A Survey of BGP Security Issues and Solutions

The Border Gateway Protocol (BGP) controls much of Internet traffic, but is vulnerable to communications interruptions and failures; finding suitable improved security measures with acceptable costs is difficult.

By Kevin Butler, *Student Member IEEE*, Toni R. Farley, Patrick McDaniel, *Senior Member IEEE*, and Jennifer Rexford, *Senior Member IEEE*

ABSTRACT | As the Internet's *de facto* interdomain routing protocol, the Border Gateway Protocol (BGP) is the glue that holds the disparate parts of the Internet together. A major limitation of BGP is its failure to adequately address security. Recent high-profile outages and security analyses clearly indicate that the Internet routing infrastructure is highly vulnerable. Moreover, the design of BGP and the ubiquity of its deployment have frustrated past efforts at securing interdomain routing. This paper considers the current vulnerabilities of the interdomain routing system and surveys both research and standardization efforts relating to BGP security. We explore the limitations and advantages of proposed security extensions to BGP, and explain why no solution has yet struck an adequate balance between comprehensive security and deployment cost.

KEYWORDS | Authentication; authorization; BGP; border gateway protocol; integrity; interdomain routing; network security; networks; routing

I. INTRODUCTION

The Internet is a global, decentralized network comprised of many smaller interconnected networks. Networks are largely comprised of end systems, referred to as hosts, and intermediate systems, called routers. Information travels through a network on one of many paths, which are selected through a routing process. Routing protocols communicate reachability information (how to locate other hosts and routers) and ultimately perform path selection. A network under the administrative control of a single organization is called an autonomous system (AS) [1]. The process of routing within an AS is called *intradomain routing*, and routing between ASes is called *interdomain routing*. The dominant interdomain routing protocol on the Internet is the Border Gateway Protocol (BGP) [2]. BGP has been deployed since the commercialization of the Internet, and version 4 of the protocol has been in wide use for over a decade. BGP generally works well in practice, and its operational simplicity and resilience have enabled it to play a fundamental role within the global Internet [3], despite providing no performance or security guarantees.

Unfortunately, the limited guarantees provided by BGP sometimes contribute to serious instabilities and outages. While many routing failures have limited impact and scope, others may lead to significant and widespread damage. One such failure occurred on 25 April 1997, when a misconfigured router maintained by a small service provider in Florida injected incorrect routing information into the global Internet and claimed to have optimal connectivity to all Internet destinations. Because such statements were not validated in any way, they were widely accepted. As a result, most Internet traffic was routed to this small Internet Service Provider (ISP). The traffic overwhelmed the misconfigured and intermediate routers, and effectively crippled the Internet for almost two hours [4]. Several similar incidents have taken place in recent years [5], including a major outage caused by ConEd [6] and an outage for the popular YouTube site (http://www.youtube.com/) caused by Pakistan Telecom [7]. In addition, "spammers" (i.e., people sending spam e-mail) sometimes introduce false information into BGP to enable them to exchange e-mail with mail servers using unallocated IP addresses that are hard to trace [8]. Introducing false information into BGP is also an effective way for an attacker to snoop on traffic en route to a

Manuscript received August 13, 2008. Current version published December 23, 2009.
K. Butler and **P. McDaniel** are with the Department of Computer Science and Engineering, Pennsylvania State University, University Park, PA 16802 USA (e-mail: butler@cse.psu.edu; mcdaniel@cse.psu.edu).
T. R. Farley is with the School of Computing and Informatics, Arizona State University, Tempe, AZ 85287 USA (e-mail: toni@asu.edu).
J. Rexford is with the Department of Computer Science, Princeton University, Princeton, NJ 08540 USA (e-mail: jrex@cs.princeton.edu).

Digital Object Identifier: 10.1109/JPROC.2009.2034031

legitimate destination, impersonate a Web site (e.g., to perform identity theft), or block access to certain sites [9].

These attacks and misconfigurations can cause anything from an inconsequential annoyance to a devastating communications failure. For example, critical applications such as online banking, stock trading, and telemedicine run over the Internet. Significant harm may arise if communication is lost at a crucial moment. As the number of critical applications on the Internet grows, so will the reliance on the underlying network infrastructure to provide reliable and secure services. Consequently, there is great interest in increasing the security of BGP, as it is essentially the glue that holds the disparate parts of the Internet together. For example, the United States government cites BGP security as part of the national strategy to secure cyberspace [10]. In addition, the Internet Engineering Task Force (IETF) working group on Secure Interdomain Routing [11] is investigating these security issues and defining practical solutions. BGP security is also a prominent topic at network operator meetings and mailing lists, such as the North American Network Operators Group (NANOG) [12].

Current research on BGP focuses on exposing and resolving both operational and security concerns. Operational concerns relating to BGP, such as scalability, convergence delay (i.e., the time required for all routers to have a consistent view of the network), routing stability, and performance, have been the subject of much effort. Similarly, much of the contemporary security research has focused on the integrity, confidentiality, authentication, authorization, and validation of BGP messages. These two fields of operational issues and security research are inherently connected. Successes and failures in each domain are informative to both communities.

This paper explores operational practice, standards activity, and ongoing research in interdomain routing security, exposing the similarities and differences in the proposed approaches to building a more secure Internet infrastructure. The next section provides a brief overview of interdomain routing and BGP. Subsequent sections examine today's security practices and longer-term solutions for secure interdomain routing.

II. BORDER GATEWAY PROTOCOL

The Internet consists of tens of thousands of Autonomous Systems (ASes) that use the Border Gateway Protocol (BGP) to exchange information about how to reach blocks of destination IP addresses (called *IP prefixes*). BGP is an incremental protocol—a BGP-speaking router sends an announcement message when a new route is available, and a withdrawal message when a route no longer exists. BGP is also a path-vector protocol, where each AS adds its AS number to the beginning of the AS path before advertising the route to the next AS. Each router selects a single preferred BGP route for each destination prefix and may apply complex policies for selecting a route and deciding whether to advertise the route to a neighboring router in another AS.

In this section, we present an overview of interdomain routing in the Internet and describe how most of BGP's security problems stem from: i) uncertainty about the relationship between IP prefixes and the AS numbers of the ASes who manage them; ii) the use of the Transmission Control Protocol (TCP) as the underlying transport protocol; and iii) the potential to tamper with route announcements in order to subvert BGP routing policy.

A. IP Prefixes and AS Numbers

An IP address is a 32-bit number, typically represented in dotted-decimal notation with a separate integer for each of the four octets.[1] Addresses are assigned to institutions in blocks of contiguous addresses, represented by the first address and a mask length. For example, the prefix 192.0.2.0/24 contains all addresses where the first three octets are 192, 0, and 2—the 256 addresses 192.0.2.0 to 192.0.2.255. Allocating addresses in blocks leads to smaller routing tables and fewer route advertisements, as most routers need only know how to direct traffic toward the block of addresses, rather than storing separate routing information for every IP address. Since prefixes have variable length, one IP prefix may be completely contained within another. For example, a router may have routing information for two prefixes 211.120.0.0/12 and 211.120.132.0/22, where the first prefix completely covers the second one. To decide how to forward a data packet, an IP router identifies the longest prefix that matches the destination IP address. For example, a packet with destination IP address 211.120.132.37 would match 211.120.132.0/22, since this prefix is more specific than 211.120.0.0/12.

Initially, institutions received address assignments directly from the Internet Assigned Numbers Authority (IANA), whose duties are currently performed by the Internet Corporation for Assigned Names and Numbers (ICANN). More recently, IANA began to delegate this responsibility to address registries responsible for different parts of the world. For example, the American Registry for Internet Numbers (ARIN) manages the IP address assignments for North America, whereas the Réseaux IP Européens (RIPE) assigns much of the address space for Europe, the Middle East, and North Africa; the Asia-Pacific Network Information Center (APNIC) assigns IP addresses in Asia and the Pacific Rim, the Latin American and Caribbean Internet Address Registry (LACNIC) distributes address space through the Latin American and Caribbean regions, and the African Internet Numbers Registry (AfriNIC) serves the African region. These regional registries can assign IP addresses directly to organizations or other registries, including national registries and ISPs

---

[1]While IPv4 addresses are predominant and the focus of our discussion, IPv6 addresses, which are 128 bits in length, are also being routed today.

**Fig. 1.** *An example of address delegation from the root (IANA) to regional and national registries.*

that may, in turn, assign smaller portions of the address block to other institutions. Fig. 1 shows an example of address delegation. Here, IANA delegates the large address block 210.0.0.0/7 to APNIC, which delegates 211.120.0.0/12 to the Japan Network Information Center (JPNIC), which in turn assigns 211.120.132.0/22 to Sony. Sony can then perform further delegation based on its organizational setup.

Autonomous Systems are assigned AS numbers (ASNs) in a similar manner, with IANA serving as the ultimate authority for delegating numbers. AS numbers from 1 to 64511 are public and have Internet-wide scope, requiring each number to correspond to a single AS. For example, Sony has been assigned AS number 2527. In contrast, some companies have multiple ASes. For example, AS 701 corresponds to the North American backbone of Verizon Business (formerly UUNET), whereas AS 702 corresponds to Verizon Business's European backbone. Public AS numbers can appear in the AS-path attribute of BGP advertisements. However, many institutions do not need a unique AS number. For example, an Autonomous System may connect to a single upstream network provider (i.e., a

provider closer to the Internet backbone) that bears sole responsibility for providing connectivity to the rest of the Internet. The customer AS may be assigned a private AS number in the range 64512–65535 for communicating via BGP with its provider. The provider's routers would then advertise the BGP routes on behalf of this customer, without including the private AS number in the path. This allows service providers to reuse the same private AS number for their own customers.

The AS that introduces a destination prefix into the global routing system—by advertising the prefix to neighboring ASes—is called the *originating* AS. In the example in Fig. 2(a), AS 6 advertises a BGP route for 12.34.0.0/16 with an AS path of "6" to its upstream provider AS 5, which adds its own AS number to the front of the AS path before sending the BGP advertisement to other neighbors like ASes 4 and 7. However, BGP does not ensure that a BGP-speaking router uses the AS number it has been allocated, or that the AS holds the prefixes it originates. A router can be configured to advertise routes into BGP with any AS number, as long as the neighboring router is configured to accept them. Similarly, a router can originate routes for any destination prefix, including very small address blocks (e.g., 211.120.132.4/30) and address blocks it does not hold. The neighboring router will accept these advertisements unless configured to do otherwise, based on prior knowledge of the acceptable prefixes or prefix lengths. This makes the routing system extremely vulnerable to misconfiguration or malicious attack.

An AS can advertise a prefix from address space unassigned by or belonging to another AS—an action known as *prefix hijacking*. Neighboring ASes receiving this announcement may select this route and direct traffic toward the wrong AS; these ASes may, in turn, advertise the BGP route to their own neighbors. In the example in Fig. 2(b), if malicious AS 1 announces 12.34.0.0/16 and all ASes select shortest-path routes, then ASes 2 and 3 mistakenly choose routes through AS 1 rather than AS 6. If the offending AS simply drops all packets destined to the hijacked addresses, the effect is called a *black hole* and the

(a)   (b)
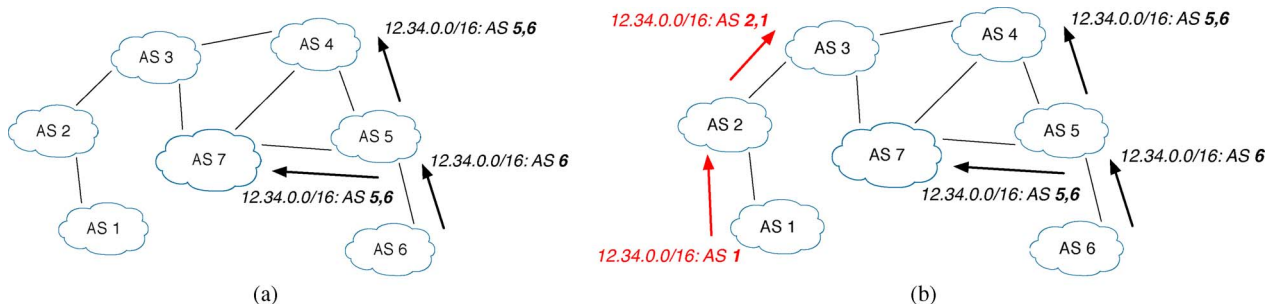
Fig. 2. Announcement of prefix 12.34.0.0/16 originating from the valid AS 6 and from a malicious AS 1. AS 2 and 3 may prefer the malicious advertisement from AS 1 because the path length will be shorter than the valid advertisements from AS 6. *(a) Regular advertisement from AS 6. (b) Malicious advertisement from AS 1.*

Fig. 3. *An example of deaggregation. Because AS 1 advertises a longer prefix for the address block 12.34.128.0/17, it will be preferred over the larger advertised block 12.34.0.0/16 even if it is invalid.*

destinations seem unreachable—at least to the parts of the Internet that believe the bogus BGP routes. If the AS decides to direct the traffic to hosts under its control, the effects can be much more severe. These hosts may *impersonate* the service provided by the legitimate, hijacked destinations; the malicious AS can then analyze the traffic these hosts receive, possibly receiving sensitive information such as passwords and credit-card numbers. In some cases, prefix hijacking can be used to perform an *interception attack*, where the AS inspects the packets (compromising the user's privacy) before forwarding them along to the legitimate destination [9].

To ensure that virtually *all* ASes direct traffic to the wrong place, the offending AS may advertise more-specific prefixes (e.g., 12.34.128.0/17 and 12.34.0.0/17) contained in the original address block. Because of the "longest prefix match" rule, IP routers would always forward packets toward the offending AS rather than the real AS that advertised the larger address block. In the example in Fig. 3, if AS 1 originates 12.34.128.0/17 and 12.34.0.0/17 into BGP, *all* other ASes would forward their traffic toward AS 1. The 1997 routing failure described in Section I is a canonical example of such "deaggregation," as the offending AS misconfigured its routers, deaggregating every prefix in their routing table and advertising the first /24 block of each of these prefixes as their own. The Pakistan Telecom attack on YouTube similarly involved announcing a smaller address block that effectively misdirected all packets meant for the YouTube site to the wrong place, where they were dropped. These were not necessarily malicious attacks, but simply innocent configuration mistakes by the network operators. A well-planned, targeted, malicious attack on BGP could do even more serious harm, be more difficult to detect, or both.

B. Using TCP as the Transport Protocol

A pair of routers exchange BGP announcement and withdrawal messages by establishing a BGP session that runs over an underlying Transmission Control Protocol (TCP) connection. The TCP connection provides the abstraction of a communication channel that reliably delivers an ordered stream of bytes, obviating the need for BGP to provide error correction or retransmission. BGP neighbors often have a direct physical connection at the IP layer. For example, a router in one AS may have a link connecting to a router in another AS, and the BGP session runs over this link. More generally, the two routers may have to communicate through an intermediate device, such as a firewall or another router; in this case, the TCP connection must traverse several IP-layer hops. In addition to having external BGP (eBGP) sessions with other ASes, a router may also have internal BGP (iBGP) sessions with other routers in the same AS. These internal sessions are used to disseminate the BGP routes learned from neighboring domains throughout the AS.

The communication channel between two BGP-speaking routers is vulnerable to attacks. To simplify the discussion of possible attacks, we consider two BGP-speaking routers Alice and Bob, and a malicious third-party, who we call Charlie. Possible attacks include:

Attacks against confidentiality: Two routers communicating over a channel may be assumed to have a modicum of confidentiality; that is, they may expect that messages they send to each other would not be seen by any other party. However, Charlie could *eavesdrop* on the message stream between Alice and Bob, in an attempt to learn policy and routing information from the two parties. While this information is not necessarily sensitive, many service providers have business relationships that can be inferred from the BGP data [13]. Allowing Charlie to infer these business relationships may be highly undesirable to Alice and Bob. These *passive* attacks are not unique to BGP, as they apply to any protocol that uses TCP for the underlying transport of messages without any additional security infrastructure.

Attacks against message integrity: Charlie can become a *man in the middle* between Alice and Bob, and tamper with the BGP messages. For example, Charlie could *insert* forged BGP messages into the message stream. These messages could introduce incorrect information into the routing system or trigger Alice or Bob to abort the session. Excessive messages could also overwhelm Alice and/or Bob, causing the routers to crash. Charlie could also selectively *delete* messages. For example, BGP speakers exchange periodic keep-alive messages to test that they can still communicate; deleting these messages would cause Alice and/or Bob to think the connection is broken, causing them to tear down the BGP session. Charlie could also *modify* the messages between Alice and Bob, leading them to have inconsistent views of the routing information. Finally, Charlie can launch a *replay* attack, where he records messages between Alice and Bob and resends them at a later time. This allows Charlie to re-assert withdrawn routes or withdraw valid ones and force traffic to routes he defines.

Denial-of-service attack: The TCP connection between Alice and Bob may itself be the object of a denial-of-service attack, even from a remote adversary that does not have

direct access to the link(s) between Alice and Bob. TCP uses a three-way handshake (SYN, SYN-ACK, and ACK) to establish the connection between Alice and Bob, and closes the connection with a FIN or RST packet. Charlie could send Bob an RST packet that convinces Bob to close the connection, even though both Alice and Bob want to continue communicating. Alternatively, Charlie could send a large number of SYN packets to Bob without completing the three-way handshake (i.e., without sending the ACK packet). This *SYN flooding* attack [14] would consume Bob's connection state memory, leaving Bob unable to perform any TCP transactions. Bob's neighbors are adversely affected as well because they eventually declare their sessions with Bob to be dead, forcing them to withdraw all of the BGP routes they learned from Bob. After coming back online, Bob announces all of these BGP routes again, forcing the neighbors to switch to new routes and advertise them to their neighbors. This *route flapping* is detrimental to all routers because it consumes processing and bandwidth resources, and also causes repeated disruptions in connectivity.[2]

Denial-of-service attacks may be implemented by attacking the physical infrastructure on which the network itself runs, and such attacks may successfully cause changes in BGP routing. Bellovin and Gansner [15] showed that through *link cutting attacks*, which can be manifested by both physically attacking a link (i.e., the "backhoe attack") and through DoS attacks that effectively swamp a link with traffic, an adversary can effectively force BGP traffic through the ASes of his choice.

In addition, the ability of Charlie to force a BGP session reset can allow the configuration of Alice or Bob to transition into a stable but undesired forwarding state, known as a BGP Wedgie [16]. If these undesired states occur, manual intervention by network operators becomes necessary to change the state. These may require cooperation of network operators across several ASes, as it is often the case that no single group of operators has a sufficiently global view of the network to implement a correct solution.

C. Routing Policy and BGP Route Attributes

ASes are not only bound by physical relationships; they are also bound by business or other organizational relationships. When an AS holder serves as a provider to another organization, there are associated contractual agreements involved. Such agreements are often defined by *service level agreements* (SLAs), which indicate the quality of service that the provider will guarantee, or *peering contracts*, which define where two ASes will connect to each other and what traffic they will carry for

each other. Therefore, for both legal and financial reasons, network operators need to be able to specify routing policies that influence which BGP routes are chosen and which neighbors can direct traffic over these routes [17]. BGP enforces routing policies, such as the ability to forward data only for paying customers, through a number of protocol features. Principal among these is the assignment of attribute values in UPDATE messages. A BGP-speaking router selects a preferred route for each destination prefix from a set of candidate routes by comparing their route attributes. Routing policies, specified in advance by human operators, influence how a route's attributes are set. For example, important BGP route attributes include:

1) **Local preference:** This value is propagated *within* an AS and is used to override shortest-path routing in favor of other policy goals. For example, local preference is commonly used to prefer routes through a paying customer over routes through other neighbors, even if the route through the customer has a longer AS path. This policy would be realized by assigning a high local-preference value (e.g., 100) if a route's next-hop AS corresponds to a customer, and a smaller value (e.g., 90) otherwise. Network operators also use the local-preference attribute to direct traffic toward less-congested connections to neighboring ASes. For example, a route through a low-capacity link to a customer may be assigned a smaller local preference (e.g., 99) than a route through a high-capacity link (e.g., 100).

2) **AS path length:** BGP is called a *path vector* algorithm because each AS adds its own AS number to a path before propagating a route to its neighbors. When multiple routes have the same (maximum) local-preference value, a route with the smallest AS-path length is chosen. When advertising a route to its neighbors, an AS may artificially inflate the length of the AS path to make the route look less attractive to other ASes; in particular, the AS may add its own AS number to the AS path *multiple times*, in a process known as *AS prepending*.

3) **Origin type:** Whether a route was learned internally within the AS versus from a source outside the AS (i.e., through an interior gateway protocol rather than an exterior gateway protocol), or from an unknown or other method of learning the route (e.g., through BGP route redistribution) is the next tie-breaking step in the BGP route-selection process. In practice, an AS may modify the origin-type attribute to influence whether a route is chosen over other routes with the same local preference or AS-path length.

4) **Multi-Exit Discriminator (MED):** Two neighboring ASes may connect to each other at multiple

---

[2]In practice, routers typically employ route-flap damping to penalize unstable BGP routes. If a neighbor continually advertises and withdraws a route for a prefix, the router eventually suppresses the route. This can cause parts of the Internet to lose connectivity to the destination prefix, even though the physical paths exist.

geographic locations; for example, two large ISPs (like AT&T and Sprint) may easily peer at a dozen or more places spread across a country. The MED attribute provides a way for one AS to influence which peering location receives the traffic sent by the neighbor. For example, an ISP may advertise a route with a MED of 0 in New York City and another route with a MED of 1 in San Francisco, to ensure that the neighbor directs all traffic through New York City (and uses the San Francisco route only if the New York route fails). The use of the MED attribute is typically specified in advance as part of the contract between the two ASes; otherwise, an AS has complete freedom to select among the alternate routes based on its own local policy goals.

BGP routers can be configured with route preferences, selective destination reporting (i.e., reporting a destination to some neighbors and not others), and rules concerning path editing [18]. The range of policies the network operators might wish to enforce is almost without bound. Policies configured in a BGP router allow it to filter the routes received from each of its neighbors (import policy), filter the routes advertised to its neighbors (export policy), select routes based on desired criteria, and forward traffic based on those routes [19], [20]. For example, a transit AS (one which allows traffic to pass through that neither originates from, nor terminates at, the AS), may have several neighbors. The BGP policy may be configured to export routes learned from paying customers to all neighbors (to ensure that the rest of the Internet can reach the ISP's customers). Yet, the ISP may not agree to carry traffic from one competitor to another, by refusing to export routes learned from one peer to another. Similarly, a small AS (such as a university campus) with two ISPs would not export routes learned from one provider to another; otherwise, the small AS would be responsible for carrying traffic between the two providers.

Unfortunately, adversaries can easily manipulate how an AS selects routes by sending BGP route announcements with bogus attributes. For example, an AS could forge the AS-path attribute by truncating the AS path (to make a route look shorter, and hence more attractive) or adding additional AS hops at the end (to make a hijacked route look like it was originated by the proper AS). An adversary AS could also remove a particular hop from the AS path to thwart policies in other ASes that try to avoid directing traffic through certain ASes (e.g., ASes known to have bad performance or to filter/modify traffic). In an even more subtle attack, the adversary may add a victim's AS number to the AS path so that, while other ASes would propagate the route throughout the Internet, the victim AS would unintentionally delete the route, thinking it contains a loop. An adversary AS may add numerous AS hops to the AS path, to increase the storage demands on routers in other ASes or even crash a router that does not allocate

sufficient memory to store the long AS path. Additionally, the adversary could attach MED values to routes, even if its neighbor has not agreed to respect MEDs, in the hope of influencing the neighbor's decisions; similarly, an AS could send routes with different origin types at different peering locations to achieve the same goal.

III. BGP SECURITY TODAY

Securing interdomain routing has been a challenge for many years. Seminal work by Perlman [21] showed that a fundamental problem in securing protocols like BGP is that routers may exhibit *Byzantine*, or faulty and possibly malicious, behavior. Consequently, a secure interdomain routing protocol must display *Byzantine robustness*; that is, in the face of malicious or faulty behavior from other hosts, all non-faulty hosts in the system should reach a decision on a particular message's contents within a finite time period (termination). This decision should be the same among all non-faulty hosts (agreement), and the message should be the one sent by the source node (validity).

Existing solutions to date largely provide only some facets of Byzantine robustness. The majority of defenses that have been implemented by ISPs to protect BGP have focused on solutions that can be implemented locally or require only limited interaction with parties outside the local administrative domain. In particular, protection of the underlying TCP connection and defensive filtering of BGP announcements are the most commonly implemented solutions, with some limited deployment of cryptographic protections between routers. However, these solutions are ultimately limited in the protections they can offer against more complex and sophisticated attacks that target BGP itself. Ultimately, a more complete view of which routes are valid is necessary for protecting against this latter class of attacks. In this section, we describe the currently-implemented solutions and levels of protection they provide, starting with an overview of the cryptographic techniques used in many of the current and proposed solutions for improving BGP security.

A. Cryptographic Techniques for BGP Security

Understanding the specific proposals and methods used for protecting BGP necessitates a familiarity with cryptographic techniques that provide the underlying security for these schemes. We provide a short discussion on cryptography used in BGP security, presenting only the concepts necessary for an understanding of current and proposed defenses.

1) Pairwise Keying: Many of the cryptographic mechanisms protecting a pair of parties rely on the existence of a *shared secret key*, often as input for a message authentication code (discussed later). The two parties agree ahead of time, often in an offline manner, on a key to be shared between themselves, and this key is then configured

manually at each end-point. This approach is limited in that maintaining shared secrets between many peer routers concurrently can be difficult; notably, the complexity of pairwise key management is $O(n^2)$ in the number of peers. Moreover, such secrets, if not replaced frequently, are subject to exposure by cryptanalysis and through churn amongst ISP operations personnel.

2) Cryptographic Hash Functions: Also known as digest algorithms, cryptographic hash functions compute a fixed-length hash value from an input text and form the basis for message authentication codes and digital signatures (discussed later). The most common hash functions currently in use are Message-Digest algorithm 5 (MD5) [22] and the Secure Hash Algorithm family, particularly SHA-1 [23]. A hash function is cryptographically sound if it is non-invertible (i.e., it is computationally infeasible to find a preimage of a hash value) and collision resistant (i.e., it is computationally infeasible to find two inputs with same output hash value). For MD5, the output is 128 bits in length. To illustrate infeasibility, consider an attempt to find a message that will map to a particular MD5 digest: with a 128-bit digest, one would require on average 2^{127} messages to find the particular message that mapped to the digest value, or 2^{64} messages to find a message that created a *collision*, a different message that maps to the same digest value.[3] The MD5 digest mechanism requires that a shared secret key be configured manually at each session end-point.

3) Message Authentication Codes: A message authentication code (MAC) is an unforgeable tag appended to a message that provides security by guaranteeing the *integrity* of a message (i.e., proof that the message was not tampered with) and *authenticity* (i.e., only a party with access to a given secret key could have generated the MAC). A MAC is generated by computing a function that takes a secret key and a message of interest as input, and outputting a tag. The party receiving the message who has knowledge of the secret key will be able to compute the same function and verify whether the generated MAC matches the one that was sent with the information. A common method of generating a MAC is the HMAC [24] variant, where a cryptographic hash function is used as the function to generate the MAC.

4) Diffie-Hellman Key Negotiation: Diffie and Hellman [25] created a method of allowing two parties with no prior knowledge of each other to share a secret key. Briefly, the exchange works through the two parties agreeing to use a common prime number and base. Each party then chooses a value unknown to the other party and performs a modular

exponentiation, where the base to the chosen value is computed modulo the agreed upon prime number. This result is passed to the other party, which performs a modular exponentiation using the received result as the base to the exponent they originally chose, and computing the result modulo the originally chosen prime number. Both parties will compute the same result, which may now be used as a shared key. Determining the exponents used to generate the final value is considered equivalent to solving the discrete logarithm problem, which is thought to be hard [26]. Intuitively, this means that seeing past messages provides no insight into how to generate or guess new keys; a general requirement for cryptosystems.

5) Public Key Infrastructure: The cryptographic techniques described to this point rely on a *shared key* between two parties. Because announcements can originate from any of the over 35 000 ASes in the Internet, being able to establish the integrity of these messages through mechanisms such as message authentication codes and digital signatures is necessary, but these rely on the establishment of keys between potentially any AS. Managing these pairwise keys between over 35 000 ASes will quickly become intractable. Key management on a global scale requires *public key cryptography*. As applied to BGP, every AS has a *public key*, distributed freely to any other AS in the Internet, and a *private key*, which is never divulged. Two ASes without *a priori* knowledge of each other can negotiate a key for secure communication with each other (e.g., through a Diffie-Hellman key exchange) if they can find the public key for the AS they wish to communicate with. *Public key infrastructure*, or PKI, provides a framework for assignment and delegation of public keys. The PKI handles requests for public keys originating from other ASes. Keys are distributed in a hierarchical manner. For example, an AS's key may be associated with its organization, which receives its key from a regional registry, which in turn receives its key from IANA (the root of the hierarchy tree as shown in Fig. 1). From that diagram we can see that we can retrieve the public key for SONY by querying IANA initially, which would direct us to APNIC and JPNIC. Currently, such an infrastructure does not exist, but there has been considerable research in the field. Notably, Seo *et al.* [27] explore a PKI for the Secure BGP protocol, discussed in detail in Section IV.

6) Public Key Cryptographic Primitives: Asymmetric, or public-key cryptography, is used extensively in many security solutions. Message confidentiality is ensured through use of *encryption*, where ciphertext is generated using the public key of the message recipient. Only the AS with the correct associated private key will be able to decrypt these messages.[4] The complementary security

---

[3]Less messages are required to find a colliding digest value because of the *birthday paradox*, which shows that for n inputs and k possible outputs that can be generated, if $n > \sqrt{k}$, there is a better than 50% chance that a pair of inputs will map to the same output.

[4]Note that in practice, a public key transaction is used to establish a symmetric encryption key between two parties, as symmetric encryption is orders of magnitude faster than public-key encryption.

parameter to confidentiality is integrity, which provides evidence that a message has not been modified in transit. Integrity is accomplished through the use of *digital signatures*, hashes of messages enciphered with the private key of the AS originating the message. To verify the message, the receiving AS requires the public key of the AS that sent the message, which can be retrieved through a PKI, and compares the hash of the received message it generates with that obtained from the decoded signature. Due to the *non-invertibility* of hash functions, it is virtually computationally infeasible to reverse the hash function and create a message that hashes to the same value. Consequently, one can verify that only the signing AS could have sent the message and that it was not altered during transmission.

7) Certificates and Attestations: The concepts of *certificates* and *attestations* are features of PKIs and as such, appear in several of the comprehensive solutions for BGP security. Attestations are proofs that an entity is authorized to advertise a particular resource, e.g., a given AS is the holder of a certain address prefix. Attestations can include information on who a resource has been delegated to (e.g., a block of addresses from a larger network block is allocated to another AS) and the parent organization that delegated the resource to the attestor (e.g., IANA is the ultimate root for all address allocation), and are signed by the attesting AS or organization. The digital signature ensures the integrity of the attestation, and one can follow the delegation chain, verifying the attestation at each link, back to the source of the original delegation. To verify the attestations, the public key of an AS is required; this information is retrieved through a PKI using certificates. Certificates contain both the public key of the requested AS and a signature attesting to the validity of the certificate, issued by a certification authority, or CA. The CA can be an ISP or a national or regional registry that issues an AS number to the organization, in which case it in turn may have a certificate signed ultimately by a root organization, typically assumed to be IANA. The root certificate is self-signed by IANA in this instance. In a similar manner to attestations, the certificate chain can be verified all the way to the root organization. To return to Fig. 1, the certificate for SONY is signed by JPNIC, which is signed by APNIC and in turn signed by IANA, such that as long as the verifier has IANA's public key, the entire chain of certificates can be verified.

B. Protecting the BGP Session Between a Pair of Routers

Protecting the connection between two BGP-speaking routers relies on both protecting the underlying TCP session and implementing defenses that protect the BGP session itself. Below, we describe methods for protecting pairwise communications between two BGP-speaking routers that provide multiple layers of protection.

1) MD5 Integrity: Recent enhancements to BGP suggest the use of a TCP extension that carries an MD5 digest [22] based MAC. An MD5 keyed digest [24] of the TCP header and BGP data is included in each packet passing between the BGP speakers. The authenticity of the packet data is ensured because the digest could have only been generated by someone who knows the secret key. A number of variants consider hashing all or part of the TCP and BGP data message using one or more keys [28], which addresses many of the problems of spoofing and hijacking inherent to TCP [29], [30].

MD5 authentication can also be used directly with TCP. Early versions of BGP included a similar authentication field which was largely unused. With the addition of MD5 MACing and sequence numbers, TCP can protect the integrity of a message (i.e., it is protected against modification) and against replay attacks. It does not protect the confidentiality of the message because there is no encryption mechanism specified. In addition, this solution requires that a shared secret be manually configured in two routers, which can place a significant operational burden on network administrators.

A recent proposal by the IETF replaces TCP-MD5 with a mechanism known as TCP-Authentication Only (TCP-AO) [31], where the MAC algorithm is not fixed as it is with MD5, but can be one of many and can be changed if found to be weak (i.e., algorithm agility). TCP-AO also provides replay protection and allows for rekeying during a TCP connection without any packet loss, if a mechanism exists to provide new keys. Through the concept of master key tuples (MKTs) in TCP-AO, unique keys can be generated and key management can be automated, which is not possible with TCP-MD5.

2) Session and Message Protection: Smith and Garcia-Luna-Aceves [32], [33] proposed five countermeasures to secure interdomain routing. These countermeasures enhance the BGP protocol by modifying both the session environment and the BGP message attributes. Two countermeasures aim to protect BGP control messages by encrypting all BGP data between peers (using a secret key shared by the peers) and adding sequence numbers to enforce a total ordering on the messages. The other three countermeasures offer protection for UPDATE messages and include the addition of an UPDATE sequence number or timestamp, addition of a new path attribute, PREDECESSOR, that identifies the last AS before the destination AS, and digital signatures (signed by the peer) of all fields in the UPDATE message whose values are fixed.

By providing encryption and authenticated sequence numbers, confidentiality and integrity of BGP may, to a degree, be protected.[5] However, this scheme relies on

[5]The authors' claim that the session encryption provides integrity is technically incorrect: encryption alone does not provide integrity. However, exploiting the vulnerabilities exposed to a lack of integrity of ciphertext is somewhat difficult in this case.

shared keys between peers. As discussed above, managing this number of keys becomes enormously complex as the number of peers scales to all routers in the Internet. Additionally, use of these extensions requires altering BGP, which is seen by many as a prohibitive barrier to adoption. There are hundreds of thousands of routers spanning thousands of organizations on the Internet. Such barriers are cited as motivation for out-of-band solutions such as IRV (discussed in Section IV).

3) Hop Integrity Protocols: Within the context of interdomain routing, *hop integrity* is the property that peers can detect any modification or replay of exchanged information. Gouda *et al.* [34] propose a suite of protocols that also provide security at the IP layer. As with the Smith approach discussed above, sequence numbers and MACs are used to ensure integrity and ordering. Gouda *et al.* extend this approach by suggesting an authenticated Diffie–Hellman style protocol that uses public key certificates to negotiate and refresh the secret keys shared by peers.

4) Generalized TTL Security Mechanism: Originally called the "BGP TTL Security Hack," the Generalized TTL Security Mechanism (GTSM) provides a method for protecting peers from remote attacks [35]. This approach builds on the premise that in the vast majority of BGP peering sessions, the two peers are adjacent to each other. (Multihop BGP sessions, where peers are more than one hop away from each other, are possible but uncommon in practice.) The time-to-live, or TTL, attribute in an IP packet is set to a value that is decremented at every hop. For example, if a packet traverses four hops from source to destination, the TTL decrements by four. Routers using GTSM set the TTL of an IP packet to its maximum value of 255. When a BGP peer receives a packet, it checks the TTL and if this value is lower than 254 (decremented by one), the packet is flagged or discarded outright. This prevents remote attacks which come from more than one hop away, as those packets will have TTLs lower than the threshold value of 254, as shown in Fig. 4.

GTSM weakly defends against attackers who are more than one hop away. It does not defend against subverted peers sending malicious information or other similar insider attacks, and it is less useful in multi-hop scenarios where BGP peers are farther than one hop away from each other. The TTL threshold can be lowered to account for how many hops away the peer is, but there will consequently be no defense against attackers the same number of hops away, as those packets will pass unfiltered. Additionally, if an attacker tunnels an IP packet by encapsulating it within another IP packet to a peer one hop away from the victim, the decapsulated packet, with a TTL set to the maximum value, will be able to evade GTSM. GTSM is simple, low cost, and generally effective against unsophisticated attackers. However, the effectiveness of the solution to mitigate motivated attackers is limited.

Hence, it can be considered protection for "off-path" versus remote attacks.

5) IPsec: Many recent proposals have suggested the use of IPsec as a mechanism for securing the BGP session. IPsec is not specific to BGP, but is a suite of protocols that provide security at the network layer [36], [37]. These protocols define methods for encrypting and authenticating IP headers and payload, and provide key management services for the maintenance of long term sessions. The Internet Key Exchange (IKE) protocol deals with the issues of dynamic negotiation of session keys [38]. The IPsec Authentication Header protocol (AH) [39] and Encapsulating Security Payload (ESP) protocol [40] implement packet-level security with differing guarantees. All of these services work in concert to establish and maintain the secret keys used to guarantee the confidentiality and authenticity of data passed over IP between two end-points. Within BGP, this is typically used to secure the BGP messages passed between peers.

IPsec is often used as the security mechanism for implementing Virtual Private Networks (VPNs) [41]. If properly configured, it provides the desirable security guarantees for peer sessions, e.g., authenticity of data, integrity, message replay prevention, and data confidentiality. IPsec sessions implement security between peers only. Hence, while they address many issues relating session-local vulnerabilities, they do little to address widespread attacks.

IPsec is increasingly becoming the dominant means of deploying secure peer communications, as it is ubiquitous, well understood, and easy to configure; it also forms the basis for the comprehensive BGP security solutions to be presented in Section IV. As shown in Table 1, out of the existing solutions, IPsec provides the most comprehensive

Fig. 4. *The Generalized TTL Security Mechanism (GTSM) in operation. Routers set the TTL on a packet to 255, which is decremented when it reaches a peer. If the router is configured such that no packets with a TTL of less than 254 will be accepted, then remote adversaries attempting to inject malicious information to a router will have their packets dropped.*

Table 1 BGP Peer Session Security Solutions—Requirements (Columns) Relate to the Guarantees Provided for AS to AS Peering Sessions.
*Note That Some Solutions Such as S-BGP Use the ESP-Null Mode; in Such a Configuration, ESP Does Not Provide Confidentiality

	Integrity	Confidentiality	Replay Prevention	DOS Prevention
MD5 Integrity [28]	yes	no	yes	no
Countermeasures [33]	yes	yes	yes	no
HOP Protocol [34]	yes	no	yes	no
GTSM [35]	no	no	no	no
IPsec (AH) [39]	yes	no	yes	yes
IPsec (ESP) [40]	yes	yes*	yes	yes

protection, including some limited protections against denial-of-service attacks (described in more detail in Section III-E). Other proposed solutions, such as the hop integrity protocols and countermeasures by Smith *et al.*, provide a subset of IPsec functionality using specialized protocols. IPsec was not widely available at the time most of these solutions were proposed. Hence, while of historical interest, it is unclear what these protocols offer that IPsec does not already more effectively provide. Solutions such as GTSM and MD5 are currently used because they are easy to implement and low cost. Clearly, these protocols serve as short-term measures, and should not be considered by anyone as long-term solutions to peer session security. Hence, ASes will and should continue to use these inexpensive countermeasures until a strong security service such as IPsec can be deployed in their environment.

C. Defensive Filtering of Suspicious BGP Announcements

Defensive routing policies are used to filter bad and potentially malicious announcements, and to manipulate potentially dangerous attributes of received routes. BGP speakers commonly filter ingress and egress routes based on route policies. Among other things, these policies filter prefixes that are documented special use address (DSUA) prefixes (e.g., loopback addresses), and *bogons* (advertisements of address blocks and AS numbers with no matching allocation data), also known as *martians*. The CIDR report keeps an updated list of bogons [42] which many organizations use to filter BGP route announcements. An AS can also filter announcements containing private AS numbers [43] or exceptionally long AS paths. In addition, routes for small subnets (i.e., smaller than a /24 block of 256 addresses) are often filtered in order to limit the size of the global routing tables.[6] ASes can also impose a hard

limit on the number of prefixes a neighboring router can announce; if the number exceeds the configured maximum, the BGP session to that neighbor is closed (and later restarted). This not only penalizes neighbors that artificially deaggregate the routes they advertise, but also prevents the AS's own routers from crashing due to running out of memory [45].

An AS is best equipped to perform more fine-grained filtering of routes advertised by its own customers, especially "stub AS" customers that do not provide transit service for others. For example, an ISP may filter a customer-learned route if the AS path contains the AS number of another large ISP, under the assumption that the customer mistakenly propagated the route learned from one provider to another. Additionally, ISPs often filter routes from customers for prefixes the customer does not own. In fact, if all ISPs filtered customer routes accurately, the global routing system would be much more secure; unfortunately, many ASes do not, and it is much more difficult for an ISP to identify invalid routes that originated several AS hops away. In addition, creating and maintaining filter lists becomes more challenging when the customer has a large, and possibly changing, set of prefixes, or serves downstream customers of its own. In addition to the logistical challenges of ensuring the lists are always accurate, the underlying router equipment may impose limits on the length of a filter list.

An AS may also "rewrite" any BGP attributes its neighbors should not be setting, as a sort of "defensive programming." For example, if an AS has not agreed to accept MEDs from a neighbor, the router could be configured to assign a MED value of "0" to all routes learned from that neighbor. Similarly, the routers could be configured to set the origin-type attribute to a single value for all routes. Some neighbors may have an agreement to tag BGP routes with a "community" attribute to control how the receiving AS should treat the route. For example, an AS could include a community attribute in the route announcement to instruct the receiving AS to prepend extra hops in the AS path, assign a lower local preference (i.e., to treat the route as a backup path), or filter the route when exporting to other ASes. If the neighbors do not have

---

[6]In fact, ASes can filter even more aggressively based on guidelines from the Regional Internet Registries (RIRs) about the minimum appropriate address-block sizes in different parts of the IP address space [44]. However, ISPs are sometimes reluctant to filter too aggressively, because of the risk of "blackholing" the traffic their customers send to the affected destinations.

an agreement to respect the community tags, the receiving AS may filter any routes containing unexpected community values, or strip the offending community attribute from the route, to prevent the neighbor from controlling how routes are selected and exported.

A policy of careful ingress and egress filtering greatly aids in maintaining security for both the local AS and its neighbors, and is the most widely deployed and effective BGP security measure. Filtering, however, is not a replacement for a strong security architecture. The filtering rules are fundamentally limited by the heuristics used, and can only remove announcements which are overtly bad. BGP routing and filtering policies and their ramifications are described in more detail by Casear and Rexford [17], while Nordström and Dovrolis discuss filtering in the context of BGP attacks [46]. In some cases, static filtering rules are not sufficient.[7] Detecting invalid route announcements is even more challenging when the offending AS is several hops away. Today's Internet is still quite vulnerable to attacks launched by ASes connected to ISPs that do not apply "best common practices" for filtering routes, or by an adversary who has compromised a router in the ISP's network.

D. Routing Registries

Despite the benefits of protective route filtering, detecting and disregarding bogus BGP routes is more challenging when the erroneous information stems from a misconfiguration or an attack several AS hops away. Having a shared, global view of "correct" routing information would make it much easier to detect invalid routes. An accurate *routing registry* [48] of, for example, prefix ownership, AS-level connectivity, and routing policies would enable security-conscious ASes to detect and discard invalid routes. ASes using a registry service insert details of their policy and topological information into the repository for other ASes to query. External applications query this data to validate received routes and policy. Registries may also be used by organizations constructing route filters. For example, an ISP's customers may register their routes in a route registry, and the ISP will use this information to construct a filter such that the only routes that are valid, and hence not filtered, are those customer routes in the registry. Additionally, valid registry information may be used to assist a transit provider in determining what filtering to perform on BGP announcements received from its neighbors.

However, to use a registry, one must first be assured that the registry itself is secure, complete, and accurate; without correct information, the route filters generated will not be accurate. Blunk *et al.* [49] propose an authentication and authorization model for providing data integrity in routing policy systems. One drawback of the registry model is that corporations often consider their routing policies and topological information to be proprietary (and are thus reluctant to share it), though measurement tools such as Rocketfuel [13] provide relatively accurate maps of an ISP's internal topology, and algorithms exist for inferring the business relationships between pairs of neighboring ASes [50]–[52]. The community-supported registry approach is also limited in that the registry itself is often untrusted; a malicious registry can manipulate the route information at will. Information in routing registries also tends to become less accurate over time because of a lack of clear incentives for organizations to maintain their information [53].

Increasingly, however, the value of maintaining and securing routing registries is gaining greater appreciation, because of the critical role they play in many proposals to secure interdomain routing. For example, accepting routing announcements from a remote, unknown source requires a level of trust in that remote system. There is no currently-practised method for determining that information received from an unknown AS is true or valid. The best immediate solution to alleviate these concerns is the implementation of authoritative registries. For example, ARIN may be queried for ownership information of an address block. However, this information is not updated with any frequency, and many of the address delegations have changed since the original ownership block was issued. Organizations may have folded into bankruptcy or merged into other companies, and hence the ultimate ownership for the space is often unknown. By systematically verifying all aspects of the address space that the regional registries delegate, there can be some degree of confidence that route advertisements and the ASes that advertise them are truly authentic.

There is a need for both authenticated registries, which can store public keys for the organizations that receive AS numbers and address space from them, and an infrastructure that enables this information to be easily found (i.e., a PKI). Many of the current proposals for securing BGP rely in large part on the implementation of routing registries and a PKI. Efforts are therefore being made to clean the existing routing registries of spurious information and to make them more complete. In addition, APNIC is currently exploring the creation of a certificate repository based on registry details holding certificates, Certificate Revocation Lists (CRLs), and related signing objects such as route origin authorizations [54] (but not routing policy information), which would form the basis for a PKI [55]. ARIN, RIPE, and LACNIC began offering these services on a trial basis in 2009. Still, creating and maintaining a

---

[7]For example, many peering contracts require a neighbor to advertise a route (for a given prefix) with the same AS-path length at all peering locations. This so-called "consistent export" requirement cannot be enforced by static filtering rules applied at individual routers. For example, the neighbor should not advertise a prefix in one location but not in another, or announce routes with different AS-path lengths (for the same prefix) at different locations. Static filtering rules cannot detect when a neighbor violates this kind of peering requirement. Instead, an AS must monitor the routes received from each neighbor to perform an AS-wide check for these kinds of violations [47].

complete, accurate registry—especially when such a registry does not already exist—is quite challenging in practice.

E. Securing Router Management

Because BGP is dependent on the underlying TCP transport protocol, which in turn is dependent on lower layers to provide framing and propagation of the information itself, it is critical to secure the physical router infrastructure, as attacks targeting lower layers of the networking stack may also affect BGP. Most notably, the *link cutting* attack of Bellovin and Gansner, described earlier, shows that by causing a physical cut in the fiber transporting network traffic (the so-called "backhoe attack") or flooding links with traffic to render them inoperable, an adversary can force data traffic to be rerouted in such a way that it passes nodes controlled by the attacker [15]. If adversaries can access the management interface of a router, for example, they can turn down interfaces or spoof BGP NOTIFICATION messages, causing the BGP session between two routing peers to be terminated.

At the physical security level, common practices include protecting physical access to data centers and network points of presence where routers are housed. Often, networking equipment can be remotely accessed and managed through the Simple Network Management Protocol (SNMP) [56] or through a remote connection to the Command Line Interface (CLI). Securing SNMP management interfaces is of critical importance, to prevent attackers from being able to remotely access and modify critical operational parameters within the equipment. Most router configuration happens through the CLI, so protecting access to this is critically important, by securing access to elevated privilege levels on routers and ensuring only secure communication to the router (e.g., through *ssh* or a VPN) is possible from remote locations. In addition, designing the network infrastructure in a way that ensures robustness is similarly important. For example, ensuring fiber diversity by having multiple fibers not laid in the same conduit provides some assurance against a fiber cut. Best common practices for defending against these and related threats are described in RFC 4778 [57].

Protocols that preserve message integrity also effectively prevent some classes of denial-of-service attacks. For example, remotely resetting a TCP connection or forcibly closing a BGP session becomes considerably more difficult when sequence numbers must be guessed and, more importantly, when digests relying on shared secrets are used. Distributed denial-of-service attacks [58] are certainly harmful to BGP operation, as flooding a link could cause timers to expire and information not to arrive. Some protocols, such as IPsec, provide limited forms of DOS prevention, but none adequately address flooding attacks. One method of defending against these attacks is to prevent the router from having to perform more processing than necessary, by segregating incoming traffic into multiple queues based on priority [59]. Messages that affect the routing process (e.g., BGP UPDATE, WITHDRAW, and NOTIFICATION messages, or other messages sent to TCP port 179—meant for BGP traffic—and addressed to the router's loopback address) would be placed in a higher-priority queue that has increased access to a router's processor, with other traffic placed in a lower-priority queue that can be processed when resources are available.

IV. BGP SECURITY SOLUTIONS

The currently-implemented security solutions that consider protection of BGP are limited in their effectiveness. Finding solutions that comprehensively defend BGP against attack is an active area of research, and we examine the numerous proposals that consider many facets of the problem. Because of the variety of issues involved, the different methodologies employed for proposed solutions, and the number of new proposals being made, a canonical categorization of solutions is difficult to achieve. We have structured our examination of security solutions as follows: we start by looking at full-scale architectures that provide origin and topology authentication, then proceed to investigate solutions that improve on or consider different facets of issues brought up in the architectural solutions. These include solutions that focus on reducing computation overhead, providing alternatives to public-key infrastructures, or considering incrementally deployable alternatives based on anomaly detection. As the field matures, more natural taxonomies may become apparent.[8]

A. BGP Security Architectures

Recent efforts within the standards bodies and in the research community have attempted to provide comprehensive architectures for BGP security. Each architecture provides an explicit threat model and suite of security services. We focus on the three most comprehensive approaches to BGP security in terms of the increasing flexibility afforded to the user: S-BGP, soBGP, and IRV. As the following sections detail, trade-offs are made by each protocol in terms of security versus deployability. We start our discussion with S-BGP.

1) S-BGP: Secure BGP (S-BGP) was the first comprehensive routing security solution targeted specifically to BGP [61]. Important elements of S-BGP, notably the notions of the PKI, have been adopted by the SIDR working group and regional registries.

S-BGP implements security by validating path attributes in BGP UPDATE messages passed between ASes through the use of digital signatures and associated public key certificates. Early work in S-BGP called for a pair of PKIs used to delegate address space and AS numbers, and

---

[8]A recent survey [60] organized BGP security solutions in terms of their use of cryptography, databases, overlay protocols, penalties, and data-plane testing.

to associate particular network elements with their parent ASes, while later work collapsed this to one hierarchy [27]. The PKI is used to authenticate address allocations through a hierarchy stretching from organizations to the providers and regional registries allocating them address space, ultimately leading to IANA (the ultimate authority for address allocation). The other functionality provided by the PKI is binding AS numbers to organizations and organizations to routers in their network, through issuance of certificates. For example, an organization's AS number is bound to a public key through a certificate. Statements made by the AS are signed using the associated private key. An entity receiving the signed data verifies it came from the AS using the certificate.

All information exchanged in S-BGP is validated using the certificates in the PKI. Address ownership, peer AS identity, path vectors, policy attributes, and control messages are all signed by the organizations or devices that create them. Because this allows receivers of the data to unambiguously authenticate the routing information, they can detect and remove forged data. However, because of the amount of data and number of possible signers, validation can be costly [62]. These and similar results have raised concerns about the feasibility of S-BGP in the Internet, and led many to seek alternative solutions.

Attestations are digitally signed statements used to assert the authenticity of prefix ownership and advertised routes. *Address attestations* claim the right to originate a prefix, and are signed and distributed *out-of-band*. An out-of-band mechanism does not directly use the BGP protocol to transmit information, instead using some external interface or service to communicate relevant data. Each address attestation is a signed statement of delegation of address space from one organization or AS to another. The right to originate a prefix is checked through the validation of a *delegation chain* from IANA to the advertising AS.

Route attestations are distributed within S-BGP in a modified BGP UPDATE message as a new attribute. To simplify, a route attestation is signed by each AS as it traverses the network. All ASes on the path sign previously attached signatures (i.e., the signatures are nested). Hence, the validator can validate not only the path, but also that *a)* the ASes were traversed in the order indicated by the path, and *b)* no intermediate ASes were added or removed by an adversary. Fig. 5 shows a simplified use of route attestations as they propagate between routers.

While S-BGP proposes the most comprehensive security guarantees of all proposals by providing full authentication of origins and the paths to destinations, there are significant barriers that hamper its adoption. A study on S-BGP deployment issues suggests that the added overhead of S-BGP countermeasures is equivalent to the CPU and memory provided by a desktop PC [63]. Thus, the hardware requirement is ostensibly minimal, although concerns have been raised over the use of time-averaged statistics. In addition, assessments of S-BGP through simulation [64]

shows that path convergence times would increase by as much as double through adoption of S-BGP, although optimizations to the protocol, such as only validating paths when they are selected as preferred, may reduce these convergence times. The substantial storage requirements for route attestations have also been noted [63].

2) Secure Origin BGP: Secure origin BGP (soBGP) seeks flexibility by allowing administrators to trade off security and protocol overhead, depending on how it is configured. In a similar manner to S-BGP, soBGP defines a PKI for authenticating and authorizing entities and organizations. The PKI manages three types of certificates. The first certificate type binds a public key to each soBGP-speaking router. A second certificate type provides details on policy, including the configured protocol parameters and local network topology. This information is stored by the soBGP router receiving the certificate, which uses the information to construct a topology database reflecting the router's view of the network. A third certificate is similar to S-BGP's address attestations in that it embodies address ownership or delegation. All information pertaining to security is transmitted in soBGP between peers via a SECURITY message, a new message type in BGP introduced by soBGP. Thus, in contrast to the out-of-band method of distributing address attestations in S-BGP, the certificates that provide origin authentication are distributed in-band in soBGP, though an out-of-band mechanism for distributing certificates binding keys to routers and topology is proposed.

soBGP routers use a topology database to validate received routes. Each AS signs and distributes its local topology (i.e., its peers) through the topology certificate to form a global database and corresponding static topology graph, of which each soBGP router should have a consistent view. The database is used to verify received routes: any UPDATE with a path that violates the AS


Fig. 5. Route attestations in S-BGP. As UPDATE messages are passed between peers, the receiving peer signs the received message before passing it to another neighbor. The result is an "onion-style" attestation that contains signatures from all routers along the path.

topology is demonstrably bad and dropped. The major difference between the approach taken by soBGP for path authentication and the one taken by S-BGP is that in S-BGP, route attestations are dynamic: they are sent with every BGP UPDATE message and the recipient of the routing information has a real-time view of the path taken by the message. By contrast, the topology graph and corresponding database used by soBGP is fundamentally static, as the topology will only change when a new policy certificate is issued; thus, a new topology may not be reflected when an UPDATE is received and the path it took may be different from the one reflected in the peer's topology database. Additional infrastructure is required to ensure that the topology updates are synchronized across all ASes. Moreover, forged paths that are "plausible," i.e., consistent with the routing topology but not actual routes, are accepted, as the soBGP topology represents all potential routes that might be advertised.

To avoid the computational overhead of validating signatures, soBGP authenticates long-term structural routing elements (such as organization relationships, address ownership, and topology) prior to participating in BGP. Authenticated data is signed, validated, and stored at the routers prior to the establishment of the BGP session, and thus their validation does not introduce significant run-time costs. Transient elements (such as paths) are locally checked for correctness, rather than validated through the PKI, e.g., adjacent ASes in the path must be reflected in the topology database; because of this, soBGP may not guard against changes to these transient elements such as mid-path disruptions, which would not be reflected in the static topology database [65].

The soBGP platform provides several deployment options [66]. One option, for example, allows the operator to choose whether to verify routes before accepting them into the routing table (placing a premium on security) or to accept routes and then verify their authenticity (placing a premium on convergence time). Another example is the option of whether to verify a route using the topology graph, or only the first hop after the origin, or to refrain from validation altogether. These options give soBGP a greater ease of deployment than S-BGP, but the number of options could introduce interoperability challenges [67]. In addition, the certificates used in soBGP are non-standard compared to the IETF PKIX certificates used in S-BGP.

3) Interdomain Route Validation: The Interdomain Route Validation (IRV) service is a receiver-driven protocol and associated architecture [68], and is the least centralized of the comprehensive solutions for securing BGP. Unlike S-BGP, IRV's operation is independent of the routing protocol. Every AS in IRV contains an IRV server. Upon reception of an UPDATE message, a receiving BGP speaker will appeal to its local IRV server for an indication of whether the received information is correct (see Fig. 6). The local IRV server determines correctness by directly

Fig. 6. *ASes running the IRV protocol query the appropriate authorities for validation of received routing data. IRV validators are independent of routers within an AS.*

querying the IRV server in the relevant AS for validation of the route information. Where validation from multiple ASes is needed, i.e., to validate a path involving multiple ASes, collections of IRV servers are queried.

The key idea of IRV is that each data item can be validated by directly querying the AS from whence it came, removing the computational and storage costs from the critical path of routing. Validation of path information is discretionary; that is, the algorithm for determining when and how an UPDATE message should be validated is chosen by each AS. The IRV server can query every AS along the path of a given update, or choose to only query a subset of the ASes based on previous associations (e.g., ASes known to provide trusted information may not be queried). Stronger guarantees can be achieved if every update is fully validated, while better performance can be maintained if the updates are checked only periodically or partially, and queries made when the results appear suspicious. Caching previous queries can also improve performance, while storing received route advertisements and withdrawals can allow for debugging and failure detection. A BGP speaker decides which data to trust, which to ignore, and which to validate via an IRV query based on local policy. IRV servers are similar to routing registries, but manage information only from the parent AS. The IRV approach may have better success than registries because the AS retains control over the data, and hence is more apt to keep it fresh, accurate, and available; however, it is reliant on the AS making accurate assertions and the IRV server not being misconfigured.

Where available, a secure underlying network layer (e.g., IPsec) or transport layer (e.g., TLS [69]) can be used to secure the communication between IRV servers (i.e., to ensure the authenticity, integrity, and confidentiality of queries and results). IRV servers can tailor responses to queries based on the requesting entity. This allows the IRV to perform access control over the routing data which

is useful in limiting the exposure of sensitive data such as policy and peering relationships.

The central limitation of IRV is that it needs a functioning network to be useful: a client indirectly uses the network to communicate with the foreign AS to query the appropriate AS IRV server. This presents problems both in bootstrapping the process and in recovering from outages. Solutions to these problems include optimistic routing (e.g., using received routes immediately and validating where possible), AS collaboration (e.g., exchanging routing data via gossip-style protocols [70]), and using static routes to IRV servers.

B. Experimental Systems

Numerous works have proposed ways to address some (but usually not all) of the security challenges inherent in interdomain routing security. Some focus on more formal properties of routing, while others explore the application of novel cryptographic structures that provide strong security guarantees. This section describes these proposals.

1) Reducing Computational Overhead: The following solutions base their security on facets of the S-BGP schemes, either with regards to address or route attestations. Each solution offers a different method of reducing the computational costs associated with attestations, while providing a similar level of security for either origin or path authentication to S-BGP itself, often by devising more efficient cryptographic proof systems.

Origin Authentication (OA) is a method of validating address ownership such that prefix hijacking and related attacks are not possible. One effort directly investigates OA by examining the design and application of OA services [71]. The semantics of address delegation are formalized, and various cryptographic structures for asserting the address block ownership and delegation are explored. In particular, the authors study cryptographic proof structures [72], [73] for carrying delegation attestations (i.e., cryptographic proofs of delegation). To simplify, a cryptographic proof structure is a structure for asserting the validity of a set of statements. The authors approximate the real IP address delegation hierarchy exhibited on the Internet by extracting the nested announcements made within the protocol. They find that the delegations are stable over time, making them ideally suited to a class of proof structures based on Merkle hash trees [72]. A simulation shows that on-line origin authentication is possible using this construction, which was previously thought to have been too computationally expensive to be feasible.

In an effort to mitigate the costs of path authentication, Hu *et al.* [74] propose a solution that uses traditional secret key cryptography to authenticate received path vectors. In their solution, each AS on an UPDATE's path shares a secret key with a previously identified validator. The originating AS computes a MAC using a shared key over a concatenation of an initial authenticator value (e.g., 0), the

path, and the fields that do not change (e.g., ORIGIN attribute, Network Layer Reachability Information (NLRI), etc.). The MAC is included in the UPDATE and propagated using BGP. Each of the subsequent ASes perform the same operation but use the received MAC as the authenticator value. This ensures that each subsequent MAC covers not only received information, but also the authenticator value of the preceding hop. Upon receiving a MAC, the destination can recursively validate all MACs using the known secret keys. In essence, this is symmetric key equivalent to the recursive signatures specified in S-BGP, where MACs are used instead of digital signatures.

Hu *et al.* extend their work in path authentication with the *Secure Path Vector* protocol (SPV) [75]. SPV implements path validation using a string of one-time signatures [76], [77] generated from a single root value. Also known as off-line signatures, one-time signatures allow the signer to perform the heavyweight cryptographic operations prior to use, making the later signing operation faster. SPV extends this approach to allow a single off-line signature to generate potentially many signatures. To simplify, in SPV, the originator of a prefix establishes a single root value used to seed the generation of one-time signature structures for each hop in the PATH. Signatures and signing material (to be used by the next hop) are forwarded to each hop in the route propagation. Receivers of the route use an initiator-generated initial validation token to verify the one-time signatures, and ultimately the path. The operation of SPV is lightweight, as hashing is used as the primary cryptographic mechanism. However, this efficiency comes at a cost; SPV is a complex protocol involving the manipulation and communication of a significant amount of state information. More generally, however, the security of SPV is in some cases based on probabilistic arguments. In particular, the authors argue that reduced exposure (in time) to forgery vulnerabilities is sufficient to mitigate attacks. While this may be acceptable for some constrained environments, it is unclear whether such arguments will be acceptable in the larger Internet. Raghavan *et al.* [78] show that over 60% of ASes are capable of forging routes in SPV with high probability, and argue that SPV is vulnerable to collusion and eavesdropping attacks. They further argue that constant-time signatures do not provide the requisite security to protect against path modification, but that signature schemes such as ESIGN [79] may provide the desired efficiency and security.

Another method of performing path authentication is suggested by Zhao *et al.* [80]. This scheme suggests the use of *signature amortization* [64], where any BGP UPDATE messages sent to the same group of peers requires only one signature for the group, rather than n signatures for n peers. The method aggregates UPDATE messages in the output buffers of a router, and builds a Merkle hash tree for all unsigned messages so that they are collectively signed with only one signature operation. Additionally, the

scheme uses previous work in adapting *aggregate signatures* to BGP [81]. Aggregate signatures allow for multiple signatures, each having been signed on a different message by a different user, to be aggregated into one signature. While signature aggregation can decrease the computational overhead of signatures, it introduces increased memory requirements. Conversely, while aggregate signatures do introduce a small computational overhead, they are space-efficient. By selecting various parameters to optimize time and space complexity, the optimal solution presented displays much faster convergence times than S-BGP with memory requirements cut by over two-thirds.

An alternate method of amortizing the costs of computation is based on considering the *reference locality* of BGP announcements [82]. The authors in this work base their cryptographic constructions on BGP updates retrieved from the Route Views data archive [83] and notice that paths are generally stable, and the number of new paths grows fairly slowly. Leveraging these facts, based on analysis of collected routing data, the authors suggest alternative path authentication mechanisms to S-BGP route attestations that maintain a similar level of security while substantially reducing the number of signature validations required. The cost in this approach is a commensurate increase in the bandwidth requirements because of the large cryptographic proof systems that are distributed. The authors claim that the constructions proposed are compatible with other solutions and could benefit from space-saving measures like the aggregate signatures proposed by Zhao *et al.* [80].

2) Alternatives to PKI: Prior to the creation of BGP version 4, Kumar and Crowcroft [84] provide an analysis of threats to interdomain routing and describe security mechanisms used in the proposed IDRP protocol [85]. Designed as a superset of BGP and EGP, IDRP is an interdomain routing path vector protocol. The protocol uses an encrypted checksum transmitted with all routing messages sent between routers. The checksum authenticates the message and is encrypted based on an algorithm agreed upon by the two routers. Additionally, authenticated timestamps and sequence numbers are provided as anti-replay mechanisms. The authors assert that malicious entities masquerading as sources will be unsuccessful in a hop-by-hop routing protocol; however, this assertion does not take prefix hijacking into consideration. The authors further assert that link level encryption is impractical due to computation cost, as is digitally signing every routing packet. While largely true at the time the authors designed the protocol (1993), this is clearly no longer the case. IDRP failed to catch on and later advances made cryptographic operations feasible. Hence, while this proposal highlights important requirements for routing security, it is not appropriate for current networks.

The *Pretty Secure BGP* (psBGP) [86] system introduces an address origin authentication service within a larger comprehensive architecture for BGP security similar to that suggested for S-BGP [27]. The central philosophy of this work is that while ASes can be managed within a PKI (because there are relatively few and the list is stable), it is not possible to manage addresses through a centralized PKI, such as those promoted by previous systems. Each AS also rates every other AS with a value that represents the amount of confidence in the trustworthiness of the foreign AS. Origin authentication is implemented in a decentralized system in which each AS creates a prefix assertion list (PAL). The PAL contains address ownership assertions of the local ASes and its peers. An origin claim is validated by checking the consistency between the PALs of peers around the advertising origin. In this way, psBGP provides a weak form of origin authentication: *any* AS can bear witness to the validity of an origin claim.[9] The assumption that no ASes will collude may be difficult to support in the general Internet. Moreover, psBGP requires that an AS place its trust in alien ASes to regulate IP addresses, most of which possess no existing relationships or often knowledge of each other. Path authentication is performed using S-BGP style signatures in combination with the rating mechanism, which in practice allows the AS to decide whether to validate all signatures or a subset, making the validation procedure more lightweight. In this manner, path authentication happens in a manner similar to IRV, with the exception being that IRV queries happen out of band while psBGP requires a change to the BGP update message. Note that a centralized PKI is necessary in psBGP for authorization of AS numbers. In addition, the deployment of PKIs by APNIC and the trial deployments by ARIN and other regional registries challenges the notion that address management is not possible through a centralized PKI.

3) Detecting and Mitigating Anomalies: The following solutions often share in common, with solutions from the previous subsection, that they are designed to be used without a PKI. They have the additional feature, however, that they are primarily based on detecting anomalies in routing or the surrounding infrastructure, and use this information to mitigate threats to routing.

An IP prefix should generally only be originated by a single AS [1]. A multiple origin AS (MOAS) conflict occurs when a prefix is simultaneously originated by more than one AS. Such events can legitimately occur in the natural course of operation where, for example, a multi-homed AS transitions between preferred routes. In some cases, however, these MOAS conflicts directly indicate prefix hijacking. A recent study of MOAS conflicts shows that potential causes included prefixes associated with exchange point addresses (which link ASes), multi-homing without BGP or with private AS numbers, and faulty

---

[9]The authors consider other modes in which *k-out-of-n* peers asserting validity are required for the origin to be accepted. However, this is only useful in weeding out highly connected colluding pairs.

configurations [87]. A proposed enhancement to BGP uses community attributes [88] to distinguish between valid and invalid MOAS conflicts [89] in response to these operational oddities. A list of ASes authorized to announce a given prefix is appended to the community attribute. This list can then be used to determine if a MOAS conflict is valid. Because the community attribute is optional and transitive, routers can drop this information without causing an error. Because they are not authenticated, the announcements can be forged or altered by malicious routers. However, the authors suggest that forged routes can be detected by flagging prefixes received with multiple, conflicting AS lists. An application of this idea is a proposal to employ path filtering based on the heuristics, such as those used for MOAS detection, to protect BGP routes to top-level DNS servers from modification, because of the importance of DNS to the network infrastructure [90]. This is possible because routes to popular destinations are found to be stable, and the DNS is highly redundant, with top-level servers distributed in both number and geography.

Kruegel *et al.* [91] consider the use of intrusion detection to identify forged origin announcements, and propose several metrics used to identify bogus announcements (e.g., strange aggregation and tracking of historical associations between prefixes and ASes). One interesting aspect of this work is its dependence on operational issues: the detection criteria are not derived from the BGP specification, but arise from the evaluation of common configurations and AS behavior. In particular, the method observes ownership over time. Any departure from normal ownership behavior (a new AS begins to announce the address, or a new MOAS occurs) is considered to be malicious and is flagged. The results show the number of incorrect alerts is relatively small, on the order of 20 per day compared to over 5 million UPDATE messages processed per day. However, the prefix ownership lists are pre-computed and not dynamic in nature, requiring rebuilding of the network model if a topology change occurs in the network.

A further extension to the work in MOAS detection is the *Prefix Hijacking Alert System* (PHAS) [92], which builds on the concept of prefix ownership. PHAS is predicated on the notion that a prefix owner is the only entity that can differentiate between real routing changes and those that take place as a result of a prefix hijacking attack. To that end, routing updates from Route Views and RIPE repositories are examined, and if there are changes to the originator of a route, the owner of that prefix is notified through email, optimally set up along multiple paths in case the common path has been hijacked. The system is incrementally deployable in that to join the system, a prefix owner need only register with the PHAS server; however, this server is also a single point of failure in the system, and if it is compromised, it could send out numerous false alarms to prefix owners. Additionally, the

system relies on the validity of entities registering their prefixes; there is no protection against an adversary making a false registration. This situation may be ameliorated if authenticated, secure registries are available. To this end, route origin authorizations (ROAs) held in the PKIs deployed by regional registries provide an effective mechanism for resolving MOAS conflicts by providing validation of route origination [93].

Another recently-proposed alerting system is *Pretty Good BGP* (PGBGP) [94]. The key insight in this work is that misconfigurations and prefix hijacking attacks could be mitigated if routers exercise a certain amount of judgement with the routes that they adopt into their routing tables. With PGBGP, an amount of state is maintained through historical routing data to determine what routes to prefixes should be considered normal. When incoming routes are received that do not adhere to these origins, they are flagged as suspicious for 24 hours, using the data from Mahajan *et al.* [95] that shows most misconfigurations and hijack attempts last for less than this amount of time. The routes are avoided while they are suspicious unless there are no suitable alternative routes. The results of this work show that this solution may often protect ASes against hijacking attacks, with some important caveats. An administrator deploying this solution must be cognizant of their business relationships with providers and customers and ensure that events such as provider changes (which result in new paths to destinations) are accounted for so that convergence is not affected; additionally, sufficiently equipped adversaries can engineer the set of routes the system is forced to accept, in a routing equivalent of the link-cutting attack by Bellovin and Gansner [15].

Hu and Mao examined prefix hijacking in greater detail and provided a mechanism for detecting prefix hijacking attacks in real time [96]. Their solution is based on fingerprinting techniques for networks and hosts. A number of criteria, including the operating systems of machines within a given prefix, and the identifier field within IP packets, TCP and ICMP timestamps, are used to characterize a particular network prefix, with information collected by probes sent to various hosts within the network of interest. If there are conflicting origin ASes advertised, which is potential evidence of a prefix hijacking attack, the collected fingerprints are compared against probes sent to all origins. Differentiation between fingerprints will provide evidence that updates have been received from different originating machines, and that a newly-advertised prefix with sufficiently different characteristics is not the original network advertising a new path, but rather an adversary attempting to hijack the prefix. This approach relies on a real-time BGP UPDATE monitor, which sends differentiating probes if prefixes are advertised from multiple locations. The availability of the monitor is critical as, if updates are delayed, the ability to collect measures, such as probing and subsequent decision making, will be compromised.

Subsequent work investigates how to optimally place route monitors within the Internet to maximize prefix hijacking detection coverage [97].

The *Whisper* protocol [98] is designed to validate the initial source of path information. The protocol does not provide explicit route authentication. Rather, it seeks to alert network administrators of potential routing inconsistencies. In its weakest form, a hash chain is used in a similar fashion to the cumulative authentication mechanism described by Hu *et al.* [74]. A random value is initially assigned to each prefix by the originator. The value is repeatedly hashed at each hop as it is propagated from AS to AS. Received paths are validated by receiving routers by comparing received hash values; if the hash values are the same, then they must have come from the same source (because they represent the same repeated application of the hash function). Stronger protocols are proposed that increase security by making the initial value more difficult to guess, using heavyweight modular exponentiation. One variant uses a construction similar to RSA [99],[10] where a random initial value is exponentiated (modulo a prime group) by the AS numbers of the ASes a route traverses. Because of the mathematical properties of the prime group, the intermediate AS values can be factored out and the result unambiguously associated with a single initial value. Another variant, using a series of hash constructions, is complicated by the fact that only the route originator can verify the route because of the non-invertibility of secure hash functions. Thus, the recipient would have to query the originator as to the veracity of the route, which is often outside of the purview of the originator's knowledge.

C. Factors Complicating Adoption of Security Solutions

BGP security is complicated by operational considerations. Interdomain routing is stressed by the continuous growth of the Internet. Around 40 000 AS numbers have already been allocated to the Regional Internet Registries, and more than 35 000 of these AS numbers have been allocated to individual institutions, with over 32 000 currently being routed. The growth in the number of ASes and the increasingly rich connectivity between them contributes to the number of routing update messages a router receives, thus adding to routing table growth, which in turn leads to scalability issues. The graph in Fig. 7 shows the number of routing prefixes advertised by BGP between 1988 and 2009, as collected by the CIDR report. The number of updates a BGP router keeps in its forwarding table has grown linearly, thus making scalability a major issue. Any security measures must take this into consideration [44], [100].

A summary of the proposed BGP security solutions is given in Table 2. Currently, the only solutions deployed in

Fig. 7. *1988–2009 routing table updates from the CIDR report (http://www.cidr-report.org/).*

wide use are the use of route filtering and some reliance on routing registries, which are only moderately effective at best. As discussed in Section III, many solutions require secure and valid route registries as a minimum for their effectiveness; for example, this information is necessary for correctly communicating address ownership and delegation, and is a necessary first condition for implementing real origin authentication solutions. Accomplishing even this goal is non-trivial because of the amount of invalid information in registries and the number of legacy allocations that exist. Ensuring the accuracy of the registries accomplishes many goals beyond security however, as ISPs can use this information to identify customers and peers, and to clarify what filtering policies are and should be. This is a necessary first step that will aid network operators and security researchers in myriad ways.

Another major concern that hampers adoption of proposed solutions is the perception within the operations community that the computational requirements (e.g., symmetric and public-key cryptography) of many current solutions will overload deployed routers, and the cost of upgrading those routers, if it is even feasible, or replacing them outright, is prohibitive. Regardless of which platform is picked, the solutions will add additional complexity, infrastructure, and cost to the network, and could potentially affect convergence [101]. However, solutions that reduce the costs of cryptography, such as those discussed in Section IV-B1, may mitigate some of these concerns. In addition, advances in cryptography may provide primitives that are more performant. For example, new and improved digital signatures may aid in the efficiency of signature-based countermeasures [102]–[105]. Forward-secure signatures [106] can preserve nonrepudiability of past signatures, a potentially important feature depending on the timeliness of revocation

---

[10]The initial published protocol inherits the *common modulus* limitation from RSA. The authors provide alternate constructions which address this problem in later versions of the paper.

Table 2 Global BGP Security Solutions—Requirements (Columns) Relate to the Guarantees Provided for Global AS Data. *In Use* Indicates Whether the Solution Is Presently in Operational Use. *Style* Indicates Whether the Solution Is Based on a Cryptographic Protocol or an Anomaly Detection Service. The Authenticity Services Include: Topology (Are Paths Conforming to the Correct Topology), Path (Are All Paths Authenticated), and Origin (Are Origins Authenticated). A System Is Strong if it Provides Cryptographic Authenticity Guarantees, and Weak if Its Received Data Is Probabilistically Authentic/Correct

Solution Definition			Security Services		
System	In Use	Style	Topo. Auth.	Path Auth.	Origin Auth.
Route Filtering [17], [46],	yes	anomaly	*weak*	*weak*	*weak*
Routing Registries [48]	yes	anomaly	*weak*	*weak*	*weak*
S-BGP [61]	no	crypto	**strong**	**strong**	**strong**
soBGP [66]	no	crypto/anomaly	**strong**	none	**strong**
IRV [68]	no	crypto/anomaly	**strong**	**strong**	**strong**
Origin Authentication [71]	no	crypto	none	none	**strong**
SPV [75]	no	crypto	**strong**	**strong**	none
Signature Amortization [80]	no	crypto	**strong**	**strong**	none
Reference Locality [82]	no	crypto	**strong**	**strong**	none
psBGP [86]	no	crypto	*weak*	**strong**	*weak*
MOAS Detection [88]	no	anomaly	none	none	*weak*
Intrusion Detection [91]	no	anomaly	none	none	*weak*
PHAS [92]	no	anomaly	none	none	*weak*
Pretty Good BGP [94]	no	anomaly	*weak*	*weak*	*weak*
Real-Time Monitoring (Hu and Mao)	no	anomaly	none	none	*weak*
Whisper [98]	no	anomaly	none	*weak*	none

throughout all ASes in the Internet. These signatures can be competitive in performance with traditional signatures if properly configured for the application [107].

Potential measures that could be implemented include more robust modeling of security protocols through formal analysis to understand the security obtained; Aiello *et al.* [71], Butler *et al.* [82], and van Oorschot *et al.* [86] explore formal semantics in their works. Understanding how to adopt these solutions and the effect of that adoption through modeling is another way to evaluate solutions; preliminary work in this area has been performed by Chan *et al.* [108]. Finally, robust simulation of the security schemes across a common testbed may help the community determine the trade-offs necessary for solution adoption and assist in the parameterization or hybridization of these schemes, (e.g., combining facets of signature amortization schemes and using them in conjunction with one or more anomaly detection schemes). Very detailed network simulators such as *ns2* [109] are often best used for simulating detailed events in a small network setting, but may be difficult to scale to a sufficient level for modeling the global network encompassed by BGP, and are not made with the protocol in mind. Simulators such as SSFNet [110] address many of the demands made by a protocol such as BGP, but fully modeling the workings of the over 30 000 ASes that

comprise the Internet may still not be feasible without extremely large-scale and highly parallelized solutions, or abstracting away details that may not be relevant to security. A common simulator and framework for deployment, such as DETER [111], may be the most appropriate method for fully evaluating solutions, along with small-scale deployments, with input from both the research and operational communities.

V. FUTURE DIRECTIONS IN BGP SECURITY RESEARCH

We turn our attention now to work that can impact how BGP security is approached, and techniques that may be used to improve aspects of BGP's operation, improving security at the same time.

A. Routing Frameworks and Policies

A study on the performance impact of incrementally deploying router-assisted services shows that choosing the right deployment strategy for a new protocol or service can mean the difference between success and failure [112]. Suggestions have been made for designing a routing architecture in large networks such that scalability requirements are met [113]. A model and middleware for routing protocols, SPHERE, decomposes routing protocols

into fundamental building blocks to support hierarchical design [114]. Another approach towards analysis of routing security is performed by Pei *et al.* [115], who suggest defining a defense framework for intra- and interdomain routing protocols. This includes classifying areas of protection into fields such as cryptographic protection schemes and semantics validation. Each of these efforts aims to provide a foundation for designing an interdomain routing security solution. Additionally, best common practices (BCPs) build resistance into BGP routing [30]. Armed with BCPs and other tools, the Internet can be made more secure by simply protecting the most connected nodes. One study shows that protecting most connected nodes provides significant security gains [116].

B. Attack Detection

Detecting attacks is an active field of research. The PAIR algorithm [117] is an approach to discover, and recover from, inconsistencies in distance-vector routing. It may be possible to employ similar techniques in a path-vector protocol such as BGP. Protocols that detect and route around faults may also yield valuable insights [118]. The ability to recover from routing attacks and failures is crucial to infrastructure reliability. One study shows that path faults in BGP can at times take up to 30 minutes to repair [119]. In certain cases, some end-to-end routing failures may not be reflected in BGP traffic at all [120]. Being able to detect attacks before they occur is clearly the best alternative, and tools such as secure traceroute [121] and AS-level traceroute [122] to detect malicious routing may aid in this effort.

C. Data Plane Protection

The solutions discussed to this point have generally focused on ensuring that BGP receives and transmits control messages, such as those related to path announcement and withdrawal, properly. These solutions protect the *control plane* of BGP. However, a largely unexamined issue is that of ensuring that packets are actually forwarded by BGP along the announced paths; in other words, protecting BGP's *data plane*. Goldberg *et al.* [123] show that even in the event that full protection of the control plane is attained, ASes may have incentives to announce one path (in the control plane) and forward traffic over another (in the data plane), and may be able to subvert the data plane unless restrictive policies or data-plane protections are in place. Preliminary efforts to protect the data plane include work by Wong *et al.* [124], which introduces a verification protocol to be run by routers such that they can verify whether data traffic is following the route that has been advertised through BGP. This solution does not require cryptographic operations in the data path, but requires offline setup and shared secrets between ASes. It also requires modification of data packets and has overhead costs: tagging 5% of packets was shown to decrease throughput by 6.1%. Providing protections in an efficient manner will continue to be an ongoing research challenge.

D. Partial Deployment

In recent years, several researchers have explored the incentives for ASes to deploy BGP security solutions and the effectiveness of partial deployments [94], [108], [125], [126]. For example, the Pretty Good BGP (PGBGP) scheme [94] discussed earlier in Section IV is quite effective when deployed by the hundred largest ASes—about 0.5% of all ASes. These gains are possible because of the Internet's topological structure, where a small number of very large ASes are responsibile for much of the path diversity. These ASes learn many routes, making it more likely that they learn at least one valid route. In addition, the choices these ASes make have a profound influence on the routes available to other ASes. As such, even ASes that do not participate in PGBGP are likely to pick a valid route propagated by these larger participating ASes. A small group of ASes can also cooperate to increase path diversity, to make it even more likely that ASes learn at least one valid route [125], [126]. In particular, participating ASes can voluntarily "deflect" each other's data packets on to alternate paths, to allow the participants to circumvent invalid paths announced by an adversary. Finally, small groups could take even more aggressive action against adversaries by cooperatively announcing each others' address space, and then deflecting traffic toward the legitimate destination [125]. Together, these techniques enable a relatively small group of (say) five ASes (with the help of one or two large ISPs) to significantly constrain the effectiveness of a BGP attack [125].

Other BGP solutions, including (but not limited to) S-BGP, soBGP, and IRV, have also considered issues of incremental deployment in detail. For example, S-BGP is incrementally deployable but requires neighboring ASes to also deploy the protocol for benefit to be obtained; in addition, within an AS, all border routers must use S-BGP to prevent routing loops within the AS [67]. soBGP considered scenarios in which some, but not all connected networks use the solution, as well as deploying parts of soBGP but not others [127].

VI. CONCLUSION

BGP has been quite successful in providing stable interdomain routing, and is surprisingly robust. It was originally thought in many circles that the ISO's Inter-domain Routing Protocol (IDRP) [85] would be the successor to BGP, but because of diminishing interest in network protocols other than IP, BGP is the only interdomain routing protocol in wide use [18]. Moreover, because of its huge installed base, BGP will continue to play a crucial role in Internet routing. As such, BGP will adapt to changing needs of its constituency. This is evident

even now: multi-protocol extensions are increasingly used to route IPv6 packets [128]. It is important to note, however, that the impact of IPv6 upon BGP, particularly with regards to security, remains unclear.

Interdomain routing security has progressed since being first investigated by Perlman, but few production environments are demonstrably more secure than they were when she began that work. Some operators are using incremental solutions that offer some protection, but comprehensive solutions have not been deployed. Notably, no solutions requiring more than lightweight cryptography have been deployed. There is a resistance in the operations community to using any sort of cryptography in networks, largely due to the costs imposed. In addition, there is resistance to a global PKI (required to deploy many of the security solutions) with a single root of trust; such issues

are problematic with PKI in general [129]. Many of these issues must be solved before effective BGP security solutions can be deployed. Because of the global impact of even minor errors in BGP configuration and operation, such deployment is increasingly imperative. Recent developments have seen the beginnings of deployments of regional PKIs by RIRs, however, and the SIDR working group's standards on deploying PKIs [130] are valuable in creating solutions to these issues.

This survey has examined the threats to BGP and proposed solutions to ensure its security. While they have not been implemented yet in practice, and while their adoption may be difficult, good progress *has* been made. In the end, a methodology for securing BGP may be one of the best ways to ensure that the Internet remains a reliable and useful vehicle for private and public communication. ∎

REFERENCES

[1] J. Hawkinson and T. Bates, *Guidelines for Creation, Selection, and Registration of an Autonomous System (AS),* RFC 1930, 1996.

[2] Y. Rekhter, T. Li, and S. Hares, *A Border Gateway Protocol 4 (BGP-4),* RFC 4271, Jan. 2006.

[3] J. Stewart, *BGP4: Inter-Domain Routing in the Internet.* Reading, MA: Addison-Wesley, 1999.

[4] R. Barrett, S. Haar, and R. Whitestone, "Routing snafu causes Internet outage," *Interactive Week,* Apr. 25, 1997.

[5] P. Boothe, J. Hiebert, and R. Bush, "How prevalent is prefix hijacking on the Internet?" in *Proc. NANOG 36,* Feb. 2006 [Online]. Available: http://www.nanog.org/mtg-0602/boothe.html

[6] Rensys Blog, *Con-Ed Steals the 'Net.* [Online]. Available: http://www.renesys.com/blog/2006/01/coned_steals_the_net.shtml

[7] Rensys Blog, *Pakistan Hijacks YouTube.* [Online]. Available: http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml%

[8] A. Ramachandran and N. Feamster, "Understanding the network-level behavior of spammers," in *Proc. ACM SIGCOMM,* Aug. 2006.

[9] H. Ballani, P. Francis, and X. Zhang, "A study of prefix hijacking and interception in the Internet," in *Proc. ACM SIGCOMM,* Aug. 2007.

[10] Department of Homeland Security, *The National Strategy to Secure Cyberspace, Feb. 2003.*

[11] Secure Inter-Domain Routing. [Online]. Available: http://www.ietf.org/html.charters/sidr-charter.html

[12] North American Network Operators Group. [Online]. Available: www.nanog.org

[13] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with Rocketfuel," *IEEE/ACM Trans. Networking,* vol. 12, no. 1, pp. 2–16, Feb. 2004.

[14] W. Eddy, *TCP SYN Flooding Attacks and Common Mitigations,* RFC 4987, Aug. 2007.

[15] S. Bellovin and E. Gansner. (2003, May). *Using Link Cuts to Attack Internet Routing.* [Online]. Available: http://www.cs.columbia.edu/smb/papers/reroute.pd

[16] T. Griffin and G. Huston, *BGP Wedgies,* RFC 4264, Nov. 2005.

[17] M. Caesar and J. Rexford, "BGP routing policies in ISP networks," *IEEE Network,* vol. 19, no. 6, pp. 5–11, Nov.–Dec. 2005.

[18] R. Perlman, *Interconnections: Bridges, Routers, Switches, and Internetworking Protocols,* 2nd ed. Reading, MA: Addison Wesley, 1999.

[19] O. Bonaventure, "Interdomain routing with BGP: Issues and challenges," in *Proceedings of the IEEE Symposium on Communications and Vehicular Technology (SCVT),* Louvain-la-Neuve, Belgium, Oct. 2002.

[20] L. Gao and J. Rexford, "Stable Internet routing without global coordination," *IEEE/ACM Trans. Networking,* Dec. 2001.

[21] R. Perlman, "Network layer protocols with Byzantine robustness," Ph.D. dissertation, Massachusetts Inst. Technol., Cambridge, Oct. 1988. MIT/LCS/TR-429.

[22] R. Rivest, *The MD5 Message-Digest Algorithm,* RFC 1321, Apr. 1992.

[23] U.S. National Bureau of Standards, *Secure Hash Standard,* FIPS PUB 180-2, Aug. 2002.

[24] H. Krawczyk, M. Bellare, and R. Canetti, *HMAC: Keyed-Hashing for Message Authentication,* RFC 2104, Apr. 1997.

[25] W. Diffie and M. Hellman, "New directions in cryptography," *IEEE Trans. Inf. Theory,* vol. IT-22, no. 6, pp. 644–654, Nov. 1976.

[26] I. F. Blake and T. Garefalakis, "On the complexity of the discrete logarithm and Diffie-Hellman problems," *J. Complexity,* vol. 20, no. 2–3, pp. 148–170, Apr.–Jun. 2004.

[27] K. Seo, C. Lynn, and S. Kent, "Public-key infrastructure for the secure border gateway protocol (S-BGP)," in *IEEE DARPA Information Survivability Conference and Exposition II,* Anaheim, CA, Jun. 2001.

[28] A. Heffernan, *Protection of BGP Sessions via the TCP MD5 Signature Option,* RFC 2385, Aug. 1998.

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

[30] B. Green, "BGP security update: Is the sky falling?" in *Proc. NANOG 25,* Jun. 2002.

[31] J. Touch, A. Mankin, and R. Bonica, *The TCP Authentication Option,* Internet Draft, Jul. 2009.

[32] B. Smith and J. Garcia-Luna-Aceves, "Securing the border gateway routing protocol," in *Global Internet '96,* London, U.K., Nov. 1996.

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

[34] M. G. Gouda, E. N. Elnozahy, C.-T. Huang, and T. M. McGuire, "Hop integrity in computer networks," in *Proc. 8th Int. Conf. Network Protocols,* Osaka, Japan, Nov. 2000.

[35] V. Gill, J. Heasley, and D. Meyer, *The Generalized TTL Security Mechanism (GTSM),* RFC 3682, Feb. 2004.

[36] S. Kent and K. Seo, *Security Architecture for the Internet Protocol,* RFC 4301, Dec. 2005.

[37] R. Thayer, N. Doraswamy, and R. Glenn, *IP Security Document Roadmap,* RFC 2411, Nov. 1998.

[38] C. Kaufman, *Internet Key Exchange (IKEv2) Protocol,* RFC 4306, Dec. 2005.

[39] S. Kent, *IP Authentication Header,* RFC 4302, Dec. 2005.

[40] S. Kent, *IP Encapsulating Security Payload (ESP),* RFC 4303, Dec. 2005.

[41] B. Gleeson, A. Lin, J. Heinanen, G. Armitage, and A. Malis, *A Framework for IP Based Virtual Private Networks,* RFC 2764, Feb. 2000.

[42] T. Bates, P. Smith, and G. Huston, "CIDR Report for 30 January 08," Mar. 2008. [Online]. Available: http://www.cidr-report.org/

[43] J. Stewart, T. Bates, R. Chandra, and E. Chen, *Using a Dedicated AS for Sites Homed to a Single Provider,* RFC 2270, Jan. 1998.

[44] S. Bellovin, R. Bush, T. Griffin, and J. Rexford. (2001, Jun.). *Slowing Routing Table Growth by Filtering Based on Address Allocation Policies.* [Online]. Available: http://www.cs.princeton.edu/~jrex/papers/filter.pdf

[45] D. Chang, R. Govindan, and J. Heidemann, "An empirical study of router response to large BGP routing table load," in *Proc. ACM SIGCOMM Internet Measurement Workshop (IMW),* Nov. 2002.

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

[47] N. Feamster, Z. M. Mao, and J. Rexford, "BorderGuard: Detecting cold potatoes from peers," in *Proc. 2004 Internet Measurement Conf.,* Taormina, Italy, Oct. 2004.

[48] 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,* RFC 1786, Mar. 1995.

[49] L. Blunk, J. Damas, F. Parent, and A. Robachevsky, *Routing Policy Specification Language Next Generation (RPSLng),* RFC 4012, Mar. 2005.

[50] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Trans. Networking*, vol. 9, no. 6, pp. 733–745, Dec. 2001.

[51] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *IEEE INFOCOM 2002*, New York, Jun. 2002.

[52] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, K. C. Claffy, and G. Riley, "AS relationships: Inference and validation," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 1, pp. 29–40, Jan. 2007.

[53] T. Griffin, Personal Communication, Jun. 2003.

[54] M. Lepinksi, S. Kent, and D. Kong, *A Profile for Route Origin Authorizations (ROAs),* Internet Draft, Jul. 2009.

[55] APNIC. (2006, Nov.). *The APNIC Resource Certification Page.* [Online]. Available: http://mirin.apnic.net/resourcecerts/

[56] D. Harrington, R. Presuhn, and B. Wijnen, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks,* RFC 3411, Dec. 2002.

[57] M. Kaeo, *Current Operational Security Pratices in Internet Service Provider Environments,* RFC 4778, Jan. 2007.

[58] J. Mirkovic and P. Reiher, "A taxonomy of DDoS attack and DDoS defense mechanisms," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 2, Apr. 2004.

[59] V. Gill, "Lack of priority queueing considered harmful," *ACM Qeue*, vol. 2, no. 8, pp. 64–69, Nov. 2004.

[60] M. O. Nicholes and B. Mukherjee, "A survey of security techniques for the border gateway protocol (BGP)," *IEEE Comm. Surveys & Tutorials*, vol. 11, no. 1, Q1 2009.

[61] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 4, Apr. 2000.

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

[63] S. Kent, C. Lynn, J. Mikkelson, and K. Seo, "Secure Border Gateway Protocol (S-BGP) real world performance and deployment issues," in *Proc. ISOC Symp. Network and Distributed System Security (NDSS)*, San Diego, CA, Feb. 2000.

[64] D. M. Nicol, S. W. Smith, and M. Zhao, "Evaluation of efficient security for BGP route announcements using parallel simulation," *Simul. Model. Prac. Theor.*, vol. 12, no. 3–4, pp. 187–216, Jul. 2004.

[65] S. Bellovin, "SBGP—Secure BGP," in *NANOG 28*, Jun. 2003.

[66] J. Ng, *Extensions to BGP to Support Secure Origin BGP (soBGP),* Internet Draft, Apr. 2004.

[67] S. Kent, "Securing the Border Gateway Protocol: A status update," in *Proc. 7th IFIP TC-6 TC-11 Conf. Communications and Multimedia Security*, Torino, Italy, Oct. 2003.

[68] 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 *Proc. ISOC NDSS'03*, San Diego, CA, Feb. 2003, pp. 75–85.

[69] T. Dierks and E. Rescorla, *The Transport Layer Security (TLS) Protocol Version 1.1,* RFC 4346, Apr. 2006.

[70] B. Baker and R. Shostak, "Gossips and telephones," *Discrete Math.*, no. 2, pp. 191–193, 1972.

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

[72] R. Merkle, "Protocols for public key cryptosystems," in *IEEE Symp. Security and Privacy*, Oakland, CA, Apr. 1980.

[73] M. Naor and K. Nissim, "Certificate revocation and certificate update," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 4, pp. 561–570, Apr. 2000.

[74] Y. Hu, A. Perrig, and D. Johnson, "Efficient security mechanisms for routing protocols," in *Proc. ISOC Network and Distributed Systems Security Symp. (NDSS)*, San Diego, CA, Feb. 2003.

[75] Y.-C. Hu, A. Perrig, and M. Sirbu, "SPV: Secure path vector routing for securing BGP," in *Proc. ACM SIGCOMM*, Portland, OR, Aug. 2004.

[76] S. Even, O. Goldreich, and S. Micali, "On-line/off-line digital signatures," *J. Cryptol.*, vol. 9, no. 1, pp. 35–67, 1996.

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

[78] B. Raghavan, S. Pranjwani, and A. Mityagin, "Analysis fo the SPV secure routing protocol: Weaknesses and lessons," *ACM SIGCOMM Comput. Commn. Rev.*, Apr. 2007.

[79] A. Fujioka, T. Okamoto, and S. Miyaguchi, "ESIGN: An efficient digital signature implementation for smart cards," in *Proc. EUROCYPT*, Brighton, U.K., Apr. 1991.

[80] M. Zhao, S. W. Smith, and D. M. Nicol, "Aggregated path authentication for efficient BGP security," in *Proc. 12th ACM Conf. Computer and Communications Security (CCS'05)*, Alexandria, VA, Nov. 2005.

[81] M. Zhao, S. W. Smith, and D. M. Nicol, "Evaluating the performance impact of PKI on BGP security," in *Proc. 4th Annu. PKI R&D Workshop*, Gaithersburg, MD, Feb. 2005.

[82] K. Butler, P. McDaniel, and W. Aiello, "Optimizing BGP security by exploiting path stability," in *Proc. 13th ACM Conf. Computer and Communications Security (CCS'06)*, Alexandria, VA, Nov. 2006.

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

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

[85] ISO, *Intermediate System to Intermediate System Inter-Domain Routing Information Exchange Protocol,* DIS 10747, Jul. 1992.

[86] P. C. van Oorschot, T. Wang, and E. Kranakis, "On inter-domain routing security and pretty secure BGP (psBGP)," *ACM Trans. Inf. Syst. Security (TISSEC)*, vol. 10, no. 3, Jul. 2007.

[87] 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 *Proc. ACM SIGCOMM Internet Measurement Workshop, 2001*, San Francisco, CA, Nov. 2001.

[88] R. Chandra, P. Traina, and T. Li, *BGP Community Attribute,* RFC 1997, Aug. 1996.

[89] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang, "Detection of invalid routing announcement in the Internet," in *IEEE DSN 2002*, Washington, DC, Jun. 2002.

[90] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. Wu, and L. Zhang, "Protecting BGP routes to top level DNS servers," *IEEE Trans. Parallel Distrib. Syst.*, vol. 14, no. 9, pp. 851–860, Sep. 2003.

[91] C. Kruegel, D. Mutz, W. Robertson, and F. Valeur, "Topology-based detection of anomalous BGP messages," in *Proc. 6th Symp. Recent Advances in Intrusion Detection (RAID)*, Sep. 2003, pp. 17–35.

[92] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A prefix hijack alert system," in *Proc. 15th USENIX Security Symp.*, Vancouver, BC, Canada, Aug. 2006.

[93] G. Huston and G. Michaelson, *Validation of Route Origination in BGP Using the Resource Certificate PKI and ROAs,* Internet Draft, Aug. 2009.

[94] J. Karlin, S. Forrest, and J. Rexford, "Autonomous security for autonomous systems," *Comput. Networks*, Oct. 2008.

[95] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *Proc. ACM SIGCOMM 2002*, Pittsburgh, PA, Aug. 2002.

[96] X. Hu and Z. M. Mao, "Accurate real-time identification of IP prefix hijacking," in *Proc. IEEE Symp. Security and Privacy*, Oakland, CA, May 2007.

[97] Y. Zhang, Z. Zhang, Z. M. Mao, Y. C. Hu, and B. M. Maggs, "On the impact of route monitor selection," in *Proc. ACM Internet Measurement Conf. (IMC)*, San Diego, CA, Oct. 2007.

[98] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. Katz, "Listen and Whisper: Security mechanisms for BGP," in *Proc. Symp. Networked Systems Design and Implementation (NSDI)*, San Francisco, CA, Mar. 2004.

[99] R. Rivest, A. Shamir, and L. M. Adelman, "A method for obtaining digital signatures and public-key cryptosystems," *Commun. ACM*, vol. 21, no. 2, pp. 120–126, Feb. 1978.

[100] G. Huston, *Commentary on Inter-Domain Routing in the Internet,* RFC 3221, Dec. 2001.

[101] C. Meyer and A. Partan, "BGP security, availability, and operator needs," in *Proc. NANOG 28*, Jun. 2003.

[102] K. Zhang, "Efficient protocols for signing routing messages," in *Proc. ISOC NDSS'98*, San Diego, CA, Mar. 1998.

[103] M. Goodrich, *Efficient and Secure Network Routing Algorithms*, provisional patent filing, Jan. 2001.

[104] D. Boneh, C. Gentry, H. Shacham, and B. Lynn, "Aggregate and verifiably encrypted signatures from bilinear maps," in *Proc. Eurocrypt 2003*, 2003, vol. LNCS 2656, pp. 416–432.

[105] D. Boneh, B. Lynn, and H. Shacham, "Short signatures from the Weil pairing," *J. Cryptol.*, vol. 17, no. 4, pp. 297–319, Sep. 2004.

[106] M. Bellare and S. Miner, "A forward-secure digital signature scheme," in *Advances in Cryptology—CRYPTO '99 Proc.*, 1999, vol. LNCS 1666, pp. 431–438.

[107] E. Cronin, S. Jamin, T. Malkin, and P. McDaniel, "On the performance, feasibility, and use of forward-secure

signatures," in *Proc. ACM CCS'03*, Washington, DC, Oct. 2003.

[108] H. Chan, D. Dash, A. Perrig, and H. Zhang, "Modeling adoptability of secure BGP protocols," in *Proc. ACM SIGCOMM 2006*, Pisa, Italy, Sep. 2006.

[109] L. Breslau, D. Estrin, K. Fall, S. Floyd, J. Heidemann, A. Helmy, P. Huang, S. McCanne, K. Varadhan, Y. Xu, and H. Yu, "Advances in network simulation," *IEEE Computer*, vol. 33, no. 5, pp. 59–67, May 2000.

[110] J. Cowie, H. Liu, J. Liu, D. Nicol, and A. Ogielski, "Towards realistic million-node Internet simulations," in *Proc. 1999 Int. Conf. Parallel and Distributed Processing Techniques and Applications (PTPTA'99)*, Las Vegas, NV, Jun. 2000.

[111] T. Benzel, R. Braden, D. Kim, C. Neuman, A. Joseph, K. Sklower, R. Ostrenga, and S. Schwab, "Experience with DETER: A testbed for security research," in *Proc. 2nd IEEE Conf. Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom 20006)*, Barcelona, Spain, Mar. 2006.

[112] X. He and C. Papadopoulos, "A framework for incremental deployment strategies for router-assisted services," in *IEEE INFOCOM 2003*, San Francisco, CA, Apr. 2003.

[113] J. Yu, *Scalable Routing Design Principles,* RFC 2791, Jul. 2000.

[114] V. Stachtos, M. Kounavis, and A. Campbell, "SPHERE: A binding model and middleware for routing protocols," in *Proc. 4th Conf. Open Architecture and Network Programming (OPENARCH 2001)*, Anchorage, AK, Apr. 2001.

[115] D. Pei, D. Massey, and L. Zhang, "A framework for resilient Internet routing protocols," *IEEE Network*, vol. 18, no. 2, pp. 5–12, Mar.–Apr. 2003.

[116] S. Gorman, R. Kulkarni, L. Schintler, and R. Stough. (2003). *Least Effort Strategies for Cybersecurity*. [Online]. Available: http://arxiv.org/ftp/cond-mat/papers/0306/0306002.pdf

[117] A. Chakrabarti and G. Manimaran, "An efficient algorithm for malicious update detection & recovery in distance vector protocols," in *IEEE Int. Conf. Communications*, Anchorage, AK, May 2003.

[118] I. Avramopoulos, H. Kobayashi, R. Wang, and A. Krishnamurthy, "Highly secure and efficient routing," in *IEEE INFOCOM 2004*, Hong Kong, Mar. 2004. PRC.

[119] C. Labovitz, A. Ahuja, R. Wattenhofer, and S. Venkatachary, "Resilience characteristics of the Internet backbone routing infrastructure," in *Proc. 3rd Information Survivability Workshop*, Boston, MA, 2000.

[120] N. Feamster, D. Andersen, H. Balakrishnan, and M. Kaashoek, "Measuring the effects of Internet path faults on reactive routing," in *Proc. ACM SIGMETRICS 2003*, San Diego, CA, Jun. 2003.

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

[122] Z. Mao, J. Rexford, J. Wang, and R. Katz, "Towards an accurate AS-level traceroute tool," in *Proc. ACM SIGCOMM 2003*, Karlsruhe, Germany, Aug. 2003.

[123] S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R. N. Wright, "Rationality and traffic attraction: Incentives for honest path announcements in BGP," in *Proc. ACM SIGCOMM*, Aug. 2008.

[124] E. L. Wong, P. Balasubramanian, L. Alvisi, M. G. Gouda, and V. Shmatikov, "Truth in advertising: Lightweight verification of route integrity," in *Proc. 26th Annu. ACM SIGACT-SIGOPS Symp. Principles of Distributed Computing (PODC 2007)*, Aug. 2007.

[125] I. Avramopoulos, M. Suchara, and J. Rexford, "How small groups can secure interdomain routing," Princeton Univ. Comput. Sci. Dept., Tech. Rep. TR-808-07, Dec. 2007.

[126] D. Wendlandt, I. Avramopoulos, D. Andersen, and J. Rexford, "Don't secure routing protocols, secure data delivery," in *Proc. ACM SIGCOMM Workshop on Hot Topics in Networking*, Nov. 2006.

[127] R. White, *Deployment Considerations for Secure Origin BGP (soBGP)*, Internet Draft, Jun. 2006

[128] T. Bates, R. Chandra, D. Katz, and Y. Rekhter, *Multiprotocol Extensions for BGP-4*, RFC 4760, Jan. 2007.

[129] C. Ellison and B. Schneier, "Ten risks of PKI: What you're not being told about public key infrastructure," *Comput. Security J.*, vol. 16, no. 1, 2000.

[130] M. Lepinski and S. Kent, *An Infrastructure to Support Secure Internet Routing*, Internet Draft draft-ietf-sidr-arch-08.txt, Jul. 2009.

ABOUT THE AUTHORS

Kevin Butler (Student Member, IEEE) is a Ph.D. candidate in Computer Science and Engineering at the Pennsylvania State University. He received his M.S. in electrical engineering from Columbia University in 2003 and his B.Sc. in electrical engineering from Queen's University at Kingston, Ontario, in 1999. Kevin's research primarily focuses on systems and storage security. He has also closely examined security and policy considerations for interdomain routing, and has investigated issues in secure hardware, privacy, and worm propagation across the Internet and in wireless networks.

Patrick McDaniel (Senior Member, IEEE) is an Associate Professor in the Computer Science and Engineering Department at the Pennsylvania State University and co-director of the Systems and Internet Infrastructure Security Laboratory. Patrick's research efforts centrally focus on network, telecommunications, and systems security, language-based security, and technical and public policy issues in digital media. Patrick was awarded the National Science Foundation CAREER Award and has chaired several top conferences in security including, among others, the 2007 and 2008 IEEE Symposium on Security and Privacy and the 2005 USENIX Security Symposium. Patrick is the editor-in-chief of the *ACM Journal Transactions on Internet Technology* (TOIT), and serves as Associate Editor of the journals *ACM Transactions on Information and System Security* and IEEE TRANSACTIONS ON SOFTWARE ENGINEERING. Prior to pursuing his Ph.D. in 1996 at the University of Michigan, Patrick was a software architect and program manager in the telecommunications industry.

Toni R. Farley is a Lecturer in the School of Computing, Informatics and Decision Systems Engineering and the School of Letters and Sciences at Arizona State University. She received B.S. and Ph.D. degrees in Computer Science from ASU. Dr. Farley joined ASU in 2001 as a Research Associate in the Information and Systems Assurance Laboratory, where she served as Assistant Director from 2003 to 2006. She has been a Lecturer since 2006. Dr. Farley is the recipient of a graduate research fellowship from AT&T Labs-Research. Her research interests include networks, security, informatics and cross-disciplinary inquiry.

Jennifer Rexford (Senior Member, IEEE) is a Professor in the Computer Science department at Princeton University. From 1996–2004, she was a member of the Network Management and Performance department at AT&T Labs-Research. Jennifer is co-author of the book *Web Protocols and Practice* (Addison-Wesley, May 2001). She served as the chair of ACM SIGCOMM from 2003 to 2007. Jennifer received her B.S.E. degree in electrical engineering from Princeton University in 1991, and her M.S.E. and Ph.D. degrees in computer science and electrical engineering from the University of Michigan in 1993 and 1996, respectively. She was the 2004 winner of ACM's Grace Murray Hopper Award for outstanding young computer professional.