A Selective Introduction to Border Gateway Protocol (BGP) Security Issues

Tao Wan P.C. van Oorschot Evangelos Kranakis

{twan, paulv, kranakis}@scs.carleton.ca
School of Computer Science, Carleton University, Ottawa, Canada.

August 1, 2005

Abstract

The Internet has become a critical communication infrastructure which we are increasingly reliant upon. As the world moves into a converged network where voice, video, and data are all transmitted over the same network, disruption of the Internet can cause more severe damage. Therefore, it is critical to protect the Internet from potential service disruption in order to ensure its continous functioning.

The Border Gateway Protocol (BGP) is the standard and only inter-domain routing protocol used on the Internet. BGP discovers and maintains routing information used for transmitting traffic across the Internet, thus, it is widely considered as a crucial component of the Internet infrastructure. Attacks on BGP can result in large scale service disruption. In this report, we study BGP security. Specifically, we study 1) the BGP protocol and its real world operations; 2) BGP security vulnerabilities and threats; and 3) BGP security mechanisms, including S-BGP from BBN, soBGP from Cisco, and psBGP from Carleton University. This report aims to provide sufficient background information for understanding BGP security issues, and to better understand the differences between existing BGP security proposals and the challenges faced in the design and practical deployment of a more secure BGP. We also provide comments regarding the role the government may play in helping to address security issues in BGP.

Contents

List of Figures

List of Tables

1 Introduction

The Internet is becoming increasingly important to our daily lives. As new exciting Internet technology and services are being developed, more and more traditional communication services are also being moved onto the Internet. As a result, we are becoming increasingly reliant on the Internet, and decreasingly tolerant of network connectivity outages. It is important to protect the Internet in order to ensure its continuous healthy operation.

However, it is well-known that the Internet is not secure, thanks to the wide spread of worms, viruses, trojans. While many people start to realize security problems caused by upper layer protocols (e.g., TCP) and software vulnerabilities (e.g., buffer overflow), less people are aware of potential damages which can be caused by exploiting security vulnerabilities of underlying Internet routing protocols.

The Internet routing infrastructure consists of a large number of intermediate systems (i.e., routers), each of which runs routing protocols for automatically discovering and maintaining routing tables. Routing tables are used for making decisions on how traffic should be forwarded over which paths to reach their ultimate destinations. If a routing table contains misinformation, wrong routing decisions will be made and traffic flow will be affected. Examples of consequences include denial of service and man-in-the-middle attacks.

In this report, we study security issues related to the Border Gateway Protocol [33], which is an IETF standard and the only inter-domain routing protocol for exchanging routing information between Autonomous Systems (ASes) on the Internet. Attacks on BGP can result in large scale service disruption, and can also be used to facilitate more sophisticated attacks against other protocols. Therefore, BGP is widely considered by security experts as one of the most important systems on the Internet which should be secured.

Unlike many other protocols whose security problems can be fixed by changing the protocols themselves, some security problems related to BGP result from deployment practices other than the BGP protocol specification itself. Thus, fixing BGP protocol vulnerabilities does not solve all BGP security problems. In addition, BGP is based on a distance vector approach in that each router computes its own routing table based on the routing tables it receives from its direct neighbors. While this approach allows propagation of good reachability information, it also facilitates propagation of misinformation. For example, one misbehaving router can poison the routing tables of many others even though they may behave correctly.

This report focuses on operational aspects of BGP which might have impact on BGP security, including IP address space allocation, AS business relationships, AS route exporting policies, and BGP route selection algorithms. We examine in detail a number of important threats against BGP which may soon be, or are already, happening on the Internet. We use examples to show step by step how a single misbehaving node can poison the routing tables of many other nodes on the network. We also show how *prefix hijacking* can be used to facilitate *advanced spamming*, *interception of password resetting messages*, and *phishing*.

We then outline a number of BGP security goals for countering identified threats. Three proposals for securing BGP (S-BGP [21, 22], soBGP [40], and psBGP [37]) are then discussed and compared against specified BGP security objectives. We suggest that psBGP has practical advantages over S-BGP and soBGP regarding IP prefix ownership verification, because it offers a distributed IP prefix registration model, i.e., each AS chooses a selected subset of its direct neighbors to endorse its prefix assertions. In other words, each AS registers its IP prefixes both in its own Prefix Assertion List (PAL) and in the PALs of a small number of direct neighbors (e.g., service providers). A prefix assertion made by X verifies successfully if it is consistent with the assertion made by one AS with which X chooses to register its prefixes. Advantages of the distributed prefix registration model used by psBGP include: 1) it distributes the difficult task of tracing IP address ownership across the Internet and thus is more scalable and practical; 2) it allows the secure inter-domain routing infrastructure to be built by ISPs more independently than a centralized approach; 3) it is resilient to a single point of failure.

The rest of the report is organized as follows. Section 2 gives a brief overview of the Internet and general routing protocols (e.g., distance vector and link state). Section 3 describes the BGP protocol and real-world operations. BGP security threats are discussed in Section 4. In Sections 5 and 6, we respectively analyze and compare three proposals (S-BGP, soBGP, and psBGP) for securing BGP. We conclude in Section 7.

2 Background - Routing Protocols

The Internet is a collection of a large number of networks operated by many Internet Service Providers (ISPs). ISPs can be classified into different tiers based on the sizes of their networks. *Tier-1 ISPs* usually have national-wide backbone networks; *tier-2 ISPs* may have a state-wide network; and *tier-3 ISPs* may have an even smaller network and usually provide Internet access to end users.

Due to the extremely large size of the Internet, a hierarchical routing approach has been adopted. Logically, the Internet consists of a number of ASes (see Figure 1), each of which consists of a number of routers under the same technical administration (e.g., using the same routing policy). An AS is identified by a 16-bit integer (this may be extended to 32 bits in the future), and usually belongs to a single ISP. For example, AS 7018 belongs to AT&T. However, one ISP may own multiple ASes. For example, UUNET owns AS 701, 702, etc.

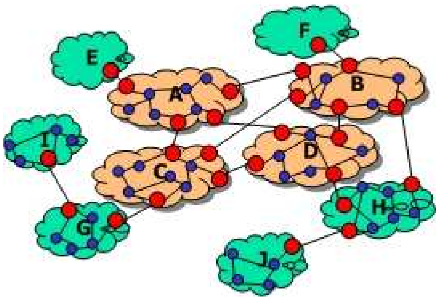

Figure 1: An example of the Internet

On the AS level, the Internet can be abstracted as a graph, where a vertex is an AS and an edge is a BGP session between two ASes. BGP is the only inter-domain routing protocol used on the Internet for exchanging reachability information between ASes. For example, in Figure 1, for a computer in AS E to communicate with another computer in AS H (assuming the path E-A-D-H is used), BGP sessions need to be established between E-A, A-D, and D-H respectively.

Within an AS, a network graph consists of *routers* (vertices) and *network links* (edges). Intra-domain routing protocols (e.g., RIP, OSPF, IS-IS) are used for exchanging reachability information *within* an AS. For example, for a computer in AS B to communicate with another non-directly connected computer located in the same AS, an intra-domain routing protocol such as RIP, is usually used to discover the path between the two computers. Such a path usually consists of a number of routers, each of which runs RIP.

There are two popular approaches used by routing protocols: *distance vector* and *link state*. In a distance vector routing protocol, each node maintains a routing table consisting of a number of *vectors*. Each vector represents a route for a particular destination in the network, and is usually measured by some distance metric (e.g., number of hops) to that destination. Each node periodically advertises its routing tables to its direct neighbors, and updates its own routing table based on the advertisements received from others. Examples of distance vector routing protocols include RIP and BGP.

In a link state routing protocol, each node advertises its link states to every other node in the network by flooding Link State Advertisements (LSAs). An LSA usually consists of a link identifier (e.g., a subnet attached to a link), state of the link, cost of the link, and neighbors of the link. Every node receives the LSAs from every other node in the network, and builds the same link state database (which is a weighted graph as each edge is associated with a cost). Each node runs Dijkstra's algorithm to compute a shortest path from itself to every other destination in the network. OSPF and IS-IS are two popular link state routing protocols.

3 BGP Protocol and Operation

In this section, we give a brief overview of the BGP protocol and its operational practice in real world deployment, including IP address allocation, AS business relationships, AS route exporting policies, and BGP route selection algorithm.

3.1 Overview of BGP

BGP is an inter-domain routing protocol based on a distance vector approach. A *BGP speaker* establishes a session over TCP with its direct neighbors, exchanges routing information with them, and updates its own routing table based on the information received from them. Unlike a simple distance vector routing protocol (e.g., RIP) where a route usually has a simple metric (e.g., number of hops), a BGP route is associated with a number of attributes and a *best route* is selected among multiple routes to the same destination based on local policy. One notable route attribute is AS_PATH, which consists of a sequence of ASes traversed by this route. Thus, BGP is often referred to as a *path vector* routing protocol.

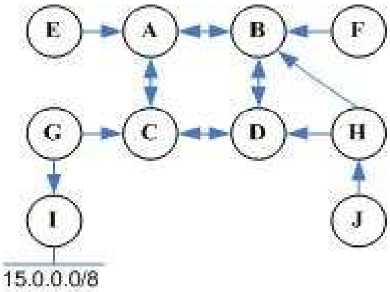

Figure 2: A BGP view of the Internet

We use Figure 2 to illustrate how BGP announcements propagate across a network. Suppose IP prefix 15.0.0.0/8 (abbreviated 15/8) is allocated to AS I (see Section 3.2.1 for IP address allocation practice). To allow other ASes to send traffic to 15/8, AS I advertises (15/8, I) to AS G. (15/8, I) is a selected portion of a *BGP update message* which consists of Network Layer Reachability Information (NLRI) and a number of attributes (e.g., AS_PATH) associated with the NLRI. In this example, NLRI is 15/8 and AS_PATH consists of AS I.

When AS G receives (15/8, I), a sequence of operations will be applied to (15/8, I), including applying route importing policies, selecting the best route, applying route exporting polices, and transforming a route (e.g., modifying the AS_PATH). We consider the simple case that (15/8, I) passes AS G's importing polices, is selected as the best route to 15/8, and passes AS G's exporting policies. AS G then transforms (15/8, I) to (15/8, G-I) by inserting its own AS number into the AS_PATH and announces the transformed route to its direct neighbor AS C.

The above process is repeated by every AS receiving the route. Eventually, J receives route (15/8, H-D-C-G-I). This route allows J to send traffic to 15/8, and J expects that its traffic will reach 15/8 via AS_PATH H-D-C-G-I. However, there is no guarantee that H-D-C-G-I will be the path traversed by traffic from J to 15/8 since each forwarding decision on the Internet is done on a hop-by-hop basis. In other words, J has no control over how other ASes will forward its traffic.

If every AS announces its IP address space through BGP, after the Internet reaches a convergence state, every other AS will have a route for reaching other ASes' IP address space. This effectively builds a routing infrastructure allowing for communications across the Internet.

3.2 BGP Operational Practice

Here we discuss some BGP operational practices which are out the scope of BGP protocol specification [33] but are nonetheless important to BGP security.

3.2.1 IP Address Allocation

The Internet Assigned Number Authority (IANA) [18] is the central authority of the whole IP address space. When the Internet was small, any organization could apply directly to IANA for a block of IP address space (or IP prefix). As the Internet grew, it became obvious that a single authority could not handle the extremely large number of IP address requests. As a result, a hierarchical structure was developed for IP address allocation.

On the top level, IANA is still the central authority of IP address space. On the second level, four Regional Internet Registries (RIRs) have been created, each of which is responsible for IP address allocation in a particular geographic location. They are: the American Registry for Internet Numbers (ARIN – www.arin.net), Reseaux IP Europeens (RIPE – www.ripe.net), Asia Pacific Network Information Centre (APNIC – www.apnic.net), and Latin American and Caribbean Internet Addresses Registry (LACNIC – www.lacnic.net).

A large ISP (e.g., tier-1) may apply for an IP address space directly from an RIR, and then delegate a portion of that address space to a downstream service provider (e.g., tier-2). A subscriber (i.e., an organization having access to the Internet but not providing Internet access service to others) may obtain IP address space directly from a tier-1 ISP or from a smaller ISP. IP address space delegation among ISPs and subscribers is mainly driven by business relationships. There is no mandated policy dictating who should get IP address space from whom. In addition, a subscriber obtaining IP address space from one ISP may buy its Internet access service from another ISP. For example, a subscriber may obtain IP address space from AT&T but connect to the Internet via Sprint.

Currently, about 180,000 IP prefixes are announced through BGP. However, it is not clear on the Internet-wide which IP prefixes have been delegated to which organizations via which ISPs. While some route registries (e.g., the Internet Routing Registries – www.irr.net) may attempt to maintain such information, it is usually out of date. The consensus is that IANA and the RIRs are responsible for initial IP address delegation, but not for keeping track of further delegation among ISPs and subscribers. To quote from a study by Atkinson and Floyd [4] on behalf of the Internet Architecture Board (IAB): "*a recurring challenge with any form of inter-domain routing authentication is that there is no single completely accurate source of truth about which organizations have the authority to advertise which address blocks*".

3.2.2 AS Business Relationships

ASes on the Internet can be roughly classified into three categories: a *stub-AS* has only one connection to other ASes; a *multihomed-AS* has more than one connection to other ASes, but is not designed to carry traffic for other ASes (e.g., for the purpose of load balance or backup); and a *transit-AS* has more than one connection to other ASes, and is designed to carry traffic for others.

Business relationships usually exist between two neighboring ASes. These are mainly derived from the cost model adopted on the Internet. The Internet has a different cost model than the traditional telephony industry in that: 1) users usually pay fixed subscription fees (e.g., a flat monthly fee) for their Internet access while paying toll voice service on a per transaction basis; 2) both the caller and the callee of an Internet transaction (e.g., a TCP connection) pay their own portion of cost, assuming each transaction incurres a certain cost; while for a voice transaction, it is usually the caller that pays the whole cost.

The cost model reflects the hierarchical structure of the Internet. At the bottom are subscribers who pay their service providers for the Internet access. Looking bottom-up, smaller service providers usually pay larger service providers for connecting through them to the Internet. At the core of the Internet, a small number of large ISPs have peer relationships and do not pay each other for accessing the others' networks. Two small ISPs may also establish a peer relationship to allow "quick" access among their customers without going through the core Internet.

To summarize, there are usually four types of AS business relationships [17, 11]: *customer-to-provider*, *provider-to-customer*, *peer-to-peer*, and *sibling-to-sibling*. A customer AS usually pays a provider AS for accessing the rest of the Internet. For example, a stub-AS is very likely a customer of the AS it connects to. Two peer ASes usually find that it is mutually beneficial to allow each other to have access to their customers. Two sibling ASes are usually owned by a common organization and allow each other to have access to the rest of the Internet.

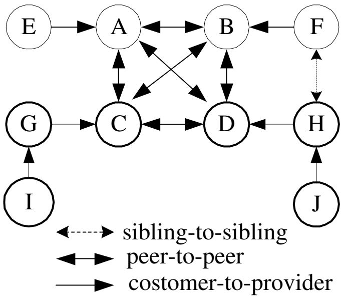

Figure 3: A simple AS topology with different types of AS relationships.

In Figure 3, ASes A, B, C, and D may attach to a Network Access Point (NAP), and form peer-to-peer relationships between each other. Each also has a direct customer, i.e., ASes E, F, G, and H are the direct customers of ASes A, B, C and D respectively. AS G has a customer AS I, and AS H has a customer AS J. AS F and H may be owned by the same ISP but are located in different geographic locations, and they form a sibling-to-sibling relationship.

3.2.3 BGP Route Exporting Policy

As discussed in Section 3.1, a BGP update message needs to pass through a number of steps before being further propagated to a next AS. One step is to check route exporting policies. While one AS might have a high degree of freedom in defining its own specific route exporting policies (e.g., for traffic engineering purposes), there are some general rules which should be obeyed. These rules are mainly derived from AS business relationships, and are summarized below (see [17, 11] for further discussion):

- *customer-to-provider*: a customer AS X of AS Y exports to Y its own routes and the routes it has learned from its customers. This exporting policy allows Y to further propagate routes destined to AS X and to X's customers (including customers' customers). The ultimate goal for X to export routes to Y is to receive from Y the traffic destined to itself or to its customers.

- *provider-to-customer*: a provider AS Y of AS X exports to X its full routing tables, including its own routes, the routes it has learned from customers, providers, peers, and siblings. This allows X to send to Y the traffic destined to the rest of the Internet.

- *peer-to-peer*: a peer AS X of AS Y exports to Y its own routes and the routes it has learned from its customers. A peer relationship is symmetric, thus Y is also a peer of X. This allows Y to send to X traffic destined to X and its customers, and vice versa.

- *sibling-to-sibling*: a sibling AS X of AS Y exports to Y its full routing tables, including its own routes, the routes it has learned from its customers, providers, peers, and siblings. A sibling relationship is symmetric, thus Y is also a sibling of X. This allows two sibling ASes to access through each other the rest of the Internet.

3.2.4 BGP Route Selection Process

The BGP specification (see §9 in [33]) defines some basic rules for selecting the most preferable route among a set of routes for a common destination. In practice, a larger set of route selection rules are usually implemented. For example, AS_PATH is not mandated to be used as part of a route selection process by the BGP specification. However, it is commonly used in practice, e.g., by Cisco IOS. Here we summarize a list of route selection rules with an order of decreased preference:

1. Select the route with the highest degree of preference. Preference values are configurable based on local policy, and are usually assigned to routes (i.e., assigning LOCAL_PREF values during the route importing process) based on the business relationship with the advertising AS. For example, a higher LOCAL_PREF value is usually assigned to routes received from a customer AS than a provider or a peer.

2. Select the route with the shortest AS_PATH if all routes have the same preference value.

3. Select the route with the lowest MULTI_EXIT_DISC (MED) among those with the same NEXT_HOP. MED is used by an advertising AS to influence which link inbound traffic will be received.

4. Select the route with the lowest cost to the NEXT_HOP of that route. The cost to the NEXT_HOP is determined by an intra-domain routing protocol, e.g., OSPF.

5. Select the route advertised by a BGP speaker with the lowest BGP identifier.

4 BGP Security Threats

In this section, we discuss a number of BGP security threats. We start with an overview of potential threat sources and malicious actions an adversary may take to attack BGP. We then focus on two serious falsification attacks.

4.1 Sources of Threats

BGP is based on TCP and IP. Thus, it is vulnerable to all threats against its underlying protocols. For example, BGP is vulnerable to a TCP Reset attack [38] which can result in significant Internet instability. BGP best practices [9] may help mitigate those threats. Here we consider threats against the BGP protocol itself.

BGP faces threats from both BGP speakers and BGP sessions (see Figure 4). For example, a BGP speaker may be compromised (e.g., by exploiting software flaws), misconfigured (mistakenly or intentionally), or unauthorized (e.g., by exploiting a BGP peer authentication vulnerability). An attacker can also set up its own BGP speaker and connect it to the Internet by purchasing connection service from a sloppy ISP (this is indeed happening on the Internet [3]). In addition, a BGP session may be compromised or unauthorized.

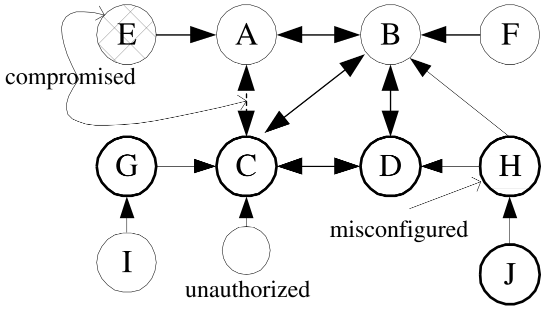

Figure 4: Sources of threats against BGP

4.2 Malicious Actions

Attacks against BGP *control messages* (see next paragraph) include, for example, *modification, insertion, deletion, exposure,* and *replaying* of messages. In this report, we focus on modification and insertion (hereafter *falsification* [5]) of BGP control messages. Deletion appears indistinguishable from legitimate route filtering. Exposure might compromise confidentiality of BGP control messages, which may or may not be a major concern [5]. Replaying is a serious threat, which can be handled by setting expiration time for a message; however it seems challenging to find an appropriate value for an expiration time.

There are four types of BGP control messages: OPEN, KEEPALIVE, NOTIFICATION, and UPDATE. The first three are used for establishing and maintaining BGP sessions with neighbors, and falsification of them will very likely result in session disruption. These messages, along with underlying transport mechanisms (e.g., TCP) can be protected by a point-to-point authentication protocol, e.g., IPsec [19]. We concentrate on falsification of BGP UPDATE messages (hereafter, we refrain from capitalizing update as UPDATE) which carry inter-domain routing information and are used for building up routing tables.

A BGP update message consists of three parts: withdrawn routes, network layer reachability information (NLRI), and path attributes (e.g., AS_PATH, LOCAL_PREF, etc.). A route should only be withdrawn by a party which had previously announced that route. Otherwise, a malicious entity could cause service disruption by withdrawing a route which is actually in service. Digitally signing BGP update messages allows one to verify if a party has the right to withdraw a route. Here we examine in detail falsification of NLRI and one of the most important route attributes – AS_PATH. Other route attributes (e.g., LOCAL_PREF, COMMUNITY, etc) are also important. However, they are either non-transitive (i.e., not propagated beyond an AS) or transitive but static (i.e., unchanged when being propagated between ASes). Thus, they are easy to protect.

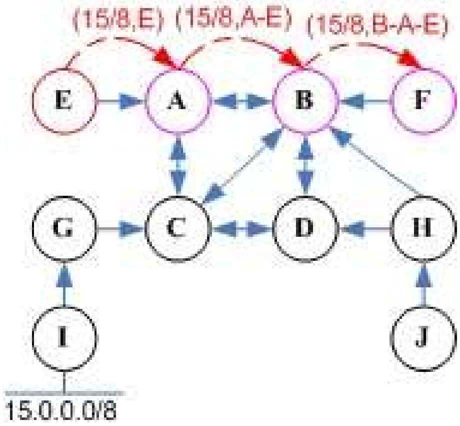

Figure 5: An AS topology with attackers

4.3 Falsification of NLRI

NLRI consists of a set of IP prefixes sharing the same characteristics as described by the path attributes. NLRI is falsified if an AS originates a prefix not owned by that AS, or aggregated improperly from other routes. Falsification of NLRI is often referred to as *prefix hijacking*, and can cause serious consequences including denial of service and man-in-the-middle (MITM) attacks.

We use Figure 5 to illustrate how an attacker controlling a BGP speaker in E (i.e., the router establishing a BGP session with A) might hijack 15/8 which is allocated to I. We assume that the network has converged on 15/8, i.e., every AS has a route to 15/8 (see Table 1).

AS	Route to 15/8	AS	Route to 15/8
A	(15/8, C-G-I)	F	(15/8, B-C-G-I)
B	(15/8, C-G-I)	G	(15/8, I)
C	(15/8, G-I)	H	(15/8, D-C-G-I)
D	(15/8, C-G-I)	I	direct route
E	(15/8, A-C-G-I)	J	(15/8, H-D-C-G-I)

Table 1: Routes to 15/8 from each AS

(AS E) This AS configures a BGP speaker under its control to advertise route (15/8, E) to A. Since 15/8 is *n*ot allocated to E (it is allocated to I), it is illegitimate for E to originate route (15/8, E). However, an attacker does not play by rules.

(AS A) After receiving (15/8, E), A may have two distinct routes to destination 15/8: (15/8, E) and (15/8, C-G-I). A will select one from them as preferable using the route selection process as described in §3.2.4. Assume that A implements a common policy in which a customer route is preferred over a provider route or a peer route. In other words, among a set of routes with the same destination prefix, the route received from a customer AS is preferred over those received from a provider or a peer AS. Thus, (15/8, E) is preferred over (15/8, C-G-I) since E is a customer of A and C is a peer of A. As a result, (15/8, E) is installed on E's routing table, and *A's routing table is poisoned*.

Since (15/8, E) is learned from A's customer, E will also re-advertise it as (15/8, A-E) to B and C (see §3.2.3 for peer-to-peer route exporting policy).

(AS C) After receiving (15/8, E), C will compare it with (15/8, G-I). Assume that C implements a common policy that a customer route is preferred over a provider route or a peer route. Since

G is a customer of C and B is a peer, (15/8, G-I) will be selected. Thus, *C's routing table is not poisoned*.

(AS B) When B receives (15/8, A-E), it will compare it with (15/8, C-G-I) assuming (15/8, C-G-I) has been received from C. Since B has a peer relationship with both A and C, the preference values assigned to the two routes might be the same. Thus, the second rule in the route selection process (cf. §3.2.4) will be applied, favoring the shorter AS_PATH. So (15/8, A-E) will be selected. *B's routing table is poisoned*.

B will also propagate (15/8, B-A-E) to F and H because they are its customers (see §3.2.3 for provider-to-customer route exporting policy). However, B will *n*ot propagate this route to C and D because they are its peers (see §3.2.3 for peer-to-peer route exporting policy).

(AS F) After receiving (15/8, B-A-E), F uses it to replace the existing route to 15/8, i.e., (15/8, B-C-G-I) without going through route selection process because in BGP, a new route will automatically replace an old one if they are received from the same source (e.g., B in this case). *F's routing table is poisoned*.

(AS H) After receiving (15/8, B-A-E) from B, F needs to compare it with (15/8, D-C-G-I). If we suppose the link H-D is a primary link and link H-B is a backup one (e.g., H-D is more cost effective than H-B), then F will assign a higher preference value to the routes received from H than those from B. AS a result, (15/8, B-A-E) is not selected. *H's routing table is n*o*t poisoned*.

After the above process, the routing tables of A, B and F are poisoned and the routing tables of G, C, D, H, J are *n*ot poisoned (see Table 2). As a result, traffic destined to 15/8 and initiated from A, B, and F will be forwarded to E, not to the real address owner I. In other words, prefix 15/8 has been *hijacked* from I from the view point of some part of the network.

AS	Route to 15/8	AS	Route to 15/8
A	(15/8, C-G-I) → **(15/8,E)**	F	(15/8, B-C-G-I) → **(15/8, B-A-E)**
B	(15/8, C-G-I) → **(15/8, A-E)**	G	(15/8, I)
C	(15/8, G-I)	H	(15/8, D-C-G-I)
D	(15/8, C-G-I)	I	direct route
E	(15/8, A-C-G-I)	J	(15/8, H-D-C-G-I)

Table 2: Routes to 15/8 from each AS after the attack

Prefix hijacking can be used to facilitate many types of attacks, including *denial of service, man-in-the-middle (MITM),* or *service hijacking* (e.g., email). While service hijacking will always deny the service of a real address holder, it also has the purpose of impersonation. Therefore, it could cause more serious consequences. Here we present three types of attacks using service hijacking: *spamming, interception of password Reset messages,* and *Phishing.* The first two attacks described here are related to email server impersonation, and the third attack is related to web server impersonation.

4.3.1 Advanced Spamming

Recently, falsification of NLRI might have been used by spammers to facilitate advanced spamming [7]. Here we describe how spammers can use prefix hijacking to bypass some email authentication mechanisms. We first give an overview of the Simple Mail Transfer Protocol (SMTP), then introduce how a sender address can be spoofed, followed by a description of a proposed email authentication mechanism. Finally, we show how to use email server hijacking to bypass email authentication.

SMTP Basics. Figure 6 illustrates the SMTP message flow between an originating SMTP server "alice.com" and a receiving SMTP server "bob.com" for delivering an email message from "x1@alice.com" to "y1@bob.com". Note the sender address specified by the SMTP command "HELO" and "MAIL FROM" can be forged to an arbitrary address if "bob.com" does not employ any authentication mechanism. This is exactly the vulnerability exploited by spammers.

Sender Address Spoofing. A spammer usually sends a large number of people unsolicited emails with spoofed sender addresses. Since SMTP does not verify the authenticity of an originating party's domain name, a spammer can use a single SMTP engine (e.g., running on a compromised PC) to send out spams with arbitrary sender addresses. Figure 7 shows how an attacker sends out spams from "attack.com" to "bob.com" using "alice.com" as the sender domain.

Email Authentication. A number of mechanisms have been proposed for fighting spams by authenticating sender addresses. Sender Policy Framework (SPF) [25] is a popular proposal which has been adopted by some organizations. SPF requires a domain running SMTP servers to publish in DNS the identities (e.g., IP addresses) of its authorized outgoing SMTP servers. An SMTP server implementing SPF can verify the authenticity of a sender address (i.e., the domain name in the MAIL FROM field) by checking the consistency between the IP address of an originating SMTP server and the IP addresses of the authorized SMTP servers published by the sender domain.

For example in Figure 7, "alice.com" publishes in DNS 15.15.2.7 as the IP address of its authorized outgoing email server. Upon receiving from "attack.com" the SMTP commands "HELO

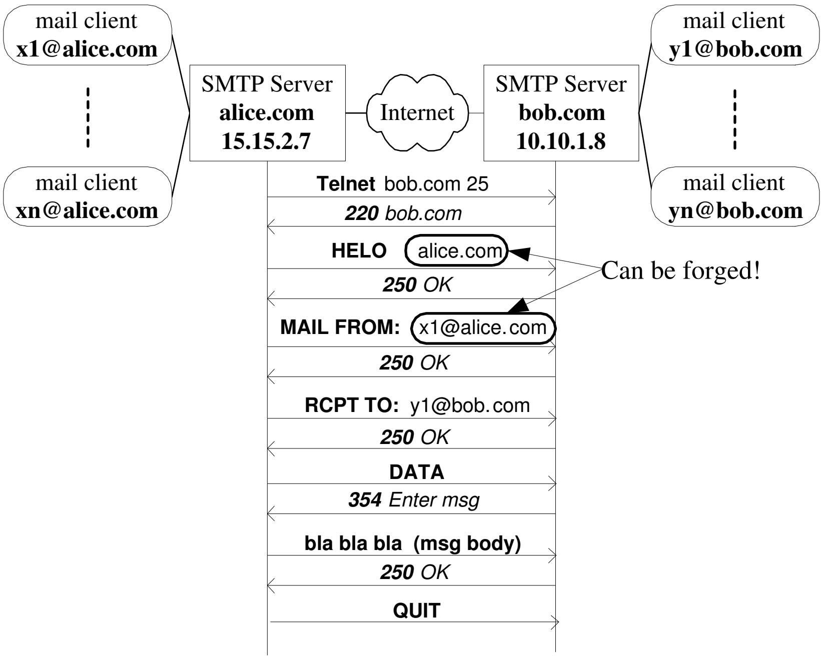

Figure 6: SMTP message flow

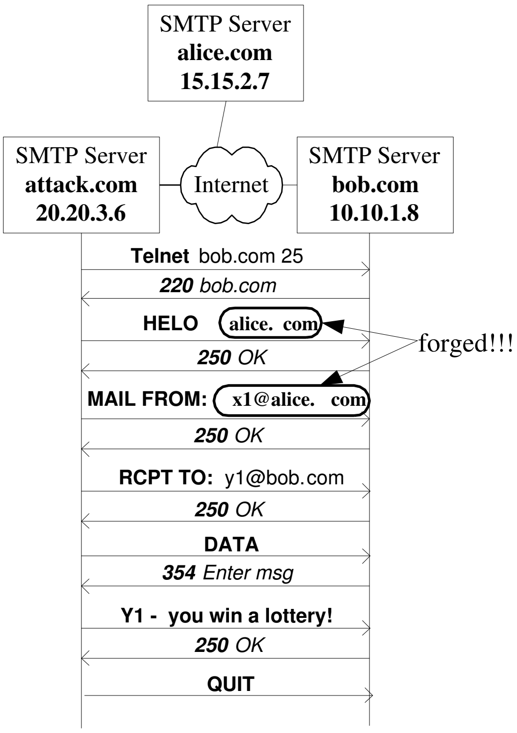

Figure 7: Spamming - sender address spoofing

alice.com" (which may be omitted by a sender) or "MAIL FROM: alice.com", the SMTP server in "bob.com" verifies the sender IP address "20.20.3.6" against the IP address of the authorized SMTP server published by "alice.com" which is "15.15.2.7"[1]. Since they are inconsistent, the SMTP server in "bob.com" will detect that this email is originated by an unauthorized party (or with a spoofed sender address), thus will NOT accept it. If every domain adopts this mechanism, it is expected that a significant amount of spams will be detected and dropped.

Defeating Email Authentication. However, authentication mechanisms such as SPF can be defeated by prefix hijacking. A spammer who wants to send out spams using the domain name "alice.com" can hijack the IP address space containing the authorized IP addresses published by "alice.com". For example, the spammer with control of a BGP speaker can announce routes for prefix 15.15.2.0/24, and set up a SMTP server with IP "15.15.2.7". This allows the spammer to use the hijacked IP address "15.15.2.7" to establish SMTP connections with "bob.com" and send out spams using 'alice.com' as the domain of the sender address. Email authentication mechanisms such as SPF will not be able to detect this type of spamming. In fact, any authentication mechanism based only on IP address can be defeated by prefix hijacking.

4.3.2 Interception of Password Reset Messages

One possible attack using prefix hijacking is to intercept password reset messages[2] for gaining illegitimate access to other people's email accounts. A traditional way of doing this is to crack the password of a victim account by either offline or online dictionary attacks. Offline dictionary attack usually requires access to the password database (e.g., /etc/passwd in Unix) which may not be possible. Online dictionary attack usually involves automatic logon retries with candidate passwords (e.g., chosen from a dictionary). Since some email service providers have adopted reverse Turing tests to defeat automatic logon retries, it becomes more difficult for online dictionary attack to succeed.

However, many email services provide "user-friendly" features to allow users to reset their passwords in the case they forget them. When a link such as "forgot your password" is clicked, a password reset message is sent to another email account (namely backup email account) associated with the account whose password has been forgot (namely primary email account). A backup email address is usually asked by many email service providers for authentication purpose such as receiving password reset message. A password reset message may contain an automatically

[1] To publish the IP addresses of authorized email servers, a domain needs to add new records, namely SPF records, into its DNS records. A verifier can then lookup DNS for an SPF record to obtain the IP address of the authorized email server for a particular domain.

[2] This attack was mentioned to us by Dan Boneh during a conversation at NDSS'05.

generated new password, or a link pointing to a page where the user can type in a new password without being asked for the old password.

The assumption made here is that a backup email address is only accessible to its owner. This assumption usually holds since an email account is usually password protected and it appears difficult to intercept an email message if an attacker does not have access to one of the following communication paths: 1) from the mail server originating a message to the mail server receiving it, and 2) from the mail client retrieving the message to the mail server storing it.

However, such an assumption will loose ground if an attacker can manipulate BGP to hijack IP prefixes. Suppose a user has a primary email address “x1@alice.com”, and the backup email address associated with this account is “x1@bob.com”. An attacker may gain access to “x1@alice.com” by performing the following steps:

1) looking up the IP address of the email server of “bob.com” (e.g., by looking up the MX record of “bob.com” in DNS), which is 10.10.1.8 (see Figure 7);

2) hijacking 10.10.1.8 by announcing a BGP route for the prefix 10.10.1/24, assuming that 10.10.1/24 is the most specific prefix on the network;

3) requesting password reset for “x1@alice.com”;

4) intercepting the password reset message sent from “alice.com” to “x1@bob.com”, e.g., by setting up an email server with the IP address 10.10.1.8. Since the IP prefix containing 10.10.1.8 has been hijacked, the password reset message will be sent to the attacker instead of the legitimate mail server of “bob.com”.

5) resetting the password for “x1@alice.com” by following instructions in the intercepted password reset message. As a result, the attacker gains access to “x1@alice.com”.

While some online service providers (e.g., Expedia) may accept requests for password resets without asking for any additional information (except the userid of the account being reset for password), many (e.g., Yahoo) do take additional steps for verifying identities. In other words, additional information is often required to show that you really are the owner of the account whose password will be reset. For example, Yahoo asks for a date of birth and a postal code, and Ebay asks for a postal code and a phone number. Gmail asks for characters in a picture for countering automatic password reset attacks, but not for identity verification. However, most information requested for countering identity theft could be obtained, e.g., by social engineering.

4.3.3 Phishing

A primary objective of *phishing* is to steal people's confidential information, e.g., credit card numbers, social insurance numbers, date of birth, home addresses, etc. so that they can be used directly or indirectly (sold to a third party) for financial benefit. A phisher usually sends out spams to a large number of people using well-known sender addresses (e.g., the email address of the security team of a well-known bank) to ask a recipient to reset its account by going to a spammer-controlled website and filling in confidential information. The link to a fraudulent website can be a numeric IP address, an irrelevant domain name, or a domain name very similar to the real one of a claimed organization. The displayed URL which a potential victim sees may also be entirely different than the URL linked to in the underlying html. However, a careful user may be able to find the discrepancy and thus avoid being fooled. The legitimate domain name or URL can also be used if its DNS record on a victim machine (i.e., the machine from which a user clicks the link) is changed (poisoned) to the IP address of the fraudulent website. Again, a careful user may still be able to notice the trick.

To use the legitimate domain or URL of a claimed organization in a phishing email without poisoning a DNS record, a phisher can hijack the IP address space of that organization and set up a fraudulent website using the IP address of the legitimate website. In this way, it will be difficult (essentially impossible) for a user to distinguish a phishing message from a real message (i.e., a message indeed sent by the organization in question). As shown in Figure 5, some ASes (more precisely the routing tables of BGP speakers in some ASes) may not be poisoned by a bogus prefix announcement, depending on their locations and relevant routing policies. Thus, users located in these ASes may go to the real website by clicking the link in a phishing email. However, some ASes may be poisoned and their users will face the risk of being phished.

4.4 Falsification of AS_PATH

There are two types of AS_PATH: AS_SEQUENCE and AS_SET. An AS_PATH of type AS_SEQUENCE consists of an ordered list of ASes traversed by the route in question. An AS_PATH of type AS_SET consists of an unordered list of ASes, sometimes created when multiple routes are aggregated. Here we focus on the security of AS_SEQUENCE. (Note: AS_SET is less widely used on the Internet. For example, as of August 1, 2004, only 23 of 17 884 ASes originated 47 of 161 796 prefixes with AS_SET.) An AS_PATH is *falsified* if an AS or any other entity illegally operates on an AS_PATH, e.g., inserting a wrong AS number, deleting or modifying an AS number on the path, etc. Since AS_PATH is used for detecting routing loops and used by route selection processes, falsification of

AS_PATH can result in routing loops or selecting routes not selected otherwise.

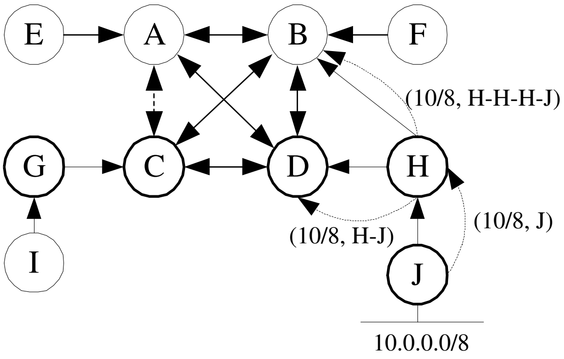

Figure 8: Changing traffic flow by AS_PATH falsification

We use Figure 8 to illustrate how an attacker might influence traffic flow by manipulating AS_PATH. Suppose AS H multi-homes with D and B; H-D is a primary link and H-B is a backup link. In the normal situation, traffic destined to AS H and H's customers (e.g., AS J) should go through link H-D. When H-D fails, H-B should then be used. To achieve this traffic engineering objective, AS H can legitimately utilize AS_PATH to influence other ASes' routing decisions. For example, AS H announces (10/8, H-J) to AS D (normal BGP operation), but (10/8, H-H-H-J) to AS B (this is a legitimate traffic engineering technique). After the network converges on 10/8, all traffic to 10/8 will be forwarded over link D-H to AS H (see Table 3).

However, B can attract traffic destined to 10/8 by announcing a route to 10/8 with a fraudulent AS_PATH, e.g., (10/8, B-J). Note the AS_PATH B-J is shorter than B-H-H-H-J which is supposed to be advertised by B. As a result, other ASes may select the route to 10/8 which goes through AS B. See Table 3 for details of route changes. To summarize, traffic flow can be changed by falsification of AS_PATH.

AS	Route to 10/8 †	AS	Route to 10/8 †
A	(10/8, D-H-J) → **(10/8,B-J)**	F	(10/8, B-H-H-H-J) → **(10/8, B-J)**
B	(10/8, H-H-H-J) → **(10/8, H-J)**	G	(10/8, C-D-H-J) → **(10/8, C-B-J)**
C	(10/8, D-H-J) → **(10/8, B-J)**	H	(10/8, J)
D	(10/8, H-J)	I	(10/8, G-C-D-H-J) → **(10/8, G-C-B-J)**
E	(10/8, A-D-H-J) → **(10/8, A-B-J)**	J	direct route

Table 3: Routes to 10/8 from each AS before and after B announces fraudulent (10/8, B-J). † - Note the "after" route as listed herein may not actually exist.

5 BGP Security Mechanisms

We first summarize a number of security goals for BGP and relate them to the BGP security threats presented in Section 4. We then discuss security mechanisms adopted by each of the three BGP security proposals (S-BGP, soBGP, and psBGP) and show how each of them achieves these security goals.

5.1 BGP Security Goals

BGP is a distributed communication protocol which faces threats from both outsiders and insiders. Outsiders include unauthorized BGP speakers and compromised links, and insiders include compromised authorized BGP speakers (see Figure 4). To prevent outsider attacks, data origin authentication which includes data integrity [26] can be used. It appears difficult, if not impossible, to prevent insider attacks since: 1) an authorized BGP speaker may run flawed software and can be compromised by an attacker by exploiting software vulnerability; 2) a legitimate person with access to an authorized BGP speaker may be malicious; and 3) an authorized BGP speaker might be misconfigured. Thus, the ultimate goal here is NOT to prevent insider attacks from happening but to contain their damages. Particularly, fraudulent BGP update messages should be detected and discarded so that routing tables of well-behaved BGP speakers are not poisoned.

We summarize five security goals for BGP (cf. [21, 22]). G1 and G2 relate to data origin authentication, G3 to data integrity, and G4 and G5 to the propriety of BGP messages. G1, G2, and G3 can prevent outsider attacks. G4 and G5 can respectively contain insider attack damages caused by falsifications of NLRI (see 4.3) and of AS_PATH (see 4.4).

G1. *(AS Number Authentication)* It must be verifiable that an entity using an AS number s_i as its own is in fact an authorized representative of the AS to which a recognized AS number authority assigned s_i.

G2. *(BGP Speaker Authentication)* It must be verifiable that a BGP speaker, which asserts an association with an AS number s_i, has been authorized by the AS to which s_i was assigned by a recognized AS number authority.

G3. *(Data Integrity)* It must be verifiable that a BGP message has not been illegally modified en route.

G4. *(Prefix Origination Verification)* It must be verifiable that it is proper for an AS to originate an IP prefix. More specifically, it is proper for AS s_i to originate prefix f if 1) f is delegated to s_i by an authoritative party; or 2) f is aggregated from a set F of prefixes such that $f \subseteq F$.

G5. *(AS Path Verification)* It must be verifiable that an AS_PATH ($p_k = [s_1, s_2, \dots, s_k]$) of a BGP route (f, p_k) is originated by s_1, and has traversed through $s_2, \dots, s_k$ in order. In addition, it must be verifiable that for all $1 \leq i \leq k$, advertising (f, p_i) to s_{i+1} by s_i does not violate s_i's route exporting policy as determined by the business relationship between s_i and s_{i+1} (cf. 3.2.3).

5.2 BGP Security Proposals

Many solutions (e.g., [35, 12, 2, 16]) have been proposed for securing BGP. Here we describe three BGP security proposals: S-BGP [21, 22], soBGP [40], and psBGP [37].

5.2.1 Secure BGP (S-BGP)

S-BGP proposes use of two strict hierarchical PKIs and other mechanisms (e.g., IPsec [19]) for securing BGP. The proposed S-BGP PKIs are parallel to the existing allocation and delegation systems for AS numbers and IP address space. A single Certificate Authority (CA) rooted at IANA/ICANN was initially proposed for S-BGP, but it evolved to multiple CAs rooted at four RIRs due to political sensibility and security considerations. We use T to denote a trusted CA (i.e., an RIR).

There are many types of certificates in S-BGP. An organization X which obtains IP address space and AS numbers directly from an RIR, will be issued the following certificates[3]:

- *Organization Public Key Certificates* – binding a public key K_x to X signed by T, denoted by $(K_x, X)_T$;
- *Address Delegation Certificates* – binding IP prefixes f_x to X signed by T, denoted by $(f_x, X)_T$;
- *AS Number Delegation Certificates* – binding an AS number (or more) s_x to X signed by T, denoted by $(s_x, X)_T$.

To participate in the inter-domain routing, X issues the following certificates or attestations:

- *Router Public Key Certificate* – binding a public key K_{r_x} to a BGP speaker r_x and an AS number s_x signed by X using K_x, denoted by $(K_{r_x}, s_x, r_x)_{K_X}$;

[3] For convenience of presentation, certificate names used here may differ from those used in the S-BGP literature.

- *Address Attestation* – binding IP prefixes f_x or a subset of f_x to an AS number (s_x) signed by X, denoted by $(f_x, s_x)_{K_x}$;
- *Route Attestation* – binding IP prefixes f_i to an AS_PATH p_j (along with other path attributes) signed by a BGP speaker r_x. For sake of simplicity, we only consider AS_PATH here, thus a Route Attestation can be denoted by $(f_i, p_j)_{K_{r_x}}$.

With all these certificates, we now show how a BGP speaker announces and verifies a route in S-BGP. Let r_x be a BGP speaker, representing AS s_x owned by organization X. Let f_x be an IP prefix allocated to X by an RIR, and assigned by X to AS s_x. We use a simple topology consisting three ASes s_x, s_y and s_y owned by organizations X, Y, and Z respectively. s_x connects to s_y which also connects to s_z. For simplicity, we assume that each AS has one BGP speaker.

Route Announcement. r_x originates and signs a route $(f_x, s_x)_{r_x}$, and forwards it to its neighboring BGP speaker r_y representing AS s_y. r_y verifies the received route (see next paragraph). If the route verification succeeds, r_y forwards the route to its neighboring BGP speaker r_z representing AS s_z. r_y needs to send to r_z:

- $(f_x, s_x)_{r_x}$ – the signed route received from r_x; and
- $(f_x, s_x\text{-}s_y)_{r_y}$ – the route with updated AS_PATH and signed by r_y.

Route Verification. Upon receiving $(f_x, s_x\text{-}s_y)_{r_y}$, r_z performs the following verifications:

- Is the first AS on the AS_PATH, s_x, authorized to originate IP prefix f_x? Prefix origin verification succeeds if there exist the following valid certificates[3]:

  $(K_x, X)_T$, $(f_x, X)_T$, $(s_x, X)_T$, $(f_x, s_x)_{K_x}$.
- Is an AS on the AS_PATH authorized by the previous AS to further propagate the route? In this example, is s_y authorized by s_x to further propagate the route? The AS_PATH s_x-s_y verifies successfully if there exists a route attestation $(f_x, s_x)_{r_x}$. Of course, it must first be verified that BGP speaker r_x has been authorized by organization X to represent AS s_x.

S-BGP is one of the earliest BGP security proposals, and is probably the most concrete one. It provides a strong guarantee of prefix origin verification and AS_PATH integrity. However, it has some drawbacks: 1) the proposed S-BGP PKIs are complex and face significant deployment challenges [4]; 2) AS_PATH verification is computational expensive; and 3) AS_PATH verification cannot detect violation of route exporting policy.

[3] For simplicity, here we do not consider IP prefix delegation among organizations. For example, X can delegate a prefix f_i which is a portion of its allocated prefix f_x to another organization Y by issuing a certification $(f_i, Y)_X$.

5.2.2 Secure Origin BGP (soBGP)

soBGP [40] proposes use of a web-of-trust model for authenticating AS public keys and a hierarchical structure for verifying IP prefix ownership. Each AS has a public key certificate, binding an AS number with a public key, signed by a "trusted" public key. To bootstrap trust, a small number of "root public key certificates" are distributed using out-of-band mechanisms. Some tier-1 ISPs and well-known authentication service providers (e.g., Verisign) are suggested to be candidates of trusted public key certificate authorities. An AS with a trusted AS public key certificate (e.g., signed by a trusted CA) may further sign a public key certificate for another AS, thus naturally forming a web-of-trust model. While a web-of-trust model has strong proponents for authenticating user public keys within the technical PGP community [42], it would appear to be less suitable for authenticating public keys of ASes which are identified by AS numbers strictly controlled by IANA; thus it is questionable if any entity other than IANA should be trusted for signing AS public key certificates.

With respect to IP prefix ownership verification, soBGP makes use of a strictly hierarchical structure similar to that of S-BGP. Prefix delegation structures might be simplified in soBGP by using ASes instead of organizations, however, it is not clear if it is practical to do so since IP addresses are usually delegated to organizations not to ASes [2]. We suggest that soBGP, like S-BGP, also faces difficulty in tracing changes of IP address ownership in a strict hierarchical way. Thus, both S-BGP and soBGP have made architectural design choices which arguably lead to practical difficulties.

5.2.3 Pretty Secure BGP (psBGP)

In [37] we present a new proposal for securing BGP, namely Pretty Secure BGP (psBGP), motivated by our analysis of the security and practicality of S-BGP and soBGP, and in essence, combining their best features. Our objective is to explore alternative policies and tradeoffs to provide a reasonable balance between security and practicality. psBGP makes use of a centralized trust model for authenticating AS numbers, and a decentralized trust model for verifying IP prefix ownership; the latter is in line with the IAB recommendations [4]. One advantage of psBGP is that apparently it can successfully defend against threats from uncoordinated, misconfigured or malicious BGP speakers in a *practical* way. The major architectural highlights of psBGP are as follows:

1) psBGP makes use of a *centralized trust model* for AS number authentication. Each AS obtains a public key certificate from one of a number of the trusted certificate authorities, e.g., RIRs, binding an AS number to a public key. We suggest that such a trust model provides best

possible authorization of AS number allocation and best possible authenticity of AS public keys. Without such a guarantee, an attacker may be able to impersonate another AS to cause service disruption.

2) psBGP makes use of a *decentralized trust model* for verifying the propriety of IP prefix ownership. Each AS creates a *prefix assertion list (PAL)* consisting of a number of bindings of an AS number and prefixes which are asserted to be originated by that AS, one such assertion for itself and one for each of its neighboring ASes. If an AS chooses not to endorse the prefix assertion of a neighboring AS, there will still be an entry for that AS but with an empty or null prefix field. A prefix ownership assertion made by an AS is *proper* if it is consistent with the assertion made by at least one of its neighbors which chooses to provide prefix endorsement. In this way, we distribute the difficult task of tracing IP address ownership across all ASes on the Internet. Assuming reasonable due diligence in tracking IP address ownership of selected subset of direct neighbors, and assuming no two ASes in collusion, a single misbehaving AS originating improper prefixes will be detected because they will cause inconsistency with prefix assertions made by its asserting peers.

6 Comparison of S-BGP, soBGP and psBGP

We compare the different approaches taken by S-BGP, soBGP, and psBGP for achieving the BGP security goals listed in §5.1. Table 4 provides a summary. We see that psBGP falls somewhere between S-BGP and soBGP in several of the security approaches and architectural design decisions, but makes distinct design choices in several others.

6.1 AS Number Authentication

Both S-BGP and psBGP use a centralized trust model for authenticating AS numbers, which is different from the web-of-trust model used by soBGP. The difference between the AS number authentication of psBGP and S-BGP is that S-BGP follows the existing structure of AS number assignment more strictly than psBGP. In S-BGP, an AS number is assigned by IANA to an organization and it is an organization that creates and signs a certificate binding an AS number to a public key (thus, a two-step chain). In psBGP, an ASNumCert is signed directly by IANA (depth=1), and is independent of the name of an organization. Thus, psBGP has less certificate management overhead than S-BGP, requiring fewer certificates. In addition, some changes in an organization X may not require revoking and reissuing the public key certificate of the AS controlled by X. For example, if X changes its name to Y but the AS number s associated with X does not change, psBGP does not need to revoke the ASNumCert $(k_s, s)_T$. However, in S-BGP, the public key certificates $(k_x, X)_T, (k_s, s)_{k_X}$ might be revoked, and new certificates $(k_y, Y)_T, (k'_s, s)_{k_y}$ might be issued.

6.2 BGP Speaker Authentication

In S-BGP, a public key certificate is issued to each BGP speaker, while both soBGP and psBGP use one common public key certificate for all speakers within one AS. Thus, soBGP and psBGP require fewer BGP speaker certificates (albeit requiring secure distribution of a common private key to all speakers in an AS).

6.3 Data Integrity

S-BGP uses IPsec for protecting BGP session and data integrity. Both soBGP and psBGP adopt this approach. TCP MD5 [15] is supported by all three proposals for backward compatibility. In addition, automatic key management mechanisms can be implemented for improving the security of TCP MD5.

6.4 Prefix Origin Verification

Both S-BGP and soBGP propose a hierarchical structure for authorization of the IP address space; however S-BGP traces how IP addresses are delegated among organizations, while soBGP only verifies IP address delegation among ASes. It appears that soBGP simplifies the delegation structure and requires fewer certificates for verification; however, it is not clear if it is feasible to do so in practice since IP addresses are usually delegated between organizations, not ASes. In psBGP, consistency checks of PALs of direct peers are performed to verify if it is proper for an AS to originate an IP prefix. Therefore, psBGP does not involve verification of chains of certificates (instead relying on offline due diligence). We note that while psBGP does not guarantee perfect security of the authorization of IP address allocation or delegation, as intended by S-BGP and soBGP, it is not clear if the design intent in the latter two can actually be met in practice.

6.5 AS_PATH Verification

Both S-BGP and psBGP verify the integrity of AS_PATH based on its definition in the BGP specification [33]. In contrast, soBGP verifies the plausibility of an AS_PATH. Thus, S-BGP and psBGP provide stronger security of AS_PATH than soBGP, at the cost of digital signature operations which might slow down network convergence. Regarding route exporting policy verification, none of them has a solution. We are currently working on a mechanism to allow psBGP to verify if an AS_PATH conforms to the route exporting policies of every AS on the path.

Goal	S-BGP	soBGP	psBGP
G1: AS Number Authentication	centralized (multiple levels)	decentralized (with trust transitivity)	centralized (depth=1)
G2: BGP Speaker Authentication	one certificate per BGP speaker	one certificate per AS	one certificate per AS
G3: Data Integrity	IPsec or TCP MD5	IPsec or TCP MD5	IPsec or TCP MD5
G4: Prefix Origination Verification	centralized (multiple levels)	centralized (multiple levels)	decentralized (no trust transitivity)
G5: AS_PATH Verification	integrity	plausibility	integrity

Table 4: Comparison of S-BGP, soBGP, and psBGP re: achieving BGP security goals.

7 Concluding Remarks

BGP is the only inter-domain routing protocol used on the Internet. It is vulnerable to a variety of attacks, and it must be secured to protect the Internet routing infrastructure, which is now clearly recognized as a critical infrastructure. There are several proposals for securing BGP. However none of them has been deployed. We suggest that psBGP combines the best features of S-BGP and soBGP, while differing fundamentally in the approach taken to verify IP prefix ownership. As no centralized infrastructure for tracing changes in IP prefix ownership currently exists, and it would appear to be quite difficult to build such an infrastructure. Thus, we suggest that the decentralized approach taken by psBGP provides significant deployment advantages.

Securing BGP and doing so in such a way that it will actually be both deployable and deployed requires collaboration among many parties, e.g., router vendors and ISPs. While many stake holders are aware of the problem, none of them has taken initiative to push it forward. One operational obstacle is that extra costs will incur from developing and deploying BGP security solutions. With the current downturn in the telecommunications industry, cost reduction has become a primary objective of many router vendors and ISPs. Thus, it appears unrealistic to expect ISPs to start to spend on deploying BGP security solutions which do not provide to them an immediate return on investment. Router vendors are not motivated either to develop BGP security solutions due to the lack of interest from ISPs.

We suggest that governments can play an important role to facilitate the development and deployment of more secure versions of BGP. While the Internet is mainly built and operated by ISPs, it is now of general public interest since most people and especially all businesses are reliant on the Internet for their daily activities. Thus, we believe it should be a government responsibility to ensure that the Internet in general, and BGP in particular, is secured, especially from a robustness and survivability perspective. As a tangible example, governments could provide funding for research and development of BGP security solutions; might encourage ISPs to deploy BGP security solutions (e.g., by subsidies or R&D tax credits or other incentives); or may even require the Internet routing infrastructure used within the government itself to employ a more secure version of BGP. We believe the latter may be particularly effective, because of the very significant spending power of the government, and its leverage over vendors of Internet infrastructure services, associated with the very large IT requirements of an organization of its size.

Acknowledgements

Partial funding for this report was provided by the CSE (Canadian Communications Security Establishment). The first author is also supported by MITACS (Mathematics of Information Technology and Complex Systems) and NCIT (National Capital Institute of Telecommunications). The second author is Canada Research Chair in Network and Software Security, and is also supported by CSE, NCIT, MITACS, an NSERC (Natural Sciences and Engineering Research Council of Canada) Discovery Grant, and the Canada Research Chairs Program. The third author is supported by MITACS and NSERC.

References

[1] C. Adams and S. Lloyd. Understanding Public-Key Infrastructure, 2^{nd} edition. Addison Wesley Professional, 2003.

[2] W. Aiello, J. Ioannidis, and P. McDaniel. Origin Authentication in Interdomain Routing. In *Proc. of the 10th ACM Conference on Computer and Communication Security (CCS'03)*, Washington, D.C., USA. October 2003.

[3] X. Ao. Report on DIMACS Workshop on Large-Scale Internet Attacks. Rutgers University, November, 2003.

[4] R. Atkinson and S. Floyd. IAB Concerns & Recommendations Regarding Internet Research & Evolution. RFC 3869, August 2004.

[5] A. Barbir, S. Murphy, and Y. Yang. Generic Threats to Routing Protocols. Internet Draft, April 13, 2004.

[6] S.M. Bellovin. Security Problems in the TCP/IP Protocol Suite. *ACM Computer Communications Review*, 19(2): 32-48, April 1989.

[7] S.M. Bellovin. Spamming, Phishing, Authentication, and Privacy. *Communications of the ACM*, 47(12), December 2004, Inside Risks.

[8] V.J. Bono. 7007 Explanation and Apology.
http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html

[9] SAFE:Best Practices for Securing Routing Protocols. 2004.
http://www.cisco.com/warp/public/cc/so/neso/vpn/prodlit/sfblp_wp.pdf.

[10] P. Ferguson and D. Senie. Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing. RFC 2827, May 2000.

[11] L. Gao. On Inferring Autonomous System Relationships in the Internet. In *Proceedings of IEEE Global Internet*, November 2000.

[12] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin. Working around BGP: An Incremental Approach to Improving Security and Accuracy in Interdomain Routing. In *Proc. of 2003 Internet Society Symposium on Network and Distributed System Security (NDSS'03)*, San Diego, USA. February 2003.

[13] R. Guida, R. Stahl, T. Bunt, G. Secrest and J. Moorcones. Deploying and Using Public Key Technology: Lessons Learned in Real Life. *IEEE Security and Privacy*, July/August 2004. pp. 67-71.

[14] C. Hedrick. Routing Information Protocol. RFC 1058. June 1988.

[15] A. Heffernan. Protecting of BGP Sessions via the TCP MD5 Signature Option. RFC 2385 (Std Track), August 1998.

[16] Y.C. Hu, A. Perrig, and M. Sirbu. SPV: Secure Path Vector Routing for Securing BGP. In *Proc. of*

SIGCOMM'04, Portland, Oregon, USA. Aug.30 - Sep.3, 2004.

[17] G. Huston. Interconnection, Peering, and Settlements (Part I & II). In *Internet Protocol Journal*, March & June 1999.

[18] http://www.iana.org/.

[19] S. Kent and R. Atkinson. Security Architecture for the Internet Protocol. RFC 2401 (Std Track), November 1998.

[20] S. Kent and R. Atkinson. IP Encapsulating Security Payload (ESP). RFC 2406 (Std Track), November 1998.

[21] S. Kent and C. Lynn, J. Mikkelson, and K. Seo. Secure Border Gateway Protocol (Secure-BGP) - Real World Performance and Deployment Issues. In *Proc. of 2000 Internet Society Symposium on Network and Distributed System Security (NDSS'00)*, San Diego, USA. February 2000.

[22] S. Kent and C. Lynn and K. Seo. Secure Border Gateway Protocol (Secure-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4): 582-592, April 2000.

[23] S. Kent. Secure Border Gateway Protocol: A Status Update. In *Proceedings of the* 7^{th} *IFIP TC-6 TC-11 Conference on Communications and Multimedia Security*, Italy, October 2-3, 2003.

[24] B. Kumar. Integration of Security in Network Routing Protocols. In *ACM SIGSAC Review*, 11(2): 18-25, Spring 1993.

[25] M. Lentczner and M. Wong. Sender Policy Framework: Authorizing Use of Domains in MAIL FROM. Internet Draft (draft-lentczner-spf-00), October 12, 2004.

[26] A.J. Menezes, P.C. van Oorschot, and S. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1996.

[27] D. Meyer. The RouteViews Project. August 2004. http://www.routeviews.org/

[28] S. Murphy. Border Gateway Protocol Security Analysis. IETF Internet Draft, draft-murphy-bgp-vuln-00.txt. November 2001.

[29] S. Murphy. BGP Security Protection. IETF Internet Draft, draft-murphy-bgp-protect-02.txt. Feburary 2002.

[30] D.M. Nicol, S.W. Smith, and M.Y. Zhao. Evaluation of efficient security for BGP route announcements using parallel simulation. *Simulation Pratice and Theory Journal*, special issue on Modeling and Simulation of Distributed Systems and Networks. June 2004.

[31] University of Oregon - Looking Glass. http://antc.uoregon.edu/route-views/

[32] R. Perlman. *Network Layer Protocols with Byzantine Robustness*. PhD thesis, Massachusetts Institute of Technology, August 1988.

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

[34] K. Seo, C. Lynn, and S. Kent. Public-Key Infrastructure for the Secure Border Gateway Protocol (S-BGP). *IEEE DARPA Information Survivability Conference and Exposition II*, 2001.

[35] B.R. Smith and J.J. Garcia-Luna-Aceves. Securing the Border Gateway Routing Protocol. In *Proceedings of Global Internet 1996.* London, UK. November 1996.

[36] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. Katz. Listen and Whisper: Security Mechanisms for BGP. In *Proc. of the First USENIX Symposium on Networked Systems Design and Implementation (NSDI'04)*, San Francisco, CA, USA. March 2004.

[37] T. Wan, E. Kranakis and P.C. van Oorschot. Pretty Secure BGP (psBGP). In *Proc. of the 12th Annual Network and Distributed System Security Symposium (NDSS'05)*, San Diego, USA. February 3-4, 2005.

[38] Slipping in the Window: TCP Reset Attacks.
http://www.osvdb.org/reference/SlippingInTheWindow_v1.0.doc.

[39] R.White, D. McPherson, and S. Sangli. *Practical BGP*. Addison-Wesley. June 2004.

[40] R. White. Securing BGP Through Secure Origin BGP (soBGP). In *The Internet Protocol Journal*, 6(3): 15-22, September 2003.

[41] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S.F. Wu, and L. Zhang. An Analysis of BGP Multiple Origin AS (MOAS) Conflict. In *ACM SIGCOMM Internet Measurement Workshop*, San Francisco, USA. Nov. 2001.

[42] P. Zimmermann. *The Official PGP User's Guide* (second printing). Cambridge, MA: MIT Press, 1995.