

# Beware of BGP Attacks

Ola Nordström and Constantinos Dovrolis  
College of Computing  
Georgia Institute of Technology  
{nalo,dovrolis}@cc.gatech.edu

## ABSTRACT

This note attempts to raise awareness within the network research community about the security of the interdomain routing infrastructure. We identify several attack objectives and mechanisms, assuming that one or more BGP routers have been compromised. Then, we review the existing and proposed countermeasures, showing that they are either generally ineffective (route filtering), or probably too heavyweight to deploy (S-BGP). We also review several recent proposals, and conclude by arguing that a significant research effort is urgently needed in the area of routing security.

## 1. INTRODUCTION

As more and more businesses and organizations become dependent on the Internet, the risks posed by malicious attacks on the Internet infrastructure become more significant. The Internet has experienced several successful large-scale attacks that caused major losses to their victims. The attacks typically target major Web servers, content providers, the DNS system, or just end-hosts [1].

The Internet routing infrastructure is also vulnerable to attacks. Because of the very nature of this infrastructure, routing attacks can affect a large number of hosts, entire networks, or even the global Internet [2]. The objectives of routing attacks can include blackholing and loss of connectivity, traffic redirection to networks controlled by adversaries, traffic subversion and data interception, or persistent routing instability [3].

An intradomain routing system operates within an Autonomous System (AS). The threat of an attack on intradomain routing is thus typically contained within a single network. The interdomain routing infrastructure, on the other hand, is based on the BGP protocol and it provides connectivity between ASs [4, 5]. In this note, we focus on the vulnerability of interdomain routing and BGP, because such attacks have the potential to affect a much larger number of users

and potentially compromise routing across the global Internet. We assume that the reader has some basic familiarity with the BGP protocol and with how it is deployed in the Internet to provide policy-based routing.

So far, there have been no major BGP routing attacks (or at least, they have not been publicly documented as malicious attacks). As a result, relatively little attention from the network research community has been placed on studying the routing infrastructure's overall susceptibility to malicious users. On the other hand, it has been shown that routing misconfigurations are quite common in practice, and they can cause the same reachability and BGP convergence problems that an attack could cause [6]. The notorious AS7007 incident on April 25 1997 was caused by a misconfigured router that flooded the Internet with incorrect advertisements, announcing AS7007 as the origin of the best path to essentially the entire Internet. As a result that AS quickly became a major traffic sink, and it disrupted reachability to many networks for several hours [7]. Similar events occurred on April 7 1998, when AS8584 announced about 10,000 prefixes it did not own, and on April 6 2001 when AS15412 announced about 5,000 prefixes it did not own [8].

In this note, we explore how an attacker might exploit the BGP protocol to compromise the interdomain routing infrastructure. Our objective is to show that BGP is vulnerable to a number of malicious attacks, and to raise awareness within the network research community about this issue. The presented attacks are relatively easy to perform as long as a hacker manages to compromise one or more BGP speakers. Note that we focus on general attacks that are allowed by the BGP protocol, rather than on specific bugs and vulnerabilities of different BGP implementations. Then, we describe the major proposed countermeasures for BGP security, namely route filtering and S-BGP, together with some more recent research proposals. Unfortunately, neither filtering nor S-BGP can prevent all the attacks that we consider. S-BGP is much more effective than filtering, but it requires major changes in the interdomain routing infrastructure, preventing, at least so far, its deployment.

The rest of this note is organized as follows. Section 2 outlines the objectives an attacker may have when targeting interdomain routing. The BGP mechanisms that enable such attacks are described in Section 3. The effectiveness of the two major countermeasures (filtering and S-BGP) are described in Section 4. A review of some recently proposed

countermeasures, still in the research phase, is given in Section 5. We conclude in Section 6.

## 2. ATTACK OBJECTIVES

In the scenarios that we consider next, the major assumption is that *a hacker has managed to compromise and take complete control of one or more BGP routers in the Internet*. This can be accomplished with password sniffers, exploiting vulnerabilities in the router’s operating system, or simply stealing passwords from the network operator.

The objectives of an attacker can include prefix blackholing, traffic redirection, traffic subversion, or creation of routing instability.

**Blackholing** occurs when a prefix is unreachable from a large portion of the Internet. Intentional blackhole routing is used to enforce private and non-allocated IP ranges. Malicious blackholing refers to false route advertisements that aim to attract traffic to a particular router and then drop it.

**Redirection** occurs when traffic flowing to a particular network is forced to take a different path and to reach an incorrect, potentially also compromised, destination. One objective of redirection attacks is that the compromised destination impersonates the true destination to receive confidential information. Another objective may be to redirect excessive amounts of traffic to a certain link or network and cause congestion collapse.

**Subversion** is a special case of redirection in which the attacker forces the traffic to pass through a certain link with the objective of eavesdropping or modifying the data. In subversion attacks the traffic is still forwarded to the correct destination, making the attack more difficult to detect.

**Instability** in interdomain routing can be caused by successive advertisements (potentially with different attributes) and withdrawals for the same network. An objective of such attacks can be to trigger route dampening in upstream routers, and thus cause connectivity outages. Another objective can be to create large increases in the volume of BGP traffic, and consequently long convergence delays [9].

## 3. ATTACK MECHANISMS

A compromised router can modify, drop, or introduce fake BGP updates. The result can be that other routers have incorrect views of the network, leading to blackholing, redirection, or instability. As explained next, the effectiveness of some attacks depends on the AS topology and on the location of the compromised router relative to the victim network.

Figure 1 shows a sample AS topology. AS1 and AS2 are stub networks that have been assigned address blocks from their provider AS3. All ASs provide transit service to their customers, which reside at the lower levels of the diagram. The horizontal lines (e.g. between routers *B-V*) represent backup links and non-transit relations between the corresponding ASs.




Figure 1: An AS topology.

**False UPDATES and prefix hijacking** are probably the most straightforward type of BGP attack. They occur when an AS announces a route that it does not have, or when an AS originates a prefix that it does not own.

Suppose that router *B* wants to subvert traffic destined to AS2. It could announce a fake route, announcing that it has a direct connection to AS2. Traffic destined to AS2 originating at AS2 and AS3, and potentially at AS1 and AS5, would still reach its destination without passing through *B* due to the shorter distance of those ASs to AS2. However, traffic from other parts of the Internet would pass through *B* due to the shorter AS-PATH presented by that compromised router.

Router *B* could also claim ownership of the address blocks originated by AS2. Routers *A* and *R* would then forward traffic destined to AS2 to *B*. In a subversion attack, *B* could then forward the traffic towards its correct destination in AS2.

In [8], Zhao et al. studied the occurrence of BGP Multiple Origin AS (MOAS) conflicts. MOAS conflicts occur when multiple ASs announce themselves as the “origin” (i.e., owner) of a particular prefix. They found that MOAS conflicts are increasingly common in the Internet. They listed several potential configuration errors that may lead to MOAS conflicts, and showed several instances where ASes have claimed false ownership of a large number of prefixes leaving the affected prefixes with partial connectivity. The study of MOAS conflicts confirms that illegitimate advertisements do occur in the Internet and, when coordinated by a malicious attack, they can cause major connectivity problems.

The effectiveness of false UPDATES is limited by the location and connectivity of the hijacked BGP speaker. Stub networks may originate MOAS conflicts or attacks, but due to their location it is unlikely that their routes will be preferred by many other networks. For example, AS1 may announce a prefix owned by AS5. However, when router *B*

receives the fake UPDATE it will not use it because the fake AS-PATH is longer than the AS-PATH of the legitimate route to AS5.

**De-Aggregation**, when used as an attack, breaks up an address block into a number of more specific (i.e., longer) prefixes. Since the BGP route selection process gives higher preference to the longest matching prefix for a given destination, the attacker can use de-aggregation to announce fake routes that will be preferred throughout the Internet over the legitimate routes to that network.

For instance, the compromised router  $B$  can de-aggregate the prefix announced by AS2 to two prefixes that are longer by one bit, while keeping the AS-PATH to AS2 the same. In that case, traffic originating anywhere in the Internet, except in AS2, and destined to AS2 would be forwarded towards router  $B$ . If AS2 owned a prefix that was aggregated with other prefixes by the provider AS3, then  $B$  could simply announce the original AS2 prefix.

Note that a compromised BGP speaker can use de-aggregation to blackhole a victim network anywhere in the Internet, regardless of the proximity between the two. In [6], Mahajan et al. found that origin misconfigurations are largely due to inadvertent de-aggregation.

**Contradictory advertisements**, meaning different routing announcements sent by the same AS to different BGP peers, is a legitimate technique for interdomain traffic engineering. We show next that it can also be used as an attack mechanism.

BGP offers a number of attributes that can be used in the route selection process to choose the most preferred path to a certain destination. For instance, a multihomed AS can send UPDATEs with a *padded* AS-PATH to one of its providers so that the link to that provider is only used if the primary link to another provider fails. Suppose that AS3 uses link  $B-M$  as is its primary connection to the global Internet, and link  $V-N$  as backup. To accomplish this policy, the AS3 border router  $N$  can “pad”, or extend, the AS-PATH of the UPDATEs going to AS5 with several repetitions of its own AS number. The AS-PATH for AS1 and AS2 to AS4 will then be  $\{AS1,AS3\}$  and  $\{AS2,AS3\}$  respectively. On the other hand, the AS-PATH of the UPDATEs sent to AS5 can be artificially padded as in  $\{AS2,AS3,AS3,AS3\}$  and  $\{AS1,AS3,AS3,AS3\}$ . This will make the path through AS5 longer and less attractive for other ASs.

Contradictory advertisements can be used by a malicious router to redirect traffic to itself or to another AS. To illustrate, the compromised router  $B$  should normally only announce the AS1 route that goes through  $\{AS1,AS3,AS4\}$ . Instead,  $B$  can propagate that route only to  $A$  indicating that it should not be announced any further, and announce the padded route that goes through AS5 to  $R$ . Effectively, this means that part of the Internet (excluding AS4) will be able to reach AS1 only through AS5. The attacker may want to do so in order to create congestion in AS5, or to redirect traffic destined to AS1 through a suboptimal, backup path.

**Update modifications** can be used by a compromised

router to redirect traffic in a way that hurts the origin AS. Suppose that AS3 uses the link  $V-N$  only for backup purposes because it is cheaper to use link  $B-M$  instead. AS4 does not advertise its AS3 route to AS5, because doing so would enable AS5 to use AS4 to reach AS3 instead of using its own link  $V-N$ . To prevent other ASs from using the link  $V-N$ , router  $N$  can pad the UPDATEs going to  $V$ , making the corresponding AS-PATH longer.

Assume now that router  $R$  is compromised, and that it wants to redirect traffic to AS3 through the more expensive link  $V-N$ .  $R$  can drop the padding in the route that includes the  $\{AS5,AS3\}$  link, and instead pad the route that includes the  $\{AS4,AS3\}$  link (or simply not announce it). This would force traffic for AS3 to take the more costly  $V-N$  route.

As long as connectivity is preserved, update modifications can be very difficult to detect. Business relationships and policies between providers are largely kept secret, and as a result the ability to detect illegitimate routing in terms of policy constraints can be difficult for a third party.

**Advertent link flapping** can be used to trigger route dampening for a victim network at an upstream router. A malicious router can advertently flap a route to a victim address block(s). This can be done by withdrawing and re-announcing the target routes at a sufficiently high rate that the neighboring BGP speakers dampen those routes. A dampened route would force the traffic to the victim AS to take a different path, enabling traffic redirection. Route dampening occurs even if the router cannot find an alternate path to the corresponding destination. The victim network, in that case, remains unreachable for the duration of the route dampening.

Let  $B$  be the malicious router, and suppose that the attacker’s goal is to trigger dampening at  $R$  for the routes to AS1.  $B$  can do so by sending to  $R$  a sequence of withdrawals for the route  $\{AS1,AS3,AS4\}$ , followed by announcements for the route  $\{AS1,AS3,AS5,AS4\}$ , followed by new announcements for the route  $\{AS1,AS3,AS4\}$ .

In [10], Mao et al. found that even a single BGP withdrawal followed by a re-announcement for a certain network can activate dampening, making that network unreachable for up to an hour. The dampening can be triggered when a single route flap forces BGP peers to consider several backup paths, causing a large number of additional withdrawals and announcements.

**Instability**, in the form of wide-scale cascading failures, can occur when a number of BGP sessions repeatedly time-out due to router reboots, link congestion, or physical link intermittent failures [11]. Instability, in the form of delayed convergence (up to several minutes), can also occur upon routing or policy changes, due to the MinRouteAdver timer and the way BGP explores alternate paths [9].

A hijacked router or, more likely, a number of hijacked routers may be able to cause the same kind of BGP instability by advertently flapping a large number of their routes. The flapping should be of appropriate frequency so that it does not trigger dampening at the upstream routers, or oth-

erwise the instability will be of limited scope. Such attacks can cause intermittent reachability and blackholing to the victim prefixes, but even worse, they may also be able to cause degraded routing performance in the global Internet.

Another type of routing attack, based on link cuts, is presented in [12]. The basic idea is that a hacker that knows the topology of a network can determine which links to disable in order to launch a blackholing, redirection, or subversion attack. We consider such attacks outside the scope of this note.

**Congestion-induced BGP session failures.** An indirect way to attack the interdomain routing infrastructure is by causing heavy congestion in links that carry BGP peering sessions. During heavy congestion, the TCP-based BGP sessions can be so slow that they are eventually aborted, causing thousands of routes to be withdrawn. When BGP sessions are brought up again, routers must exchange full routing tables, creating large spikes of BGP traffic and significant routing convergence delays. For instance, it is possible that the Code Red and Nimda worms of 2001 affected interdomain routing in that way. The two worms compromised thousands of hosts within a few hours, and they caused persistent congestion in several network links. A report published by Renesys showed that during the adverse effects of the worms, BGP traffic “exploded” by a factor of 25 [13]. It is interesting to note that the worms did not target BGP; the aborted BGP sessions were simply a side effect that probably even the attackers had not predicted. More recently, the results of [13] were questioned by [14], arguing that over 40% of the observed BGP updates during the worm attacks were an artifact of the measurement infrastructure used by the Renesys study. Nevertheless, the results of both [14] and [13] agree that BGP implementations are quite vulnerable to congestion and stressful network conditions.

## 4. MAJOR COUNTERMEASURES

In the current Internet, the possibility of BGP attacks and misconfigurations has been so far mostly dealt with “Best Common Practice” (BCP) documents from router vendors. BCPs typically recommend practical measures to prevent a router from being hijacked, and to avoid fake or incorrect advertisements from being accepted by a router.

For instance, the “BGP TTL Security Hack” (BTSH) protects against hackers that attempt to hijack a BGP session without controlling either of the two speakers [15]. The basic idea is to set the IP header TTL field to a value that allows those BGP packets to reach the receiving router only if the latter is exactly one hop away from the sender. Obviously, BTSH is not effective in multi-hop BGP sessions; currently, however, most external-BGP sessions are between adjacent routers.

To protect against spoofed messages and TCP connection hijacking, BGP sessions are often protected using the TCP MD5 signature option [16]. Another security feature is to perform Unicast Reverse Path Filtering (Unicast-RPF), examining whether the received BGP messages have the source address of the peering BGP speaker. Also, to protect against “out-of-memory” attacks some vendors provide a *MaxPrefixLimit* feature that terminates a peering session if a peer

advertises too many prefixes.

Even though the previous countermeasures can prevent hackers from hijacking a BGP session, they are unable to deal with attacks from a compromised BGP router. Next, we describe the two major solutions for these kinds of attacks, namely routing filtering and S-BGP. We note that routing filtering is already used in different capacities around the Internet mostly to enforce various routing policies. S-BGP was developed by BBN during 1997-2000, but it has not been widely deployed yet [17, 18].

### 4.1 Route filtering

Currently, the main use of route filtering is to enforce business relationships between ASs. Filtering works by creating Access Control Lists of prefixes or ASs which are then used by a router when it sends or receives UPDATES. Outgoing UPDATES pass through *egress* filters allowing operators to control which routes are announced to peers. In Figure 1, AS4 may not want to provide transit service to AS5. To enforce this policy, router *B* can install an egress filter that only allows the routes owned by AS4, or by its customers, to reach *V*.

*Ingress* filters, on the other hand, are applied to incoming UPDATES and they can be used to check the validity of the received routes. Specifically, some ISPs use routing filters to verify that the origin AS of a route truly owns the corresponding prefix, and to prevent prefix hijacking. To construct such filters, ISPs are supposed to know the owner of each address block from Internet Routing Registries (IRRs). From their side, IRRs are supposed to be updated and consistent with each other. When this is the case, this type of filtering can be used to verify origin AS announcements and to prevent MOAS conflicts.

Unfortunately, the IRR databases are often not well-maintained and updated, and ISPs do not query them frequently enough. Labovitz et al. [19] reported that “due to the technical and contractual difficulties of maintaining filter lists for the large number of routes advertised by peers, most providers resort to trusting their peers to send only valid information”. In Europe, even though ISPs use extensive filtering, the fact that the corresponding IRR databases are not always up-to-date often causes correct routes to be rejected [6].

A filter-based verification of routing updates is against, in some sense, the dynamic nature of the Internet. Complete routing filters require a global knowledge of the AS topology and of the business relations between ASs. Changes in either the topology or the routing policies can be inconsistent with installed filters, causing either correct routes to be dropped or fake routes to be allowed.

Note that filtering would be much simpler in a completely hierarchical Internet, where the AS-level topology is a tree, and the only relation between two ASs is that of a customer-provider pair. In that case, the ingress filter of an AS would just check that the routes received by downstream ASs match the prefixes allocated to those customers and that the origin AS is correct. In reality, however, the AS topology is far from being a tree. Instead, the AS graph

consists of many multihomed nodes at the edges and dense transit, peer, or backup relations at the core. To make things worse, the policies between ASs are typically kept confidential even from IRRs.




Figure 2: A tree AS topology (A) vs a mesh AS topology with several backup links (B).

Figure 2-a represents a tree topology where both egress and ingress filtering can be easily constructed. *A* can use static ingress filters accepting routes from *B* only if their origin is any of the customers *B*, *D*, or *F*. Figure 2-b presents a quite different, and more realistic scenario, with several backup/peering links between ASs. *A* cannot filter announcements received from *B* as in the tree topology, because the latter is now also connected through backup links to *C*, *G* and *E*. In practice, it would be impossible to predict the routes that remote networks can generate as a result of arbitrary link failures or policy changes.

## 4.2 S-BGP

Secure BGP (S-BGP) was designed by researchers at BBN as an extension to BGP with the objective to protect BGP from erroneous or malicious UPDATES [17, 18]. S-BGP adds strong authorization and authentication capabilities to BGP based on public-key cryptography. S-BGP makes three major additions to BGP. First, it introduces a Public Key Infrastructure (PKI) in the interdomain routing infrastructure to authorize prefix ownership and validate routes. Second, a new transitive attribute is introduced to BGP updates. That attribute ensures the authorization of routing UPDATES, and prevents route modifications from intermediate S-BGP speakers. Third, all routing messages can be secured using IPSec, if routing confidentiality is a requirement.

Two key features of S-BGP are *Address Attestations* and *Route Attestations*. An Address Attestation (AA) is generated by the owner of a prefix, and it is used by S-BGP routers to verify that the origin AS is indeed authorized to advertise that address block. Route Attestations (RAs), on the other hand, are added by S-BGP routers in UPDATES, authorizing a neighboring AS to propagate the route contained in that UPDATE. S-BGP uses a PKI infrastructure to authorize AAs and RAs. The private keys are stored in S-BGP speakers, while the public keys are made available by a hierarchical PKI infrastructure.




Figure 3: S-BGP attestations at work.

RAs are daisy-chained as an UPDATE flows through a sequence of S-BGP routers. Each S-BGP router along the path is required to validate the integrity of an UPDATE before signing it and re-advertising it to its neighbors. Figure 3 shows an UPDATE flowing from the origin AS speaker *A* to router *B* and so on. The following 6 steps illustrate the protocol's operation.

1. *A* generates an RA for the prefix *P* indicating *B* as the next-hop for that route.
2. *A* sends the UPDATE, including the RA, to *B*.
3. *B* validates the signature in the RA using the public key of *A*.
4. *B* also verifies the AA for *P* (fetched offline) checking that *A* is the true owner of that prefix.
5. *B* verifies that *B* is the next-hop in the RA.
6. *B* generates two new RAs for its peers *C* and *D*, includes each RA in a different UPDATE, and forwards the two UPDATES to *C* and *D*.

If widely deployed, S-BGP would be a radical step towards securing the routing infrastructure. A deployment obstacle, however, is that it requires the presence of a hierarchical PKI infrastructure and distribution system, trusted by all participating ISPs. Another obstacle is that S-BGP is quite cryptographically intensive, requiring each UPDATE to be verified and signed by each S-BGP router (or by each participating AS) it goes through. This performance overhead can be unacceptable upon initialization (or reboot) of a BGP peering session due to the large number of routes that would be then received in a short time interval. Another implementation issue is that routers may need a large memory space (about 20 MB per peer) to store the public keys needed for route attestations. The space requirement can be significant for a speaker with tens of peers.

Aggregation is an additional problem for S-BGP. Route aggregation provides a means to coalesce several prefixes into a larger address block, thus reducing the number of UPDATES generated by a BGP speaker. S-BGP however, requires that all UPDATES be signed by the prefix owner. An upstream router performing aggregation would not be the owner of all the constituent prefixes. Also, S-BGP cannot prevent "collusion attacks". Such attacks are possible when two compromised routers fake the presence of a direct link between them. For the rest of the Internet, it then appears as if those two ASs are connected.

Unfortunately S-BGP has not been deployed so far [20]. It is unclear at this point whether this is due to the technical complexities of the protocol, due to the large overhead involved in the transition from BGP to S-BGP or in the establishment of a PKI infrastructure, or simply because the risk of routing attacks is not considered as significant by ISPs.

### 4.3 Effectiveness of route filtering and S-BGP

The following table shows the effectiveness of filtering and S-BGP against each of the routing attacks we considered in Section 3.

Attack	Filtering	S-BGP
False Updates	Partial	Secure
De-Aggregation	Possible	Secure
Contradictory Advertisements	Possible	Possible
Update Modifications	Possible	Secure
Advertent Link Flapping	Possible	Possible
Instability	Possible	Possible

To prevent false advertisements, both filtering and S-BGP enable a router to reject unauthorized incoming routes. The route attestation feature of S-BGP allows a router to reject any invalid routes or fake prefix-origin claims by a malicious router. Filtering, on the other hand, is mostly used today to reject fake prefix-origin claims close to the edge of the AS hierarchy, rather than in the core. Filtering could also be effective in detecting invalid routes if the IRR routing databases contained updated topology and policy information. To the extent of our knowledge, this is not widely the case however.

De-aggregation is not possible with S-BGP because of the authentication that S-BGP provides. Filtering typically verifies only the origin AS for an announced prefix, and so it is possible for an attacking router to de-aggregate the prefix of another AS.

Contradictory advertisements can be performed by sending certain UPDATES to some peers while sending other UPDATES to others. When a prefix is announced, the origin AS has no control over how far that announcement will be propagated because it does not know or control the policies of its peers. So, such attacks are possible with both filtering and S-BGP.

Update modifications can be detected using S-BGP due to the authentication provided by that protocol. On the other hand, update modifications are generally possible with filtering.

Malicious link flapping, or routing instability, cannot be avoided by filtering or S-BGP. An AS has no knowledge or control over how the routes that it announces propagate through the Internet.

## 5. RECENT RESEARCH PROPOSALS

In this section we review some recent proposals, still at the research stage, for providing interdomain routing security. We first note that the IETF created recently the Routing

Protocol Security Requirements (rpsec) working group [21]. The main objectives of the working group are to document general threat models for routing systems, provide a list of security requirements for routing systems, and analyze the vulnerabilities of specific protocols, such as BGP and OSPF. There are also a few (now expired) Internet Drafts that discuss interdomain routing security risks, such as the excellent BGP vulnerability analysis by Murphy [22].

**Secure Origin BGP:** soBGP is a lightweight alternative to S-BGP, mostly proposed by researchers at Cisco Systems [23]. soBGP aims to authenticate two aspects of routing information. First, soBGP validates that an AS is authorized to originate a given prefix. Second, soBGP attempts to verify that an AS advertising a prefix has at least one valid (in terms of policy and topology) path to that destination. soBGP is based on the use of three certificate types. The *Entity Certificate* is used to establish the identity and public key of an AS. The *Authorization Certificate* authenticates the assignment and delegation of IP address blocks, and it is used to verify prefix ownership. The *Policy Certificate* authenticates per-AS or pre-prefix policies and AS connectivity information, and it is used to verify the validity of a route. Instead of relying on a hierarchical PKI infrastructure, soBGP uses a Web-of-Trust model to validate certificates, relying on the existing relations between ISPs.

**MOAS-based conflict detection:** In [24], the authors proposed a protocol enhancement to BGP with the objective of detecting false route announcements and prefix hijacking. Rather than using route authentication techniques, the proposal of [24] is based on the observation that the AS-level Internet topology is typically densely connected. The premise is that it would be difficult for an attacker to completely block correct routing information. An alarm for a potentially false route advertisement can be triggered whenever a router detects a MOAS conflict. It is important to note however that some MOAS conflicts are valid, for instance, those resulting from multihoming. The authors of [24] propose the use of a BGP community attribute as a simple way to attach a list of the valid originating ASs to a route.

**Interdomain Routing Validation IRV:** Another form of path and origin verification, called Inter-domain Routing Validation (IRV), has been proposed by Goodell et al. in [25]. IRV separates the authentication component from the BGP protocol, moving the former to a separate companion protocol called IRV. With IRV, each AS deploys one or more IRV servers. When a BGP router receives an UPDATE, the corresponding IRV server contacts the IRV server of each ASs in the AS-PATH to verify both the origin and the routing path of the received UPDATE.

**Other proposals:** Securing distance vector protocols from malicious routers was the subject of [26] and [27]. BGP is a path vector protocol, however, and it allows arbitrary policy-based routing filters. A BGP security enhancement was proposed in [28]. The main innovation in that work was to augment UPDATE messages with a “predecessor” attribute, identifying the AS prior to the destination AS. Using this information, the authenticity and integrity of the entire path can be established. In an attempt to reduce the

computation overhead of protocols (such as S-BGP) that are based on public-key cryptography, the protocols of [29] are based on symmetric cryptography. A “secure traceroute” protocol was proposed in [30] with the objective to identify a router that causes routing problems and poor performance, rather than secure the routing protocol itself. More recently, a “Listen and Whisper” protocol has been proposed as an alternative to BGP enhancements that require a PKI infrastructure [31]. The “Listen” component of the protocol probes the data plane to detect whether advertised routes to different destinations actually work. The “Whisper” component uses cryptographic functions and routing redundancy to detect false routing advertisements.

## 6. DISCUSSION

The objective of this note is to raise awareness within the broader network research community about the security of the interdomain routing infrastructure. We identified several attack objectives and mechanisms, assuming that one or more BGP routers have been compromised. Then, we reviewed the existing and proposed countermeasures, arguing that they are either generally ineffective (filtering), or probably too heavyweight (S-BGP). The full extent of BGP’s susceptibility to attacks is difficult to evaluate, as major attacks of this type have not happened yet.

The nature of BGP gives ASs considerable latitude in determining which routes to modify, forward, or reject. Furthermore, the best-effort service model that the Internet is based on gives routers significant flexibility in choosing routing paths and forwarding behavior. This implies that there is a class of routing attacks that cannot be avoided simply because they do not necessarily constitute malicious behavior. For instance, a BGP speaker can install a *null* route for a set of prefixes, and drop all the traffic destined to those networks. A BGP router can also choose to deny forwarding routing updates containing certain prefixes or ASNs, or to simply withdraw those routes. This routing flexibility raises a fundamental question: which are the security requirements for an interdomain routing protocol? Is it reasonable to only require that packets reach their correct destination “most of the time”, or should we impose stricter requirements regarding the path(s) that traffic goes through?

Instead of securing BGP, another option would be to construct an entirely different interdomain routing protocol. Such a proposal is made in [32]. Obviously, replacing BGP at this point would be an enormously expensive and difficult task, given the wide deployment of BGP and the investment in operator expertise. Additionally, even if it was possible to replace BGP, it is unclear how to design a better interdomain routing system that can do everything BGP does well, and at the same time fix anything that BGP does not do well (including security).

We close this note noting that there may exist a *fundamental tradeoff between the resilience and security of a routing protocol* and that this tradeoff needs to be explored. A resilient routing protocol adapts quickly to changes and it is always able to restore connectivity as long as there is an alternate path. This is accomplished, however, by trusting the reachability information that other nodes provide, reducing the protocol’s security against Byzantine failures.

## 7. ACKNOWLEDGMENTS

The authors are grateful to the anonymous referees and to the CCR Editor, John Wroclawski, for their many constructive comments.

## 8. REFERENCES

- [1] A. Chakrabarti and G. Manimaran, “Internet Infrastructure Security: A Taxonomy,” *IEEE Network*, vol. 16, no. 6, pp. 13–21, Nov. 2002.
- [2] P. Papadimitratos and Z. J. Haas, “Securing the Internet Routing Infrastructure,” *IEEE Communications Magazine*, Oct. 2002.
- [3] R. J. Perlman, “Network Layer Protocols with Byzantine Robustness,” Ph.D. dissertation, Department of Electrical Engineering and Computer Science, MIT, 1988.
- [4] Y. Rekhter and T. Li, *A Border Gateway Protocol 4 (BGP-4)*, Mar. 1995, RFC 1771.
- [5] J. W. Stewart, *BGPv4: Inter-Domain Routing in the Internet*. Addison-Wesley, 1999.
- [6] R. Mahajan, D. Wetherall, and T. Anderson, “Understanding BGP Misconfiguration,” in *Proceedings of ACM Sigcomm*, Aug. 2002, pp. 3–16.
- [7] V. J. Bono, “7007 Explanation and Apology,” <http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html>, Apr. 1997.
- [8] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, “An Analysis of BGP Multiple Origin AS (MOAS) Conflicts,” in *Proceedings of ACM Internet Measurement Workshop*, Nov. 2001.
- [9] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, “Delayed Internet Routing Convergence,” in *Proceedings of ACM SIGCOMM*, 2000.
- [10] Z. M. Mao, R. Govindan, G. Varghese, and R. H. Katz, “Route Flap Damping Exacerbates Internet Routing Convergence,” in *Proceedings of ACM SIGCOMM*, Oct. 2002.
- [11] E. G. Coffman, Z. Ge, V. Misra, and D. Towsley, “Network Resilience: Exploring Cascading Failures within BGP,” in *Proceedings of Allerton Conference on Communications, Computing, and Control*, 2001.
- [12] S. M. Bellovin and E. R. Gansner, “Using Link Cuts to Attack Internet Routing,” ATT Research, Tech. Rep., 2004.
- [13] J. Cowie, A. Ogielski, B. J. Premore, and Y. Yuan, “Internet Worms and Global Routing Instabilities,” in *Proceedings of SPIE conference, Vol. 4868*, July 2002.
- [14] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, “Observation and Analysis of BGP Behavior under Stress,” in *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, Nov. 2002.

- [15] V. Gill, J. Heasley, and D. Meyer, "The BGP TTL Security Hack (BTSH)," Presentation at NANOG-27 meeting, Oct. 2002.
- [16] A. Heffernan, *Protection of BGP Sessions via the TCP MD5 Signature Option*, Aug. 1998, IETF RFC 2385.
- [17] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 4, pp. 582–592, Apr. 2000.
- [18] S. Kent, C. Lynn, J. Mikkelsen, and K. Seo, "Secure Border Gateway Protocol (S-BGP) - Real World Performance and Deployment Issues," in *Proceedings of Symposium on Network and Distributed Systems Security*, Feb. 2000.
- [19] C. Labovitz, R. Wattenhofer, S. Venkatachary, and A. Ahuja, "The Impact of Internet Policy and Topology on Delayed Routing Convergence," in *Proceedings of IEEE INFOCOM*, 2001.
- [20] D. Meyer and A. Partan, "BGP Security, Availability, and Operator Needs," Presentation at NANOG-28 meeting, June 2003.
- [21] T. Tauber and R. White, *Routing Protocol Security Requirements*, Mar. 2004, RPSEC IETF working group.
- [22] S. Murphy, *BGP Security Vulnerabilities Analysis*, June 2003, IETF Internet Draft (draft-ietf-idr-bgp-vuln-00.txt).
- [23] R. White, "Securing BGP: soBGP," <ftp-eng.cisco.com/sobgp/index.html>, Sept. 2003.
- [24] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, "Detection of Invalid Routing Announcement in the Internet," in *Proceedings of International Conference on Dependable Systems and Networks (DSN)*, June 2002.
- [25] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin, "Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing," in *Proceedings of Symposium on Network and Distributed Systems Security*, Feb. 2003.
- [26] A. Chakrabarti and G. Manimaran, "An Efficient Algorithm for Malicious Update Detection and Recovery in Distance Vector Protocols," in *Proceedings of IEEE International Conference on Communications (ICC)*, May 1997.
- [27] B. R. Smith, S. Murthy, and J. Garcia-Luna-Aceves, "Securing Distance-Vector Routing Protocols," in *Proceedings of Symposium on Network and Distributed Systems Security*, Feb. 1997.
- [28] B. R. Smith and J. Garcia-Luna-Aceves, "Efficient Security Mechanisms for the Border Gateway Routing Protocol," *Computer Communications Journal*, vol. 21, no. 3, pp. 203–210, 1998.
- [29] Y.-C. Hu, A. Perrig, and D. B. Johnson, "Efficient Security Mechanisms for Routing Protocols," in *Proceedings of Symposium on Network and Distributed Systems Security*, Feb. 2003.
- [30] V. Padmanabhan and D. R. Simon, "Secure Traceroute to Detect Faulty or Malicious Routing," in *Proceedings of Hot Topics in Networks (HotNets) Workshop*, Oct. 2002.
- [31] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz, "Listen and Whisper: Security Mechanisms for BGP," in *Proceedings of Networked Systems Design and Implementation (NSDI)*, Mar. 2004.
- [32] D. Zhu, M. Gritter, and D. R. Cheriton, "Feedback Based Routing," in *Proceedings of Hot Topics in Networks (HotNets) Workshop*, Oct. 2002.