routing protocol inside an AS. In this case, internal and external BGP must be strictly separated. Misconfiguration can lead to internal routes being exported to the outside and announced to other autonomous systems. These internal prefixes may be more specific than externally routed ones and are therefore preferred by other ASes.

This behavior is similar to a route hijack, also known as prefix hijack, which occurs when an AS originates a prefix it is not the holder of. If the address space belongs to another AS, this can have severe consequences as the illegitimately announced route may be selected by other ASes across the Internet due to a shorter path or a more specific prefix. [1], [5]

The distinction between BGP route leaks and hijacks is not very sharp in the literature. While Wijchers and Overeinder [24] draw a clear line between the terms, the RFC on route leak classification cites resources about prefix hijacks and lists hijacking incidents as examples for route leaks [23]. Both incidents can happen by accidental misconfiguration or with malicious intent [5], [23].

4. Security Mechanisms

BGP uses the Transmission Control Protocol (TCP) [25], which offers some basic error correction and retransmission mechanisms [11]. However, BGP provides no confidentiality or integrity protection for its messages [5]. By default, route advertisements are not authenticated, so autonomous systems can announce arbitrary prefixes, no matter if they are the legitimate holders [12]. There are various approaches to make BGP more secure and reliable. We will analyze the most popular ones regarding their security properties, performance, practicability, and adoption.

4.1. Internet Routing Registry

The Internet Routing Registry (IRR), proposed in 1995, was one of the first attempts to make BGP routing

 doi: 10.2313/NET-2025-11-3_19

more transparent and secure [26]. The IRR is a set of public databases where ASes can upload routing information on prefixes they hold using the Routing Policy Specification Language (RPSL) [27]. Other ASes can then fetch the data and create BGP route filters [1].

As of May 2025, there exist 18 IRR databases[1] hosted by companies as well as Regional and National Internet Registries. Many of the databases mirror other IRRs to provide exhaustive data sets. The different IRR providers do not use standardized authentication and validation mechanisms, resulting in varying quality of database records [9]. RIRs can authenticate their members and approve ownership of resources before publishing routing information for a certain address space. Other routing registries lack the ability to perform this kind of authorization but may publish routing information on an address space anyway. [12]

Malicious actors exploited these circumstances in the past. Du et al. [9] describe a case where attackers were able to hijack a prefix belonging to Amazon and added a manipulated route object to the ALTDB routing registry. They were thereby able to reroute customers of a company that used Amazon's cloud resources to a phishing page and stole cryptocurrency worth $235,000. This is not an isolated incident. Over the period of 1.5 years, Du et al. analyzed RADB, which is the largest IRR database holding over 1.5 million route objects. By investigating consistency across IRR databases and checking if database entries match BGP announcements, they detected 34,199 irregular entries and classified 6,373 of them as suspicious.

Proposals to improve IRR's security did not resolve all vulnerabilities or were simply not implemented on a broad scale [1], [28], [29]. IRR's main issue is the lack of global uniformity, making route validation unreliable and error-prone. While it was a great advance in terms of routing transparency in the 90s, it does not meet the expectations of BGP security today. [12]

4.2. Resource Public Key Infrastructure

The Resource Public Key Infrastructure (RPKI) [30] is an out-of-band system that enables resource holders to cryptographically prove their identity and digitally sign statements on the intended use of their resources. The infrastructure uses X.509 certificates [31] with two extensions that bind AS numbers and IP address prefixes to the holder of the certificate's private key [32]. The hierarchical structure of RPKI is based on five trust anchors, with each Regional Internet Registry operating its own root Certificate Authority (CA). AS administrators can request resource certificates over the member portal of their respective RIR. This procedure is called hosted RPKI, as certificate creation, publication, and key rollover are all done by the RIR. Although this process is relatively simple and convenient for most members, it is not sufficient for all of them. Some organizations must host their own child CA in order to delegate RPKI administration to their customers. This operating mode is called delegated RPKI and is supported by all RIRs except for AFRINIC (as of May 2025). [7]

Using their resource certificate, autonomous systems can sign Route Origin Authorizations (ROAs) [33]. A ROA authorizes an AS to originate a certain IP address space. Apart from the prefix and ASN, a ROA contains the maximum prefix length the specified AS is allowed to advertise. The issued ROAs are typically published by the RIRs. Other ASes can then fetch the data sets, verify the ROA signatures and derive validated ROA payloads (VRPs). The verification of incoming BGP announcements, known as route origin validation (ROV) [34], can return three different results:

- **Valid**: announcement covered by a VRP
- **Invalid**: announcement from unauthorized AS or announcement violates the VRP's maximum prefix length attribute
- **NotFound**: announced prefix not (fully) covered by any VRP

Given that RPKI is still far from universal adoption, announcements returning the *NotFound* status should be accepted for now. [7], [12]

The deployment of RPKI started in 2011 and was impaired by many configuration errors in the early adoption stages. Chung et al. [7] evaluated RPKI statistics that were collected from 2011 to 2019. Throughout 2011, they observed 48.92% of VRP-covered announcements to be invalid. This value drastically improved over the years, settling between 2% and 5% in 2019. During the whole measurement period, around half of the invalid announcements were caused by too specific prefixes. Many of the published ROAs were missing the MaxLength attribute, indicating misconfiguration as a cause. In recent years, the general adoption has improved.

The ratio of unique IPv4 prefix-origin pairs (combinations of IP prefix and origin AS number) covered by VRPs has increased from 12% in early 2019 to over 56% in May 2025 according to the NIST RPKI Monitor[2]. The IPv6 coverage shows similar numbers, recording almost 58% validatable IPv4 prefix-origin pairs in May 2025. While there is less historic data available for IPv6, the observable trend is similar to IPv4 [7].

RPKI only takes full effect if all ASes systematically drop invalid announcements. Some major service providers, such as AT&T, are already enforcing this policy. [7] RPKI's route origin validation prevents prefix hijacks but does not target route leaks as defined in Section 3.3. An AS can restrict route origination for its resources, but the path an announcement takes is unknown [8].

4.3. Border Gateway Protocol Security

Border Gateway Protocol Security, better known as BGPsec [35], is a BGP extension that relies on RPKI's resource certificates to provide cryptographic path validation. BGPsec replaces the AS_PATH attribute with a BGPsec_PATH attribute inside the BGP update messages. The new attribute contains digitally signed path information. Each AS that an announcement traverses signs the previous path, its own AS number, and the number of the AS the announcement will be propagated to. Every router

1. https://irr.net/registry

2. https://rpki-monitor.antd.nist.gov

 doi: 10.2313/NET-2025-11-3_19

in the announcement chain verifies the signatures and can detect path forgery. [10]

The main downside of BGPsec is its poor adoption due to the high entry barrier. Most ASes would need to replace their hardware to support BGPsec, without deriving a direct benefit from the investment [1]. If one router along the announcement path has no BGPsec capabilities, the whole path immediately becomes invalid. There is no *NotFound* state like in RPKI, so partial BGPsec deployment has little utility. [36] Aside from that, the extension generates significant computational overhead slowing down the BGP operations by a factor of 70, as shown by Kim and Kim [37]. There have been optimization efforts to reduce CPU overhead and memory consumption, but no optimization algorithm could be implemented at scale yet. While there is some room for improvement, well-performing BGPsec requires specialized hardware accelerators. [10]

4.4. Autonomous System Provider Authorization

Autonomous System Provider Authorization (ASPA) [38] is a draft that has been discussed by the Internet Engineering Task Force (IETF) over the last years. ASPA objects are part of the RPKI, similar to ROAs [7]. By generating and signing an ASPA object, an AS can specify a list of provider ASes. Using these statements, each hop along a path can be identified as

- **Provider**,
- **Not Provider**, or
- **No Attestation** if no ASPA can be retrieved from the customer.

Given this information, we can make plausibility checks and detect implausible paths potentially caused by route leaks. ASPA explicitly supports incremental deployment. If some ASes make no attestation, the plausibility check succeeds if no other policy violations are detected, no matter how sparsely the feature is deployed. ASPA is capable of detecting (accidental) route leaks and even protects against some forms of prefix hijacks. The wider the adoption of ASPA, the harder it becomes for attackers to perform unrecognized route hijacks. However, Route Origin Authorization is still the most reliable method of hijacking prevention. ASPA's deployment model is very similar to ROAs. Cryptographic operations are handled by certificate authorities, while routers can handle verification without much computational overhead.

RIPE officials stressed ASPA's significance at the "RIPE 90" meeting in May 2025. It is yet to be standardized and is planned to be implemented in RIPE's hosted RPKI in summer 2025. ARIN and APNIC also intend to test their implementations in 2025. [39]

4.5. Other Approaches

Apart from the already presented security mechanisms, there were many other proposals in the past, most of them having no relevance nowadays. One of the first approaches was Secure-BGP (S-BGP) [2], which is conceptually very similar to RPKI. It uses an out-of-band public key infrastructure to authenticate prefix announcements and provide route origin validation. Since S-BGP's adoption was obstructed by performance issues, Secure Origin BGP (soBGP) [3] was developed at Cisco Systems in 2003. It also relies on a PKI and is characterized by a lower-overhead implementation, but uses proprietary certificates and, hence, lacks interoperability.

Further proposals were made in the following years, but none of them was realized at a large scale until RPKI was standardized in 2012 [1]. More recent research by Hari and Lakshman [40] suggests blockchain-based BGP security, omitting the PKI's central root of trust. However, the scalability and performance of a blockchain-based solution have not yet been demonstrated with a practicable implementation.

5. Conclusion and Outlook

When assessing the features, performance, and adoption of BGP security mechanisms, we can conclude that efficient solutions have already been found, but advancements are impeded by a heterogenous system topography and rigid legacy systems.

The resource public key infrastructure provides a solid base that allows developments like Route Origin Authorization and Autonomous System Provider Authorization to build upon. BGPsec is a proven solution for path validation, which was well integrated into BGP and RPKI but is constrained by computational overhead. This substantial entry barrier slows down the adoption process. The most important advancement in the upcoming years will be the further increase of RPKI adoption. The objective should be a full coverage allowing for strict ROA policy enforcement without locking out autonomous systems from global routing. Even though ASPA is still an IETF draft, it can be seen as a promising technique. As soon as the standardization is finished, it will profit from a low entry barrier, given the persisting RPKI and ASPA's excellent support for incremental deployment.

References

[1] A. Mitseva, A. Panchenko, and T. Engel, "The state of affairs in BGP security: A survey of attacks and defenses," *Computer Communications*, vol. 124, pp. 45–60, 2018. [Online]. Available: https://doi.org/10.1016/j.comcom.2018.04.013

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

[3] R. White, "Securing BGP Through Secure Origin BGP," 2003.

[4] M. O. Nicholes and B. Mukherjee, "A survey of security techniques for the border gateway protocol (BGP)," *IEEE Communications Surveys and Tutorials*, vol. 11, no. 1, pp. 52–65, 2009.

[5] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, 2010. [Online]. Available: https://dx.doi.org/10.1109/JPROC.2009.2034031

[6] G. Huston, M. Rossi, and G. Armitage, "Securing BGP — A Literature Survey," *IEEE Communications Surveys and Tutorials*, vol. 13, no. 2, pp. 199–222, 2011.

[7] T. Chung, E. Aben, T. Bruijnzeels, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, R. v. Rijswijk-Deij, J. Rula, and N. Sullivan, "RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins," in *Proceedings of the Internet Measurement Conference*, ser. IMC '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 406–419. [Online]. Available: https://doi.org/10.1145/3355369.3355596

[8] N. Rodday, I. Cunha, R. Bush, E. Katz-Bassett, G. D. Rodosek, T. C. Schmidt, and M. Wählisch, "The Resource Public Key Infrastructure (RPKI): A Survey on Measurements and Future Prospects," *IEEE Transactions on Network and Service Management*, vol. 21, no. 2, pp. 2353–2373, 2024. [Online]. Available: https://doi.org/10.1109/TNSM.2023.3327455

[9] B. Du, K. Izhikevich, S. Rao, G. Akiwate, C. Testart, A. C. Snoeren, and k. claffy, "IRRegularities in the Internet Routing Registry," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, ser. IMC '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 104–110. [Online]. Available: https://doi.org/10.1145/3618257.3624843

[10] M. Abdelhafez and Y. Fadlalla, "BGPsec Deployment Challenges and Optimization Efforts," in *2024 IEEE 22nd Student Conference on Research and Development (SCOReD)*, 2024, pp. 277–281. [Online]. Available: https://doi.org/10.1109/SCOReD64708.2024.10872667

[11] Y. Rekhter, S. Hares, and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, Jan. 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4271

[12] A. Band, "RPKI Documentation," https://rpki.readthedocs.io, 2018, [Online; accessed 13-May-2025].

[13] R. Housley, J. Curran, G. Huston, and D. R. Conrad, "The Internet Numbers Registry System," RFC 7020, Aug. 2013. [Online]. Available: https://www.rfc-editor.org/info/rfc7020

[14] M. Athira, L. Abrahami, and R. G. Sangeetha, "Study on network performance of interior gateway protocols — RIP, EIGRP and OSPF," in *2017 International Conference on Nextgen Electronic Technologies: Silicon to Software (ICNETS2)*, 2017, pp. 344–348.

[15] K. Lougheed and Y. Rekhter, "Border Gateway Protocol (BGP)," RFC 1105, Jun. 1989. [Online]. Available: https://www.rfc-editor.org/info/rfc1105

[16] D. L. Mills, "Exterior Gateway Protocol formal specification," RFC 904, Apr. 1984. [Online]. Available: https://www.rfc-editor.org/info/rfc904

[17] C. Kunzinger, "OSI INTER-DOMAIN ROUTING PROTOCOL (IDRP)," Internet Engineering Task Force, Internet-Draft draft-kunzinger-idrp-ISO10747-01, Nov. 1994. [Online]. Available: https://datatracker.ietf.org/doc/draft-kunzinger-idrp-ISO10747/01/

[18] J. A. Hawkinson and T. J. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)," RFC 1930, Mar. 1996. [Online]. Available: https://www.rfc-editor.org/info/rfc1930

[19] K. Butler, T. Farley, P. McDaniel, and J. Rexford, "A survey of BGP security," *ACM*, vol. 5, pp. 1–35, 2004.

[20] L. Gao and J. Rexford, "Stable Internet routing without global coordination," in *Proceedings of the 2000 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '00. New York, NY, USA: Association for Computing Machinery, 2000, p. 307–317. [Online]. Available: https://doi.org/10.1145/339331.339426

[21] S. Y. Qiu, P. D. McDaniel, and F. Monrose, "Toward Valley-Free Inter-domain Routing," in *2007 IEEE International Conference on Communications*, 2007, pp. 2009–2016.

[22] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '02. New York, NY, USA: Association for Computing Machinery, 2002, p. 3–16. [Online]. Available: https://doi.org/10.1145/633025.633027

[23] K. Sriram, D. Montgomery, D. R. McPherson, E. Osterweil, and B. Dickson, "Problem Definition and Classification of BGP Route Leaks," RFC 7908, Jun. 2016. [Online]. Available: https://www.rfc-editor.org/info/rfc7908

[24] B. Wijchers and B. Overeinder, "Quantitative Analysis of BGP Route Leaks," NLnet Labs, Nov. 2014. [Online]. Available: https://ripe69.ripe.net/presentations/157-RIPE-69-Routing-WG.pdf

[25] W. Eddy, "Transmission Control Protocol (TCP)," RFC 9293, Aug. 2022. [Online]. Available: https://www.rfc-editor.org/info/rfc9293

[26] T. Bates, E. Gerich, L. Joncheray, J.-M. Jouanigot, D. Karrenberg, M. Terpstra, and J. Yu, "Representation of IP Routing Policies in a Routing Registry (ripe-81++)," RFC 1786, Mar. 1995. [Online]. Available: https://www.rfc-editor.org/info/rfc1786

[27] C. Villamizar, T. J. Bates, C. Alaettinoglu, D. Meyer, M. Terpstra, D. Karrenberg, and E. P. Gerich, "Routing Policy Specification Language (RPSL)," RFC 2280, Jan. 1998. [Online]. Available: https://www.rfc-editor.org/info/rfc2280

[28] E.-y. Kim, L. Xiao, K. Nahrstedt, and K. Park, "Secure Interdomain Routing Registry," *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 2, pp. 304–316, 2008.

[29] G. Siganos and M. Faloutsos, "Analyzing BGP policies: methodology and tool," in *IEEE INFOCOM 2004*, vol. 3, 2004, pp. 1640–1651 vol.3.

[30] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," RFC 6480, Feb. 2012. [Online]. Available: https://www.rfc-editor.org/info/rfc6480

[31] S. Boeyen, S. Santesson, T. Polk, R. Housley, S. Farrell, and D. Cooper, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," RFC 5280, May 2008. [Online]. Available: https://www.rfc-editor.org/info/rfc5280

[32] D. C. W. L. Jr., K. Seo, and S. Kent, "X.509 Extensions for IP Addresses and AS Identifiers," RFC 3779, Jun. 2004. [Online]. Available: https://www.rfc-editor.org/info/rfc3779

[33] M. Lepinski, D. Kong, and S. Kent, "A Profile for Route Origin Authorizations (ROAs)," RFC 6482, Feb. 2012. [Online]. Available: https://wwfw.rfc-editor.org/info/rfc6482

[34] R. Bush, "Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI)," RFC 7115, Jan. 2014. [Online]. Available: https://www.rfc-editor.org/info/rfc7115

[35] M. Lepinski and K. Sriram, "BGPsec Protocol Specification," RFC 8205, Sep. 2017. [Online]. Available: https://www.rfc-editor.org/info/rfc8205

[36] R. Lychev, S. Goldberg, and M. Schapira, "BGP security in partial deployment: is the juice worth the squeeze?" *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 4, p. 171–182, Aug. 2013. [Online]. Available: https://doi.org/10.1145/2534169.2486010

[37] K. Kim and Y. Kim, "Comparative analysis on the signature algorithms to validate as paths in bgpsec," in *2015 IEEE/ACIS 14th International Conference on Computer and Information Science (ICIS)*, 2015, pp. 53–58.

[38] A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, and K. Sriram, "BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects," Internet Engineering Task Force, Internet-Draft draft-ietf-sidrops-aspa-verification-22, Mar. 2025. [Online]. Available: https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-verification/22/

[39] T. Bruijnzeels, "RPKI Functionality Roadmap," May 2025. [Online]. Available: https://ripe90.ripe.net/wp-content/uploads/presentations/110-RIPE-NCC-RPKI-Features-2025.pdf

[40] A. Hari and T. V. Lakshman, "The Internet Blockchain: A Distributed, Tamper-Resistant Transaction Framework for the Internet," in *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*, ser. HotNets '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 204–210. [Online]. Available: https://doi.org/10.1145/3005745.3005771