Reviewing a historical Internet vulnerability: why isn't BGP more secure and what can we do about it?

Cecilia Testart

August 2018

Abstract

The Border Gateway Protocol (BGP) plays a crucial role in today's communications as it is the inter-domain routing protocol that holds together the Internet, providing the path for IP packets to flow between networks across the globe operated by different providers. Although the first version of BGP was published in 1989 and its lack of security mechanisms has been known since then, BGP remains vulnerable to attacks that can cause large scale outages or can be used for other malicious purposes on the Internet, such as traffic sniffing or spam sending. Moreover, the lack of security has not prevented the surge of new applications that run on top of the Internet, making tampering with BGP increasingly attractive. As an example, BGP hijacking was used to steal at least $\$83,000$ worth of cryptocurrency in 2014, and again more recently in April 2018. Thus, securing BGP is key to increasing the overall security of the Internet ecosystem. This paper offers a historical review of the different ideas put forward to secure inter-domain routing and what happened to these ideas along the way, noting if they were implemented and are in use, the impact they had on other proposals, and other characteristics explaining the difficulty of securing BGP. This study analyzes 10 BGP extensions focused on BGP availability, 7 BGP extensions and best practices focused on securing BGP communication and routing information, and 11 security proposals coming from the research and industry communities. It examines where the ideas came from, the implicit trust delegation and the residual vulnerabilities of proposals. Even though performance and incentives of specific security solutions have been largely discussed, most proposals have not even been implemented, limiting the overall security improvement of BGP. Reviewing the full life cycle of the proposed ideas, the trusted actors and mechanisms and their requirements gives insight into why adoptions rate are so limited. In fact, there is a

remarkable lack of consensus on what needs to be secured or validated, and the approach to be taken, preventing solutions to get critical support to move their deployment forward. Additionally, no BGP security mechanism, even the most narrow one, has been easily implemented and deployed. However, there are security best practices that certainly improve local and overall BGP security. Since no solution comes without costs and all proposal have opponents in network operation community, it may be possible that secure routing should be provided as a separate service from routing, and other entities could offer such solutions.

1 Introduction

1.1 Background

The Border Gateway Protocol (BGP) is the inter-domain routing protocol that holds together the Internet, providing the necessary information to route traffic between networks across the globe operated by different providers. However, BGP lacks internal security mechanisms to protect communication between peers and to verify the correctness of routing information. Unfortunately, many events along the years have shown that BGP operation can easily be disrupted by an intended attacker [1] or unintentional configuration mistake [2–4]. More recently, there have been BGP attacks targeting specific services and applications that run on the public Internet. For instance, BGP attacks can have a significant impact in the operation of cryptocurrencies [5, 6]. Thus the relevance of understanding how to move forward to secure BGP.

There have been many proposals to secure BGP from the IETF [7–11] and from the industry and academic community [12–22]. Previous work has focus on evaluating security proposals performance and efficiency [23, 24], their limitation and advantages concerning their security guarantees [25], the techniques they use to secure BGP [26], the dynamics of their architecture [27], and analyze their incremental security benefit in deployment and fully deployed. [28, 29] . Additionally, the work of the Secure Inter-Domain working group at the IETF has been studied in detailed [30]. This work is based on the analysis of the life-cycle of ideas to secure BGP, their delegation of trust, residual vulnerabilities, mechanism used and their requirements, coupled with the study of BGP development.

1.2 Methodology

This work selected proposals to secure BGP from RFC documents describing BGP extensions and literature survey papers. The proposal considered are specifically focused on improving BGP and not any generic routing algorithm, although they may use aspects from those works. The proposals that had some traction were selected. For IETF proposal, it meant that RFCs were updated, mentioned in BGP protocol updates or discussed in operational practice documents. For industrial and academic proposals, it meant that those proposals were mentioned in many literature surveys, performance or scalability studies.

This work also considers BGP extensions to improve its availability as availability is one of the classic dimensions of information security. Additionally, since those extensions are all deployed and in use, they provide a good example on what can be achieved in changing a protocol such as BGP.

Information about the life-cycle and motivations of proposals was inferred from the main document or accompanying documents of the proposals and from studies analyzing different aspects of the proposals.

1.3 Paper outline

This paper is organized as follow: Section 2 describes how BGP works and its evolution. Section 3 reviews the IETF proposals to secure BGP and their life cycle. Section 4 reviews proposals to secure BGP from academia and industry. The proposals life-cycle and other characteristics are compared and discussed in section 5. Section 6 concludes the paper.

2 BGP and BGP development

The Border Gateway Protocol (BGP) was developed in the late 1980s for border routers (gateways) to exchange inter-domain routing reachability information with neighboring networks. The first Internet Engineering Task Force (IETF) [31] Request For Comment (RFC) formalizing the standard was published in June 1989 [32]. Since then five new versions and updates of the protocol have been published alongside more than 15 protocol extensions that add different capabilities to the base protocol, totaling over 50 standard track RFCs related to BGP. The next section describes the overall functioning of BGP-4, the current version of the protocol.

2.1 How BGP works

In BGP, Autonomous Systems (AS) exchange network reachability information with neighbors. An AS is a group of routers operated by a single administration, providing all routers in the AS the same coherent routing plan. An Internet Service Provider (ISP), also called network operator, may have one or more ASes under its control. Regional growth, mergers and acquisition continually change AS ownership and increase the number of ASes in the public Internet. Figure 1 depicts an example topology with 4 ASes: AS1 has one router ($R1$), AS2 has 3 routers ($R2$, $R3$ and $R4$) and AS3 ans AS4 have also one router each ($R5$ and $R6$ respectively). A router that runs BGP is called a *BGP speaker*. The end points of links running BGP are *BGP peers* —$R1$ and $R2$ are BGP peers in figure 1. Two BGP peers establish a *BGP session* and start exchanging routing information. The BGP session runs on top of the Transport Control Protocol (TCP). After opening the session, BGP speakers send each other the whole list of routes to IP addresses reachable from their network. This list contains the AS paths to different prefixes —the reachability information— and is called the Routing Information Base (RIB). ASes choose which routes they want to advertise as available through their own network to each BGP peers. After sending the initial RIB, BGP peers exchange incremental changes to inform about routes that became active and routes that became inactive.

To identify a network, BGP uses an IP address prefix and a prefix length, such as 18.23.0.0/16, where 18.23.0.0 is the IP prefix and the /16 is the prefix length. When a BGP speaker advertises a prefix to one of its peers, it comes with a number of *attributes*. One of the most relevant attributes is the *AS Path*, which is a list of AS numbers corresponding to the path a network announcement has followed since its origination by the source AS. If a speaker is in an AS originating a route to a network, the AS path will only consist of that AS number. For instance, in figure 1, AS2 originates the announcement of 137.29.138.0/24, AS3 originates 18.23.0.0/16 and AS4 originates 203.70.0.0/16. However, if a speaker is in an AS re-advertising a route to a network received by another peer, it will prepend its AS number at the beginning of the route. In figure 1, AS2 is re-advertising 18.23.0.0/16 and 203.70.0.0/16 to AS1, prepending its AS number in front of the AS path.

Other attributes such as the *Local Preference* and *Multi-Exit Discriminator (MED)* attributes are metrics that are used to specify a degree of preference for the route. The BGP speaker receiving these attributes may decide to use or not the values of these attributes in its route selection decision process. Many new attributes have been added over time to BGP to add features and functionalities to the protocol.

The network reachability information received from a peer —routes composed of an IP prefix and an AS path— is considered valid by a BGP speaker until the peer

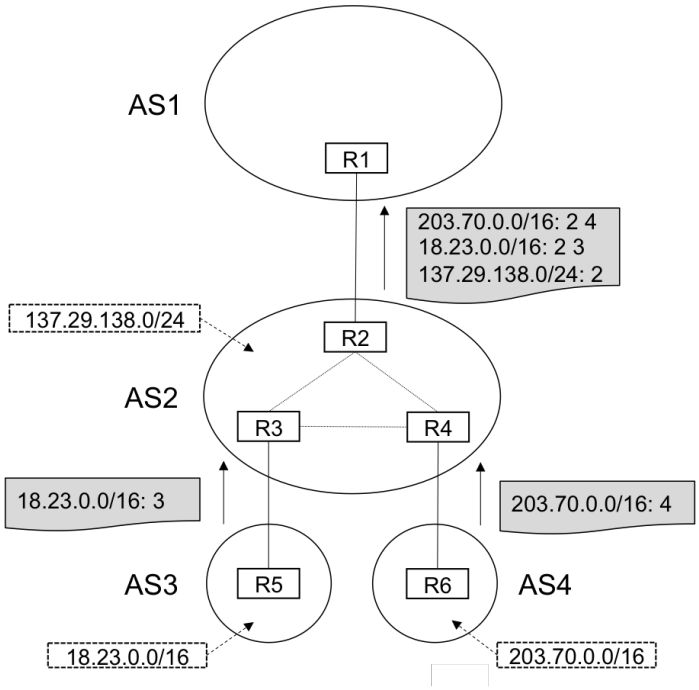

Figure 1: Example topology of 4 Autonomous Systems (ASes). Ovals represent ASes, boxes represent routers (BGP speakers) within ASes. Lines indicate BGP sessions. Prefixes used by ASes are shown in dashed boxes. Grey boxes have examples of BGP announcements with prefix, prefix length and AS path going from ASes 2, 3 and 4 to AS1.

advertises an update to the same route or a withdrawal of the route, or the BGP session is closed. The closure of a BGP session causes both speakers to erase all information learned from each other.

With the network availability information received from its peers, a BGP speaker runs a route selection process to select the routes it will use and which subset of these routes it will advertise to each of its other peers based on their relationship. The BGP route selection process is based solely on route attributes —including the AS path— and AS routing policies.

In the route selection process, if there is only one route to a specific network —a specific prefix with a given prefix length— that route is selected and its information is copied to the router forwarding table. However, if there are two or more routes to a network, the BGP speaker uses tie-break rules to select which route it will use. The first tie-break rule is the degree of preference of the route computed in the route selection process. The second tie break rule is the shortest AS path: the smallest count of AS numbers found in the AS path attribute from the route. The rest of the tie break rules are based on other attributes of routes and the BGP peer the routes came from.

The routes selected for use by a BGP speaker are then copied to its routing table. In the routing table, each reachable prefix is associated with an IP address of the AS where traffic to the prefix should be sent next. When a data packet is received by a BGP speaker, it will be sent to the IP address found in the routing table next to the IP prefix with the longest prefix match between this IP address prefix and the destination IP address of the packet. This means that the packet will be forwarded to the more specific and smallest network in the routing table that includes the destination IP address of the data packet.

Summarizing, BGP is an algorithm allowing BGP peers to learn available networks reachability information in a distributed way, without requiring a central entity to sort the information or network structure. Additionally, BGP offers the possibility to apply AS-level policy routing decisions based on the destination network of packets and its route attributes.

2.2 BGP evolution

From 1989 to nowadays, there have been four versions of BGP, with BGP-4 having had two major updates [32–37]. From the beginning, the main goal of BGP has been to exchange network reachability information between ASes. Figure 2 is a timeline of RFCs documents with the main BGP protocol documents on the left side and the extensions and their respective updates on the right side. As it can be seen, ever since the first formalization, BGP has been in constant evolution. Either the main protocol itself or the many extensions have been modified or created to accommodate the evolving usage of BGP and the evolution of the interconnection of networks forming the public Internet.

A major change to the protocol was made in version 4. In previous versions, networks were advertised according to the hierarchical class system. Instead, BGP-4 supports classless inter-domain routing, which means available networks are identified by an IP address prefix and a prefix length. It also meant that longest prefix match —the more specific network match— became the base behavior for forwarding data packets.

Additionally, the first versions of BGP (BGP-1, BGP-2 and BGP-3) had the option of including authentication data in messages [32–34]. However, no specific authentication mechanism was ever mentioned or formalized in RFCs and in BGP-4, the authentication option was deprecated because it was not being used [36].

According to BGP RFCs, the fundamental priorities for the development of the main BGP protocol are:

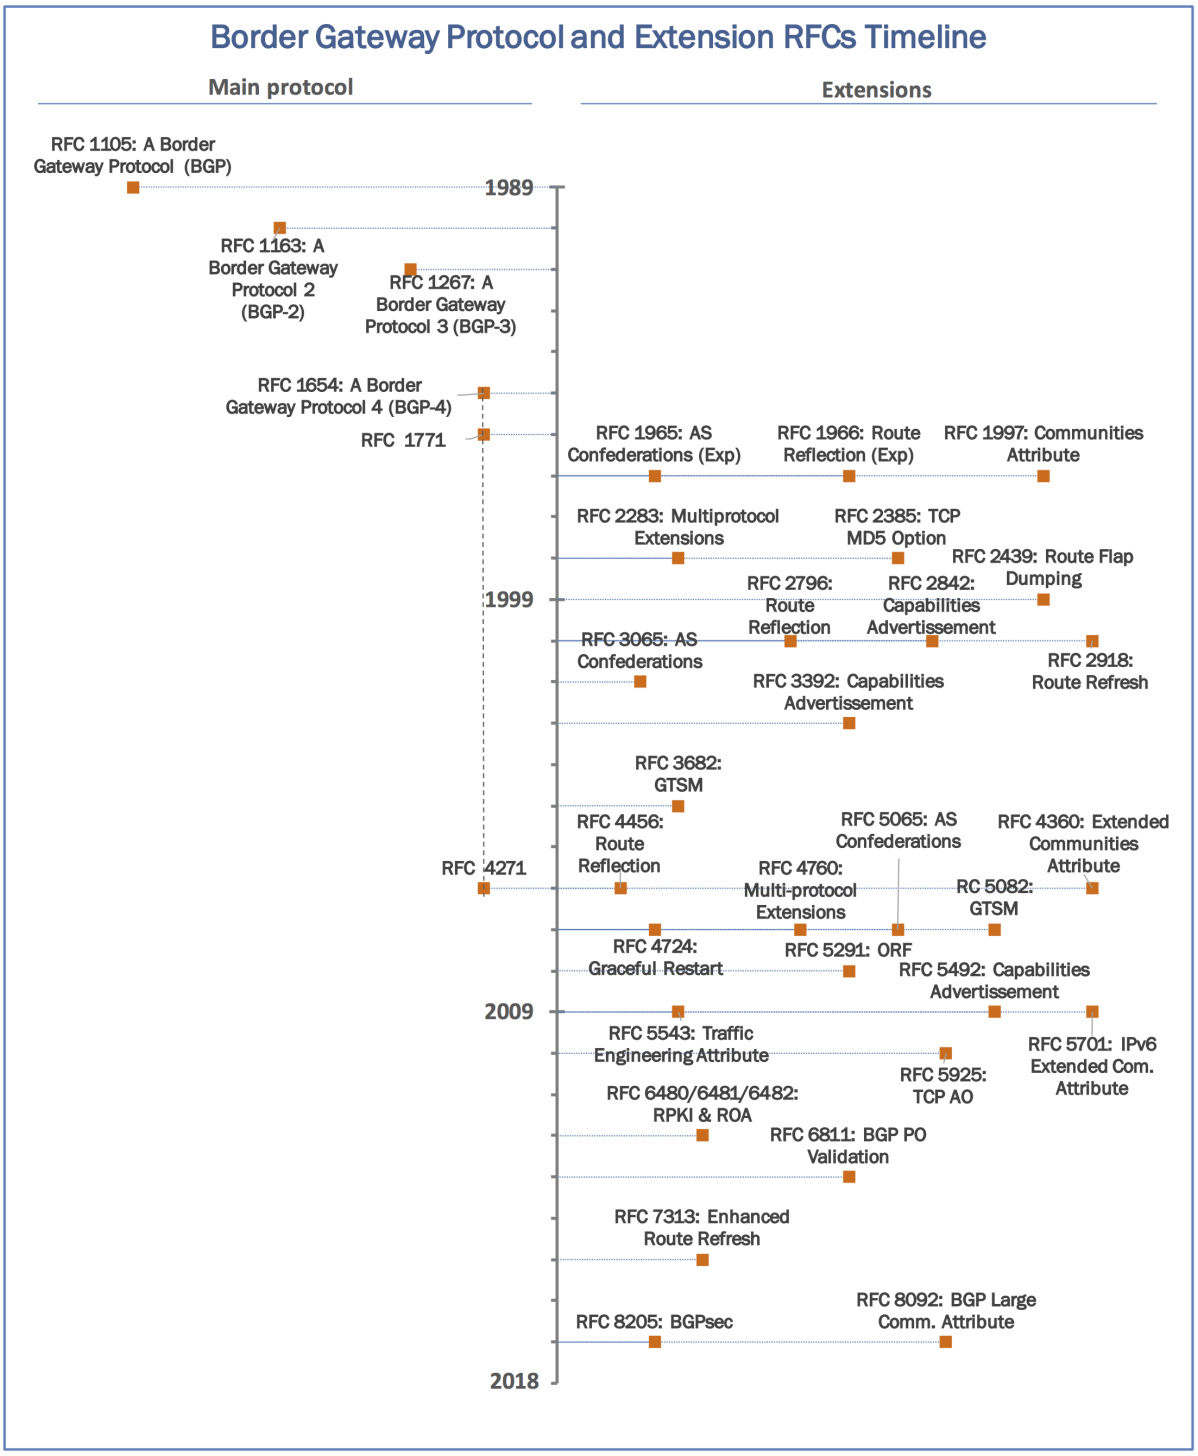

Figure 2: Timeline of standard track IETF RFCs of the Border Gateway Protocol (on the left side) and extension (on the right side). Horizontal lines indicate RFC documents year of publication

- Ability to enforce destination based AS-level policies: each AS needs to be able to enforce the policy of choice concerning destinations reachable through its network.
- Scalability and efficiency: BGP needs to be able to handle the increasing number of prefixes advertised in the public Internet without using too much traffic.
- Dynamic routing while limiting convergence time: BGP needs to accommodate frequent changes to reachability information without taking too long in its route selection process.
- Identification of routing loops: BGP needs to provide a mean to prune routes from the route selection process if a loop is identified to prevent endless looping of data packets.
- Limiting manual configuration of routing policies: BGP needs to allow for automated decisions based on general policies configurations.
- Flexibility for complex and creative routing policies: BGP has to be flexible to accommodate new developments to support policy routing.

When the fourth version of BGP was introduced in 1994, BGP had been deployed in different networking environments and many independent interoperable implementations existed. In 2006, all inter-domain routing was done using BGP.

Since BGP-4 was introduced, more than 15 extensions have been developed and standardized in IETF RFC documents. Extensions are depicted in the right side of figure 2. These extensions add new capabilities to the base protocol. Most BGP extensions were designed to improve and extend BGP operation, to facilitate policy management, or to provide some level of security to the protocol.

However, BGP extensions developed with one motivation in mind are usually detrimental to the other motivations. For instance, all BGP Community Attribute extensions[1] were developed to support and facilitate extensive and creative policy routing management between ASes, even though they add vulnerabilities. Indeed, an attacker could use communities attributes to influence route selection of other ASes.Conversely, the addition of the security extension TCP MD5 digest negatively impacts BGP operation and availability since it adds processing time and can limit the ability of a peer to reset the connection after a crash where key information was lost.

[1]BGP Communities Attribute, BGP Extended Communities Attribute, IPv6 Address Specific BGP Extended Community Attribute and BGP Large Communities Attribute extensions.

2.3 BGP vulnerability analysis

As experience with BGP-4 accumulated, awareness of its shortcomings and vulnerabilities developed. The main aspect of security discussed and addressed in BGP RFCs is availability. There is an understanding that BGP has to provide as much availability as possible, including in the case of an attack or unintended failure. Consequently, to increase availability, most BGP extensions aim at reducing management complexity and manual configuration of BGP speakers —which are prone to unintended failures— and routing instability, given the increasing complexity of network topologies and routing policies.

Other aspects of security, such as integrity and correctness of routing information, are discussed in RFCs specifically addressing BGP or routing protocols security vulnerabilities. A detailed BGP vulnerability analysis, was published in 2006 [38]. Although the document mainly discusses vulnerabilities linked to the mechanics of BGP sessions and message exchange, it mentions that vulnerabilities related to the algorithmic behavior of BGP arise from three main reasons:

1. the lack of authentication of BGP messages failing to provide integrity, authentication and dynamic validation to BGP data in BGP messages: BGP messages can be spoofed, modified, deleted or relayed by a peer or an outsider.
2. the lack of authority validation of network (prefix) announcements: any BGP speaker can announce a route to any prefix.
3. the lack of authentication validation of AS path and path attributes: AS path and attributes can be modified along the hops of an announcement.

Indeed, when a BGP speaker receives a BGP message, the speaker trusts the message is coming from a legitimate peer if the identification information in the message simply matches that of a known peer. In addition, if the BGP speaker processes the routing information, it is trusting that the route prefixes in the message were legitimately announced, and that all ASes in the path of the announcement have properly prepended their AS number in the AS path and modified path attributes. As captured in RFC 4360, " [a network] operator who is relying on the information carried in BGP must have a transitive trust relationship back to the source of the information" [39]. However, this is usually not the case and as more networks have connected to the Internet, it has become increasingly difficult.

The BGP protocol does not include any mechanism to verify the correctness of the information send with respect to the guidelines and norms of the Internet community. The Internet is composed of interconnected networks. In contrast, there is a hierarchy of organizations that are involved in IP address (number resources of the Internet) allocation and delegation. The Internet Assigned Numbers Authority (IANA) is at the top of the address allocation process and allocates big portions of the IP address space

for IPv4 and IPv6 to the five Regional Internet Registries (RIRs) covering different geographical regions: ARIN for North America, LACNIC for South America and the Caribbean, RIPE for Europe and Middle East, APNIC for the Asia-Pacific region and AFRINIC for Africa. These regional registries then allocate smaller chunks of IP addresses to Local Internet Registries (LIR) or directly to Internet Service Providers (ISPs). LIR and ISPs in turn delegate IP addresses to other (smaller) ISPs and networks. Therefore, most legitimate IP network prefixes that have been legitimately allocated have followed a specific allocation path in the hierarchy.

On top of the vulnerabilities described above, there are attack vectors related to the mechanics of BGP communication. They can be divided in three main categories:

1. Vulnerabilities related to BGP messages and finite state machine: There are many ways a BGP message can end up closing a BGP session triggering the deletion of the routing information learned by BGP peers.
2. Vulnerabilities related to the transport layer security: Since BGP runs on top of TCP, it inherits TCP vulnerabilities in a transport-level security mechanism is not used.
3. Vulnerabilities related to the infrastructure where BGP is running: the configuration, operation and management of the infrastructure supporting BGP can impact the correct operation of BGP

3 IETF efforts to secure BGP

At the IETF, there has been many efforts to improve BGP operation, including a couple of working groups addressing routing security. The next sections quickly summarize security additions to BGP considered in this study, including BGP extensions focused on availability.

3.1 BGP extensions for better availability

Since availability is one of the dimensions of information security [40], extensions to BGP addressing availability were studied. For these extensions, the threat model is operational requirements, changes or errors causing BGP session closure or other routing perturbations.

With the goal of reducing operational complexity, BGP session downtime or routing instabilities, at least ten extensions to BGP have been developed. The BGP extensions considered in this paper are shown in table 1.

BGP Extension	Main motivation	BGP change
Autonomous Systems Confederation [41–43]	Reduce managing complexity by reducing the number of internal BGP sessions needed in an AS.	Modifies message field
Route Reflection [44, 45]	Reduce managing complexity by reducing the number of internal BGP sessions needed in an AS.	Adds route attribute
Community Attributes [39, 46–48]	Reduce configuration requirements for policy routing.	Adds route attribute
Route Refresh [49]	Reduce BGP sessions closures and downtime.	Adds new message type
Capabilities Advertisement [50]	Reduce BGP opening negotiation time before the BGP session is up.	Adds message field
Graceful Restart [51]	Reduce routing perturbation from BGP session closures due to expected operational procedures.	Modifies use of message
Route Flap Dumping [52]	Reduce routing perturbations due to route oscillations.	Adds local structure for storing route flap information

Table 1: BGP protocol extensions addressing availability. The Community Attributes include: BGP Communities Attribute [46], BGP Extended Communities Attribute [39], IPv6 Address Specific BGP Extended Community Attribute [47] and BGP Large Communities Attribute extensions [48].

Table 1 summarizes the main motivation behind these extensions and the changes introduced to BGP with the extension. The ideas behind these extensions came from network operators' experience with BGP. Indeed, many of the RFCs describing them were published after these extensions had been in use for a while. These extensions were introduced along BGP development (see figure 2 , and many have been updated a few times along the years to add improvements based on their usage experience. The last update of the main BGP protocol included changes to reflect the usage of AS confederation, Route Reflector and Route Refresh extensions [37].

Moreover, all of the extensions in table 1 modify BGP behavior. All extensions except for the Route Flap dumping also modify BGP message formats or introduce a new message. Nonetheless, in spite of requiring these changes, all of these extensions are implemented and used.

BGP Extension	Main Motivation	Status
TCP-MD5 [7]	Transport-level security instead of BGP level-security	Obsoleted but in use.
TCP-AO [9]	Transport-level security stronger than TCP-MD5 with re-keying option.	Implemented, limited use.
GTSM [8, 53]	Simple transport-level security hack if not using more sophisticated solutions.	Implemented, limited use.

Table 2: Main motivations of proposal to secure BGP

3.2 BGP extensions for transport-level security

The first BGP extensions to address routing information security focused on securing the transport level layer. Three BGP extensions are considered in this category and are shown in table 2.

For these extensions, the threat model considered is an attacker able to spoof BGP messages. This had been considered since the early development of BGP. Indeed, as mentioned in section 2.2, early BGP version included an authentication option deprecated in BGP-4. One of the arguments against the use of this option was that if the TCP layer was not secured, BGP would still be vulnerable to spoofing attacks. This was the main motivation for the first of these extensions, TCP-MD5.

The TCP-MD5 option was introduced in 1998. It does not involve changes to BGP, but rather uses a TCP option for carrying an MD5 digest that is used to verify TCP packets integrity using a password known to both ends of a BGP session.

In 2004, the Generalized TTL Security Mechanism (GTSM) was introduced simple hack to increase BGP transport level security without the burden of TCP-MD5 configuration. Again, this extension does not modify the BGP protocol itself but rather uses TCP features to allow BGP speakers to verify that a received BGP message was sent by a router a hop away.

In 2010, the TCP Authentication Option (TCP-AO) obsoleted the TCP MD5 option as it considered a weak mechanism. Nowadays all BGP implementations are required to support TCP-AO [9]. TCP-AO uses a scheme similar to TCP MD5 but with a stronger message authentication mechanism and a re-keying option to update secret keys without manual configuration.

All of these extensions are currently implemented and deployed. Their use is varied among ISPs. In 2007, TCP-MD5 was still not much in use and GTSM was

Best practice	Main motivation	Status
Route Filtering [55]	Extend use of filters to security (also used for policy routing)	In use, varied levels of implementation, best current practice.
IRRs [57]	Regional routing registries were in use in Europe for supporting routing correct operation. RPSL leveraged RIPE development.	In use, varied levels of implementation, best current practice.

Table 3: Main motivations of proposal to secure BGP

not included in all vendor versions of BGP [54]. Additionally, GTSM use is limited to simple topologies. Finally, in 2015, eventhough TCP-MD5 had been obsoleted and replaced by TCP-AO, is was more widely deployed and used than TCP-AO [55]. To make matters worse, many ISPs had never changed TCP-MD5 password since they started to use it years ago, [56]. The most recent Best Current Practice for BGP security recommends the use of GTSM in direct peering links and that TCP-AO should be preferred to TCP-MD5 when implemented. However, operators are recommended to consider the trade-off of applying TCP-level security [55].

3.3 Security best practices

RFCs documents have formalized security measures that network operator use to limit BGP vulnerabilities. In February 2015, the IETF published a set of best current practices (BCP) for BGP operations and security, BCP 194 [55]. Two security best practices are considered in this review. Their motivation and current state is shown in table 3.

Route filters are mechanisms used to discard routes either received from or to be forwarded to a BGP peer based on the IP prefix address, the AS path or other attribute of the routes. They are widely use for enforcing routing policies. Filters can also be used to verify that routes containing information known to be wrong are not considered in the route selection process and to prevent propagating such invalid announcements.

The IRR are a distributed public registry of routing information for Internet networks. The IRR use a Routing Policy Specification Language (RPSL) allowing router configurations to verify the validity of BGP routing information to be generated from routing information in the registries [58]. Additionally, routing policies in RPSL have

the possibility to assemble globally in a routing registry. In addition, security mechanisms for authentication and authorization of additions and changes were introduced in RPSL and the registries [59].

Route filters and the IRR are old ideas. Policy filtering has been in place since the NSFNET was the backbone of the Internet [60] and their use can be extended to provide security. Registries that inspired the IRR were already in use by the RIPE community [61] to generate filters to validate routing information [58]. However, many network operators do not use this best practices because of their management overhead and the risk of mistakes disrupting traffic. Nonetheless, especially in Europe, some ISPs require their peers and customers to have their information up-to-date in an IRR registry to use it for route filtering. More recently, in 2014, almost 50 network operators signed a document of *Mutually Agreed Norms for Routing Security (MANRS)* describing actions that ISPs should take to increase routing security and the use of route filters and the IRR are two of them [62]. Still, in 2015 filtering and the use of the IRR was inconsistent among ISPs [55], even though it has been shown that using prefix filtering with origin validation techniques provides comparable security to origin and AS path validation while AS path validation mechanism are still in deployment [29], which is currently the case.

3.4 IETF SIDR security solutions

In 2006, the IETF started a the Secure Inter-Domain Routing (SIDR) working group [63] to address routing information vulnerabilities in BGP. The group developed two solutions: Resource Public Key Infrastructure (RPKI) and Route Origin Authorization for prefix validation, and BGPsec for AS path validation. Both are considered in this study and are described in the next sections. Since this solutions are thoroughly described in RFC documents, more detail about them is given.

3.4.1 RPKI and ROAs for prefix origin validation in BGP

The infrastructure developed by the IETF to secure prefix origin relationship is called Resource Public Key Infrastructure (RPKI) and was published in 2012 [10, 64, 65]. It provides the means to verify the legitimate allocation of IP prefix blocks to a resource holder following the current hierarchical allocation system for IP address prefix and AS numbers, and the authorization from this resource holder to an AS to originate traffic to all or part of the IP prefix block it was allocated.

It is based on a Public Key Infrastructure (PKI) to "provide cryptographically verifiable attestations" of number resources allocation [10]. This attestation takes the

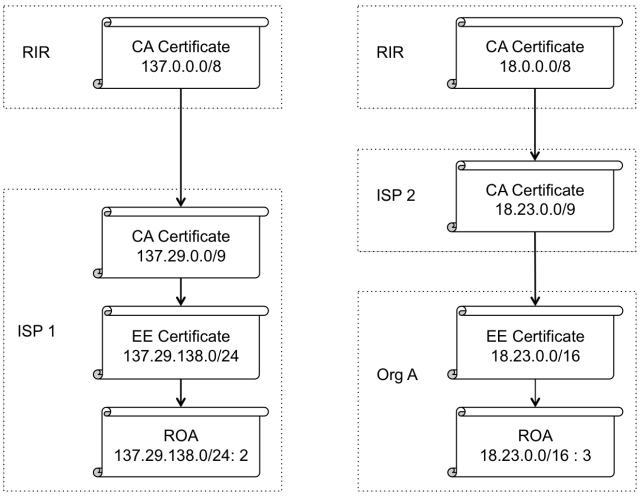

Figure 3: Example of possible hierarchical allocations of IP address prefix.

form of certificates and authorizations: Certificate Authority (CA) certificates attest the validity of the allocation of IP address blocks and AS numbers; the End-Entity (EE) certificates provide the transit authority over these IP blocks, and Route Origin Authorizations (ROAs) link an IP prefix with the number of the AS originating the route. Finally, the distributed repository system is used to store and make available signed objects, such as ROAs.

Figure 3 shows two simplified examples of allocation hierarchies. On the left, the hierarchy has only one large ISP under the RIR, which was previously allocated the block 137.0.0.0/8. The RIR allocates a large block 137.29.0.0/9 to ISP 1, which issues its own CA certificate to then issue an EE certificate for a smaller block 137.29.138.0/4 and sign the ROA linking the EE prefix 137.29.138.0/24 to the originating AS 2. On the right, the hierarchy shows a large ISP which was allocated a big block 18.23.0.0/9 by the RIR. ISP 2 issues a CA certificate and then an EE certificate for a smaller block 18.23.0.0/16 to another organization, Org A. Org A uses the EE certificate to sign the ROA linking the prefix 18.23.0.0/16 to the originating AS 3.

With the current RPKI infrastructure, a BGP speaker needs to have access to cache storage of prefix origin information from validated ROAs to verify the originating AS of a route prefix received from BGP peers. BGP speakers can include the validation result in the route selection process. The level of priority it gets in the route selection process is the matter of local AS routing policies.

The RPKI infrastrucure and its use does not directly introduce any change in BGP, the prefix origin validation is taken into account through router configuration. However, ROAs could be included in Update messages in future version of BGP.

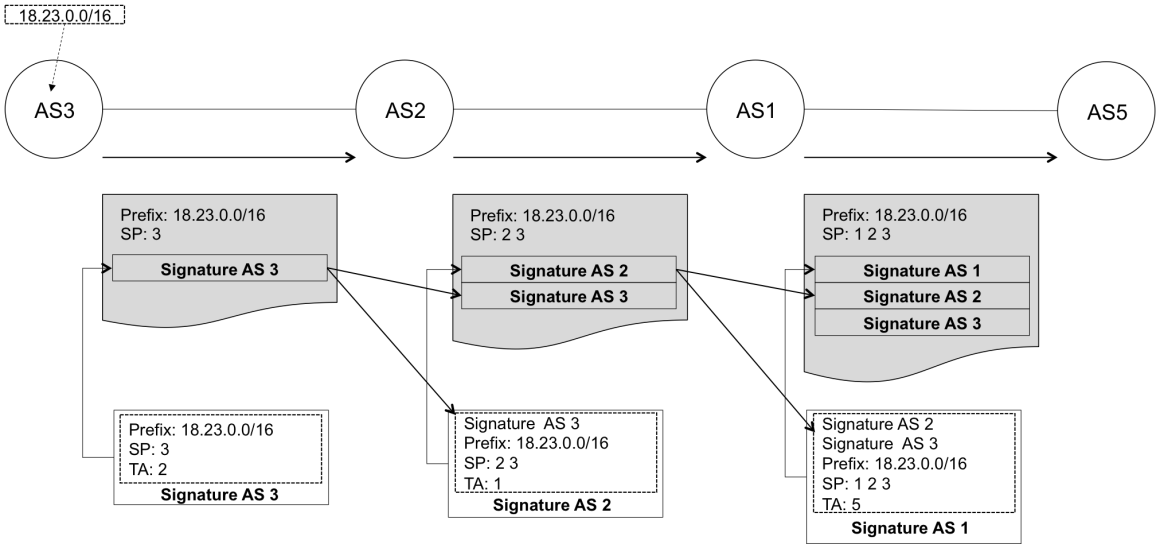

Figure 4: Simplified representation of BGPsec announcements of a route traversing three ASes. The Secure Path (SP) and the Signature list are prepended newer elements as the message is forwarded to the next AS. The AS signature is computed using the Target AS (TA), the Secure Path (SP) and all signatures by previous ASes in the Path.

3.4.2 BGPsec for AS path validation

The mechanism to provide BGP path security developed by the IETF is called BGPsec and it was publish in September 2017 [11, 66]. BGPsec has the means to give assurance to ASes receiving a BGPsec message that the ASes listed in the BGPsec path have explicitly authorized the advertisement of the route prefix to the subsequent AS in the path. The BGPsec protocol and its operation relies on the RPKI infrastructure for the dissemination of signed objects used in the validation of the AS path.

BGPsec introduces a new path attribute to include the digital signature of all ASes in the path, binding the prefix and BGPsec path along the way. Figure 4 shows a simplified representation of BGPsec Update messages of a route with prefix 18.23.0.0/16 originated by AS 3 and forwarded to AS 2 and then re-advertised to AS 1 and then AS 5. Each AS in the path prepends its AS number in the new attribute along with its signature. Each AS along the path computes its signature using an increasing amount of information as it includes all previous signatures and path segments.

For validating a BGPsec message with routing information, all included signatures need to be validated. This can use significant computation resources, especially for long AS paths. BGP speakers can perform the signatures validation by themselves or

Security solution	Main Motivation	Status
RPKI & ROAs	Inspired by S-BGP and other proposals using PKI with IANA/ICANN as root of trust and IETF requirements including backward compatibility.	Implemented and deployed. Properly used by less than 10% of all Internet prefix.
BGPsec	Inspired by S-BGP and other proposal and IETF requirements including backward compatibility and non exposure of AS information.	Standardized in 2017. Efforts underway to improve performance.

Table 4: Main motivations of proposal to secure BGP

receive the data form a trusted source using for instance the RPKI-Router Protocol [67].

BGPsec speakers can advertise route announcements even if its verification failed. After validating the BGPsec Update message, it is expected that the BGPsec router will include this result in the route selection process. However, this is left as a matter of local AS policy.

3.5 SIDR solutions life-cycle

It took the SIDR working group six years to publish the RPKI and ROA standards and another 5 more years to publish BGPsec. Table 4 summarizes these solutions main motivation and last know status.

The RPKI, ROAs and BGPsec were clearly influenced by proposals outside the IETF. In particular, both solution draw aspects from S-BGP [13], discussed in section 4. Note that the authors of S-BGP actively participate in the SIDR working group [63]. S-BGP was the first proposal to present a Public Key Infrastructure paralleling the current IP address and AS number allocation hierarchy with ICANN/IANA at the top. Many other proposals followed the idea as can be seen in table 9, although not all PKI structure proposed had the same hierarchical structure. S-BGP was also the first proposal to include signatures from BGP speakers from the different ASes in the path of a route announcement in the announcement itself. Some other proposals also include ASes signatures along the path of a route announcement with varying schemes.

Nonetheless, the SIDR designs had to accomplish a set of requirements, which are

reviewed in detail in [30]. One of the most relevant requirements, and also its pitfalls, is that SIDR designs had to be backward compatible with BGP-4. Therefore, in both solution, the priority given to the result of prefix origin and path validation are a matter of AS local policies.

Additionally, both SIDR solution are still in deployment phase. The ROAs currently in the RPKI repositories represent less than 10% of the prefix announced in BGP [68].In comparison, in 2013, 87% of prefix announcements were originated by ASes that had registered their IP block information in one registry part of the IRRs [69]. Recent work has extended the RPSL language to include RPKI object, adding a new channel for the dissemination of ROAs [70]. And the development of the [67] will facilitate the management of router cache storage of validated ROAs used to verify prefix origin in BGP.

BGPsec is even earlier in the cycle of adoption, with the standard only published in September 2017, although it was much talked about during its development and the concept has been discussed for more than a decade [71]. In addition, the designers of BGPsec recognize that it may be a long time before BGPsec is widely adopted, in particular because of its CPU and memory requirement [72]. There are efforts underway to optimize BGPsec performance [73]. Nonetheless, BGPsec has strong opponents who consider unacceptable the trade-offs of implementing it and disagree with the need to secure AS path as defined in BGPsec [74].

4 Other proposals to secure BGP

There has been much work looking to secure BGP outside the IETF. This study considered proposals that were specifically design to secure BGP and had some traction in the IETF, industry or academic community. Generic algorithm to secure routing are not considered.

4.1 Proposals description

The main aspects of considered proposals are quickly describe in the following list in chronological order.

- Securing BGP AS Path with Predecessor information: In 1996, Smith and Garcia-Luna-Aceves published one of the first set of measure to protect BGP [12]. Their proposal is based on cryptographic mechanisms to provide confiden-

tiality to BGP messages and changes to BGP messages protect and verify route announcements.

- The Secure Border Gateway Protocol (S-BGP): S-BGP was developed by a group at BBN Technologies and was published in 2000 [13]. S-BGP protocol also considers the uses of a cryptographic mechanism to protect BGP messages integrity and confidentiality, IPsec [75]. Additionally, it relies on a Public Key Infrastructure (PKI) that follows the the existing hierarchy for IP address and AS number assignments, and the use of certificates for the delegation of such resources. The authors propose the use of a new path attribute to send AS signature in route announcements that would be verified by each BGP speaker along the path.

- Hop Integrity for routing security: In the early 2000's, the Hop Integrity protocol was developed by Gouda *et al.* in collaboration with IBM research Labs [14]. It uses a cryptographic mechanism through the support of two new protocol layers that need to be added in the protocol stack of routers to provide integrity check and exchange new secret keys smoothly.

- Origin lists for false origin detection: In 2002, Zhao *et al.* proposed the use of a list containing the origin information of prefix to check for false origin announcements [15]. This list would include the multiple ASes that can legitimately originate a specific IP address prefix and would be sent along announcements using of an existing BGP optional attribute.

- Secure Origin BGP (soBGP): In the early 2000s, the secure origin BGP (soBGP) protocol was developed by a group mostly within Cisco System and was published in 2003 [16]. The soBGP protocol considers the use of cryptographic mechanisms to secure BGP communication. It also relies on a PKI infrastructure, but the root of trust for IP allocation is a "small number of well-known entities" that would then issue certificates to other ISPs and organizations, forming a "web of trust" for allocating IP prefixes [16]. soBGP also requires AS to publish a list of BGP peers and the set of policies that the origin AS would like to apply to the route announcement of an IP prefix. A new type of BGP message would be used to disseminate certificates between BGP peers.

- The Interdomain Route Validation (IRV) Protocol: Also in the early 2000's, a group for the ATT research lab developed the IRV Protocol and published it in 2003 [17]. It is based on a decentralized query systems that connects ASes and is used to verify routing information from BGP. ASes may host or designate an IRV database to speak authoritatively about their network status and routing information.

- The Secure Path Vector (SPV) Protocol: Researchers from UC Berkeley and Carnegie Mellon University published the Secure Path Vector (SPV) Protocol in

2004 [18]. SPV protocol is based on a series of cryptographic mechanisms that authenticate ASes, allows ASes to authorize route announcements and ensures message freshness through certificates expiration. SVP also relies on hierarchical certificate structure equivalent to the Public Key Infrastructure used in S-BGP [13] for allocating IP prefixes.

- Listen and Whisper: In 2004, researchers at UC Berkeley presented the two mechanisms *Listen* and *Whisper* to improve BGP security [19]. The Whisper protocol is a monitor system based on a signature scheme that is included in BGP announcements. If a BGP speaker receives two routes to the same prefix origin it can verify if the signatures are consistent and choose the route coming from the best behaving AS according to a penalty metric. If a BGP speaker wants to verify that a prefix is reachable through a specific route, the Listen protocol can be used.

- Pretty secure BGP (ps-BGP): In 2004, the Pretty Secure BGP (psBGP) was presented by researchers from Carleton University based on the analysis of S-BGP and soBGP [20]. psBGP also uses IPsec [75]. To authenticate AS numbers, psBGP relies on a centralized PKI like S-BGP [13]. To validate the allocation of IP prefix, each AS creates a list of the prefixes it originates and additional lists for the prefixes its peers originate. psBGP also requires each AS in the path to append its signature of the route information to be sent in the route announcement.

- External Security Monitors (ESM) to secure BGP: In 2006 researchers from Cornell University published a mechanism to use an overlay network of ESMs to monitor and secure BGP traffic, verifying the correct modification of the AS path at each hop and check origin authentication certificates [21]. The authors propose the use of a decentralized PKI infrastructure called Grassroots [76], where ASes are able to directly issue their certificates for their prefixes. Certificates would be send over the ESM network.

- In 2006, Qiu and Gao [22] published Hi-BGP after studying previous proposals to secure BGP. Like soBGP [16], Hi-BGP relies on a "web-of-trust" PKI infrastructure to issue IP prefix ownership certificates, requires the use of transport-level security or encryption, and introduces a new type of BGP message to send certificates. However, Hi-BGP asks AS to publish full and accurate routing information including prefix ownership, AS links and AS relationships to verify routing information in BGP.

4.2 Life-cycle of proposals

The main motivations that guided the BGP security proposals studied in this paper are considerably different. From the need to protect BGP communication to developing from scratch secure monitors to verify the correct operation of BGP speakers, the main reasons why protocol designers considered their solution to be a good one are very varied. This probably stemmed from the fact that there is no agreement between these proposals about what needs to be secure, how it should be secure or what are acceptable trade-offs. Some solutions such as Listen and Whisper and the use of origin lists only address conflicting routes, while other solution address all route announcements.

In addition, there is a clear influence of the first proposals in the latter proposals. Many take elements of the first proposal, improve some aspects or clearly oppose some principle and provide an alternative. For instance, S-BGP proposed a hierarchical PKI infrastructure to validate allocations of IP prefix, which other solutions later also included. However, soBGP proposed a more decentralized PKI structure with a small group of trusted entities at the top. And ESM proposed an even more decentralized PKI structure where each AS would issue its own certificates and only in case of conflict would those certificate need to look for attestations.

Furthermore, usually the designers motivation guided one feature of the solution proposal and to cover more security aspects of BGP, designer used elements from earlier proposals or other protocol already standardized. For example, SPV designers focused on developing a more efficient mechanism to validate AS Path than the one from S-BGP but use a centralized PKI infrastructure like S-BGP for validating prefix origin.

Finally, none of these proposals were fully implemented or used, and only a handfull of them are still discussed in BGP security related works. However, many of these works have influenced the SIDR developments described in section 3.4 and some can potentially have influenced current BGP monitoring services, either by ISPs or other entities.

5 Discussion

The security solutions considered in this paper involve specific mechanisms to secure BGP, the inter-domain routing protocol used in the Internet. Although generic algorithm for network and routing security are not part of the review, some proposals include specific applications of such algorithm, taking into account how BGP works.

Security Solution	Main Motivation	Status
Peer-to-Peer	First overall security proposal	Not implemented.
Encryption and Predecessor Information	motivated by BGP main vulnerabilities.	Not implemented.
S-BGP	First proposal with prefix origin validation using a PKI infrastructure, motivated by the use of cryptographic mechanisms for improving BGP vulnerabilities.	Closed project, one implementation, mentioned in recent review.
Hop Integrity	Transport-level security with easy re-keying.	Not implemented.
Origin Lists	Light-weight solution with validation only in case conflict.	Not implemented.
soBGP	Distributed root of trust for validating IP prefix allocation, AS transparency for security.	Implementation guidelines exist, mentioned in recent review.
IRV	Direct routing information querying between ASes, no changes to BGP.	Not implemented, mentioned in recent review.
SPV	Efficient validation of signatures for AS path validation.	Not implemented.
Listen and Whisper	Light-weight solution with validation only in case conflict paired with data plane security.	Not implemented.
psBGP	Improvements over S-BGP and soBGP.	Not implemented, mentioned in recent review.
ESMs	Delegate security to blank-state hardware and software developed from scratch for secure routing.	Not implemented as proxy, unknown if used as monitoring tool.
Hi-BGP	Prevent hijacking attacks.	Not implemented.

Table 5: Main motivations of proposal to secure BGP

The first proposals to improve BGP security focus on improving BGP availability (see section 3.1). Although availability is often disregarded in security discussion, is one of the fundamental objectives of information security along with confidentiality and integrity [40]. In the routing context, availability is one of the major objectives. All the BGP extensions focusing on solving problems that impact availability are widely deployed and used. Indeed, many of them have been updated to introduce small changes based on usage experience, which confirms the relevance of availability for network operators.

The rest of the security proposals for BGP studied in this paper focus on the correctness and integrity of BGP and BGP routing information. The rest of the discussion section focus on those proposals.

From the study of their life-cycle, it is clear that no security proposal has been quickly implemented and deployed, no matter its motivation and design. The solutions focusing on transport-level security provide a stark example: the obsoleted TCP-MD5 is still in use, and TCP-AO and GTSM only have a limited use although they were standardized 8 and 14 years ago respectively. And it took at least 10 years for TCP-MD5 to go beyond a limited usage.

Additionally, the main motivation behind the developments evidence the lack of agreement of what needs to be secured and what secured means in the context of BGP. Some proposals are concerned to provide security in some cases, such as when there is a conflict between two routes, whereas other cover all possible cases. In addition, some proposals defined prefix origin validation as being with respect to certificate and attestations of IP address block allocation and transit authorization, whereas others defined it with respect to ISPs own declaration of resources or consensus among what ISPs consider to be correct. This is related to the trust delegation implied by the different proposals and their residual vulnerabilities. Tables 6 and 7 summarize the proposal's trust delegation and residual vulnerabilities.

The trust delegation analysis reveals that proposals have different trusted actors. For example, the Peer-to-Peer Encryption and Predecessor Information proposal uses a topology built from BGP route announcements to secure announcements in BGP, i.e. uses BGP to secure itself. In contrast, the proposals relying on a PKI structure to verify the allocation of IP prefix have an additional source of routing information to verify BGP routing information. However, the root of trust of the additional routing information varies. Some trust the current organizations in charge of IP allocation, whereas others choose to trust ISPs or a selected group of ISPs.

Equally important, all proposals imply some level of trust in other ASes. Indeed, since validation of routing information or configuration of secure communication mechanisms depend on local router configuration, if those configuration are not

Security Solution	Trust delegation
TCP-MD5	Secret key configuration and management of BGP peers.
TCP-AO	Master key configuration of BGP peers.
GTSM	Next hop BGP peers.
Route Filtering	Other ASes filters.
IRRs	Data entered by ISPs, registry holder security mechanisms, other ASes filters.
RPKI & ROAs	RPKI organizations and repositories, other ASes security policies.
BGPsec	RPKI organizations and repositories, other ASes security policies.
Peer-to-Peer Encryption and Predecessor Information	Topology built from routing announcements in BGP, other ASes security policies.
S-BGP	PKI organizations and repositories, other ASes security policies
Hop Integrity	BGP Peers
Origin Lists	ASes prefix origin lists and other ASes route conflict resolution.
soBGP	Well-known entities acting as root in 'web of trust' (PKI), ASes routing information and security policies
IRV	ASes query responses and security policies.
SPV	PKI organizations and repositories, other ASes security policies
Listen and Whisper	Other ASes route conflict resolution.
psBGP	PKI organizations for AS numbers, AS prefix list and peer prefix origin lists.
ESMs	Certificate conflict resolution in decentralized scheme, other ASes security policies.
Hi-BGP	Well-known entities acting as root in "web of trust" (PKI), ASes routing information and security policies

Table 6: Trust delegation of studied proposals to secure BGP.

Security solution	Residual vulnerabilities
TCP-MD5	Weak protection especially if key never changed. Routing information not validated.
TCP-AO	Routing information not validated.
GTSM	Complex topologies cannot use GTSM, Routing information not validated.
Route Filtering	Depends on filters used, vulnerable to resourceful attackers.
IRRs	Depends on information available and filters used.
RPKI & ROAs	AS path not validated, RPKI vulnerabilities.
BGPsec	Prefix origin not validated, RPKI vulnerabilities.
Peer-to-Peer Encryption and Predecessor Information	Complex AS path hijacking, ASes collution, prefix allocation not validated.
S-BGP	Route leaks, PKI vulnerabilities.
Hop Integrity	Routing information not validated.
Origin Lists	AS path not validated, prefix allocation not validated, false but not conflicting origin accepted.
soBGP	Mis-behaving actors in web-of-trust. Route leaks.
IRV	IRV database impersonation, prefix allocation not validated.
SPV	Route leaks, PKI vulnerabilities.
Listen and Whisper	Prefix origin not validated, false but not conflicting routes are accepted.
psBGP	PKI vulnerabilities, colluding ASes, prefix allocation not validated.
ESMs	ASes collusion, ESM vulnerabilities.
Hi-BGP	Mis-behaving actors in web-of-trust.

Table 7: Residual vulnerabilities of studied proposals to secure BGP.

properly done, or not done at all, then no security was provided by the mechanism. Nonetheless, an AS cannot verify that its peers perform the appropriate validations on routing information received in BGP. And if those validations are not performed, the AS legitimate prefix and path are not protected, as peers could select a wrong route announcement and send traffic the wrong way.

Since some proposals use a hierarchical PKI for AS number allocation, the certificates involved in this delegation could be removed for mis-behaving ASes or ASes not following security practices. No proposal studied considered this possibility.

The proposed solutions have different residual vulnerabilities when all actors involved in the solution act accordingly, entailing different risks. Solutions relying on a PKI structure are vulnerable to the misbehavior of the entities involved in the structure and security failures of the dissemination channels. The proposals that do not use such a structure, do not verify the proper allocation of IP addresses and thus only work in case of conflict and even then the wrong announcement can be chosen.

To bring more insights on what happened to the ideas of the studied proposals, tables 8 and 9 summarize the proposals requirements and cryptographic mechanisms used.

The studied solutions to secure BGP have different requirements to be able to yield their security goals. Most solution proposals (11 out of 17) use additional infrastructure to disseminate routing information to verify routing information in BGP messages. The IRR were first solution to propose an out-of-band method to verify data in BGP. The most frequent additional infrastructures to verify BGP information are Public Key Infrastructures, which are used for the issuance and validation of certificates and attestations.

While all solutions require actions to be taken by BGP routers or network administrators when a verification is invalid —not considering a route announcement if the originating AS is not the legitimate prefix holder for instance—, there are six solutions that involve changes to the BGP protocol. Some require the addition of a new message format while others add a new field in BGP Update messages. Eight proposals require the deployment of a Public Key Infrastructure (PKI) to support the issuance of certificates and dissemination of public key for validation. In addition, six solutions necessitate or potentially necessitate additional hardware, either as a storage resource, to compute validations or to monitor the network. Finally, three solutions demand that ASes share their relationship information to check the validity of the AS path of a route, some taking care that the information is not publicly available but distributed on a need-to-know bases.

Any of the requirements shown in table 8 may give rise to strong opponents to

Security Solution	Additional hardware	Additional infrastructure	BGP protocol changes	AS relationship exposure
TCP-MD5	No	No	No	No
TCP-AO	No	No	No	No
GTSM	No	No	No	No
Route Filtering	No	No	No	No
IRRs	No	Yes - Registries	No	No
RPKI & ROAs	No	Yes - PKI	No	No
BGPsec	Potentially	Yes - PKI	Yes	No
Peer-to-Peer Encryption and Predecessor Information	No	No	Yes	No
S-BGP	Potentially	Yes - PKI	No	No
Hop Integrity	No	No	No*	No
Origin Lists	No	No	No+	No
soBGP	Potentially	Yes - PKI	Yes	Yes
IRV	Yes	Yes	No	Limited
SPV	No	Yes - PKI	Potentially	No
Listen and Whisper	No	Yes	Yes	No
psBGP	No	Yes - PKI	Yes	No
ESMs	Yes	Yes - PKI	No	No
Hi-BGP	Yes	Yes - PKI	Yes	Yes

Table 8: BGP security proposal requirements.

the adoption of the proposed BGP security mechanisms, as they directly impact the usual operation of network operators or other networking related organizations such as the RIR. Indeed, designers of proposals that do not have a certain requirement such as modifying BGP, usually believe that requirement is a prohibitive cost that will deter deployment. The proposals focused on TCP security do not need any additional hardware, infrastructure, change to BGP or exposes ASes routing information. Nonetheless, as mentioned earlier, their implementation and deployment has taken a long time.

Most of the studied proposal to secure BGP (12 out of 17) use a cryptographic mechanism, either to attest, authenticate or protect routing information. Of these solutions, eight use certificates and public keys, eight use cryptographic signatures and five use encryption, with seven solution using a mix of them. The first proposal to consider a cryptographic mechanism to secure BGP communication was the work by Smith and Garcia-Luna-Aceves [12], who proposed the use of peer-to-peer encryption for protecting confidentiality and integrity of BGP messages. Since then, many

Security Solution	Certificates	Digital signature	Encryption
TCP-MD5	No	Yes	No
TCP-AO	No	Yes	No
GTSM	No	No	No
Route Filtering	No	No	No
IRRs	Potentially	No	No
RPKI & ROAs	Yes	No	No
BGPsec	Yes	Yes	No
Peer-to-Peer Encryption and Predecessor Information	No	Yes	Yes
S-BGP	Yes	Yes	Yes
Hop Integrity	No	No	Yes
Origin Lists	No	No	No
soBGP	Yes	No	No
IRV	No	Potentially	Potentially
SPV	Yes	Yes*	No
Listen and Whisper	No	Yes	No
psBGP	Yes	Yes	Yes
ESMs	Yes	No	Yes
Hi-BGP	Yes	Yes	Potentially

Table 9: BGP security proposal mechanisms used.

proposal incorporate a mechanism with similar goals, but not all. The protection of BGP communication and message integrity is another focus of disagreement. Not all designer of proposals agree that protection is worth the performance cost. Indeed, even the Best Current Practice for BGP security publish by the IETF are undefined with respect to the trade-offs of such protection [55].

It comes to no surprise that there is no BGP security proposal without costs, and most of the cost is carried by network operators. All solution involve some level of modification of BGP speaker configuration or use devices acting as proxy that need extensive configuration, increasing the risk of configuration mistakes impacting availability.

One valid question is if it is the role of network operators to provide secure routing. If network operators and ISPs provide routing but not secure routing, there are still ways to increase routing security, for instance by monitoring BGP routing information. If some threat is detected by the monitoring system, an alarm could be raised

and action taken depending on the level of exposure. In many settings other than routing, security is provided by external monitoring that is usually outsourced. In routing, the same can be done. In fact, there are companies such as ThousandEyes [77] that offer BGP monitoring services. ThousandEyes deploys its own monitors located across the Internet, to gather BGP messages and analyze routing information for their customers. In this setting, the security proposals based on monitoring could be used. As an example, when Amazon Web Services (AWS) where the target of a BGP hijacking attack in April 2018, ThousandEyes detected it in their system [6]. Companies like Akamai [78] and Cloudflare [79] are also well positioned to offer routing security services to their customers as part of their services, since they have visibility of such a large part of the public Internet.

Another possibility is that secure routing be offered as an additional service by ISPs and network operators, for customers willing to pay the premium. Customers wanting better security guarantees for the traffic to and from their network could request an ISP an additional level of security for their routing. ISP and network operators cloud convey this by following all BGP operational security best practices [55] and specifically monitoring routes for these customers.

In this case,the ISP security effort would be pretty narrow in scope, limiting misconfiguration mistakes. In addition, the extra cost incurred by the ISP is directly passed to the customers.

If secure routing is not necessarily the role of network operators, customers would have to pay for securing their routing, which they might not be willing to do. Regulation might give incentives to customers providing critical services to end-user that run on top of the public Internet, such as online banking, to secure their routing through monitoring, best practices and updating their routing information in the IRRs, ideally including a valid ROA.

6 Conclusion

This paper studied the lifetime of proposals to secure BGP, the inter-domain routing protocol used in the public Internet. These proposals come from the IETF, industry and the academia. Extensions done to BGP focused on availability were also considered as availability is an aspect of information security and a major goal of routing protocols. Indeed, solutions that improve BGP and routing availability are deployed and used, whereas almost no solution that can negatively impact availability is currently in use.

In fact, no solution focusing on the correctness of BGP and BGP routing infor-

mation has been quickly implemented and deployed. As an example, it took at least 10 years to for the TCP level security mechanism TCP-MD5 to go beyond limiter usage, and it is a much simpler mechanism than most proposals to secure BGP.

Additionally, the main motivations of proposals make evident the lack of agreement with respect to what securing BGP means: should BGP communication be protected or routing information validated? Validated by whom? Just in the case of conflict or all the time? Light-weight solutions propose means to detect conflict easily and use BGP routing information to decide the correct answer, whereas other solutions propose to block route announcements whose validity cannot be verified using the public information of Public Key Infrastructure in charge of IP addresses and AS number allocations.

At the time of writing, the best way to secure BGP routing is through the use of BGP security best practices: applying route filters, using and updating routing information available in the Internet Routing Registries (IRRs), and monitoring.

However, secure routing may not need to be the default service offered by Internet Service Providers (ISPs), it can be an additional service offered by ISPs or external companies. In fact, there are companies that provide BGP security monitoring services based on their own deployed infrastructure. Content-delivery network (CDN) companies could also provide BGP security services to their customers based on their network infrastructure. For organizations providing critical services relying on the public Internet, regulation cloud create the incentives for them to use routing security services.

References

[1] O. Nordström and C. Dovrolis, "Beware of BGP attacks," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 2, pp. 1–8, 2004.

[2] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *ACM SIGCOMM Computer Communication Review*, vol. 32. ACM, 2002, pp. 3–16.

[3] N. Feamster, J. Jung, and H. Balakrishnan, "An empirical study of bogon route advertisements," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 1, pp. 63–70, 2005.

[4] H. Ballani, P. Francis, and X. Zhang, "A study of prefix hijacking and interception in the Internet," in *ACM SIGCOMM Computer Communication Review*, vol. 37. ACM, 2007, pp. 265–276.

[5] M. Apostolaki, A. Zohar, and L. Vanbever, "Hijacking Bitcoin: Routing Attacks on Cryptocurrencies," in *2017 IEEE Symposium on Security and Privacy (SP)*. San Jose, CA, USA: IEEE, May 2017, pp. 375–392. [Online]. Available: http://ieeexplore.ieee.org/document/7958588/

[6] A. Naik, "Anatomy of a BGP Hijack on Amazon's Route 53 DNS Service," Apr. 2018. [Online]. Available: https://blog.thousandeyes.com/amazon-route-53-dns-and-bgp-hijack/

[7] A. Heffernan, "RFC 2385: Protection of BGP Sessions via the TCP MD5 Signature Option," Aug. 1998. [Online]. Available: https://tools.ietf.org/html/rfc2385

[8] D. Meyer, J. Heasley, and V. Gill, "RFC 3682: The Generalized TTL Security Mechanism (GTSM)," Feb. 2004. [Online]. Available: https://tools.ietf.org/html/rfc3682

[9] J. Touch, A. Mankin, and R. P. Bonica, "RFC 5925: The TCP Authentication Option," Jun. 2010. [Online]. Available: https://tools.ietf.org/html/rfc5925

[10] M. Lepinski, R. Barnes, and S. Kent, "RFC 6480: An Infrastructure to Support Secure Internet Routing," Feb. 2012. [Online]. Available: https://tools.ietf.org/html/rfc6480

[11] M. Lepinski and K. Sriram, "RFC 8205: BGPsec Protocol Specification," Sep. 2017. [Online]. Available: https://tools.ietf.org/html/rfc8205

[12] B. R. Smith and J. J. Garcia-Luna-Aceves, "Securing the border gateway routing protocol," in *Global Telecommunications Conference, 1996. GLOBECOM '96. 'Communications: The Key to Global Prosperity*, Nov. 1996, pp. 81–85.

[13] 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.

[14] M. G. Gouda, E. N. Elnozahy, C.-T. Huang, and T. M. McGuire, "Hop integrity in computer networks," *IEEE/ACM Transactions on Networking*, vol. 10, no. 3, pp. 308–319, Jun. 2002.

[15] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, "Detection of invalid routing announcement in the Internet," in *Proceedings International Conference on Dependable Systems and Networks*, 2002, pp. 59–68.

[16] R. White, "Securing BGP Through Secure Origin BGP - The Internet Protocol Journal - Volume 6, Number 3," *The Internet Protocol Journal*, vol. 6, no. 3, Sep. 2003. [Online]. Available: https://www.cisco.com/c/en/us/about/press/internet-protocol-journal/back-issues/table-contents-25/securing-bgp-sobgp.html

[17] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. D. McDaniel, and A. D. Rubin, "Working around BGP: An Incremental Approach to Improving Security and Accuracy in Interdomain Routing." in *ISOC Symposium on Network and Distributed Systems Security*, vol. 23, 2003, p. 156.

[18] Y.-C. Hu, A. Perrig, and M. Sirbu, "SPV: Secure path vector routing for securing BGP," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, pp. 179–192, 2004.

[19] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz, "Listen and Whisper: Security Mechanisms for BGP," in *1st Symposium Networked System Design and Implementation,*, 2004, p. 14.

[20] T. Wan, E. Kranakis, and P. C. van Oorschot, "Pretty Secure BGP, psBGP." in *Proceedings of the 2005 ISOC Symposium on Network and Distributed Systems Security*, San Diego, 2005.

[21] P. Reynolds, O. Kennedy, E. G. Sirer, and F. B. Schneider, "Using External Security Monitors to Secure BGP," 2006.

[22] J. Qiu and L. Gao, "Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol," p. 12, 2006.

[23] M. Zhao, S. W. Smith, and D. M. Nicol, "The performance impact of BGP security," *IEEE network*, vol. 19, no. 6, pp. 42–48, 2005.

[24] D. M. Nicol, S. W. Smith, and M. Zhao, "Evaluation of efficient security for BGP route announcements using parallel simulation," *Simulation Modelling Practice and Theory*, vol. 12, no. 3-4, pp. 187–216, Jul. 2004. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S1569190X04000383

[25] K. Butler, T. 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, Jan. 2010. [Online]. Available: http://ieeexplore.ieee.org/document/5357585/

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

[27] G. Huston, M. Rossi, and G. Armitage, "Securing BGP — A Literature Survey," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 2, pp. 199–222, 2011. [Online]. Available: http://ieeexplore.ieee.org/document/5473881/

[28] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, "How secure are secure interdomain routing protocols?" *Computer Networks*, vol. 70, pp. 260–287, Sep. 2014. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S1389128614001984

[29] R. Lychev, M. Schapira, and S. Goldberg, "Rethinking security for internet routing," *Communications of the ACM*, vol. 59, no. 10, pp. 48–57, Sep. 2016. [Online]. Available: http://dl.acm.org/citation.cfm?doid=3001840.2896817

[30] M. Siddiqui, D. Montero, R. Serral-Gracià, X. Masip-Bruin, and M. Yannuzzi, "A survey on the recent efforts of the Internet Standardization Body for securing inter-domain routing," *Computer Networks*, vol. 80, pp. 1–26, Apr. 2015. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S1389128615000286

[31] IETF, "Home." [Online]. Available: /

[32] Y. Rekhter and K. Lougheed, "RFC 1105: A Border Gateway Protocol (BGP)," Jun. 1989. [Online]. Available: https://tools.ietf.org/html/rfc1105

[33] ——, "RFC 1163: A Border Gateway Protocol (BGP-2)," Jun. 1990. [Online]. Available: https://tools.ietf.org/html/rfc1163

[34] Y. Rekhter and T. Li, "RFC 1267: A Border Gateway Protocol 3 (BGP-3)," Oct. 1991. [Online]. Available: https://tools.ietf.org/html/rfc1267

[35] ——, "RFC 1654: A Border Gateway Protocol 4 (BGP-4)," Jul. 1994. [Online]. Available: https://tools.ietf.org/html/rfc1654

[36] ——, "RFC 1771: A Border Gateway Protocol 4 (BGP-4)," Mar. 1995. [Online]. Available: https://tools.ietf.org/html/rfc1771

[37] S. Hares, Y. Rekhter, and T. Li, "RFC 4271: A Border Gateway Protocol 4 (BGP-4)," Jan. 2006. [Online]. Available: https://tools.ietf.org/html/rfc4271

[38] S. Murphy, "RFC 4272: BGP Security Vulnerabilities Analysis," Jan. 2006. [Online]. Available: https://tools.ietf.org/html/rfc4272

[39] Y. Rekhter and S. R. Sangli, "RFC 4360: BGP Extended Communities Attribute," Feb. 2006. [Online]. Available: https://tools.ietf.org/html/rfc4360

[40] D. D. Clark and D. R. Wilson, "A Comparison of Commercial and Military Computer Security Policies," in *1987 IEEE Symposium on Security and Privacy*. Oakland, CA, USA: IEEE, Apr. 1987, pp. 184–184. [Online]. Available: http://ieeexplore.ieee.org/document/6234890/

[41] P. Traina, "RFC 1965: Autonomous System Confederations for BGP," Jun. 1996. [Online]. Available: https://tools.ietf.org/html/rfc1965

[42] D. McPherson, P. Traina, and J. G. Scudder, "RFC 3065: Autonomous System Confederations for BGP," Feb. 2001. [Online]. Available: https://tools.ietf.org/html/rfc3065

[43] D. McPherson and J. G. Scudder, "RFC 5065: Autonomous System Confederations for BGP," Aug. 2007. [Online]. Available: https://tools.ietf.org/html/rfc5065

[44] T. Bates, "RFC 1966: BGP Route Reflection An alternative to full mesh IBGP," Jun. 1996. [Online]. Available: https://tools.ietf.org/html/rfc1966

[45] R. Chandra, E. Chen, and T. Bates, "RFC 2796: BGP Route Reflection - An Alternative to Full Mesh IBGP," Apr. 2000. [Online]. Available: https://tools.ietf.org/html/rfc2796

[46] P. Traina, "RFC 1997: BGP Communities Attribute," Aug. 1996. [Online]. Available: https://tools.ietf.org/html/rfc1997

[47] Y. Rekhter, "RFC 5701: IPv6 Address Specific BGP Extended Community Attribute," Nov. 2009. [Online]. Available: https://tools.ietf.org/html/rfc5701

[48] J. Snijders, I. Bagdonas, K. Patel, J. Heitz, and N. Hilliard, "RFC 8092: BGP Large Communities Attribute," Feb. 2017. [Online]. Available: https://tools.ietf.org/html/rfc8092

[49] E. Chen, "RFC 2918: Route Refresh Capability for BGP-4," Sep. 2000. [Online]. Available: https://tools.ietf.org/html/rfc2918

[50] R. Chandra and J. G. Scudder, "RFC 2842: Capabilities Advertisement with BGP-4," May 2000. [Online]. Available: https://tools.ietf.org/html/rfc2842

[51] R. Fernando, S. R. Sangli, Y. Rekhter, E. Chen, and J. G. Scudder, "RFC 4724: Graceful Restart Mechanism for BGP," Jan. 2007. [Online]. Available: https://tools.ietf.org/html/rfc4724

[52] C. Villamizar, R. Govindan, and R. Chandra, "RFC 2439: BGP Route Flap Damping," Nov. 1998. [Online]. Available: https://tools.ietf.org/html/rfc2439

[53] P. Savola, V. Gill, C. Pignataro, D. Meyer, and J. Heasley, "RFC 5082: The Generalized TTL Security Mechanism (GTSM)," Oct. 2007. [Online]. Available: https://tools.ietf.org/html/rfc5082

[54] M. Kaeo, "RFC 4778: Operational Security Current Practices in Internet Service Provider Environments," Jan. 2007. [Online]. Available: https://tools.ietf.org/html/rfc4778

[55] G. Doering, J. Durand, and I. Pepelnjak, "RFC 7454: BGP Operations and Security," Feb. 2015. [Online]. Available: https://tools.ietf.org/html/rfc7454

[56] L. Zheng, K. Patel, and M. Jethanandani, "RFC 6952: Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide," May 2013. [Online]. Available: https://tools.ietf.org/html/rfc6952

[57] Merit, "IRR - Internet Routing Registry." [Online]. Available: http://www.irr.net/

[58] T. Bates, D. Meyer, D. Karrenberg, M. Terpstra, E. Gerich, and C. Alaettinoglu, "RFC 2280: Routing Policy Specification Language (RPSL)," Jan. 1998. [Online]. Available: https://tools.ietf.org/html/rfc2280

[59] S. Murphy, C. Villamizar, C. Alaettinoglu, and D. M. Meyer, "RFC 2725: Routing Policy System Security," Dec. 1999. [Online]. Available: https://tools.ietf.org/html/rfc2725

[60] J. W. Stewart, *BGP4: Inter-domain Routing in the Internet*. Addison Wesley, 1999.

[61] R. NCC, "RIPE Network Coordination Centre." [Online]. Available: https://www.ripe.net

[62] MANRS, "MANRS." [Online]. Available: https://www.manrs.org/

[63] IETF, "Secure Inter-Domain Routing (sidr) - Documents." [Online]. Available: https://datatracker.ietf.org/wg/sidr/documents/

[64] G. Huston, G. Michaelson, and R. Loomans, "RFC 6481: A Profile for Resource Certificate Repository Structure," Feb. 2012. [Online]. Available: https://tools.ietf.org/html/rfc6481

[65] M. Lepinski, D. Kong, and S. Kent, "RFC 6482: A Profile for Route Origin Authorizations (ROAs)," Feb. 2012. [Online]. Available: https://tools.ietf.org/html/rfc6482

[66] R. Bush, "RFC 8207: BGPsec Operational Considerations," Sep. 2017. [Online]. Available: https://tools.ietf.org/html/rfc8207

[67] R. Bush and R. Austein, "RFC 8210: The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1," Sep. 2017. [Online]. Available: https://tools.ietf.org/html/rfc8210

[68] N. I. for Santards and Technology, "RPKI Deployment Monitor." [Online]. Available: https://rpki-monitor.antd.nist.gov/

[69] A. Khan, H.-c. Kim, T. Kwon, and Y. Choi, "A comparative Study on IP Prefixes and their Origin ASes in BGP and the IRR," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 3, pp. 16–24, 2013. [Online]. Available: http://dl.acm.org/citation.cfm?id=2500101

[70] R. Kisteleki and B. Haberman, "RFC 7909: Securing Routing Policy Specification Language (RPSL) Objects with Resource Public Key Infrastructure (RPKI) Signatures," Jun. 2016. [Online]. Available: https://tools.ietf.org/html/rfc7909

[71] R. White and B. Akyol, "RFC 5123: Considerations in Validating the Path in BGP," Feb. 2008. [Online]. Available: https://tools.ietf.org/html/rfc5123

[72] K. Sriram, "RFC 8374: BGPsec Design Choices and Summary of Supporting Discussions," Apr. 2018. [Online]. Available: https://tools.ietf.org/html/rfc8374

[73] V. K. Sriram and D. Montgomery, "Design and analysis of optimization algorithms to minimize cryptographic processing in BGP security protocols," *Computer Communications*, vol. 106, pp. 75–85, Jul. 2017. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0140366417303365

[74] R. White, "BGPsec and Reality," Oct. 2017. [Online]. Available: https://rule11.tech/bgpsec-and-reality/

[75] R. Atkinson and S. Kent, "RFC 2401: Security Architecture for the Internet Protocol," Nov. 1998. [Online]. Available: https://tools.ietf.org/html/rfc2401

[76] Y.-C. Hu, D. McGrew, A. Perrig, B. Weis, and D. Wendlandt, "(R)Evolutionary Bootstrapping of a Global PKI for Securing BGP," p. 6, 2006.

[77] ThousandEyes, "Network Intelligence Software." [Online]. Available: https://www.thousandeyes.com/

[78] Akamai, "Cloud Delivery, Performance, and Security | Akamai." [Online]. Available: https://www.akamai.com/

[79] Cloudflare, "Cloudflare - The Web Performance & Security Company." [Online]. Available: https://www.cloudflare.com/