A Survey of BGP Security

KEVIN BUTLER
Systems and Internet Infrastructure Labratory
Pennsylvania State University
TONI FARLEY
Arizona State University
PATRICK MCDANIEL
Systems and Internet Infrastructure Labratory
Pennsylvania State University
and
JENNIFER REXFORD
Princeton University

The Border Gateway Protocol (BGP) is the *de facto* interdomain routing protocol of the Internet. Although the performance BGP has been historically acceptable, 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 adequately address security. Recent outages and security analyses clearly indicate that the Internet routing infrastructure is highly vulnerable. Moreover, the design and ubiquity of BGP has frustrated past efforts at securing interdomain routing. This paper considers the vulnerabilities of existing interdomain routing and surveys works relating to BGP security. The limitations and advantages of proposed solutions are explored, and the systemic and operational implications of their design considered. We centrally note that no current solution has yet found an adequate balance between comprehensive security and deployment cost. This work calls not only for the application of ideas described within this paper, but also for further introspection on the problems and solutions of 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

---

This work was performed while Farley and Butler were interns at AT&T Labs.
Authors' addresses: T. Farley, Information and Systems Assurance Laboratory, Arizona State University, 1711 S. Rural Road, Goldwater Center, Tempe, AZ 85287, USA; email: toni@asu.edu. K. Butler and P. McDaniel, Systems and Internet Infrastructure Laboratory, Pennsylvania State University, 344 Information Sciences and Technology Building, University Park, PA 16802, USA; email: {butler, mcdaniel}@cse.psu.edu.

1. 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) [Hawkinson and Bates 1996]. 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) [Rekhter and Li 1995]. 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 works well in practice, and its simplicity and resilience have enabled it to play a fundamental role within the global Internet [Stewart 1999]. However, BGP has historically provided few performance or security guarantees.

The limited guarantees provided by BGP often contribute to global instability and outages. While many routing failures have limited impact and scope, others 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 Virginia 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 ISP. The traffic overwhelmed the misconfigured and intermediate routers, and effectively crippled the Internet for almost two hours [Barrett et al. 1997].

Loss of connectivity on the Internet can be manifested as anything from an inconsequential annoyance to a devastating communications failure. For example, today's Internet is home to an increasing number of critical business applications, such as online banking and stock trading. Significant financial harm to an individual or institution can arise if communication is lost at a critical time (such as during a time-sensitive trading session). As the number of critical applications on the Internet grows, so will the reliance on it to provide reliable and secure services. Because of the increased importance of the Internet, there is much more interest in increasing the security of its underlying infrastructure, including BGP. Such assertions are not novel: the United States government cites BGP security as part of the national strategy for securing the Internet [Department of Homeland Security 2003].

Current research on BGP focuses on exposing and resolving operational and security concerns. Operational concerns relating to BGP, such as scalability, convergence time (the time required for all routers to have a consistent view of the network), route stability, and performance, have been the subject of much effort. Similarly, much of the contemporary security research has focused on the integrity, authentication, confidentiality, authorization, and validation of BGP data. 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 current research in interdomain routing security, exposing the similarities and differences in proposed approaches to building a more secure Internet. The next section provides a brief overview of interdomain routing and BGP. Subsequent sections examine current research addressing BGP and interdomain routing security issues.

2. OVERVIEW OF INTERDOMAIN ROUTING

The autonomous systems that collectively comprise the Internet are controlled by individual organizations. They vary in size, from large national and multinational networks owned by corporations and governments, to small networks servicing a single business or school. The *lingua franca* of the Internet is the Internet Protocol (IP) [Postel 1981], allowing communication between disparate networks. 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. These ASes are often Internet Service Providers (ISPs), providing connectivity to the global Internet for their customers. The relationship between stub, multihomed and transit ASes is illustrated in Figure 2. ISPs can form *peering* relationships with each other, where they mutually forward their customer traffic over common links.

2.1 Routing within and between Autonomous Systems

Within an AS, routers communicate with each other through the process of intradomain routing. This is accomplished using an interior gateway protocol (IGP) such as the Routing Information Protocol (RIP) [Malkin 1994], the Open Shortest Path First protocol (OSPF) [Moy 1998], and the Intermediate System to Intermediate System protocol (IS-IS) [Callon 1990]. ASes communicate routing information via an external gateway protocol (EGP). The *de facto* standard EGP in use on the Internet is BGP version 4, which has obsoleted previous versions and the original ARPANET EGP protocol [Mills 1984]. While other interdomain routing protocols and architectures exist (e.g., [Alaettinoglu and Shankar 1995] and [Castineyra et al. 1996]), we restrict our discussion to BGP. However, many 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 and by employing it, 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. BGP peers are often directly connected at the IP layer; that is, there are no intermediate nodes between them. This is not necessary for operation, as peers can form a *multi-hop* session, where an intermediate router that does not run BGP passes protocol messages to the peer. This is a less commonly seen configuration.

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 routers that communicate using EBGP, which are connected to routers

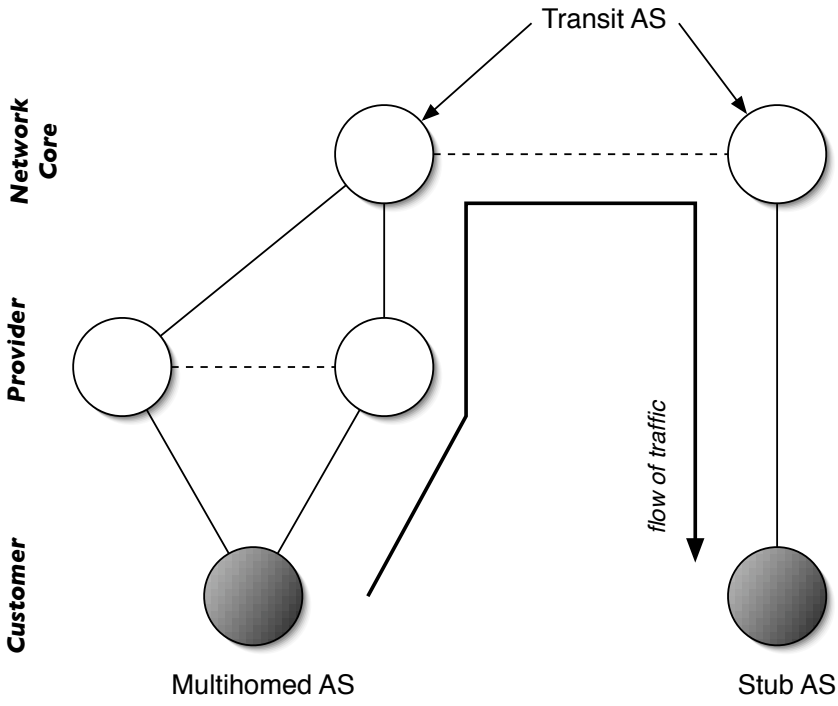

Fig. 1. Multihomed and stub ASes connect to providers who "transit" their traffic. Transit ASes forward traffic toward their destination as indicated by available BGP route information. Dashed lines in the figure indicate a peering relationship between ASes.

in different ASes, are called border routers. [1] The relationships between ASes and BGP peers are shown in Figure 2.

2.2 BGP Routing

There are currently more than 17,500 ASes in the Internet [CIDR 2004]. Each AS *originates* one or more *prefixes* representing the addresses assigned to hosts and devices within its network. 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.

BGP peers constantly exchange Network Layer Reachability Information (NLRI) — the set of known prefixes and paths for all destinations in the Internet — via UPDATE messages. Each AS advertises the prefixes it is originating to its peers. Additionally, all ASes update their routing tables based on their neighbors' NLRI, and forward the received information information to each of their other neighbors. This *flooding* process ensures that all ASes are informed of the reachability of all

[1] Routers were originally referred to as gateways, which is how the border gateway protocol got its name.

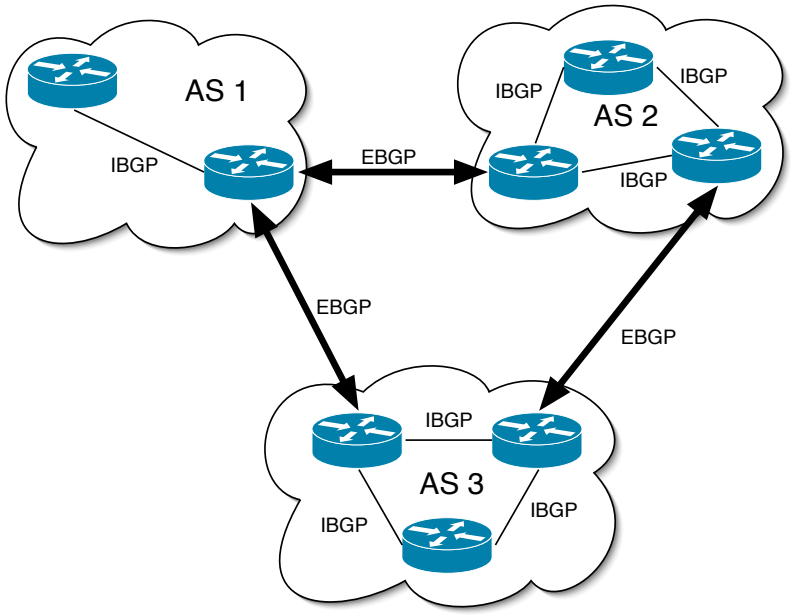

Fig. 2. BGP is used to by routers in different ASes to communicate. Two routers form a BGP session, and are peers with each other. Within an AS, routers communicate via an internal gateway protocol and form a logical mesh of IBGP links, while EBGP is used between ASes.

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.

BGP is a path vector protocol. ASes establish a *AS path* for each advertised prefix during the flooding process. The paths are vectors of ASes that packets must traverse to reach the originating AS. Path vectors are stored in a routing table and shared with neighbors via BGP. It is ultimately this information that is used to forward individual packets toward their destination.

All address ownership is the result of prefix delegation between the Internet Corporation for Assigned Names and Numbers (ICANN), regional and national registries, and organizations. ICANN and its predecessors[2] originally delegated blocks of IP addresses directly to organizations, but more recently began to delegate to address registries around the world. For example, the American Registry for Internet Numbers (ARIN) manages the IP address space delegation in North America. The *Réseaux IP Européens* (RIPE) delegates much of address space in Europe, the Middle East, and North Africa, and the Asia-Pacific Network Information Centre (APNIC) delegates IP space in Asia and the Pacific Rim. These regional registries

[2]The US Department of Commerce selected ICANN to administer the IP address space in 1993. This role was originally held by the Internet Assigned Numbers Authority (IANA), which still administers some IP namespaces (e.g., AS numbers).

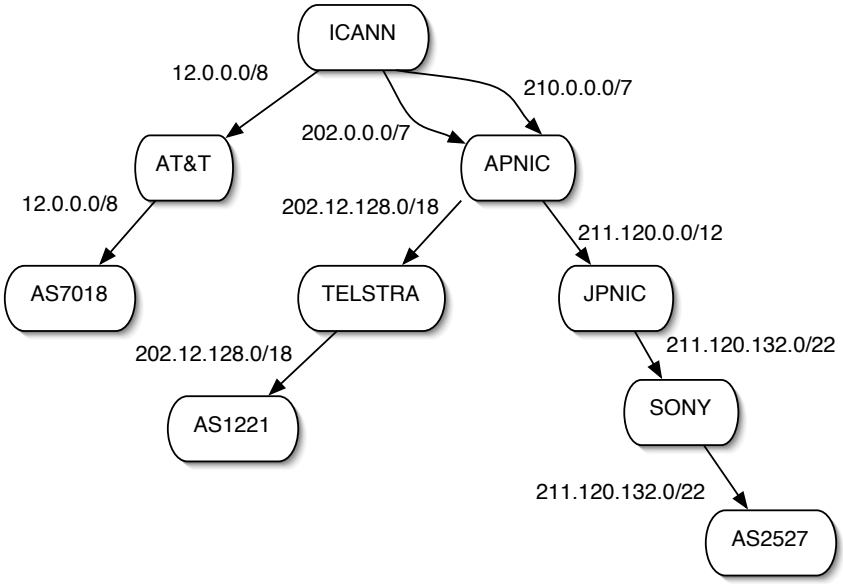

Fig. 3. A sample address delegation graph for a small part of the IPv4 address space. The address space is administered by ICANN, and hence all delegation flows from that organization.

directly delegate prefixes to organizations, or in some cases, further delegate to national registries (e.g., the Japan Network Information Center (JPNIC)), who in turn can delegate to local registries. Most networks and enterprises, however, are delegated address space from their ISPs, such as AT&T or Sprint. Once can visualize current IP address space ownership as a tree emanating from ICANN, as illustrated in Figure 3.

ASes are assigned an AS number (ASN) in a similar manner, with ICANN being the ultimate authority for delegating numbers. ASNs are used to identify the AS, and can be public or private. Public ASNs appear in BGP path vectors and are globally visible. Private ASNs can be assigned by an ISP to a customer that does not want to administer its own globally visible AS but wants to perform BGP peering with the provider, to gain benefits such as traffic engineering over multiple links.

2.3 Routing Policy

ASes are not only bound by physical relationships; they are also bound by business or other organizational relationships. When an AS owner 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. Therefore, for legal and financial reasons, it is necessary to be able to enforce SLAs at the routing policy level. BGP enforces routing policies, such as the ability to forward data only for paying customers [Halabi 2000] through a number of protocol features. Principal

among these is the assignment of attribute values in UPDATE messages.

The range of policies one 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 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]. For example, a transit AS may have several peers. The BGP policy may be configured to only allow routes to transit the network if they come from peers who have signed a contract with the organization allowing transit service. 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 [Perlman 1999]. Setting policy often involves techniques to bias BGP's route selection algorithm. For example, one of the most significant criteria BGP uses for path selection is the length of an AS path vector. This length can be modified by an organization repeatedly adding its AS number to a path, in order to discourage its use (a technique known as padding or prepending).

BGP has had success as a policy-based interdomain routing protocol. The flexibility with which polices can be specified and enforced has enabled ISPs and other organizations to fine tune their interaction, which has helped to support a more reliable and predictable Internet. In the next section, we discuss the security issues that have concerned users of BGP since its introduction.

3. A THREAT MODEL FOR BGP

The Internet was designed to enable communication between largely trusted parties. Likewise, BGP was designed to enable interdomain routing within and between trusted networks. However, commercial interests and new user communities, while essential to the growth of the Internet, have changed the nature of the network; hence, assumptions of trust present in the Internet's original design no longer hold. This is particularly true of routing — the loose collaborations that BGP was designed for are fundamentally different from interactions in the current environment. Note that changing models of trust have led to problems in other areas of the Internet. For example, the proliferation of spam [Cranor and LaMacchia 1998] is a direct result of the failure of the open model upon which electronic mail is based to be resilient to malicious entities wishing to exploit the medium for financial or other gains.

3.1 Attacks Between Peers

In order to take full stock of BGP's vulnerabilities, it is instructive to consider a threat model. This provides an outline of the sort of attacks that are desirable to prevent, and characterizes the ability of adversaries to attack the protocol. Consider the minimal case of BGP operation; that is, there are two routers communicating information to each other over a shared channel. Let us call these two parties Alice and Bob, the classical names of communicating parties in security literature. There are three potentially malicious entities in this case. Alice could be malicious, as could Bob. The channel that they communicate over could also be subverted by a malicious third-party, who we call Charlie. (If both Alice and Bob are malicious, the protocol is of course doomed to failure – routing only works if at least some entities are good.) Alice or Bob could be malicious entities, either by choice or unwittingly,

due to subversion by an external attacker (i.e., following router compromise). The following considers the attacks possible within this limited scenario.

3.1.1 *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 between each other will not be seen by any other parties. As we previously described, however, the channel over which they communicate may have been subverted by a third party. Alice and Bob's messages between each other could be possibly observed by the attacker, Charlie. Charlie could be *eavesdropping* 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 always sensitive, many service providers and large organizations have business relationships (e.g., undisclosed peering arrangements) that can be inferred by the BGP traffic [Spring et al. 2002]. These relationships are often considered confidential trade secrets, and having an eavesdropper determine them, perhaps for corporate espionage purposes, is highly undesirable. These *passive* attacks are not unique to BGP, but are true of any protocol that uses TCP as an underlying transport without additional security infrastructure (e.g., session hijacking [Traina 1995]).

3.1.2 *Attacks Against Message Integrity.* An additional risk occurs if Charlie, the attacker, does not merely passively listen to updates, but becomes an active, unseen part of the communications channel. Charlie can become a *man in the middle* between Alice and Bob, and tamper with BGP messages. One method of tampering is *message insertion*, where Charlie inserts forged BGP messages into the message stream. This can have the effect of introducing incorrect routing information. It can also force the connection between Alice and Bob to shut down, as erroneous BGP messages will abort the session. Charlie can also affect the message stream through *message deletion*, where he selectively removes messages. BGP relies on keep-alive messages being periodically sent, and if they are not received, the connection will be closed. Another method of tampering is *message modification*, where Charlie intercepts a message in flight and alters its contents before forwarding it. Finally, Charlie can launch a *replay* attack, where he records messages between Alice and Bob 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's routers, forcing them to crash and go offline.

3.1.3 *Session Termination.* A consequence of modifying messages is the ability to terminate a BGP session. The following example demonstrates how an attacker takes advantage of the protocol's state machine model. Events received by BGP speakers cause their internal state to change, causing them to expect certain messages and react to them in a different manner. For example, if Alice and Bob are setting up a BGP session, Alice sends Bob an OPEN message and transitions into the OpenSent state. When Bob receives this message, he responds with an OPEN message. Upon reception of this message, Alice changes to the OpenConfirm state. When the session has been completely set up, both Alice and Bob are in the Established state, the state that BGP regularly operates in. If the attacker Charlie inserts an OPEN message at this point, the session between Alice and Bob

will be closed, because it violates the expected input. Another way to close the session is by forging a NOTIFICATION message, which indicates an error has occurred. When either Alice or Bob receives this message, they will terminate the BGP session. The BGP state machine [Rekhter and Li 1995] introduces several vulnerabilities [Murphy 2004]. For example, the state machines require that the protocol be reset following any fault. As detailed in the following sections, such features can be exploited to decrease the stability or availability of the Internet.

3.2 Larger Scale Attacks

BGP is a distributed protocol run by hundreds of thousands of routers. Hence, there are many points at which an adversary can mount an attack. Moreover, each autonomous system is indirectly connected to every other AS in the Internet. Adversaries can affect routers and networks far removed from their peers by exploiting this scale and interconnectedness. The form and results of these attacks is considered in the following sections.

3.2.1 *Fraudulent Origin Attacks.* An autonomous system can advertise incorrect information through BGP UPDATE messages passed to routers in neighboring ASes. A malicious AS can advertise a prefix originated from another AS and claim that it is the originator, a process known as *prefix hijacking*. Neighboring ASes receiving this announcement will believe that the malicious AS is the prefix owner and route packets to it. The real originator of the AS will not receive the traffic that is supposed to be bound for it. If the malicious AS chooses to drop all the packets destined to the hijacked addresses, the effect is called a *black hole*. This attack makes the hijacked addresses unavailable. Note that the outage outwardly looks like any other kind of outage, and is often difficult to diagnose. If the malicious AS chooses to forge all addresses in a block using hosts and devices within its control, the affect may be much more severe. Unless properly authenticated using some other security service, one can impersonate all of the services and resources of the hijacked address space. The malicious AS can then analyze the traffic it receives, possibly retrieving sensitive information such as passwords.

One particularly virulent method of spreading false information is through *prefix deaggregation*. This occurs when the announcement of a large prefix is fragmented or duplicated by a collection of announcements for smaller prefixes. BGP performs *longest prefix matching*, whereby the longest mask associated with a prefix will be the one chosen for routing purposes. For example, if the prefixes `12.0.0.0/8` and `12.0.0.0/16` are advertised, the latter prefix, which corresponds to a more specific portion of the address block, will be chosen. Deaggregation harms the performance of BGP and indirectly the network by increasing the size of BGP tables and flooding the network with redundant, and sometimes incorrect updates.

If an AS falsely claims to be the origin of a prefix and the update has a longer prefix than others currently in the global routing table, it will have fully hijacked that prefix. Not only will neighboring routers believe this update, but they will flood the false update to its peers. This flooding eventually propagates the attack throughout the Internet.

3.2.2 *Subversion of Path Information.* Another method that a malicious AS can use to spread misinformation is to tamper with the path attributes of an UPDATE

message. As previously mentioned, BGP is a path vector protocol, and routing to destinations is performed based by sending packets through the series of ASes denoted in the path string. An AS can modify the path it receives from other ASes by inserting or deleting ASes from the path vector, or changing the order of the ASes, in order to create routing delays or to allow the malicious AS to alter network traffic patterns. By altering attributes in an UPDATE message, such as the multi-exit discriminator (used to suggest a preferred route into an AS to an external AS) or the community attribute (used to group routes with common routing policies), traffic engineering and routing policy can be undermined.

Another potent attack alters the paths to transit a malicious AS. In addition to correctly transiting the data, the malicious AS eavesdrops on application traffic of the originating AS. Such data, if not properly secured, could expose an enormous amount of information about the activities of the victim.

3.3 Denial of Service

Many of the attacks above can be considered denial of service attacks. Black holing a route, for example, causes denial of service for that prefix, and subverting the path can also lead to service delays or denials. For example, a sufficiently long route can cause the time-to-live (TTL) of a packet to be exceeded. In the two peer case, denial of service has also been considered by a remote attacker using erroneous or false BGP messages to shut down a connection. Since BGP uses TCP as a transport protocol, it is subject to TCP attacks as well. For example, the TCP RST attack can cause a remote attacker to be able to reset a TCP connection between two BGP peers. Additionally, TCP is vulnerable to the SYN flood attack, where the three-way handshaking process is initiated but never completed (the attacker never acknowledges the open handshake). The victim will run out of connection state memory[3] and either be unable to perform any TCP transactions or crash altogether. These attacks are harmful enough to the individual routers, but become even more consequential when the distributed case is considered. If a router goes offline, then when it comes back online, its routing table will need to be recreated, and it re-announces all of the prefixes it is originating, a process known as a *table reset*. The neighboring routers dump their BGP tables to the peer that has just come online so that it has full data for making its routing decisions. Sifting through this information places a considerable computational burden on the router, and delays processing of normal traffic. If the router is continually knocked offline, the routes it advertises will disappear and reappear in peer routing tables. This is called *route flapping* and is detrimental to all routers, as extra computation and reconfiguration of routes becomes necessary if this happens often. In order to lower the burden, unstable routes are often penalized through a process called *route dampening*. Neighboring routers will ignore advertisements from the router for an increasing amount of time, depending on how often the route flapping occurs. Suppression of these routes can be a highly effective denial of service attack.

Attacks against the underlying protocols and links will also deny service to BGP.

[3]A finite amount of memory is set aside for connection state in most implementations of TCP. How a particular device responds to the exhaustion of this resource is implementation dependent, but many simply reboot the device.

Examples of these include Internet Control Message Protocol (ICMP) magnification attacks such as Smurf [Baltatu et al. 2000], where ping packets are spoofed with the source address of the victim and directed at broadcast destinations, which can generate many more responses towards the victim. With enough nodes participating, the links to the victim can become saturated and not allow any other traffic, including BGP keep-alive messages, through, forcing a session termination. Additionally, physical attacks against the underlying network circuits or the routers themselves can influence BGP's behavior. For example, Bellovin and Gansner [2003] showed how an adversary could arbitrarily alter traffic routing by (only) severing links between BGP speakers.

3.4 Misconfiguration

The effects of misconfiguration are often the same as an attack. BGP is complex to configure, and even minor errors can create widespread damage. 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, between 200 and 1200 prefixes, equivalent to 0.2-1% of all prefixes in the global routing table, are misconfigured. It also identifies two forms of misconfigurations that can be globally visible:

(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).

An example of router misconfiguration that led to widespread damage occurred in October 2002 with the Internet service provider WorldCom. [Slater III 2002]. Improper filtering rules added to a router caused the routing tables of WorldCom's internal infrastructure to become flooded with external routing data; in other words, the routers within the AS were subject to much more data than they should have been. Faced with this additional burden, 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. As the routers came back after crashing, they were flooded with the routing table information by their neighbors. The flood of information would again overwhelm the routers and cause them to crash. This process of route flapping served to destabilize not only the surrounding network, but the entire Internet.

Malicious prefix deaggregation can allow adversaries to take over a prefix by advertising a more specific prefix block. The canonical example occurred in 1997, when misconfigured routers in the Florida Internet Exchange (AS7007) deaggregated every prefix in their routing table and started advertising the first /24 block of each of these prefixes as their own. A /24 block is the smallest prefix generally allowed to be advertised by BGP, and because of its specificity, routers trying to reach those addresses would choose the small /24 blocks first. This caused backbone networks throughout North America and Europe to crash, as AS 7007 was overwhelmed by a crush of traffic and the routes it advertised started flapping [Misel 1997]. This was not a malicious attack, but a mere error made by the network operators. Consider that a well-planned, targeted, malicious attack on BGP could do very serious harm to the network infrastructure.

3.5 Limitations of BGP

Murphy [2004] suggests that there are three primary limiting factors of BGP that lead to the vulnerabilities described in the proceeding sections:

—BGP does not protect the integrity, freshness and origin authentication of messages. Integrity ensures that a message has not been tampered with, freshness ensures that the recipient has actually received a new message, not one that has been replayed, and origin authentication refers to the verification that the originator of the update message is not fraudulent.

—BGP does not validate an AS's authority to announce reachability information. This is related to path subversion, as an AS can currently announce that it has the shortest path to a destination by forging the path vector, even if it is not part of the destination path at all.

—BGP does not ensure the authenticity of the path attributes announced by an AS. Altering the path attributes is another way that a malicious AS can impair or manipulate the routing infrastructure.

Moreover, analyses of BGP of the end-to-end behavior of Internet show that that routing can and often does experience substandard, and even *broken* behavior. Broken behavior is often manifest as IP packets being grossly misrouted. For example, Paxson [1999] reports that packets originating 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. 2001].

3.6 Consequences of Attacks

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, and other undesirable outcomes can result from an attack. Attacks can be used in concert to amplify their effect or to enable further malicious activity. The generic consequences of routing threats are further discussed in [Barbir et al. 2003]. Examples of these consequences include the disclosure of confidential information, deceptive or incorrect information introduced into the network through message modification, the disruption of network activity through denial of service attacks, and the usurping of router services and functions.

Consider the ramifications of a dysfunctional routing system under attack. An individual router is subject to being overloaded with information, knocked offline or taken over by an attacker. An autonomous system can have its traffic blackholed or otherwise misrouted, and packets to or from it can be grossly delayed or dropped altogether. Malfunctioning ASes harm their peers by forcing them to recalculate routes and alter their routing tables. As the misconfiguration examples have shown, these events can disrupt international backbone networks and have the potential to bring a large part of the Internet to a standstill. From the individual level of an organization's traffic being stolen to the worldwide scale of IP traffic being globally subverted, the threats against BGP are a matter of grave concern to anybody reliant on the Internet.

4. BGP SECURITY SOLUTIONS

BGP security is an active area of research. Because this activity is relatively new, no solution have been universally deployed in the Internet. Network administrator currently mitigate some attacks by implementing local countermeasures. The following section reviews the tools used in the Internet to protect BGP. The subsequent sections describe proposal architectures and countermeasures for BGP security.

4.1 BGP Security Today

Protecting the TCP connection is an easy way to mitigate attacks on BGP sessions. A popular and inexpensive countermeasure against attacks on TCP is the use of message authentication codes (MACs). Recent enhancements to BGP suggest the use of a TCP extension that carries an MD5 digest [Rivest 1992] based MAC. An MD5 keyed digest [Krawczyk et al. 1997] 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 be 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 [Heffernan 1998], which addresses many of the problems of spoofing and hijacking inherent to TCP [in the TCP/IP Protocol Suite 1989; Green 2002].

Known more generally as cryptographic hash algorithms, digest algorithms compute a fixed-length hash value from an input text. The 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.[4]

The MD5 digest mechanism requires that a *shared secret key* be configured manually at each session end-point. This approach is limited in that maintaining shared secrets between potentially thousands of routers concurrently is immensely difficult. Moreover such secrets, if not replaced frequently, are subject to exposure by cryptanalysis.

4.1.1 *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 [Kent and Atkinson 1998c; Thayer et al. 1998]. 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 IPsec Internet Security Association and Key Management Protocol (ISAKMP) defines a framework for key management and negotiating security services [Maughan et al. 1998] while the Internet Key Exchange (IKE) protocol

[4]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.

deals with the issues of dynamic negotiation of session keys [Harkins and Carrel 1998]. The IPsec Authentication Header protocol (AH) [Kent and Atkinson 1998a] and Encapsulating Security Payload (ESP) protocol [Kent and Atkinson 1998b] implement packet level security with differing guarantees. All of these services work in concert to establish and maintain the secret keys used 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) []. If properly configured 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.

4.1.2 *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 [Gill et al. 2004]. 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.

4.1.3 *Defensive routing policies.* 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. The policies filter prefixes that are documented special use addresses (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 [CIDR 2004] which many organizations use to filter announcements. Filtering is also used to removing conflicting announcements. For example, announcements containing private ASes [Stewart et al. 1998] or from unexpected downstream ASes are automatically dropped by some BGP speakers.

A policy of careful ingress and egress filtering greatly aids in maintaining security for both the local AS and its neighbors, and is widely held to be the most widely deployed and effective BGP security measure. Filtering is not a replacement for a strong security architecture. The filtering rules are fundamentally limited by the the heuristics it represent, and can only remove announcements which are overtly bad.

BGP attributes are another potential vehicle for an attack. For example, MEDs

can be used by an adversary to control the egress point of an AS. Rexford et. al. show how this vulnerability is used to force an AS to perform cold-potato routing [Feamster et al. 2004]. The community string is an equally dangerous attribute. These strings are used as internal tags to indicate how the route should be treated, and are hence be abused by an adversary to influence the propagation and selection of routes. Other attributes such as "origin type" are used in the route selection process, and also may be misused. Routers frequently defend against all these attacks but clearing or validating the attribute value, e.g., clearing MEDs and community strings, or zeroing the origin type values.

4.1.4 *Routing Registries.* A route registry is a centralized repository of routing policy information [Bates et al. 1995]. 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. However, to use a registry, one must first be assured that the registry itself is secure. Villamizar et al. [1999] 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 peering data, policies and routes to be proprietary information (and are thus reluctant to sharing it), though tools such as Rocketfuel [Spring et al. 2002] provide accurate maps of internal topology, and algorithms exist for inferring customer and peering relationships [Gao 2001; Subramanian et al. 2002]. The community-supported registry approach is also limited in that the registry itself is often untrusted; a malicious registry manipulate the route information at will. Information in routing registries also tends to decay quickly because of a lack of clear incentives for organizations to maintain their information [Griffin 2003].

4.2 BGP Security Architectures

Recent efforts within the standards bodies and in research community have attempted to provide comprehensive architectures for BGP security. Each architectures provide an explicit threat model and suite of security services. The following sections consider several of these architectures.

4.2.1 *S-BGP.* Secure BGP (S-BGP) was the first comprehensive routing security solution targeted specifically to BGP [Kent et al. 2000]. The S-BGP protocol and its associated architecture are currently under consideration for standardization by the Internet Engineering Task Force (IETF), the organization that provides Internet standards. Implementations of S-BGP exist, and its authors are actively experimenting with its use in operational networks.

A primary element of S-BGP is its use of *public key certificates* to communicate authentication data. Public key certificates bind cryptographic information to an identity such as an organization. Anyone in possession of the public key certificate can validate information *digitally signed* with the private key associated with the public key. As the name would imply, the public key is widely distributed, and the private key is kept private [Rivest et al. 1978]. A *public key infrastructure* (PKI) is a system for issuing, authenticating and distributing certificates.

S-BGP implements security by validating the data passed between ASes using public key certificates. S-BGP supports a pair of PKIs used to delegate address

space and AS numbers, as well as to associate particular network elements with their parent ASes [Seo et al. 2001]. One PKI is used to authenticate address allocations through a hierarchy stretching from organizations to the providers and regional registries allocating them space, ultimately leading to ICANN (the ultimate authority for address allocation). The second PKI is used to bind AS numbers to organizations and organizations to routers in their network. This is accomplished through issued 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 verify it came from the AS using the certificate. Because of the properties of the underlying cryptography, no adversary could have generated the signature, and hence it could have only come the signing AS.

All data received by a AS in S-BGP is validated using the certificates in the dual PKIs. Address ownership, peer AS identity, path-vectors, policy attributes, and control messages are all signed (and sometimes counter-signed) by the organizations or devices that create them. Because this allows the receiver of the data to unambiguously authenticate the routing information, they can easily detect and remove forged data. However, because of the amount of data and number of possible signers, validation can be extremely costly [Nicol et al. 2002]. 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 choose 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 ICANN to the advertising AS.

Route attestations are distributed within S-BGP in a modified BGP UPDATE message as a new attribute. To simplify, route attestations are signed by each AS as it traverses the network. All signatures on the path sign previously attached signatures (e.g., are nested). Hence, the validator can validate not only the path, but can validate that *a)* path was traversed the ASes in the order indicated by the path, and *b)* no intermediate ASes were added or removed by an adversary. Figure 4 shows a simplified use of route attestations as they propagate between routers.

4.2.2 *Secure Origin BGP.* Secure origin BGP (soBGP) seeks flexibility by allowing administrators to trade off security and protocol overhead using protocol parameters. Similar 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 selected protocol parameters and local network topology. 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 new

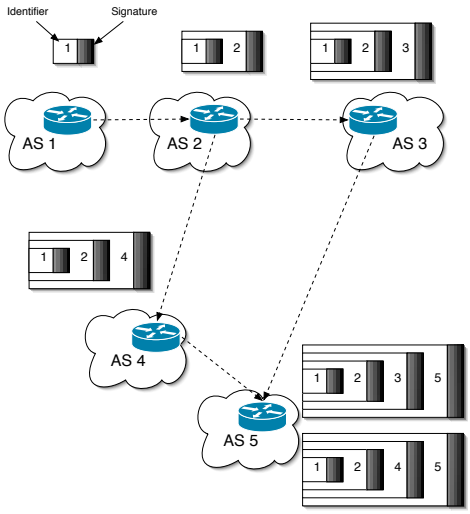

Fig. 4. Route attestations in S-BGP. As UPDATE messages are passed between peers, the receiving peer signs the received message before passing it to another neighbor. The result is an "onion-style" attestation that contains signatures from all routers along the path.

SECURITY BGP message.

soBGP routers use the topology database to validate received routes. Each AS signs and distributes its local topology (e.g., peers) through the topology certificate. The certificates from ASes form the global topology database. The database is used to sanity check received routes: any UPDATE with a path that violates the AS topology is demonstrably bad and dropped. Kruegel et al. [2003] extend this approach by using other heuristics in detecting anomalous paths (e.g., multiple entrances into core ASes, strange geographic routes, etc.).

Validating signatures is a computationally expensive operation. soBGP tries to mitigate this cost in the presence of limited resources by authenticating 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 cost. 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.

4.2.3 *Interdomain Route Validation.* 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 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 Figure 5). The local IRV server determines correctness by directly querying the IRV server in the relevant AS for validation of the route information. Where validation from multiple ASes is needed, i.e., to

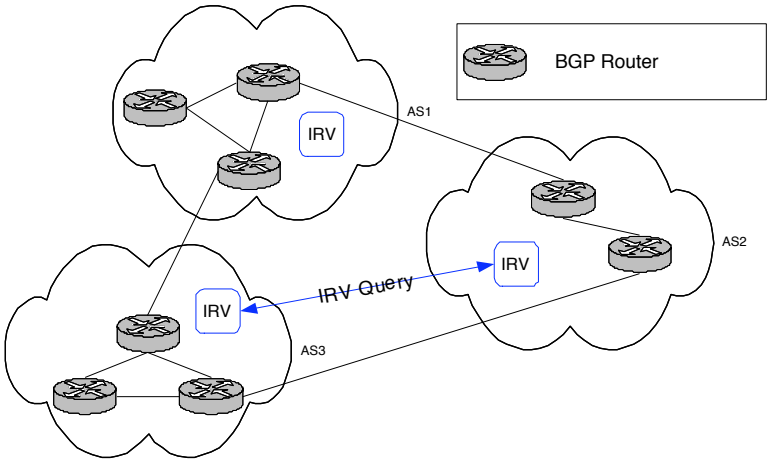

Fig. 5. ASes running the IRV protocol query the appropriate authorities for validation of received routing data.

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. A BGP speaker decides which data to trust, which to ignore, and which to validate via IRV query based on local policy. Hence, the amount of effort expended in validating data is exactly as is required. IRV servers are similar to route registries, but manage information only from the parent AS. IRV approach is arguably more likely to be successful than registries because the AS retains control over the data, and hence is more likely to keep it fresh, accurate, and available.

Where available, a secure underlying network layer (e.g., IPsec) or transport layer (e.g., TLS [Dierks and Allen 1999]) is 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 algorithm for determining when and how an UPDATE message should be validated is chosen by each AS. 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 suspecious (as determined by heursitics). Caching previous queries can also improve performance, while storing received route advertisements and withdrawals can be used for debugging and failure detection.

4.3 Experimental Systems

A number of works have sought solutions to the myriad security issues in inter-domain 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 sketches a number of these works in broad detail. We being in the following subsection by reviewing the some of the early works in BGP security.

4.3.1 *Early Approaches.* Radia Perlman thesis [Perlman 1988] was the first significant work addressing routing security. The dissertation is particularly notable for examining *Byzantine behavior* within routing protocols. Byzantine behavior occurs when any routing element exhibits arbitrarily faulty or malicious behavior. For a protocol to provide security in this environment, it 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). Perlman develops a link-state protocol that satisfies the properties for Byzantine robustness. Perlman's link-state solution does not effectively scale to networks the size of the Internet, and hence is not suitable for interdomain routing. However, some conceptual elements of Byzantine robustness are present in almost every proposed definition of BGP security. For example, assessing the validity (as defined by Perlman) of received routes and policy is the central goal of the three architectures defined in the preceding section.

Smith and Garcia-Luna-Aceves [1998] 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.

Smith and Garcia-Luna-Aceves's countermeasures are similar to those provided by S-BGP, where the session encryption and sequencing provides confidentiality and ordering, the peer signatures guarantee authenticity of the full BGP path. However, 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.

4.3.2 *IDRP Countermeasures.* Prior to the creation of BGP version 4, Kumar and Crowcroft [1993] provided an analysis of threats to interdomain routing and described security mechanisms used in then proposal IDRP protocol [ISO 1992]. Designed as a superset of BGP and EGP, IDRP is a interdomain routing path vector protocol. The protocol uses a encrypted checksum transmitted with all routing messages transmitted 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, however, that malicious entities masquerading as

sources will be unsuccessful in a hop-by-hop routing protocol, neglecting to consider prefix hijacking. The authors further asserted 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 made cryptographic operations feasible. Hence, while this proposal highlighted important requirements for routing security, it is not appropriate for current networks.

4.3.3 *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. [2000] propose a suite of protocols that also provide security at the IP layer. As with the Smith approach discussed above, sequence numbers and message MACs are used to ensure integrity and ordering. Gouda et. al. extend this approach by suggesting a Diffie-Hellman [Diffie and Hellman 1976] style protocol[5] that uses public key certificates to negotiate and refresh the secret keys shared by peers. Due to is wide deployment and flexibility, IPsec has supplanted this proposal as the way to perform hop integrity.

4.3.4 *MOAS Detection and Mitigation.* An IP prefix should generally only be originated by a single AS [Hawkinson and Bates 1996]. A multiple origin AS (MOAS) conflict occurs when a prefix is simultaniously originated by more than on 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 showed potential causes included 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]. An enhancement BGP was proposed that uses community attributes [Chandra et al. 1996] to distinguish between valid and invalid MOAS conflicts [Zhao et al. 2002] in responses 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 an 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 advertisements can be forged or altered by malicous routers. However, the authors suggest that forged routes can be detected by flagging prefixes recieved with multiple, conclicting AS lists.

4.3.5 *Origin Authentication.* Origin Authentication (OA) is a method of validating address ownership such that prefix hijacking and related attacks are not possible. One effort directly investigates origin authentication (OA) by examining the semantics, design and application of OA services [Aiello et al. 2003]. 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 [Merkle 1980; M. Naor and

[5]Diffie-Hellman protocols use public key cryptography to negotiate shared secrets between parties over an untrusted media, e.g., a public network. This protocol and its variants are the most widely used protocol for performing *key negotiation*.

K. Nissim 1998] 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 delegation hierarchy by extracting the nested announcements made within the protocol. They found that the delegations were very stable over time (see Section 5.2.1). This made them ideally suited to a class of proof structures based on Merke hash trees [Merkle 1980]. A simulation shows that on-line origin authentication is possible using this construction, a feat which was previously thought to have been too computationally expensive to be feasible.

The Pretty Secure BGP (psBGP) system introduces an addresses origin authentication service within a larger comprehensive architecture for BGP security [Wan et al. 2005]. ASes are validated in psBGP using a PKI similar to that suggested for S-BGP [Seo et al. 2001]. Path authentication is performed using an optimized version of S-BGP introduced by Nicol et al. [2002]. The central philosophy of their work is that while ASes can be managed within a PKI (because their 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. 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 s a very weak form of origin authentication: *any* AS can bear witness to the validity of an origin claim.[6] The assumption that any two of 18,00 ASes will not collude is seen as somewhat difficult to support in the general Internet. Moreover, psBGP requires an AS place its trust in the alien ASes to regulate IP addresses, most of whom you have no relationship or often knowledge.

Kruegel et al. [2003] consider the use of intrusion detection to identify forged origin announcements, and discover several metrics used to identify bogus announcements (e.g., strange aggregation, 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.

4.3.6 *Path Authentication.* Hu et al. [2003] proposed 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 indentified validator known as the *destination AS*. 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, NLRI, etc.). The MAC is included in the UPDATE and propagated using BGP. Each of

[6]Tan et. al. consider other modes in which it may require *k-out-of-n* peers asserting validity for the origin to be accepted. However, this is only useful in weeding out highly connected colluding pairs.

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 preceeding hop. Upon receiving the MAC, the destination recursively validates the 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. The destination AS can validate all the MACs because it knows all the secret keys.

Note that there are many destination ASes. Creating shared secrets between all ASes is difficult, and possibly operationally impossible in the current Internet (e.g., with n ASes, there are $O(n^2)$ shared keys required with the naive solution). The authors suggest that such costs can be mitigated by using a protocol similar to TESLA [Perrig et al. 2002] that provides public key semantics using symmetric key cryptography. Specifically, the TESLA construction allows a single secret key operation to behave like an unforgeable signature in that it simultaneously authenticates the source of a message to many receivers. To simplify, the TESLA protocol creates and transmits MACs and keys known only to the sender. The key is *released* (e.g., transmitted) by the sender at some predetermined point of time in the future. The important point is that if the receiver can guarantee that it received the MAC before the key was released, then it knows that the advertisement is authentic (because no adversary could have generated the MAC without the key). The timed key release approach that TESLA is based on was originally suggested as a possible security solution of link state routing by Cheung [1997] in their specification of the Optimistic Link State Verification protocol.

4.3.7 *SPV.* Hu et.al. extended their work in path authentication in the Secure Path Vector Protocols(SPV) [Hu et al. 2004]. SPV implements path validation using a string of one-time signatures [S. Even and O. Goldreich and S. Micali 1996; C. Wong and S. Lam 1999] generated from a single root value. Also known as off-line signatures, one-time signatures allows the signer to perform the heavyweight cryptographic operations prior to use, and the later signing operation is very fast. 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 initiator generated an initial validation token to verify the one-time signatures, and ultimately the path. The operation of SPV is extremely lightweight, where hashing is used as the primary cryptographic mechanisms. However, this efficiency comes at a cost; SPV is a very complex protocol involving the manipulation and communication of a significant amount of state. 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 are appropriate in the larger Internet.

4.3.8 *Whisper.* The Whisper protocols [Subramanian et al. 2004] are designed to validate the initial source of path information. The protocol do 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 Hu's cumulative authentication mechanism described above. A random value is initially assigned to each prefix by the originator. The value is repeatedly hashed at each hop as it is propgated 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 applicaiton of the hash function). Stronger protocols are proposed that increase security by making the initial value less knowable using heavyweight modular exponentiation. One variant uses a construction similar to RSA [Rivest et al. 1978][7], 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.

4.3.9 *Listen.* Also introduced by Subramanian in the Whisper paper [Subramanian et al. 2004], the Listen protocol is also used to identify inconsistent route advertisements. The Listen service monitors TCP traffic flows and determines 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, monitoring for ACK packets is not important. If a threshold of hosts in a remote prefix do not respond, the protocol assumes that the route is not verifiable, and is flagged as possibly being black-holed, misconfigured, or possibly hijacked.

5. EVALUATING BGP SECURITY

We now consider the degree to which the solutions presented in the preceding section address the threat model defined in section 3. Principally, we consider the the limitations and trade offs of these solutions and draw general conclusions about the applicability of each in current and future BGP environments.

5.1 Defenses Against Peer Attacks

Summarized in Table I, we begin by considering the features and limitation of the proposed BGP session security solutions. Recall that peer attacks include both passive activity, such as eavesdropping, and actively malicious activities, such as modifying BGP messages. Both forms of attacks are mitigated by IPsec, which introduces authenticated sequence numbering, distribution of shared keys between peers, and encryption. IPsec is assumed to be the underlying network mechanism with S-BGP, soBGP, and IRV (the latter can also use TLS). The IPsec AH mode protects against replay attacks through the use of sequence numbers, and protects message integrity by calculating a message authentication code using a hashing function such as MD5 or SHA-1. The IPsec ESP mode provides AH data integrity

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

	Integrity	Confidentiality	Replay Prevention	DOS Prevention
IPsec (ESP) [1998b]	yes	yes	yes	yes
IPsec (AH) [1998a]	yes	no	yes	yes
MD5 Integrity [1998]	yes	no	yes	no
HOP Protocol [2000]	yes	no	yes	no
GTSM [2004]	no	no	no	no
Smith .et al. [1998]	yes	yes	yes	no

Table I. BGP peer session security solutions - requirements (columns) relate to the guarantees provided for the AS to AS peering sessions.

and authenticity in a similar manner to AH, and additionally introduces further defenses against eavesdropping, e.g,. confidentiality. The hop integrity protocols proposed by Gouda et al. [2000] duplicate the services of IPsec: Diffie-Hellman key negotiation, data integrity, and data authentication are provided.

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 both two routers. Unlike the IPsec IKE protocol which dynamically negotiates secret keys, manual configuration of MD5 keys can place significant operational burden on network administrators.

Two of the countermeasures specified by Smith and Garcia-Luna-Aceves [1996] protect the confidentiality and integrity of BGP through the encryption and authenticated sequence numbering; however, use of these extensions require 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.

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 very limited.

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 are certainly harmful to BGP

Solution Definition			Security Services		
System	In Use	Style	Topo. Auth.	Path Auth.	Origin Auth.
Route Filtering	yes	anomaly	*weak*	*weak*	*weak*
Route Registries	yes	anomaly	*weak*	*weak*	*weak*
S-BGP	no	crypto	**strong**	**strong**	**strong**
soBGP	no	crypto/anom.	**strong**	none	**strong**
IRV	no	crypto/anom.	**strong**	**strong**	**strong**
Origin Auth. (Aiello et. al.)	no	crypto	none	none	**strong**
Path Auth. (Hu et. al.)	no	crypto	**strong**	**strong**	none
SPV	no	crypto	**strong**	**strong**	none
Listen	no	anomaly	none	none	*weak*
Whisper	no	anomaly	none	*weak*	none

Table II. Global BGP security solutions - requirements (columns) relate to the guarantees provided for global AS data. *Deployed* 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). We a system is strong if it provides authenticity guarantees, and weak if it received data is probabilistically authentic/correct.

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.

The prohibitive favorite solution for BGP peer session security has become IPsec. At this point, IPsec is ubiquitous, well understood, and easy to configure. Proposed solutions such as the Hop protocol and Smith .et.al countermeasures provide a subset of IPsec functionality using specialized protocols. IPsec was not widely available at the time most 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 can be deployed in their environment, i.e., IPsec.

5.2 Defending Against Larger-Scale Attacks

The most damaging attacks on BGP are those that manipulate prefix origins and path vectors. We summarize the various BGP security solutions that address these threats in Table II and examine their effectiveness in the following subsections.

5.2.1 *Defenses Against Fraudulent Origin Attacks.* Heuristic mechanisms to origin authentication are attractive because they require little cooperation from foreign ASes. For example, the MOAS attribute extension helps identify MOAS conflicts possibly indicative of prefix hijacking. The mechanism is limited in that it does not provide security, but simply reflects markings by a (potentially malicious) AS. The Listen protocol is also attractive in that it unilaterally detects origin misuse.

However, it does not provide timely or authoritative information and requires a significant amount of state. As such, it is unlikely to gain favor within the operational community.

Prefix hijacking has been a focus of several BGP security efforts. S-BGP uses address attestations within an authenticating PKI to ensure that organizations have the authority to originate their advertised prefixes. In this system, a recipient of a BGP UPDATE message traces the address delegations from the organizational level to ICANN (the root issuer of address space) in order to prove the legitimacy of an advertisement. soBGP uses similar certifications to provide authorization of an address delegation.

The address attestations introduced by Aiello et al. [2003] are appropriate for both S-BGP and soBGP. In addition to identifying optimizations, Aiello et al. evaluated the stability and structure of address management and uncovered two important facts. First, address space is very stable: 70-90% of the address prefixes advertised on the Internet did not change origin over the 6 month experiment. Second, the delegation is very dense: 80% of the of address space is delegated by 16 organizations, and 90% is delegated by 122 organizations. This stability and density demonstrates that attestations are long lived and (most) come from a few organizations. This indicates that proof systems such as those defined by S-BGP and soBGP are likely to be successful in practice.

Origin authentication is addressed today by route filtering; each AS identifies the set of prefixes that it knows the proper origins for. This is highly effective in managing misconfiguration, but will not be effective under concerted attack. The long term solution is for the operational community to regain control of the address space. That is, we must introduce an infrastructure for the delegation and maintenance of the address space. The authors of the S-BGP, soBGP and later the origin authentication study have identified the appropriate and efficient structures and mechanisms for communicating address ownership and delegation correctly. It is now incumbent on the Internet community to solving the operational challenge; finding and certifying the ownership of the IP address space. It is unclear how the resulting certificates will be managed. Protocols such as S-BGP, soBGP, and IRV each have operational trade-offs that will be evaluated more when address ownership data becomes available.

5.2.2 *Defending Against Path Manipulation.* Globally acceptable solutions to path authentication have thus far been illusive. S-BGP is costly both in terms computation and storage. For example, Nicol et.al. assessed the cost of implementing S-BGP in via simulation [Nicol et al. 2002]. They found that the early versions of the protocol were often costly on a global scale: path convergence times would double in some cases. However, the simulation did not include recent optimizations (i.e., validate paths only when they are selected as best paths), and may not be a perfect reflection of S-BGP performance. The storage costs have also been cited as prohibitive: in a normal environment, S-BGP attestations are likely to exhaust most or all of the available router storage resources.

The IRV system tries to remove the computational and storage costs from the critical path of routing. In IRV, validation of path information is discretionary. 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 to be queried). All of this occurs between IRV servers, and not routers. Hence, costs are controllable by the AS, and resources demands are largely external to the routers. 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 appropriarate AS IRV server. This present problems both in bootstrapping the process and in recovering from outages. Solutions to this problem include optimistic routing (e.g., use received routes immediately and validate possible), AS collaboration (e.g, exchange routing data via gossip-style protocols [B. Baker and R. Shostak 1972]), and using static routes to IRV servers.

In soBGP, there is a static topology graph created through dissemination of policy certificates. This approach allows routers to compare the path of a received route with the contents of their topology databases. The route is penalized or dropped if this information is inconsistent. The limitation of this approach is that it is fundementally static. Additional infratructure is required to ensure that the topology updates are synchronized accross all ASes. Moreover, forged paths that are consistient with the routing topology will be accepted.

The cumulative authentication mechanism described by Hu et al. [2003] guarantees path integrity. An AS with position n in an update's path vector cannot manipulate the $n-1$ ASes preceding it, because reversibility is cryptographically prevented. This AS itself could be malicious, but it can only forge the contribution it makes to the path. Key management in this protocol is also problematic. Manual configuration does not scale, and the TESLA approach requires synchronized timing information, which is difficult to scale. Hence, it does not appear that the solution is likely to be feasible on the Internet. Similarly, the complexity and security model of SPV make it an unlikely candidates for wide adoption.

Whisper relies on RSA style constructions to provide path integrity. However, the construction as described is susceptible to the *common modulus* RSA attack, which allows an attacker to determine the two large prime numbers that form the basis of the cryptographic security (which relies on the product of the large primes being extremely difficult to factor) and hence, to compromise the entire construction. The authors have created a new draft that proposes using a hash construction, but only the route originator can verify the route because of the non-invertibiity of secure hash functions. Thus, the recipient would have to query the originator as to the veracity of the route.

None of the solutions for BGP path validation that have been proposed have provided appropriate tradeoffs between security, resource usage, and deployability. Recent moves within the standards community signal that there is a movement to begin considering entirely new mechanisms and solutions. Until such efforts result in more appropriate solutions, ASes will continue to huerstically vet received path using filtering and other mechanisms.

5.3 Evaluating the Security Schemes

- incremental deployment
  - SBGP - costs (computation, storage)

Many protocols offer a degree of protection against attack, with S-BGP offering

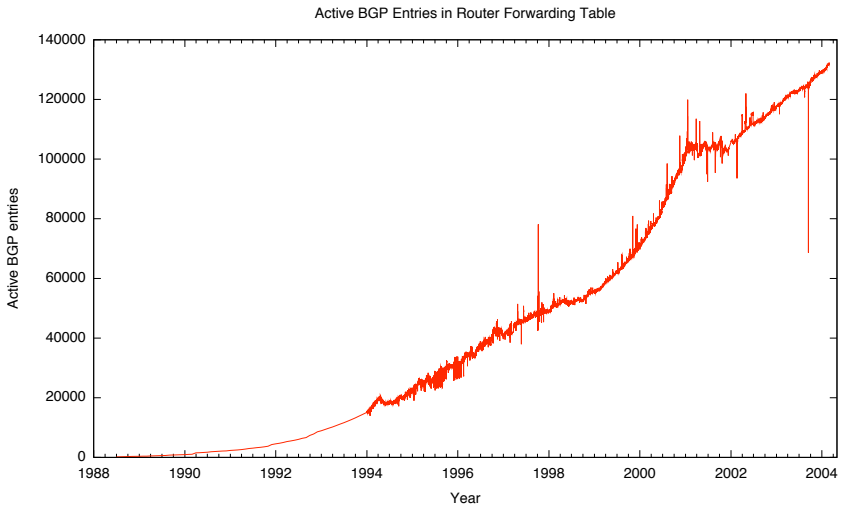

Fig. 6. 1988-2003 routing table updates from the CIDR report (http://www.cidr-report.org/)

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.

BGP security is complicated by operational considerations. 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 6 shows BGP updates from the CIDR report for 1988 to 2003. 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 consider these scalability issues [Huston 2001; Bellovin et al. 2001].

S-BGP has been touted as the leading candidate to provide comprehensive BGP security, but there are significant drawbacks that must be addressed before it can be effectively deployed. 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 of time-averaged statistics. The load in routers is not uniformly distributed, as Internet traffic is bursty in nature [Uhlig and Bonaventure 2001]. It has also been claimed that S-BGP will cause administrative delays [Meyer and Partan 2003]. Recent work by Nicol et al. [2002] proposes aggressive caching and pre-computing part of the DSA signature operation, with results that very closely approximate regular BGP performance.

The soBGP platform provides several deployment options and the ability to be incrementally deployed [White 2002]. 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 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 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, and under certain circumstances, 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 laid the groundwork for more comprehensive frameworks such as S-BGP and soBGP. By themselves, though, the five countermeasures do not provide the critical service of origin authentication, and they require changes to BGP without the commensurate benefits provided by the framework protocols. The IRV solution performs all security signaling out of band, allowing it to be deployed without modifying BGP at all.[8] However, IRV requires more analysis of infrastructure requirements and operational semantics to be a viable security alternative.

The cumulative authentication mechanisms for path integrity rely on a full mesh of shared secrets, making its use impractical in the Internet. Ingress filtering makes for good policy, but does not provide comprehensive support against BGP attacks. Similarly, MOAS conflict detection can be useful to prevent origin-based attacks, but is limited in the other defenses it can supply. Finally, the Listen and Whisper protocols are meant to be easily deployable without major infrastructure changes, but are limited in the amount of security they can provide, and weaknesses with Whisper's constructions severely limits its usefulness.

6. 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.

6.1 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 [He and Papadopoulos 2003]. Suggestions have been made for designing a routing architecture in large networks

[8]The ability to distribute Entitycerts and Policycerts out of band has recently been added to soBGP [White 2002].

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]. Another approach towards analysis of routing security is performed by Pei et al. [2003], 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 [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]. Finally, an overview of route filtering and S-BGP as countermeasures to BGP attacks is given in [Nordström and Dovrolis 2004].

6.2 Cryptographic Constructions

Future BGP security research can exploit new cryptographic constructions to efficiently and securely protect the routing fabric. For example, 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].

6.3 Attack Detection

Detecting attacks is an active field of research. The PAIR algorithm [Chakrabarti and Manimaran 2003] 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 [Avramopoulos et al. 2004]. 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 [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 traceroute [Padmanabhan and Simon 2002] and AS-level traceroute [Mao et al. 2003] to detect malicious routing may aid in this effort.

7. 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 Interdomain Routing Protocol (IDRP) [ISO 1992] 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 [Perlman 1999]. 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 evident even now: multi-protocol extensions are increasingly used to route IPv6 packets [Bates et al. 2000].

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. Solving the issue of BGP security is very difficult because of the scale and complexity of the Internet. Every network in the world communicating with other organizations through the Internet uses BGP, and errors in configuration and operation can have a global impact. 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 to securing BGP may be one of the best way to ensure that the Internet remains a reliable and useful vehicle for private and public communication.

REFERENCES

Aiello, W., Ioannidis, J., and McDaniel, P. 2003. Origin authentication in interdomain routing. ACM CCS'03, Washington, DC, USA.

Alaettinoglu, C. and Shankar, A. U. 1995. The viewserver hierarchy for interdomain routing: Protocols and evaluation. *IEEE Journal on Selected Areas in Communications 13,* 8 (Oct.), 1396–1410.

Avramopoulos, I., Kobayashi, H., Wang, R., and Krishnamurthy, A. 2004. Highly secure and efficient routing. IEEE INFOCOM 2004, Hong Kong, PRC.

B. Baker and R. Shostak. 1972. Gossips and Telephones. *Discrete Mathematics* 2, 191–193.

Baltatu, M., Lioy, A., Maino, F., and Mazzocchi, D. 2000. Security issues in control, management and routing protocols. *Computer Networks (Amsterdam, Netherlands: 1999) 34,* 6, 881–894. Elsevier Editions, Amsterdam.

Barbir, A., Murphy, S., and Yang, Y. 2003. Generic threats to routing protocols. Internet Draft.

Barrett, R., Haar, S., and Whitestone, R. 1997. Routing snafu causes Internet outage. *Interactive Week*.

Bates, T., Gerich, E., Joncheray, L., Jouanigot, J.-M., Karrenberg, D., Terpstra, M., and Yu, J. 1995. Representation of IP routing policies in a routing registry. RFC 1786.

Bates, T., Rekhter, Y., Chandra, R., and D, K. 2000. Multiprotocol extensions for BGP-4. RFC 2858.

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. 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.

C. Wong and S. Lam . 1999. Digital Signatures for Flows and Multicasts. *IEEE/ACM Transactions on Networking 7,* 4 (August), 502–513.

Callon, R. 1990. Use of OSI IS-IS for routing in TCP/IP and dual environments. RFC 1195.

Castineyra, I., Chiappa, N., and Steenstrup, M. 1996. The Nimrod routing architecture. RFC 1992.

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. RFC 1997.

Cheung, S. 1997. An efficient message authentication scheme for link state routing. In *Proceedings of the 13th Annual Computer Security Applications Conference*. IEEE Computer Society, 90–98.

CIDR. 2004. CIDR report for 21 February 04. http://www.cidr-report.org/.

Cranor, L. and LaMacchia, B. 1998. Spam! *Communications of the ACM 41,* 8 (Aug.), 74–83.

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.

Department of Homeland Security. 2003. The national strategy to secure cyberspace.

Dierks, T. and Allen, C. 1999. The TLS protocol version 1.0. RFC 2246.

Diffie, W. and Hellman, M. 1976. New Directions in Cryptography. *IEEE Transactions on Information Theory IT-22,* 6 (November), 644–654.

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, USA.

Feamster, N., Mao, Z. M., , and Rexford, J. 2004. BorderGuard: Detecting Cold Potatoes from Peers. *to appear*.

Gao, L. 2001. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking 9,* 6, 733–745.

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

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. ISOC NDSS'03, San Diego, CA, USA, 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, Osaka, Japan.

Green, B. 2002. BGP security update: Is the sky falling? NANOG 25.

Griffin, T. 2003. *personal communication*.

Griffin, T. and Wilfong, G. 1999. An analysis of BGP convergence properties. ACM SIGCOMM 1999, Cambridge, MA, USA.

Halabi, B. 2000. *Internet Routing Architectures*, Second ed. Cisco Press.

Harkins, D. and Carrel, D. 1998. The Internet Key Exchange. RFC 2409.

Hawkinson, J. and Bates, T. 1996. Guidelines for creation, selection, and registration of an autonomous system (AS). RFC 1930.

He, X. and Papadopoulos, C. 2003. A framework for incremental deployment strategies for router-assisted services. IEEE INFOCOM 2003, San Francisco, CA, USA.

Heffernan, A. 1998. Protection of BGP sessions via the TCP MD5 signature option. RFC 2385.

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, USA.

Hu, Y.-C., Perrig, A., and Sirbu, M. 2004. SPV: Secure Path Vector Routing for Securing BGP. In *ACM SIGCOMM*. ACM.

Huston, G. 2001. Commentary on inter-domain routing in the Internet. RFC 3221.

Huston, G. 2003. BGP AS number exhaustion. NANOG 28.

in the TCP/IP Protocol Suite, S. P. 1989. Steven m. bellovin. *Computer Communications Review 2,* 19, 32–48.

ISO. 1992. Intermediate System to Intermediate System Inter-Domain Routeing Information exchange protocol. DIS 10747.

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. 1998a. IP Authentication Header. RFC 2402.

Kent, S. and Atkinson, R. 1998b. IP Encapsulating Security Payload. RFC 2406.

Kent, S. and Atkinson, R. 1998c. Security architecture for the Internet Protocol. RFC 2401.

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 (S-BGP). *IEEE Journal on Selected Areas in Communications 18,* 4 (Apr.).

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

Kruegel, C., Mutz, D., Robertson, W., and Valeur, F. 2003. Topology-based detection of anomalous BGP messages. In *Proceedings of the 6th Symposium on Recent Advances in Intrusion Detection (RAID)*. 17–35.

Kumar, B. and Crowcroft, J. 1993. Integrating security in inter-domain routing protocols. *ACM SIGCOMM Computer Communication Review 23,* 5 (Oct.), 36–51.

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, Anchorage, AK, USA.

Lonvick, C. 2003. RADIUS attributes for soBGP support. Internet Draft.

M. Naor and K. Nissim. 1998. Certificate Revocation and Certificate Update. In *Proceedings of the 7th USENIX Security Symposium*. San Antonio TX USA, 217 – 228.

Mahajan, R., Wetherall, D., and Anderson, T. 2002. Understanding BGP misconfiguration. ACM SIGCOMM 2002, Pittsburgh, PA, USA.

Malkin, G. 1994. RIP version 2. RFC 1723.

Mao, Z., Rexford, J., Wang, J., and Katz, R. 2003. Towards an accurate AS-level traceroute tool. ACM SIGCOMM 2003, Karlsruhe, Germany.

Maughan, D., Schertler, M., Schneider, M., and Turner, J. 1998. Internet Security Association and Key Management Protocol (ISAKMP). RFC 2408.

Merkle, R. 1980. Protocols for public key cryptosystems. IEEE Symposium on Research in Security and Privacy, Oakland, CA.

Meyer, C. and Partan, A. 2003. BGP security, availability, and operator needs. NANOG 28.

Mills, D. 1984. External Gateway Protocol formal specification. RFC 904.

Minoli, D. and Schmidt, A. 1999. *Internet Architectures*. John Wiley & Sons, New York, NY.

Misel, S. 1997. "Wow, AS7007!". Merit NANOG Archive.
http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html.

Moy, J. 1998. OSPF version 2. RFC 2178.

Murphy, S. 2004. BGP security vulnerabilities analysis. Internet Draft.

Nicol, D., Smith, S., and Zhao, M. 2002. Efficient security for BGP route announcements. Dartmouth Computer Science Technical Report TR-2003-440.

Nordström, O. and Dovrolis, C. 2004. Beware of BGP attacks. *Computer Communications Review 34,* 2 (Apr.), 1–8.

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, USA.

Paxson, V. 1999. End-to-end Internet packet dynamics. *IEEE/ACM Transactions 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. Nov.

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.

Perrig, A., Canetti, R., Tygar, J. D., and Song, D. 2002. The TESLA broadcast authentication protocol. *RSA CryptoBytes 5,* Summer.

Postel, J. 1981. Internet Protocol. RFC 791.

Rekhter, Y. and Li, T. 1995. A Border Gateway Protocol 4 (BGP-4). RFC 1771.

Rivest, R. 1992. The MD5 Message-Digest Algorithm. RFC 1321.

Rivest, R., Shamir, A., and Adelman, L. M. 1978. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM 21,* 2 (Feb.), 120–126.

S. Even and O. Goldreich and S. Micali. 1996. On-Line/Off-Line Digital Signatures. *Journal of Cryptology 9,* 1, 35–67.

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, Anaheim, CA, USA.

Slater III, W. 2002. The Internet outage and attacks of October 2002. http://www.isoc-chicago.org/Internetoutage.pdf.

Smith, B. and Garcia-Luna-Aceves, J. 1996. Securing the Border Gateway Routing Protocol. Global Internet '96, London, UK.

Smith, B. and Garcia-Luna-Aceves, J. 1998. Efficient security mechanisms for the border gateway routing protocol. *Computer Communications 21,* 3, 203–210.

Spring, N., Mahajan, R., and Wetherall, D. 2002. Measuring ISP Topologies with Rocketfuel. ACM SIGCOMM 2002, 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, AK, USA.

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

Stewart, J., Bates, T., Chandra, R., and Chen, E. 1998. Using a Dedicated AS for Sites Homed to a Single Provider. *Internet Engineering Task Force*. RFC 2270.

Subramanian, L., Agarwal, S., Rexford, J., and Katz, R. 2002. Characterizing the Internet hierarchy from multiple vantage points. IEEE INFOCOM 2002, 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.

Thayer, R., Doraswamy, N., and Glenn, R. 1998. IP Security Document Roadmap. RFC 2411.

Traina, P. 1995. Experience with the BGP-4 protocol. RFC 1773.

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. 1999. Routing policy system security. RFC 2725.

Wan, T., Kranakis, E., and van Oorschot, P. C. 2005. Pretty Secure BGP (psBGP). In *Proc. Network and Distributed Systems Security 2005*. Internet Society (ISOC), San Diego, CA.

White, R. 2002. Deployment considerations for secure origin BGP (soBGP). Internet Draft.

Yu, J. 2000. Scalable routing design principles. RFC 2791.

Zhang, K. 1998. Efficient protocols for signing routing messages. ISOC NDSS'98, San Diego, CA, USA.

Zhang, Y., Duffield, N., Paxson, V., and Shenker, S. 2001. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW '2001)*. San Francisco, California, USA.

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, USA.

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, San Francisco, CA, USA.