The Internet Protocol Journal

April 2022 Volume 25, Number 1

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

You can download IPJ back issues and find subscription information at:
www.protocoljournal.org

ISSN 1944-1134

From The Editor

I have spent some time in recent months studying the history and development of the world-wide telephone network. Broadly speaking, the telephone network has evolved in two directions away from the traditional system of interconnected public and private telephone switches and their associated hard-wired telephones. First, starting in the mid-1980s we saw the introduction of mobile devices and networks, eventually leading to what we refer to as "smartphones" today. Secondly, many of the traditional telephone networks have been augmented or completely replaced by numerous systems that employ *Voice-over Internet Protocol* (VoIP) technologies. In spite of the differences in technologies, it is still possible to place and receive voice calls to *telephone numbers,* an addressing system that has proved remarkably resilient to growth and technological evolution since its introduction some 130 years ago. (The first commercial telephone exchange was installed in 1892 in La Porte, Indiana.) Unlike IP addresses, telephone numbers are not fixed-length, nor are they managed by a single global entity, but for such a system to work we do rely on a unique set of country codes and numerous interconnection agreements, thus there are some similarities to the way the Internet operates.

I have also been reading numerous recent postings to the "internet-history" e-mail list, operated by The Internet Society. If you're interested in hearing from Internet pioneers such as Vint Cerf, Brian Carpenter, Noel Chiappa, Jack Haverty, and many others, this list is a great place to start. You can find further details here:
`https://elists.isoc.org/mailman/listinfo/internet-history`

In our previous issue, Geoff Huston presented Part 1 of "A Survey on Securing Inter-Domain Routing." He described the design and operation of the *Border Gateway Protocol* (BGP), the threat model, and the requirements from a security framework for BGP. In this issue, Geoff concludes the survey by looking at the various proposals to add security to the routing environment and evaluates 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.

—Ole J. Jacobsen, Editor and Publisher
`ole@protocoljournal.org`

A Survey on Securing Inter-Domain Routing

Part 2 – Approaches to Securing BGP

by Geoff Huston, APNIC

The *Border Gateway Protocol* (BGP) is the inter-domain routing protocol of the Internet, and after some thirty years of operation it is now one of the more venerable of the core protocols on the Internet. 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 evaluate 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.

Securing the Operation of BGP Sessions

BGP uses a long-held *Transmission Control Protocol* (TCP) session, and you can use the same approaches to secure any TCP session[1] 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.

Generalized TTL Security Mechanism

The *Generalized TTL Security Mechanism* (GSTM), originally described in [2] and updated in [3], 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 *Time To Live* (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 the receiver can discard all inbound TCP packets within this session with a TTL of 254 or less.

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.

You can use this GTSM approach 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.

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* (DoS) attacks if the processing load of the cryptographic functions to detect bogus packets is sufficiently high. The target still has to process bogus packets 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 the 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 *Message-Digest 5* (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 combination 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, and it is not 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 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.

TCP Authentication Option

A somewhat different approach—the TCP *Authentication Option* (AO)[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 or not a key ID has been appended to the *Message Authentication Code* (MAC). The message-digest algorithm in this case is specified as HMAC-MD5-96, although you can use other algorithms if you configure them 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* before using this mechanism. This database contains a description of the supported TCP connections, the key set, the MAC algorithm, and MAC length.

IPSec

Internet Protocol Security (IPsec) is a suite of protocols that operate at the IP level of the protocol stack; these protocols secure 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 sessions. This implementation is one of public/private key cryptography, and it can ensure the confidentiality and integrity of all IP messages passed between two hosts. You can use IPsec to secure BGP sessions, and it provides greater levels of assurance than MD5 offers.

However, IPsec is not widely used in the public Internet for the purpose of securing BGP sessions[7,9], 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 relates to the complications for rekeying *Internet Key Exchange* (IKE)/IPsec sessions, and the observation that processing load to detect bogus packets is considerably higher with IPsec than with MD5. Using IPsec for BGP exposes a DoS 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 degradation of other concurrent router functions.

More Options

As was observed in Part 1 of this survey, there are many alternatives here, including *Transport Layer Security* (TLS)[80] and *Quick UDP Internet Connections* (QUIC)[81], but more choice is not a substitute for better quality.

These session-level encryption approaches that applications use provide no better answer to dynamic rekeying, and they 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. 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 reduces considerably the radius of potential eavesdroppers and attackers.

Securing the Integrity of Routing Information Passed in BGP

One of the earlier recognised works that addressed routing security was the 1988 study on *Byzantine Robustness* by Radia Perlman[9]. If failure or malicious behaviour on the part of one or more entities in the system occurs, 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 sporadically, requesting validation or credentials from a remote system when necessary.

Although the link-state approach described in this paper does not exactly match the inter-domain 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 *Autonomous System* (AS) prior to the destination AS for the current route and proposed digitally signing all fixed fields in the UPDATE message. The predecessor attribute constructs 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, because partial adoption would create gaps in any assurance that a predecessor attribute could provide. Their approach was similar to the earlier *Interdomain Routing Protocol* (IDRP) work[12]. IDRP eschewed the use of TCP and included a reliable flow-controlled transport into the IDRP protocol, also including numerous 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, *Network Layer Reachability Information* (NLRI) *Origin AS Verification*, used the *Domain Name System* (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 *Domain Name System Security Extensions* (DNSSEC)-signed DNS record and DNSSEC validation[14], presumably, although this work was contemporaneous with DNSSEC and did not use 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 that the BGP speaker used was located within an address prefix that was being validated prior to local acceptance of the route associated with 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 third-party validation. 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, dating back to the routing work of the early 1990s in the *Routing Arbiter* project that was part of the US *National Science Foundation Network* (NSFNET), and a project coordinated under the auspices of the RIPE IRR 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 a response to the need described in RFC 1787[84] for improving global consistency by allowing providers to 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 that an AS provides to the IRR 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, 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 current information sits along with contingency information that may never be actually used.

Although 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 if they 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. In addition, although 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 looks like 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 the enrollment practice that governs the authority to generate such signatures is robust. Given such a practice, the consequent observation is that whether or not this digital signature framework is placed into the DNS via a DNSSEC framework[15] or into a framework of X.509 certificates and an associated *Public Key Infrastructure* (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 the database being queried is incomplete.

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

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:

- No protocol UPDATE messages have 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
- The peer is authorised to advertise information on behalf of the peer AS

In addition, for every prefix and its originating AS, the prefix must be a validly allocated prefix, and the prefix "right-of-use" holder must have authorised the advertisement of the prefix and 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*. Produced by an address holder, an *address attestation* 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; it 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.

Figure 1 shows inter-operation and information exchange between sBGP elements.

Figure 1: sBGP

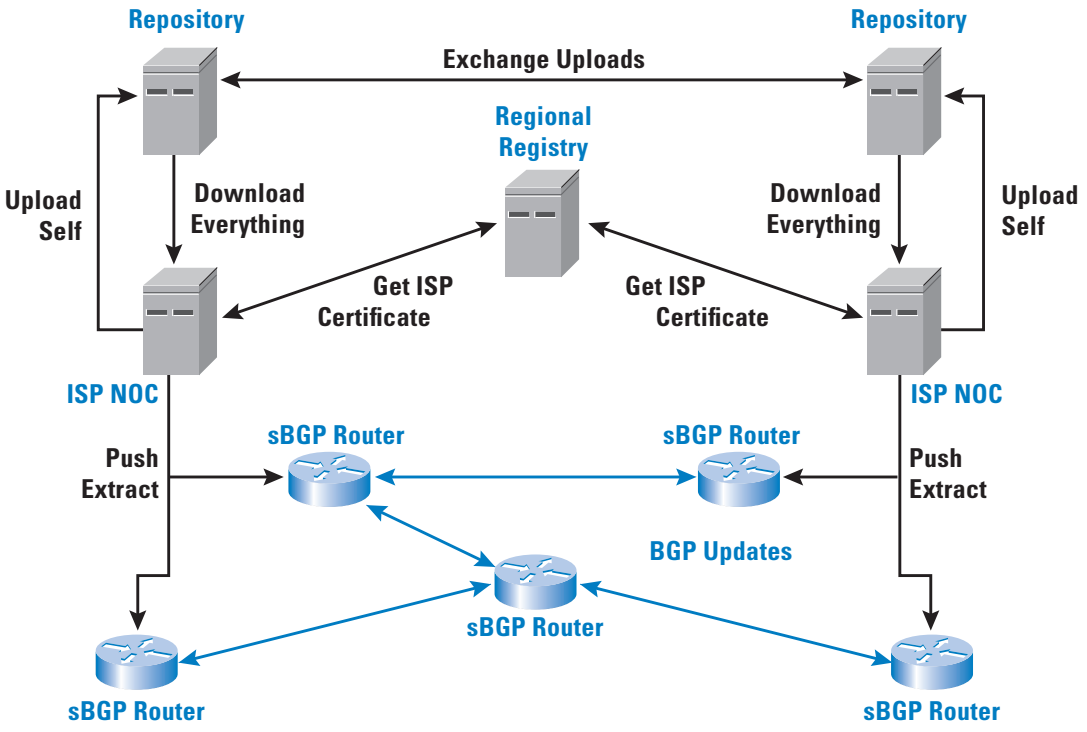

Certificates for each Internet Service Provider (ISP) are issued by the regional registries. The ISPs exchange public keys through special repositories. The keys are pushed to the sBGP routers, which validate the BGP UPDATE messages.

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.

Numerous significant issues have been 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].

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 that soBGP posed 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 AS Path validation is quite different from that of 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 the means of validation of these certificates. At the time of its development, 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 that was derived from the established AS and address distribution framework. This consideration appears to have been pragmatic at the time, because no such PKI existed then, and it was unclear if the various address registries were in a position to undertake this type of role of administering such a specialised PKI in any case. This situation left open the problem of how to establish trust anchors for validation of these signed objects, a rather significant deficiency in the validation framework of soBGP.

Instead of sBGP 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 local peer ASes, signed by the AS 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: soBGP Peering Network

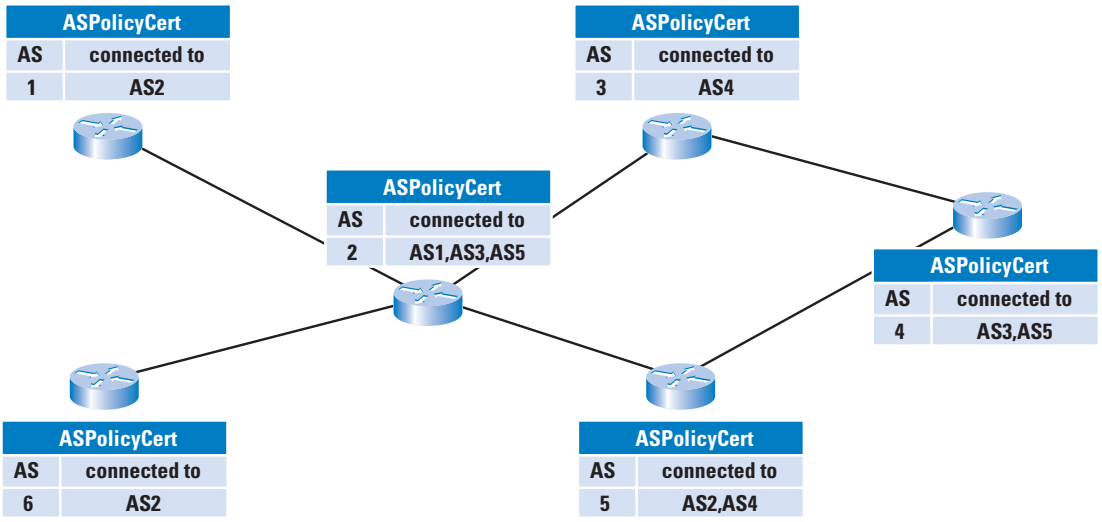

The ASPolicyCert is a self-signed certificate containing routing policies. An UPDATE message originating at AS4 would necessarily take the path {AS4,AS5,AS2,AS1} instead of {AS4,AS3,AS2,AS1} because the connection between AS2 and AS3 would not be regarded as valid.

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, because the derivation of authority to speak about addresses is very unclear in this model, but this absence occurred because the protocol was developed prior to the completion of the work on the *Resource PKI* (RPKI).

It is clear that soBGP could be readily adapted to use the RPKI as its trust and authority framework.

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

Pretty Secure BGP

Pretty Secure BGP (psBGP)[22] put forward the proposition that the proposals relating to the authentication 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 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 it may no longer be accurate at the time the relying party accesses it.

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

Inter-domain Route Validation

All of the approaches to securing the semantics of BGP described in this section so far entail changes to the operation of BGP itself and operate most effectively in an environment of universal deployment. In practical terms this scenario is unlikely, 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 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 about 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 that the IRRs already use, 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 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 IRV server of the AS 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 a set of AS peer connections that correspond to the AS sequence exists. Here the validation test is that each AS in the sequence is currently advertising this prefix to the next AS in sequence.

This IRV architecture has numerous 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. Whether or not the query and response are preconditions to the local acceptance of a BGP route is also unspecified. 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 true that the IRV associated with the originating AS may be reachable only 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 was 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 compared to an alternative approach of further augmentation to the existing IRR framework.

Figure 3: AS Path Verification Using IRV

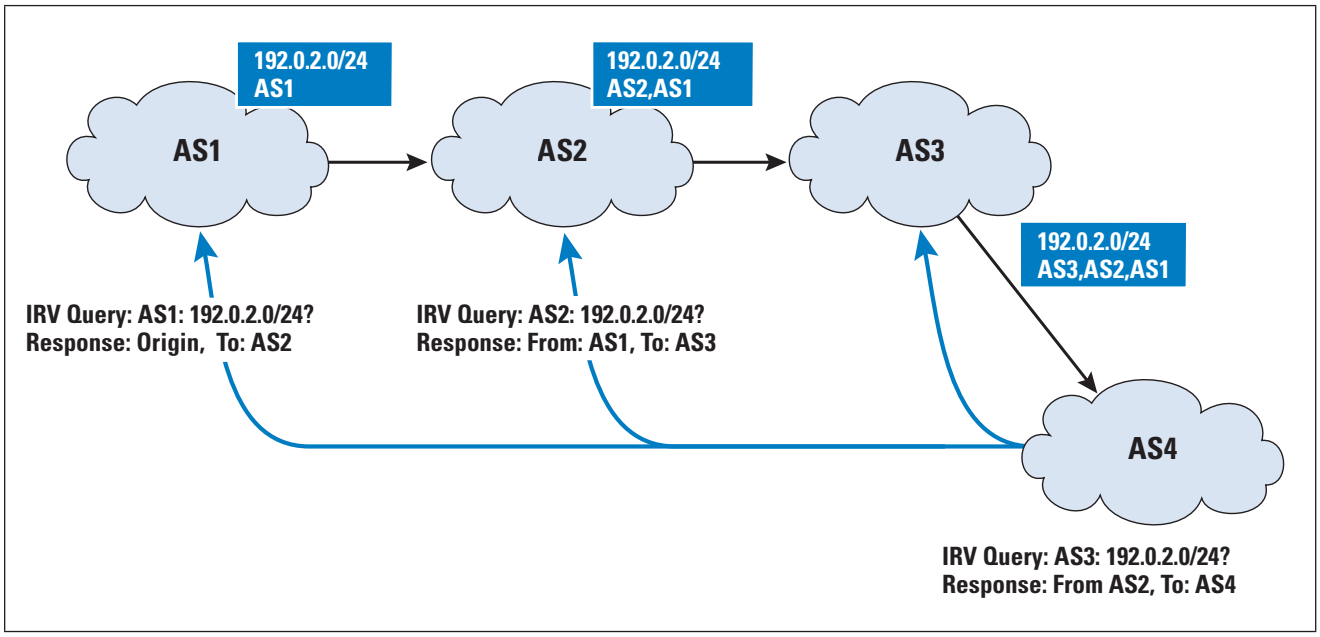

A similar approach within the IRR framework that integrates the concept of an address and AS PKI could make provision for signed responses in a way that allows the IRR client to authenticate that the response is accurate, current, and contains information that has been digitally signed by the AS or prefix 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].

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 influence local route-selection decisions
- *Denial of Service*, where the attack attempts to overwhelm the intended victim's resources
- *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 multiparty collusion. SPV takes the approach of tree-authenticated hash values and applies it specifically to AS Path validation as an alternative to the nested digital signature structure proposed as the AS Path validation mechanism of sBGP. The SPV study paper claims significantly improved processor performance using this technique, based on the difference in computational complexity for asymmetric cryptography from symmetric cryptography as used in hash functions.

This proposal falls into the category of proposals that call 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 requirement is the weakest part of the approach in terms of performance evaluation, because 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, obviates any need for a BGP re-advertisement 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.

Signature Amortisation and Aggregate Signatures

If the signature load of sBGP is the problem, then how can this load be reduced? Numerous papers have addressed this question.

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 places a topology description as a vector in an equivalent of an AS connectivity attestation that 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 the sBGP 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 are sent as the signed UPDATE to each peer. This technique improves upon [35], which uses bilinear maps instead of Merkle hash trees.

Exploiting Path Stability

You also can mitigate the validation overhead 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 1-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 numerous 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 *Minimum Route Advertisement Interval* (MRAI) timer. The third part of the approach was to exploit the stability of path advertisements to amortise cryptographic operations over many validations 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 that 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 such a sudden change might indicate an attack to the routing system. pgBGP does not provide completely automated security, because 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. You can deploy this proposal incrementally, and it imposes little overhead on the routing system. It is a method to mitigate effects of an attack to the routing system, not an effective mechanism for prevention of such attacks.

Detecting Prefix Hijacking

One special case of routing attacks that is considered a major threat and evokes high interest in the research community is *prefix hijacking*. A considerable amount of research has been undertaken to provide security against this single form of attack. The approaches describe possible methods of detecting prefix hijacking[38,39,40,41], as well as complete systems and implementations of prefix hijacking detection in order to possibly react to 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 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.

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 addressed in [50], where a study of BGP update dynamics showed 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 *Least Recently Used* (LRU) cache-replacement algorithm.

When distance-vector algorithms react to a change in prefix reachability, many 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 depend 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* (ISPs) and noted that these extended convergence times and larger path exploration events occurred at lower levels of the tiering hierarchy. It 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 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 the BGP route-selection process and noted that in the absence of the *Multiple Exit Discriminator* (MED) and *Route Reflectors*, 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.

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 ensuring 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 subject to not only attacks that try to subvert the routing system, but also 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, not only on a next-hop basis as *Stealth Probing*[58] intends to, because carriers might span over multiple ASes and synthesise false routing information that spans multiple AS hops.

Proposed approaches focus mainly 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 Fatih approach[60] instead focuses on using traffic summary func tions, and comparing their results with those of other routers, allowing detection of ASes that provide anomalous values. These traffic summary functions seem to be prone to inaccuracy because of a variety of applications running on routers that 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 that try to ensure data plane/control plane consistency. Like other data-plane security solutions, this approach seems infeasible, because 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, and is achieved by probing data paths by 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 provides comprehensive data-plane security. Authenticity of a data path from start to end could easily be forged by two ASes deploying tunnels between them, and thus disabling the possibility to effectively verify the data path by a third party.

IETF Activity – RPKI, ROV, BGPSec, and ASPA

Following numerous 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 *Internet Engineering Steering Group* (IESG) chartered the *Secure Inter-Domain Routing Working Group* (SIDR) in 2006[64]. The charter for this effort presented some problems, 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, 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 decision turned out to be very important because 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.

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 RPKI framework this information is published as a signed *Route Origin Authorization* (ROA)[65,66].

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

Many additional implications are associated with publishing a ROA. The first is that no other AS has permission to announce that prefix when a cryptographically valid ROA exists in the RPKI system. If the prefix holder wishes to authorize multiple ASes to originate a route for this prefix, the prefix holder must generate multiple ROAs, meaning 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 is more specific than the prefix listed in the ROA. You can use a *MaxLength* attribute of a ROA to define a range of more specific prefix lengths than a ROA permits. 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 that is unaware of the permission.

The RPKI framework has no symmetric instrument 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.

One more important component of the ROV framework is the *RPKI to Router Protocol* (RTR)[67]. This protocol allows you to remove a crypto engine 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, which 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.

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 nontransitive 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 that 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 overview, 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 design decisions is available in [69].

Figure 4: BGPsec Handling of AS Path Signature Structure

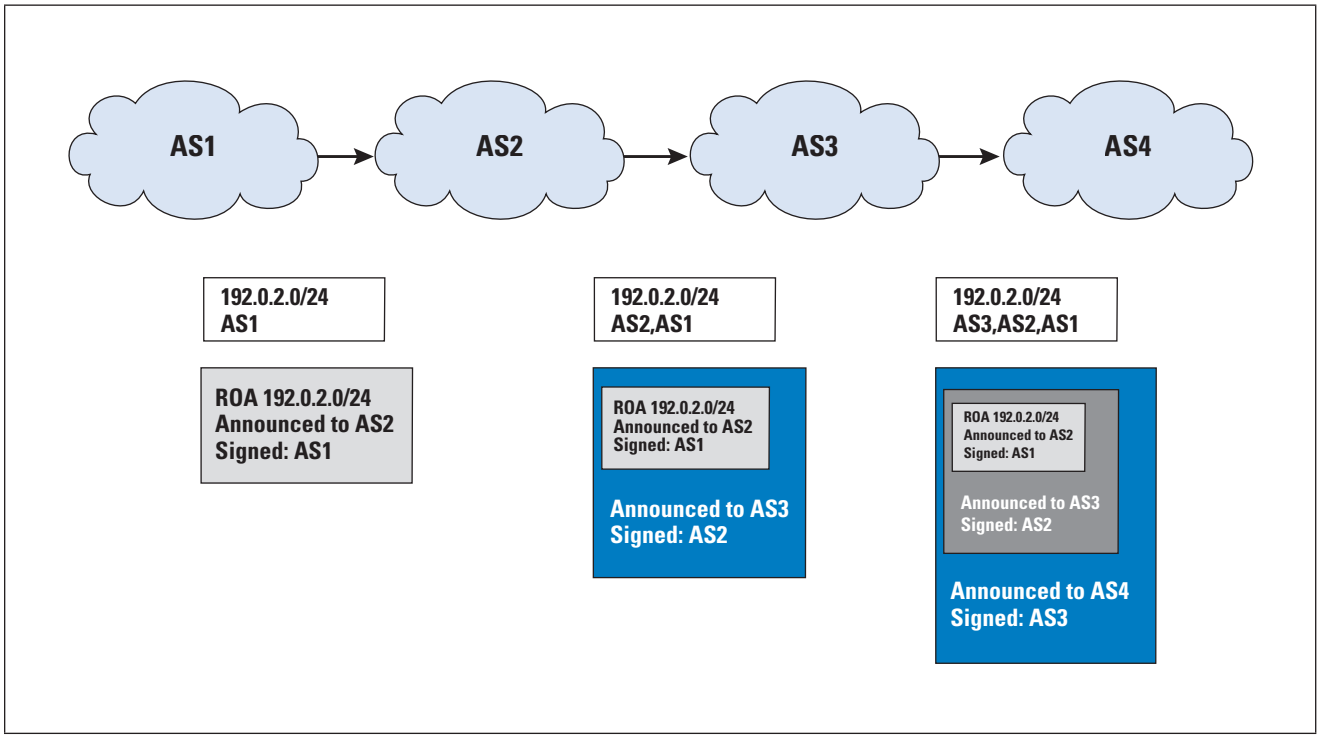

Stepwise AS Path validation cannot tolerate AS Sets in this approach, nor AS Confederation Sets, nor sets that are in the process of being deprecated in response to this limitation[88]. In a similar vein, BGP Route Reflectors require special processing, as do private AS numbers.

This design approach has numerous consequences.

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, meaning that if 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 load is a feasible impost on the routing infrastructure of the Internet, 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 they also hold secure private keys to perform signing in real time 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, there is no protection against spurious WITHDRAW messages, no ability to ascertain the alignment of the route object with the forwarding state of the network, 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, a more extreme view says that BGPsec cannot achieve any of the security properties because of the fundamental design principles of BGP and BGPsec. In one research paper[70], it is asserted that routes can still be hijacked in BGPsec, and routing loops can still appear. The authors of the paper 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.

Autonomous System Provider Authorization

The issue with the overall SIDR approach to BGP security is that if BGPsec is impractical, we cannot rely on ROV alone. All a determined routing attacker would need to do is tack on the originating AS to a synthesised AS Path and then could place any AS sequence 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 ROV offers 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.

Many proposals to address this shortfall have been made. 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 two of these Tier 1 ASes should appear in any AS Path, and if there are two 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 many common forms of route leaks, and its implementation is quite simple and very lightweight.

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, a proposal to use RPKI-signed AS adjacency attestations as a response to the issues with BGPsec has been made.

The proposal has a slight twist, however, which is different from soBGP in that an element of routing policy is also used in the *Autonomous System Provider Authorization* (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 (because 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 that describes the inter-AS topology[73]. The property that the ASPA proposal uses 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, and then a set of provider-to-customer (down) AS pairs. Any AS sequence that contains a down and then an up (or a "valley") represents a customer AS leaking routes learned from one provider to another (Figure 5).

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

Figure 5: ASPA and Route Leaks

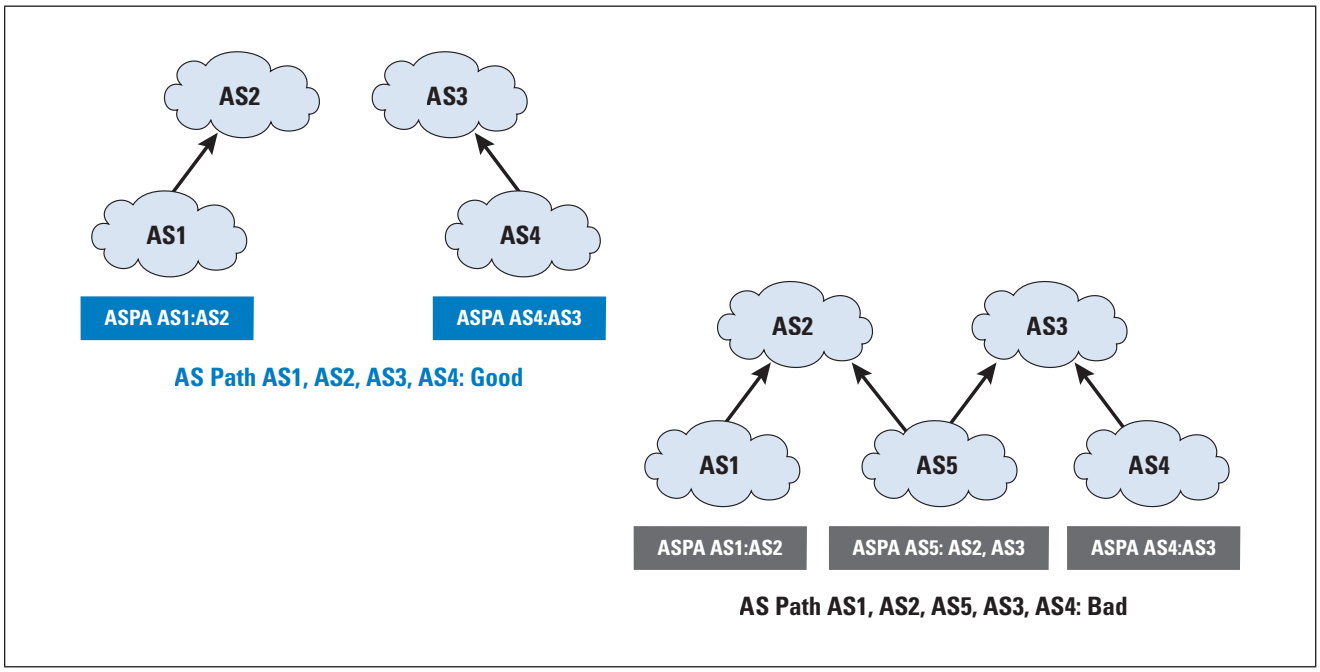

ASPA still provides some benefit, even in scenarios of partial deployment. After an AS issues an ASPA, a routing attacker can include this AS in a synthetic AS Path attribute only if it also includes an adjacent provider AS, and the synthetic AS pair can be inserted in the "front part" of the AS Path (customer-to-provider) only 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 you can use 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 3 years the work remains a study item in the *SIDR Operations* (SIDROPS) Working Group of the IETF. Part of the issue here is that the SIDROPS Working Group has had its collective attention diverted away from the issues of BGP security mechanisms and AS Path validation and has taken on the role of the RPKI operational maintenance working group. In addition, in the area of RPKI operations the topic that presently takes 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 reopen 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.

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 can provide clear indications when convergence to a stable routing state is 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 that BGP uses vulnerable to attack, and the messages that BGP uses vulnerable to alteration that would disrupt the network routing system, 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 overwhelm a BGP speaker and force its reset, because BGP is a method of directly accessing the processing unit of a router and a saturation attack can cause processor and memory overload. Other attacks are aimed at altering the forwarding state of a router, 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, and 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 it 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 not only the AS path as presented in a BGP route advertisement matches the BGP propagation path taken by the prefix advertisement, and also that the current forwarding state of the network 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 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.

Following 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, if 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 answer to this question is that securing one element of a 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 choice 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 deploying 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 the Route Origination Validation offer can mitigate some forms of routing incidents but are ineffectual against other forms of route manipulation that preserve the origination data. Practically, is this mitigation enough? Or do we need to also deploy some mechanism that allows detection of various forms of 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 the BGPsec 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 to consider 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?

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" grew from under 20,000 distinct prefixes to some 1,000,000 distinct prefixes by mid-2021, with 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.

Because of 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 it 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 make the case that the cost of a change to improve this situation will be money and effort widely spent.

References

[0] Geoff Huston, "A Survey on Securing Inter-Domain Routing, Part 1 – BGP: Design, Threats and Security Requirements," *The Internet Protocol Journal*, Volume 24, No. 3, October 2021.

[1] Steven Michael Bellovin, "Security problems in the TCP/IP Protocol Suite," ACM SIGCOMM *Computer Communication Review*, Volume 19, Issue 2, April 1, 1989.

[2] Vijay Gill, John Heasley, and David Meyer, "The Generalized TTL Security Mechanism (GTSM)," RFC 3682, February 2004.

[3] Carlos Pignataro, Pekka Savola, David Meyer, Vijay Gill, and John Heasley, "The Generalized TTL Security Mechanism (GTSM)," RFC 5082, October 2007.

[4] Andy Heffernan, "Protection of BGP Sessions via the TCP MD5 Signature Option," RFC 2385, August 1998.

[5] Ronald L. Rivest, "The MD5 Message-Digest Algorithm," RFC 1321, April 1992.

[6] Marcus Leech, "Key Management Considerations for the TCP MD5 Signature Option," RFC 3562, July 2003.

[7] Ronald P. Bonica, Allison Mankin, and Joe Touch, "The TCP Authentication Option," RFC 5925, June 2010.

[8] Karen Seo and Stephen Kent, "Security Architecture for the Internet Protocol," RFC 4301, December 2005.

[9] Radia Perlman, "Network Layer Protocols with Byzantine Robustness," MIT Doctoral Thesis, August 1988.
https://dspace.mit.edu/handle/1721.1/14403

[10] Bradley R. Smith and Jose Joaquin Garcia-Luna-Aceves, "Securing the border gateway routing protocol," in *Proceedings of GLOBECOM '96, 1996 IEEE Global Telecommunications Conference*, November 1996.

[11] Bradley R. Smith and Jose Joaquin Garcia-Luna-Aceves, "Efficient Security Mechanisms for the Border Gateway Routing Protocol," *Computer Communications*, Volume 21, No. 3, 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] Tony Bates, Randy Bush, Tony Li, and Yakov Rekhter, "DNS-based NLRI origin AS verification in BGP," Internet Draft, work in progress, July 1998.
https://datatracker.ietf.org/doc/html/draft-bates-bgp4-nlri-orig-verif-00

[14] Scott Rose, Matt Larson, Dan Massey, Rob Austein, and Roy Arends, "DNS Security Introduction and Requirements," RFC 4033, March 2005.

[15] Lutz Donnerhacke and Wouter Wijngaards, "DNSSEC protected routing announcements for BGP," Internet Draft, work in progress, May 2008.
https://datatracker.ietf.org/doc/html/draft-donnerhacke-sidr-bgp-verification-dnssec/

[16] Stephen Kent, Charles Lynn, and Karen Seo, "Secure Border Gateway Protocol (SBGP)," *IEEE Journal on Selected Areas in Communications*, Volume 18, Issue 4, April 2000.

[17] Karen Seo, Charles Lynn, and Stephen Kent, "Public-key infrastructure for the Secure Border Gateway Protocol (S-BGP)," *Proceedings DARPA Information Survivability Conference and Exposition II. DISCEX'01*, August 2002.

[18] Stephen Kent, Charlie Lynn, Joanne Mikkelson, and Karen Seo, "Secure Border Gateway Protocol (S-BGP) – Real World Performance and Deployment Issues," in *Proceedings of the 7th Annual Network and Distributed System Security Symposium*, February 2000.

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

[20] Russ White, "Securing BGP Through Secure Origin BGP," *The Internet Protocol Journal*, Volume 6, No. 3, September 2003.

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

[22] Paul C van Oorschot, Tao Wan, and Evangelos Kranakis, "On interdomain routing security and pretty secure BGP (psBGP)," ACM *Transactions on Information System Security*, Volume 10, Issue 3, July 2007.

[23] Philip R. Zimmermann, *The Official PGP User's Guide*, ISBN-13: 978-0262740173, MIT Press, 1995.

[24] Geoff Huston, "Exploring Autonomous System Numbers," *The Internet Protocol Journal*, Volume 9, No. 1, March 2006.

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

[26] Tony Bates, Elise Gerich, Laurent Joncheray, Jean-Michel Jouanigot, Daniel Karrenberg, Marten Terpstra, and Jessica Yu, "Representation of IP Routing Policies in a Routing Registry (ripe-81++)," RFC 1786, March 1995.

[27] David Kessens, Tony Bates, Cengiz Alaettinoglu, David Meyer, Curtis Villamizar, Marten Terpstra, Daniel Karrenberg, and Elise Gerich, "Routing Policy Specification Language (RPSL)," RFC 2622, June 1999.

[28] Joao Damas, Andrei Robachevsky, Larry Blunk, and Florent Parent, "Routing Policy Specification Language next generation (RPSLng)," RFC 4012, March 2005.

[29] Robert Kisteleki and Jos Boumans, "Securing RPSL Objects with RPKI Signatures," Internet Draft, work in progress, October 2008.
https://datatracker.ietf.org/doc/html/draft-kisteleki-sidr-rpsl-sig

[30] Yih-Chun Hu, Adrian Perrig, and Marvin Sirbu, "SPV: Secure Path Vector Routing for Securing BGP," ACM SIGCOMM *Computer Communication Review,* Volume 34, Issue 4, October 2004.

[31] Barath Raghavan, Saurabh Panjwani, and Anton Mityagin, "Analysis of the SPV Secure Routing Protocol: Weaknesses and Lessons," ACM SIGCOMM *Computer Communication Review*, Volume 37, Issue 2, April 2007.

[32] David M. Nicol, Sean W. Smith, and Meiyuan Zhao, "Efficient Security for BGP Route Announcements," TR-2003-440, Dartmouth College, Computer Science, 2003.

[33] Meiyuan Zhao, Sean W. Smith, and David M. Nicol, "Aggregated Path Authentication for Efficient BGP security," in *CCS '05: Proceedings of the 12th ACM Conference on Computer and Communications Security*, November 2005.

[34] Ralph C. Merkle, "Protocols for Public Key Cryptosystems," *IEEE Symposium on Security and Privacy*, April 1980.

[35] Dan Boneh, Craig Gentry, Ben Lynn, and Hovav Shacham, "Aggregate and verifiably encrypted signatures from bilinear maps," in *Advances in Cryptology - EUROCRYPT 2003*, Lecture Notes in Computer Science, Volume 2656, Springer Verlag, January 2003.

[36] Kevin Butler, Patrick McDaniel, and William Aiello, "Optimizing BGP Security by Exploiting Path Stability," in *CCS '06: Proceedings of the 13th ACM conference on Computer and Communications Security*, October 2006.

[37] Josh Karlin, Stephanie Forrest, and Jennifer Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," in *ICNP '06: Proceedings of the 2006 IEEE International Conference on Network Protocols*, IEEE Computer Society, November 2006.

[38] Jian Qiu, Lixin Gao, Supranamaya Ranjan, and Antonio Nucci, "Detecting Bogus BGP Route Information: Going Beyond Prefix Hijacking," in *Third International Conference on Security and Privacy in Communications Networks and the Workshops — SecureComm 2007*, September 2007.

[39] Changxi Zheng, Lusheng Ji, Dan Pei, Jia Wang, and Paul Francis, "A Light-Weight Distributed Scheme for Detecting IP Prefix Hijacks in Real-Time," ACM SIGCOMM *Computer Communication Review*, Volume 37, Issue 4, October 2007.

[40] Xin Hu and Z. Morley Mao, "Accurate Real-time Identification of IP Prefix Hijacking," in *SP '07: Proceedings of the 2007 IEEE Symposium on Security and Privacy*, IEEE Computer Society, May 2007.

[41] Noor Hadi Hammood and Bahaa Al-Musawi, "Using BGP Features Towards Identifying Type of BGP Anomaly," in *Proceedings of 2021 International Congress of Advanced Technology and Engineering (ICOTEN)*, July 2021.

[42] Christopher Kruegel, Darren Mutz, William Robertson, and Fredrik Valeur, "Topology-Based Detection of Anomalous BGP Messages," in *Recent Advances in Intrusion Detection*, Lecture Notes in Computer Science, Volume 2820, Springer Verlag, February 2003.

[43] Mohit Lad, Daniel Massey, Dan Pei, Yiguo Wu, Beichuan Zhang, and Lixia Zhang, "PHAS: A Prefix Hijack Alert System," in *USENIX-SS'06: Proceedings of the 15th conference on USENIX Security Symposium*, USENIX Association, July 2006.

[44] E-yong Kim, Klara Nahrstedt, Li Xiao, and Kunsoo Park, "Identity-Based Registry for Secure Interdomain Routing," in *ASIACCS '06: Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security*, March 2006.

[45] Zheng Zhang, Ying Zhang, Y. Charlie Hu, and Zhuoqing Morley Mao, "Practical defenses against BGP prefix hijacking," in *CoNEXT '07: Proceedings of the 2007 ACM CoNEXT Conference*, December 2007.

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

[47] Xiaoliang Zhao, Dan Pei, Lan Wang, Dan Massey, Allison Mankin, S. Felix Wu, and Lixia Zhang, "An analysis of BGP multiple origin AS (MOAS) conflicts," in *IMW '01: Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, November 2001.

[48] Xiaoliang Zhao, Dan Pei, Lan Wang, Dan Massey, Allison Mankin, S. Felix Wu, and Lixia Zhang, "Detection of invalid routing announcement in the Internet," in *Proceedings. International Conference on Dependable Systems and Networks*, IEEE, December 2002.

[49] Zheng Zhang, Ying Zhang, Y. Charlie Hu, Zhuoqing Morley Mao, and Randy Bush, "ISPY: Detecting IP Prefix Hijacking on My Own," ACM SIGCOMM *Computer Communication Review*, Volume 38, Issue 4, October 2008.

[50] Geoff Huston, "Measures of self-similarity of BGP updates and implications for securing BGP," in *Proceedings of the 8th International Conference on Passive and Active Network Measurement (PAM 2007)*, Volume 4427, Springer Verlag, April 2007.

[51] Ricardo Oliveira, Beichuan Zhang, Dan Pei, Rafit Izhak-Ratzin, and Lixia Zhang, "Quantifying Path Exploration in the Internet," in *IMC '06: Proceedings of the 6th ACM SIGCOMM Conference on Internet measurement*, October 2006.

[52] Jaideep Chandrashekar, Zhenhai Duan, Zhi-Li Zhang, and Jeff Krasky, "Limiting Path Exploration in BGP," in *Proceedings of the IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies*, March 2005.

[53] Tony Li and Geoff Huston, "BGP Stability Improvements," Internet Draft, work in progress, June 2007. **https://datatracker.ietf.org/doc/html/draft-li-bgp-stability-01.txt**

[54] Geoff Huston, Mattia Rossi, and Grenville Armitage, "A Technique for Reducing BGP Update Announcements through Path Exploration Damping," in *IEEE Journal on Selected Areas in Communications*, Volume 28, Issue 8, October 2010.

[55] Nick Feamster and Jennifer Rexford, "Network-Wide Prediction of BGP Routes," *IEEE/ACM Transactions on Networking*, Volume 15, Issue 2, April 2007.

[56] Dan Wendlandt, Ioannis C. Avramopoulos, David G. Andersen, and Jennifer Rexford, "Don't Secure Routing Protocols, Secure Data Delivery," in *Proceedings of the 5th ACM Workshop on Hot Topics in Networks (Hotnets-V)*, November 2006.

[57] Zhuoqing Morley Mao, Jennifer Rexford, Jia Wang, and Randy H. Katz, "Towards an accurate AS-level traceroute tool," in *SIGCOMM '03: Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, August 2003.

[58] Ioannis C. Avramopoulos, and Jennifer Rexford, "Stealth probing: efficient data-plane security for IP routing," in *ATEC '06: Proceedings of the Annual Conference on USENIX '06 Annual Technical Conference*, USENIX Association, January 2006.

[59] Venkata N. Padmanabhan and Daniel R. Simon, "Secure Traceroute to Detect Faulty or Malicious Routing," ACM SIGCOMM *Computer Communication Review*, Volume 33, Issue 1, January 2003.

[60] Alper Tugay Mızrak, Yu-Chung Cheng, Keith Marzullo, and Stefan Savage, "Fatih: Detecting and isolating malicious routers," in *DSN '05: Proceedings of the 2005 International Conference on Dependable Systems and Networks*, IEEE Computer Society, July 2005.

[61] Lakshminarayanan Subramanian, Volker Roth, Ion Stoica, Scott Shenker, and Randy H. Katz, "Listen and Whisper: Security Mechanisms for BGP," in *NSDI'04: Proceedings of the 1st Symposium on Networked Systems Design and Implementation*, USENIX Association, March 2004.

[62] Edmund L. Wong, Praveen Balasubramanian, Lorenzo Alvisi, Mohamed G. Gouda, and Vitaly Shmatikov, "Truth in Advertising: Lightweight Verification of Route Integrity," in *PODC '07: Proceedings of the Twenty-sixth Annual ACM Symposium on Principles of Distributed Computing*, August 2007.

[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] Matt Lepinski, Derrick Kong, and Stephen Kent, "A Profile for Route Origin Authorizations (ROAs)," RFC 6482, February 2012.

[66] Geoff Huston and George Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infra structure (PKI) and Route Origin Authorizations (ROAs)," RFC 6483, February 2012.

[67] Rob Austein and Randy Bush, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1," RFC 8210, September 2017.

[68] Matthew Lepinski, "BGPsec Protocol Specification," RFC 8205, September 2017.

[69] Kotikalapudi Sriram, "BGPsec Design Choices and Summary of Supporting Discussions," RFC 8374, April 2018.

[70] Qi Li, Yih-Chun Hu, and Xinwen Zhang, "Even Rockets Cannot Make Pigs Fly Sustainably: Can BGP be Secured with BGPsec?" in *Proceedings of the 2014 Network and Distributed System Security (NDSS) Symposium*, February 2014.

[71] Job Sniders, "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] Alexander Azimov, Eugene Bogomazov, Randy Bush, Keyur Patel, and Job Snijders, "Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization," February 2021, Internet Draft, work in progress, **https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification**

[73] Lixin Gao, "On Inferring Autonomous Relationships in the Internet," *IEEE/ACM Transactions on Networking*, Volume 9, Issue 6, December 2001.

[74] Robert Lychev, Sharon Goldberg, and Michael Shapira, "BGP Security in Partial Deployment: Is the Juice Worth the Squeeze?" in *SIGCOMM '13: Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, August 2013.

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

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

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

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

[79] Brian Weis, "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] Eric Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.3," RFC 8446, August 2018.

[81] Jana Iyengar and Matin Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport," RFC 9000, May 2021.

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

[83] Dave Mitchell, Larry J. Blunk, Danny McPherson, Shane Amante, and Eric Osterweil, "Considerations for Internet Routing Registries (IRRs) and Routing Policy Configuration," RFC 7682, December 2015.

[84] Yakov Rekhter, "Routing in a Multi-provider Internet," RFC 1787, April 1995.

[85] Sandy Murphy, Curtis Villamizar, Cengiz Alaettinoglu, and David M. Meyer, "Routing Policy System Security," RFC 2725, December 1999.

[86] Richard Steenbergen, "Examining the validity of IRR Data," NANOG 44, October 2006. **https://archive.nanog.org/meetings/nanog44/presentations/Tuesday/RAS_irrdata_N44.pdf**

[87] Blaine Christian and Tony Tauber, "BGP Security Requirements," Internet Draft, work in progress, November 2008. **https://datatracker.ietf.org/doc/html/draft-ietf-rpsec-bgpsecrec-10**

[88] Warren Kumari, "Recommendation for Not Using AS_SET and AS_CONFED_SET in BGP," RFC 6472, December 2011.

GEOFF HUSTON, B.Sc., M.Sc. A.M., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for building the Internet within the Australian academic and research sector in the early 1990s. He is author of numerous Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005. He served on the Board of Trustees of the Internet Society from 1992 until 2001. At various times Geoff has worked as an Internet researcher, an ISP systems architect, and a network operator. E-mail: **gih@apnic.net**

Fragments

FCC Launches Inquiry To Reduce Cyber Risks

On February 25, 2022 the *Federal Communications Commission* (FCC) Chairwoman Jessica Rosenworcel shared with her colleagues a proposed action to help protect America's communications networks against cyberattacks. Earlier that week, the Department of Homeland Security warned U.S. organizations at all levels that they could face cyber threats stemming from the Russia-Ukraine conflict. The proposal would begin an inquiry into the vulnerabilities of the Internet's global routing system.

If adopted by a vote of the full Commission, this action, called a *Notice of Inquiry*, would begin a proceeding by seeking public comment on vulnerabilities threatening the security and integrity of the *Border Gateway Protocol* (BGP), which is central to the Internet's global routing system. The inquiry would also examine the impact of these vulnerabilities on the transmission of data through email, e-commerce, bank transactions, interconnected *Voice-over Internet Protocol* (VoIP), and 911 calls—and how best to address these challenges.

BGP is the routing protocol used to exchange reachability information among independently managed networks on the Internet. BGP's initial design, which remains widely deployed today, does not include explicit security features to ensure trust in this exchanged information. As a result, a bad network actor may deliberately falsify BGP reachability information to redirect traffic. Russian network operators have been suspected of exploiting BGP's vulnerability to hijacking in the past. "BGP hijacks" can expose Americans' personal information, enable theft, extortion, and state-level espionage, and disrupt otherwise-secure transactions.

Working with its federal partners, the Commission has urged the communications sector to defend against cyber threats, while also taking measures to reinforce the nation's readiness and to strengthen the cybersecurity of vital communications services and infrastructure, especially in light of Russia's actions inside of Ukraine. Chairwoman Rosenworcel also recently shared with her colleagues a Notice of Proposed Rulemaking that would begin the process of strengthening the Commission's rules for notifying customers and federal law enforcement of breaches of *Customer Proprietary Network Information* (CPNI). The inquiry under consideration would build on those efforts. For more information, visit: `https://www.fcc.gov`

APNIC Announces "hybrid" APNIC 54 Conference in Singapore

The *Asia Pacific Network Information Centre* (APNIC) is pleased to announce that APNIC 54 will include a face-to-face event in Singapore in September 8–15, 2022. The conference will provide full online participation support so all attendees—online or in-person—receive the best possible conference experience. The *Asia Pacific Regional Internet Governance Forum* (APrIGF) and the *Asia Pacific School on Internet Governance* (APSIG) intend to co-locate their 2022 meetings with APNIC 54.

This will be the first time an APNIC conference has included a face-to-face component since APNIC 49 in March 2020 held in Melbourne, Australia, in conjunction with the *Asia Pacific Regional Internet Conference on Operational Technologies* (APRICOT).

As previously announced, the usual "rotation" of the location of APRICOT and APNIC conferences has been suspended since the start of the pandemic, but a decision to restart it will be considered by APNOG and APNIC in the coming months. When the conference rotation restarts, the first face-to-face APRICOT will be held in Manila, Philippines. More information about APNIC 54, including the venue, dates, online participation options, partner meetings and other details can be found here: **https://conference.apnic.net/54/**

Our Privacy Policy

The *General Data Protection Regulation* (GDPR) is a regulation for data protection and privacy for all individual citizens of the *European Union* (EU) and the *European Economic Area* (EEA). Its implementation in May 2018 led many organizations worldwide to post or update privacy statements regarding how they handle information collected in the course of business. Such statements tend to be long and include carefully crafted legal language. We realize that we may need to provide similar language on our website and in the printed edition, but until such a statement has been developed here is an explanation of how we use any information you have supplied relating to your subscription:

- The mailing list for *The Internet Protocol Journal* (IPJ) is entirely "opt in." We never have and never will use mailing lists from other organizations for any purpose.
- You may unsubscribe at any time using our online subscription system or by contacting us via e-mail. We will honor any request to remove your name and contact information from our database.
- We will use your contact information only to communicate with you about your subscription; for example, to inform you that a new issue is available, that your subscription needs to be renewed, or that your printed copy has been returned to us as undeliverable by the postal authorities.
- We will never use your contact information for any other purpose or provide the subscription list to any third party other than for the purpose of distributing IPJ by post or by electronic means.
- If you make a donation in support of the journal, your name will be listed on our website and in print unless you tell us otherwise.

Check your Subscription Details!

If you have a print subscription to this journal, you will find an expiration date printed on the back cover. For the last couple of years, we have "auto-renewed" your subscription, but now we ask you to log in to our subscription system and perform this simple task yourself. The subscription portal is here: **https://www.ipjsubscription.org/**

Thank You!

Publication of IPJ is made possible by organizations and individuals around the world dedicated to the design, growth, evolution, and operation of the global Internet and private networks built on the Internet Protocol. The following individuals have provided support to IPJ. You can join them by visiting **http://tinyurl.com/IPJ-donate**

Kjetil Aas
Fabrizio Accatino
Michael Achola
Martin Adkins
Melchior Aelmans
Christopher Affleck
Scott Aitken
Jacobus Akkerhuis
Antonio Cuñat Alario
Nicola Altan
Shane Amante
Marcelo do Amaral
Matteo D'Ambrosio
Selva Anandavel
Jens Andersson
Danish Ansari
Finn Arildsen
Tim Armstrong
Richard Artes
Michael Aschwanden
David Atkins
Jac Backus
Jaime Badua
Bent Bagger
Eric Baker
Santosh Balagopalan
William Baltas
David Bandinelli
Benjamin Barkin-Wilkins
Feras Batainah
Michael Bazarewsky
David Belson
Hidde Beumer
Pier Paolo Biagi
Tyson Blanchard
John Bigrow
Orvar Ari Bjarnason
Axel Boeger
Keith Bogart
Mirko Bonadei
Roberto Bonalumi
Lolke Boonstra
Julie Bottorff Photography
Gerry Boudreaux
L de Braal
Kevin Breit
Thomas Bridge
Ilia Bromberg
Václav Brožík
Christophe Brun
Gareth Bryan
Ron Buchalski
Paul Buchanan
Stefan Buckmann
Caner Budakoglu
Darrell Budic
BugWorks
Scott Burleigh
Chad Burnham
Jon Harald Bøvre
Olivier Cahagne
Antoine Camerlo
Tracy Camp
Ignacio Soto Campos
Fabio Caneparo
Roberto Canonico
David Cardwell
Richard Carrara
John Cavanaugh
Lj Cemeras
Dave Chapman
Stefanos Charchalakis
Greg Chisholm
David Chosrova
Marcin Cieslak
Lauris Cikovskis
Guido Coenders
Brad Clark
Narelle Clark
Horst Clausen
Joseph Connolly
Steve Corbató
Brian Courtney
Beth and Steve Crocker
Dave Crocker
Kevin Croes
John Curran
André Danthine
Morgan Davis
Jeff Day
Julien Dhallenne
Freek Dijkstra
Geert Van Dijk
David Dillow
Richard Dodsworth
Ernesto Doelling
Michael Dolan
Eugene Doroniuk
Karlheinz Dölger
Michael Dragone
Joshua Dreier
Lutz Drink
Aaron Dudek
Dmitriy Dudko
Andrew Dul
Joan Marc Riera Duocastella
Pedro Duque
Holger Durer
Mark Eanes
Andrew Edwards
Peter Robert Egli
George Ehlers
Peter Eisses
Torbjörn Eklöv
Y Ertur
ERNW GmbH
ESdatCo
Steve Esquivel
Jay Etchings
Mikhail Evstiounin
Bill Fenner
Paul Ferguson
Ricardo Ferreira
Kent Fichtner
Armin Fisslthaler
Michael Fiumano
The Flirble Organisation
Gary Ford
Jean-Pierre Forcioli
Susan Forney
Christopher Forsyth
Andrew Fox
Craig Fox
Fausto Franceschini
Valerie Fronczak
Tomislav Futivic
Laurence Gagliani
Edward Gallagher
Andrew Gallo
Chris Gamboni
Xosé Bravo Garcia
Osvaldo Gazzaniga
Kevin Gee
Greg Giessow
John Gilbert
Serge Van Ginderachter
Greg Goddard
Tiago Goncalves
Ron Goodheart
Octavio Alfageme Gorostiaga
Barry Greene
Jeffrey Greene
Richard Gregor
Martijn Groenleer
Geert Jan de Groot
Christopher Guemez
Gulf Coast Shots
Sheryll de Guzman
Rex Hale
Jason Hall
Darow Han
Handy Networks LLC
James Hamilton
Stephen Hanna
Martin Hannigan
John Hardin
David Harper
Edward Hauser
David Hauweele
Marilyn Hay
Headcrafts SRLS
Hidde van der Heide
Johan Helsingius
Robert Hinden
Asbjørn Højmark
Damien Holloway
Alain Van Hoof
Edward Hotard
Bill Huber
Hagen Hultzsch
Kauto Huopio
Kevin Iddles
Mika Ilvesmaki
Karsten Iwen
David Jaffe
Ashford Jaggernauth
Thomas Jalkanen
Martijn Jansen
Jozef Janitor
John Jarvis
Dennis Jennings
Edward Jennings
Aart Jochem
Nils Johansson
Brian Johnson
Curtis Johnson
Richard Johnson
Jim Johnston
Jonatan Jonasson
Daniel Jones
Gary Jones
Jerry Jones
Michael Jones
Amar Joshi
Javier Juan
David Jump
Anders Marius Jørgensen
Merike Kaeo
Andrew Kaiser
Christos Karayiannis
Daniel Karrenberg
David Kekar
Stuart Kendrick
Robert Kent
Jithin Kesavan
Jubal Kessler
Shan Ali Khan
Nabeel Khatri
Dae Young Kim
William W. H. Kimandu
John King
Russell Kirk
Gary Klesk
Anthony Klopp
Henry Kluge
Michael Kluk
Andrew Koch
Ia Kochiashvili
Carsten Koempe
Richard Koene
Alexader Kogan
Matthijs Koot
Antonin Kral
Robert Krejčí
Mathias Körber
John Kristoff
Terje Krogdahl
Bobby Krupczak
Murray Kucherawy
Warren Kumari
George Kuo
Dirk Kurfuerst
Darrell Lack
Andrew Lamb
Richard Lamb
Yan Landriault
Edwin Lang
Sig Lange
Markus Langenmair
Fred Langham
Tracy LaQuey Parker
Alex Latzko
Jose Antonio Lazaro Lazaro
Rick van Leeuwen
Simon Leinen

Robert Lewis
Christian Liberale
Martin Lillepuu
Roger Lindholm
Link Light Networks
Chris and Janet Lonvick
Sergio Loreti
Eric Louie
Adam Loveless
Josh Lowe
Guillermo a Loyola
Hannes Lubich
Dan Lynch
David MacDuffie
Sanya Madan
Miroslav Madić
Alexis Madriz
Carl Malamud
Jonathan Maldonado
Michael Malik
Tarmo Mamers
Yogesh Mangar
Bill Manning
Harold March
Vincent Marchand
Normando Marcolongo
Gabriel Marroquin
David Martin
Jim Martin
Ruben Tripiana Martin
Timothy Martin
Carles Mateu
Juan Jose Marin Martinez
Ioan Maxim
David Mazel
Miles McCredie
Brian McCullough
Joe McEachern
Alexander McKenzie
Jay McMaster
Mark Mc Nicholas
Olaf Mehlberg
Carsten Melberg
Kevin Menezes
Bart Jan Menkveld
Sean Mentzer
William Mills
David Millsom
Desiree Miloshevic
Joost van der Minnen
Thomas Mino
Rob Minshall
Wijnand Modderman-Lenstra
Mohammad Moghaddas
Roberto Montoya
Charles Monson
Andrea Montefusco
Fernando Montenegro
Joel Moore
John More
Maurizio Moroni
Brian Mort
Soenke Mumm
Tariq Mustafa
Stuart Nadin
Michel Nakhla
Mazdak Rajabi Nasab
Krishna Natarajan
Naveen Nathan
Darryl Newman
Thomas Nikolajsen
Paul Nikolich
Travis Northrup
Marijana Novakovic
David Oates
Ovidiu Obersterescu
Tim O'Brien
Mike O'Connor
Mike O'Dell
John O'Neill
Jim Oplotnik
Packet Consulting Limited
Carlos Astor Araujo Palmeira
Alexis Panagopoulos
Gaurav Panwar
Manuel Uruena Pascual
Ricardo Patara
Dipesh Patel
Alex Parkinson
Craig Partridge
Dan Paynter
Leif Eric Pedersen
Rui Sao Pedro
Juan Pena
Chris Perkins
Michael Petry
Alexander Peuchert
David Phelan
Derrell Piper
Rob Pirnie
Marc Vives Piza
Jorge Ivan Pincay Ponce
Victoria Poncini
Blahoslav Popela
Andrew Potter
Eduard Llull Pou
Tim Pozar
David Raistrick
Priyan R Rajeevan
Balaji Rajendran
Paul Rathbone
William Rawlings
Mujtiba Raza Rizvi
Bill Reid
Petr Rejhon
Robert Remenyi
Rodrigo Ribeiro
Glenn Ricart
Justin Richards
Rafael Riera
Mark Risinger
Fernando Robayo
Gregory Robinson
Ron Rockrohr
Carlos Rodrigues
Magnus Romedahl
Lex Van Roon
Marshall Rose
Alessandra Rosi
David Ross
William Ross
Boudhayan Roychowdhury
Carlos Rubio
Rainer Rudigier
Timo Ruiter
RustedMusic
Babak Saberi
George Sadowsky
Scott Sandefur
Sachin Sapkal
Arturas Satkovskis
PS Saunders
Richard Savoy
John Sayer
Phil Scarr
Gianpaolo Scassellati
Elizabeth Scheid
Jeroen Van Ingen Schenau
Carsten Scherb
Ernest Schirmer
Philip Schneck
Peter Schoo
Dan Schrenk
Richard Schultz
Timothy Schwab
Roger Schwartz
SeenThere
Scott Seifel
Yury Shefer
Yaron Sheffer
Doron Shikmoni
Tj Shumway
Jeffrey Sicuranza
Thorsten Sideboard
Greipur Sigurdsson
Fillipe Cajaiba da Silva
Andrew Simmons
Pradeep Singh
Henry Sinnreich
Geoff Sisson
John Sisson
Helge Skrivervik
Terry Slattery
Darren Sleeth
Richard Smit
Bob Smith
Courtney Smith
Eric Smith
Mark Smith
Tim Sneddon
Craig Snell
Job Snijders
Ronald Solano
Asit Som
Ignacio Soto Campos
Evandro Sousa
Peter Spekreijse
Thayumanavan Sridhar
Paul Stancik
Ralf Stempfer
Matthew Stenberg
Martin Štěpánek
Adrian Stevens
Clinton Stevens
John Streck
Martin Streule
David Strom
Colin Strutt
Viktor Sudakov
Edward-W. Suor
Vincent Surillo
Terence Charles Sweetser
T2Group
Roman Tarasov
David Theese
Douglas Thompson
Kerry Thompson
Lorin J Thompson
Fabrizio Tivano
Peter Tomsu Fine Art Photography
Joseph Toste
Rey Tucker
Sandro Tumini
Angelo Turetta
Michael Turzanski
Phil Tweedie
Steve Ulrich
Unitek Engineering AG
John Urbanek
Martin Urwaleck
Betsy Vanderpool
Surendran Vangadasalam
Ramnath Vasudha
Philip Venables
Buddy Venne
Alejandro Vennera
Luca Ventura
Scott Vermillion
Tom Vest
Peter Villemoes
Vista Global Coaching & Consulting
Dario Vitali
Jeffrey Wagner
Don Wahl
Michael L Wahrman
Laurence Walker
Randy Watts
Andrew Webster
Tim Weil
Jd Wegner
Westmoreland Engineering Inc.
Rick Wesson
Peter Whimp
Russ White
Jurrien Wijlhuizen
Derick Winkworth
Pindar Wong
Makarand Yerawadekar
Phillip Yialeloglou
Janko Zavernik
Bernd Zeimetz
Muhammad Ziad Ziayuddin
Tom Zingale
Jose Zumalave
Romeo Zwart
廖明沂.

Follow us on Twitter and Facebook

@protocoljournal

https://www.facebook.com/newipj

Call for Papers

The *Internet Protocol Journal* (IPJ) is a quarterly technical publication containing tutorial articles ("What is...?") as well as implementation/operation articles ("How to..."). The journal provides articles about all aspects of Internet technology. IPJ is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. In addition to feature-length articles, IPJ contains technical updates, book reviews, announcements, opinion columns, and letters to the Editor. Topics include but are not limited to:

- Access and infrastructure technologies such as: Wi-Fi, Gigabit Ethernet, SONET, xDSL, cable, fiber optics, satellite, and mobile wireless.
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance.
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, troubleshooting, and mapping.
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, cloud computing, and quality of service.
- Application and end-user issues such as: E-mail, Web authoring, server technologies and systems, electronic commerce, and application management.
- Legal, policy, regulatory and governance topics such as: copyright, content control, content liability, settlement charges, resource allocation, and trademark disputes in the context of internetworking.

IPJ will pay a stipend of US$1000 for published, feature-length articles. For further information regarding article submissions, please contact Ole J. Jacobsen, Editor and Publisher. Ole can be reached at `ole@protocoljournal.org` or `olejacobsen@me.com`

Supporters and Sponsors

Supporters

Internet Society

CISCO

Diamond Sponsors

Your logo here!

Ruby Sponsors

ICANN

Sapphire Sponsors

APNIC FOUNDATION

Emerald Sponsors

Akamai

amsix

APRICOT Asia Pacific Regional Internet Conference on Operational Technologies

COMCAST

DE-CIX

donuts inc

EQUINIX

Google

JPRS

JUNIPER NETWORKS

lacnic

LinkedIn

linx

netskope

NSRC Network Startup Resource Center

NTT Communications

RIPE NCC RIPE NETWORK COORDINATION CENTRE

TEAM CYMRU

VERISIGN

WIDE PROJECT

Corporate Subscriptions

AFRINIC The Internet Numbers Registry for Africa

DNS-OARC

IWL

IPXO

ISC Internet Systems Consortium

Limelight NETWORKS

qa cafe

For more information about sponsorship, please contact `sponsor@protocoljournal.org`

The Internet Protocol Journal
Link Fulfillment
7650 Marathon Dr., Suite E
Livermore, CA 94550

CHANGE SERVICE REQUESTED

The Internet Protocol Journal is published quarterly and supported by the Internet Society and other organizations and individuals around the world dedicated to the design, growth, evolution, and operation of the global Internet and private networks built on the Internet Protocol.

Email: ipj@protocoljournal.org
Web: www.protocoljournal.org