The ISP Column
A column on various things Internet

July 2021
Geoff Huston

A Survey on Securing Inter-Domain Routing

Part 2 – Approaches to Securing BGP

'a'The Border Gateway Protocol (BGP) is the Internet's inter-domain routing protocol, and after some thirty years of operation BGP is now one of the more venerable of the Internet's protocols. One of the major ongoing concerns related to BGP is its lack of effective security measures, and as a result the routing infrastructure of the Internet continues to be vulnerable to various forms of attack.

In Part 1 we looked at the design of BGP, the threat model and the requirements from a security framework for BGP. In Part 2 we will look at the various proposals to add security to the routing environment and also review the current state of the effort in the Internet Engineering Task Force (IETF) to provide a standard specification of the elements of a secure BGP framework.

The approaches to securing BGP can be further classified in the same fashion as the security requirements: securing the operation of BGP and securing the integrity of the BGP data.

1. Securing the Operation of BGP Sessions

BGP uses a long-held TCP session and the same approaches to securing any TCP session [1] can be used in the context of a BGP session. These approaches fall into two categories: those that simply attempt to protect the TCP session from disruption via injection of spurious traffic, and those that also attempt to protect the TCP session from eavesdropping and alteration by encrypting the payload.

1.1 Generalized TTL Security Mechanism

The Generalized TTL Security Mechanism (GSTM) was originally described in [2] and updated in [3] and is based on the observation that the overall majority of BGP peering sessions are established between routers that are directly connected. The technique is to configure each BGP IP packet to be sent with a TTL field value in the IP header of 255, and for the BGP receiver to discard all packets with an inbound TTL of less than a set threshold value. For a direct connection the inbound TTL value should be 255, so all inbound TCP packets with within this session with a TTL of 254 or less can be discarded by the receiver.

The motivation for this approach is that spoofing of the TTL field in an IP header is challenging for an unassisted remote attacker. This TTL packet filter is a lightweight defensive measure intended to add some protection to the BGP session from efforts to intrude into the session using remote attacks. This GTSM approach can be used for multi-hop BGP peer sessions as well as directly connected BGP sessions, but it is not all that robust in terms of its security properties because of the additional variables introduced with TTL changes due to routing changes and the potential to mask the conventional TTL behaviour with tunnelling techniques.

1.2 TCP MD5 Signature Option

A more robust approach to protecting the TCP session is through the use of cryptographic protection of the TCP session. While these crypto approaches can be highly resilient to intrusion attempts, they also expose the

BGP speaker to potential denial of service attacks if the processing load of the cryptographic functions to detect bogus packets is sufficiently high. Bogus packets still have to be processed by the target just to ascertain that they are bogus.

The TCP MD5 Signature Option [4] uses message authentication codes, which are a class of cryptographic hash algorithms applied to messages of arbitrary length that produce a "message digest" of the message, intended to protect the integrity of a message. The desired property of a message digest is that it is infeasible to generate two messages that have the same message digest value, and equally infeasible to generate a new message that has a particular message digest value. The MD5 algorithm [5] is intended for digital signature applications where a message digest is generated over the combination of a message and a secret shared key value. The message and the digest value can be transmitted openly, and the receiver can use a local copy of the secret key and apply the message digest algorithm to the combination of the received message and the key. If the digest value matches the received value, then the receiver can be assured that the message has not been altered in transit, and that the message was generated by a party who also has knowledge of the key.

The TCP MD5 Signature option is a TCP extension where each TCP segment contains a TCP option that contains the 128-bit MD5 digest of the combination of the TCP pseudo header, the TCP segment payload excluding TCP options, and a connection-specific key. This establishes a cryptographically secure signature of the packet. Without knowing the key, it is very challenging to construct a TCP segment with a valid signature, nor is it readily possible to alter the packet without causing the signature to be invalidated. The receiver calculates the MD5 digest across the received data, using a locally held copy of the key, and rejects the segment if the digest value fails to match that provided in the packet. In the context of BGP the TCP session is resistant to various forms of intrusion attack unless the attacker has knowledge of the shared secret key value. The TCP MD5 specification does not specify how the shared key is passed between the two BGP speakers, nor how the key value can be changed during the session. This latter problem is significant, in that continued use of a key weakens its integrity, and it is conventionally advised that MD5 session keys should be changed every 90 days or so in this type of use context [6]. With a mechanism for in-band key change, this advice implies the need for a BGP session reset every 90 days or so, which is counter to conventional operational practice in BGP where sessions are held up for as long as possible. Even with tools such as BGP graceful restart, deliberate BGP session resets are generally avoided in the operational community.

1.3 TCP Authentication Option

A somewhat different approach, the TCP Authentication Option [7], uses a Message Authentication Field in the place of the MD5 message digest, where the final bit of the length field of the option determines whether a key ID has been appended to the Message Authentication Code or not. The message digest algorithm in this case is specified as HMAC-MD5- 96, although other algorithms can be used if configured in advance. This approach relies on a similar form of out-of-band provisioning as the original MD5 approach, where each end of the conversation must configure a TCP Security Association Database in advance of the use of this mechanism. This database contains a description of the supported TCP connections, the key set, the MAC algorithm, and MAC length.

1.4 IPSEC

IPSEC is a suite of protocols that operate at the IP level of the protocol stack that secures all communications between two endpoints [8]. The functionality of IPSEC includes methods for protection of IP packet headers, methods for protection and encryption of IP payloads and key management services that allow key rollover during long held sessions. This is an implementation of public/private key cryptography and can ensure the confidentiality and integrity of all IP messages passed between two hosts. IPSEC can be used to secure BGP sessions, and it provides greater levels of assurance than can be derived from MD5.

However, IPSEC is not widely used in the public Internet for the purpose of securing BGP sessions [79], and no generally accepted profile of IPSEC for BGP has been standardised so far, with earlier efforts along these lines not progressing within the standards process. The perceived problem with IPSEC relate to the complications for rekeying IKE/IPSEC sessions, and the observation that processing load to detect bogus packets is considerably higher with IPSEC than MD5. This exposes a denial-of-service attack where a stream of bogus IPSEC packets directed at a BGP speaker may be capable of exercising the processor into a fully saturated mode of operation, causing other concurrent router functions to be degraded.

1.5 More Options

As was observed in Part 1 of this servey, there are many alternatives here, including TLS [80] and QUIC [81], but more choice is not a substitute for better quality. These session-level encryption approaches used by applications provide no better answer to dynamic rekeying and follow a now well-established Internet tradition of adding more options to divert attention from the observation that the common fundamental problems are inadequately addressed many or either all such options! The design goal of such application-level session approaches is protection for transient short-duration sessions, while the vulnerabilities associated with long-held BGP sessions are somewhat different.

The best advice today is that a combination of TCP AO and GTSM is as good as it gets at present. However, it's also highly desirable to avoid multi-hop BGP wherever possible and directly attach the two BGP speakers. That way the radius of potential eavesdroppers and attackers is reduced considerably.

2. Securing the Integrity of Routing Information passed in BGP

2.1. Early Work: 1988 - 2000

One of the earlier recognised works that addressed routing security was the 1988 study on *Byzantine Robustness* by Radia Perlman [9]. In the event of failure or malicious behaviour on the part of one or more entities in the system, all correctly operating entities should reach a mutually consistent decision regarding the validity of each message in finite time. This study was in the area of link-state protocol design, and the work described a protocol that satisfied the properties for *Byzantine Robustness*. It categorised route validation in three approaches:

- *Bound* or just in time — validation occurs the same moment a route is announced, and appropriate measures are taken immediately. Credentials must be available immediately.
- *Unbound* or just in case — validation occurs only if a new router takes part in the system. Credentials are retrieved on arrival of this router.
- *Interrogative* or just too late — validation occurs on a sporadic base, requesting validation or credentials from a remote system when necessary.

While the link-state approach described in this paper does not exactly match the interdomain routing environment, the concept of validation of routing information is a consistent theme in all BGP security architectures.

Subsequent work by Smith and Garcia-Luna-Aceves [10], [11], published in 1996, attempts to address session security by modifying the BGP protocol. This work proposed the protection of BGP control messages using message encryption at the BGP level, with session keys exchanged at BGP session establishment time. It also proposed the addition of a message sequence number to protect against replay attacks and message removal. This approach also proposed a predecessor path attribute that indicated the AS prior to the destination AS for the current route and proposed digitally signing all fixed fields in the UPDATE message. The predecessor attribute is used to construct a means of validation of the AS Path attribute. These proposed changes to the BGP protocol required comprehensive adoption and deployment in order to be effective, as partial adoption would create gaps in any assurance that a predecessor attribute could provide. Their approach was similar to the earlier IDRP work [12]. IDRP eschewed the use of TCP and included a reliable flow-controlled transport into the IDRP protocol, also including a number of message integrity protection options.

A contemporary proposal to the Smith and Garcia-Lunes-Aceves proposal for securing BGP was based on leaving the BGP protocol unchanged but augmenting the BGP data flow with access to credential information. This additional information was intended to allow a BGP speaker to confirm the authenticity of origination information in BGP UPDATE messages by validating the binding of address prefixes to originating ASes [13]. This proposal, NLRI origin AS verification, used the DNS as the distribution mechanism for origination information, where a BGP speaker could perform a DNS query to validate the prefix size and authorised originating AS information contained in a BGP route object. Informally, it was intended to allow a DNS query to answer the question: "Which ASes have been authorised by the address holder to originate a route for this prefix?" The proposed framework assumed that the reverse DNS space was securely associated with the holder of the address prefix, and the DNS response was verifiable (using a DNSSEC-signed DNS record and DNSSEC validation [14], presumably, although this work was contemporaneous with DNSSEC and did not make use of it in this proposal). This proposal assumed that the performance of DNS queries was within the same order of timescale as the propagation of BGP messages within BGP. It also assumed that there was no circularity, where a DNS recursive resolver or authoritative name server used by the BGP speaker was located within an address prefix that was being validated prior to local acceptance of the route associated that that prefix.

The DNS delegation hierarchy would need to be precisely aligned to the address allocation framework, so that the zone administrator of each of these origination authentication zones was in fact the duly delegated holder of the addresses, and this alignment should, preferably, be capable of being validated by third parties. Meeting these requirements would create a digital signature hierarchy embedded in the DNS that would be aligned to the address allocation framework.

The Internet Routing Registry (IRR) [82] pre-dates most other efforts in this space, and dates back to the routing work of the early 1990's in the Routing Arbiter project that was part of the US NSFNET, and a project coordinated under the auspices of RIPE in Europe. The IRR objective was to provide a set of routing policy databases populated by the ASes themselves that described the addresses that they intended to announce in the routing system and the routing policies that they intended to apply to these announcements [83]. The Routing Registry was to a response to the need described in RFC 1787 [84] for improving global consistency by allowing providers share routing policies. Each participating AS submits policy data, encoded using the Routing Policy Specification Language (RPSL) [27] [28]. Clients may use the registry to determine the stated policies for a particular AS, including what ASes (and possibly prefixes) are suitable for import or export, potentially using the data to populate filter sets on their BGP feeds. Additional information provided to the IRR by an AS could include policy concerning the configuration of BGP communities and the policy responses associated with particular community settings.

However, the utility of the IRR for securing routing is quite limited. First, the IRR does not provide information about current routes, but only about potential routes. Some potential routes may be legal according to the IRR, but undesirable from a more global point of view. Next, the IRR has many security vulnerabilities concerning the integrity of registry contents and authorization of changes to the registry. There is no intrinsic authority model that constrains which party can publish data about addresses and ASes in an IRR. Moreover, some policy information concerning agreements between peering ASes is not intended for broader public distribution and the IRRs did not normally implement any form of limited disclosure rules. Efforts to improve the controls over the authority framework in registries and access frameworks [85] never really gained traction. The IRR system is a misnomer, in that there is not a single IRR but many IRRs. The contents of these IRRs are not necessarily mutually consistent and there is no clear way to resolve any such conflicts. Not only is there no authority model ensuring that only authorised parties may publish routing policy data about their own address prefixes and ASes, but there is also no way to describe the intended lifetime of the information. Old information that is no longer current or relevant sits alongside current information, and this sits along with contingency information that may never be actually used. While the overall approach of providing an out-of-band commentary on routing, enumerating all the cases of authorised (or valid) route objects has been a useful tool for many operational environments, IRR tools are only truly useful in the context of being able to detect and filter routing anomalies if the information is verifiable and authentic, current and complete. In other words, IRRs are most useful it that are carefully and continuously managed, and the accuracy and usefulness of the information rapidly declines if the information in the registry is neglected. Our experience with IRRs suggests that it would be somewhat foolhardy to automatically apply IRR data to populate route filters, given the risks of incorrect outcomes, both positive and negative, and while there have been good counter examples in some operational communities, the broader judgement for IRRs being capable of supporting a robust whole-of-internet role for route integrity is somewhat negative [86].

It appears that the common requirements in this space appear to relate to *authenticity*, *currency* and *completeness*.

Digital signatures can provide strong assurance related to authenticity and currency of information, assuming that there is robust enrolment practice than governs the authority to generate such signatures. Given such a practice, the consequent observation is that whether this digital signature framework is placed into the DNS, via a DNSSEC framework [15], or placed into a framework of X.509 certificates and an associated PKI is, at one level, an isomorphic transform of the same information. The issue of the choice of DNS (and DNSSEC) or X.509 certificates (and certificate-based validation) is then an issue of the performance requirements of these systems.

Completeness is a more challenging requirement. The identification of invalid routing information in the partial adoption case of this approach is unclear. When a query to an information source has a negative response, it is unclear whether the route object that was the basis of the query is not valid (such as a bogus prefix, or a bogus AS), or whether the database being queried is incomplete.

Let's now move forward in time to review some more recent proposals to secure BGP.

2.2 Secure BGP

Secure BGP (sBGP) [16] offered a relatively complete approach to securing the BGP protocol by placing digital signatures over the address and AS Path information contained in routing advertisements and defining an associated PKI for validation of these signatures.

sBGP defines the "correct" operation of a BGP speaker in terms of a set of constraints placed on individual protocol messages, including ensuring that:

- all protocol UPDATE messages have not been altered in transit between the BGP peers,
- the UPDATE messages were sent by the indicated peer,
- the UPDATE messages contain more recent information than has been previously sent to this BGP speaker from the peer,
- the UPDATE was intended to be received by this BGP speaker, and
- that the peer is authorised to advertise information on behalf of the peer AS.

In addition, for every prefix and it's originating AS, the prefix must be a validly allocated prefix, and the prefix's "right-of-use" holder must have authorised the advertisement of the prefix and must have authorised the originating AS to advertise the prefix.

The basic security framework proposed in sBGP is that of digital signatures, X.509 certificates and PKIs to enable BGP speakers to verify the identities and authorisation of other BGP speakers, AS administrators and address prefix owners.

The verification framework for sBGP requires a PKI for address allocations, where every address assignment is reflected in an issued certificate [17]. This PKI provides a means of verification of a "right-of-use" of an address. A second PKI maps the assignment of ASes, where an AS number assignment is reflected in an issued certificate, and the association between an AS number and a BGP speaking router is reflected in a subordinate certificate. In addition, sBGP proposes the use of IPSEC to secure the inter-router communication paths.

sBGP also proposes the use of *attestations*. An *address attestation* is produced by an address holder and authorises a nominated AS to advertise itself as the origin AS for a particular address prefix. A *route attestation* is produced by an AS holder and attests that a BGP speaker is an authorised member of that AS and that it has received a specified route. The address and AS PKIs, together with these attestations, allow a BGP speaker to verify the origination of a route advertisement and verify that the AS path as specified in the BGP UPDATE is the path taken by the routing UPDATE message via the sequence of nested route attestations.

Inter-operation and information exchange between sBGP elements is shown in Figure 1.

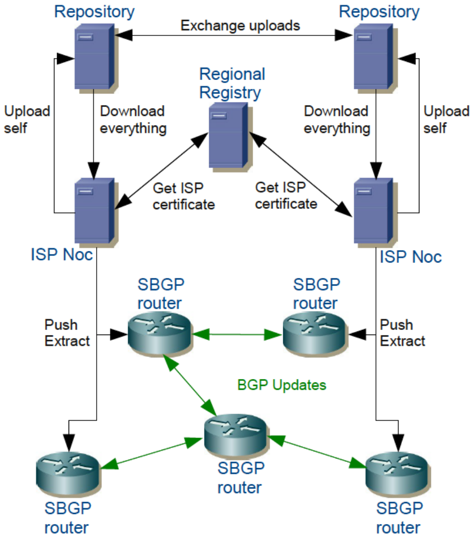

Figure 1 - sBGP: Certificates for each ISP are issued by the regional registries. The ISPs exchange public keys through special repositories. The keys are pushed to the sBGP routers which validate the BGP UPDATE messages

sBGP proposed to distribute the *address attestations* and the set of certificates that compose the two PKIs via conventional distribution mechanisms outside of BGP messages. For *route attestations* it is necessary to pass these attestations via path attributes of the BGP UPDATE message, as an additional attribute of the UPDATE message.

There are a number of significant issues that were identified with sBGP including the computation burden for signature generation and validation, the increased load in BGP session restart, the issue of piecemeal deployment and the completeness of route attestations, and the requirement that the BGP UPDATE message has to traverse the same AS sequence as that contained in the UPDATE message [18], [19].

2.3 Secure Origin BGP

Secure Origin BGP (soBGP) [20] [21] was a response to some of the significant issues that were raised with the sBGP approach, particularly relating to the update processing load when validating the chain of router attestations and the potential overhead of signing every advertised UPDATE with a locally generated router attestation.

The validation questions posed by soBGP also included the notion of an explicit authorisation from the address holder to the originating AS to advertise the prefix into the routing system. soBGP's AS path validation is quite different from sBGP, in that soBGP attempted to validate that the AS path, as presented in the UPDATE message, represents a feasible inter-AS path from the BGP speaker to the destination AS. This feasibility test is a weaker validation condition than validating that the UPDATE message actually traversed the AS path described in the message.

soBGP avoids the use of a hierarchical PKI that mirrors the AS number distribution framework and nominates the use of a web of trust, or a reputation mechanism, as means of validation of these certificates. At the time no Address or AS PKI had been devised or deployed, so this web of trust approach was a pragmatic response to this critical omission. soBGP uses the concept of an *AuthCert* to bind an address prefix to an originating AS. This *AuthCert* is not signed by the address holder, but by a private key that is bound to an AS via an *EntityCert.* soBGP deliberately avoided the use of a PKI which was derived from the established AS and address distribution framework. This appears to have been a pragmatic consideration at the time, as no such PKI existed at the time, and it was unclear if the various address registries were in a position to undertake such a role of administering such a specialised PKI in any case. This left open the issue of how to establish trust anchors for validation of these signed objects, which was a rather significant deficiency in the validation framework of soBGP.

Instead of sBGP's route attestations, soBGP used the concept of an *ASPolicyCert* as the foundation for constructing the data for testing the feasibility of a given AS Path. An *ASPolicyCert* contained a list of the AS's local peer ASes, signed by the AS's private key. An AS peering was considered valid only if both ASes list each other in their respective *ASPolicyCerts*. Figure 2 depicts a possible soBGP peering network.

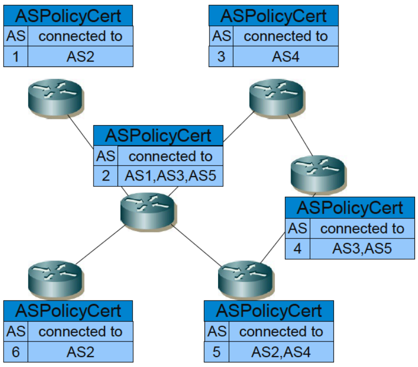

Figure 2 - The ASPolicyCert is a self signed certificate containing routing policies. An UPDATE message originating at AS4 would necessarily take the Path {AS4,AS5,AS2,AS1} instead of {AS4,AS3,AS2,AS1} becuase the connection between AS2 and AS3 would not be regarded as valid.

The overall approach proposed in soBGP represented a different set of design trade-offs to sBGP, where the amount of validated material in a BGP UPDATE message is reduced. This approach was intended to reduce the processing overhead for validation of UPDATE messages. In soBGP each local BGP speaker assembles a validated inter-AS topology map as it collects *ASPolicyCerts*, and each AS Path in UPDATE messages is then checked to see if the AS sequence matches a feasible inter-AS path in this map. soBGP proposed to use BGP itself to flood *ASPolicyCerts* through the network, using a new BGP message type (a *Security* Message) for this function.

The use of *Web of* Trust and the avoidance of a hierarchical PKI for the validation of *AuthCerts* and *EntityCerts* could be considered a weakness in this approach, as the derivation of authority to speak about addresses is very unclear in this model, but this absence was a result of the protocol being developed prior to the completion of the work on the RPKI. It is clear that soBGP could be readily adapted to use the RPKI as its trust and authority framework.

soBGP's use of BGP itself to flood the security credentials though the network represented an interesting approach to the problem of distributing such credentials, but it also raised some at the time unanswered questions relating to partial deployment scenarios. Interest in continuing work on soBGP waned in the early 2000's, most likely in recognition that there was an inadequate level of operator demand to sustain the development effort.

2.4. Pretty Secure BGP

Pretty Secure BGP (psBGP) [22] put forward the proposition that the proposals relating to the authentication of the use of an address in a routing context must either rely on the use of signed attestations that need to be validated in the context of a PKI or rely on the authenticity of information contained in Internet Routing Registries.

The weakness of routing registries is that the commonly used access controls to the registry are insufficient to validate the accuracy or the current authenticity of the information that is represented as being contained in a route registry object. The information may have been accurate at the time the information was entered into the registry, but this may no longer be the case at the time the information is accessed by a relying party.

The psBGP approach was also motivated by the proponent's opinion that a PKI could not be constructed in a deterministic manner because of the indeterminate nature of some forms of address allocations. This led to the assertion that any approach that relied on trusted sources of comprehensive information about prefix assignments and the identity of current right-of-use holders of address space was not a feasible proposition.

Accordingly, psBGP rejected the notion of a hierarchical PKI that could be used to validate assertions about addresses and their use.

Interestingly, although psBGP rejected the notion of a hierarchical address PKI, psBGP assumed the existence of a centralised trust model for AS numbers and the existence of a hierarchical PKI that allowed public keys to be associated with AS numbers in a manner that could be validated in the context of this PKI. This exposed a basic inconsistency in the assumptions that lie behind psBGP, namely that a hierarchical PKI for ASes aligned to the AS distribution framework was assumed to be feasible, but a comparable PKI for addresses was not. Given that the same distribution framework has been used for both resources in the context of the Internet, it is unclear why this distinction between ASes and addresses was necessary or even appropriate.

psBGP used a rating mechanism similar to that used by PGP [23], but in this case the rating was used for prefix origination. An AS asserted the prefixes it originated and also could list the prefixes originated by its AS peers in signed attestation. The ability of an AS to sign an attestation about prefixes originated by a neighbour AS allowed a psBGP speaker to infer AS neighbour relationship from such assertions, allowing the local BGP speaker to construct a local model of interAS topology in a fashion analogous to soBGP. One of the critical differences between psBGP and soBGP was the explicit inclusion of the *strict* AS Path validation test, namely that it was a goal of psBGP to allow a BGP speaker to verify that the BGP UPDATE message traversed the same sequence of ASes as is asserted in the AS Path of the UPDATE message. The AS path validation function relies on a sequence of nested digital signatures of each of the ASes in the AS Path for trusted validation, using a similar approach to sBGP. However, psBGP allowed for partial path signatures to exist, mapping the validation outcome to a confidence level rather than a more basic sBGP model of accepting an AS path only if the AS Path in the BGP UPDATE message was completely verifiable.

The essential approach of psBGP was the use of a reputation scheme in place of a hierarchical address PKI, but the value of this contribution was based on accepting the underlying premise that a hierarchical PKI for addresses was infeasible. It is also noted that the basis of accepting inter-AS ratings in order to construct a local trust value was based on accepting the validity of an AS trust rating, which, in turn, was predicated upon the integrity of the AS hierarchical PKI. psBGP appeared to be needlessly complex and bears much of the characteristics of making a particular solution fit the problem, rather than attempting to craft a solution within the bounds of the problem space.

The use of inter-AS cross certification with prefix assertion lists introduces considerable complexity in both the treatment of confidence in the assertions and in the resulting assessment of the reliability of the verification of the outcome. psBGP does not consider the alternate case where the trust model relating to addresses is based on a hierarchical PKI that mirrors the address distribution framework. In such a case the calculation of confidence levels would be largely unnecessary. The major contribution of psBGP relates to the case of partial deployment of a security solution in relation to AS Path validation, with the calculation of a confidence rating in the face of partial security information.

2.5. Inter-domain Route Validation

The approaches to securing the semantics of BGP described in this section so far all entail changes to the operation of BGP itself and operate most effectively in an environment of universal deployment. In practical terms this is an unlikely scenario, and the experience with the uptake of modifications to BGP that supported 32-bit AS number values suggests that the public Internet has considerable inertia and is very resistant to adopting changes to BGP [24]. In such a system as large as the public Internet, long term piecemeal deployment is a far more likely scenario.

The approach proposed with Inter-domain Route Validation (IRV) [25] is not to modify the BGP protocol in any way, but to define a companion information distribution protocol. The intent here was to attempt to provide legacy compatibility and incremental deployment capability. The IRV approach replaced the concept of simultaneously feeding both routing information and associated credentials in BGP with the concept of moving the provision of credentials into a query response framework where the receiver of a route object can query the originating AS as to the authenticity of a received route object, or request additional information relating to the object in a similar fashion to the information contained in an Internet Routing Registry (IRR)

[26]. In IRV each AS is responsible for providing an IRV server capable of providing authoritative responses relating to prefixes originated by this AS. IRV is envisaged as being used to provide routing policy information, using the Routing Policy Specification Language (RPSL) [27], [28] structure already used by the Internet Route Registries (IRRs), community configuration information, contact information, a local view of the routing system in terms of received route advertisements and withdrawals and route updates that have been sent to neighbouring ASes.

Assuming that there is a way to reliably query a per-AS IRV server, and receive a response that can be validated, then AS origination validation in the IRV framework is a case of querying the originating AS's IRV server with the origination query for the prefix in question and verifying the response. In a similar fashion AS Path validation is a case of querying each AS's IRV server in the AS path, confirming that an advertisement was received from the previous AS in the AS Path, and that an advertisement has been sent to the next AS in the AS path (Figure 3). This approach is midway between a strict AS Path test that validates that the UPDATE message was passed along the AS sequence described in the AS Path, and AS Path plausibility that validates that there is a set of AS peer connections that correspond to the AS sequence. Here the validation test is that each AS in the sequence is currently advertising this prefix to the next AS in sequence.

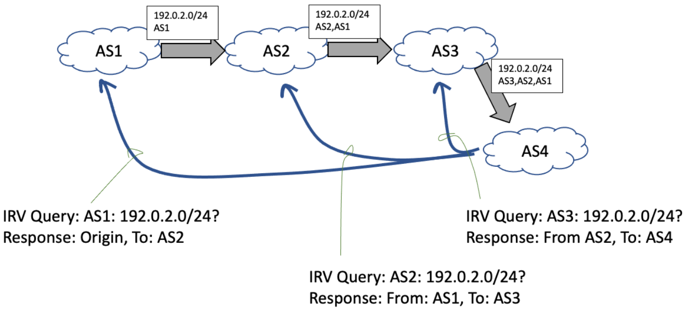
192.0.2.0/24
AS1
192.0.2.0/24
AS2,AS1
AS1
AS2
AS3
192.0.2.0/24
AS3,AS2,AS1
AS4
IRV Query: AS1: 192.0.2.0/24?
Response: Origin, To: AS2
IRV Query: AS3: 192.0.2.0/24?
Response: From AS2, To: AS4
IRV Query: AS2: 192.0.2.0/24?
Response: From: AS1, To: AS3

Figure 3 – AS Path Verification using IRV

This IRV architecture has a number of issues that are not completely specified, including IRV discovery, IRV query redirection, authentication of queries and responses, selective responses, transport layer protection and imposed overheads. It is unclear how an IRV response is to be validated, and how the relying party can verify that the received response originated from the IRV server of the AS in question, that the response has not been altered in any way, and that the response represents the actual held state in the queried AS. A similar concern lies in the estimation of additional overhead associated with performing a query to each AS in the AS Path for every received BGP UPDATE. It is also unspecified whether the query and response is a precondition to the local acceptance of a BGP route or not. While making validation of a route a precondition for acceptance of a route would appear to offer a more robust form of security, it is also the case that the IRV associated with the originating AS may only be reachable via the prefix being advertised, in which case the IRV would be unreachable until the route is accepted. It is also unclear to what extent the additional information that the IRV could provide would be useful within strict real time constraints.

The IRV approach is essentially an extension of the IRR concept that further decentralises the publication point of routing information to individual ASes. It extends the IRR in a manner that is intended to provide adequate assurance that received responses are responses to the original query, that the response has been formed by the authoritative IRV for an AS, that the response is complete and has not been altered in any way, and that the response is an accurate representation of the state of the remote AS, using DNS-style chained look-ups. What is unclear here is whether this decentralisation has superior performance and security properties to an alternative approach of further augmentation to the existing IRR framework.

A similar approach within the IRR framework that integrates the concept of an address and AS PKI could make provision for signed responses in a way that allows the IRR client to authenticate that the response is accurate, current, and contains information that has been digitally signed by the AS or prefix holder. In such a model of publication the relying party is able to validate the authenticity of the IRR object independently of the manner in which the object was published or the manner in which it has been retrieved [29].

2.6 Secure Path Vector routing for securing BGP

Secure Path Vector routing for securing BGP (SPV) is another proposal that explores the feasibility of using symmetric cryptographic operations to secure the AS path in BGP UPDATE messages [30] using hash chains and trees. The SPV study identified the following classes of path attacks:

- *forgery* where false paths are associated with routes in order to influence local route selection decisions,
- *modification* where the path is altered in order to hide the UPDATE from a target AS or in order to influence local route selection decisions,
- *denial of service* where the attack attempts to overwhelm the intended victim's resources, and
- *worm-holing* where colluding adversaries assert false AS-to-AS links.

The first two classes are attacks via BGP, whereas the second two could be more accurately classified as attacks on the routing system itself through multi-party collusion. SPV takes the approach of tree-authenticated hash values and applies this specifically to AS Path validation as an alternative to the nested digital signature structure proposed as the AS Path validation mechanism of sBGP. The paper claims significantly improved processor performance using this technique, based on difference in computational complexity for asymmetric cryptography from symmetric cryptography as used in hash functions.

This proposal falls into the category of proposals that calls for changes to the operation of the BGP protocol. In this case the significant change is the requirement that all routes must be re-advertised to peers within a fixed time interval. This is the weakest part of the approach in terms of performance evaluation, as much of the leverage in terms of scaling BGP, is based on the use of a reliable transport protocol for BGP messages which, in turn, obviated any need for a BGP readvertisement function. The need to regularly re-advertise the entire routing table to all peers has some adverse implications in terms of the performance of the protocol and its scaling capabilities.

SPV also assumes that the originating AS has knowledge of the private key associated with an address, as distinct from the more logical approach that an originating AS need only be able to produce an authority from the address allowing the AS to originate the advertisement. This approach, while efficient on processing speed, requires more storage, a higher level of time synchronisation, higher update rates within the BGP protocol, coupled with some form of loose time synchronisation and complex key pair distribution. It has also been observed [31] that SPV does not sufficiently protect against route forgery and eavesdropping or collusion attacks.

2.7 Signature Amortisation and Aggregate Signatures

If the signature load of sBGP is the problem, then how can this load be reduced? This question has been studied in a number of papers.

It may be possible to amortise the cost of signature validation over many messages [32]. The technique signs a subset of the connected topology over which an UPDATE flows and placing a topology description as a vector in an equivalent of an AS connectivity attestation which is flooded to all relying parties. The AS Path signing can then be generalized such that the same vector is reproduced in the signed data, with the AS neighbours who were passed the UPDATE messages marked in the bit vector. All AS neighbours can now receive the same UPDATE.

Related work [33] combines the time-efficient approach of signature amortisation with space-efficient techniques of aggregate signatures to propose a set of constructions for aggregated path authentication that improve on sBGP's requirements for processing throughput and memory space.

Aggregate signatures apply to a collection of UPDATE messages that are to be sent to a peer. Instead of signing each UPDATE separately, the UPDATE messages are hashed into a Merkle hash tree [34] and the root of the tree is signed, and the UPDATE and the root of the hash tree is sent as the signed UPDATE to each peer. This technique improves upon [35] which uses bi-linear maps instead of Merkle hash trees.

2.8 Exploiting Path Stability

Mitigating the validation overhead can also be achieved by caching validation outcomes and reapplying the outcome if the same update information is received within the cache lifetime. A study by Butler, McDaniel, and Aiello [36] noted that across a one-month period less that 2% of advertised prefixes were advertised using more than 10 paths and less than 0.06% of prefixes were advertised with more than 20 paths.

Their paper proposed combining a number of approaches to reduce the AS Path validation workload. The first was the use of hash chains and signature aggregation, where a BGP speaker sends all local viable paths to its peers along with the tokens that represent hash chain anchors, allowing route change to be represented by an authentication token that can be validated by hash operations. The second part of the approach was to use Merkle hash trees to sign across a set of UPDATE messages that are queued awaiting the MRAI Timer. The third part of the approach was to exploit the stability of path advertisements to amortise cryptographic operations over many validations, achieved by caching the cryptographic proofs.

The paper asserted that simulations point to a reduction of the computational costs by as much as 97% over existing approaches using this approach.

Another approach, termed *pretty good BGP* (pgBGP) [37], analyses path stability over a longer period of time and builds a local database which is then consulted in order to detect anomalous routes. The idea is that origin ASes usually do not suddenly change over time for certain prefixes, and that such a sudden change might indicate an attack to the routing system. pgBGP does not provide completely automated security, as it does not eliminate any route advertisements, but rather puts them into quarantine for 24 hours (similar to route flap damping), giving operators the time to decide how to classify the event. This proposal can be incrementally deployed and imposes little overhead on the routing system. It is a method to mitigate effects of an attack to the routing system, and not an effective mechanism for prevention of such attacks.

2.9 Detecting Prefix Hijacking

One special case of routing attacks which is considered a major threat and evokes high interest in the research community is prefix hijacking. There has been a considerable amount of research undertaken in order to provide security against this single form of attack. The approaches describe possible methods of detecting prefix hijacking [38], [39], [40] as well as complete systems and implementations of prefix hijacking detection in order to possibly react on the attack. These systems [42], [43], [44], [45] rely on existing external route monitoring databases like Route Views [46] or need special routing registries to be deployed to detect prefix hijacking. The quality of such prefix hijack detection systems is strongly dependent on the quality of the route databases, all of which have some level of perspective bias given that all views of the BGP routing system are relative to the location of the collector.

Another method to detect prefix hijacking is to look for multiple origin AS (MOAS) [47], [48], which can be either a sign of multi-homing an AS or a sign of bogus route announcements, thus prefix hijacking.

A different approach is presented for iSPY [49], which tries to detect prefix hijacking by continuously probing known transit ASes in order to detect whether the prefix owned by the probing AS has been hijacked through a path change in the routing fabric to reach the address prefix.

2.10 Secure BGP and BGP Dynamics

If securing BGP is a case of applying cryptographic operations to BGP UPDATE messages, then the other approach to reducing the security overhead is to exploit the dynamic behaviour of these messages.

The BGP update pattern is studied in [50] where in a study of BGP update dynamics it was shown that a cache of 10,000 prefix and AS Path validation outcomes, or less than 5% of the total number of distinct routed entries, would achieve a cache rate of between 30% to 50% using a simple LRU cache replacement algorithm.

When distance vector algorithms react to a change in prefix reachability a number of UPDATE messages are generally observed before the routing system reaches a stable state. A study of BGP convergence across the global Internet concluded that the severity of path exploration and the convergence speed depends on the relative positions of the event origin and the observer [51]. This study aligned the originator and the observer in terms of the "tiering" of Internet Service Providers and noted that this extended convergence times and larger path exploration events occurred at lower levels of the tiering hierarchy. It was hypothesised that the richer inter-connectivity that was typically prevalent at such lower levels in the tiering hierarchy was a major contributing factor here. Fail-over and new route announcements converge in similar times, while route withdrawals have far longer convergence times.

A similar study on BGP's path exploration characteristics proposed modifications to the BGP UPDATE message intended to identify and limit the path exploration behaviour of BGP [52]. If a significant level of update load is related to path exploration and a significant level of AS Path security overhead is related to validation of short-term transient routing states associated with path exploration, then another direction in terms of reducing security overheads is to limit path exploration behaviour. An approach to do so by selective damping of BGP updates that are characteristic of BGP path exploration following a withdrawal at source is described in Path Exploration Damping [53], [54].

Further study of BGP update behaviour has explored the level of determinism that exists in BGP's route selection process and noted that in the absence of the Multiple Exit Discriminator (MED) and route reflectors, then the process can be considered to be a deterministic one [55]. The paper suggests some refinements to BGP that could achieve a similar outcome to MEDs and route reflectors while preserving the deterministic route selection property. The question this paper raises is that most security proposals view AS Path validation as an "after the event" activity because of the assumed lack of predictability in BGP. This paper questions this basic assumption and raises the possibility of path security as a provisioning activity, which, in turn raises some interesting performance optimisations for BGP path security as a provisioning exercise rather than a reactive task.

2.11 Securing the Data Plane

Securing BGP is not only a matter of securing the control plane, but also of securing the data plane [56] and to make sure that the status of the forwarding table is consistent with the advertised BGP routing information.

A study by Mao et al. [57] showed that up to 8% of the paths advertised through the control plane, do not match the actual paths in the data plane. The data plane is not only subject to attacks which try to subvert the routing system, but also subject to synthetic BGP announcements from network operators that could enable the theft of carriage capacity. It is therefore necessary to provide security for the whole data path, and not only on a Next Hop basis as Stealth Probing [58] intends to, as carriers might span over multiple ASes and synthesise false routing information that spans multiple AS hops.

Proposed approaches mainly focus on probing the full data-path through packet injection, trying to detect and isolate malicious routers. In "secure traceroute" [59] a modified traceroute is used to control which path data packets actually take and compares it to the actual AS path of the routing table, effectively detecting malicious ASes. Secure traceroute comes with the overhead of a PKI and related key exchange and no chance for piecemeal deployment.

The Faith approach [60] instead focuses on using traffic summary functions, and comparing their results with those of other routers, allowing to detect ASes which provide anomalous values. These traffic summary functions seem to be prone to inaccuracy due to a variety of applications running on routers which might alter the packet flow and their application appears infeasible in routers with very high packet volumes.

The solution proposed as Listen and Whisper [61] tries to detect inaccuracies in the data plane (the Listen part) but focuses also on control plane security (the Whisper part) and aims to provide an almost complete BGP security solution, combining both parts. Compared to sBGP, Listen and Whisper should be classified as a "just too late" solution for BGP security, like many solutions which try to ensure data plane - control plane consistency. Like other data plane security solutions, this approach seems infeasible, as it tries to detect data plane anomalies by analysing individual TCP flows, and scaling this approach to the high speed core of the Internet presents some practical challenges.

An approach that aims towards high performance and possible partial deployment is described in [62]. Its focus is to ensure that the data path always conforms to the announced AS path, which is achieved by probing data paths through injecting tagged IP packets, or by using IP options. Similar to pgBGP, it leaves the decision of which action to take towards a malicious router to the network operator and builds up a small database to detect possible malicious routers. It deploys the roles of verifiers and provers on certain ASes, with the verifier being an AS that wants to verify a certain route, and the prover being an AS that helps the verifier in the process by replying on probe data.

Even though all these approaches intend to provide a certain level of data plane security, and also a certain level of control plane security, none provide comprehensive data plane security. Authenticity of a data path from start to end could easily forged by two ASes deploying tunnels between them, and thus disabling the possibility to effectively verify the data path by a third party.

3 IETF Activity – RPKI, RoV, BGPSEC and ASPA

Following a number of efforts to make progress in this area, the IETF charted a Routing Protocol Security Requirements Working Group (RPSEC) in 2002 to develop a common set of security requirements for routing protocols. The activity concluded in 2009. In terms of the study of inter-domain security requirements the work stalled on some fundamental and evidently irreconcilable disagreements over the issue of the requirements for AS Path security [63] [87], and the BGP-related working drafts from the RPSEC Working Group were never published as RFCs.

Based on the initial RPSEC work on security of route origination, the IESG chartered the Secure Inter-Domain Routing Working Group (SIDR) in 2006 [64]. The charter for this effort presented some issues, in that it was stalled in assuming security requirements for AS Path validation and had to await results from the RPSEC activity. Given that RPSEC was unable to agree on a requirement for AS Path security then the initial work in SIDR was concentrated on securing the origination of routing information, rather than its propagation through the inter-domain space. Notably in retrospect, SIDR was also constrained from making any changes to the BGP protocol, implying that any security framework applied to the operation of BGP was to be positioned as an overlay rather than a basic change to the BGP protocol itself. This turned out to be a very important decision as it precluded some design decisions that would turn out to be critical for the SIDR design work.

The initial SIDR products were a collection of specifications that described a profile for a PKI for IP addresses and AS numbers (the RPKI), as well as a model for publication and maintenance of local cache, discussed earlier in Part 1 of this survey. From this foundation the SIDR Working Group moved on to Route Origination Validation.

3.1 Route Origination Validation

Route Origination Validation (ROV) builds upon the earlier work in the Routing Registry effort, where a prefix holder is able to publish information as to how an address prefix is to be announced into the routing system by nominating the AS number(s) that are permitted to originate a routing announcement for the prefix. In the ROV framework. In the RPKI framework this information is published as a signed Route Origin Authorization (ROA) [65] [66].

A ROA is signed by a prefix holder, and denotes a permission given by the address prefix holder for an AS to originate a route.

There are a number of additional implications associated with publishing a ROA. The first is that no other AS has permission to announce that prefix when there is a cryptographically valid ROA extant in the RPKI system. If the prefix holder wishes to authorizer multiple ASes to originate a route for this prefix, then the prefix holder must generate multiple ROAs. This means that an address holder can declare that a prefix should not be routed at all by issuing a ROA that provides a permission to AS0. Secondly, the ROA denies permission for any AS to originate a prefix that his more specific than the prefix listed in the ROA. There is a MaxLength attribute of a ROA that may be used to define a range of more specific prefix lengths that are permitted by a ROA. Thirdly, there is no acknowledgement of the ROA on the part of the AS. A prefix holder may publish a ROA providing a permission to an AS who is unaware of the permission.

There is no symmetric instrument in the RPKI framework relating to the AS holder. An AS holder does not have the ability to issue a signed attestation that lists all the prefixes that it intends to originate in the routing system.

There is one more important component of the ROV framework, namely the RPKI to Router protocol (RTR) [67]. This protocol allows a crypto engine to be removed from a router and operate on a dedicated platform. The result of this local processing of ROA data is expressed in the form of a filter list, and this filter list is implemented as a shared state between a RTR server and one or more RTR client routers. This mechanism offloads most of the RPKI overheads from the router, and leaves just a residual filtering function on the router.

3.2 BGPsec

THE SIDR working group commenced work on an extension to BGP what would allow validation of the AS Path attribute in 2011, and the standard track specification of BGPsec was published in 2017 [68].

Unlike ROV, BGPsec is not implemented in an off-router mode but is implemented through the definition of non-transitive BGP AS Path attributes. These attributes carry the digital signatures produced by the AS that propagates a BGP UPDATE message. These signatures, signed by the AS, provide confidence that every AS listed in the AS Path attribute has handled the propagation of this prefix, that the order in the AS Path is the exact order of propagation of the UPDATE message through the inter-domain routing space, and each AS listed has explicitly authorised the propagation of an UPDATE message to its eBGP peer.

BGPsec appears to be solidly based on the concepts described in the earlier sBGP work [8]. In essence each eBGP speaker generates a digital signature that covers the information it received (including that digital signature) and the AS number to whom this UPDATE is to be sent (Figure 4).There is a wealth of detail behind this simple summary, but it can be summarised by the observation that this mechanism ties the AS Path in the UPDATE message to the sequence of ASes that handled the propagation of the route object. A detailed exposition of BGPsec's design decisions can be found in [69].

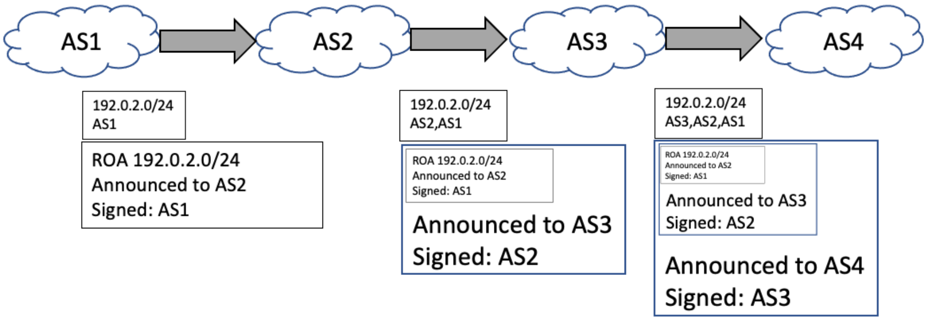

Figure 4 – BGPsec handling of AS Path Signature structure

Stepwise AS Path validation cannot tolerate AS Sets in this approach, nor AS Confederation Sets, and that are in the process of being deprecated in response to this limitation [88]. In a similar vein BGP Route Reflectors require special processing, as do private AS numbers.

There are a number of consequences of this design approach.

The first, and perhaps the most important consequence, is that piecemeal incremental deployment is simply not possible in BGPsec. When an UPDATE is passed from a BGPsec BGP speaker to a non-BGPsec BGP speaker all BGPsec attributes are lost. This means that it the UPDATE is further propagated to a BGPsec BGP speaker the initial BGPsec information is unavailable. In today's Internet the consequences of this highly constrained deployment scenario are prohibitive factors for adoption.

This approach also places a high crypto processing load on BGPsec-aware BGP speakers. There is some scepticism that this is a feasible impost on the Internet's routing infrastructure, and this scepticism guided the design of the ROV RTR approach. However, for BGPsec not only are routers expected to process the BGPsec messages, but also hold secure private keys to perform signing on the fly for outgoing UPDATE messages.

Thirdly, while this approach can provide some assurance regarding the "correct" operation of the BGP protocol and can detect efforts to tamper with update messages but there is no protection against spurious WITHDRAW messages, no ability to ascertain the alignment of the route object with the network's forwarding state and no protection of alignment of the UPDATE with the policy state. In other words, route leaks can still occur in BGPsec.

In summary, BGPsec represents a relatively high overhead to pay for a limited set of assurances and a limited protective capability. Furthermore, there is a more extreme view that BGPsec cannot achieve any of the security properties due to the fundamental design principles of BGP and BGPsec. In one research paper [70] it is asserted that in BGPsec, routes can still be hijacked, and routing loops can still appear. The paper's authors hope to stimulate further dialog to rethink the fundamental tenets of BGP and BGPsec designs by publishing their analysis of the observed shortcomings of BGPsec.

3.3 Autonomous System Provider Authorization (ASPA)

The issue with the overall SIDR approach to BGP security is that if BGPsec is impractical then we cannot rely on ROV alone. All a determined routing attacker need do is tack on the originating AS to a synthesised AS Path and any AS sequence can be placed in the AS Path attribute of a synthetic route.

ROV represents a substantial effort to get the infrastructure deployed, but without any form of AS Path protection the level of protection offered by ROV is minimal at best. The conclusion is that ROV needs to be accompanied by some form of AS Path validation if it is to be useful.

There have been a number of proposals to address this shortfall. An interesting approach is Peer Lock [71], which is based on the observation that the *core* of the routed Internet is a small set of *Tier 1* ASes, and no customer of an AS should be announcing a route where the AS Path includes any of these Tier 1 networks. Secondly, no more than 2 of these Tier 1 ASes should appear in any AS Path, if there are 2 such ASes in the AS Path they should be adjacent. This approach does not necessarily catch much in the way of deliberate efforts to generate a synthetic AS Path, but it can be effective in catching a number of common forms of route leaks, and its implementation is quite simple and very light weight.

Can we do better?

In what appears to be a replay of the situation from around 2000 when soBGP was proposed as a lighter weight response to the crypto load associated with sBGP in the area of AS Path validation, there has been a proposal to use RPKI-signed AS adjacency attestations as a response to the issues with BGPsec.

There is a slight twist on this, however which different from soBGP, in that there is an element of routing policy that is also used in the ASPA proposal [72]. Instead of an AS listing its adjacent ASes in the inter-domain routing space and requiring both ASes to list each other as BGP neighbours before accepting the AS adjacency as valid, the ASPA framework requires an AS to list only its adjacent ASes that act in a transit provider role to the issuing AS. Given that a common criticism of BGPsec, sBGP and soBGP was that these proposals were incapable of identifying route leaks (as route leaks represent a violation of route policy as distinct from a violation of the BGP protocol itself) ASPA provides a means of identifying such route leaks.

The ASPA relationship is a graph fragment in the directed graph which describes the inter-AS topology [73]. The property used by the ASPA proposal is described as "valley-free" AS Paths. All AS Paths can be characterised by zero or more paired relationships from Customer-to-Provider (up), zero or one Peer-to-Peer relationships (flat) and zero or more Provider-to-Customer relationships. (down). In other words, all viable AS Paths are a sequence of customer-to-provider (up) AS pairs, then a peer AS pair, then a set of provider-to-customer (down) AS pairs. Any AS sequence that contains a down then an up (or a "valley) represents a customer AS leaking routes learned from one provider to another (Figure 5).

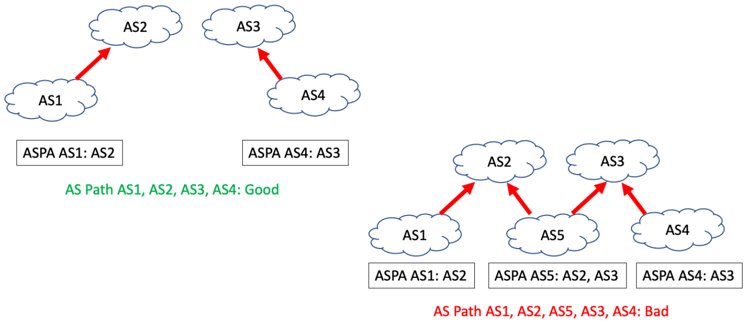

Figure 5 – ASPA and Route leaks

ASPA requires any AS that issues an ASPA object has to comply with the constraint that the providers listed in an AS's ASPA are the complete set of providers for that AS.

ASPA is still provide some benefit even in scenarios of partial deployment. Once an AS issues an ASPA than a routing attacker can only include this AS in a synthetic AS Path attribute if it also includes an adjacent provider AS, and the synthetic AS pair can only be inserted in the "front part" of the AS Path (Customer-to-Provider) if the order is preserved, and in the "back part" of the AS Path (Provider-to-Customer) in reverse order. Like soBGP, the use of ASPAs does not necessarily prevent the synthesis of AS Paths by a routing attacker, but it limits what can be used to make such synthetic paths, and the greater the use of ASPAs the more it becomes the case that the only AS Paths that can be synthesised are viable BGP AS Paths in any case. soBGP termed this constraint *AS Path Plausibility*, and the same condition applies to ASPA.

It's evidently still early days for ASPA and after three years the work remains a study item in the SIDROPS Working Group of the IETF. Part of the issue here is that the SIDROPS Work Group has turned its attention to the operational aspects of the operation of the RPKI and has taken on the role of the RPKI operational maintenance working group and has had its collective attention diverted from the issues of BGP security mechanisms and AS Path validation. And in the area of RPKI operations the topic that is taking up the Working Group's attention is not the PKI itself, but the ongoing ramifications of the original design decision to use an out-of-band client-pull credential distribution mechanism for RPKI distribution. The emerging observation is that this original design choice is sufficiently flawed that the efforts in the working group to adjust the parameters of this distribution system will in all likelihood be unable to adequately address the operational issues that accompany scaling up the use of the RPKI credential system.

It may be productive at this point in time to re-open the question of how to use BGP itself to perform a *just-in-time* push based distribution of BGP security credentials, but within the structure of the IETF it is difficult for an operationally focussed working group to perform protocol development work. However, it's an equally difficult ask for the IETF to reopen a protocol design effort on BGP security so soon after the closure of the original SIDR effort. The protracted and painful saga of the DNSSEC development effort in the IETF is one that many participants in the IETF are unwilling to repeat for BGP security.

4. Open Questions on Securing BGP

It appears to some observers that no current solution to routing security has found an adequate balance between appropriate security and acceptable deployment overhead [74] [75], and that's an observation that I can agree with. We are just not there yet.

Current research on BGP performance is focused on topics related to scalability, convergence times, stability and consistency, while the questions on security research have been focused on the integrity, authenticity, authority and verifiability of routing information. These two fields of research are inherently connected, in that a more stable routing system that was able to provide clear indications when convergence to a stable routing state had been achieved is believed to also provide clear indications of when verification of routing information is appropriate.

In exploring the threat model for BGP it is noted that BGP was designed to support inter-domain routing between trusted networks, while today's networks operate in a looser confederation that does not exhibit the same mutual trust properties. Not only are the TCP sessions used by BGP vulnerable to attack, and the messages used by BGP vulnerable to alteration in order to disrupt the network's routing system, but the integrity of the operation of BGP is also threatened by misconfiguration, where incorrect information is injected into the routing system unintentionally, and by router vulnerabilities where a compromised routing system can exploit its trusted role and intentionally inject false information into the routing system.

Some of these attacks are intended to cause a BGP speaker to be overwhelmed and reset, as BGP is a method of directly accessing a router's processing unit and a saturation attack can cause processor and memory overload. Other attacks are aimed at altering the router's forwarding state, generating an incorrect or unintended forwarding state for one or more prefixes. Other forms of attack are aimed at causing a BGP speaker to become unstable and thereby disrupt the forwarding function and impact on applications. A BGP session that is being continually reset will cause large local traffic bursts as neighbouring BGP speakers continually resend their routing tables upon each reset, but the continued instability will trigger a flap damping response in other BGP speakers.

The factors that contribute to these vulnerabilities include a lack of BGP message integrity checks, an as yet partial ability to check the authority of an originating AS to actually originate an advertisement for a prefix, and an inability to verify the accuracy, completeness and authenticity of as path attributes of a routing advertisement. The use of the RPKI to support address attestations, as in ROAs, provides a very robust means of detecting incorrect origin route objects, as long as the RPKI itself is accurately aligned to the address distribution framework and as long as the RPKI is generally, if not universally, used.

In contrast, robust solutions to the problem of AS Path authentication have been elusive so far. BGPsec provides a robust method of path validation but has been assessed to be significantly expensive in terms of processor and memory cost, and also detrimental to BGP convergence times and requires comprehensive adoption to be effective. Efforts to substitute AS Path plausibility in place of actual AS Path validity, as is the case with ASPA, offer a different level of robustness that appears to be more practically achievable.

The study of approaches to securing BGP has raised several questions about the behaviour of inter-domain routing and the most effective approach to securing BGP. These questions include consideration of security topics and raise the issue of whether it is possible to secure the routing information to the extent that the routing information being presented is tightly aligned to the associated forwarding state [76].

- Is it possible to secure this association of routing information to the chained forwarding state? Can a BGP speaker validate that the AS path as presented in a BGP route advertisement not only matches the BGP propagation path taken by the prefix advertisement, but that the network's current forwarding state to reach the address prefix is aligned to this AS path and this alignment can be validated? To put is simply, can a router validate that a route matches the forwarding path? This question is not one that is directly addressed within any of the current set of inter-domain routing security measures.

- A related issue concerns the overheads of securing BGP and the scaling properties of BGP. Is BGP too monolithic a protocol even before adding security capabilities? BGP simultaneously performs the functions of exchanging reachable prefixes, maintaining an inter-domain network topology, binding prefixes to paths, and implementing routing policy. Would inter-domain routing be more scalable if these functions were to be performed by separate protocols? Adding security and authentication within BGP, as in the sBGP model, increases the complexity of the protocol and may diminish its long-term prospects for scalability across ever larger and denser inter-domain topologies. At the same time, using a separate mechanism to flood security credentials in a manner that is entirely distinct from BGP itself, as used in the Route Origination Validation framework, becomes a source of additional operational complexity and potential vulnerability, even though the BGP protocol itself is unaltered.

There are several practical and some more fundamental questions relating to securing BGP.

- The first is a practical question relating to the inevitable design trade-off between the level of security and the performance overheads of processing security credentials. The question concerns what aspects of securing BGP should be considered essential and what is simply desirable, but not essential. Our level of understanding as to what aspects of BGP performance and load are critical for the robust operation of network applications and what are not so critical appears to be less than comprehensive. The impact of performance trade-offs in BGP in terms of time to converge, the size of the routing space, the router memory and processing load and scaling capability are not well understood to the extent that there is a commonly accepted answer here.

- The next question is whether verification of the correct operation of the BGP protocol is sufficient, or whether the policy intent of the routing environ is equally critical. For example, is a stub network were to leak the routes it learned from one transit network to another transit network this route leak would, in the normal situation, be regarded as contrary to routing policies, but there is no violation of the BGP protocol itself. If we want to also include alignment to routing policies then the question arises as to how such policies are to be expressed, who has the authority to express them, and how BGP speakers reconcile local routing policies with external routing policies when the policies differ.

- The next question is whether securing the operation of the BGP protocol (securing the control plane) is sufficient in and of itself to adequately mitigate the vulnerabilities in the overall routing system, or whether it is also necessary to include mechanisms that extend the security model to validate that the routing information represents current forwarding state in each routing element in the network (securing the data plane). One perspective on this is that securing one element of system with multiple components does not necessarily address the underlying vulnerabilities of the entire system. The more common outcome is that such work exposes the residual vulnerabilities in other components, and that an effective security system needs to address all components of the routing system. While it may be possible for a BGP speaker to be able to validate that the originating AS did indeed originate the prefix advertisement and that the AS path accurately represents the propagation path of this advertisement through the network, that is not the basic question in terms of the properties of the overall system.

- The more basic question here is whether a BGP speaker can verify that if it decides to forward a packet on the next hop along a path indicated by the routing system as the optimal path to a destination is this indeed the optimal local choice and does this next hop decision pass the packet "closer" to the destination address?

- If a comprehensive security framework is proving to be elusive in terms of deployment considerations, then could a less comprehensive approach offer acceptable outcomes? Many security frameworks demonstrate a profile of diminishing returns, where the incremental cost of deployment of additional security capabilities increases, while the incremental benefit in terms of risk mitigation decreases. In the case of securing BGP could an approach of reducing the security credential generation and validation workload, through reducing the amount or timeliness of validated information, represent an acceptable trade-off? We see a practical form of this question today, where the capabilities offered by Route Origination Validation can mitigate some forms of routing incidents but are ineffectual against other forms of route manipulation that preserve the origination data. Practically, is this enough? Or do we need to also deploy some mechanism that allows detection of various forms as AS Path manipulation? A similar question relates to the comparison of the earlier soBGP and sBGP models. Is *Path Plausibility* sufficient? Did the mechanisms of soBGP exercise sufficient levels of constraint such that any synthesised path is close enough to a viable network path that the difference is of little consequence from a security perspective? This question is being replayed today when we consider the relative merits of the ASPA approach against the heavier weight of BGPsec's fully signed AS Path attribute.

- A final question here concerns the practicalities of deployment. The Internet is now far too large to sustain the concept of a flag day for deployment of any technology, and it is not possible to assume that a technology would be universally adopted without a protracted period of piecemeal deployment as part of a transitional interval. Indeed, as the Internet continues to grow and the diversity within the Internet increases, the anticipated transitional periods become indefinite, and piecemeal deployment becomes a continuing factor rather than a temporary transitional factor. The questions this exposes include whether it is even possible to deploy high integrity security using partial deployment scenarios, or whether the BGP protocol is too incomplete in terms of its information distribution properties to allow robust validation of the intended forwarding state? Does securing forwarding imply carrying additional information relating to the routing and forwarding state coupling in addition to routing that would be entirely impractical in a partial deployment scenario?

5. Conclusions

BGP has proven surprisingly resilient in terms of its longevity of useful operational life, despite early predictions of its imminent demise in favour of IDRP [12]. BGP-4 has routed the inter-domain Internet since late 1993 and the number of routed elements for the IPv4 Internet 'default-free zone' has grown from under 20,000 distinct prefixes to some 1,000,000 distinct prefixes by the middle of 2021, and a further 130,000 prefixes in the IPv6 network [10]. Despite the changes in the IPv4 address infrastructure due to exhaustion of the registry free pools the growth in the number of routing IPv4 prefixes appears to continue unabated, and together with the continued deployment of IPv6, these numbers are expected to continue to rise in the coming years.

Due to its extensibility and large installed base, BGP-4 will likely remain the only inter-domain routing protocol in the foreseeable future for the Internet (although the term "foreseeable is prudently measured in units of years and perhaps not in decades). So far BGP has not changed in any substantive manner, including in its security properties.

There is ample evidence from reports of use of unregistered addresses [77] or of "routing incidents" [78] that BGP is the subject of various forms of accidental inattention and possibly deliberate forms of abuse. Current efforts at mitigation of these forms of abuse appear in the inter-domain routing space to be less than fully adequate and the ease with which unauthorised or bogus route objects can be injected into the inter-domain routing system remains a continuing threat issue for the security, stability and utility of the Internet. We appear to be getting very comfortable in operating a network that experiences a continuing stream of routing incidents, both intentional and unintentional, and the longer this situation persists the more we are resigned to just accept this as the status quo for the Internet and place the onus on applications and content distribution systems to defend themselves from routing attack. Like many unintended outcomes it's not the outcome we would prefer to have, nor is it necessarily the optimal outcome in terms of collective cost and benefit, but it's the outcome many of us have simply accepted. All change comes at a price, and the more we resign ourselves to operating networks in the face of a poorly secured routing system the greater the effort required to justify the case that the cost of a change to improve this situation will be money and effort widely spent.

References

[1] Bellovin, S., *Security problems in the TCP/IP protocol suite*, SIGCOMM Computer Communications Review, vol. 19, no. 2, pp. 32–48, 1989.

[2] Gill, V., Heasley, J., and Meyer, D., *The Generalized TTL Security Mechanism (GTSM)*, RFC 3682, DOI 10.17487/RFC3682, February 2004, https://www.rfc-editor.org/info/rfc3682

[3] Gill, V., Heasley, J., Meyer, D., Savola, P., and Pignataro, C., *The Generalized TTL Security Mechanism (GTSM)*, RFC 5082, DOI 10.17487/RFC5082, October 2007. https://www.rfc-editor.org/info/rfc5082

[4] Heffernan, A., *Protection of BGP Sessions via the TCP MD5 Signature Option*, RFC 2385, DOI 10.17487/RFC2385, August 1998. https://www.rfc-editor.org/info/rfc2385

[5] Rivest, R., *The MD5 Message-Digest Algorithm*, RFC 1321, DOI 10.17487/RFC1321, April 1992. https://www.rfc-editor.org/info/rfc1321

[6] Leech, M., *Key Management Considerations for the TCP MD5 Signature Option*, RFC 3562, DOI 10.17487/RFC3562, July 2003. https://www.rfc-editor.org/info/rfc3562

[7] Touch, J., Mankin, A., and Bonica, R., *The TCP Authentication Option*, RFC 5925, DOI 10.17487/RFC5925, June 2010. https://www.rfc-editor.org/info/rfc5925

[8] Kent, S. and Seo, K., *Security Architecture for the Internet Protocol*, RFC 4301, DOI 10.17487/RFC4301, December 2005. https://www.rfc-editor.org/info/rfc4301

[9] Perlman, R., *Network Layer Protocols with Byzantine Robustness*, Technical Report, M.I.T., 1988. http://www.lcs.mit.edu/ publications/specpub.php?id=997

[10] Smith, B., and Garcia-Luna-Aceves, J., *Securing the border gateway routing protocol*, Global Telecommunications Conference, 1996. GLOBECOM '96, Nov 1996, pp. 81–85.

[11] Smith, B., and Garcia-Luna-Aceves, J., *Efficient security mechanisms for the border gateway routing protocol*, Computer Communications, vol. 21, no. 3, pp. 203–210, March 1998.

[12] *Protocol for Exchange of Inter-Domain Routing Information Among Intermediate Systems to Support Forwarding of ISO 8473 PDUs*, ISO/IEC 10747, October 1994. https://standards.globalspec.com/std/9960/iso-iec-10747

[13] Bates, T., Bush, R., Li, T., and Rekhter, Y., *DNS-based NLRI origin AS verification in BGP*, Internet Draft, Jul. 1998. http://tools.ietf.org/html/draft-bates-bgp4-nlri-orig-verif-00

[14] Arends, R., Austein, R., Larson, M., Massey, D., and Rose, S., *DNS Security Introduction and Requirements*, RFC 4033, DOI 10.17487/RFC4033, March 2005. https://www.rfc-editor.org/info/rfc4033

[15] Donnerhacke L., and Wijngaards, W., *DNSSEC protected routing announcements for BGP*, May 2008. Internet Draft: https://tools.ietf.org/html/draft-donnerhacke-sidr-bgp-verification-dnssec-04

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

[17] Seo, S., Lynn, C., and Kent, S., *Public-key infrastructure for the Secure Border Gateway Protocol (S-BGP)*, DARPA Information Survivability Conference & Exposition II, 2001. DISCEX '01. Proceedings, vol. 1, 2001, pp. 239–253 vol.1.

[18] Kent, S., Lynn, C., Mikkelson, J., and Seo, K., *Secure Border Gateway Protocol (S-BGP) – Real World Performance and Deployment Issues*, 7th Annual Network and Distributed System Security Symposium (NDSS'00), Feb 2000, pp. 103–116.

[19] Zhao, M., Smith, S., and Nicol, D., *Evaluating the Performance Impact of PKI on BGP Security*, Internet 2 4th Annual PKI R&D Workshop, April 2005. https://nvlpubs.nist.gov/nistpubs/Legacy/IR/nistir7224.pdf

[20] White, R., *Securing BGP Through Secure Origin BGP*, The Internet Protocol Journal, vol. 6, no. 3, Sep 2003.

[21] White, R., *Architecture and Deployment Considerations for Secure Origin BGP (soBGP)*, Internet Draft, June 2006. https://datatracker.ietf.org/doc/html/draft-white-sobgp-architecture-02

[22] Oorschot, P. v., Wan, T., and Kranakis, E., *On interdomain routing security and pretty secure BGP (psBGP)*, ACM Transactions on Information System Security, vol. 10, no. 3, p. 11, 2007.

[23] Zimmermann, P. R., *The official PGP user's guide*, Cambridge, MA, USA: MIT Press, 1995.

[24] Huston, G., *Exploring Autonomous System Numbers*, The Internet Protocol Journal, vol. 9, no. 1, Mar 2006.

[25] Goodell, G., Aiello, W., Griffin, T., Ioannidis, J., McDaniel, P., and Rubin, A., *Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing*, Proceedings of Internet Society Symposium on Network and Distributed System Security (NDSS 03), February 2003.

[26] Bates, T., Gerich, E., Joncheray, L., Jouanigot, J-M., Karrenberg, D., Terpstra, M., and Yu, J., *Representation of IP Routing Policies in a Routing Registry (ripe-81++)*, RFC 1786, DOI 10.17487/RFC1786, March 1995. https://www.rfc-editor.org/info/rfc1786

[27] Alaettinoglu, C., Villamizar, C., Gerich, E., Kessens, D., Meyer, D., Bates, T., Karrenberg, D., and M. Terpstra, M., *Routing Policy Specification Language (RPSL)*, RFC 2622, DOI 10.17487/RFC2622, June 1999. https://www.rfc-editor.org/info/rfc2622

[28] Blunk, L., Damas, J., Parent, F., and Robachevsky, A., *Routing Policy Specification Language next generation (RPSLng)*, RFC 4012, DOI 10.17487/RFC4012, March 2005. https://www.rfc-editor.org/info/rfc4012

[29] R. Kisteleki and Boumans, J., *Securing RPSL Objects with RPKI Signatures*, Oct. 2008. Internet Draft: https://tools.ietf.org/id/draft-kisteleki-sidr-rpsl-sig-00.txt

[30] Hu, Y.-C., Perrig, A., and Sirbu, M., *SPV: secure path vector routing for securing BGP*, SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications. New York, NY, USA: ACM, 2004, pp. 179–192.

[31] Raghavan, B., Panjwani, S., and Mityagin, A., *Analysis of the SPV secure routing protocol: weaknesses and lessons*, SIGCOMM Computer Communications Review, vol. 37, no. 2, pp. 29–38, 2007.

[32] Nicol, D., Smith, S., and Zhao, M., *Efficient Security for BGP Route Announcements*, TR-2003-440, Tech. Rep., 2003.

[33] Zhao, M., Smith, S., and Nicol, D., *Aggregated path authentication for efficient BGP security*, CCS '05: Proceedings of the 12th ACM conference on Computer and communications security. New York, NY, USA: ACM, 2005, pp. 128–138.

[34] Merkle, R., *Protocols for Public Key Cryptosystems*, IEEE Symposium on Security and Privacy, vol. 0, p. 122, 1980.

[35] Boneh, D., Gentry, C., Lynn, B., and Shacham, H., *Aggregate and verifiably encrypted signatures from bilinear maps*, Advances in Cryptology - EUROCRYPT 2003, vol. 2656. Springer Berlin / Heidelberg, January 2003, p. 641.

[36] Butler, K., McDaniel, P., and Aiello, W., *Optimizing BGP security by exploiting path stability*, CCS '06: Proceedings of the 13th ACM conference on Computer and Communications Security. New York, NY, USA: ACM, 2006, pp. 298–310.

[37] Karlin, J., Forrest, S., and Rexford, J., *Pretty Good BGP: Improving BGP by cautiously adopting routes*, ICNP '06: Proceedings of the Proceedings of the 2006 IEEE International Conference on Network Protocols. Washington, DC, USA: IEEE Computer Society, 2006, pp. 290–299.

[38] Qiu, J., Gao, J., Ranjan, S., and Nucci, A., *Detecting bogus BGP route information: Going beyond prefix hijacking*, Security and Privacy in Communications Networks and the Workshops, 2007. SecureComm 2007. Third International Conference on, Sept. 2007, pp. 381–390.

[39] Zheng, C., Ji, L., Pei, D., Wang, J., and Francis, P., *A light-weight distributed scheme for detecting IP prefix hijacks in real-time*, SIGCOMM Computer Communications Review, vol. 37, no. 4, pp. 277–288, 2007.

[40] Hu, X., and Mao, M., *Accurate real-time identification of IP prefix hijacking*, SP '07: Proceedings of the 2007 IEEE Symposium on Security and Privacy. Washington, DC, USA: IEEE Computer Society, 2007, pp. 3–17.

[42] Kruegel, C., Mutz, D., Robertson, W., and Valeur, F., *Topology-Based Detection of Anomalous BGP Messages*," Recent Advances in Intrusion Detection, vol. 2820. Springer Berlin / Heidelberg, February 2003, pp. 17–35.

[43] Lad, M., Massey, D., Pei, D., Wu, Y., Zhang, B., and Zhang, L., *PHAS: a prefix hijack alert system*, USENIX-SS'06: Proceedings of the 15th conference on USENIX Security Symposium. Berkeley, CA, USA: USENIX Association, 2006.

[44] Kim, E., Nahrstedt, K., Xiao, L., and Park, K., *Identity-based registry for secure interdomain routing*, ASIACCS '06: Proceedings of the 2006 ACM Symposium on Information, computer and communications security. New York, NY, USA: ACM, 2006, pp. 321–331.

[45] Zhang, Z., Zhang, Y., Hu, Y., and Mao, M., *Practical defenses against BGP prefix hijacking*, CoNEXT '07: Proceedings of the 2007 ACM CoNEXT conference. New York, NY, USA: ACM, 2007, pp. 1–12.

[46] *University of Oregon Route Views Project.* http://www.routeviews.org

[47] Zhao, X., Pei, D., Wang, L., Massey, D., Mankin, A., Wu, S., and Zhang, L., *An analysis of BGP multiple origin AS (MOAS) conflicts*, IMW '01: Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement. New York, NY, USA: ACM, 2001, pp. 31–35.

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

[49] Zhang, Z., Zhang, Y., Hu, Y., Mao, M., and Bush, R., *ISPY: detecting IP prefix hijacking on my own*, in SIGCOMM '08: Proceedings of the ACM SIGCOMM 2008 conference on Data communication. New York, NY, USA: ACM, 2008, pp. 327–338.

[50] Huston, G., *Measures of self-similarity of BGP updates and implications for securing BGP*, Proceedings of the 8th International Conference on Passive and Active Network Measurement (PAM 2007), vol. 4427. Heidelberg, DE: Springer-Verlag Berlin, April 2007, pp. 1–10.

[51] Oliveira, R.,. Zhang, B., Pei, D,. Izhak-Ratzin, R., and Zhang, L., *Quantifying path exploration in the Internet*, IMC '06: Proceedings of the 6th ACM SIGCOMM conference on Internet measurement. New York, NY, USA: ACM, 2006, pp. 269–282

[52] Chandrashekar, J., Duan, Z., Zhang, Z., and Krasky, J., *Limiting path exploration in BGP*, INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE, vol. 4, March 2005, pp. 2337–2348 vol. 4.

[53] Li, T., Huston, G., *BGP Stability Improvements*, June 2007. Internet Draft: https://datatracker.ietf.org/doc/html/draft-li-bgp-stability-01.txt

[54] Huston, G., Rossi, M., and Armitage, G., *A Technique for Reducing BGP Update Announcements through Path Exploration Damping*, in *IEEE Journal on Selected Areas in Communications*, vol. 28, no. 8, pp. 1271-1286, October 2010. DOI: 10.1109/JSAC.2010.101005

[55] Feamster, N., and Rexford, J., *Network-Wide Prediction of BGP Routes*, Networking, IEEE/ACM Transactions on, vol. 15, no. 2, pp. 253–266, April 2007.

[56] Wendlandt, D., Avramopoulos, I., Andersen, D., and Rexford, J., *Don't Secure Routing Protocols, Secure Data Delivery*, Proc. 5th ACM Workshop on Hot Topics in Networks (Hotnets-V), Irvine, CA, Nov. 2006.

[57] Mao, M., Rexford, J., Wang, J., and Katz, R., *Towards an accurate AS-level traceroute tool*, SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications. New York, NY, USA: ACM, 2003, pp. 365–378.

[58] Avramopoulos, I., and Rexford, J., *Stealth probing: efficient data-plane security for IP routing*, in ATEC '06: Proceedings of the annual conference on USENIX '06 Annual Technical Conference. Berkeley, CA, USA: USENIX Association, 2006, pp. 25–25.

[59] Padmanabhan, V., and Simon, D., *Secure traceroute to detect faulty or malicious routing*, SIGCOMM Computer Communications Review, vol. 33, no. 1, pp. 77–82, 2003.

[60] Mizrak, A., *Fatih: Detecting and isolating malicious routers*, DSN ’05: Proceedings of the 2005 International Conference on Dependable Systems and Networks. Washington, DC, USA: IEEE Computer Society, 2005, pp. 538–547.

[61] Subramanian, L., Roth, V., Stoica, I., Shenker, S., and Katz, R., *Listen and whisper: security mechanisms for BGP*, NSDI’04: Proceedings of the 1st conference on Symposium on Networked Systems Design and Implementation. Berkeley, CA, USA: USENIX Association, 2004, pp. 10–10.

[62] Wong, E., Balasubramanian, P., Alvisi, L., Gouda, M., and Shmatikov, V., *Truth in advertising: lightweight verification of route integrity*, PODC ’07: Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing. New York, NY, USA: ACM, 2007, pp. 147–156.

[63] *IETF Routing Protocol Security Requirements Working Group*, https://datatracker.ietf.org/wg/rpsec/about/

[64] *IETF Secure Inter-Domain Routing Working Group*, https://datatracker.ietf.org/wg/sidr/about/

[65] Lepinski, M., Kent, S., and Kong, D., *A Profile for Route Origin Authorizations (ROAs)*, RFC 6482, DOI 10.17487/RFC6482, February 2012. https://www.rfc-editor.org/info/rfc6482

[66] Huston, G. and Michaelson, G., *Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)*, RFC 6483, DOI 10.17487/RFC6483, February 2012. https://www.rfc-editor.org/info/rfc6483

[67] Bush, R. and Austein, R., *The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1*, RFC 8210, DOI 10.17487/RFC8210, September 2017. https://www.rfc-editor.org/info/rfc8210

[68] Lepinski, M., and K. Sriram, K., Eds., *BGPsec Protocol Specification*, RFC 8205, DOI 10.17487/RFC8205, September 2017. https://www.rfc-editor.org/info/rfc8205

[69] Sriram, K., Ed., *BGPsec Design Choices and Summary of Supporting Discussions*, RFC 8374, DOI 10.17487/RFC8374, April 2018. https://www.rfc-editor.org/info/rfc8374

[70] Q. Li, Y.C. Hu, and Zhang, X., *Even Rockets Cannot Make Pigs Fly Sustainably: Can BGP be Secured with BGPsec*”, SENT’14, February 2014. https://www.research-collection.ethz.ch/bitstream/handle/20.500.11850/86415/1/eth-8844-01.pdf

[71] Sniders, J., *Practical everyday BGP filtering with AS_PATH filters: Peer Locking*, NANOG 56, June 2016. https://archive.nanog.org/sites/default/files/Snijders_Everyday_Practical_Bgp.pdf

[72] Azimov, A., et. al., *Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization*, February 2021. Internet Draft: https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification

[73] Gao, L., *On Inferring Autonomous Relationships in the Internet*, IEEE/ACM Transactions on Networking, 9 (6), 733-745, 2001

[74] Lychev, R., Goldberg, S., and Shapira, M., *BGP security in partial deployment: is the juice worth the squeeze?*, SIGCOMM '13: Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM, August 2013 Pages 171–182 https://doi.org/10.1145/2486001.2486010

[75] Testart, C., *Reviewing a Historical Internet Vulnerability: Why Isn't BGP More Secure and What Can We Do About it?*, *TPRC 46: The 46th Research Conference on Communication, Information and Internet Policy 2018*, September 2018.

[76] Feamster, N., Balakrishnan, H., and Rexford, J., *Some Foundational Problems in Interdomain Routing*, 3rd ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets), San Diego, CA, November 2004.

[77] Huston, G., *The CIDR Report*, https://www.cidr-report.org/as2.0/#Bogons

[78] *MANRS Observatory*, https://observatory.manrs.org/#/overview

[79] Weis, B., *Why IPsec and BGP don't play well together in real networks*, Security Area Working Group presentations, IETF 66, July 2006. https://www.ietf.org/proceedings/66/slides/saag-2.pdf

[80] Rescorla, E., *The Transport Layer Security (TLS) Protocol Version 1.3*, RFC 8446, DOI 10.17487/RFC8446, August 2018. https://www.rfc-editor.org/info/rfc8446

[81] Iyengar, J., and M. Thomson, M., Eds., *QUIC: A UDP-Based Multiplexed and Secure Transport*, RFC 9000, DOI 10.17487/RFC9000, May 2021. https://www.rfc-editor.org/info/rfc9000

[82] *Internet Routing Registry*, http://www.irr.net

[83] McPherson, D., Amante, S., Osterweil, E., Blunk, L., and Mitchell, D., *Considerations for Internet Routing Registries (IRRs) and Routing Policy Configuration*, RFC 7682, DOI 10.17487/RFC7682, December 2015. https://www.rfc-editor.org/info/rfc7682

[84] Rekhter, Y., *Routing in a Multi-provider Internet*, RFC 1787, DOI 10.17487/RFC1787, April 1995. https://www.rfc-editor.org/info/rfc1787

[85] Villamizar, C., Alaettinoglu, C., Meyer, D., and Murphy, S., *Routing Policy System Security*, RFC 2725, DOI 10.17487/RFC2725, December 1999. https://www.rfc-editor.org/info/rfc2725

[86] Steenbergen, R., *What's wrong with IRR*, NANOG 44, October 2006. https://archive.nanog.org/meetings/nanog44/presentations/Tuesday/RAS_irrdata_N44.pdf

[87] Christian, B., and Tauber, T., Eds., *BGP Security Requirements*, Internet Draft, November 2008. https://datatracker.ietf.org/doc/html/draft-ietf-rpsec-bgpsecrec-10

[88] Kumari, W. and Sriram, K., *Recommendation for Not Using AS_SET and AS_CONFED_SET in BGP*, BCP 172, RFC 6472, DOI 10.17487/RFC6472, December 2011. https://www.rfc-editor.org/info/rfc6472

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net