

A Survey of BGP Security Issues and Solutions

TONI FARLEY, PATRICK MCDANIEL and KEVIN BUTLER

AT&T Labs Research

BGP is the *de facto* protocol enabling interdomain routing in the Internet. Although BGP has proven to be generally stable, there are mounting concerns about its ability to meet the needs of the rapidly evolving Internet. A central limitation of BGP is its failure to address security. The design and ubiquity of BGP have complicated past efforts at securing interdomain routing. This paper surveys works relating to BGP security. We explore the limitations and advantages of proposed solutions, and consider the systemic and operational implications of their design. We centrally note that no current solution has yet found a perfect balance between comprehensive security and deployment cost. Recent BGP-related outages and security analyses clearly indicate that the current Internet routing infrastructure is highly vulnerable. Our investigation calls not only for application of ideas and approaches described within this paper, but also for further introspection on the problems and solutions for BGP security.

Categories and Subject Descriptors: C.2.0 [Computer-Communication Networks]: General—Security and Protection; C.2.2 [Computer-Communication Networks]: Network Protocols—Routing protocols; C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks—Internet

General Terms: Security

Additional Key Words and Phrases: authentication, authorization, BGP, border gateway protocol, integrity, interdomain routing, network security, networks, routing

1. INTRODUCTION

Information on the Internet is sent via IP packets, which follow a path of routers from their source to their destination. Routers are collectively responsible for maintaining paths, or routes, to all reachable destinations on the Internet. Reachability information is shared between routers by routing protocols. As traffic is received at a router, it is forwarded based on the reachability information stored in the router's forwarding table, and other information stored in the packet's header.

Routers on the Internet use an interdomain routing protocol called the Border Gateway Protocol (BGP) to share routing information. BGP has been around since the commercialization of the Internet and is widely deployed, maintained and researched. BGP works well in practice. However, it does not provide performance or security guarantees.

Problems with interdomain routing lead to poor performance and routing vulnerabilities. While many routing problems are mere annoyances, there have also been documented routing failures of significant impact. One such outage occurred

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0000-0000/20YY/0000-0001 \$5.00

April 25, 1997:

A misconfigured router maintained by a small Virginia service provider injected an incorrect routing map into the global Internet. This map indicated that the Virginia company's network provided optimal connectivity to all Internet destinations. Internet providers that accepted this map automatically diverted all of their traffic to the Virginia provider. The resulting network congestion, instability, and overload of Internet router table memory effectively shut down most of the major Internet backbones for up to two hours. Incorrect published contact information for operations staff, and lack of procedures for inter-provider coordination exacerbated the problem. [Barrett et al. 1997]

Loss of connectivity on the Internet may manifest itself as anything from an inconsequential annoyance to a devastating communication failure, depending on the breadth and scope of the lost connections. For example, today's Internet is home to an increasing number of critical business applications, such as online banking and stock trading, that can cause financial harm to an individual or institution if communication is lost at a critical time (such as during a time-sensitive trading session). As the number of time-sensitive applications on the Internet grows, so will our reliance on the Internet to provide us with reliable and secure services.

Current research involving BGP focuses on exposing and resolving issues related to interdomain routing operation and security. While security is the main focus of this paper, there are operational concerns relating to BGP, such as scalability, slow convergence, route stability and performance. Security research largely focuses on the integrity, authentication, confidentiality, authorization and validation of BGP speakers and the communications between them.

This paper explores current research in interdomain routing security and exposes the similarities and differences in proposed approaches to building a more secure Internet. The next section gives a brief overview of interdomain routing and BGP. Subsequent sections expose current research addressing BGP and interdomain routing security issues.

2. ROUTING VULNERABILITIES AND CONSEQUENCES

The Internet is a network of networks [Perlman 1999; Stewart 1999; Minoli and Schmidt 1999; Tanenbaum 2003]. Networks share information, in the form of IP packets, via routers. A group of routers under the same administrative control is considered an autonomous system (AS) [Stewart 1999; Hawkinson and Bates 1996]. There are three types of ASes: stub, multihomed, and transit. Stub ASes are communication endpoints, with connections to the rest of the Internet only made through a single upstream provider. Multihomed ASes are similar to stub ASes, but possess multiple upstream providers. Transit ASes have connections to multiple ASes and allow traffic to flow through to other ASes, even if the traffic does not originate or terminate within them (i.e. Internet Service Providers).

Within an AS, routers communicate with each other through the process of *intradomain* routing. This is accomplished using an internal gateway protocol (IGP), examples of which include the Routing Information Protocol (RIP) [Hedrick 1988] [Malkin 1994], the Open Shortest Path First IGP (OSPF) [Moy 1998], and

the Intermediate System to Intermediate System protocol (IS-IS) [Callon 1990]. To communicate between ASes, routers perform *interdomain* routing, using an external gateway protocol (EGP). The *de facto* standard EGP in use on the Internet is the Border Gateway Protocol version 4 (BGP4) [Rekhter and Li 1995], which has obsoleted previous versions and the original NSFNET EGP protocol [Mills 1984]. While other interdomain routing protocols exist, we restrict ourselves to BGP. However, some of the issues related to interdomain routing are independent of the protocol in use.

A router running the BGP protocol is known as a BGP *speaker*. BGP speakers communicate across TCP and become *peers* or *neighbors*. TCP is a reliable connection-oriented protocol. By employing TCP, BGP does not need to provide error correction at the transport layer [Minoli and Schmidt 1999]. Each pair of BGP neighbors maintains a *session*, over which information is communicated. A BGP speaker's neighbor is one *hop*¹ away, thus the term *per hop* refers to the relationship between BGP neighbors. BGP peers within the same AS (internal peers) communicate via internal BGP (IBGP). External BGP (EBGP) is used between speakers in different ASes (external peers). The relationships between ASes and BGP peers are shown in Figure 1.

There are currently more than 16,500 ASes in the Internet [CIDR 2004]. Each AS *originates* one or more address *prefixes*. A prefix is a representation for a block of IP addresses. Prefixes are expressed as "prefix / # most significant bits". For example, the prefix 192.68.0.0/16 has 16 significant bits and thus represents all of the IP addresses between 192.68.0.0 and 192.68.255.255 inclusive. Each AS establishes a path for the prefixes advertised by BGP. To simplify, the paths are vectors of ASes that any packets must traverse to reach the IP address. The last AS in the path is the *origin* of that address and its parent prefix. These vectors are stored in a routing table and shared with neighbors via BGP.

BGP peers constantly exchange Network Layer Reachability Information (NLRI), e.g., known paths and prefixes, via UPDATE messages. Each peer updates its routing tables based on its neighbors' NLRI, and forwards that information to its other neighbors. This *flooding* process ensures that all ASes are informed of the reachability of all prefixes. For as long as the session is active, peers use UPDATE messages to inform each other of routing table changes, which include the addition of new routes and withdrawal of old ones.

An AS may and often should receive multiple paths to a single prefix. BGP uses a complex algorithm to select which of these paths to use to forward and advertise to its neighbors. Policy communicated in UPDATE messages as well as local configuration may influence this process. However, in the absence of mitigating policy and subject to several other factors, BGP will select the shortest path (as measured in hops).

ASes are not only bound by physical relationships; they are also bound by business relationships. When an AS owner serves as a provider to another business entity, there are associated contractual agreements involved. Therefore, it is necessary to be able to enforce these agreements at the routing policy level. BGP

¹In other contexts the term hop often denotes directly connected router to router communication. Throughout, we use the term hop to denote direct AS to AS neighbor communication.

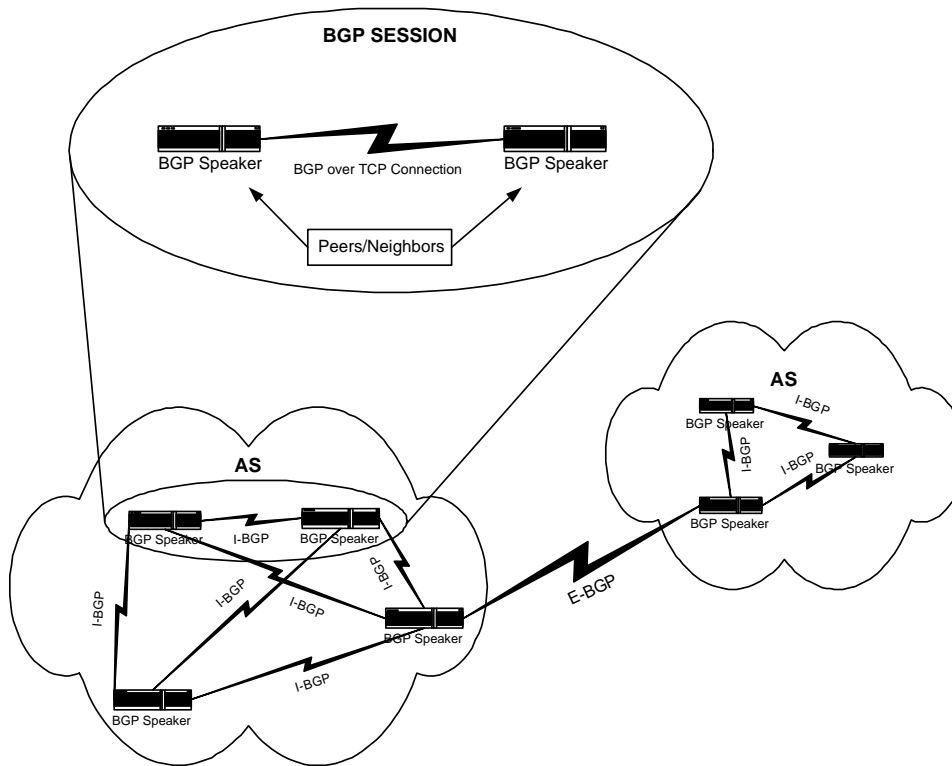


Fig. 1. I-BGP is used within an AS, while E-BGP is used between ASes.

enforces routing policies, such as the ability to forward data only for paying customers [Tanenbaum 2003].

Policies configured in a BGP router allow it to filter the routes received from each of its peers (import policy), filter the routes advertised to its peers (export policy), select routes based on desired criteria, and forward traffic based on those routes [Bonaventure 2002]. An example BGP policy restricts a speaker to only advertise transit routes to peers with whom it has a contract with to provide such service. BGP routers can be configured with route preferences, selective destination reporting (i.e., report a destination to some neighbors and not others), and rules concerning path editing [Perlman 1999]. Setting path preferences usually involves path editing, such as adding AS numbers to a path to discourage its use (a technique known as padding). These aspects of the protocol enable BGP to adhere to desired policies.

Although BGP has had success as a policy-based interdomain routing protocol, there are a number of issues that suggest that the Internet may have evolved beyond its current incarnation. In the next two sections, we discuss the security issues that have concerned users of BGP since its introduction.

3. INTERDOMAIN ROUTING SECURITY

Interest in BGP grew tremendously during the 1990s [Stewart 1999]. Prior to that, few had thought deeply about routing security [Perlman 1988]. In 1995, RFCs 1771 and 1772, describing BGP4 and its application in the Internet, were published [Rekhter and Li 1995; Rekhter and Gross 1995]. Since this time, a number of issues have emerged related to using BGP for interdomain routing. Li reports issues related to the scalability, slow convergence, instability, and efficiency of interdomain routing [Li 2003]. In this survey, we focus on security related issues and defer to other sources for discussions of these and other operational concerns.

BGP messages are subject to modification, deletion, forgery, and replay [Murphy 2003]. These exploits can be caused by malicious intent as well as faulty or misconfigured BGP routers. Moreover, bogus messages can originate from malicious sources or accidentally misconfigured peers. The effects of misconfiguring a BGP router can be similar to those of an attack. An analysis of BGP misconfigurations suggests that better router design could prevent most occurrences [Mahajan et al. 2002]. This study found that in the course of a day, 200-1200 prefixes, equivalent to 0.2-1% of the global routing table size, are misconfigured.

Mahajan et al. identify two areas of globally visible misconfigurations in BGP:

- (1) A router exports a route it should have filtered (export misconfiguration).
- (2) An AS accidentally injects a prefix into the global BGP tables (origin misconfiguration).

In October 2003, a seemingly small misconfiguration of a router caused widespread outages [W. Slater 2002]. Improper filtering rules added to a router caused the routing tables of WorldCom's internal infrastructure to become flooded with external routing data. The internal routers became overloaded and crashed repeatedly. This caused prefixes and paths advertised by these routers to disappear from routing tables and reappear when the routers came back online. This repeated advertisement and withdrawal of prefixes, known as *route flapping*, served to destabilize the surrounding network.

Another adverse effect of misconfiguration can be de-aggregation. This occurs when the announcement of a large prefix is fragmented or duplicated by a collection of announcements for smaller prefixes. De-aggregation harms the performance of BGP and indirectly the network by increasing the size of BGP tables and flooding the network with redundant UPDATES.

Malicious BGP packet manipulation can introduce errors in routing tables. Murphy suggests that this is due to three primary security related limitations of BGP. These are:

- BGP does not protect the integrity, freshness and source authentication of messages.
- BGP does not validate an AS's authority to announce reachability information.
- BGP does not ensure the authenticity of the path attributes announced by an AS.

As evidenced by the growth and apparent resilience of the Internet, BGP appears to work well in practice. However, recent analyses of BGP of the end-to-end behav-

ior of Internet show that that routing can and often does experience substandard, and even *broken* behavior [Paxson 1996; 1997; 1999]. Broken behavior is often manifest as IP packets being grossly misrouted. For example, Paxson reports packets that originated in the US and destined for London were erroneously routed through Israel. Moreover, subsequent studies show that the problems have not improved with time [Zhang et al. 2000].

3.1 BGP Vulnerability

Vulnerabilities provide an open door for attacks on the Internet. Currently, interdomain routing is vulnerable to a number of specific attacks [Murphy 2003]. These threats manipulate the three distinct types of BGP communication: control messages when setting up a session, or reachability updates and error messages throughout the duration of a session. The following describes and highlights the effect of these attacks:

- **Eavesdropping:** An adversary passively listens to data on the wire. This gives the adversary access to sensitive policy and route information being forwarded between ASes. Note that interdomain routing information is not widely viewed as sensitive. However, because it may expose the existence and details of commercial relationships, organizations often desire that exchanged peering policy be kept confidential.
- **Replay:** An adversary records messages and resends them to the original recipient. This approach can be used to confuse the routing protocols by re-asserting withdrawn routes or withdrawing valid ones. When sent in bulk, these messages can overwhelm the victim routers, causing a denial of service attack.
- **Message insertion:** An adversary inserts forged messages into a BGP session. These messages can erroneously terminate BGP sessions between peers or inject bad routing data. While BGP does not directly protect against this, its transport protocol, TCP, provides limited protection. TCP uses sequence numbers to preserve the ordering of packets [J 1981]. Because sequence numbers are often unpredictable, an adversary with limited abilities will find it difficult to insert forged BGP messages. Of course, adversaries who can eavesdrop or hijack the BGP session can trivially inject forged messages.
- **Message deletion:** An adversary intercepts and deletes a message passed between BGP peers. Deleted BGP UPDATE messages can lead to inaccurate routing tables. Again, TCP provides limited protection against this kind of attack.
- **Message modification:** An adversary removes messages from a BGP session, modifies them, and reinserts them. Like message insertion, this also leads to inaccurate routing (possibly across compromised links) and/or the breaking of peering relationships, resulting in routing failures.
- **Man-in-the-middle:** An adversary inserts itself between two peers and poses as the sender to the receiver and vice versa (see Figure 2). The threat against BGP in this type of attack is similar to that of message insertion, deletion and modification. Because BGP does not provide authentication of sources, it is particularly vulnerable to this type of attack.




Fig. 2. Man-in-the-middle attacks exploit BGP's failure to authenticate session end-points.

- Denial of service: An adversary floods the victim with resource request in an attempt to degrade or eliminate the availability of that resource. In BGP, the victim router is flooded with messages. This flood a routing table with bogus or unnecessary routes, causing the table size to exceed its capacity. Self deaggregation is a kind of denial of service, where an AS announces prefixes that should rightfully be aggregated, thereby unnecessarily advertising more specific prefixes [Aiello et al. 2003].

Attacks may be passive or active. An attack is passive if the adversary does not perform any overt (and often externally observable) act. For example, eavesdropping is the canonical passive attack. This type of attack may represent the entirety of an intrusive activity, or merely the reconnaissance phase of a later attack. Active attacks occur when the adversary attempts to directly manipulate the interdomain routing protocol. Rescorla and Korver advise writers of security specifications to classify attacks and countermeasures as being/addressing passive and active attacks [Rescorla and Korver 2003]. As described above, With the notable exception of eavesdropping, all attacks listed above are active.

The consequences of these attacks are as diverse as their approach. BGP sessions can be prematurely severed, networks and ASes can be made unreachable, the address space can become fragmented, or any number of other undesirable things can result from an attack. These problems can be used in concert to amplified bad behavior or enable further malicious activity. For example, Bellovin et al. showed how an adversary can route traffic through subverted elements by severing links between BGP speakers [Bellovin and Gansner 2003].

Barbir, Murphy, and Yang further view the problems in routing security in terms of generic consequences [Barbir et al. 2003]. They show how the consequences spring from the vulnerabilities and limitations of the routing infrastructure through which they can be realized. These include, for example:

- Disclosure: eavesdropping, deliberate exposure, traffic analysis
- Deception: message insertion, deletion, modification, man-in-the-middle
- Disruption: replay, DoS/DDoS
- Usurpation: accroach a router's services and/or functions

The vulnerabilities and consequences presented here illustrate the need for secure interdomain routing. No matter how its limitations are classified, it is evident that the current BGP protocol lacks the facilities necessary to secure the future Internet.

4. SECURITY IN INTERDOMAIN ROUTING

The protocols that the Internet is built upon were designed to enable communication between largely trusted parties. BGP was designed to address interdomain routing within those trusted networks. While essential to the growth of the Internet, commercial interests and new user communities have changed the nature of the originally assumed network. This is particularly true of routing. Thus, the environment for which BGP was designed for is fundamentally different than the current Internet.

As discussed in the preceding section, the fundamental shift in the nature of the Internet enables and amplifies the effect of malicious behavior or accidental misconfiguration. This has led to a call for greater Internet infrastructure security [Perlman 1988; Hares 2003]. Many note that this security infrastructure must encompass routing, of which BGP is essential ingredient [Green 2002]. One approach consists of defining a defense framework for intra- and interdomain routing protocols, classifying areas of protection into fields such as cryptographic protection schemes and protocol semantics checking [Pei et al. 2003]. The remainder of this paper considers the security of BGP. Current research efforts in this area can be classified as addressing *hop integrity*, *origin authentication* and *path validation*.

Hop integrity investigates techniques for securing the communication between BGP-speaking peers. Solutions for hop integrity typically address the integrity and authenticity of the BGP session. An integrity mechanism validates that the data passed between peers is not modified in any way (e.g., not altered, augmented, deleted or replayed). An authentication mechanism validates the identity of the sending peer.

An origin authentication mechanism validates an BGP speaker's right to assume an AS number of advertise a range of addresses. In short, this latter service authenticates the use of address space. The IPv4 address space used on the Internet is delegated to ASes through a hierarchical network of issuing authorities and delegating organizations [Aiello et al. 2003]. Origin authentication asks the question: "Is the advertising AS authorized to be the origin of the (range of addresses represented by) prefix 120.40.0.0/16?" Solutions to origin authentication vary widely, and are the subject of active investigation.

A path validation solution ensures that a received path is topologically valid and authentic. In a BGP UPDATE message, each announced prefix has an associated AS path to that prefix. This ensures that the path reflects real and usable connectivity between ASes. Furthermore, it validates determine that each AS in the path did indeed advertised its part of the path.

The following sections considers architectures and solutions that address various aspects of these services. We begin by considering the three leading candidates for BGP security, S-BGP, soBGP, and the IRV system.

4.1 Secure BGP (S-BGP)

Secure BGP (S-BGP) is a comprehensive solution to BGP security [Kent et al. 2000]. It attempts to address the majority of security issues defined in the preceding section. The S-BGP protocol and its architecture are currently under consideration for standardization by the IETF. Implementations exist, and its authors have been

experimenting with its use in real networks.

S-BGP implements security by extending the existing routing infrastructure and BGP protocol. The central element of S-BGP is a pair of public key infrastructures (PKIs) used to delegate address space and AS numbers, as well as associate particular network elements with their parent ASes. The first PKI is used to authenticate address allocation. The second PKI manages AS assignment and router associations through a combination of three certificates: AS number and organization's public key, AS number and its public key, and AS number and router information (Domain Name System (DNS) name, id, public key).

Attestations are signed statements of delegation or identity also managed within the PKIs. The attestations are signed by a BGP speaker or other authority using a private key which is associated with the public key certificate. A new BGP path attribute is introduced to carry attestations in UPDATE messages. Receivers of attestations use the PKI to validate the signature, and hence the authenticity of the signed statements.

S-BGP provides hop integrity by mandating peers that communicate via the IPsec security protocol suite. IPsec provides IP layer integrity, authentication, and optionally confidentiality services [Kent and Atkinson 1998]. The IPsec authentication and key management services are also used. All hop integrity requirements are achieved by simply applying these services. Note that other solutions (e.g., IRV) acknowledge the ease of this approach and recommend its use.

Origin authentication is implemented by S-BGP through address attestations [Seo et al. 2001]. These attestations are statements of delegation between the address authorities and the ASes that use them. Organizations and address authorities explicitly delegate blocks of address space to other entities by creating and signing attestations. Receivers of the attestation in BGP updates validate the delegation by validating the associated signature.

Path validation is implemented through route attestation. Created by the ASes while passing routes, these attestations sign advertisements. The signed route attestation includes path and potentially policy information including in the update. Each AS signs the attestation augmented with its own path and policy information. A speaker is able to validate the authenticity and integrity of every AS on a path from source to nearest neighbor by validating received attestations.

4.2 Interdomain Route Validation (IRV) Service

The Interdomain Route Validation (IRV) service is a receiver-driven protocol and associated architecture [Goodell et al. 2003]. Unlike S-BGP, IRV's operation is independent of the routing protocols. The authors of IRV use an external validation service to allow for incremental deployment – independent communities can implement IRV without affecting the operation or knowledge of other, non-IRV, ASes. The authors of IRV further argue that much of the suboptimal behavior of BGP is the result of the many functions added to the protocol, but not considered by its original designers. Hence, adding new features may further exacerbate these issues.

IRV uses a validator model. The IRV server in an AS queries IRV servers in other ASes for validation of received routing information. Upon reception of an UPDATE message, a receiving BGP speaker will appeal to its local IRV service for an indication of whether the received information is correct (see Figure 3). Where




Fig. 3. ASes running the IRV protocol query the appropriate authorities for validation of received routing data.

deemed necessary, the IRV will query another AS's IRV using an arbitrary query language. The query transaction is executed over a secure transport (e.g., IPsec, TLS/SSL). Because the IRV queries sources directly over a secure transport, it does not incur the signature costs of S-BGP style attestation generation or validation.

Each AS is responsible for choosing an algorithm that determines when an UPDATE messages should be validated. Upon deciding a message is suspicious, the AS can query all of the relevant ASes to verify the authenticity and accuracy of the contents.

Hop integrity is provided by the IRV through the secure transport. Integrity and authenticity are provided inasmuch as an IRV validates statements from one hop away. However, this is not provided on a per message basis (unless of course the receiving IRV chooses to authenticate every message it receives).

An origin is authenticated in IRV in a similar manner to how sources are authenticated. An IRV can perform a top down query of the address authority and delegating ASes/organizations. By corroborating data with the sending and neighboring IRVs, this method has the ability to use multiple feedback sources to determine the authenticity of an origin AS. The methodology behind the IRV protocol stresses the importance of robust origin authentication, which includes validating the association between an organization and its AS (AS authentication), along with the authentication of the source of an announcement.

A path is validated by querying each AS in the path. The path is deemed valid if the ASes acknowledge transmission of the path. This operation may consume many resources or take considerable time. Such queries should be performed by an external service, and it may not be desirable to put the transaction into the critical path of route selection. Each AS must decide how to handle the validation delay, either optimistically, by acting on the new information, or pessimistically, by delaying its use. The former policy may lead to the use of forged routes, and the latter may delay use of correct ones.

4.3 Secure Origin BGP (soBGP)

Secure origin BGP (soBGP) proposes an extension to BGP [Ng 2002]. soBGP adds small security enhancements to the existing BGP protocol. The primary mechanism of soBGP is the new SECURITY message type. The SECURITY message is used by BGP speakers to share certificates and attestations. The data of these messages are signed by the sender and allows the receiver to validate the public key bindings, policy, or routing data. In essence, all security data relevant to soBGP is communicated through the SECURITY message. Hop integrity is not provided by soBGP, but the authors acknowledge the availability and necessity of secure transport protocols like IPsec for BGP sessions.

soBGP provides three types of certificates/attestations transported by the SECURITY message: Entity, Policy and Authorization. The entity certificate is used to verify the existence of an entity (i.e., source) within a routing system. The policy certificate provides information about an AS, which can be used to validate its authenticity. The authorization certificate provides information about an AS's authority to announce an address. This latter certificate is used to provide origin authentication.

Path validation is provided through path verification and AS validations. The verification of a path is supported through the creation of a path database. To build this database, each speaker announces the ASes that it is connected to via *attached AS* fields in the policy certificate, one field identifying attached transit ASes and another identifying nontransit ASes. From this information, speakers can build a path database of all possible paths to a prefix. As each prefix is processed, these databases can be queried to confirm that the associated path is valid. Furthermore, the origin and second hop (the second entry in the AS path) information can be validated as the message travels to its destination by checking that the second hop is connected to the AS. The entire path is validated by recursively applying second-hop validation. However, this information only serves to illustrate that the path is valid, not that the incoming packet actually arrived via the specified route.

4.4 Smith's Countermeasures

Predating much of the other work in BGP security, Smith et al. proposed five countermeasures to secure interdomain routing [Smith and Garcia-Luna-Aceves 1996; Smith et al. 1997; Smith and Garcia-Luna-Aceves 1998]. These solutions propose to enhance the existing BGP protocol with new and modified attributes. Two countermeasures aim to protect BGP control messages: encryption of all BGP messages between peers, and addition of message sequence numbers. The other three countermeasures offer protection for UPDATE messages and include the following: 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 of all fields in the UPDATE message whose values are fixed.

Encrypting all messages requires the use of a session key between peers. This allows for hop integrity between the peers, and provides authenticity for the messages. Using a sequence number also provides integrity, as it makes replay attacks more difficult. Digitally signing the unchanging UPDATE information is also a

method of providing integrity.

Origin authentication is not presented as a feature of the countermeasures, but path validation is provided through use of the PREDECESSOR attribute. Digitally signing the destination's predecessor information provides some assurance of the validity of the second hop. This idea of using the penultimate AS for validation has not gone unnoticed; validation of the second hop is an option in protocols such as soBGP. Digitally signing the fields in an UPDATE message ensures attribute freshness and the verification of path information, in conjunction with the predecessor information. However, the path validation does not extend beyond the penultimate AS, meaning that past that point, a malicious router can tamper with path and policy information.

4.5 Hop Integrity Approaches

Received information has data integrity if one can validate that it has not been modified in transit. If data integrity is not provided, an adversary may modify the information in any number of ways. For example, an adversary may alter the AS path so that data is routed across a subverted link. Source, or origin, authentication is a subtly different property that guarantees the identity of the sending peer (e.g., the sending peer is in fact who he or she claims to be). Data integrity is contingent upon source authentication, as one needs to know what was uttered is accurate before one can attribute it to a particular entity. Source authentication is discussed in detail in the following section, while this section concentrates on approaches that deal with hop and path integrity.

BGP uses TCP as its transport protocol, which itself is carried over IP datagrams. As a result, BGP is vulnerable to attacks that can be mounted on TCP and IP, including the spoofing of IP packets and session hijacking [Murphy 2003; Traina 1995]. To guard against spoofing, recent enhancements to BGP include the use of a TCP extension for carrying an MD5 digest [Heffernan 2002; 1998]. The MD5 digest mechanism requires manually configured keys, or a *shared secret*, at both ends of the BGP session. An MD5 keyed digest [Krawczyk et al. 1997] of the TCP header and BGP data is included each packet passing between the BGP speakers. A number of variants that have considered hashing all or part of the TCP and BGP data message using one or more keys [Heffernan 2002; Chunzhe et al. 2003; Przygienda 1997]. This simple and easily-implemented solution addresses many of the vulnerabilities associated with peer communication in BGP. In particular, it addresses many of the problems of spoofing and hijacking inherent to TCP [Green 2002].

The lack of periodic rekeying and the cost of manual configuration are undesirable features of MD5 digests. A static shared secret be installed on the router, unlike IPSec, which allows the dynamic negotiation of shared secrets [Murphy 2003; Harkins and Carrel 1998]. A static secret is more challenging to implement because it requires that each router be pre-configured with the shared secrets of each of its peers. Heffernan (2002) describes that MD5 may be vulnerable to attack through collisions, and suggests another cryptographic hashing algorithm, such as SHA-1, could be used instead. Finally, many consider the optional nature of MD5 digests in BGP to be undesirable [Murphy 2003; Green 2002], although the most recent draft of the BGP protocol requires implementations to support MD5 authentication.

tion [Rekhter and Li 2003].

Gouda et al. proposed a suite of protocols that provide peer source authentication, data integrity through detection of modified messages, and protection against replay attacks [Gouda et al. 2000]. Two peers communicate over a secret exchange protocol. They send each other requests for secret keys using public-key encryption, and when the setup round is complete, each receives a copy of the other's secret key to set up a symmetric session (note that this is a key for secret-key cryptography, not their peer's private key used in public-key cryptography). Because symmetric keys are much faster, these are what are used for communication between the two peers, until a given amount of time elapses and new keys are exchanged. This eliminates the need to manually set up shared secrets on each router. There is a weak integrity protocol that provides protection against message manipulation. Messages are signed with a message digest function, keyed with the sender's shared key. The receiving router, which has a copy of the secret shared key, uses it to verify the message using the same digest algorithm. The strong integrity protocol adds sequence numbering to prevent messages from being replayed.

4.6 Origin Authentication

Origin Authentication (OA) is a method of validating address ownership. It addresses what is potentially the most dangerous problem currently facing BGP, because of the protocol's inherent vulnerabilities. A misconfigured router that originates incorrect route information, or even information relating to an AS it does not own, can cause major black hole effects throughout the entire Internet [Misel html]. One effort directly investigates origin authentication (OA) by looking at the semantics, design and application of OA services [Aiello et al. 2003]. A formalization of the semantics of address delegation is performed, and different proof structures for carrying delegation attestations are shown. These include simple attestations that are generated for each update (a model used by S-BGP), an authenticated delegation list that contains all delegations made by an organization and is signed once, an authenticated delegation tree based on a Merkle hash tree that results in proof size growing logarithmically rather than linearly, and an authenticated delegation dictionary based on a balanced 2-3 search tree. Using an efficient structure such as an authenticated delegation tree, Aiello et al. show that it is feasible to provide on-line, in-band origin authentication, which had previously been thought to be too computationally expensive to perform in this fashion.

One of the earliest attempts at providing authentication of address space delegation was a DNS-based approach to verify NLRI [Bates et al. 1998]. In this approach, a new resource record would be added to the DNS, called the Autonomous System RR. It consists of an AS number that a given address prefix would be delegated to, and a decimal representation of the prefix length of the addresses to be allocated. An ORIGIN field would be repurposed for reporting the address prefix owned by the organization identified in the DNS zone file. The NS and CNAME fields would also be used for delegating parts of a prefix allocated on a non-octet boundary and non-octet allocations, respectively. There was a circular dependency, however, as the routing system was required to transport DNS information required to secure it. In addition, the DNS database is vulnerable to forgery and cache poisoning [Bellovin 1995]. The servers are currently more hardened to attack and play a critical role

in the Internet infrastructure, and there has been a proposal to protect the BGP routes to these servers [Wang et al. 2003]. It is possible that the stability of the root servers and the relatively static routes to reach them could make them useful as part of a security infrastructure.

A particular origin authentication problem is a Multiple Origin AS (MOAS) conflict. In general, each IP prefix should only be associated with one origin AS [Hawkinson and Bates 1996]. Any BGP path to a prefix should end in the AS that originates the prefix. A MOAS conflict occurs when a prefix appears to belong to multiple ASes. While there are some valid cases for prefixes to appear as being originated from multiple sources, this is generally a typical indicator of *traffic hijacking*, intentionally or otherwise. A study of MOAS conflicts showed that they occurred as a result of [Zhao et al. 2001].

MOAS conflicts can result from intrusive activity. By ignoring these conflicts, we allow an opportunity for traffic to be *hijacked*. However, there are valid cases where MOAS conflicts may occur, resulting in false positives if used for intrusion detection. A recent study of MOAS conflicts showed potential causes to include prefixes associated with exchange point addresses (which link ASes), multi-homing without BGP or with private AS numbers, and faulty configurations [Zhao et al. 2001].

A proposal was made to add enhancements to BGP, using the community attribute ([Chandra et al. 1996]) to distinguish between valid and invalid MOAS conflicts [Zhao et al. 2001; Zhao et al. 2002]. 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 an MOAS conflict is valid. Because the community attribute is optional and transitive, however, routers can drop this information without causing an error.

4.7 Path Validation

An effort aimed at securing path vector protocols could be applied to BGP for path validation [Hu et al. 2003]. In this proposal, a cumulative authentication mechanism is employed that authenticates the list of routers on a path. In terms of BGP security, this translates to validating the ASes in the AS path of a BGP UPDATE message.

The mechanism works with symmetric cryptographic techniques. Each AS on an UPDATE's path shares a secret key with the destination AS. A well-known value, such as 0, is used as the path authenticator, and a message authentication code (MAC) is computed using a concatenation of the authenticator and the fields in the UPDATE message that do not change (e.g. ORIGIN attribute, NLRI, withdrawn routes, etc.). The MAC is a keyed hash using the shared secret as a key and a cryptographic hash algorithm, such as MD5 or SHA-1. The next AS on the path uses its shared secret value with the destination AS to compute a MAC over the previous MAC value, and so on until the destination is reached. Knowing the path traversed through the AS_PATH attribute, the destination can verify the MAC by using the secret key it shares with each AS and generating each MAC in sequence from the origin. Since the resulting hash chain is comprised of one-way functions, an AS cannot delete or modify the path earlier in sequence, as that would be computationally infeasible. For example, for packet p and destination AS Z , if the

update originates in AS W and traverses AS X and Y , the MACs would be created with shared keys WZ , XZ , and YZ respectively, and if the value 0 is chosen to as the path authenticator, the packet would have the form $h_{YZ}(h_{XZ}(h_{WZ}(0\|p)\|p)\|p)$, when it arrived at AS Z , where $h_{ab}(x)$ is the hashing operation with shared secret key ab used to generate a MAC and $\|$ is the concatenation operator.

While symmetric keys are orders of magnitude faster than using public and private keys, this scheme relies on a mechanism for distributing and negotiating pairwise shared keys between every AS. The destination must also be known *a priori* so that the correct set of pairwise shared keys are used.

4.8 Listen and Whisper

Another scheme that does not rely on a public key infrastructure for dealing with security vulnerabilities is the combination of the Listen and Whisper security mechanisms [Subramanian et al. 2004]. The protocols do not seek to provide perfect security, but rather to alert network administrators if routing inconsistencies are found, particularly misconfigurations and malicious adversaries acting alone.

Whisper deals with control plane anomalies, including propagating false AS origin information or a fake path.. There are two modes of operation: Weak Split Whisper (WSW) and Strong Split Whisper (SSW). WSW uses a hash chain, similar to the cumulative authentication mechanism used by Hu et al (2003). A difference is that WSW does not provide any real authentication, using the hash chain only to determine the number of hops (or ASes in this case) have been traversed by the update, and does not provide much protection.SSW replaces a hash chain with a construction similar to RSA. A large number N is created as the product $p \times q$ where p and q are large primes. A prime generator g is chosen and a random number z , then the value $g^z \bmod N$ is computed and used as the value sent by the origin AS as part of a 2-tuple $(N, g^z \bmod N)$. The origin and subsequent ASes on the path use the value of their AS number as the exponent z , so that if the origin's AS is Q and the update passes through ASes W and X on its way to destination Y , the value of the tuple when it reaches Y will be $(N, g^{zQWX} \bmod N)$. The authors define some mathematical functions that reduce the effects of the commutative property of multiplication, the ability to derive ASes through factoring and the ability to add an AS to the path without being detected. If the value at the destination does not match what is computed, all ASes on the path are considered suspect and a penalty is assigned to each of them. ASes with higher penalty values are more likely to be originating faulty routing information.

Listen is a simpler protocol, used to sound alerts if there are data plane attacks such as inconsistent route advertisements as a result of misconfiguration or attack. It acts as a sniffer, monitoring TCP traffic flows and determining if hosts in remote prefixes are reachable. If a TCP SYN packet is observed, followed by a DATA packet, the connection is considered to be complete. Since forward and reverse traffic can follow different paths, looking for the ACK is not important. If a certain percentage of hosts in a remote prefix do not respond, the protocol assumes that the route is not verifiable, and may be blackholed or otherwise misconfigured. Actively dropping a subset of packets and listening for retransmission, or checking that packets are not retransmitted at an abnormally high rate, are ways to counter active attackers that are generating SYN and DATA packets without a corresponding

ACK.

Listen provides some assurances of data integrity by monitoring for black holes, and Whisper provides some origin and path assurances, though not to the extent of more comprehensive solutions such as S-BGP. While they provide extra heuristics when dealing with adversaries who collude to provide faulty information or with the vast amount of active probe traffic (problems that affect Whisper and Listen, respectively), they are not meant to be able to find the crux of problems, only to raise an alert that they have been found.

4.9 Supplemental Resources for BGP Security

Solutions to one problem are often confounded by another. This section looks at some efforts that complement and provide foundations for other interdomain routing security efforts.

One research approach to providing security against attacks is to first define attack scenarios. Attack trees provide a tool for defining common atomic attack goals and subsequent attack scenarios [Convery et al. 2003] in a clearly defined and easily understandable fashion.

A routing registry that stores routing policy information is another potential tool for adding security to interdomain routing [Bates et al. 1993; Bates et al. 1995]. Information about AS policies and routes can be stored and accessed by other ASes to learn information about another AS, validate their place in the AS topology, and even build a picture of the interdomain routing paths in the Internet. To use a registry, one must first be assured that the registry itself is secure. One study proposes an authentication and authorization model for providing data integrity in routing policy systems [Villamizar et al. 1998]. One drawback of the registry model is that corporations often consider their peering data, policies and routes to be proprietary information, though tools such as Rocketfuel [Spring et al. 2002] provide accurate maps of internal topology, and algorithms exist for inferring customer and peering relationships [Subramanian et al. 2002].

Many security techniques involve the use of digital signatures. New and improved signatures may aid in the efficiency of signature-based countermeasures [Goodrich 2001; Boneh et al. 2003]. One study also suggests an efficient, low cost protocol for signing routing messages [Zhang 1998]. One area of particular interest is the field of forward-secure digital signatures [Bellare and Miner 1999], where the public key of a digital signature is fixed but the private key, used for signing, changes with time. This ensures that if the key is compromised, messages from the past cannot be forged, thus preserving non-repudiability of past signatures. Recent work has shown that forward-secure signatures can have performance figures competitive with traditional signatures if properly configured for the application [Cronin et al. 2003].

BGP routers often are configured to filter ingress routes, which come into the router, and egress routes, which leave the router [Green 2002; Gouda et al. 2000]. This filtering includes disallowing prefixes that are documenting special use addresses (DSUA) prefixes, 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 [CIDR 2004]. A policy of careful ingress and egress filtering greatly aids in maintaining security for both the local AS and

	Hop Integrity	Origin Authentication	Path Validation
Lynn, Kent and Seo (S-BGP)	Yes	Yes	Yes
Goodell, et al (IRV)	Yes	Yes	Yes
Ng, White and Lorwick (soBGP)		Yes	Yes
Smith, et al (Countermeasures)	Yes		Yes
Gouda, et al (Shared secrets)	Yes		
Przygienda & Heffernan & Chunzhe (MD5)	Yes		
Zhao, et al (MOAS Lists)		Yes	
Bates, et al (DNS-based)		Yes	
Hu, et al (hash chains)			Yes

Fig. 4. Summary of interdomain routing security efforts.

its neighbors.

5. EVALUATING STRATEGIES AND RESPONDING TO ATTACKS

Figure 4.

From the summary in Figure 4, it is apparent that many protocols offer a degree of protection against attack, with S-BGP offering the most comprehensive solution. One may question why this scheme is not already in place on the Internet. We consider the costs and potential difficulties associated with implementing a BGP security solution architecture.

Designing security into BGP is tricky, complicated by the need to avoid aggravating existing operational problems and designing new ones. Interdomain routing is stressed by the continuous growth of the Internet. Around 30,000 AS numbers have already been assigned. Due to the increasing number of ASes, There are predictions that if current trends continue, the AS number space will be exhausted by as early as 2009 [Huston 2003]. This growth 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 Figure 5 shows BGP updates from the CIDR report for 1994 to 2003. Scalability problems must be considered when adding security measures to BGP [Huston 2001; Bellovin et al. 2001].

5.1 Evaluating the Major Security Schemes

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 or failure [He and Papadopoulos 2003]. Suggestions have been made for designing a routing architecture in large networks such that scalability requirements are met [Yu 2000]. A model and middleware for routing protocols, SPHERE, decomposes routing protocols into fundamental building blocks to support hierarchical design [Stachtos et al. 2001]. Each of these efforts aims to provide a foundation for designing an interdomain routing security solution.

A study on S-BGP deployment issues finds that the added overhead of S-BGP countermeasures is equivalent to the CPU and memory provided by a desktop PC [Kent et al. 2000]. Thus, the hardware requirement is ostensibly minimal, although concerns have been raised over the use time-averaged statistics. The load




Fig. 5. 1994-2003 routing table updates from the CIDR report (<http://www.cidr-report.org/>)

in routers is not uniformly distributed, as Internet traffic is bursty in nature [Uhlig and Bonaventure 2001; Leland et al. 1993]. It has also been claimed that S-BGP will cause administrative delays [Meyer and Partan 2003].

The soBGP platform provides several deployment options and the ability to be incrementally deployed [White 2002]. These options give it a greater ease of deployment than S-BGP, but the number of options could yield issues with interoperability [Kent 2003]. Further work on soBGP defines RADIUS attributes to support its provisioning [Lonvick 2003], but the author admits that using RADIUS is a suboptimal solution in the absence of a better alternative. Furthermore, soBGP may not guard against mid-path disruptions [Bellovin 2003]

Regardless of which platform is picked, the solutions will add additional complexity, infrastructure, and cost to the network, and could potentially affect convergence [Meyer and Partan 2003]. BGP convergence is a major issue, as it has been shown that the protocol may not converge at all [Griffin and Wilfong 1999; Labovitz et al. 2001]. It is possible, though, that advances such as in-band origin authentication [Aiello et al. 2003] could make either proposal easier to manage once deployed.

The countermeasures developed by Smith allowed for the development of S-BGP and soBGP, but they do not provide the necessary origin authentication, and require changes to BGP, without being as comprehensive in scope as either of the two forementioned protocols. The IRV solution takes care all security signalling

out of band, a feature that has been deployed in soBGP, but requires more analysis of infrastructure requirements and operational semantics to be a viable security alternative. Finally, the Listen/Whisper protocol set is meant to be easily deployable without major infrastructure changes, but they are limited in the amount of security they can provide.

5.2 Responding to Attacks

On September 18, 2001, the Code Red/Nimda attack was correlated with a 30% increase in BGP UPDATE messages [Wang et al. 2002]. The behavior of BGP under this attack is analyzed by Wang et al., who finds that over 40% of those updates are attributable to BGP measurement settings. There is clearly room to improve the protocol's performance under stress. Best common practices (BCPs) build resistance into BGP routing [Green 2002]. 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 [Gorman et al. 2003].

Detecting attacks is an active field of research. A scalable algorithm, *PAIR*, detects router attacks by attaching predecessor and pathsum metrics to UPDATE messages [Chakrabarti and Manimaran 2003]. The predecessor metrics can be used by a receiving router to build a path tree and calculate pathsum and hoplength metrics for each node. These metrics can then be compared against the pathsum and hoplength message in the actual update for accuracy. The ability of the Internet to recover from 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 [Labovitz et al. 2000]. In certain cases, some end-to-end routing failures may not be reflected in BGP traffic at all [Feamster et al. 2003]. Being able to detect attacks before they occur is clearly the best alternative and tools such as secure tracroute [Padmanabhan and Simon 2002] to detect malicious routing may aid in this effort.

6. CONCLUSION

BGP has been quite successful in providing relatively stable interdomain routing. Enhancements to the protocol, such as TCP MD5 Signatures, serve to add much needed security measures. This survey exposes areas where it is commonly believed that BGP still needs improvements in security.

BGP has been surprisingly robust. It was originally thought in many circles that the ISO's Interdomain Routing Protocol (IDRP) would be the successor to BGP, but because of diminishing interest in network protocols other than IP, BGP is the one interdomain routing alternative [Perlman 1999]. BGP is being used with IPv6 as well, so it will continue to play a crucial role for many years in Internet routing. While moving towards more complex solutions and public key infrastructures seems like a lot of work, it may be the best way to ensure that the Internet stays reachable and secure in the years to come.

REFERENCES

- AIELLO, W., IOANNIDIS, J., AND MCDANIEL, P. 2003. Origin authentication in interdomain routing. ACM CCS, Washington, DC.

- BARBIR, A., MURPHY, S., AND YANG, Y. 2003. Generic threats to routing protocols. IETF Draft.
- BARRETT, R., HAAR, S., AND WHITESTONEO, R. 1997. Routing snafu causes internet outage. *Interactive Week*.
- BATES, T., BUSH, R., LI, T., AND REKHTER, Y. 1998. Dns-based nlri origin as verification in bgp. IETF Draft.
- BATES, T., GERICH, E., JONCHERAY, L., JOUANIGOT, J.-M., KARREBERG, D., TERPSTRA, M., AND YU, J. 1995. Representation of ip routing policies in a routing registry. IETF RFC 1786.
- BATES, T., JOUANIGOT, J.-M., KARREBERG, D., LOTHBERG, P., AND TERPSTRA, M. 1993. Representation of ip routing policies in the ripe database. RIPE-81.
- BELLARE, M. AND MINER, S. 1999. A forward-secure digital signature scheme. Vol. LNCS 1666. *Advances in Cryptology - CRYPTO '99 Proceedings*, 431–438.
- BELLOVIN, S. 1995. Using the domain name system for system break-ins. Fifth Usenix Security Symposium, Salt Lake City, UT, USA.
- BELLOVIN, S. 2003. Sbgp - secure bgp. NANOG 28.
- BELLOVIN, S., BUSH, R., GRIFFIN, T., AND REXFORD, J. 2001. Slowing routing table growth by filtering based on address allocation policies. <http://www.research.att.com/jrex/>.
- BELLOVIN, S. AND GANSNER, E. 2003. Using link cuts to attack internet routing. Draft: <http://www.research.att.com/smb/papers/index.html>.
- BONAVENTURE, O. 2002. Interdomain routing with bgp: Issues and challenges. IEEE SCVT2002, Louvain-la-Neuve, Belgium.
- BONEH, D., GENTRY, C., SHACHAM, H., AND LYNN, B. 2003. Aggregate and verifiably encrypted signatures from bilinear maps. Vol. LNCS 2656. *Eurocrypt 2003*, 416–432.
- CALLON, R. 1990. Use of osi is-is for routing in tcp/ip and dual environments. IETF RFC1195.
- CHAKRABARTI, A. AND MANIMARAN, G. 2003. An efficient algorithm for malicious update detection & recovery in distance vector protocols. IEEE Intl. Conf. on Communications, Anchorage, AK, USA.
- CHANDRA, R., TRAINA, P., AND LI, T. 1996. Bgp community attribute. IETF RFC1997.
- CHUNZHE, H., QIULIN, D., HUI, N., AND DEFENG, L. 2003. Bgp sessions protection via md5 authentication. IETF Draft.
- CIDR. 2004. Cidr report for 21 february 04. <http://www.cidr-report.org/>.
- CONVERY, S., COOK, D., AND FRANZ, M. 2003. An attack tree for the border gateway protocol. IETF Draft.
- CRONIN, E., JAMIN, S., MALKIN, T., AND MCDANIEL, P. 2003. On the performance, feasibility, and use of forward-secure signatures. ACM CCS'03, Washington, DC, USA.
- FEAMSTER, N., ANDERSEN, D., BALAKRISHNAN, H., AND KAASHOEK, M. 2003. Measuring the effects of internet path faults on reactive routing. ACM SIGMETRICS 2003, San Diego, CA.
- GOODELL, G., AIELLO, W., GRIFFIN, T., IOANNIDIS, J., MCDANIEL, P., AND RUBIN, A. 2003. Working around bgp: An incremental approach to improving security and accuracy of interdomain routing. Internet Society Network and Distributed Systems Security 2003, San Diego, California, 75–85.
- GOODRICH, M. 2001. Efficient and secure network routing algorithms. provisional patent filing.
- GORMAN, S., KULKARNI, R., SCHINTLER, L., AND STOUGH, R. 2003. Least effort strategies for cybersecurity. <http://arxiv.org/ftp/cond-mat/papers/0306/0306002.pdf>.
- GOUDA, M. G., ELNOZAHY, E. N., HUANG, C.-T., AND MCGUIRE, T. M. 2000. Hop integrity in computer networks. Eighth International Conference on Network Protocols.
- GREEN, B. 2002. Bgp security update: Is the sky falling? NANOG 25.
- GRIFFIN, T. AND WILFONG, G. 1999. An analysis of bgp convergence properties. ACM SIGCOMM 1999.
- HARES, S. 2003. Bgp attack trees: Real world examples. NANOG 28.
- HARKINS, D. AND CARREL, D. 1998. The Internet Key Exchange. *Internet Engineering Task Force*. RFC 2409.
- HAWKINSON, J. AND BATES, T. 1996. Guidelines for creation, selection, and registration of an autonomous system (as). IETF RFC 1930.

- HE, X. AND PAPADOPOULOS, C. 2003. A framework for incremental deployment strategies for router-assisted services. IEEE INFOCOM.
- HEDRICK, C. 1988. Routing information protocol. IETF RFC1058.
- HEFFERNAN, A. 1998. Protection of bgp sessions via the tcp md5 signature option. IETF RFC 2385.
- HEFFERNAN, A. 2002. Protection of bgp sessions via the tcp md5 signature option. IETF Draft.
- HU, Y., PERRIG, A., AND JOHNSON, D. 2003. Efficient security mechanisms for routing protocols. Internet Society Network and Distributed Systems Security 2003, San Diego, CA.
- HUSTON, G. 2001. Commentary on inter-domain routing in the internet. IETF RFC3221.
- HUSTON, G. 2003. Bgp as number exhaustion. NANOG 28.
- J, P. 1981. Transmission Control Protocol - DARPA Internet Protocol Program Specification. *Internet Engineering Task Force*. RFC 793.
- KENT, S. 2003. Securing the border gateway protocol: A status update. Seventh IFIP TC-6 TC-11 Conference on Communications and Multimedia Security, Torino, Italy.
- KENT, S. AND ATKINSON, R. 1998. Security architecture for the internet protocol. IETF RFC2401.
- KENT, S., LYNN, C., MIKKELSON, J., AND SEO, K. 2000. Secure border gateway protocol (s-bgp) real world performance and deployment issues. ISOC Symposium on Network and Distributed System Security.
- KENT, S., LYNN, C., AND SEO, K. 2000. Secure border gateway protocol (secure-bgp). *IEEE Journal on Selected Areas in Communications* 18, 4 (April).
- KRAWCZYK, H., BELLARE, M., AND CANETTI, R. 1997. HMAC: Keyed-Hashing for Message Authentication. *Internet Engineering Task Force*. RFC 2104.
- LABOVITZ, C., AHUJA, A., WATTENHOFER, R., AND VENKATACHARY, S. 2000. Resilience characteristics of the internet backbone routing infrastructure. Third Information Survivability Workshop, Boston, MA.
- LABOVITZ, C., AHUJA, A., WATTENHOFER, R., AND VENKATACHARY, S. 2001. The impact of internet policy and topology on delayed routing convergence. IEEE INFOCOM 2001.
- LELAND, W., TAQQ, M., WILLINGER, W., AND WILSON, D. 1993. On the self-similar nature of ethernet traffic. ACM SIGCOMM, San Francisco, CA, USA.
- LI, W. 2003. Inter-domain routing: Problems and solutions. State University of New York at Stony Brook.
- LONVICK, C. 2003. Radius attributes for sobgp support. IETF Draft.
- MAHAJAN, R., WETHERALL, D., AND ANDERSON, T. 2002. Understanding bgp misconfiguration. ACM SIGCOMM 2002.
- MALKIN, G. 1994. Rip version 2. IETF RFC1723.
- MEYER, C. AND PARTAN, A. 2003. Bgp security, availability, and operator needs. NANOG 28.
- MILLS, D. 1984. External gateway protocol formal specification. IETF RFC904.
- MINOLI, D. AND SCHMIDT, A. 1999. *Internet Architectures*. John Wiley & Sons, New York, NY.
- MISEL, S. 1998. <http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html>. "wow, as7007!". Merit NANOG Archive.
- MOY, J. 1998. Ospf version 2. IETF RFC2178.
- MURPHY, S. 2003. Bgp security vulnerabilities analysis. IETF Draft.
- NG, J. 2002. Extensions to bgp to support secure origin bgp (sobgp). IETF Draft.
- PADMANABHAN, V. AND SIMON, D. 2002. Secure traceroute to detect faulty or malicious routing. ACM SIGCOMM Workshop on Hot Topic in Networks (HotNets-I), Princeton, NJ.
- PAXSON, V. 1996. End-to-end routing behavior in the internet. ACM SIGCOMM 1996.
- PAXSON, V. 1997. Measurements and analysis of end-to-end internet dynamics. Ph.D. thesis, University of California at Berkeley, Berkeley, California.
- PAXSON, V. 1999. End-to-end internet packet dynamics. *IEEE/ACM Transaction on Networking* 7, 3 (June), 277–292.
- PEI, D., MASSEY, D., AND ZHANG, L. 2003. A framework for resilient internet routing protocols. Tech. rep., UCLA. November.

- PERLMAN, R. 1988. Network layer Protocols with Byzantine Robustness. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA. MIT/LCS/TR-429.
- PERLMAN, R. 1999. *Interconnections: Bridges, Routers, Switches, and Internetworking Protocols, 2nd Edition*. Addison Wesley, Reading, MA.
- PRZYGIENDA, T. 1997. Bgp-4 md5 authentication. IETF Draft.
- REKHTER, Y. AND GROSS, P. 1995. Application of the border gateway protocol in the internet. IETF RFC 1772.
- REKHTER, Y. AND LI, T. 1995. A border gateway protocol 4 (bgp-4). IETF RFC 1771.
- REKHTER, Y. AND LI, T. 2003. A border gateway protocol 4 (bgp-4). IETF Draft.
- RESCORLA, E. AND KORVER, B. 2003. Guidelines for writing rfc text on security considerations. IETF Draft.
- SEO, K., LYNN, C., AND KENT, S. 2001. Public-key infrastructure for the secure border gateway protocol (s-bgp). IEEE DARPA Information Survivability Conference and Exposition II.
- SMITH, B. AND GARCIA-LUNA-ACEVES, J. 1996. Securing the border gateway routing protocol. Global Internet'96, London, UK, 20–21.
- SMITH, B. AND GARCIA-LUNA-ACEVES, J. 1998. Efficient security mechanisms for the border gateway routing protocol. *Computer Communications* 21, 3, 203–210.
- SMITH, B., MURTHY, S., AND GARCIA-LUNA-ACEVES, J. 1997. Securing distance vector routing protocols. Internet Society Symposium on Network and Distributed System Security, San Diego, CA.
- SPRING, N., MAHAJAN, R., AND WETHERALL, D. 2002. Measuring isp topologies with rocketfuel. ACM SIGCOMM, Pittsburgh, PA, USA.
- STACHTOS, V., KOUNAVIS, M., AND CAMPBELL, A. 2001. Sphere: A binding model and middleware for routing protocols. Fourth Conference on Open Architecture and Network Programming (OPENARCH 2001), Anchorage, Alaska.
- STEWART, J. 1999. *BGP4: Inter-Domain Routing in the Internet*. Addison-Wesley, Reading, MA.
- SUBRAMANIAN, L., AGARWAL, S., REXFORD, J., AND KATZ, R. 2002. Characterizing the internet hierarchy from multiple vantage points. IEEE INFOCOM, New York, NY, USA.
- SUBRAMANIAN, L., ROTH, V., STOICA, I., SHENKER, S., AND KATZ, R. 2004. Listen and whisper: Security mechanisms for bgp. First Symposium on Networked Systems Design and Implementation, San Francisco, CA, USA.
- TANENBAUM, A. 2003. *Computer Networks, fourth edition*. Prentice Hall, New Jersey.
- TRAINA, P. 1995. Experience with the bgp-4 protocol. IETF RFC1773.
- UHLIG, S. AND BONAVENTURE, O. 2001. Understanding the long-term self-similarity of internet traffic. 2nd International Workshop of Quality of Future Internet Services, Coimbra, Portugal.
- VILLAMIZAR, C., ALAETTINOGLU, C., MEYER, D., MURPHY, S., AND ORANGE, C. 1998. Routing policy system security. IETF Draft.
- W. SLATER, I. 2002. The internet outage and attacks of october 2002. <http://www.isoc-chicago.org/internetoutage.pdf>.
- WANG, L., ZHAO, X., PEI, D., BUSH, R., MASSEY, D., MANKIN, A., WU, S., AND ZHANG, L. 2003. Protecting bgp routes to top level dns servers. *IEEE Transactions on Parallel and Distributed Systems*, 851–860.
- WANG, L., ZHAO, Z., PEI, D., BUSH, R., MASSEY, D., MANKIN, A., WU, S., AND ZHANG, L. 2002. Observation and analysis of bgp behavior under stress. IMW workshop 2002.
- WHITE, R. 2002. Deployment considerations for secure origin bgp (sobgp). IETF Draft.
- YU, J. 2000. Scalable routing design principles. IETF RFC2791.
- ZHANG, K. 1998. Efficient protocols for signing routing messages. NDSS, 1998.
- ZHANG, Y., PAXSON, V., AND SHENKER, S. 2000. The stationarity of internet path properties: Routing, loss, and throughput. Tech. rep., ACIRI. May.
- ZHAO, X., MANKIN, A., MASSEY, D., PEI, D., WANG, L., WU, S., AND ZHANG, L. 2001. Validation of multiple origin ascs conflicts through bgp community attribute. IETF Draft.
- ZHAO, X., PEI, D., WANG, L., MASSEY, D., MANKIN, A., WU, S., AND ZHANG, L. 2002. Detection of invalid routing announcement in the internet. IEEE DSN 2002, Washington DC.

ZHAO, X., PEI, D., WANG, L., MASSEY, D., MANKIN, A., WU, S. F., AND ZHANG, L. 2001. An analysis of bgp multiple origin as (moas) conflicts. ACM SIGCOMM Internet Measurement Workshop, 2001.