Securing BGP - A Literature Survey

Geoff Huston, Mattia Rossi, Grenville Armitage
Swinburne University of Technology
Melbourne, Australia
gih@apnic.net, {mrossi,garmitage}@swin.edu.au

***Abstract*—The Border Gateway Protocol (BGP) is the Internet's inter-domain routing protocol. One of the major concerns related to BGP is its lack of effective security measures, and as a result the routing infrastructure of the Internet is vulnerable to various forms of attack. This paper examines the Internet's routing architecture and the design of BGP in particular, and surveys the work to date on securing BGP. To date no proposal has been seen as offering a combination of adequate security functions, suitable performance overheads and deployable support infrastructure. Some open questions on the next steps in the study of BGP security are posed.**

***Index Terms*—BGP security, Inter-domain routing security, routing, BGP, Computer Network Protocols.**

I. Introduction

The Internet is a decentralised collection of interconnected component networks. These networks are composed of end hosts (who originate and/or receive IP packets, and are identified by IP addresses) and active forwarding elements (routers) whose role is to pass IP packets through the network. The routing system is responsible for propagating the relative location of addresses to each routing element, so that routers can make consistent and optimal routing decisions in order to pass a packet from its source to its destination. Routing protocols are used to perform this information propagation.

The Internet's current routing system is divided into a two-level hierarchy. At one level is intra-domain routing, used by the set of autonomous routing systems operating within each component network. At the other level is a single inter-domain routing system that maintains the inter-autonomous system connectivity information that straddles these component networks. A single inter-domain routing protocol, the Border Gateway Protocol (BGP) [1], has provided inter-domain routing services for the Internet's disparate component networks since the late 1980's [2]. Given the central role of routing in the operation of the Internet, BGP is one of the critical protocols that provide security and stability to the Internet [3].

BGP's underlying distributed distance vector computations rely heavily on informal trust models associated with information propagation to produce reliable and correct results. It can be likened to a hearsay network — information is flooded across a network as a series of point-to-point exchanges, with the information being incrementally modified each time it is exchanged between BGP speakers. The design of BGP was undertaken in the relatively homogeneous and mutually trusting environment of the early Internet. Consequently, its approach to information exchange was not primarily designed for robustness in the face of various forms of negotiated trust or overt hostility on the part of some routing actors.

Hostile actors are a fact of life in today's Internet. The Internet is a significant public communications utility operated by a disparate collection of service providers, together with a relatively unclear distinction between the roles of service providers and customers. It is quite reasonable to characterise today's Internet environment as one where both customers and service providers[1] are potentially hostile actors, and where trust must be explicitly negotiated rather than assumed by default. This environment is no longer consistent with the inter-domain trust framework assumed by BGP, and BGP's operational assumptions relating to trust are entirely inappropriate today.

Today's inter-domain routing environment remains a major area of vulnerability [3]. BGP's mutual trust model involves no explicit presentation of credentials, no propagation of instruments of authority, nor any reliable means of verifying the authenticity of the information being propagated through the routing system. Hostile actors can attack the network by exploiting this trust model in inter-domain routing to their own ends. An attacker can easily transform routing information in ways that are extremely difficult for any third party to detect. For example, false routing information may be injected, valid routing information removed or information altered to cause traffic redirection [6], [7], [8], [9]. This approach can be used to prevent the correct operation of applications, to conduct fraudulent activities, to disrupt the operation of part (or even all) of the network in various ways. The consequences range from relatively inconsequential (minor degradation of application performance due to sub-optimal forwarding paths) through to catastrophic (major disruption to connectivity and comprehensive loss of any form of cohesive Internet) [10], [11], [4].

Current research on BGP is predominately focused on two major themes — scaling, and resistance to subversion of integrity [12].

[1]There are instances of cyber attacks where the suspected attacker is a state or nation rather than an individual. e.g. the YouTube "incident" [4] was an example of a service provider's actions resulting in the corruption of the global routing system. It has also been determined, that providers can be hostile towards other providers, by attempting carriage theft as described in [5]

Accepted 25 Jan 2010 for 2nd issue, 2011, of IEEE Communications Surveys and Tutorials (final version copyright IEEE)

Scaling the routing system is an operational concern regarding the ability of the inter-domain routing system to cope with a larger domain of discourse. This encompasses increasing numbers of objects to be managed, using an increasingly expressive language to represent route objects and path attributes, increased frequency with which objects advertise changes in their state, more paths across the inter-domain environment, and more frequent dynamic changes to the attributes of inter-domain paths [13], [14].

Resisting subversion of integrity requires that a BGP speaker (an entity participating in the exchange of inter-domain routing information) has:

- Sufficient information at hand to verify the authenticity and completeness of the information being provided via the inter-domain routing system, and
- The ability to generate authoritative information such that others may verify the authenticity of information that this speaker is passing into the inter-domain routing system.

A key question is whether further information can be added into the inter-domain routing environment such that attempts to pervert, remove or withhold routing information may be readily and reliably detected. Any proposed scheme(s) must also be evaluated for their impact on the scaling properties of BGP [15].

It is assumed in this paper that the reader is familiar with routing concepts and basic security concepts including hash algorithms, public key cryptography, and public key infrastructure. An introduction to routing concepts and extended explanations of various routing protocols can be found in [16]. An introduction to security concepts can be found in [17].

This paper surveys the current research in BGP routing security. In section II we begin by examining the architecture of the Internet's routing system. Section III provides a detailed summary of BGP itself, and section IV discussing the primary threats against BGP. Section V provides a wide-ranging review of the major approaches to providing security in inter-domain routing and the various refinements to these approaches. Section VI reflects on some open questions in BGP security and the paper concludes in section VII.

II. The Architecture of IP Routing

The Internet has been designed using a modular or decoupled framework, where inter-dependencies between distinct functional components are minimised and inter-module interfaces are clearly defined. The concepts of Internet Protocol (IP) [18] addresses, packet forwarding, routing and routing protocols are treated as being mutually distinct and having well-defined modes of interaction and dependencies. Mutually consistent and coherent interaction between these components results in the Internet's service of end-to-end packet delivery.

An IP address indicates identity, rather than location, of an addressed host. The address provides no indication of how to direct a packet through the network in order to reach the addressed host. An address distribution system may impose some locational structure on addresses (which may further result in some numerically adjacent address values being topologically adjacent) but such a property is not statically encoded into the address itself. It is the role of the *routing system* (or *control plane*) to propagate the dynamic binding of addresses to locations, and the role of the *forwarding system* (or *data plane*) to use these bindings in order to deliver addressed packets to the correct locations [16], [19].

IP forwarding is a local autonomous action within each IP routing element. Packets passing between interfaces of a routing element are *forwarded*, with the choice of outgoing interface guided by local information contained in a *forwarding table*. Forwarding tables encode rules indicating the next routing element (the next hop) to which a packet should be sent based on the address to which the packet is ultimately destined. End-to-end packet forwarding across the Internet relies on mutually consistent population of forwarding tables that are maintained in every routing element.

The IP routing system's primary role is to propagate address location information so that routers across the Internet may properly populate (and update) their local forwarding tables. The routing system is a distributed collection of processes that participate in self-learning information exchanges through the operation of *routing protocols*. Self-learning routing systems operate on a peer-to-peer level rather than through a structured hierarchy of information dissemination, and can be characterised informally as a set of point-to-point information exchanges of the form: "You tell me everything you think I should know, and I'll tell you everything I think you should know." Each routing protocol's objective is to support a distributed computation that produces consistent best path outcomes in the forwarding tables of all IP routing elements.

The Internet's routing system is a structured two-level hierarchy [20]. At the bottom level we have routing elements grouped into *Autonomous Systems* (ASes) [21]. Each AS represents a collection of routing elements sharing a common administrative context. Internally, an AS is an interconnected network with a coherent routing structure and a single consistent path metric framework that allows for a consistent interpretation of path comparison. Routing within ASes is known as *intra-domain routing*, and handled by Interior Gateway Protocols (IGPs). While the Routing Information Protocol (RIP) [22] was widely deployed in the 1980's, it is more common to see the Open Shortest Path First (OSPF) protocol [23] and the Intermediate System to Intermediate System (ISIS) protocol [24] deployed as IGPs today. Inter-AS connections form the second level of the Internet's routing hierarchy. The Border Gateway Protocol (BGP) [1] is the sole inter-AS (or *inter-domain*) routing protocol operating in today's Internet.

BGP is a path vector form of distance vector routing protocol. Routers who run BGP are known as *BGP speakers*. Each BGP speaker communicates with other BGP speakers, termed variously BGP *peers* or *neighbours*. Like other distance vector routing protocols, a BGP speaker receives *route objects* from all BGP routing neighbours. Each route object is a logical information block that contains an *address prefix* that describes a contiguous set of address values and a set of *attributes* that provide additional routing information that has been associated with the address prefix. One of the critical attributes for the operation of the BGP protocol is the attribute of an *AS Path*.

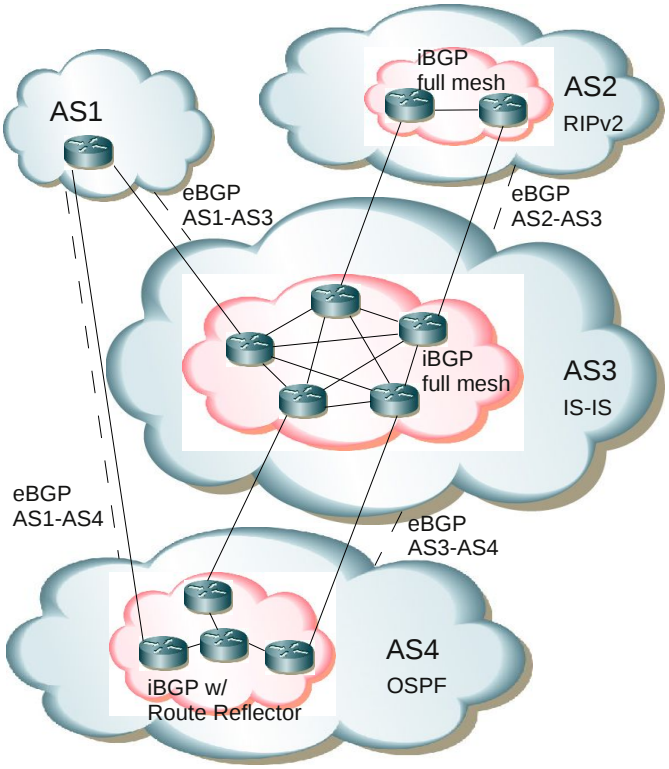

Fig. 1. **An example routing topology**: BGP is the sole routing protocol used for inter-AS peering (eBGP). A single AS can be multihomed within the same eBGP session (AS2-AS3 and AS3-AS4) or use multiple eBGP sessions for the purpose (AS1-AS3/AS1-AS4 and AS4-AS3/AS4-AS1). Each AS deploys its own IGP (RIPv2, IS-IS, OSPF) to route between internal subnets (not depicted) and uses iBGP to connect the BGP boundary routers of the AS internally in a full mesh or using Route Reflectors (see Section III-D)

This attribute is an ordered enumeration of AS values that form the path of ASes from the AS that originated this route object (*origin AS*) to the current AS. The number of elements in the path is the *AS Path length*. Where a BGP speaker is presented with multiple paths to the same address prefix from a number of peers, the BGP speaker selects the "best" path to use by minimising a distance metric across all the possible paths. The distance metric used by BGP speakers is the AS Path length. This BGP-selected route object is used to populate the local forwarding table. The BGP speaker then assembles a new route object by taking the locally selected route object, attaching locally significant attributes and adding its own AS value to the route object's AS path vector. This route object is then announced to all BGP peers.

Each AS may have more than one *exterior* connection to one or more other ASes [25]. Such inter-AS BGP connections are termed *eBGP* sessions. Within an AS BGP speakers exchange route objects between each other, also using BGP. The variant of BGP behaviour that supports this intra-AS routing exchange is termed an *iBGP* session[2]. An example of the various modes of peering sessions between BGP speakers is shown in Figure 1.

[2]iBGP should not be mistaken as a separate IGP. It is still BGP and does not obsolete the need for IGPs as discussed in section III-D

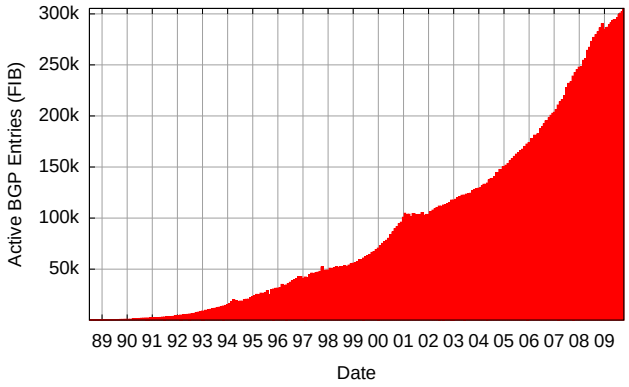

Fig. 2. The growth of the Internet's Inter-domain Routing system [30]

III. The Design and Operation of BGP

BGP has undergone a number of refinements over its operational life. BGP was originally described in RFC1105, in June 1989 [26], allowing the Internet's inter-domain architecture to move on from a constrained architecture of a "core" and attached "stub" domains into a framework of peer routing domains without any central "core". BGP-2 was described in RFC1163, in June 1990 [27], and BGP-3 was described in RFC1267 in October 1991 [28]. The current version, BGP-4, was first deployed within the Internet in 1993. The RFC describing this protocol, RFC1771 [29], was published in March, 1995, and subsequently refined with the publication of RFC4271 in January 2006 [1]. The protocol has been stable for some years now. Across the deployment lifetime of BGP-4 the Internet has grown from an average of 20000 distinct routing entries in 1993 to some 300000 routing entries in 2009 [30]. The growth of the size of the Internet's routing table over time is shown in Figure 2.

A. BGP and TCP

BGP is not a link-level topology maintenance protocol. This has allowed BGP to use the IP transport protocol TCP [31] as a reliable transport protocol to support the protocol's transactions across a BGP peer session. Essentially, BGP assumes the existence of a functional IP forwarding environment at the link level.

TCP manages reliable message delivery and flow control between the BGP peers, and allows BGP to operate across end-to-end logical connections whether they reside on the same sub-net, the same LAN, or across an Internet. There is no requirement for BGP speakers to be connected on a common media connection, and the choice of TCP allows this flexibility of connectivity by requiring only that a BGP peering session is supported by an IP network.

The TCP stream is divided into messages using BGP-defined markers, where each message is between 19 and 4096 octets in length [19]. The use of a reliable transport platform implies that BGP need not explicitly confirm receipt of a protocol message. This removes much of the protocol overhead seen in other routing protocols that sit directly on top of a media level connection. There are no message identifiers, no message number initiation protocol, no explicit acknowledgement of messages nor any provision to manage lost, re-ordered

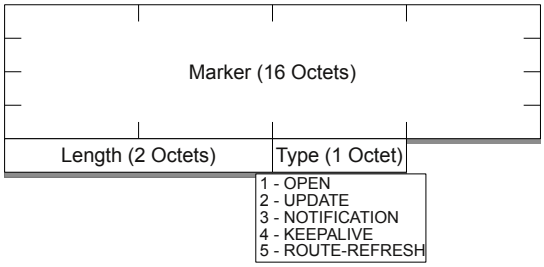

Fig. 3. BGP Common Header Message Format

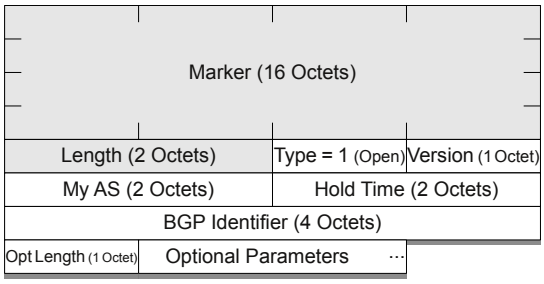

Fig. 4. BGP OPEN Message Format

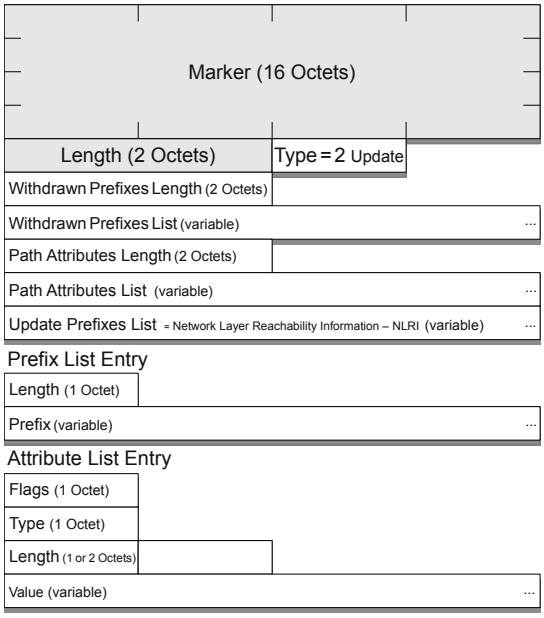

Fig. 5. BGP UPDATE message format

or duplicated messages. These functions are managed by TCP and are therefore unnecessary for BGP to also support. The use of a reliable transport protocol also obviates the need for BGP to periodically refresh the routing state by re-flooding the entire routing information set between BGP speakers. After the initial exchange of routing information, a pair of BGP routers exchange only incremental changes to routing information.

B. BGP Messages

As TCP is a stream protocol rather than a record-oriented protocol, BGP uses record marking within the TCP stream to delineate logical protocol units, or messages with a 16-byte marker as the BGP message delimiter. The marker is followed by a 2-byte length and a 1-byte type field, making the minimum BGP message size 19 bytes. The repertoire of defined messages are: an OPEN message to start a BGP session, an UPDATE message to exchange reachability information, a NOTIFICATION message, which is used to convey a reason code prior to termination of the BGP session, KEEPALIVE messages, used to confirm the continued availability of the BGP peer, and ROUTE-REFRESH request messages to request a resend of the routing information. The Common BGP header message format is indicated in Figure 3.

BGP uses an explicit OPEN message to commence a BGP peering session. This message exchange confirms the identity of the BGP speakers and includes the option for a capability negotiation to understand what optional or extended capabilities are supported by each BGP speaker. A session is active only when both BGP speakers have sent their OPEN messages and neither has rejected the others offered capabilities through a NOTIFICATION response. The BGP OPEN message format is indicated in Figure 4.

Once the session is active, BGP operates via the exchange of UPDATE messages. Each UPDATE message contains a set of address prefixes that are unreachable (withdrawals), followed by a set of common route object attributes, and a set of address prefixes that share this set of attributes (announcements). The withdrawn prefixes are those prefixes where the local BGP speaker sees no reachability, and now wants to withdraw a previous advertisement of reachability. No routing attributes are associated with these withdrawn prefixes. The announced prefixes are those prefixes where the local BGP instance has an updated view of the reachability of a prefix that was previously withdrawn or unannounced, or has an updated view of the routing attributes of the locally selected "best" route for a prefix. BGP may group multiple prefixes together in a single UPDATE message, but can only do so if all the updated prefix share a common set of attributes. The withdrawn prefix set or the announced prefix set may be empty, but not both, within an UPDATE message. The layout of the BGP UPDATE message is shown in Figure 5.

The NOTIFICATION message is used to convey the nature of an error condition prior to the closing of the underlying TCP session.

A relatively recent addition to BGP was proposed in 2000 [32]. This is the ROUTE-REFRESH BGP message, which requests the BGP peer to re-send its set of advertised route objects to this BGP speaker.

C. AS Path Attribute

BGP binds together the concept of network address blocks and autonomous systems into a path vector-based routing technology. Every route object represented within a BGP-4 route database contains an address prefix and an associated path vector of AS values. BGP does not indicate the precise path a packet should following within an AS, nor does it maintain a complete map of the topology of the Internet at a link-by-link level. BGP uses a level of abstraction which views the Internet as a set of per-AS routing domains, and the role of BGP is to maintain a routing map of the network at this AS level, associating every reachable address prefix with an AS transit path from the current location to the address prefix's originating AS.

One of the most important route object attributes in BGP is the AS Path attribute. As address prefix reachability information traverses the Internet in the form of individual route objects in BGP, this BGP routing information is augmented by the list of ASes that have been traversed thus far, forming the AS Path attribute. Each BGP speaker adds its own AS value to the route object's AS Path attribute when passing the route object through an eBGP session. This AS Path attribute allows straightforward suppression of the looping of routing information, using the simple algorithm that a local AS will reject any forwarded route object that already contains its own AS in the AS Path attribute. Also the length of the AS Path vector forms the BGP route metric. A local BGP system, when attempting to select one of a number of potential route objects that refer to the same address prefix, will, in the absence of any local policy directive, prefer the route object with the shortest AS Path length.

In addition to undertaking the role of path metric and loop detector, the AS Path attribute serves as a versatile mechanism for policy-based routing, where a local AS can alter the default preferences for route selection based on local policy settings coupled with pattern matching rules to be performed on the AS Path.

D. iBGP and eBGP

BGP is intended to provide a mechanism for one AS to exchange routes with another, and BGP sessions that connect two different ASes are termed *eBGP* sessions. In a simple stub AS configuration, there is a single exterior boundary router that supports all the AS's eBGP sessions. The interior routing protocol typically directs a default route to this boundary point.

However, if the external connections for an AS are terminated in separate boundary routers, and the AS has a internal requirement to pass routes learnt from one eBGP session to the other, the destination routes and associated path attributes must be passed between the two boundary routers. Using a redistribution of the BGP routes into an IGP to perform this transfer will cause the learnt eBGP path attributes to be discarded within the IGP. Instead, an internal BGP peering session between the two boundary routers is configured, allowing a full transfer of all BGP route attributes between the two BGP speakers in the same AS. Such an internal BGP session is termed an *iBGP* session.

The AS path vector construct is inadequate to detect routing loops that may arise across the iBGP sessions within the AS, so there is a simple restriction on iBGP that addresses this potential for loop creation: routes learnt via an iBGP peer session are not advertised to other iBGP peers. The corollary of this constraint is that every iBGP router must form an iBGP peering session with every other iBGP router within the AS. That is, all BGP speakers within an AS must directly peer with all other BGP speakers within the AS as shown in Figure 1.

This requirement for an $O(N^2)$ peering mesh leads to one of the major scaling issues with autonomous systems and BGP. This mesh of BGP peering sessions can generate a message update load that potentially exceeds the processing capabilities of the component routers. The most effective method to mitigate this iBGP load is to introduce BGP Route Reflectors [33], which dilute the strict requirement for a complete mesh of peering sessions by explicitly permitting iBGP route propagation. The distinction here is that a Route Reflector performs iBGP route redistribution, using a new BGP attribute, the ORIGINATOR_ID, to perform loop detection in the iBGP domain.

E. BGP Route Selection Process and Routing Policies

A BGP speaker may receive two or more announcements for the same address prefix from different peers. The "best" announcement is selected as the locally used announcement, and this announcement is the one that is announced to its BGP peers. BGP defines an ordered sequence of comparisons to determine which route object is selected by the local BGP speaker:

- Select the route object with the highest value for LOCAL-PREF attribute value
- Select the route object shortest AS_PATH attribute length
- Select the lowest MULTI_EXIT_DISCRIMINATOR attribute value
- Select the minimum IGP cost to the NEXT_HOP address given in the route object
- Select eBGP over iBGP-learnt routes
- If iBGP select the lowest BGP Identifier value.

Although a network administrator's usually employs routing policies depending on his needs [34], [35], within the generic BGP route selection process the highest priority selection rule is that a route for a more specific address prefix is to be preferred over that of a covering prefix

IV. The BGP Threat Model

One approach to providing a taxonomy for threats in routing in general, and BGP in particular, is to view a BGP peer session as a conversation between two BGP speakers and pose a number of questions relating to this conversation. These questions are:

- How do we talk?
  The manner in which the BGP session between the BGP speakers is secured such that the conversation is not altered, disrupted or hijacked and is protected from unauthorised eavesdropping.
- Whom am I talking to?
  Verifying the identity of the other party and verifying that they are authorised to speak for the routing entity that they purport to represent.
- What are you saying?
  Verifying the authenticity and completeness of the routing information being passed in the BGP session.
- Should I believe you?
  Verifying that the routing information actually represents the state of the forwarding system.
- How recent is your information and is it still valid?
  Verifying for how long routing information is valid.

Each of these security questions can be further deconstructed to a set of specific objectives, as well as recognising a set of specific threats.

A. Securing the BGP session

A BGP session between two routers is assumed to have some level of integrity at the session transport level. BGP assumes that the messages sent by one party are precisely the same messages as received by the other party, and assumes that the messages have not been altered, reordered, have spurious messaged added into the stream or have messages removed from the conversation stream in any way.

As with any long-held TCP session, the BGP peer session is vulnerable to eavesdropping, session reset, session capture, message alternation and denial of service attacks via conventional TCP attack vectors.

The threat at the session level is that a third party may attempt to break into the TCP session, and alter the BGP message flow. One form of threat is by injection, where a third party eavesdrops on the conversation and injects spurious messages into the BGP session. Eavesdropping allows the attacker to have knowledge of the TCP sequence numbers, thereby making injection a trivial task. Even if the attacker is not able to eavesdrop the BGP session it is still possible to attempt to guess the current sequence number. While this is often impractical in the case of injecting data into the session, if all that is to be injected is a TCP Reset, then the sequence number guess only has to sit within the current TCP window in order to be recognised as a valid reset TCP message [36].

Another form of threat is by active intermediation where a third party sits on the wire between the two BGP speakers and intercepts all traffic in both directions. In this case the third party has complete control of the BGP message stream and can perform any form of message alteration. A variation of this form of threat is by session hijacking, where the third party intrudes upon an active BGP session and injects its own traffic into the message stream that allows the third party to take over the session and masquerade as one of the parties to the BGP session.

As timing is important in the overall performance of BGP another form of attack at the session level is to delay messages. While the content of the messages are unaltered, the implicit timing signals within the message stream are altered by this form of intervention, potentially causing the local BGP speaker to behave differently and fall out of sync with its routing peers. For example, it is possible to exercise various forms of local suppression of routes by altering the timing of propagation of BGP messages.

Another form of attack is a replay attack, where older BGP messages are replayed into a hijacked TCP session. One form of this replay attack could be to replay a pair of messages that withdraw and then announce the same address prefix. Route Flap Damping (RFD) [37], [38] is a widespread defensive BGP configuration that monitors the frequency of BGP updates for a given prefix from each peer, and if the update rate exceeds a locally set threshold the peer's advertisement of this prefix will be locally suppressed for a damping interval. The replay of updates could be used to trigger an RFD response in the remote BGP speaker [39]. If a route is fully dampened through RFD, updates for this prefix will not be advertised by the BGP speaker for a damping interval (commonly 60 minutes), possibly causing a route to be disrupted within that time frame. Another form of replay attack is to replay a route for a previously withdrawn prefix, possibly in conjunction with some form of *prefix hijack*[3] attack.

Another form of threat is by withholding traffic. BGP uses keepalive timers to determine remote end "liveness". By intercepting and withholding all messages for the hold down timer interval, a third party can force the BGP session to be terminated and reset. This causes the entire route set to be re-advertised upon session resumption so that repeated attacks of this form can be an effective form of denial of service for BGP.

It is also possible to undertake a saturation attack on a BGP speaker by sending it a rapid stream of invalid TCP packets. In this case the processing capability of the BGP speaker is put under pressure, and the objective of the attack is to overwhelm the BGP speaker and cause the BGP session to fail and be reset. This is particularly problematic if the BGP session uses MD5 or IPSEC as session protection protocols, as the cryptographic function overhead also applies to the injected packets[4], increasing the processing overhead on these spurious injected packets.

The underlying aspect of the BGP protocol is that BGP itself has no enforced minimum level of message protection. BGP messages are, by default, placed into the TCP stream without encryption or additional message wrapping of message sequencing. Any threat that is applicable to long held TCP sessions applies to this default mode of BGP operation.

B. Verifying BGP Identity

BGP sessions commence by passing the local AS to the remote end of the session in the sent OPEN message, and receiving the remote end's AS in the received OPEN message.

BGP itself does not verify these asserted AS identities, and it is theoretically possible for a remote party to masquerade as another party and assert an identity in BGP that cannot be directly verified by the other party, or by any third party that subsequently receives this routing information. Most BGP implementations provide a level of protection against this threat by applying a constraint that the local BGP speaker will only initiate a peer session with a configured remote IP address, and reject all other TCP connection attempts, and furthermore will not complete the BGP OPEN message exchange if the AS in the OPEN message does not match the AS number associated with the remote end IP address in the configuration.

This approach places a heavy reliance on the out-of-band process of BGP configuration, and if an attacker can compromise or take control over BGP equipment connected to the Internet or use social engineering to convince a network administrator to configure incorrect information into the BGP equipment then it is possible to masquerade as a different party in BGP and potentially inject incorrect information into the routing system [40].

[3]Further explained in Section IV-C and Figure 6

[4]Each packet's cryptographic checksum has to be calculated in any case, in order to determine that the packet is bogus.

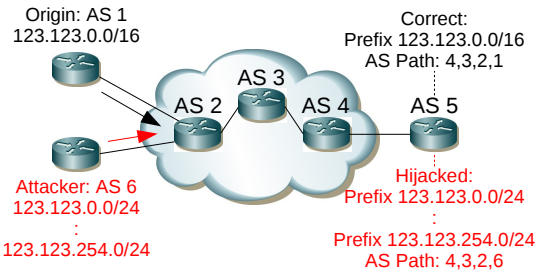

Fig. 6. **Prefix Hijacking**: An attacker (AS 6) could inject 255 /24 prefixes in the routing system in order to hijack the /16 prefix of AS 1. AS 5 will usually prefer routes belonging to the more specific prefix

The real question here is: "Are you really who you claim to be?" Here is it necessary for the BGP speaker to be able to confirm the validity of the peer's claim to be speaking for an AS.

C. Verifying BGP Information

The objective here is that of verifying the authenticity and completeness of the routing information being passed in the BGP session.

The intention of BGP is that a local BGP speaker provides to all its BGP peers a complete feed of its locally selected route objects. Once a session is opened with a remote BGP speaker the local BGP instance believes everything it is told without further qualification, and the threat is that a BGP peer can deliberately feed false information to the local BGP instance, which BGP itself will be unable to detect as false.

This could be in the form of suppression of routing information, or in the form of alteration of the route object that is being passed, or the invention of spurious route objects. The BGP speaker could be asserting that an AS Path is genuine when it reflects an artificial path, or that it has the authority to originate an advertisement for a prefix when, in fact, no such authority exists. This is also known as *prefix hijacking* and is depicted in Figure 6.

A BGP speaker may preserve all the attributes of a route object, but alter the prefix set to be the equivalent collection of more specific prefixes. The deliberate alteration of routing information can cause the local BGP instance to make an incorrect choice of a local best path and also cause the local BGP instance to propagate this incorrect information to its neighbours. Not only could the BGP speaker be passing incorrect attributes for an address prefix in order to bias the local route selection process, it could also be providing incorrect information regarding the prefix itself. The prefix that is the subject of the route object could be a prefix that has never been allocated and should not be legitimately routed, or the prefix is an aggregate address prefix that spans both allocated and unallocated address space. *Prefix hijacking* is a major threat to the integrity of the BGP routing infrastructure [41], and is now the subject of scrutiny from the public media [9].

The fundamental weakness here is that BGP provides no explicit means of verifying the authenticity of the address prefixes that are listed in a BGP UPDATE message, nor in the authenticity of the attributes of the prefix, including the origination information and the AS Path vector. The threat here is that by deliberately altering this information the local BGP speaker can be induced to make incorrect route selection decisions and thereby make incorrect forwarding decisions for IP traffic.

A known common problem illustrative of exploiting this vulnerability is operational misconfiguration [42], which could result in propagating more specific routes, and other forms of route leakage or withholding that may impact on the routing decisions made by other BGP speakers. This form of verification of intentionality by a remote BGP speaker is far more challenging — while these forms of security mechanisms are intended to verify that the received information matches the original information that was passed into the routing system, they are incapable of verifying that such information is consistent with the true intent of the original author of the information.

D. Verifying Forwarding Paths

The overall intention of the BGP protocol is to distribute the current binding of address to location such that individual routers can make accurate judgements about how to populate their local forwarding tables and hence make optimal local decisions for each packet that passes along the shortest path to its ultimate destination.

BGP does not provide any ability for a local BGP speaker to validate that the route advertisements it receives from a BGP peer actually represent the current state of the peer's forwarding system.

The threat model here is that a bad actor in the routing system may make a different forwarding decision to that being advertised in the routing system. This can represent a subversion of local policies, theft of carriage capacity, deliberate denial of service, the potential to eavesdrop on a conversation or to support the interception and alteration of application level transactions. Even a completely secured control plane does not avert such vulnerabilities [5].

E. The Consequences of Attacks on the Routing System

The ability to alter the routing system provides a broad array of potential consequences [6]. The consequences fall into a number of broad categories, which are briefly described here.

1) *The ability to eavesdrop.*
   The forwarding system can be altered so as to pass all traffic to a class of destination addresses via a certain path. This allows the attacker to attempt to pass all such traffic through an eavesdropping location prior to conventional delivery. In such a case the parties my not be aware that an eavesdropping attack is taking place.
2) *Denial of service.*
   The simplest form of a denial of service is where traffic to an address prefix is passed to a point where it is then discarded. Routing loops also are a form of denial of service, where not only will the traffic to a destination address prefix never reach its intended destination, but the traffic will be held in the loop for the life of the packet TTL field. For sufficiently short loops the potential exists for the loop to act as a link load amplifier, where the traffic on the loop is several

times the traffic load being addressed to the affected destination address prefix.

3) *The potential to masquerade.*
Subversion of routing allows sites to masquerade as other sites; The routing system will misdirect the traffic to the masquerading site. The consequences of such an attack can vary from the specific, where a particular site is targeted, to the more generic where authoritative DNS servers are the subject of the masquerading attack and the DNS responses are believed as being authentic. In this case if the masquerading occurs at the level of the root of the DNS hierarchy incorrect information can be provided to any query, allowing for the attack to then be extended to any site.
4) *The ability to steal addresses and obscure identity.*
Routing an unallocated address is subtly different to routing an already allocated address. Here the consequence is not displacement of traffic forwarding to incorrect locations in the network, but the assertion of the existence of addresses and forwarding paths to those addresses that should not exist in the network in the first place. The consequence is the ability to use addresses on the network that have no allocation registration information associated with them, allowing the originator of the routing attack some degree of ease to mount an anonymous attack at the application level. Such forms of attack have been observed to be associated with SPAM and botnet controllers where anonymity of the attack coordinator is desired [43].

V. Securing BGP

The vulnerabilities of BGP arise from four fundamental weaknesses in the BGP and the inter-domain routing environment [6]. These are:

- No mechanism to protect the integrity, freshness and source authenticity of BGP messages.
- No mechanism to verify the authenticity of an address prefix and an AS origination of this prefix in the routing system.
- No mechanism to verify the authenticity of the attributes of a BGP UPDATE message.
- No mechanism to verify that the local cache RIB information is consistent to the current state of the forwarding table.

The other pragmatic observation about BGP security is that it appears that by far the most straightforward form of attack is to obtain control and configuration access to a deployed router[5] and use this compromised platform as the base for launching attacks on the routing system [45], [46]. In the face of such an encompassing attack on the control instruments of the routing system, BGP session-level security needs to be placed in some perspective [47]. It is not possible to prevent routers from attempting to generate false information as long as routers themselves are in a position to be compromised. The consequent vulnerability on the routing system, as distinct from a narrower view of BGP, is that there is no mechanism that limits the extent to which a misbehaving routing element can make inaccurate claims about reachability in the routing system.

[5]It appears that some routers accessible via standard interfaces like SSH and Telnet deployed *default passwords* [44] and still might do so today

A. The Security Toolset

The available tools for securing BGP start at the level of the BGP peer session, and encompass the tools that are used to protect the TCP session and the two ends of the TCP session. The TCP protection mechanisms include the generalized TTL security mechanism [48], [49], which is intended to limit the effective radius of potential attack on the session to hosts that lie on or within the worst case hop count radius between the two BGP speakers, and host-level defences against TCP SYN attacks [50].

There are two tools to protect the BGP TCP session from external disruption that use the approach of cryptographic protection of the BGP connection. These are the use of IPSEC at the IP level [51] and the TCP MD5 signature option at the TCP session level [52], [53]. While the MD5 signature option has some potential weaknesses when compared with IPSEC [6], MD5 is considered preferable to no form of TCP protection at all. The decision between IPSEC and MD5 includes consideration of key rollover capability, where no standard key rollover mechanism exists in MD5 [54], and the cryptographic processing load, where the load of IPSEC processing is significantly higher than MD5 processing. Requiring cryptographic validation also exposes a potential denial of service threat where a BGP speaker is flooded with invalid messages each of which must be cryptographically processed before being detected as invalid and discarded [55].

In addition to message integrity protection provided by transparent session level protection mechanisms, the tools to provide protection of the integrity of BGP messages relate to the use of digital signatures [56] to provide a set of credentials that allow relying parties to verify the correctness of the information carried as the message payload in BGP. The reason for the use of digital signatures as opposed to an integrity check using some form of shared secret is due to the observation that the number and identities of all eventual recipients of the information is not known in advance, and non-repudiation is desirable [57]. It is also the case that the verification of the contents of a message is not only a test of whether the message has been altered in any way during its transit between BGP speakers, but a test of whether the message represents correct origination information and correct operation of the processing of the message during the process of message propagation. This implies a need to establish a means of verification of information where the author of any security credentials relating to origination and propagation are not necessarily known to the relying party that is attempting the validation operation. This typically invokes a form of validation that relies upon third party trust, where the relying party is attempting to build a testable chain of trust between its trust anchor and the party or action that is the subject of the verification operation.

This implies that the use of digital signatures is accompanied by the requirement to verify such digital signatures. This,

in turn, involves some form of mechanism to authenticate the public key that is associated with an address prefix or an AS number, and validate this association. A common approach to this is via X.509 certificates and a Public Key Infrastructure [58], [59], where verification of a digital signature entails a test of the authenticity and current validity of the associated certificate that describes the public key of the address or AS number holder in the context of a structured set of signed relationships between certificate issuers and subjects [60].

B. Security Requirements

The primary requirements for securing BGP are securing the data payload of the BGP protocol and securing the semantics of the payload.

The security requirements for the data payload are such that the data received by a BGP speaker can be cryptographically verified to have been sent by the BGP peer, that the data is not a replay of previously transmitted data, and that no data has been removed from the transmission [55]. There is no strict requirement for encryption of the payload, as the routing information being exchanged is not intrinsically confidential to the two parties involved.

The security requirements for the semantics of the payload concern specifically some selected fields (transitive attributes) of the BGP UPDATE message. The BGP speaker must be able to verify that the advertised prefix is valid, and that the originating AS has been duly authorised by the legitimate right-of-use holder for that prefix. The BGP speaker should also be able to validate that the AS Path in the UPDATE represents a valid inter-AS transit path through the network in terms of inter-AS topology and AS transit policies, and that the prefix reachability information has been propagated along the reverse inter-AS Path [55]. It is noted that route withdrawals and non-transitive announcement attributes are local, and thus do not need to be transitively protected in a similar fashion to route origination and the AS Path attribute of announcements. Local attributes can be adequately protected by BGP peer session protection.

The associated requirements for a secure inter-domain routing system is that the additional use of security credentials and verification of routing information should not alter the temporal properties of the BGP protocol, and that authentication of the security credentials should occur in the same time frame as the BGP message processing operation.

It is also a requirement that piecemeal incremental deployment should be feasible [61], [62], [63]. A secure operational mode should be a capability negotiation with each BGP peer, with the ability to support backward compatibility with those peers who do no recognise such a capability. It seems to be a good idea to start deployment of BGP security on the most connected nodes and incrementally deploy it towards least connected nodes [64]. Additionally it puts the question on how a party which uses security credentials deals with information arriving from a peer which does not use any security credentials. Having no security credentials does not necessary mean that the information is wrong. Importantly, in these piecemeal deployment scenarios there should be some incremental benefit of piecemeal deployment to those actors who choose to supply such security credentials and to those who chose to validate routing information using these credentials. It is necessary that deployment of security within the whole BGP routing system is made appealing for the large variety of participants [65]

A routing system, secure or otherwise, should never make route selections that include routing loops. It is preferred that in a fully secured environment a secure routing system would be able to converge on best paths that are either identical to or no worse than an unsecured BGP speaker would select, assuming that such paths can be validated in a secure environment. In an environment of partial adoption of secure routing systems it is recognised that a BGP speaker may use local preference settings that prefer sub-optimal paths that have preferred security credentials over unsecured paths.

The trust model of routing appears to involve two forms of trust. The first is a trust environment related to the public network and the legitimacy of use of a public address and the legitimacy of use of a public AS number. It is necessary to be able to verify that a particular party has the right to use these number resources in a public context. The closest fit in the form of a trust model for verification of this assertion of right of use is a public authority that can provide authoritative information on the distribution of these numbers. This approach leads to a rooted hierarchy model of trust, where the trust anchor is this public authority. The second form is a trust environment in private contexts, where the use of an address or AS number is bounded by a specific context of use, and the trust in an assertion of a right of use is one made in the context of this bounded environment. In this environment there is no clear ability to use public authorities as a trust anchor and other means of trust that may involve reputation or web of trust concepts may be appropriate. A general security approach to BGP should be able to encompass that diversity of deployment environments and the corresponding diversity of authority models.

C. Approaches to Securing BGP

The major contribution to this area of study is the secure BGP (sBGP) proposal [66], which is the most complete contribution to date. However, the assumptions relating to the environment in which sBGP must operate, particularly in terms the performance capability of routing systems appear to be beyond the capabilities of routers used in today's Internet [67]. A refinement of this approach, soBGP [68], is an attempt to strike a pragmatic balance between the security processing overhead and the capabilities of deployed routing systems and security infrastructure, where the requirements for AS Path verification are relaxed and the nature of the related Public Key Infrastructure (PKI) is altered to remove the requirement for a strict hierarchical address PKI that precisely mirrors the address distribution framework. Another refinement of the sBGP model, psBGP [69], represents a similar effort at crafting a compromise between security and deployed capability through the crafting of a trust rating for assertions based on assessment of confidence in corroborating material. Another proposal, IRV [70], takes a different direction

and extends the existing model of Internet Route Registries (general repositories of routing information, connectivity and routing policies), into per-AS route registries. It replaces the augmentation of the BGP protocol with security credentials that are a common aspect of the previously noted proposals, with a form of query-based retrieval of credentials as an out-of-band function structured as an adjunct to the operation of BGP. The motivation with this approach is that any effort to change the installed based with new software and potentially more capable hardware is not an attractive proposition and one approach is to make the security function an incremental overlay on the existing routing infrastructure.

There is also a considerable body of work that refines each of these approaches, based either on refinements in the cryptographic functions that attempt to provide comparable security but with a reduced processing load, or in refinements to the application of the security function based on observed BGP behaviour, or even to modify the operation of the BGP protocol to reduce the security overhead by deliberately reducing the BGP message load. There has been done additionally a high amount of research in providing at least some level of security in the current BGP routing systems, with solutions which aim less to be complete security solutions but to be easily deployable instead.

These various approaches to securing BGP will now be reviewed in more detail.

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

1) Securing the operation of BGP: BGP is a long held TCP session and the same approaches to securing any TCP session [71] can be used in the context of a BGP session. These approaches fall into two categories: those that simply attempt to protect the TCP session from disruption via injection of spurious traffic, and those that also attempt to protect the TCP session from eavesdropping and alteration by encrypting the payload.

The *Generalized TTL Security Mechanism* (GSTM) was originally described in [48] and updated in [49] and is based on the observation that the overall majority of BGP peering sessions are established between routers that are directly connected. The technique is to configure each BGP IP packet to be sent with a TTL field value in the IP header of 255, and for the BGP receiver to discard all packets with an inbound TTL of less than a set threshold value. For a direct connection the inbound TTL value should be 255, so all inbound TCP packets with within this session with a TTL of 254 or less can be discarded by the receiver. The motivation for this approach is that spoofing of the TTL field in an IP header is challenging for an unassisted remote attacker, and this TTL packet filter is a lightweight defensive measure to protect the BGP session from efforts to intrude upon the session from remote attacks. Figure 7 illustrates the idea.

This GTSM approach can be used for multi-hop BGP peer sessions as well as directly connected BGP sessions, but it is not as 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 [49].

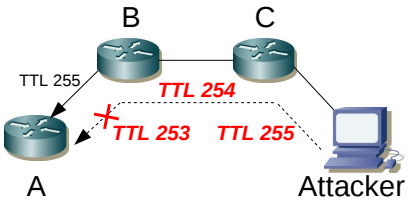

Fig. 7. **Generalized TTL Security Mechanism (GSTM)**: A distant attacker's injected packets will always arrive with a TTL of 254 or less, and thus be discarded.

A more robust approach to protecting the TCP session is through the use of cryptographic protection of the TCP session. While these approaches can be highly resilient to intrusion attempts they also expose the BGP speaker to potential denial of service attacks if the processing load of the cryptographic functions to detect bogus packets is sufficiently high.

The *TCP MD5 Signature Option* [53] uses message authentication codes, which are a class of cryptographic hash algorithms applied to messages of arbitrary length that produce a "message digest" of the message, intended to protect the integrity of a message. The desired property of a message digest is that it is infeasible to generate two messages that have the same message digest value, or to generate a new message that has a particular message digest value. The MD5 algorithm [52] is intended for digital signature applications where a message digest is generated over the combination of a message and a secret shared key value. The message and the digest value can be transmitted openly, and the receiver can use a local copy of the secret key and apply the message digest algorithm to the combination of the received message and the key. If the digest value matches the received value then the receiver can be assured that the message has not been altered in transit, and that the message was generated by a party who also has knowledge of the key. The TCP MD5 Signature option is a TCP extension where each TCP segment contains a TCP option that contains the 128 bit MD5 digest of the combination of the TCP pseudo header, the TCP segment payload excluding TCP options, and a connection-specific key. This establishes a cryptographically secure signature of the packet. Without knowing the key, it is very challenging to construct a TCP segment with a valid signature, nor is it readily possible to alter the packet without causing the signature to be invalidated. The receiver calculates the MD5 digest across the received data, using a locally held copy of the key, and rejects the segment if the digest value fails to match that provided in the packet. In the context of BGP the TCP session is resistant to various forms of intrusion attack unless the attacker has knowledge of the shared secret key value.

The TCP MD5 specification does not specify how the shared key is passed between the two BGP speakers, nor how the key value can be changed during the session. This latter problem is significant, in that continued use of a key weakens its integrity, and it is conventionally advised that keys should be changed every 90 days or so in this type of use context [72]. 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. This has lead to some further work in supporting MD5 key rollover in active sessions. The simplest approach is to continue with the use of an out-of-band key management mechanism and allow a number of keys to be considered active a any point in time. As there is no key identifier field in the TCP MD5 Signature Option the receiver simply has to attempt to use every key to determine if a segment passes the MD5 check, starting with the one that succeeded for the previous segment [73]. This approach increases processing load for each bogus packet, as all keys need to be checked before rejecting a packet as failing the MD5 signature, exposing denial of service vulnerabilities.

Rather than have the receiver undertake a key search by repeated application of the MAC algorithm looking for a key match, it would be more efficient for the receiver if the sender were to provide a key identifier in the manner of an authentication option. This would allow the receiver to identify the key that was used to create the MAC for a given message without performing an expensive key search. This is considered to be useful with both manual and automatic key management [74].

However it is not only the extended use lifetime of keys that weakens the MD5 approach. The MD5 algorithm itself has been the subject of some recent concern regarding its robustness [75], [76], [77], [78]. Other algorithm choices for a stronger Message Authentication Code (MAC) include HMAC-SHA1 [79], UMAC [80] and HMAC-MD5-96 [81], [82].

These two concerns have prompted a proposal to extend the TCP MD5 options with both a key identifier and an algorithm identifier, allowing the sender the ability to specify which key to use as well as the message digest algorithm [83]. A further enhancement to this approach uses automated key generation and selection. The shared secret in this case is a Key Encrypting Key, and this key is only used to encrypt the MAC key that is passed to the other party in a TCP Enhanced Authentication Option [84]. A somewhat different approach, the TCP Authentication Option [85], proposes to use a Message Authentication Field in the place of the MD5 message digest, where the final bit of the length field of the option determines whether a key ID has been appended to the Message Authentication Code or not. The message digest algorithm in this case is specified as HMAC-MD5-96, although other algorithms can be used if configured in advance. This approach relies on a similar form of out-of-band provisioning as the original MD5 approach, where each end of the conversation has to configure a TCP Security Association Database in advance of the use of this mechanism. This database contains a description of the supported TCP connections, the key set, the MAC algorithm and MAC length.

IPSEC is a suite of protocols that operate at the IP level of the protocol stack that secures all communications between two hosts [51], [86]. The functionality of IPSEC includes methods for protection of IP packet headers, methods for protection and encryption of IP payloads and key management services that allow key rollover during long held sessions. This is an implementation of public/private key cryptography and can ensure the confidentiality and integrity of all IP messages passed between two hosts. IPSEC can be used to secure BGP sessions, and it provides greater levels of assurance than can be derived from MD5 [87].

However, IPSEC is not widely used in the public Internet for the purpose of securing BGP sessions, 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 [88]. The perceived problem with IPSEC is that the processing load to detect bogus packets is considerably higher with IPSEC than MD5 [88]. This exposes a denial of service attack where a stream of bogus IPSEC packets directed at a BGP speaker may be capable of exercising the processor into a fully saturated mode of operation, causing other concurrent router functions to be degraded.

2) Securing the integrity of BGP Data: One of the earlier recognised works that addressed routing security was the 1988 study on *Byzantine Robustness* [89] by Perlman. In the event of failure or malicious behaviour on the part of one or more entities in the system, all correctly operating entities should reach a mutually consistent decision regarding the validity of each message in finite time. This study was in the area of link-state protocol design, and the work described a protocol that satisfied the properties for Byzantine Robustness. It categorised route validation in 3 approaches:

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

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

The work by Smith and Garcia-Luna-Aceves [90], [91] attempts to address session security by modifying the BGP protocol. This work proposed the protection of BGP control messages using message encryption at the BGP level, with session keys exchanged at BGP session establishment time. It also proposed the addition of a message sequence number to protect against replay attacks and message removal. This approach also proposed a predecessor path attribute that indicated the AS prior to the destination AS for the current route, and proposed digitally signing all fixed fields in the UPDATE message. The predecessor attribute is used to construct a means of validation of the AS Path attribute. These proposed changes to the BGP protocol required comprehensive adoption and deployment in order to be effective. This approach was similar to the earlier IDRP work [92], which eschewed the use of TCP and included a reliable flow controlled transport into the IDRP protocol, also including a number of message integrity protection options.

A contemporary proposal to the Smith and Garcia-Lunes-Aceves proposal for securing BGP was based on leaving the BGP protocol unchanged, but augmenting the BGP data flow with access to credential information that allowed a BGP speaker to confirm the authenticity of origination information in BGP UPDATE messages by validating the binding of address prefixes to originating ASes [93]. This work proposed using the DNS as the distribution mechanism for origination information, where a BGP speaker could perform a DNS query to validate the origination information provided in a BGP UPDATE message. The proposed mechanism, Inter-domain Route Validation (IRV)[6] [70], operated in a similar way to the reverse DNS pointer space where an IP address is mapped to a domain name. In this case an address prefix was mapped by a DNS Autonomous System Resource Record (RR) to an AS number and a prefix length. The proposal called for a new DNS zone in the "in-addr.arpa" zone, namely "bgp.in-addr.arpa", populated by NS RRs, CNAME RRs and AS RRs. A relying party could query the DNS with a BGP address prefix and the AS RR response would indicate the originating AS that was associated with that prefix and the authorised prefix length.

One issue with this approach was that it attempted to resolve the issue of implicit trust in BGP to provide reliable and authentic information relating to origination by instead placing implicit trust in the DNS to provide authentic information relating to DNS queries in the corresponding responses. This proposal required the use of DNSSEC [94] to allow for these DNS RRs to be digitally signed, so that the DNS response could be validated as authentic. The DNS delegation hierarchy would needed to be precisely aligned to the address allocation framework, so that the zone administrator of each of these origination authentication zones was in fact the duly delegated holder of the addresses, and this alignment should, preferably, be capable of being validated by third parties. Meeting these requirements would create a digital signature hierarchy embedded in the DNS that would be aligned to the address allocation framework. The consequent observation is that whether this digital signature hierarchy is placed into the DNS, via a DNSSEC framework [95], or placed into a framework of X.509 certificates and an associated PKI (as proposed by IRV) is, at one level, an isomorphic transform of the same information. The issue of the choice of DNS or X.509 certificates is then an issue of the performance requirements of these systems. Also, this approach relates only to verification of route origination, while the verification of the AS Path is not addressed in this framework. The identification of invalid routing information in the partial adoption case of this approach is unclear. When a DNS query receives no response at all, it is unclear whether the routing information is incorrect, or whether the DNS information is incomplete in terms of the appropriate interpretation of the outcome by the relying party.

Limitations to the DNS structure itself are probably the reason why the approach has never been deployed. An entire address block might be sub-allocated to a DNS sub-level

[6]Further explained in Section V-C2d

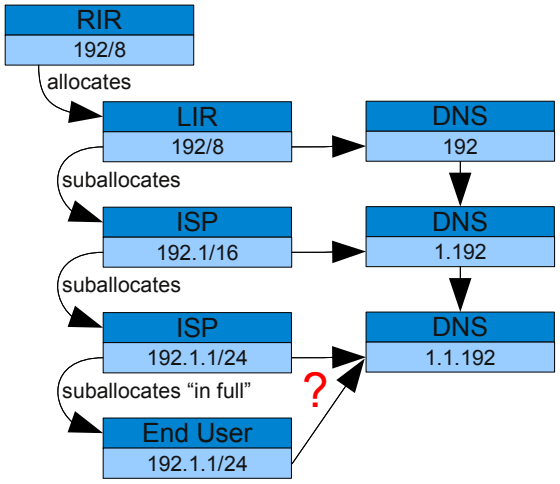

Fig. 8. **DNS+BGP**: Using DNS for origin authentication is not feasible because of DNS protocol limitations. In this example, who is responsible for 192.1.1/24 ?

(allocated "in full"), causing a prefix to refer to the DNS entry that did the "allocation in full" instead of the legitimate owner of the prefix. This shortcoming shown in Figure 8 is caused by the DNS protocol definition, and could only be solved with modifications to DNS.

Subsequent studies have concentrated on securing the semantics of BGP messages, and, in particular, approaches to allow a BGP speaker a means of validating that the origination information is authentic and that the accumulated routing information, as represented in the AS PATH, is an authentic record of the transit path of the routing information through the inter-AS network.

a) sBGP: Secure BGP (sBGP) [66], represents one of the major contributions to the study of inter-domain routing security, and offers a relatively complete approach to securing the BGP protocol by placing digital signatures over the address and AS Path information contained in routing advertisements and defining an associated PKI for validation of these signatures.

sBGP defines the "correct" operation of a BGP speaker in terms of a set of constraints placed on individual protocol messages, including ensuring that all protocol UPDATE messages have not been altered in transit between the BGP peers, that the UPDATE messages were sent by the indicated peer, the UPDATE messages contain more recent information than has been previously sent to this BGP speaker from the peer, the UPDATE was intended to be received by this BGP speaker, and that the peer is authorised to advertise information on behalf of the peer Autonomous System. In addition, for every prefix and its originating AS, the prefix must be a validly allocated prefix, and the prefix's "right-of-use" holder must have authorised the advertisement of the prefix and must have authorised the originating AS to advertise the prefix.

The basic security framework proposed in sBGP is that of digital signatures, X.509 certificates and PKIs to enable BGP speakers to verify the identities and authorisation of other BGP speakers, AS administrators and address prefix owners. The verification framework for sBGP requires a PKI for address allocation, where every address assignment is reflected in an issued certificate [96]. This PKI provides a

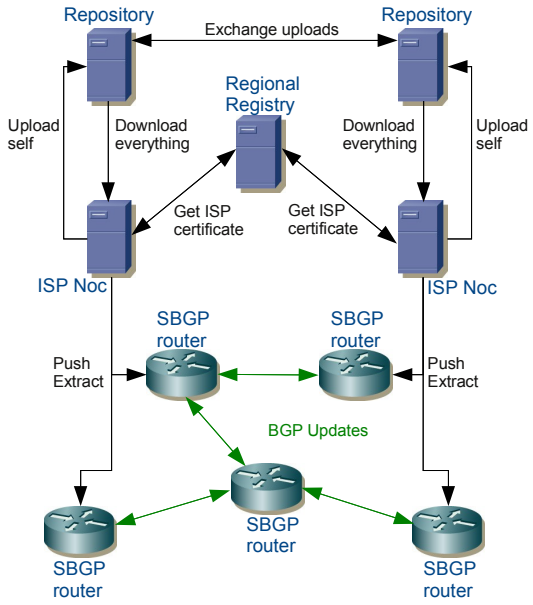

Fig. 9. **sBGP**: Certificates for each ISP are issued by the regional registries. The ISPs exchange public keys through special repositories. The keys are pushed to the sBGP routers which validate the BGP UPDATE messages

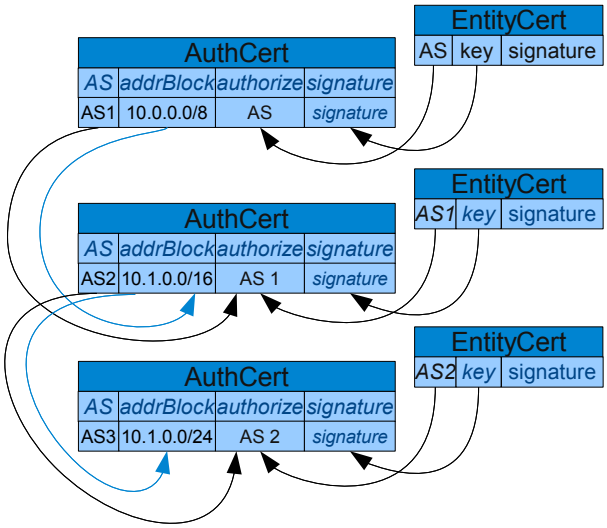

Fig. 10. **soBGP**: An EntityCert is bound to an AS and signs an AuthCert which then binds to an Address.

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

sBGP also proposes the use of attestations. An "address attestation" is produced by an address holder, and authorises a nominated AS to advertise itself as the origin AS for a particular address prefix. A "route attestation" is produced by an AS holder and attests that a BGP speaker is an authorised member of that AS and that it has received a specified route.

The address and AS PKIs, together with these attestations, allow a BGP speaker to verify the origination of a route advertisement and verify that the AS path as specified in the BGP UPDATE is the path taken by the routing UPDATE message via the sequence of nested route attestations.

Inter-operation and information exchange between sBGP elements is shown in Figure 9.

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

There is a number of significant issues that 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 [67], [97], [98].

b) soBGP: Secure Origin BGP (soBGP) [68] is a response to some of the significant issues that have been 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 [99].

The validation questions posed by soBGP also includes the notion of an explicit authorisation from the address holder to the originating AS to advertise the prefix into the routing system. The AS path validation is quite different from sBGP, however, in that soBGP attempts 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 uses the concept of an EntityCert to bind an AS to a public key. soBGP avoids the use of a hierarchical PKI that mirrors the AS number distribution framework and nominates the use of a web of trust, or a reputation mechanism, as means of validation of these certificates. 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. Figure 10 illustrates the interactions between EntityCert and AuthCert. The explicit avoidance of reliance on the established AS and address distribution framework and any form of associated PKI as the derivation of a trust hierarchy may have been a pragmatic consideration in the design of this approach, but it leaves open the issue of how to establish trust anchors for validation of these signed objects. This is a rather significant deficiency in the validation framework of soBGP.

Instead of sBGP's route attestations, soBGP uses the concept of an ASPolicyCert as the foundation for constructing the data for testing the feasibility of a given AS Path. An ASPolicyCert contains a list of the AS's local peer ASes, signed by the AS's private key. An AS peering is considered valid if both ASes list each other in their respective ASPolicyCerts. Figure 11 depicts a possible soBGP peering network.

The overall approach proposed in soBGP represents a different set of design trade-offs to sBGP, where the amount of validated material in a BGP UPDATE message is reduced. This

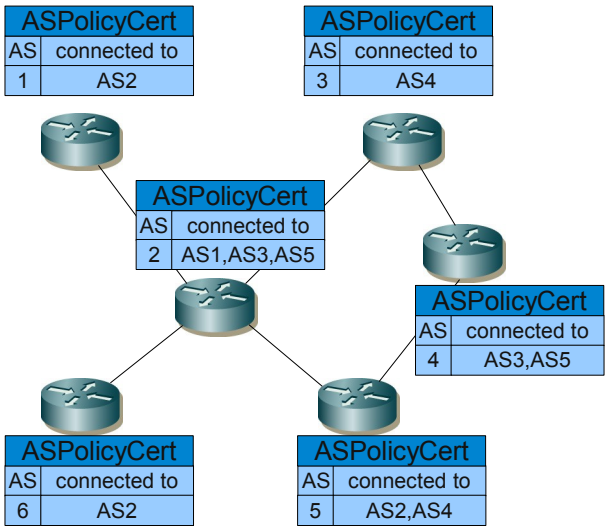

Fig. 11. **soBGP**: The PolicyCert is a self signed certificate containing routing policies. An UPDATE message originating at AS4 would necessarily take the Path {AS4,AS5,AS2,AS1} instead of {AS4,AS3,AS2,AS1} as the connection between AS2 and AS3 would not be regarded as valid.

can 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. The avoidance of a hierarchical PKI for the validation of AuthCerts and EntityCerts could be considered a weakness in this approach, as the derivation of authority to speak about addresses is very unclear in this model.

c) psBGP: Pretty Secure BGP (psBGP) [69] puts forward the proposition that the proposals relating to the authentication of the use of an address in a routing context must either rely on the use of signed attestations that need to be validated in the context of a PKI, or rely on the authenticity of information contained in Internet Routing Registries. The weakness of routing registries is that the commonly used access controls to the registry are insufficient to validate the accuracy or the current authenticity of the information that is represented as being contained in a route registry object. The information may have been accurate at the time the information was entered into the registry, but this may no longer be the case at the time the information is accessed by a relying party. The psBGP approach is also motivated by the proponent's opinion that a PKI could not be constructed in a deterministic manner because of the indeterminate nature of some forms of address allocations. This leads 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 rejects the notion of a hierarchical PKI that can be used to validate assertions about addresses and their use.

Interestingly, although psBGP rejects the notion of a hierarchical address PKI, psBGP assumes the existence of a centralised trust model for AS numbers and the existence of a hierarchical PKI that allows public keys to be associated with AS numbers in a manner that can be validated in the context of this PKI. This exposes a basic inconsistency in the assumptions that lie behind psBGP, namely that a hierarchical PKI for ASes aligned to the AS distribution framework is assumed to be feasible, but a comparable PKI for addresses is 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 is necessary or even appropriate.

psBGP uses a rating mechanism similar to that used by PGP [100], but in this case the rating is used for prefix origination. An AS asserts the prefixes it originates and also may 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 allows 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 (illustrated in Figure 11.

One of the critical differences between psBGP and soBGP is the explicit inclusion of the "strict" AS Path validation test, namely that it is 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. psBGP allows 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 is completely verifiable.

The essential approach of psBGP is the use of a reputation scheme in place of an hierarchical address PKI, but the value of this contribution is based on accepting the underlying premise that a hierarchical PKI for addresses is infeasible. It is also noted that the basis of accepting inter-AS ratings in order to construct a local trust value is based on accepting the validity of an AS trust rating, which, in turn, is predicated upon the integrity of the AS hierarchical PKI.

psBGP appears to be needlessly complex and bears much of the characteristics of making a particular solution fit the problem, rather than attempting to craft a solution within the bounds of the problem space. The use of inter-AS cross certification with prefix assertion lists introduces considerable complexity in both the treatment of confidence in the assertions and in the resulting assessment of the reliability of the verification of the outcome. psBGP does not consider the alternate case where the trust model relating to addresses is based on a hierarchical PKI that mirrors the address distribution framework. In such a case the calculation of confidence levels would be largely unnecessary.

The major contribution of psBGP relates to the case of partial deployment of a security solution in relation to AS Path validation, where the calculation of a confidence rating in the face of partial security information may be of some utility.

d) Inter-domain Route Validation (IRV): The approaches to securing the semantics of BGP described so far all entail changes to the operation of BGP itself, and operate most effec-

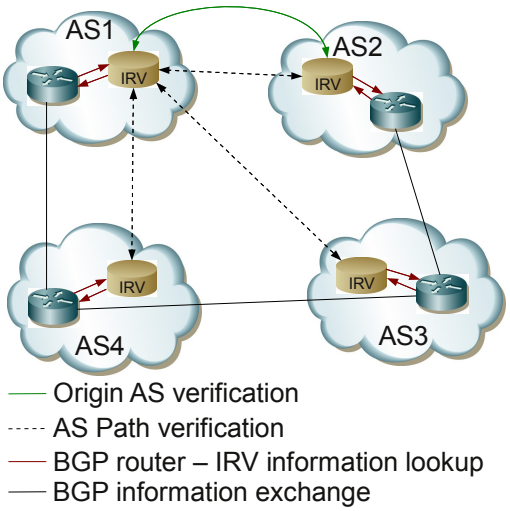

Fig. 12. **Inter-domain Route Validation (IRV)**: IRV adds an additional control plane to the BGP control plane. Each AS has to implement an IRV server exchanging information about Origin and AS Path authenticity. Routers which receive BGP updates can query the IRV server of the own AS for validity of the routes.

tively in an environment of universal deployment. In practical terms this is an unlikely scenario, and the current experience with the uptake of a revised version of BGP that supports 32-bit AS number values suggests that the public Internet has considerable inertia and is very resistant to adopting changes to BGP [101]. In such a system as large as the public Internet, long term piecemeal deployment is a more likely scenario.

The approach proposed with IRV [70] is not to modify the BGP protocol in any way, but to define a companion information distribution protocol. The intent here is to attempt to provide legacy compatibility and incremental deployment capability. The IRV approach replaces the concept of simultaneously feeding both routing information and associated credentials in BGP with the concept of moving the provision of credentials into a query response framework where the receiver of a route object can query the originating AS as to the authenticity of a received route object, or request additional information relating to the object in a similar fashion to the information contained in an Internet Routing Registry (IRR) [102]. As Figure 12 shows, 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) [103], [104] structure already used by the Internet Route Registries (IRRs) [105], 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 AS's IRV server in the AS path, confirming that an advertisement was received from the previous AS in the AS Path, and that an advertisement has been sent to the next AS in the AS path. This approach is midway between the strict AS Path test of sBGP that validates that the UPDATE message was passed along the AS sequence described in the AS Path, and the soBGP AS Path feasibility that validates that there is a set of AS peer connections that correspond to the AS sequence. Here the validation test is that each AS in the sequence is currently advertising this prefix to the next AS in sequence.

This IRV architecture has a number of issues that are not completely specified, including IRV discovery, IRV query redirection, authentication of queries and responses, selective responses, transport layer protection and imposed overheads.

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

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

A similar approach within the IRR framework that integrates the concept of an address and AS PKI could make provision for signed responses in a way that allows the IRR client to authenticate that the response is accurate, current, and contains information that has been digitally signed by the AS or prefix holder. In such a model of publication the relying party is able to validate the authenticity of the IRR object independently of the manner in which the object was published or the manner in which it has been retrieved [106].

e) Chained Hash Functions: Symmetric cryptographic techniques such as message authentication codes (MACs) or cryptographic hash functions, have been measured to be 3 to 4 orders of magnitude faster than asymmetric cryptographic functions for digital signatures [107]. As the cost of the asymmetric cryptographic functions in authentication of AS Path information is seen as being a prohibitive factor for the

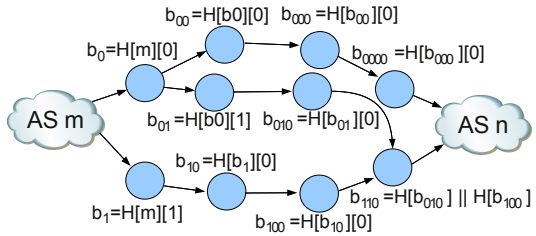

Fig. 13. **Chained Hash Functions**: Each transversed AS adds it's own hash on top of the received hash, creating a hash chain.

deployment prospects of sBGP, there has been some interest in evaluating approaches that substitute symmetric cryptographic processing in parts of the sBGP security architecture.

This observation about the relative performance factors of symmetric and asymmetric cryptographic functions can be used to reduce the processing load of applying cryptographic operations to a sequence of data objects when an ordered relationship between the objects can be determined in advance. One way hash chains are the result of repeated iteration of a hash function, starting with an initial seed value, and validation of an object within a one way hash chain relies on knowledge of the earlier value in the chain. The approach of tree-authenticated hash values relies on constructing a reverse hierarchical sequence of hashes over successive pairs of values, and then distributing the root hash value. Given a key set of size N, any key can then be authenticated using $log_2 N$ hash functions applied to a vector of $log_2 N$ values, which bounds the authentication workload [107].

These techniques can be used as a cumulative authentication mechanism to authenticate the list of ASes on the path in a BGP UPDATE, preventing removal or reordering of the ASes in the list. The mechanism uses only a single authenticator for the AS Path, and then uses pair-wise hashes for each predecessor-successor AS pair, and hash tree authentication of these pairs to authenticate the AS path [107]. Figure 13 illustrates the idea.

This approach of substituting a combination of symmetric cryptographic operations and explicit relationships between objects can also be applied to the address origination function using the hierarchical relationships that are part of the address distribution framework. One such approach to origin authentication [108] has analysed the semantics, dynamic behaviour, design and costs of origin authentication in inter-domain routing. The semantics of address delegation are formalised and various cryptographic structures for asserting address ownership and delegation are explored, with attention to cryptographic proof structures. The address delegation structure is observed to be static and dense, which makes for an effective cryptographic proof structure using largely static relationships between objects. This approach approximates the delegation hierarchy by extracting the nested announcements made within the protocol, and found evidence that address delegations were very stable over time. This property makes the associated address ownership assertions ideally suited to a class of proof structures based on Merkle hash trees [109]. A simulation of this approach showed that on-line real time origin authentication was possible using this construction with a limited form of caching, an outcome which was previously thought to have been too computationally expensive to be feasible [107].

f) SPV: 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 [110] as a further extension to hash chains and trees..

The SPV study identified the following classes of path attacks: "forgery" where false paths are associated with routes in order to influence local route selection decisions, "modification" where the path is altered in order to hide the UPDATE from a target AS or in order to influence local route selection decisions, "denial of service" where the attack attempts to overwhelm the intended victim's resources, and "worm-holing" where colluding adversaries assert false AS-to-AS links. The first two classes are attacks via BGP, whereas the second two could be more accurately classified as attacks on the routing system itself through multi-party collusion.

SPV takes the approach of tree-authenticated hash values and applies this specifically to AS Path validation as an alternative to the nested digital signature structure proposed as the AS Path validation mechanism of sBGP. The paper claims significantly improved processor performance using this technique, based on difference in computational complexity for asymmetric cryptography from symmetric cryptography as used in hash functions [111], [112].

This proposal falls into the category of proposals that calls for changes to the operation of the BGP protocol. In this case the significant change is the requirement that all routes must be re-advertised to peers within a fixed time interval. This is the weakest part of the approach in terms of performance evaluation, as much of the leverage in terms of scaling BGP, is based on the use of a reliable transport protocol for BGP messages which, in turn, obviated any need for a BGP 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. Raghvan et al. [113] also argue that SPV does not sufficiently protect against route forgery and eavesdropping or collusion attacks.

g) Signature Amortisation and Aggregate Signatures: Another approach is intended to amortise the cost of signature validation over many messages [114]. The technique signs a subset of the connected topology over which an UPDATE flows and placing a topology description as a vector in an equivalent of an AS connectivity attestation which is flooded to all relying parties. The AS Path signing can then be generalized such that the same vector is reproduced in the signed

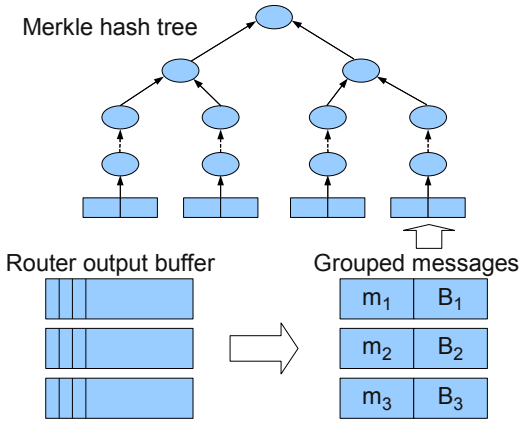

Fig. 14. **Signature Amortisation and Aggregate Signatures**:Signatures are applied to a group of UPDATE messages of the router output queue, using a Merkle hash tree.

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 [115] combines the time-efficient approach of signature amortisation with space-efficient techniques of aggregate signatures to propose a set of constructions for aggregated path authentication that improve on sBGP's requirements for processing throughput and memory space. Aggregate signatures apply to a collection of UPDATE messages that are to be sent to a peer. Instead of signing each UPDATE separately, the UPDATE messages are hashed into a Merkle hash tree [109] and the root of the tree is signed, and the UPDATE and the root of the hash tree is sent as the signed UPDATE to each peer as shown in Figure 14. This technique improves upon [116] which uses bi-linear maps instead of Merkle hash trees.

Even though there is ongoing research about aggregate signatures [117], [118] this approach appears to rely on a misinterpretation of the semantics of the Minimum Route Advertisement Interval Timer (MRAI Timer). The MRAI Timer is intended to prevent a BGP speaker from advertising an UPDATE for the same prefix to the same BGP peer within an MRAI Timer interval. Any such UPDATE that refers to a prefix for which an UPDATE has been sent within that time interval is to be held by the sender until the timer expires, at which time the prefix may be advertised to the peer. The authors assert that the MRAI Timer semantics is to hold all BGP UPDATE messages to the same peer until an MRAI Timer expires, at which time all queued UPDATE messages are released to the peer. They use this to then generate an aggregate signature across the collection of held UPDATE messages. While this update queueing behaviour describes the operation of some BGP implementations, other implementations set the MRAI Timer interval to zero, effectively bypassing the MRAI Timer completely, while others implement the timer precisely according to the BGP specification on a per peer per prefix level, in which case aggregation of UPDATE messages does not occur.

h) Exploiting Path Stability: Mitigating the validation overhead can also be achieved by caching validation outcomes and reapplying the outcome if the same update information is received within the cache lifetime. A study by Butler, McDaniel, and Aiello [119] noted that across a one month period less that 2% of advertised prefixes were advertised using more than 10 paths and less than 0.06% of prefixes were advertised with more than 20 paths. They proposed combining a number of approaches to reduce the AS Path validation workload. The first was the use of hash chains and signature aggregation, where a BGP speaker sends all local viable paths to its peers along with the tokens that represent hash chain anchors, allowing route change to be represented by an authentication token that can be validated by hash operations. The second was to use Merkle hash trees to sign across a set of UPDATE messages that are queued awaiting the MRAI Timer. The third part of the approach was to exploit the stability of path advertisements to amortise cryptographic operations over many validations, achieved by caching the cryptographic proofs. The authors assert that their simulations point to a reduction of the computational costs by as much as 97% over existing approaches using this approach. The same comments relating to the precise interpretation of the MRAI timer apply to this study, and it is unclear that the same results would be obtained if the MRAI timer were implemented on a per-prefix basis rather than a per-peer basis.

A recent study, pgBGP [120], analyses path stability over a longer period of time and builds a local database which is then consulted in order to detect anomalous routes. The idea is that origin ASes usually do not suddenly change over time for certain prefixes, and that such a sudden change might indicate an attack to the routing system. pgBGP does not provide completely automated security, as it does not eliminate any route advertisements, but rather puts them into quarantine for 24 hours (similar to route flap damping), giving operators the time to decide how to classify the event. This proposal can be incrementally deployed and imposes little overhead on the routing system. It is a method to mitigate effects of an attack to the routing system, and not an effective mechanism for prevention. A more recent study about pgBGP [121], shows the positive effects such an interim solution could have for the global Internet if deployed on a minimum amount of ASes.

i) The Threat of Prefix Hijacking: With the current BGP infrastructure missing adequate security mechanisms, attacks to the BGP systems are becoming more common, and have reached the interest of mainstream public media [9]. One special case of routing attacks which is considered a major threat and evokes high interest in the research community is *prefix hijacking*. An increasing amount of research is undertaken in order to provide security against this single form of attack. The approaches describe possible methods of detecting prefix hijacking [122], [123], [124], [125] as well as complete systems and implementations of prefix hijacking detection in order to possibly react on the attack. These systems [126], [127], [128], [129] rely on existing external route monitoring databases like Routeviews [130] or need special routing registries to be deployed [128] to detect prefix hijacking. The quality of such prefix hijack detection systems is strongly dependent on the quality of the databases [131], which usually are not deployed for this purpose.

Another method to detect prefix hijacking is to look for

multiple origin AS (MOAS) [132], [133], which can be either a sign of multi-homing an AS or a sign of bogus route announcements, thus prefix hijacking.

A different approach is presented for iSPY [134], 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. While the method to detect prefix hijacking has been proven successful and deployable with little overhead, it is unclear what measures a prefix owner can take to regain routing control over the prefix.

While prefix hijacking is seen as a major area of vulnerability in the current routing system, it is not the only possible attack, and many endpoint ASes seem already to be quite resilient to prefix hijacking attacks [135]. The quality of prefix hijack detection is also questionable, and strongly dependent on the database used [136] and it is also questionable whether it is worth the effort to build specific databases [137] in order to provide such *interrogative* route validation, which is mostly used for a small group of threats.

j) Secure BGP and BGP Dynamics: If securing BGP is a case of applying cryptographic operations to BGP UPDATE messages then the other approach to reducing the security overhead is to exploit the dynamic behaviour of these messages. The BGP update pattern is studied in [138] where in a study of BGP update dynamics it was shown that a cache of 10,000 prefix and AS Path validation outcomes, or less than 5% of the total number of distinct routed entries, would achieve a cache rate of between 30% to 50% using a simple LRU cache replacement algorithm.

When distance vector algorithms react to a change in prefix reachability a number of UPDATE messages are generally observed before the routing system reaches a stable state. A study of BGP convergence across the global Internet concluded that the severity of path exploration and the convergence speed depends on the relative positions of the event origin and the observer [139]. This study aligned the originator and the observer in terms of the "tiering" of Internet Service Providers and noted that this extended convergence times and larger path exploration events occurred at lower levels of the tiering hierarchy. It was hypothesised that the richer inter-connectivity that was typically prevalent at such lower levels in the tiering hierarchy was a major contributing factor here. Fail-over and new route announcements converge in similar times, while route withdrawals have far longer convergence times.

A similar study on BGP's path exploration characteristics proposed modifications to the BGP UPDATE message intended to identify and limit the path exploration behaviour of BGP [140]. 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.

Further study of BGP update behaviour has explored the level of determinism that exists in BGP's route selection process, and noted that in the absence of the Multiple Exit Discriminator (MED) and route reflectors, then the process can be considered to be a deterministic one [141]. 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.

3) Securing the Data Plane: Securing BGP is not only a matter of securing the control plane, but also of securing the data plane [142] and to make sure that the status of the forwarding table is consistent with the advertised BGP routing information. A study by Mao et al. [143] shows that up to 8% of the paths advertised through the control plane, do not match the actual paths in the data plane. The data plane is not only subject to attacks which try to subvert the routing system, but also subject to synthetic BGP announcements from network operators that could enable the theft of carriage capacity [5]. It is therefore necessary to provide security for the whole data path, and not only on a nexthop basis as *Stealth Probing* [144] intends to, as carriers might span over multiple ASes and synthesise false routing information that spans multiple AS hops. As ensuring security on a per packet basis for inter-domain routers which process billions of packets per second is quite irrational, approaches mainly focus on probing the full data-path through packet injection, trying to detect and isolate malicious routers.

In [145] a modified traceroute (*secure 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 [146], [147] instead focuses on using *traffic summary functions*, and comparing their results with those of other routers, allowing to detect ASes which provide anomalous values. These *traffic summary functions* seem to be prone to inaccuracy due to a variety of applications running on routers which might alter the packet flow and their application appears infeasible in routers with data traffic loads of a scale up to billions of packets per second.

The solution proposed as *Listen and Whisper* [148] 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* has to be classified as a "just too late" solution for BGP security, like many solutions which try to ensure data plane - control plane consistency. Like other data plane security solutions, this approach seems infeasible, as it tries to detect data plane anomalies by analysing TCP flows, which can be millions per second on heavily trafficked routes.

An approach that aims towards high performance and possible partial deployment is described in [149]. It's focus is to ensure that the data path always conforms to the announced AS path, which is achieved by probing data paths through

injecting tagged IP packets, or by using IP options. Similar to pgBGP, it leaves the decision of which action to take towards a malicious router to the network operator and builds up a small database to detect possible malicious routers. It deploys the roles of *verifiers* and *provers* on certain ASes, with the *verifier* being an AS that wants to verify a certain route, and the *prover* being an AS that helps the *verifier* in the process by replying on probe data.

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

4) State of BGP Security: No current solution to routing security has found an adequate balance between appropriate security and acceptable deployment overhead [65]. 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 [150]. These two fields of research are inherently connected, in that a more stable routing system that was able to provide clear indications when convergence to a stable routing state had been achieved is believed to also provide clear indications of when verification of routing information is appropriate.

In exploring the threat model for BGP it is noted that BGP was designed to support inter-domain routing between trusted networks, while today's networks operate in a looser confederation that does not exhibit the same mutual trust properties. Not only are the TCP sessions used by BGP vulnerable to attack, and the messages used by BGP vulnerable to alteration in order to disrupt the network's routing system, but the integrity of the operation of BGP is also threatened by misconfiguration, where incorrect information is injected into the routing system unintentionally, and by router vulnerabilities where a compromised routing system can exploit its trusted role and intentionally inject false information into the routing system. Some of these attacks are intended to cause a BGP speaker to be overwhelmed and reset, as BGP is a method of directly accessing a router's processing unit and a saturation attack can cause processor and memory overload. Other attacks are aimed at altering the router's forwarding state, generating an incorrect or unintended forwarding state for one or more prefixes. Other forms of attack are aimed at causing a BGP speaker to become unstable and thereby disrupt the forwarding function and impact on applications. A BGP session that is being continually reset will cause large local traffic bursts as neighbouring BGP speakers continually resend their routing tables upon each reset, but the continued instability will trigger a flap damping response in other BGP speakers [6].

The factors that contribute to these vulnerabilities include a lack of BGP message integrity checks, an inability 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.

System	Type	Reference Implementation	Deployed
GTSM	session security only	Yes[7]	Yes
sBGP	crypto	Yes[8]	No
soBGP	crypto/anomaly	No[9]	No
psBGP	crypto	No	No
IRV	crypto/anomaly	No	No
SPV	crypto	No	No
pgBGP	anomaly	Yes[10]	Yes
iSPY	anomaly	No	No
PHAS	anomaly	No[11]	No
Secure Traceroute	crypto	No	No
Fatih	anomaly	No	No
Listen & Whisper	crypto/anomaly	No	No

TABLE I
DEPLOYMENT AND IMPLEMENTATION STATUS OF BGP SECURITY APPROACHES

In terms of message integrity, heuristic mechanisms that can assess confidence levels in the authenticity of origination assertions are attractive simply because they do not require concerted action on the part of all BGP speakers, although the outcomes are such that incorrect routing information cannot be reliably detected in all cases. The extent to which such mechanisms are useful in the face of informed attack is limited, in that an informed attack would normally be expected to exploit the weaknesses in such heuristic approaches, negating the overall value of the effort [150].

On the other hand, use of a PKI to support address attestations, as in sBGP, provides a very robust means of detecting incorrect origin route objects, as long as the PKI itself is accurately aligned to the address distribution framework and as long as the PKI is universally used [151]. The most effective approach for securing origination information in BGP appears to be for the operational community to regain control of the address space, and it is now necessary to solve the operational challenge of certifying the ownership of the IP address space [152], [153].

In contrast, robust solutions to the problem of AS path authentication have been elusive so far. sBGP provides a robust method of path validation, but has been assessed to be significantly expensive in terms of processor and memory cost, and also detrimental to BGP convergence times and requires comprehensive adoption to be effective. Efforts to mitigate these costs through IRV query approaches, or substituting path feasibility in place of actual path validity, as is the case with soBGP do not appear to be adequately robust. It is also likely that sBGP will still be subject to attacks at data plane level [154]. None of the solutions for BGP path validation that have been proposed have provided appropriate trade-offs between security, resource usage, and deployability. As Table I shows, of all proposals, only few have been implemented and it does not appear that any of them are actually deployed in real production environments.

[7]Patch for the Quagga Software Routing Suite version 0.99.11 [155]
[8]Dead project [156]
[9]Guidelines exist [157]
[10]Patch for the Quagga Software Routing Suite version 0.99.9 [158]
[11]Maybe in a near future through SHASAM [137]

VI. Some Open Questions on Securing BGP

The study of approaches to securing BGP has raised a number of 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 [12]. Is it possible to secure this association such that any relying party can validate that the AS path as presented in a BGP UPDATE not only matches the path taken by the prefix advertisement, but that the network's current forwarding state to reach the address prefix is aligned to this AS path and this alignment can be validated? This question is not one that is directly addressed within any of the current set of inter-domain routing security proposals.

The related issue concerns the overheads of securing BGP and the scaling properties of BGP. In 2001 one of the authors (Huston) [159] argued that BGP may already be 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. The commentary argued that inter-domain routing might be more scalable if these functions where performed by separate protocols [160], [159]. Adding security and authentication to BGP, as sBGP does, increases the complexity of the protocol and may diminish its long term prospects for scalability across ever larger and denser inter-domain topologies [161].

At present there are a number of practical and a number of 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 associated with BGP UPDATE messages. The question concerns what aspects of securing BGP should be considered essential and what is considered to be 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 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 actually represents current forwarding state in each routing element in the network (securing the data plane). One perspective on this is that securing one element of system with multiple components does not necessarily address the underlying vulnerabilities of the entire system. The more common outcome is that such work exposes the residual vulnerabilities in other components, and that an effective security system needs to address all components of the routing system. While it may be possible for a BGP speaker to be able to validate that the originating AS did indeed originate the prefix advertisement and that the AS path accurately represents the propagation path of this advertisement through the network, that is not the basic question in terms of the properties of the overall system. The more basic question is can a BGP speaker verify that if it makes a decision to forward a packet on the first hop along a path indicated by the routing system as the optimal path to a destination is this indeed the optimal local choice and does this first hop decision lead along a path to the destination address?

If a comprehensive security framework is proving to be elusive in terms of deployment considerations then could a less comprehensive approach offer acceptable outcomes? Many security frameworks demonstrate a profile of diminishing returns, where the incremental cost of deployment of additional security capabilities increases, while the incremental benefit in terms of risk mitigation decreases. In the case of securing BGP could an approach of reducing the security credential generation and validation workload, through reducing the amount or timeliness of validated information, represent an acceptable trade-off?

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

VII. Conclusion

BGP has proven surprisingly resilient in terms of its longevity of useful operational life [162], despite early predictions of its imminent demise in favour of IDRP [21]. BGP version 4 has routed the inter-domain Internet since late 1993 and the number of routed elements has grown from under 20,000 distinct prefixes to in excess of 300,000 distinct prefixes at any point in time by the middle of 2009 [30] and with the need for more IP addresses and the parallel deployment of IPv4 and IPv6 using the BGP multi-protocol extension [163], these numbers are destined to grow even more rapidly. Due to its extensibility and large installed base, BGP version 4 will also most likely remain the only inter-domain routing protocol in the near to mid-term future.

Across this period BGP has not changed in any substantive manner, including in its security properties. Some operators

use MD5 protection on BGP sessions, particularly in the context of exchange point configurations where the potential for attack at the session level is considered to be higher, but the overall picture of BGP security is largely unchanged. This is in spite of ample evidence from inadvertent misconfiguration through to reports use of unregistered addresses [164] or of research on hostile application level traffic [43] that BGP is abused in various ways. Current efforts at mitigation of these forms of abuse appear to be inadequate and the ease with which unauthorised or bogus route objects can be injected into the inter-domain routing system remains a significant issue for the security, stability and utility of the Internet.

There have been a number of approaches proposed that would provide significantly greater levels of assurance that what is being routed is precisely what was intended to be routed, but these approaches all appear to rely on the availability of a security infrastructure that does not exist today. The most obvious omission in today's environment appears to be a PKI for addresses and ASes that would allow anyone to verify a digitally signed attestation relating to addresses and their use [152]. With such a PKI it would then be possible to improve the situation regarding the security of addresses and their advertisement into the inter-domain routing system.

VIII. Acknowledgements

This work has been made possible in part by a grant from the Cisco University Research Program Fund at Community Foundation Silicon Valley.

References

[1] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271 (Draft Standard), Internet Engineering Task Force, Jan. 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4271.txt

[2] Y. Rekhter, "Experience with the BGP Protocol," RFC 1266 (Informational), Internet Engineering Task Force, Oct. 1991. [Online]. Available: http://www.ietf.org/rfc/rfc1266.txt

[3] Office of the President of the United States, "Priority II: A National Cyberspace Security Threat and Vulnerability Reduction Program," 2004. [Online]. Available: http://www.us-cert.gov/reading_room/cyberspace_strategy.pdf

[4] M. A. Brown, "Pakistan hijacks YouTube," Renesys Blog, Feb 2008. [Online]. Available: http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml

[5] S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R. N. Wright, "Rationality and traffic attraction: incentives for honest path announcements in BGP," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 267–278, 2008.

[6] S. Murphy, "BGP Security Vulnerabilities Analysis," RFC 4272 (Informational), Internet Engineering Task Force, Jan. 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4272.txt

[7] A. Barbir, S. Murphy, and Y. Yang, "Generic Threats to Routing Protocols," RFC 4593 (Informational), Internet Engineering Task Force, Oct. 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4593.txt

[8] H. Ballani, P. Francis, and X. Zhang, "A study of prefix hijacking and interception in the Internet," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 265–276, 2007.

[9] K. Zetter, "Revealed: The Internet's Biggest Security Hole," Wired Magazine - ThreadLevel, Aug 2008. [Online]. Available: http://www.wired.com/threatlevel/2008/08/revealed-the-in/

[10] V. J. Bono, "7007 Explanation and Apology," Apr. 1997. [Online]. Available: http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html

[11] T. Underwood, "Con-Ed Steals the 'Net'," Renesys Blog, Jan 2006. [Online]. Available: http://www.renesys.com/blog/2006/01/coned-steals-the-net.shtml

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

[13] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB Workshop on Routing and Addressing," RFC 4984 (Informational), Internet Engineering Task Force, Sep. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4984.txt

[14] G. Huston, "The BGP Report for 2005," June 2006. [Online]. Available: http://www.potaroo.net/ispcol/2006-06/bgpupds.html

[15] M. Nicholes and B. Mukherjee, "A survey of security techniques for the Border Gateway Protocol (BGP)," *Communications Surveys and Tutorials, IEEE*, vol. 11, no. 1, pp. 52–65, Quarter 2009.

[16] C. Huitema, *Routing in the Internet*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1995.

[17] C. Kaufman, R. Perlman, and M. Speciner, *Network Security: Private Communication in a Public World, 2nd Edition*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 2002.

[18] J. Postel, "Internet Protocol," RFC 791 (Standard), Internet Engineering Task Force, Sep. 1981, updated by RFC 1349. [Online]. Available: http://www.ietf.org/rfc/rfc791.txt

[19] B. Halabi, *Internet Routing Architectures*. Cisco Press, 1997.

[20] B. Donnet and T. Friedman, "Internet topology discovery: a survey," *Communications Surveys and Tutorials, IEEE*, vol. 9, no. 4, pp. 56–69, Quarter 2007.

[21] J. Hawkinson and T. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)," RFC 1930 (Best Current Practice), Internet Engineering Task Force, Mar. 1996. [Online]. Available: http://www.ietf.org/rfc/rfc1930.txt

[22] C. Hedrick, "Routing Information Protocol," RFC 1058 (Historic), Internet Engineering Task Force, Jun. 1988, updated by RFCs 1388, 1723. [Online]. Available: http://www.ietf.org/rfc/rfc1058.txt

[23] J. Moy, "OSPF Version 2," RFC 2328 (Standard), Internet Engineering Task Force, Apr. 1998. [Online]. Available: http://www.ietf.org/rfc/rfc2328.txt

[24] "Intermediate System to Intermediate System intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)," ISO/IEC 10589:2002, 2002. [Online]. Available: http://standards.iso.org/ittf/PubliclyAvailableStandards/c030932_ISO_IEC_10589_2002(E).zip

[25] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier, "Network topologies: inference, modeling, and generation," *Communications Surveys and Tutorials, IEEE*, vol. 10, no. 2, pp. 48–69, Quarter 2008.

[26] K. Lougheed and Y. Rekhter, "Border Gateway Protocol (BGP)," RFC 1105 (Experimental), Internet Engineering Task Force, Jun. 1989, obsoleted by RFC 1163. [Online]. Available: http://www.ietf.org/rfc/rfc1105.txt

[27] K. Lougheed and Y. Rekhter, "Border Gateway Protocol (BGP)," RFC 1163 (Historic), Internet Engineering Task Force, Jun. 1990, obsoleted by RFC 1267. [Online]. Available: http://www.ietf.org/rfc/rfc1163.txt

[28] K. Lougheed and Y. Rekhter, "Border Gateway Protocol 3 (BGP-3)," RFC 1267 (Historic), Internet Engineering Task Force, Oct. 1991. [Online]. Available: http://www.ietf.org/rfc/rfc1267.txt

[29] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771 (Draft Standard), Internet Engineering Task Force, Mar. 1995, obsoleted by RFC 4271. [Online]. Available: http://www.ietf.org/rfc/rfc1771.txt

[30] G. Huston, "BGP Routing Table Resource Pages," June 2009. [Online]. Available: http://bgp.potaroo.net

[31] J. Postel, "Transmission Control Protocol," RFC 793 (Standard), Internet Engineering Task Force, Sep. 1981, updated by RFCs 1122, 3168. [Online]. Available: http://www.ietf.org/rfc/rfc793.txt

[32] E. Chen, "Route Refresh Capability for BGP-4," RFC 2918 (Proposed Standard), Internet Engineering Task Force, Sep. 2000. [Online]. Available: http://www.ietf.org/rfc/rfc2918.txt

[33] T. Bates, R. Chandra, and E. Chen, "BGP Route Reflection - An Alternative to Full Mesh IBGP," RFC 2796 (Proposed Standard), Internet Engineering Task Force, Apr. 2000, obsoleted by RFC 4456. [Online]. Available: http://www.ietf.org/rfc/rfc2796.txt

[34] T. Griffin and G. Huston, "BGP Wedgies," RFC 4264 (Informational), Internet Engineering Task Force, Nov. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4264.txt

[35] F. Wang and L. Gao, "On inferring and characterizing internet routing policies," in *IMC '03: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*. New York, NY, USA: ACM, 2003, pp. 15–26.

[36] A. Ramaiah, R. Stewart, and M. Dalal, "Improving TCP's Robustness to Blind In-Window Attacks," Nov. 2008. [Online]. Available: http://tools.ietf.org/html/draft-ietf-tcpm-tcpsecure-11

[37] C. Villamizar, R. Chandra, and R. Govindan, "BGP Route Flap Damping," RFC 2439 (Proposed Standard), Internet Engineering Task Force, Nov. 1998. [Online]. Available: http://www.ietf.org/rfc/rfc2439.txt

[38] P. Smith and C. Panigl, "RIPE Routing Working Group Recommendations on Route-flap Damping," ripe-378, May 2006, obsoletes: ripe-229, ripe-210, ripe-178. [Online]. Available: http://www.ripe.net/docs/ripe-378.html

[39] K. Sriram, D. Montgomery, O. Borchert, O. Kim, and D. Kuhn, "Study of BGP peering session attacks and their impacts on routing performance," *Selected Areas in Communications, IEEE Journal on*, vol. 24, no. 10, pp. 1901–1915, Oct. 2006.

[40] O. Nordström and C. Dovrolis, "Beware of BGP attacks," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 2, pp. 1–8, 2004.

[41] P. Boothe, J. Hiebert, and R. Bush, "How prevalent is prefix hijacking on the internet?" February 2006. [Online]. Available: http://www.nanog.org/meetings/nanog36/presentations/boothe.pdf

[42] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *SIGCOMM '02: Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*. New York, NY, USA: ACM, 2002, pp. 3–16.

[43] A. Ramachandran and N. Feamster, "Understanding the network-level behavior of spammers," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 4, pp. 291–302, 2006.

[44] K. J. Houle and G. M. Weaver, "Trends in Denial of service Attack Technology," Oct. 2001. [Online]. Available: http://www.cert.org/archive/pdf/DoS_trends.pdf

[45] B. Kumar, "Integration of security in network routing protocols," *SIGSAC Rev.*, vol. 11, no. 2, pp. 18–25, 1993.

[46] B. Kumar and J. Crowcroft, "Integrating security in inter-domain routing protocols," *SIGCOMM Comput. Commun. Rev.*, vol. 23, no. 5, pp. 36–51, 1993.

[47] G. Huston, "Securing Routing - An ISP's Perspective," February 2005. [Online]. Available: http://www.potaroo.net/ispcol/2005-02/route-sec.html

[48] V. Gill, J. Heasley, and D. Meyer, "The Generalized TTL Security Mechanism (GTSM)," RFC 3682 (Experimental), Internet Engineering Task Force, Feb. 2004, obsoleted by RFC 5082. [Online]. Available: http://www.ietf.org/rfc/rfc3682.txt

[49] V. Gill, J. Heasley, D. Meyer, P. Savola, and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)," RFC 5082 (Proposed Standard), Internet Engineering Task Force, Oct. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc5082.txt

[50] W. Eddy, "TCP SYN Flooding Attacks and Common Mitigations," RFC 4987 (Informational), Internet Engineering Task Force, Aug. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4987.txt

[51] S. Kent and R. Atkinson, "Security Architecture for the Internet Protocol," RFC 2401 (Proposed Standard), Internet Engineering Task Force, Nov. 1998, obsoleted by RFC 4301, updated by RFC 3168. [Online]. Available: http://www.ietf.org/rfc/rfc2401.txt

[52] R. Rivest, "The MD5 Message-Digest Algorithm," RFC 1321 (Informational), Internet Engineering Task Force, Apr. 1992. [Online]. Available: http://www.ietf.org/rfc/rfc1321.txt

[53] A. Heffernan, "Protection of BGP Sessions via the TCP MD5 Signature Option," RFC 2385 (Proposed Standard), Internet Engineering Task Force, Aug. 1998. [Online]. Available: http://www.ietf.org/rfc/rfc2385.txt

[54] M. Behringer, "BGP Session Security Requirements," Internet-Draft (Informational), Aug. 2007. [Online]. Available: http://tools.ietf.org/html/draft-behringer-bgp-session-sec-req-02

[55] B. Christian and T. Tauber, "BGP Security Requirements," Internet-Draft (Informational), Nov. 2008. [Online]. Available: http://tools.ietf.org/html/draft-ietf-rpsec-bgpsecrec-10

[56] B. Schneier, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, P. Sutherland, Ed. New York, NY, USA: John Wiley & Sons, Inc., 1995.

[57] S. Murphy, "BGP Security Analysis," Nov. 2001. [Online]. Available: http://tools.ietf.org/html/draft-murphy-bgp-secr-04

[58] R. Housley, W. Polk, W. Ford, and D. Solo, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," RFC 3280 (Proposed Standard), Internet Engineering Task Force, Apr. 2002, obsoleted by RFC 5280, updated by RFCs 4325, 4630. [Online]. Available: http://www.ietf.org/rfc/rfc3280.txt

[59] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," RFC 5280 (Proposed Standard), Internet Engineering Task Force, May 2008. [Online]. Available: http://www.ietf.org/rfc/rfc5280.txt

[60] C. Lynn, S. Kent, and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers," RFC 3779 (Proposed Standard), Internet Engineering Task Force, Jun. 2004. [Online]. Available: http://www.ietf.org/rfc/rfc3779.txt

[61] X. He, C. Papadopoulos, and P. Radoslavov, "A framework for incremental deployment strategies for router-assisted services," in *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies. IEEE*, vol. 2, March-3 April 2003, pp. 1488–1498 vol.2.

[62] M. Suchara, I. Avramopoulos, and J. Rexford, "Securing BGP incrementally," in *CoNEXT '07: Proceedings of the 2007 ACM CoNEXT conference*. New York, NY, USA: ACM, 2007, pp. 1–2.

[63] J. Rexford and J. Feigenbaum, "Incrementally-deployable security for interdomain routing," in *Conference For Homeland Security, 2009. CATCH '09. Cybersecurity Applications & Technology*, March 2009, pp. 130–134.

[64] S. Gorman, R. Kulkarni, L. Schintler, and R. Stough, "Least effort strategies for cybersecurity," 2003. [Online]. Available: http://arxiv.org/pdf/cond-mat/0306002

[65] H. Chan, D. Dash, A. Perrig, and H. Zhang, "Modeling adoptability of secure BGP protocols," in *SIGMETRICS '06/Performance '06: Proceedings of the joint international conference on Measurement and modeling of computer systems*. New York, NY, USA: ACM, 2006, pp. 389–390.

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

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

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

[69] P. v. Oorschot, T. Wan, and E. Kranakis, "On interdomain routing security and pretty secure BGP (psBGP)," *ACM Trans. Inf. Syst. Secur.*, vol. 10, no. 3, p. 11, 2007.

[70] 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.

[71] S. M. Bellovin, "Security problems in the TCP/IP protocol suite," *SIGCOMM Comput. Commun. Rev.*, vol. 19, no. 2, pp. 32–48, 1989.

[72] M. Leech, "Key Management Considerations for the TCP MD5 Signature Option," RFC 3562 (Informational), Internet Engineering Task Force, Jul. 2003. [Online]. Available: http://www.ietf.org/rfc/rfc3562.txt

[73] S. Bellovin, "Key Change Strategies for TCP-MD5," RFC 4808 (Informational), Internet Engineering Task Force, Mar. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4808.txt

[74] S. Bellovin and R. Housley, "Guidelines for Cryptographic Key Management," RFC 4107 (Best Current Practice), Internet Engineering Task Force, Jun. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4107.txt

[75] S. Bellovin and E. Rescorla, "Deploying a new Hash Algorithm," in *NIST Cryptographic Hash Workshop*, October 2005. [Online]. Available: http://csrc.nist.gov/groups/ST/hash/documents/Bellovin_new-hash.pdf

[76] B. Burr, "NIST Cryptographic Standards Status Report," Internet 2 5th Annual PKI R&D Workshop, April 2006. [Online]. Available: http://middleware.internet2.edu/pki06/proceedings/burr-nist_crypto_standards.ppt

[77] X. Wang and H. Yu, "How to Break MD5 and Other Hash Functions," in *EUROCRYPT*, 2005, pp. 19–35. [Online]. Available: http://dx.doi.org/10.1007/11426639_2

[78] S. Bellovin and A. Zinin, "Standards Maturity Variance Regarding the TCP MD5 Signature Option (RFC 2385) and the BGP-4 Specification," RFC 4278 (Informational), Internet Engineering Task Force, Jan. 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4278.txt

[79] D. Eastlake 3rd and T. Hansen, "US Secure Hash Algorithms (SHA and HMAC-SHA)," RFC 4634 (Informational), Internet

Engineering Task Force, Jul. 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4634.txt

[80] T. Krovetz, "UMAC: Message Authentication Code using Universal Hashing," RFC 4418 (Informational), Internet Engineering Task Force, Mar. 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4418.txt

[81] H. Krawczyk, M. Bellare, and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication," RFC 2104 (Informational), Internet Engineering Task Force, Feb. 1997. [Online]. Available: http://www.ietf.org/rfc/rfc2104.txt

[82] C. Madson and R. Glenn, "The Use of HMAC-MD5-96 within ESP and AH," RFC 2403 (Proposed Standard), Internet Engineering Task Force, Nov. 1998. [Online]. Available: http://www.ietf.org/rfc/rfc2403.txt

[83] R. Bonica, B. Weis, S. Viswanathan, A. Lange, and O. Wheeler, "Authentication for TCP-based Routing and Management Protocols," Feb. 2007. [Online]. Available: http://tools.ietf.org/html/draft-bonica-tcp-auth-06

[84] B. Weis, C. Apanna, D. McGrew, and A. Ramaiah, "Automated key selection extension for the TCP Enhanced Authentication Option," Oct. 2007. [Online]. Available: http://tools.ietf.org/html/draft-weis-tcp-auth-auto-ks-03

[85] J. Touch, A. Mankin, and R. Bonica, "The TCP Authentication Option," Mar. 2009. [Online]. Available: http://tools.ietf.org/html/draft-ietf-tcpm-tcp-auth-opt-04

[86] R. Thayer, N. Doraswamy, and R. Glenn, "IP Security Document Roadmap," RFC 2411 (Informational), Internet Engineering Task Force, Nov. 1998. [Online]. Available: http://www.ietf.org/rfc/rfc2411.txt

[87] D. Ward, "Securing BGPv4 using IPsec," Jan. 2002. [Online]. Available: http://tools.ietf.org/html/draft-ward-bgp-ipsec-00

[88] "HP-UX IPSec Performance and Sizing White Paper," December 2005. [Online]. Available: http://www.docs.hp.com/en/6092/ipsecperf/ipsecperf.pdf

[89] R. Perlman, "Network Layer Protocols with Byzantine Robustness," Tech. Rep., 1988. [Online]. Available: http://www.lcs.mit.edu/publications/specpub.php?id=997

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

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

[92] "Protocol for exchange of inter-domain routeing information among intermediate systems to support forwarding of ISO 8473 PDUs," ISO/IEC 10747:1994, 1994. [Online]. Available: http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=21417

[93] T. Bates, R. Bush, T. Li, and Y. Rhekter, "DNS-based NLRI origin AS verification in BGP," Jul. 1998. [Online]. Available: http://tools.ietf.org/html/draft-bates-bgp4-nlri-orig-verif-00

[94] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security Introduction and Requirements," RFC 4033 (Proposed Standard), Internet Engineering Task Force, Mar. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4033.txt

[95] L. Donnerhacke and W. Wijngaards, "DNSSEC protected routing announcements for BGP," May 2008. [Online]. Available: http://tools.ietf.org/html/draft-donnerhacke-sidr-bgp-verification-dnssec-04

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

[97] M. Zhao and D. Nicol, "Evaluating the Performance Impact of PKI on BGP Security," Internet 2 4th Annual PKI R&D Workshop, April 2005. [Online]. Available: http://middleware.internet2.edu/pki05/proceedings/zhao-sbgp.pdf

[98] M. Zhao, S. Smith, and D. Nicol, "The performance impact of BGP security," *Network, IEEE*, vol. 19, no. 6, pp. 42–48, Nov.-Dec. 2005.

[99] S. T. Kent, "Securing the border gateway protocol: A status update," in *Seventh IFIP TC-6 TC-11 Conference on Communications and Multimedia Security*, Torino, 2003.

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

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

[102] T. Bates, E. Gerich, L. Joncheray, J.-M. Jouanigot, D. Karrenberg, M. Terpstra, and J. Yu, "Representation of IP Routing Policies in a Routing Registry (ripe-81++)," RFC 1786 (Informational), Internet Engineering Task Force, Mar. 1995. [Online]. Available: http://www.ietf.org/rfc/rfc1786.txt

[103] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg, and M. Terpstra, "Routing Policy Specification Language (RPSL)," RFC 2622 (Proposed Standard), Internet Engineering Task Force, Jun. 1999, updated by RFC 4012. [Online]. Available: http://www.ietf.org/rfc/rfc2622.txt

[104] L. Blunk, J. Damas, F. Parent, and A. Robachevsky, "Routing Policy Specification Language next generation (RPSLng)," RFC 4012 (Proposed Standard), Internet Engineering Task Force, Mar. 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4012.txt

[105] "Internet Routing Registry." [Online]. Available: http://www.irr.net

[106] R. Kisteleki and J. Boumans, "Securing RPSL Objects with RPKI Signatures," Oct. 2008. [Online]. Available: http://tools.ietf.org/id/draft-kisteleki-sidr-rpsl-sig-00.txt

[107] Y.-C. Hu, A. Perrig, and D. Johnson, "Efficient Security Mechanisms for Routing Protocols," in *Proceedings of Internet Society Symposium on Network and Distributed System Security (NDSS 03)*, February 2003.

[108] P. McDaniel, W. Aiello, K. Butler, and J. Ioannidis, "Origin authentication in interdomain routing," *Comput. Netw.*, vol. 50, no. 16, pp. 2953–2980, 2006.

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

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

[111] C. K. Wong and S. Lam, "Digital signatures for flows and multicasts," *Networking, IEEE/ACM Transactions on*, vol. 7, no. 4, pp. 502–513, Aug 1999.

[112] S. Even, O. Goldreich, and S. Micali, "On-line/off-line digital signatures," in *CRYPTO '89: Proceedings on Advances in cryptology*. New York, NY, USA: Springer-Verlag New York, Inc., 1989, pp. 263–275.

[113] B. Raghavan, S. Panjwani, and A. Mityagin, "Analysis of the SPV secure routing protocol: weaknesses and lessons," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 2, pp. 29–38, 2007.

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

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

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

[117] A. Boldyreva, C. Gentry, A. O'Neill, and D. H. Yum, "Ordered multisignatures and identity-based sequential aggregate signatures, with applications to secure routing," in *CCS '07: Proceedings of the 14th ACM conference on Computer and communications security*. New York, NY, USA: ACM, 2007, pp. 276–285.

[118] A. Boldyreva, C. Gentry, A. O'Neill, and D. H. Yum, "New multiparty signature schemes for network routing applications," *ACM Trans. Inf. Syst. Secur.*, vol. 12, no. 1, pp. 1–39, 2008.

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

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

[121] J. Karlin, S. Forrest, and J. Rexford, "Autonomous security for autonomous systems," *Comput. Netw.*, vol. 52, no. 15, pp. 2908–2923, 2008.

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

[123] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis, "A light-weight distributed scheme for detecting ip prefix hijacks in real-time," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 277–288, 2007.

[124] X. Hu and Z. M. Mao, "Accurate real-time identification of IP prefix hijacking," in *SP '07: Proceedings of the 2007 IEEE Symposium on*

Security and Privacy. Washington, DC, USA: IEEE Computer Society, 2007, pp. 3–17.

[125] X. Hu and Z. Mao, "Accurate Real-time Identification of IP Prefix Hijacking," in *Security and Privacy, 2007. SP '07. IEEE Symposium on*, May 2007, pp. 3–17.

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

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

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

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

[130] T. U. of Oregon, "University of Oregon Route Views Project." [Online]. Available: http://www.routeviews.org

[131] Y. Zhang, Z. Zhang, Z. M. Mao, C. Hu, and B. MacDowell Maggs, "On the impact of route monitor selection," in *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. New York, NY, USA: ACM, 2007, pp. 215–220.

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

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

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

[135] M. Lad, R. Oliveira, B. Zhang, and L. Zhang, "Understanding resiliency of internet topology against prefix hijack attacks," in *DSN '07: Proceedings of the 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. Washington, DC, USA: IEEE Computer Society, 2007, pp. 368–377.

[136] K. Sriram, O. Borchert, O. Kim, P. Gleichmann, and D. Montgomery, "A Comparative Analysis of BGP Anomaly Detection and Robustness Algorithms," *Conference For Homeland Security, Cybersecurity Applications & Technology*, vol. 0, pp. 25–38, 2009.

[137] C. S. N. S. Group, "Simple Hijack Alert System And Monitor (SHASAM) is Coming!" 2009. [Online]. Available: http://phas.netsec.colostate.edu/

[138] G. 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)*, vol. 4427. Heidelberg, DE: Springer-Verlag Berlin, April 2007, pp. 1–10.

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

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

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

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

[143] Z. M. Mao, J. Rexford, J. Wang, and R. 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*. New York, NY, USA: ACM, 2003, pp. 365–378.

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

[145] V. N. Padmanabhan and D. R. Simon, "Secure traceroute to detect faulty or malicious routing," *SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 1, pp. 77–82, 2003.

[146] A. T. Mizrak, "Fatih: Detecting and isolating malicious routers," in *DSN '05: Proceedings of the 2005 International Conference on Dependable Systems and Networks*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 538–547.

[147] Y.-C. Cheng, "Detecting and isolating malicious routers," *IEEE Trans. Dependable Secur. Comput.*, vol. 3, no. 3, pp. 230–244, 2006, student Member-Mizrak, Alper Tugay and Member-Marzullo, Keith and Member-Savage, Stefan.

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

[149] E. L. Wong, P. Balasubramanian, L. Alvisi, M. G. Gouda, and V. 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*. New York, NY, USA: ACM, 2007, pp. 147–156.

[150] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, Jan. 2010.

[151] S. Wilson, "Public key superstructure 'it's PKI Jim, but not as we know it!'," in *IDtrust '08: Proceedings of the 7th symposium on Identity and trust on the Internet*. New York, NY, USA: ACM, 2008, pp. 72–88.

[152] M. Parameswaran, X. Zhao, A. B. Whinston, and F. Fang, "Reengineering the Internet for better security," *Computer*, vol. 40, no. 1, pp. 40–44, 2007.

[153] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," Jul. 2009. [Online]. Available: http://tools.ietf.org/html/draft-ietf-sidr-arch-08

[154] D. Pei, L. Zhang, and D. Massey, "A framework for resilient Internet routing protocols," *Network, IEEE*, vol. 18, no. 2, pp. 5–12, Mar-Apr 2004.

[155] N. Hilliard, "RFC 5082 GTSM for Quagga bgpd," Nov. 2008. [Online]. Available: http://lists.quagga.net/pipermail/quagga-dev/2008-November/006116.html

[156] B. S. LLC, "Secure BGP Prototype Software," 2003. [Online]. Available: http://www.ir.bbn.com/sbgp/src/S-BGP-1.0.html

[157] J. Ng, "Extensions to BGP to Support Secure Origin BGP (soBGP)," Apr. 2004. [Online]. Available: http://tools.ietf.org/html/draft-ng-sobgp-bgp-extensions-02

[158] J. Karlin, "Pretty Good BGP - Quagga Reference Implementation," Jul. 2008. [Online]. Available: http://lists.quagga.net/pipermail/quagga-dev/2008-July/005574.html

[159] G. Huston, "Commentary on Inter-Domain Routing in the Internet," RFC 3221 (Informational), Internet Engineering Task Force, Dec. 2001. [Online]. Available: http://www.ietf.org/rfc/rfc3221.txt

[160] G. Huston, "Scaling Inter-Domain Routing - A View Forward," *The Internet Protocol Journal*, vol. 4, no. 4, Dec 2001.

[161] D. Meyer and A. Partan, "BGP security, availability and operator needs," June 2003. [Online]. Available: http://www.nanog.org/meetings/nanog28/presentations/meyer.pdf

[162] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz, "BGP routing dynamics revisited," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 2, pp. 5–16, 2007.

[163] T. Bates, R. Chandra, D. Katz, and Y. Rekhter, "Multiprotocol Extensions for BGP-4," RFC 4760 (Draft Standard), Internet Engineering Task Force, Jan. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4760.txt

[164] T. Bates, P. Smith, and G. Huston, "The CIDR Report," last accessed: June 2009. [Online]. Available: http://www.cidr-report.org/as2.0/

Geoff Huston holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is currently the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He is a longstanding member of the Internet Engineering Task Force, a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of the Internet Society from 1992 until 2001.

Mattia Rossi holds a B.Eng. and a M.Sc. (Dipl.Ing.) from the Leopold-Franzens-Universitaet Innsbruck, Austria. He is currently working as Research and Development Engineer at the Centre for Advanced Internet Architectures at Swinburne University of Technology, Melbourne, Australia. He has been involved in transport layer and network security research, and is performing BGP, routing and network layer related research since 2008.

Grenville Armitage earned a B.Eng (Elec)(Hons) in 1988 and a PhD in electronic engineering in 1994, both from the University of Melbourne, Australia. He is currently Professor of Telecommunications Engineering and Director of the Centre for Advanced Internet Architectures at Swinburne University of Technology, Melbourne, Australia. He authored Quality of Service In IP Networks: Foundations for a Multi-Service Internet (Macmillan Technical Publishing, April 2000) and co-authored Networking and Online Games - Understanding and Engineering Multiplayer Internet Games (John Wiley Sons, UK, April 2006). Professor Armitage is also a member of ACM and ACM SIGCOMM.