

On Interdomain Routing Security and Pretty Secure BGP (psBGP)

P. C. VAN OORSCHOT, TAO WAN, and EVANGELOS KRANAKIS
Carleton University

It is well known that the Border Gateway Protocol (BGP), the IETF standard interdomain routing protocol, is vulnerable to a variety of attacks, and that a single misconfigured or malicious BGP speaker could result in large-scale service disruption. In this paper, we present *Pretty Secure BGP (psBGP)*—a proposal for securing BGP, including an architectural overview, design details for significant aspects, and preliminary security and operational analysis. psBGP differs from other security proposals (e.g., S-BGP and soBGP) in that it makes use of a single-level PKI for AS number authentication, a decentralized trust model for verifying the propriety of IP prefix origin, and a rating-based stepwise approach for AS_PATH (integrity) verification. psBGP trades off the strong security guarantees of S-BGP for presumed-simpler operation, e.g., using a PKI with a simple structure, with a small number of certificate types, and of manageable size. psBGP is designed to successfully defend against various (nonmalicious and malicious) threats from uncoordinated BGP speakers, and to be incrementally deployed with incremental benefits.

Categories and Subject Descriptors: C.2.6 [**Computer-Communication Networks**]: Internet-working—*Security*

General Terms: Security, Reliability, Standardization

Additional Key Words and Phrases: BGP, trust, interdomain routing, secure routing protocols, authentication, certificates, public-key infrastructure

ACM Reference Format:

van Oorschot, P. C., Wan, T., and Kranakis, E. 2007. On interdomain routing security and pretty secure BGP (psBGP). *ACM Trans. Info. Syst. Secur.* 10, 3, Article 11 (July 2007), 41 pages. DOI = 10.1145/1266977.1266980 <http://doi.acm.org/10.1145/1266977.1266980>

1. INTRODUCTION AND MOTIVATION

The Internet routing infrastructure consists of a number of autonomous systems (ASes), each of which consists of a number of routers under a single

Preliminary versions of this work appeared as Wan et al. [2005] and Wan [2006]. The second author is currently with Nortel, Canada.

Authors' address: P. C. van Oorschot, Tao Wan, and Evangelos Kranakis, Carleton University, Ottawa, Canada; email: {paulv, twan, kranakis}@scs.carleton.ca.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org. © 2007 ACM 1094-9224/2007/08-ART11 \$5.00 DOI 10.1145/1266977.1266980 <http://doi.acm.org/10.1145/1266977.1266980>

technical administration (e.g., sharing the same routing policy). The Border Gateway Protocol (BGP) [Rekhter and Li 1995] is the IETF standard interdomain routing protocol for exchanging reachability information between ASes on the Internet. Each network layer destination is identified by an IP prefix representing a range of IP addresses. An AS announces its IP prefixes via BGP to its direct neighbors, which may further propagate the prefix announcement to their neighbors. A remote AS receiving such announcement may build routes for forwarding traffic destined to the addresses within the address range specified by the announced prefixes.

One critical question with BGP is the following: which AS has a right to announce a given IP prefix? The current version of BGP does not have any mechanism to verify the propriety of IP prefix origin, i.e., if the originating AS indeed holds a prefix (allocated) or is authorized by the actual holder of the prefix (delegated). This opens a serious security hole, which allows one AS to announce IP prefixes allocated or delegated (hereafter *assigned*) to any other ASes. This is commonly referred to as *prefix hijacking*. Examples of consequences include denial of service (i.e., legitimate user traffic cannot get to its ultimate destination) and man-in-the-middle attacks (i.e., legitimate user traffic is forwarded through a router under the control of an adversary). Warnings about attacks exploiting routing vulnerabilities were given circa 1988 by Perlman [1988], and by Bellovin [1989]; and such attacks have recently reportedly been carried out by spammers [Bellovin 2004].

Many proposals [Kent et al. 2000; Goodell et al. 2003; White 2003; Aiello et al. 2003] have been made for improving BGP security and, in particular, for verifying if an AS has the right to announce a given IP prefix. There are two main approaches: (1) building centralized routing registries storing information about address space assignments, e.g., the Internet Routing Registry (IRR) [IRR 2005], to facilitate the containment of fraudulent route announcements, e.g., by filtering; and (2) building a strict hierarchical public key infrastructure (PKI) in parallel to the existing IP address assignment structure (e.g., S-BGP [Seo et al. 2001; Lynn et al. 2003]). While these two approaches may differ in many ways, e.g., protecting a database itself versus protecting individual objects in the database, they both typically require a large-scale PKI to provide strong security or to meet some operational requirements (e.g., multihoming).

IRR needs to perform identity authentication to verify if an entity requesting to make changes to the routing database is authorized to do so. Currently in IRR, PGP [Zimmermann 1995] is used for public key authentication. However, this authentication is done using a sender's email address when an object is first created and, thus, is vulnerable to email spoofing [Zsako 1999]. As a result, a global PKI or something equivalent, appears to be required to provide stronger guarantees. More seriously, there are no controls in place to ensure that the information asserted by a user is accurate, even though the user can be authenticated. S-BGP makes use of a hierarchical tree structure for address assignment, rooted at Regional Internet Registries (RIRs). For each consecutive pair of nodes on the address assignment chain, the first node (an organization) on the chain assigns a subset of its own address space to the second. While an organization obtaining its address space from its Internet Service Providers

(ISPs) may not need to appear on an address delegation chain (i.e., need not be issued relevant certificates), it will need a certificate to do multihoming (i.e., connecting to two independent ISPs) or to connect to another ISP different from the one it obtained the address space from. Since these common operational practices must be supported [Villamizar et al. 1999], it implies that many organizations not running BGP may also need to be involved in the S-BGP PKI, resulting in the challenging requirement of a large-scale (essentially global) PKI. In addition, it appears difficult to build a centralized PKI for verifying IP address assignment given the complexity, if not impossibility, of tracing how the existing IP address space is assigned, and tracing all changes of IP address assignments. This is in part because of the large number of prefixes in use and the large number of organizations involved. In particular, many IP addresses were given out before the existing hierarchical address allocation structures were in place. Therefore, it might not be possible to construct address assignment chains for them [DHS 2005]. Fundamentally, all these approaches assume trusted authoritative sources of all prefix assignments. We suggest that such an assumption might not be realistic, or at least it requires a large-scale infrastructure to support, which appears difficult to realize.

1.1 Contributions

In this paper, we present a new BGP security proposal—pretty secure BGP (psBGP), based on our preliminary overviews [Wan et al. 2005; Wan 2006]. psBGP includes defenses against falsification of BGP UPDATE messages, and a new approach for verifying the propriety of prefix origin by cross-checking information from multiple, ideally independent, sources. Specific psBGP security goals are outlined in Section 2.3. psBGP is based on the following assumptions: (1) trusted authorities of prefix assignments on the Internet may not always be available; (2) some entities may have partial knowledge of prefix assignments; and (3) corroboration of information from different sources can increase confidence in the assessment of that information. In particular, RIRs are the trusted authority of initial prefix allocations, and some ASes might have partial knowledge of prefix assignments of their direct neighbors. We note that while psBGP makes use of corroboration for increasing confidence in prefix assertions, it does not prevent the use of a centralized PKI for prefix delegations. If such an infrastructure (e.g., [Kent 2006]) or part of it does exist, it can also be used by psBGP in constructing AS prefix graphs (see Section 4.1), in which case the corroboration approach by psBGP can be used for authenticating prefixes not accommodated by the centralized PKI, e.g., legacy address space.

1.2 psBGP Highlights

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 several trusted certificate authorities (i.e., 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. Authentication

is usually the first step toward authorization. Without such a guarantee, an attacker may be able to impersonate another AS and, thus, be able to announce prefixes assigned to the impersonated AS.

2. psBGP makes use of a rating mechanism for flexibility in balancing security and practicality in prefix origin and AS_PATH verification.
3. psBGP makes use of a *decentralized trust model* for verifying the propriety of IP prefix assignment. Each AS periodically issues a digitally signed *Prefix Assertion List (PAL)* consisting of a number of bindings of an AS number and (zero or more) IP prefixes, one such binding for itself and one for each of its neighbors. An assertion made by an AS s_i regarding its own prefixes (*prefix assertion*) lists all prefixes assigned to s_i . An assertion made by s_i for a neighboring AS s_j (*prefix endorsement*) may list all or a subset of the prefixes assigned to s_j . An *AS prefix graph* (see Section 4.3) is built independently by each AS s_i , based on the *PALs*, which s_i has received from other ASes, and s_i 's ratings of those ASes. An AS prefix graph is then used for evaluating the trustworthiness and preference of a prefix origin by an AS, in conjunction with its local configurable parameters (e.g., its trust in those ASes involved in a prefix assertion and trust thresholds). In this way, the difficult task of tracing IP address assignments is distributed across ASes on the Internet.
4. psBGP modifies the S-BGP digital signature approach with a rating mechanism and a stepwise approach for verifying AS_PATH integrity. Each AS computes a weight for an AS_PATH, based on ratings of the ASes digitally signing the path, and determines whether or not to accept the path based on local parameters. This approach allows an upgrading path to countering increased threats, as recommended in Bellovin et al. [2005].

Our design is inspired by the referral model widely used in social society for increasing confidence in the truth of a piece of information when an authoritative source of truth regarding that information is not available.¹ For example, a job applicant is usually required to provide reference letters to allow cross-checking the applicant statements on his quality and background. A reference letter should be from an individual who has closely worked with the applicant, e.g., a former supervisor. Similarly in psBGP, each AS should obtain endorsement for its prefix assertions from some ASes, which are likely to have, or likely to be reliable sources for, knowledge of its prefix assignment, e.g., a direct neighbor with which it has a business relationship. An AS choosing to endorse a prefix assertion made by a neighboring AS should carry out some form of due diligence (or other means to increase accountability) to increase confidence in the correctness of that assertion, i.e., to increase its own confidence that the asserted prefix is, indeed, assigned to the asserting AS. The security assurances of this aspect of psBGP are directly related to the quality of such due diligence, which will impose extra work on BGP operators; this is the price to pay for increased security.

¹In this sense (and regarding stepwise integrity, see Section 3.5), there is some similarity to IRV [Goodell et al. 2003].

Table I. Notation

\mathbb{S}, s_i	\mathbb{S} is the set of all AS numbers; currently $\mathbb{S} = \{1, \dots, 2^{16}\}$. $s_i \in \mathbb{S}$ is an AS number
\mathbb{P}, f_i	\mathbb{P} is the set of all IP addresses. $f_i \subseteq \mathbb{P}$ is an IP prefix specifying a range of IP addresses $f_i = f_j \cup f_k$, if the IP addresses specified by f_i equal those by f_j and f_k combined
T	an authority with respect to \mathbb{S} and \mathbb{P} , e.g., $T \in \{x x \text{ is an RIR}\}$
p_k	$p_k = [s_1, s_2, \dots, s_k]$ is an AS_PATH; s_1 is the first AS inserted onto p_k
m	$m = (f_1, p_k)$ is a BGP route (a selected part of a BGP UPDATE message)
$N(s_i)$	s_i 's neighbors, i.e., the set of ASes with which s_i establishes a BGP session on a regular basis. A given AS s_i may have many BGP speakers, each of which may establish BGP sessions with speakers from many other ASes. $N(s_i)$ is the set of all other such ASes
$k_A, \overline{k_A}$	A's public and private keys, respectively
$\{m\}_A$	digital signature on message m generated with A's private key $\overline{k_A}$
$(k_A, A)_{\overline{k_B}}$	a public key certificate binding k_A to A, signed using $\overline{k_B}$, verifiable using k_B
$(f_i, s_i)_A$	an assertion made by A that f_i is assigned to s_i

In addition to the benefits derived from being incrementally deployable, psBGP is *lightweight*—it uses a PKI which has a simple structure, a small number of certificate types, and is of manageable size, while remaining *effective*—it is designed to successfully defend against selected threats from uncoordinated, misconfigured or malicious BGP speakers.

1.3 Organization

The rest of the paper is organized as follows. Section 2 defines notation, overviews BGP, discusses BGP threats, and summarizes BGP security goals. psBGP is presented in Sections 3 and 4. Security and operational analysis of psBGP is given in Sections 5 and 6, respectively. A brief review of related work is given in Section 7. We conclude in Section 8.

2. BACKGROUND: BGP SECURITY THREATS AND GOALS

After defining notation, we give a brief overview of relevant aspects of BGP, discuss BGP security threats, and summarize five security goals for BGP, for later use in the paper.

2.1 Notation

A and B denote entities (e.g., an AS or a BGP speaker). X or Y denotes an assertion, which is any statement. An assertion may be *proper* or *improper*. We avoid use of the term *true* or *false*, since, in BGP, it is not always clear that a statement is 100% factual or not. An assertion is proper if it conforms to the rules (e.g., psBGP rules) governing the related entity making that assertion. Table I defines some of the notation used in this paper.

2.2 Selective Overview of BGP

Conceptually, a routing network can be abstracted as a graph, where a vertex is a router and an edge is a network link. If a network consists of a small (e.g., several) or medium (e.g., tens or hundreds) number of routers and they are under a common administrative domain, a single routing protocol can be used for exchanging and maintaining routing information in that network. Since there are a large number of routers (e.g., exceeding hundreds of thousands)

on the Internet and they are administrated by many different organizations, a hierarchical routing approach has been chosen for better organizational and administrative control and error containment, as well as scalability. Internet routing protocols can be classified as *intra-domain* (used within an AS) or *inter-domain* (used between ASes).

BGP is an interdomain routing protocol based on a *distance vector* approach. A BGP speaker establishes a session over TCP with each of its direct neighbors, exchanges routes with them, and builds routing tables based on the routing information received from them. Unlike a simple distance vector routing protocol (e.g., RIP [Hedrick 1988]) where a route has a simple metric (e.g., number of hops), a BGP route is associated with a number of attributes and routes are selected based on local routing policy. One notable route attribute is AS_PATH, which consists of the sequence of ASes traversed by the route that is being propagated. BGP is often considered a *path vector* routing protocol.

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 redundancy); and a *transit-AS* has more than one connection to other ASes and is designed to carry traffic for others.

While a stub-AS may have only one BGP speaker, a multihomed or a transit-AS often has more. A BGP session between two BGP speakers located within two different ASes is often referred to as external-BGP (eBGP), and a BGP session between two BGP speakers within a common AS is often referred to as internal-BGP (iBGP). An eBGP speaker actively exchanges routing information with an external neighbor by importing and exporting BGP routes. An iBGP speaker only helps propagate routing updates to other BGP speakers within a common AS; it does not make any changes to a routing update.

A BGP session between two different ASes usually implies one of the following four types of business relationship [Gao 2000]: *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. Two peer ASes usually find 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.

2.3 Attacks on BGP

BGP faces attacks from both BGP speakers and BGP sessions. A misbehaving BGP speaker may be misconfigured (mistakenly or intentionally), compromised (e.g., by exploiting software flaws), or unauthorized (e.g., by exploiting a BGP peer authentication vulnerability). A BGP session may be compromised or unauthorized. We focus on attacks against BGP control messages without considering those against data traffic (e.g., malicious packet dropping [Just et al. 2003]). Attacks against BGP control messages include, for example, modification, insertion, deletion, exposure, and replaying of messages. In this paper, we focus on modification and insertion (hereafter *falsification* [Barbir et al. 2004]) of BGP control messages; deletion, exposure, and replaying can be addressed by a point-to-point authentication protocol, e.g., IPsec [Kent and Atkinson 1998a].

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 peers; falsification of them will very likely result in session disruption. As mentioned by Hu et al. [2004], they can be protected by IPsec [Kent and Atkinson 1998a]. In psBGP, we concentrate on falsification of BGP UPDATE messages (and hereafter, refrain from capitalizing UPDATE), which carry interdomain 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). As commonly agreed [Hu et al. 2004], 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. Further discussion is beyond the scope of the present paper.

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 neither held by that AS nor authorized by the holder of that prefix or aggregated improperly from other routes. Examples of consequences include denial-of-service and man-in-the-middle attacks. 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 currently being propagated. An AS_PATH of type AS_SET consists of an unordered list of ASes, sometimes created when multiple routes are aggregated. 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. 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. Some other path attributes (e.g., community, Multi_Exit_Disc [Rekhter and Li 1995]) may also need protection, but many of these are usually only used between two neighbors and not globally transitive. Thus, damage resulting from attacking them is relatively contained. In psBGP, we focus on countering falsification of NLRI and AS_PATH, which can result in large-scale service disruption.

We assume there are multiple noncolluding misbehaving ASes (but see Section 4.2) in the network, which may have their own legitimate cryptographic keying materials.

2.4 BGP Security Goals

We seek to design secure protocol extensions to BGP, which can resist the threats as discussed above, i.e., primarily falsification of BGP update messages. As with most other secure communication protocols, BGP security goals must include data origin authentication and data integrity. In addition, verification of the propriety of BGP messages is required to resist falsification attacks. Specifically, the propriety of NLRI and AS_PATH should be verified. Most likely, all verification will be performed by a BGP speaker online, but possibly by an operator off-line, which is not discussed in the present paper.

We summarize five security goals for BGP (cf. [Kent et al. 2000], also see [Wan et al. 2005; Wan 2006]), for reference later in Sections 3, 4, 5.1, and 7. G1 and G2 relate to data origin authentication and are separated for the sake of clarity, G3 to data integrity, and G4 and G5 to the propriety of BGP control messages. These five security goals address a large number of serious attacks against BGP. Thus, it is highly desirable for any serious BGP security proposal to achieve them. However, these alone should not be considered as sufficient for BGP security, since other attacks (e.g., unauthorized route withdrawal) remain (see Section 2.3).

- G1. (*AS Number Authorization*) 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 Authorization*) 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 control message has not been illegally modified in a point-to-point BGP session.
- G4. (*AS Path Verification*) It must be verifiable that an AS_PATH ($p_k = [s_1, s_2, \dots, s_k]$) of a BGP route m being propagated consists of a sequence of ASes traversed by m in the specified order, i.e., m originated from s_1 and has traversed s_2, \dots, s_k in order.
- G5. (*Prefix Origin Authentication*) It must be verifiable that it is proper for an AS to originate an IP prefix. It is *proper* for AS s_1 to originate prefix f_1 if (1) f_1 is indeed held by s_1 (prefix allocation); (2) s_1 is authorized by the holder of f_1 (prefix delegation); or (3) s_1 is assigned (allocated or delegated) a set F_1 of prefixes; s_1 has received a set of routes with a set F_2 of prefixes; and f_1 is aggregated from F_1, F_2 , or both, such that $\forall f_x \subseteq f_1, f_x \subseteq F_1 \cup F_2$ (prefix aggregation).²

3. PRETTY SECURE BGP (PSBGP)

psBGP makes use of a centralized trust model for authorizing (and verifying the authorized use of) AS numbers and authenticating AS public keys. Five RIRs are the root-trusted certification authorities (CAs) and can cross-sign each other's public key certificates. Another option would be to have a single CA rooted at IANA. However, for political and availability reasons [Seo et al. 2001], we recommend multiple-rooted CA's (with cross-certification) over a single one. In psBGP, each AS s is issued an intermediate CA public key certificate (ASNumCert), signed by one of the RIRs (say T), denoted by $(k_s, s)_{\overline{k_T}}$. Such an AS creates and signs two end-entity certificates, SpeakerCert and a Session-Cert binding two different public keys to s , respectively, and a *prefix assertion list* (*PAL*). The latter, pal_s , is an ordered list: the first assertion is for s itself and the rest are endorsements by s for each of s 's neighbors ordered by AS

²If f_1 is not assigned to s_1 and $\exists f_x \subseteq f_1$ such that $f_x \not\subseteq F_1 \cup F_2$, then s_1 *overclaims* IP prefixes, which is a type of falsification.

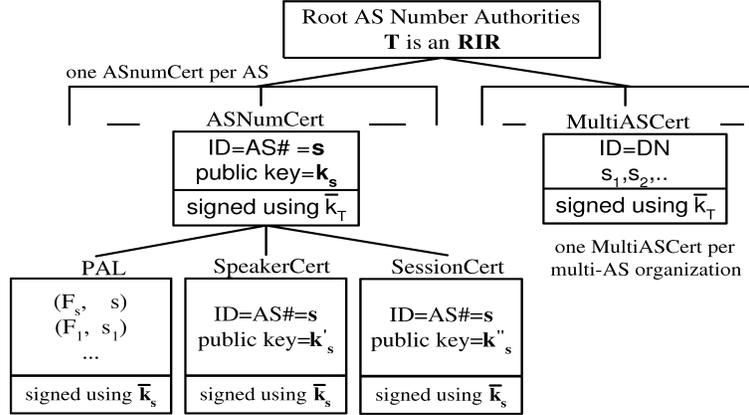


Fig. 1. psBGP certificate structure.

number. Figure 1 illustrates the certificate structure used in psBGP. In what follows, we start with a description of a rating mechanism used by an AS in determining its confidence in an AS_PATH or a prefix assertion. We next describe psBGP with respect to the above five security goals: G1–G4 here and G5 in Section 4.

3.1 A Rating Mechanism

In psBGP, each AS s_i rates every other AS s_j with a value in $[0, 1]$, denoted by $r_i(s_j)$, representing s_i 's confidence or belief in s_j 's trustworthiness, i.e., in an assertion made by s_j , such as a digitally signed AS_PATH or a prefix assertion or endorsement. $r_i(s_j) = 0$ or 1 , respectively, indicates s_i fully distrusts or trusts s_j . When there is no ambiguity, we omit the subscript on r in $r_i(s_j)$.

While each AS has freedom in determining how to rate other ASes, we suggest the following guidelines: an RIR should be fully trusted (i.e., rated 1); a direct neighbor might be expected, in many cases, to be more trustworthy than a remote AS; and a majority of ASes should be neutrally trusted, e.g., rated 0.5. We next present a method [Wan et al. 2004] for computing the confidence value in a statement, which is consistent among a set of assertions made by a group of ASes (a *corroborating group*), based on one's ratings of those ASes. We consider two types of consistency in psBGP: *path consistency* and *prefix consistency*. The former is regarding the consistency among a set of digital signatures over an AS_PATH (see Definitions 1 and 2 in Section 3.5). The latter is regarding the consistency of a prefix assertion and a prefix endorsement (see Definition 4 in Section 4.1).

Let s_1, \dots, s_n be a group of ASes, which independently produce a set of consistent assertions a_{s_1}, \dots, a_{s_n} . Let $\lambda_{s_1, \dots, s_n}$, abbreviated by $\lambda_{[1..n]}$, denote a common subset that can be derived from each of the above n consistent assertions. The precise meaning of $\lambda_{[1..n]}$ depends on the type of consistency in question. In prefix consistency, if a_{s_1} is a prefix assertion $(f_1, s_1)_{s_1}$ and a_{s_2}, \dots, a_{s_n} are prefix endorsements $(f_1, s_1)_{s_2}, \dots, (f_1, s_1)_{s_n}$, then $\lambda_{[1..n]}$ represents a prefix assignment of s_1 , i.e., s_1 is assigned a prefix f_1 . In path consistency, if $a_{s_1} = \{f_1, [s_1,$

$s_2]]_{s_1}, \dots, a_{s_n} = \{f_1, [s_1, \dots, s_n, s_{n+1}]]_{s_n}$ are digital signatures present with a BGP route $m = (f_1, p_n = [s_1, \dots, s_n])$, then λ_{s_1, s_2} represents a statement that p_n contains a path segment $[s_1, s_2]$, λ_{s_2, s_3} represents a statement that p_n contains a path segment $[s_2, s_3]$, and so on. We next show how an AS s_i computes a confidence value or a belief in $\lambda_{[1..n]}$, denoted $b(\lambda_{[1..n]})$, based on s_i 's ratings of s_1, \dots, s_n in the corroborating group. By definition, s_i 's rating of s_j , $1 \leq j \leq n$, represents s_i 's confidence in the assertion a_j made by s_j or any subset λ_{s_i} derived from a_j , i.e., $b(\lambda_{s_j}) = b(a_{s_j}) \triangleq r(s_j)$. $b(\lambda_{[1..n]})$ is defined as:

$$b(\lambda_{[1..n]}) = \begin{cases} r(s_1) & \text{if } n = 1 \\ r(s_2) + [1 - r(s_2)] \cdot r(s_1) & \text{if } n = 2 \\ r(s_n) + [1 - r(s_n)] \cdot b(\lambda_{[1..(n-1)]}) & \text{if } n \geq 3 \end{cases} \quad (1)$$

Consistent with Dempster–Shafer theory [Dempster 1967; Shafer 1976] of belief reasoning, properties of Eq. (1) include: (i) endorsement from a fully distrusted AS (i.e., rated 0) does not increase one's confidence; (ii) endorsement from a fully trusted AS (i.e., rated 1) increases one's confidence to maximum (i.e., 1); and (iii) if no AS in the corroborating group is fully distrusted or trusted (i.e., the rating is $0 < r < 1$), one's confidence increases, but never reaches maximum.

For later cross-reference, Algorithm 1 describes how to increase one's confidence in $\lambda_{[1..(n-1)]}$ when an additional endorsement is obtained, e.g., from s_n . Algorithm 2 describes how to reduce one's confidence in $\lambda_{[1..n]}$ when (without loss of generality) s_n 's endorsement is withdrawn.

Algorithm 1 Adding new endorsement from AS s_n

1: **INPUT:** $b(\lambda_{[1..(n-1)]}), r(s_n)$
2: **OUTPUT:** $b(\lambda_{[1..n]})$
3: $t \leftarrow r(s_n) + [1 - r(s_n)] \cdot b(\lambda_{[1..(n-1)]})$
4: **return**(t)

Algorithm 2 Removing existing endorsement from AS s_n

1: **INPUT:** $b(\lambda_{[1..n]}), r(s_n)$
2: **OUTPUT:** $b(\lambda_{[1..(n-1)]})$
3: $t \leftarrow \frac{b(\lambda_{[1..n]}) - r(s_n)}{1 - r(s_n)}$
4: **return**(t)

3.2 AS Number Authorization in psBGP (G1)

Following S-BGP [Seo et al. 2001], psBGP makes use of a centralized PKI for AS number authorization, with five root certificate authorities (CAs), corresponding to the five existing RIRs. When an organization B applies for an AS number, besides supplying documents currently required, B additionally supplies a public key and should be required to prove possession of the corresponding private key [Seo et al. 2001; Adams and Lloyd 2003]. When an AS number is granted

Table II. AS Number Dynamics from January 1 to August 1, 2004

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug
Start of month	16,554	16,708	16,879	17,156	17,350	17,538	17,699	17,884
Removed during month	153	137	155	174	138	179	164	N/A
Added during month	307	308	432	368	326	342	349	N/A

to B by an RIR or by its subordinate registries, an intermediate CA public key certificate (ASNumCert) is also issued, signed by the issuing RIR, binding the public key supplied by B to the granted AS number. An AS number s is called *certified* if there is a valid ASNumCert $(k_s, s)_{\overline{k_T}}$, binding s to a public key k_s signed by one of the RIRs, T .

The proposed PKI for authorizing AS numbers is practical for the following reasons. (1) The roots of the proposed PKI are the existing trusted authorities of the AS number space, removing a major trust issue, which is one of the most difficult parts of a PKI: the root of a PKI must have control over the name space involved in that PKI. Thus, RIRs are the natural and logical AS number certificate authorities. We acknowledge that nontrivial (but feasible) effort might be required for implementing such a PKI. (2) The number of ASes on the Internet and its growth rate are relatively manageable (see Table II), based on the RouteViews dataset [RouteViews 2005]. Considering there are five RIRs, the overhead of managing ASNumCerts should certainly be manageable, given that larger PKIs are currently commercially operational [Guida et al. 2004].

To verify the authenticity of an ASNumCert, an AS must have the trusted public key (or verifiable certificate) of the signing RIR. These few root-trusted public key certificates can be distributed using *out-of-band* mechanisms to all ASes. ASNumCerts can be distributed with BGP update messages. An ASNumCert should be revoked when the corresponding AS number is no longer used or is reassigned to another organization. ASNumCerts can be revoked through any standard means, e.g., a certificate revocation list (CRL) [Housley et al. 1999] (cf. [Ma et al. 2006]), which can be distributed using out-of-band mechanisms, e.g., a repository. To summarize, we assume that every AS has the public key certificates of RIRs and can obtain the ASNumCerts of any other ASes if and when necessary.

In discussion related to various proposals for securing BGP, there is much debate in the BGP community on the architecture for authenticating the public keys of ASes, particularly on the pros and cons of using a strict hierarchical trust model versus a distributed trust model, e.g., a web-of-trust model [Zimmermann 1995]. We make use of a strict hierarchical trust model (with depth of one) for authorizing AS numbers and authenticating their public keys to provide a strong guarantee of security. Therefore, it would appear to be difficult for an attacker to spoof an AS in psBGP as long as it cannot obtain the private key corresponding to the public key of an ASNumCert signed by an RIR, or the signing key of an RIR. In contrast, a web-of-trust model does not provide such a guarantee. Other issues that arise with a web-of-trust model include:

trust bootstrapping, trust transitivity, and vulnerability to a single misbehaving party [Maurer 1996; Reiter and Stubblebine 1997].

3.3 BGP Speaker Authorization in psBGP (G2)

An AS may have one or more BGP speakers. A BGP speaker must be authorized by an AS to represent that AS to establish a BGP session with a BGP speaker in another AS. In psBGP, an AS with a certified ASNumCert issues an operational end-entity public key certificate shared by all BGP speakers within the AS, namely, SpeakerCert. A SpeakerCert is signed using the private key of the issuing AS, corresponding to the public key in the AS's ASNumCert (see Figure 1). A SpeakerCert is an assertion made by an AS that a BGP speaker with the corresponding private key is authorized to represent that AS. SpeakerCerts can be distributed with BGP update messages.

We consider three design choices for BGP speaker authentication: (a) each BGP speaker has a distinct key pair and is issued a unique public key certificate; (b) group signatures (e.g., see Boneh et al. [2004]) are used, i.e., each BGP speaker has a unique private key but shares a common public key and public key certificate with other speakers in the same AS; or (c) all BGP speakers in a given AS share a common public-private key pair. We propose the latter primarily for its operational simplicity. Choice (a) provides stronger security, in theory, but requires more certificates and discloses BGP speaker identities, which may introduce competitive security concerns [White et al. 2004]. Choice (b) again provides stronger security in theory, requires the same number of certificates, and does not disclose BGP speaker identities, but involves a more complex system, which we believe significantly reduces its chances of being commercially accepted and securely deployed.

The private keys corresponding to the public keys of a SpeakerCert and SessionCert are, respectively, used for signing BGP update messages and establishing secure connections with neighbors (see Section 3.4). Therefore, they would most likely be stored in the communication device associated with a BGP speaker. In contrast, since the private key corresponding to the public key of an ASNumCert is only used for signing a SpeakerCert, a SessionCert, and a *PAL*, it need not be stored in a BGP speaker. Thus, compromising a BGP speaker, at most, discloses the private keys of a SpeakerCert and a SessionCert, requiring revocation and reissuing of them, without impact on an ASNumCert. This separation of ASNumCerts from SpeakerCerts and SessionCerts provides a more conservative design (from a security viewpoint) and distributes from RIRs to ASes (or their delegated certificate service providers) the workload of certificate revocation and reissuing resulting from BGP speaker compromises. While ASNumberCerts and SpeakerCerts need to be distributed to every other ASes, e.g., via BGP update messages, a SessionCert need only be distributed to direct neighbors, e.g., via IKEv2 [Kaufman 2005]. In summary, an ASNumCert must be revoked if the corresponding AS number is reassigned or the corresponding private key is compromised; a SpeakerCert or SessionCert must be revoked if a BGP speaker in that AS is compromised, or for other reasons (e.g., if the private key is lost).

3.4 Data Integrity in psBGP (G3)

To protect data integrity, BGP sessions between neighboring ASes must be protected. Following S-BGP and soBGP, psBGP uses IPsec Encapsulating Security Payload (ESP) [Kent 2005] with null encryption for protecting BGP sessions. Since many existing BGP speakers implement TCP MD5 [Heffernan 1998] with manual key configurations for protecting BGP sessions, it must be supported by psBGP as well. In psBGP, automatic key management techniques, e.g., IKEv2 [Kaufman 2005], can be implemented to improve the security of TCP MD5 as each BGP speaker has a unique public–private key pair for BGP session security.

3.5 AS_PATH Verification in psBGP (G4)

Regarding “AS_PATH security,” different security solutions of BGP define it differently. In S-BGP, the security of an AS_PATH is interpreted as follows: for every pair of ASes on the path, the first AS authorizes the second to further advertise the prefix associated with this path. In soBGP [White 2003], it is defined as the plausibility of an AS_PATH, i.e., if an AS_PATH factually exists on the AS graph (whether or not that path was actually traversed by an update message in question is irrelevant).

Since AS_PATH is used by the BGP route selection process, greater assurance of the integrity of an AS_PATH increases the probability that routes are selected based on proper information. Without strong guarantees of AS_PATH integrity, an attacker may be able to modify an AS_PATH in a such way that it is still plausible in the AS graph and is also more favored (e.g., with a shorter length) by recipient ASes than the original path. In this way, a recipient AS may be misled to favor a falsified route over correct routes, possibly influencing traffic flow. Thus, in our view, it is not sufficient to verify only the existence/nonexistence of an AS_PATH if greater assurance of the integrity of an AS_PATH can be provided at acceptable cost.

We choose the S-BGP approach combined with the rating mechanism described in Section 3.1 to determine dynamically (at run-time) the number of digital signatures on an AS_PATH to be verified. We first give the definition of *path consistency*, then present how to calculate a confidence value in an AS_PATH.

Definition 1 (Path Consistency). Let $m = (f_1, p_k = [s_1, \dots, s_k])$ be a BGP route, and $sig_i = \{f_1, p_i\}_{s_i}$ be a digital signature generated by a psBGP-enabled BGP speaker in s_i , $1 \leq i \leq k$, where $\{p_i\}_{s_i} = [s'_1, \dots, s'_{i+1}]$ is the path signed by s_i . $\{p_i\}_{s_i}$ is consistent with p_k if $\{p_i\}_{s_i}$ consists of the first $i + 1$ ASes on p_k (i.e., $s'_1 = s_1, \dots, s'_{i+1} = s_{i+1}$) when $1 \leq i \leq k - 1$, or consists of p_k appended by another AS s_{k+1} when $i = k$.

Definition 2 (Signed-Path Consistency). Let $m = (f_1, p_k = [s_1, \dots, s_k])$ be a BGP route, and $sig_i = \{f_1, p_i\}_{s_i}$, $sig_j = \{f_1, p_j\}_{s_j}$ the digital signatures generated by two psBGP-enabled ASes s_i and s_j , $1 \leq i, j \leq k$, on p_k . $\{p_i\}_{s_i}$ and $\{p_j\}_{s_j}$ are consistent if they both are consistent with p_k .

Two consistent-signed paths $\{p_i\}_{s_i}$ and $\{p_j\}_j$ contain common subset λ_{s_i, s_j} . For example, if $\{p_2\}_{s_2} = [s_1, s_2, s_3]$, $\{p_4\}_{s_4} = [s_1, s_2, s_3, s_4, s_5]$, λ_{s_2, s_4} could be an assertion that p_k contains the path segment $[s_2, s_3]$, since both s_2 and s_4 assert it in their signed path. As a result, one may expect the belief in λ_{s_2, s_4} will increase, which may further contribute to the belief in p_k in some way. However, the definition of path confidence in psBGP is more restrictive. In psBGP, the belief in p_k , $b(p_k)$, is defined as the sum of the belief of each assertion that p_k contains a two-AS path segment $[i, i + 1]$, $1 \leq i \leq k - 1$, divided by the total number of those path segments $k - 1$.

Definition 3 (Path Confidence). Let $m = (f_1, p_k = [s_1, \dots, s_k])$ be a BGP route and $\lambda_{s_i, s_{i+1}}$ be the assertion that p_k contains a two-AS path segment $[s_i, s_{i+1}]$. The belief in p_k is defined as: $b(p_k) = \frac{1}{k-1} \sum_{i=1}^{i=k-1} b(\lambda_{s_i, s_{i+1}})$.

The belief in the assertion $\lambda_{s_i, s_{i+1}}$ that p_k contains a two-AS path segment $[s_i, s_{i+1}]$ is obtained exclusively from the signed paths by s_i and s_{i+1} (i.e., $\{p_i\}_{s_i}$, $\{p_{i+1}\}_{s_{i+1}}$ since two ASes have authority over the path segment between themselves. The signed path by another AS, e.g., s_{i+2} , may also contain $[s_i, s_{i+1}]$, but it does not contribute to the belief in $\lambda_{s_i, s_{i+1}}$, since s_{i+2} apparently does not have authority over $[s_i, s_{i+1}]$ and its signed path may be dependent on the path signed by s_i or s_{i+1} .

If one AS on $[s_i, s_{i+1}]$ is non-psBGP-enabled and does not digitally sign its path, the belief in $\lambda_{s_i, s_{i+1}}$ is then solely derived from the signed path of the other AS. If neither of them has signed the path, i.e., $\{p_i\}_{s_i}$ and $\{p_{i+1}\}_{s_{i+1}}$ are null, there is no evidence to believe $\lambda_{s_i, s_{i+1}}$. In this case, $b(\lambda_{s_i, s_{i+1}})$ is set to 0.

In psBGP, a minimum of two digital signatures must be verified if two or more are present on an AS_PATH p_k . The exact number of digital signatures to be verified depends on a verifying AS s_{k+1} 's ratings of the ASes, which have signed p_k , and a local configurable confidence threshold $\theta_{k+1} \geq 0$. Verification of p_k starts from the digital signature generated by the last AS s_k on p_k , and moves toward the first AS s_1 . Upon a digital signature sig_i verifying successfully, i.e., sig_i is valid and $\{p_i\}_{s_i}$ is consistent with p_k , the belief in the assertion $\lambda_{s_i, s_{i+1}}$ ($1 \leq i \leq k-1$) that p_k contains $[s_i, s_{i+1}]$ is recomputed (using Algorithm 1) and the current belief in p_k is updated (see Definition 3). If $b(p_k)$ is no less than s_{k+1} 's confidence threshold θ_{k+1} , i.e., $b(p_k) \geq \theta_{k+1}$, then p_k is accepted. Otherwise, more digital signatures are verified (see Algorithm 3) until:

1. one digital signature verification fails, in which case p_k is rejected; or
2. $b(p_k) \geq \theta_{k+1}$, in which case p_k is accepted; or
3. all digital signatures present on p_k have been verified successfully, in which case p_k is accepted regardless of $b(p_k)$.

Examining Algorithm 3 (line 5), note that if θ_{k+1} is set to a value higher than 1, then, since $0 \leq b(p_k) \leq 1$, $b(p_k)$ will always be less than θ_{k+1} . $i \geq 1$ remains true until all digital signatures are verified. Thus, to always verify all digital signatures present on any received AS_PATH for maximal assurance of path integrity, s_{k+1} can set $\theta_{k+1} > 1$ (e.g., $\theta_{k+1} = 1.1$).

Algorithm 3 AS_PATH Verification (by s_{k+1})

```

1: GLOBAL: threshold  $\theta_{k+1}$ ;  $s_{k+1}$ 's trust ratings  $r(s_1), \dots, r(s_k)$ 
2: INPUT:  $k, p_k = [s_1, \dots, s_k]; sig_1, \dots, sig_k$ 
3: OUTPUT: ACCEPT or REJECT the AS_PATH  $p_k$ 
4:  $i \leftarrow k; n \leftarrow 0; b \leftarrow 0$  /*  $b$  represents  $b(p_k)$  */
5: while  $i \geq 1$  and  $(b < \theta_{k+1}$  or  $n < 2)$  do
6:   if  $sig_i = \phi$  then
7:      $x \leftarrow 0$  /*  $s_i$  has no contribution to belief in  $\lambda_{s_{i-1}, s_i}$  or  $\lambda_{s_i, s_{i+1}}$  */
8:   else if  $sig_i$  fails verification then
9:     return(REJECT)
10:  else
11:     $n \leftarrow n + 1; x \leftarrow r(s_i)$ 
12:  endif
13:  if  $i = k$  then
14:     $b_2 \leftarrow 0; b_1 \leftarrow x$  /* initial belief in  $\lambda_{s_{k-1}, s_k}$  */
15:  else if  $2 \leq i \leq k - 1$  then
16:     $b_2 \leftarrow \text{Algorithm1}(x, b_1)$  /* final belief in  $\lambda_{s_i, s_{i+1}}$  */
17:     $b_1 \leftarrow x$  /* initial belief in  $\lambda_{s_{i-1}, s_i}$  */
18:  else if  $i = 1$  then
19:     $b_2 \leftarrow \text{Algorithm1}(x, b_1)$  /* final belief in  $\lambda_{s_1, s_2}$  */
20:  endif
21:   $b(p_k) \leftarrow b(p_k) + \frac{b_2}{k-1}$  /* update belief in  $p_k$  */
22:   $i \leftarrow i - 1$ 
23: return(ACCEPT)

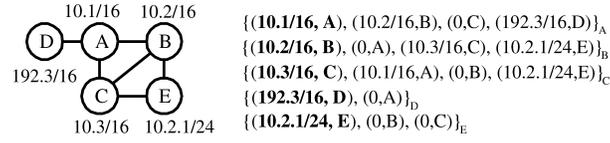
```

If $\theta_{k+1} = 0$, $b(p_k) < \theta_{k+1}$ is always false; once two digital signatures have been verified successfully; $n < 2$ remains false. Thus, no additional digital signature will be verified. Such a configuration meets the minimal requirement by psBGP and achieves maximal efficiency. For $0 < \theta_{k+1} \leq 1$, the number of digital signatures on an AS_PATH to be verified depends on s_{k+1} 's rating of each signing AS on the path.

Such configuration flexibility is in line with the recommendation that “a good initial solution is one that can easily be upgraded to handle increased threats” [Bellovin et al. 2005]. For example, an AS with constrained hardware resources (e.g., CPU) can choose to verify fewer digital signatures on an AS_PATH by setting a lower threshold, while other ASes may choose to verify more or all digital signatures on a signed AS_PATH to achieve a higher assurance of AS_PATH integrity.

We refer to psBGP AS_PATH verification as *stepwise integrity*, which allows confidence ratings on AS_PATH integrity to be formed, based on local parameters, and without requiring all ASes on the AS_PATH to digitally sign the path, or verify all digital signatures present. In contrast, the S-BGP AS_PATH verification approach provides *full integrity*, but requires full adoption of S-BGP by all ASes on the path and verification of all digital signatures present. We acknowledge that the benefit from verifying a partially secured AS_PATH depends on whether or not a non-psBGP-enabled BGP speaker on the path has sufficient memory to store and forward digital signatures (cf. Section 4.1.2).

This stepwise integrity is comparable to the approach taken by IRV [Goodell et al. 2003]. In IRV, one can choose to verify a subset of or the complete AS_PATH, based on the query results, returned from other IRVs and local

Fig. 2. A small AS graph with IP prefixes and *PALs* (0 denotes ϕ).

parameters, e.g., based on psBGP's rating mechanisms. The difference is that IRV adopts an out-of-band approach, which does not require any change to existing BGP implementation but incurs extra query and response messages, while in psBGP, change to BGP is required but information needed for AS_PATH verification is either carried within a BGP update message or stored locally.

4. PREFIX ORIGIN AUTHENTICATION IN PSBGP (G5)

We start with descriptions of *PALs* and MultiASCerts and then introduce how to build from them an *AS prefix graph*. We then describe how psBGP uses an AS prefix graph to verify the propriety of prefix origin per G5 in Section 2.3.

4.1 Prefix Assertion Lists (*PALs*)

Facing the difficulty of building a centralized infrastructure for tracing changes in IP address assignments (recall Section 1), psBGP uses a *decentralized* approach for verifying the propriety of a prefix assertion by cross-checking its consistency with endorsements from the neighbors of the asserting AS.

In psBGP, each AS s_i creates and signs an ordered *prefix assertion list* (pal_i), consisting of a number of tuples of the form $(prefixes, AS\#)$, i.e., $pal_i = \{(F_i, s_i), (F_1, s_1), \dots, (F_n, s_n)\}_{s_i}$, where for the components (F_j, s_j) , $1 \leq j \neq i \leq n$, $s_j \in N(s_i)$, and $s_j < s_{j+1}$. The first tuple (F_i, s_i) is an assertion by s_i of prefixes F_i , which include prefixes owned by s_i itself and prefixes authorized by their respective owners for s_i to originate (referred to as *prefix assertions*); the rest are ordered by AS number, and are assertions by s_i of prefixes assigned to each of s_i 's neighbors (referred to as *prefix endorsements*). If s_i chooses not to endorse any prefix for a neighbor s_j or has no information of s_j 's prefix assignments, s_i simply declares null in its prefix endorsement for s_j . Thus, $(F_j, s_j)_{s_i}$ ($F_j = \phi$) simply asserts that s_j is a direct neighbor of s_i (see Figure 2). If s_i is not willing to disclose that s_j is a direct neighbor, s_i can leave out from pal_i the prefix endorsement for s_j .

Definition 4 (Prefix Consistency). Let $(f_i, s_i)_{s_i}$ be a prefix assertion by s_i and $(f'_i, s'_i)_{s_j}$ a prefix endorsement by s_j . $(f'_i, s'_i)_{s_j}$ is consistent with $(f_i, s_i)_{s_i}$, denoted by $(f'_i, s'_i)_{s_j} \doteq (f_i, s_i)_{s_i}$, if they are regarding the prefix assignment of the same AS, i.e., $s'_i = s_i$, and f'_i is equal to or a superset of f_i , i.e., $f'_i \supseteq f_i$.

Inferred from Definition 4, $(f'_i, s'_i)_{s_j}$ is not consistent with $(f_i, s_i)_{s_i}$, if (1) they are regarding the prefix assignment of different ASes; (2) they have null mutual intersection, i.e., $f'_i \cap f_i = \phi$; or (3) f'_i is a proper subset of f_i , i.e., $f'_i \subset f_i$. In case 3, while f'_i and f_i do share a common subset, which is f'_i , they are not considered consistent in psBGP for the sake of simplicity of AS prefix graph

maintenance. In psBGP, prefix consistency is checked between a prefix assertion and an endorsement, but not between two prefix endorsements.

While an AS is free to decide for which neighbors it provides prefix endorsements and from which to solicit prefix endorsements for itself, we recommend that a provider AS endorse prefixes for a customer AS, possibly becoming a part of an existing service agreement, which includes not only physical network connectivity but now also prefix endorsements. Two neighboring ASes with a peer relationship have freedom to decide how one will endorse prefix assertions made by the other. Prefix endorsements between two peering ASes might be *asymmetric*; in the extreme case, AS s_i may endorse all prefixes assigned to a peering AS s_j , while s_j endorses no prefix assigned to s_i . It is important to allow such flexibility. In the core of the Internet, one AS may peer with many others, some of which may be assigned a large number of prefixes. It would be unrealistic to expect an AS to have full knowledge of all prefixes assigned to such a peer. However, an AS might be able to establish a certain level of confidence in a subset of the prefixes assigned to some of its neighbors. Thus, an AS can distribute such positive (albeit partial) evidence to facilitate other ASes to make a better decision in prefix origin authentication. It is an AS's own responsibility and in its own interest to ensure that its assigned prefixes are endorsed by some of its neighbors or by an RIR.

4.1.1 *Due Diligence.* As a new requirement in psBGP, each AS is responsible for carrying out some level of offline due diligence: for the safety of that AS and of the whole Internet, to increase its confidence that the prefixes it endorses for a direct neighbor are indeed assigned to that AS. We suggest the effort required for this is both justifiable and practical, since two neighboring ASes usually have a business relationship (e.g., a traffic agreement) with each other, allowing some level of offline direct interactions and the establishment of some level of trust. For example, s_i may ask a neighboring AS s_j to show proof that s_j , in fact, holds prefix f_j or is authorized by the holder of f_j to announce f_j . An AS may also ask a senior official of the neighboring AS organization to provide a formal letter asserting the organization's prefix claim. Publicly available information about IP address allocation and delegation may also be helpful. We note that while prefix endorsements may be linked to the reputation of an issuing AS, they are not intended to create any legal liability on the issuing AS (if this is viewed as a practical concern, it might be made an explicit term of agreement to participation in psBGP).

4.1.2 *Propagating PALs in Update Messages.* A PAL can be distributed along with BGP update messages in a newly defined optional and transitive path attribute. A non-psBGP-enabled BGP speaker, which does not understand these newly defined attributes, need not process them but must propagate them. Thus, PALs travel through non-psBGP-enabled BGP speakers to reach psBGP-enabled ones. Each psBGP-enabled BGP speaker can then construct and update its AS prefix graph from received PALs (see Section 4.3). This mechanism assumes that a non-psBGP-enabled BGP speaker has sufficient memory to store and forward PALs (see Section 6.3.1). If some non-psBGP-enabled BGP

speakers cannot meet memory requirements, noncontiguous deployment of psBGP may cause problems. Thus, the above assumption may rightly be viewed as questionable for present-day routers, many of which might have limited memory, e.g., 256 MB. However, as widely agreed, the deployment of a BGP security proposal, like psBGP, would be gradual. Thus, memory burden incurred on a non-psBGP-enabled router might be moderate until a large number of ASes on the Internet have deployed psBGP. One might optimistically hope that if and when a proposal like psBGP might eventually be widely adopted, e.g., in 5 or more years, memory availability at routers would accommodate this (cf. Kent [2003]). *PALs* could alternatively be distributed through out-of-band mechanisms, e.g., security repositories (cf. [Kent et al. 2000]).

4.2 Multiple-AS Certificate (MultiASCerts)

Ideally, two *PALs* issued by two neighboring ASes are based on independent data sources, and, consequently, with high probability (in the absence of collusion), a prefix erroneously asserted by one AS will not be endorsed by any of its neighbors. However, there are some organizations owning multiple ASes and it is a common practice for a multi-AS organization to use a single centralized database for generating router configurations for all of its owned ASes. Thus, it is possible that *PALs* issued by two neighboring ASes owned by a common organization would also be created from a single centralized database. If a prefix is erroneously entered into such a database, it might end up with a pair of erroneous, yet consistent, prefix assertion and endorsement, introducing a single point of failure. We recommend that “best practice” in psBGP requires that an AS obtain prefix endorsement from another AS owned by a different organization. As a recommended BGP local policy, an AS should ignore a prefix endorsement by s_j for s_i if both s_i and s_j are known to be owned by a common organization.

To facilitate the distribution of the knowledge of AS ownership by a multi-AS organization, psBGP makes use of a new certificate, namely MultiASCert (recall Figure 1), which binds a list of ASes owned by a common organization to the name of that organization, and is signed by an RIR. Prefix endorsements by s_j for s_i should be ignored if s_i and s_j appear on a MultiASCert. In this way, human errors by a multi-AS organization regarding a prefix that is assigned to another psBGP-enabled AS and endorsed by an independent neighboring AS will not result in service disruption of that prefix in psBGP (see Section 4.4.1).

4.3 AS Prefix Graph

We introduce as a new concept the *AS prefix graph*, which contains information about *AS connectivity*, *AS prefix assignments* (or prefix-AS bindings) and *ratings* of AS prefix assignments. An AS prefix graph, constructed by each AS s_c , is an attributed graph $G_c = (V, E, H)$, where $V = \{s_i\}$ is a set of AS numbers, $E = \{e_{ij}\}$ is a set of edges (BGP sessions) with e_{ij} connecting s_i to s_j , and $H: V \rightarrow W$ is a function mapping each AS s_i to a set of three-dimensional variables, which specifies the IP prefixes asserted by s_i , and supporting evidence; we call $H(s_i)$ the *APAS set* (associated prefixes and support) for s_i . More

precisely, $H(s_i) = \{(f_x, b_x, C_x)\}$, where $f_x \subseteq \mathbb{P}$ is an IP prefix, $b_x \in [0, 1]$ represents s_c 's confidence that f_x is assigned to s_i and C_x is a list of ASes asserting and endorsing the prefix assignment (f_x, s_i) . We next present how each psBGP-enabled AS constructs and updates its own AS prefix graph based on the *PALs* and MultiASCerts it has received.

4.3.1 AS Prefix Graph Construction. An AS prefix graph is initialized to null before the BGP speaker receives any *PAL* (e.g., when it first connects to the Internet). All BGP speakers within an AS build their own AS prefix graph independently. An AS s_c builds its AS prefix graph $G_c = (V, E, H)$ from the first *pal_i* received from each s_i on the Internet by performing the following tasks: (a) adding s_i and all of its declared neighbors to V ; (b) adding to E an edge from s_i to each of its declared neighbors; (c) updating $H(s_i)$ for each of the prefixes asserted by s_i ; (d) updating $H(s_j)$ for each of the prefixes asserted by $s_j \in N(s_i)$ and endorsed by s_i . (See Algorithm 4 for the details and Section 4.3.3 for an example).

4.3.2 AS Prefix Graph Update. Here we describe how to update an AS prefix graph from a newly received *pal'_i* which replaces an existing *pal_i* that has been previously used to construct or update an AS prefix graph. The prefix AS bindings in *pal_i* and *pal'_i* can be divided into three categories: *removed*, *unchanged*, and *added*. A removed prefix-AS binding appears in *pal_i* but not in *pal'_i*; an unchanged one appears in both; and a newly added one appears in *pal'_i* but not in *pal_i*. Updating an AS prefix graph is performed in two phases (see Algorithm 5 for details) as follows:

Algorithm 4 AS Prefix Graph Construction (for AS s_c)

```

1: GLOBAL:  $G_c = (V, E, H)$ ; existing PALs;  $\{r_c(s_i)|s_i$  is an AS on the Internet}
2: INPUT: pali
3: OUTPUT: updated AS prefix graph  $G_c$ 
4: /*  $F_i, N(s_i)$  are prefixes and neighbors asserted by  $s_i$  for itself in pali respectively */
5:  $V \leftarrow V \cup s_i$ ;  $H(s_i) \leftarrow \phi$ 
6: for each  $f_x \in F_i$  do
7:    $(f_x, b_x, C_x) \leftarrow (f_x, r(s_i), \{s_i\})$ 
8:   for each  $s_j \in N(s_i)$  do
9:      $V \leftarrow V \cup s_j$ ;  $E \leftarrow E \cup e_{ij}$ 
10:    for each prefix endorsement  $(f, s)_{s_j}$  in palj do
11:      /* recall Definition 4 */
12:      if  $(f, s)_{s_j} \doteq (f_x, s_i)_{s_i}$  and  $s_i, s_j$  are not in a common MultiASCert then
13:         $b_x \leftarrow \text{Algorithm1}(b_x, r(s_j))$ ;  $C_x \leftarrow C_x \cup s_j$ 
14:       $H(s_i) \leftarrow H(s_i) \cup (f_x, b_x, C_x)$ ;
15: for each  $s_j \in N(s_i)$  do
16:   retrieve APAS set  $H(s_j) = \{(f_y, b_y, C_y)\}$ 
17:   for each  $(f_y, b_y, C_y) \in H(s_j)$  do
18:     for each prefix endorsement  $(f, s)_{s_i}$  in pali do
19:       if  $(f, s)_{s_i} \doteq (f_y, s_j)_{s_j}$  and  $s_i, s_j$  are not in a common MultiASCert then
20:          $b_y \leftarrow \text{Algorithm1}(b_y, r(s_i))$ ;  $C_y \leftarrow C_y \cup s_i$ 
21:        $H(s_j) \leftarrow H(s_j) \cup (f_y, b_y, C_y)$ 
22: return

```

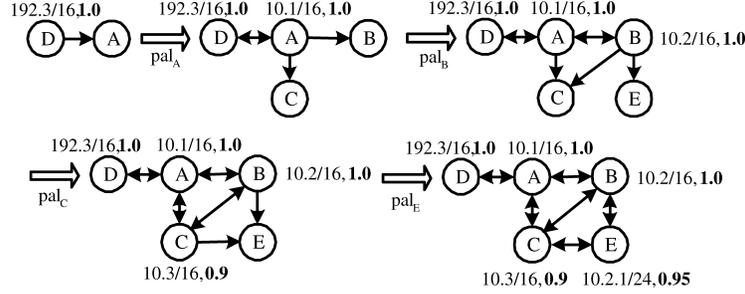


Fig. 3. Construction of an AS prefix graph by AS D (see Example 1).

1. *Removing prefix-AS bindings.* If a removed prefix-AS binding is an assertion, $(f_x, s_i)_{s_i}$, made by s_i for itself, it is simply removed from the graph. If it is an endorsement, $(f_y, s_j)_{s_i}$, by s_i for $s_j \in N(s_i)$, the confidence in s_j 's assertion of f_y must be updated (using Algorithm 2).
2. *Adding prefix-AS bindings.* If an added prefix-AS binding is an assertion, $(f_x, s_i)_{s_i}$, made by s_i for itself, a confidence value must be computed for $(f_x, s_i)_{s_i}$ (using Algorithm 1). If it is a prefix endorsement, $(f_y, s_j)_{s_i}$, and $(f_y, s_j)_{s_j}$ exists in the graph, the confidence in $(f_y, s_j)_{s_j}$ must be updated (using Algorithm 1).

4.3.3 *Example 1.* Figure 3 illustrates Algorithm 4 for an AS D . Assume D fully trusts its service provider A (i.e., $r(A) = 1$), and partially trusts the other ASes ($r(B) = r(E) = 0.5$, $r(C) = 0.8$). The AS prefix graph is constructed based on the following *PALs* received by D in order (here we focus on the construction of the APAS set):

$$\begin{aligned}
 pal_D &= \{(192.3/16, D), (\phi, A)\}_D \\
 pal_A &= \{(10.1/16, A), (10.2/16, B), (\phi, C), (192.3/16, D)\}_A \\
 pal_B &= \{(10.2/16, B), (\phi, A), (10.3/16, C), (10.2.1/24, E)\}_B \\
 pal_C &= \{(10.3/16, C), (10.1/16, A), (\phi, B), (10.2.1/24, E)\}_C \\
 pal_E &= \{(10.2.1/24, E), (\phi, B), (\phi, C)\}_E
 \end{aligned}$$

1. D starts from pal_D issued by itself and updates the graph as: $V = \{D, A\}$; $E = \{e_{DA}\}$; and $H(D) = \{(192.3/16, 1.0, \{D\})\}$. After receiving pal_A , D initializes $H(A)$ to $\{(10.1/16, 1.0, \{A\})\}$ (Algorithm 4 (line 7)). Since A endorses D 's prefix assertion, $H(D)$ is updated to $\{(192.3/16, 1.0, \{D, A\})\}$. While A also endorses B 's prefix assertion, no action is taken at this time, since D has not received pal_B .

Algorithm 5 AS Prefix Graph Update (for AS s_c)

- 1: **GLOBAL:** $G_c = (V, E, H)$; existing *PALs*; $\{r_c(s_i) | s_i \text{ is an AS on the Internet}\}$
- 2: **INPUT:** pal'_i
- 3: **OUTPUT:** updated AS prefix graph G_c
- 4: /* $N(s_i)$ is the set of neighbors asserted by s_i for itself in pal'_i */
- 5: /* Removing prefix-AS bindings */
- 6: **for** each prefix assertion $(f_x, s_i)_{s_i}$ in pal'_i that is not in pal'_i **do**
- 7: retrieve the APAS set $H(s_i) = \{(f_x, b_x, C_x)\}$

```

8:  $H(s_i) \leftarrow H(s_i) - (f_x, b_x, C_x)$  /* set subtraction */
9: for each prefix endorsement  $(f_y, s_j)_{s_i}$  in  $pal_i$  that is not in  $pal'_i$  do
10:   retrieve the APAS set  $H(s_j) = \{(f_y, b_y, C_y)\}$ 
11:   if  $H(s_j) \neq \emptyset$  and  $s_i \in C_y$  then
12:      $b_y \leftarrow \text{Algorithm2}(b_y, r(s_i)); C_y \leftarrow C_y - s_i$ 
13:   for each  $s_j$  in  $N(s_i)$  that is not in  $N(s_i)'$  do
14:      $E \leftarrow E - e_{ij}$ 
15:   if  $s_j$  has no neighbor or prefix assignment in  $G_c$  then
16:      $V \leftarrow V - s_j$ 
17: /* Adding prefix-AS bindings */
18: for each  $s_j$  in  $N(s_i)'$  that is not in  $N(s_i)$  do
19:    $V \leftarrow V \cup s_j; E \leftarrow E \cup e_{ij}$ 
20: for each prefix assertion  $(f_x, s_i)_{s_i}$  in  $pal'_i$  that is not in  $pal_i$  do
21:    $(f_x, b_x, C_x) \leftarrow (f_x, r(s_i), \{s_i\})$ 
22:   for each  $s_j \in N(s_i)'$  do
23:     for each prefix endorsement  $(f, s)_{s_j}$  in  $pal_j$  do
24:       if  $(f, s)_{s_j} \doteq (f_x, s_i)_{s_i}$  and  $s_i, s_j$  are not in a common MultiASCert then
25:          $b_x \leftarrow \text{Algorithm1}(b_x, r(s_j)); C_x \leftarrow C_x \cup s_j$ 
26:    $H(s_i) \leftarrow H(s_i) \cup (f_x, b_x, C_x)$ 
27: for each  $s_j \in N(s_i)'$  do
28:   for each prefix endorsement  $(f, s_j)_{s_i} \in pal'_i$  that is not in  $pal_i$  do
29:     retrieve APAS set  $H(s_j) = \{(f_y, b_y, C_y)\}$ 
30:     for each  $(f_y, b_y, C_y) \in H(s_j)$  do
31:       if  $(f, s_j)_{s_i} \doteq (f_y, s_j)_{s_j}$  and  $s_i, s_j$  are not in a common MultiASCert then
32:          $b_y \leftarrow \text{Algorithm1}(b_y, r(s_i)); C_y \leftarrow C_y \cup s_i$ 
33: return

```

-
2. After receiving pal_B , D initializes $H(B) = \{(10.2/16, 0.5, \{B\})\}$. Since A endorses $(10.2/16, B)$, $\text{Algorithm1}(0.5, 1.0)$ is called to update D 's confidence in $(10.2/16, B)$, and $H(B)$ is updated to $\{(10.2/16, 1.0, \{B, A\})\}$.
 3. After receiving pal_C , D initializes $H(C) = \{(10.3/16, 0.8, \{C\})\}$. Since B endorses $(10.3/16, C)$, $\text{Algorithm1}(0.8, 0.5)$ is called to update D 's confidence in $(10.3/16, C)$ to 0.9, and $H(C)$ is updated to $\{(10.3/16, 0.9, \{C, B\})\}$. Since C endorses A 's prefix assertion, $\text{Algorithm1}(1.0, 0.8)$ is called to update D 's confidence in $(10.1/16, A)$, which does not change, since it already has maximal value 1.0 (see above). $H(A)$ is updated to $\{(10.1/16, 1.0, \{A, C\})\}$.
 4. After receiving pal_E , D initializes $H(E) = \{(10.2.1/24, 0.5, \{E\})\}$. Since B endorses $(10.2.1/24, E)$, $\text{Algorithm1}(0.5, 0.5)$ is called to update D 's confidence in $(10.2.1/24, E)$ to 0.75. Since C also endorses $(10.2.1/24, E)$, $\text{Algorithm1}(0.75, 0.8)$ is called to further update D 's confidence in $(10.2.1/24, E)$ to 0.95. As a result, $H(E)$ is updated to $\{(10.2.1/24, 0.95, \{E, B, C\})\}$.

4.4 Prefix Origin Authentication

Here we describe how to perform prefix origin authentication using an AS prefix graph.

4.4.1 Verification of Prefix Assignment. Two configurable thresholds, denoted by α_i (*sufficient confidence*) and β_i (*sufficient claimants*), are used by each psBGP-enabled AS s_i for verifying the propriety of prefix assignments. α_i is a threshold defining a sufficient confidence level by s_i in a prefix-AS binding

before it can be considered proper. β_i defines a sufficient number of ASes, which assert and endorse a prefix-AS binding before the binding can be considered proper by s_i . In other words, a prefix-AS binding (f_j, s_j) is verified as proper by s_i if s_i 's confidence in (f_j, s_j) is at least α_i , or (f_j, s_j) is asserted by s_j and endorsed by at least $\beta_i - 1$ other ASes. More specifically, a nonaggregated route $(f, [s_j, \dots])$ originated by a psBGP-enabled AS s_j is verified by another psBGP-enabled AS s_i as *proper* if (a) there exists $(f_x, b_x, C_x) \in H(s_j)$; (b) $b_x \geq \alpha_i$ or $|C_x| \geq \beta_i$; and (c) $f \subseteq f_x$. Algorithm 6 specifies this explicitly.

Algorithm 6 Verification of Prefix Assignment (by an AS s_i)

```

1: GLOBAL:  $G_i = (V, E, H); \alpha_i; \beta_i$ 
2: INPUT: The BGP route  $m = (f_j, p = [s_j, \dots])$ 
3: OUTPUT: ACCEPT or REJECT  $s_j$ 's origin of  $f_j$ 
4: retrieve the APAS set  $H(s_j) = \{(f_x, b_x, C_x)\}$  from  $G_i$ 
5: for each  $(f_x, b_x, C_x) \in H(s_j)$  do
6:   if  $(b_x \geq \alpha_i$  or  $|C_x| \geq \beta_i)$  and  $f_j \subseteq f_x$  then
7:     return(ACCEPT)
8: return(REJECT)

```

α_i and β_i are independent and, in conjunction, provide extensive flexibility. $\alpha_i = 1$ allows s_i to immediately accept a prefix assertion by a fully trusted AS (i.e., without any neighbor endorsement), while prefix assertions made by partially trusted ASes require endorsements from a sufficient number of neighbors. α_i and β_i can also be configured such that only one or neither takes effect. For example, $\alpha_i > 1$ and $\beta_i \geq 1$ allows β_i to always take precedence, since the maximum confidence in a prefix assertion is 1. $b_i = \infty$ and $0 < \alpha_i \leq 1$ has the opposite effect. $\alpha_i = 0$ and $\beta_i = 0$ emulate the existing nonsecured BGP behavior (i.e., any prefix originated by any AS is considered as proper).

During the early stages of psBGP deployment, when only a small number of ASes have deployed psBGP, we recommend $\beta_i = 1$ for each psBGP-enabled AS s_i . In other words, a psBGP-enabled AS s_i allows another psBGP-enabled AS s_j to originate a prefix f_j if f_j is asserted in pal_j even it is not endorsed by any neighbor. This reflects the reality that early psBGP adopters might not have any psBGP-enabled neighbors and it offers some level of assurance (albeit limited). For example, a compromised BGP speaker within a psBGP-enabled AS s_j cannot be used to hijack prefixes assigned to other ASes unless keying material required for issuing pal_j is also compromised. In addition, the existence of a public statement about an assertion provides some assurance, in that this might carry some weight in legal dispute or affect business reputation. (See Section 6.1.2 for more discussion on incremental benefits and Section 5.2.3 on limitations of psBGP.)

After a majority of ASes have deployed psBGP, we recommend $\beta_i = 2$, i.e., a psBGP-enabled AS s_i allows another psBGP-enabled AS s_j to originate a prefix f_j only if f_j is asserted in pal_j and is endorsed by one of s_j 's neighbors. $\beta_i = 2$ is resilient to some errors resulting from a single AS. For example, if s_j mistakenly asserts a prefix f in pal_j and announces f via BGP, this would not result in service disruption of the legitimate owner of f as long as s_j 's assertion of f is not endorsed by any neighbor. However, $\beta_i = 2$ remains vulnerable to two-party

collusion. More generally, $\beta_i = k \geq 2$ resists collusion by $k - 1$ parties. Larger β_i renders a stronger assurance in the propriety of a prefix assignment, but trades off performance and results in higher maintenance overhead (see Section 6.3.4).

4.4.2 Verification of Prefix Aggregation. Suppose AS s_1 is assigned a set of prefixes F_1 . When receiving a set of routes with a set of prefixes F_2 , the BGP specification [Rekhter and Li 1995] allows s_1 to aggregate F_2 into a single prefix f_g to reduce routing information to be stored and transmitted. We call f_g an *aggregated prefix*. s_1 can aggregate F_2 into f_g if one of the following conditions holds: (1) $\forall f_i \subseteq f_g, f_i \subseteq F_1$; or (2) $\forall f_i \subseteq f_g, f_i \subseteq F_1 \cup F_2$.

In case (1), s_1 must be assigned a set of prefixes F_1 , which is a superset of the aggregated prefix f_g . Most likely, f_g is one of the prefixes assigned to s_1 , i.e., $f_g \in F_1$. This type of aggregation is sometimes referred to as *prefix reorigination*. From a routing perspective, prefix reorigination does not have any effect, since traffic destined to a more specific prefix will be forwarded to the reoriginating AS and then forwarded to the ultimate destination from there. From a policy enforcement perspective, prefix reorigination does have an effect, since the AS_PATH of an aggregated route is different from any of the AS_PATHs of the routes to be aggregated. Since AS_PATH is used by the route selection process, changing AS_PATH has an impact on route selections. From a security perspective, prefix reorigination is no different than normal prefix origination, since the aggregated prefix is either the same as, or a subset of, the prefix assigned by the aggregating AS. Therefore, f_g can be verified using the mechanism in Section 4.4.1.

In case (2), s_1 is not assigned the whole address space of the aggregated prefix f_g . Therefore, f_g cannot be verified in the same way as for prefix reorigination. To facilitate verification of the propriety of route aggregation by a receiving AS, psBGP imposes a new requirement: the routes to be aggregated must be supplied by the aggregating AS along with the aggregated route. This approach is essentially similar to that taken by S-BGP. Transmission of routes to be aggregated incurs additional network overhead, which is something BGP tries to reduce. However, we view such additional overhead to be relatively insignificant given that modern communication networks generally have high bandwidth and BGP control messages account for only a small fraction of subscriber traffic. The main purpose of route aggregation is to reduce the size of routing tables, i.e., reducing storage requirements; note that this is preserved by psBGP.

4.5 Route Selection Algorithm

In standard BGP, when a BGP speaker receives two valid routes with the same destination prefix, a route selection process is invoked to determine which is preferable. In what follows, a prefix-AS binding of a route means the binding of the prefix and the AS that originates that route. psBGP adds two new rules: one gives preference to a route whose prefix-AS binding has more neighbor endorsements, and the other to a route whose prefix-AS binding is rated higher. These two new rules are added into the fourth and fifth places in BGP route selection algorithm [Rekhter and Li 1995] to preserve existing traffic engineering

practices, which usually employ *local_pref*, *as_path* and *med* (*mult_exit_disc*). Note that the higher-numbered rule is followed if the lower-numbered rules result in a tie.

1. Select the route with a higher degree of preference, i.e., a higher *local_pref* value.
2. Select the route with a shorter *as_path*.
3. Select the route with a lower *med* value if they have the same *next_hop*.
4. *Select the route whose prefix-AS binding is endorsed by more neighbors.*
5. *Select the route whose prefix-AS binding is rated higher.*
6. Select the route with a lower intradomain routing cost to the *next_hop*.

Ongoing work [Retana and White 2002] suggests to allow customer-defined rules to be inserted anywhere in the standard BGP route selection algorithm. If this is implemented in psBGP, customers with high security requirement can choose to move psBGP-related rules up to an appropriate decision point, e.g., as rules 1 and 2.

We do not expect the proposed changes to BGP route selection process will have material impact on route convergence, since they are placed near the bottom of the process. It is likely that the route selection process will end after the first three rules have been applied. However, if one chooses to move the psBGP-related rules up, there will be some effect on route convergence. The actual effect depends on a number of factors, including the number of ASes adopting such changes, and their locations on the Internet, among others. Further study is needed to obtain quantitative results of the effect.

5. SECURITY ANALYSIS OF PSBGP

We first analyze psBGP against the listed security goals from Section 2. We then discuss how psBGP counters selected BGP threats.

5.1 Meeting Specified Security Goals

The analysis below clarifies how the proposed psBGP mechanisms meet the specified goals and by what line of reasoning and assumptions. While we believe that mathematical “proofs” of security may often be based on flawed assumptions or models (e.g., see [Koblitz and Menezes 2004]) that fail to guarantee “security” in any real-world sense, they are nevertheless very useful, e.g., for finding security flaws, for precisely capturing protocol goals, and for reducing ambiguity, all of which increase confidence. We thus provide outlines of such formalized reasoning, as a complement to alternative methods of increasing confidence.

PROPOSITION 1. *psBGP provides AS number authentication (G1).*

Proof Outline: For an AS number s to be certified, psBGP requires an AS-NumCert $(k_s, s)_{\overline{k_T}}$. Since T (i.e., an RIR) controls s , and is the trusted guardian of AS numbers (by assumption), any assertion made by T about s is proper. Thus $(k_s, s)_{\overline{k_T}}$ is proper. In other words, s is an AS number certified by T , and

k_s is a public key associated with s certified by T . More formally,³ (T controls s) \wedge $(k_s, s)_{\overline{k_T}} \Rightarrow (k_s, s)$ is a proper binding.

PROPOSITION 2. *psBGP provides BGP speaker authentication (G2).*

Proof Outline: For a BGP speaker g to be accepted as an authorized representative of an AS s , psBGP requires an ASNumCert $(k_s, s)_{\overline{k_T}}$, a SpeakerCert $(k'_s, s)_{\overline{k_s}}$, and evidence that g possesses $\overline{k'_s}$. By Proposition 1, $(k_s, s)_{\overline{k_T}}$ establishes that s is an AS number certified by T and k_s is a public key associated with s certified by T . Similarly, $(k'_s, s)_{\overline{k_s}}$ establishes that k'_s is a public key associated with s certified by s . Evidence that g possesses $\overline{k'_s}$ (i.e., an appropriate digital signature generated by g using $\overline{k'_s}$) establishes that g is authorized by s to represent s . Thus, the Proposition is established. More formally, (T controls s) \wedge $(k_s, s)_{\overline{k_T}} \Rightarrow (k_s, s)$ is a proper binding; (k_s, s) is proper \wedge $(k'_s, s)_{\overline{k_s}} \Rightarrow (k'_s, s)$ is proper binding; (k'_s, s) is proper \wedge g possesses $\overline{k'_s} \Rightarrow g$ is authorized by s .

PROPOSITION 3. *psBGP provides data integrity (G3).*

Proof Outline: psBGP uses the IPsec ESP [Kent and Atkinson 1998b] with null encryption for protecting BGP sessions, and relies upon IPsec ESP for data integrity. Thus, this provides data integrity in practice, to the extent that one can rely on practical implementations of IPsec ESP.

PROPOSITION 4. *psBGP provides assurance of AS_PATH authentication (G4).*

Proof Outline: Let $m_k = (f_1, p_k)$ be a BGP route, where $p_k = [s_1, \dots, s_k]$, and m_k is originated or forwarded by a BGP speaker in s_k . For simplicity, we refer to an AS instead of a BGP speaker within that AS. In psBGP, the integrity of p_k implies that m_k has traversed the exact sequence of s_1, \dots, s_k . We next use induction on path length to show that psBGP provides AS_PATH integrity when all ASes on an AS_PATH are psBGP-enabled and the verifying AS chooses to verify all digital signatures on the path, followed by discussion of other cases.

1. If $k = 1$, psBGP requires that for s_2 to accept m_1 , s_2 must receive a valid digital signature $sig_1 = \{f_1, [s_1, s_2]\}_{s_1}$, which serves as a signed assertion that s_1 originated m_1 (and advertised it to s_2).
2. Assume when $k = n \geq 2$, there exist digital signatures sig_1, \dots, sig_n which assert that m_n indeed traversed the exact sequence of s_1, \dots, s_n . When $k = n + 1$, we need to show that m_{n+1} has traversed from s_n to s_{n+1} and exited s_{n+1} . $sig_n = \{f_1, [s_1, \dots, s_n, s_{n+1}]\}_{s_n}$ asserts that s_n forwards m_n to s_{n+1} . psBGP requires that s_{n+1} digitally signs m_{n+1} by generating a digital signature $sig_{n+1} = \{f_1, [s_1, \dots, s_{n+1}, s_{n+2}]\}_{s_{n+1}}$, which serves as the evidence that m_{n+1} is advertised by s_{n+1} to another AS s_{n+2} . In summary, sig_n asserts that m_n traversed from s_n to s_{n+1} , and sig_{n+1} asserts that m_n is transformed by s_{n+1} to m_{n+1} which traversed through s_{n+1} to another AS. Thus, the above three steps establish Proposition 4 when all ASes on an AS_PATH are psBGP-enabled and the verifying AS verified all digital signatures on the path.

³Here we adapt BAN-like notation, modified for our purpose (cf. [Burrows et al. 1989; Gaarder and Snekenes 1991; Gligor et al. 1991]).

Partial AS_PATH integrity. If an AS chooses not to always verify all digital signatures on the path (i.e., setting $\theta < 1$, or some digital signatures are missing; see Algorithm 3 and Section 3.5), full integrity of the path is not guaranteed. For example, let $p_k = [s_1, \dots, s_j, \dots, s_k]$. If an AS only verifies the digital signatures generated by ASes from s_j to s_k , only the integrity of that the path segment is protected. The path from s_1 to s_{j-1} can be falsified if all ASes from s_j to s_k are in collusion. As another example, consider the route $m = (f, [s_1, s_2, s_3, s_4])$ with only s_2 psBGP-enabled. The digital signature generated by a well-behaved s_2 , $\{f, [s_1, s_2, s_3]\}_{s_2}$, covers the path $[s_1, s_2, s_3]$. In other words, a malicious AS cannot compromise the integrity of $[s_1, s_2, s_3]$, but it can insert any non-psBGP-enabled AS after s_3 or modify s_4 to another non-psBGP-enabled AS. In addition, $[s_1, s_2, s_3]$ can be removed or replaced as a whole with other non-psBGP-enabled ASes.

We next establish Proposition 5. As discussed in Section 3.1, psBGP uses a rating mechanism to provide the flexibility to allow an AS to fully trust an AS or an RIR, thus accepting their prefix assertions without requiring additional endorsements. We recommend that no AS should be fully trusted unless there is strong reason to do so. In the rest of our analysis, we assume that a verifying AS s_i does not immediately trust any other AS s_j . In other words, s_i rates every other AS s_j with a value lower than its confidence threshold, i.e., $r_i(s_j) < \alpha_i$. Before presenting Proposition 5, we establish two Lemmas.

LEMMA 1. *Assuming that no two ASes are in collusion (A1),⁴ then psBGP with threshold $\beta = 2$ provides reasonable⁵ assurance of prefix assignment verification, i.e., a prefix assignment that is verified as proper is, with reasonable assurance, proper.*

Proof Outline: Consider the BGP route $m = (f_x, [s_i, \dots])$. For f_x to be verified as assigned to s_i , psBGP requires that for some f_i :

(R1) prefix assertion $(f_i, s_i)_{s_i}$ exists; **(R2)** $(f'_i, s_i)_{s_j} \doteq (f_i, s_i)_{s_i}$ exists for $s_j \in N(s_i)$; **(R3)** s_i, s_j do not appear in a common MultiASCert; and **(R4)** $f_x \subseteq f_i$

R1, R2, and R3 establish that f_i is assigned to s_i , and R4 shows that f_x is a subset of f_i . Suppose f_i is not assigned to s_i but is verified as such (i.e., R1–R4 are met). For this statement to be true, the following statements must be true: $(f_i, s_i)_{s_i}$ is improper; and $(f_i, s_i)_{s_j}$ is improper. Since $(f_i, s_i)_{s_i}$ and $(f_i, s_i)_{s_j}$ are improper and consistent, s_i and s_j either share a common false data source (H1) or they are considered in collusion (H2). R3 reduces the likelihood of H1 and H2 is ruled out by assumption A1. Thus, the statement that f_i is not assigned to s_i but is verified as such is, with reasonable assurance, not true. In other words, if f_i is not assigned to s_i , it will, with reasonable assurance, not be verified as such. Equivalently, if f_i is verified as assigned to s_i , it is, with reasonable assurance, assigned to s_i . This establishes Lemma 1.

⁴See Section 5.2.3 for discussion of examples where this collusion assumption (A1) may not hold.

⁵By reasonable, we mean to emphasize that our claim is relative to our threat model and assumptions (e.g., see Section 5.2.3); we cannot claim absolute security (which we do not believe exists in the real world).

LEMMA 2. *psBGP provides reasonable assurance of IP prefix aggregation verification.*

Proof Outline: Let f_g be a prefix aggregated by AS s_x from a set of routes $\{m_i = (f_i, p_i) | p_i = [s_i, \dots]\}$ received by s_x . psBGP requires that for f_g originated by s_x to be verified as proper, s_x must either own a prefix f_x such that $f_g \subseteq f_x$ (verified by Lemma 1), or provide evidence that s_x has, in fact, received $\{m_i\}$ and $f_g \subseteq \cup\{f_i\}$. Valid digital signatures from each AS on p_i can serve as evidence that s_x has received $\{m_i\}$ (see Proposition 4). If $f_g \subseteq \cup\{f_i\}$, then s_x aggregates f_g properly. If s_x cannot provide the required evidence, s_x 's aggregation of f_g is verified as improper. This establishes Lemma 2.

PROPOSITION 5. *psBGP provides reasonable assurance of IP prefix origination authentication (G5), i.e., an AS s_i 's origination of a prefix f is, with reasonable assurance, verified as proper if f is assigned to s_i or is aggregated properly by s_i from a set of routes received by s_i .*

Proof Outline: Lemma 1 allows prefix assignment verification and Lemma 2 allows prefix aggregation verification, thus establishing Proposition 5.

The above results (Propositions 1–5) establish the psBGP security properties, as summarized by Theorem 1 (cf. Section 2.3).

THEOREM 1 (PSBGP SECURITY PROPERTIES). *psBGP achieves the following five security goals: AS number authentication (G1), BGP speaker authentication (G2), data integrity (G3), AS_PATH authentication (G4), and prefix origin authentication (G5).*

5.2 Countering Selected BGP Threats

We first consider how psBGP detects false prefix originations and next discuss how psBGP reacts to possible new threats arising from proposed security mechanisms in psBGP itself. We also discuss some attack scenarios not addressed by psBGP.

5.2.1 *Detecting False Prefix Origin.* We consider three cases in which an AS may originate routes for a prefix, which is actually assigned to another AS.

5.2.1.1 *Malicious Attack.* A malicious AS may hijack a prefix from another AS to attract its traffic. An AS is considered malicious if one or more BGP speakers within that AS are compromised, or the administrator in the AS that controls BGP software and configuration intentionally misbehaves. psBGP can detect prefix hijacking, since a malicious AS will be unable to obtain from its neighbors or a trusted authority (e.g., an RIR) endorsements for the hijacked prefix.

5.2.1.2 *Router Malfunction.* A router may mistakenly deaggregate prefixes (e.g., due to software problems) and announce more specific ones. Deaggregating another AS's prefix is referred to as *foreign deaggregation*; deaggregating one's own prefix is referred to as *self-deaggregation*. Foreign deaggregation has the same external behavior as prefix hijacking and, thus, can be detected.

Self-deaggregation appears to be equivalent to the announcement of a subset of the prefix assigned to an AS and, thus, is treated as legitimate.

5.2.1.3 Database Misconfiguration. Many ISPs use automatic scripts to generate router configurations from a centralized database containing information of prefix assignments. If a prefix is erroneously entered into such database (e.g., due to human error), automatically generated configurations will instruct a router, which might be functioning correctly, to originate a prefix, which it is not authorized to announce.

Database misconfiguration will not result in successful prefix hijacking if the erroneous database is not used by *any* neighboring AS to generate its *PAL*. In other words, if the information used by all endorsing ASes for generating *PALs* is independent of the misconfigured database containing erroneous prefixes, origin of those prefixes will result in verification failures, since there will not exist a prefix endorsement consistent with the false prefix assertion. However, an ISP may have multiple ASes and use a single centralized database for generating both router configurations and *PALs* for its own ASes. Thus, it is possible that an erroneous prefix assertion made by one AS gets endorsement from another AS owned by the same ISP. This scenario is addressed in psBGP with MultiASCerts (Section 4.2). More specifically, an endorsement from s_i for a prefix assertion made by s_j is not used if both s_i and s_j are owned by the same organization, in which case they should both appear on a MultiASCert under a common organization.

5.2.2 Countering False *PALs*. We now discuss how psBGP reacts to erroneous *PALs* that contain false assertions or endorsements. These might potentially introduce new vulnerabilities arising from the proposed enhancements, as a result of malice or human error.

5.2.2.1 Erroneous Prefix Assertions. An AS s_i erroneously asserting the ownership of a prefix through its own *PAL* will not result in service disruption of the legitimate owner of that prefix as long as none of s_i 's neighbors endorses its assertion.

5.2.2.2 Erroneous Prefix Endorsements. An AS s_i erroneously endorsing s_j for a prefix, which is not asserted by s_j , will not result in any service disruption, since such an endorsement will not be used by any AS when it verifies s_j 's prefix assertions. If s_i is the only endorsing neighbor for s_j , or more generally, $\forall s_i \in N(s_j), s_i$ issues $(f'_j, s_j)_{s_i}$ inconsistent with $(f_j, s_j)_{s_j}$, then $(f_j, s_j)_{s_j}$ will be verified as *improper* by other ASes, even if it is actually correct. This is the case when misbehaving ASes form a network cut from s_j to any part of the network. It appears difficult, if not impossible, to counter such an attack; however, we note that even if such a denial-of-service attack could be prevented, many other techniques beyond the control of BGP could also be used to deny the routing service of s_j , e.g., link cuts [Bellovin and Gansner 2003], filtering, or packet dropping. Note that a prefix assertion made by s_i about a remote AS s_k , i.e., $s_i \notin N(s_k)$, will not be checked when s_k 's prefix assertions are verified, because

s_i is not a neighbor of s_k . Thus, a misbehaving AS is unable to mislead other ASes about the prefix ownership of a nonneighboring AS.

5.2.3 Limitations of psBGP. We now discuss some limitations of psBGP. First, it is subject to human error if a psBGP-enabled AS s_i sets threshold $\beta_i = 1$ (e.g., during the early stage of psBGP deployment on the Internet). For example, if an AS uses a common database for generating BGP speaker configuration and for issuing *PALs*, a prefix erroneously entered into such a database can result in service disruption. Second, psBGP is subject to k -party collusion if $\beta_i = k \geq 2$. Suppose $\beta_i = 2$, which is the recommended configuration (see Section 4.4.1) for each psBGP-enabled AS s_i . If an attacker controls two ASes that are owned by two different organizations (i.e., they do not appear on a common MultiASCert), it is possible for the attacker to generate two erroneous yet consistent *PALs*. This is equivalent to the case that the *PALs* issued by two different ASes are, in fact, based on a single data source; thus corroborating these two dependent *PALs* does not yield additional benefit. As a result, psBGP security can be defeated. To successfully launch such an attack, an adversary needs to: (a) set up two organizations and manage to obtain an AS number from an RIR for each of them; (b) compromise the private keys used by two independent ASes for signing their *PALs*; or (c) set up one organization and manage to obtain an AS number from an RIR and compromise the private key used by another AS for signing its *PAL*. We suggest that these attacks would present nontrivial (albeit not insurmountable) practical difficulty to an adversary. Moreover, additional mechanisms can be implemented to detect and mitigate the effect of these collusion attacks. For example, one can implement a policy to favor a prefix endorsed by more ASes or by an AS, which has been verified to hold a larger address space of that prefix. If a collusion attack does indeed occur and is detected, this could be reflected within the rating system by lowering the rating of the colluding ASes (giving them less credibility) with the intent of making it harder for them to launch new attacks in the future.

6. OPERATIONAL ANALYSIS OF PSBGP

Here we analyze some operational and performance issues of psBGP.

6.1 Deployment Analysis of psBGP

We first argue that the effort involved in deploying psBGP is reasonable (relative to alternatives) and next discuss incremental benefits from psBGP deployment.

6.1.1 Reasonable Deployment Effort. To deploy psBGP, an AS needs to: upgrade its BGP speakers to support psBGP; issue a single SpeakerCert for all of its own BGP speakers and a unique SessionCert for each of them; distribute the corresponding private keys securely to its speakers; and periodically issue an appropriate *PAL*. Upgrading BGP speakers can be done in a similar manner as upgrading existing router software, although this may require to add more memory (cf. Sections 4.1.2 and 6.3.1). Issuing a SpeakerCert (e.g., in X.509v3 format) requires some level of knowledge of public key certificates. However,

many people responsible for BGP operations might have already acquired similar knowledge, e.g., from the use of PGP [Zsako 1999]; in any case, we acknowledge that additional effort will always be involved in setting up a new system. For example, personnel familiar with PGP may still need to spend some time studying the X.509v3 certificate format. Issuing a *PAL* requires carrying out a certain level of due diligence (Section 4.1.1) in improving an AS's confidence in the prefixes assigned to a (typically) small number of selected neighbors. We expect such effort is reasonable, since two direct neighbors usually have established service agreements allowing some level of direct interaction. Such effort is also justifiable (in our opinion) considering potential security benefit to the Internet as a whole. Overall, all of this work can be done independently by an AS without requiring authorization from other ASes (e.g., an upstream ISP). In other words, psBGP can be deployed from the bottom up, mirroring the growth model of the Internet.

6.1.2 Incremental Deployability. As with the deployment of almost any other large-scale security system, it is unrealistic to expect psBGP to be deployed by all ASes simultaneously or to be deployed at different times, but turned on at the same time. It is expected that if adopted, a small number of ASes will deploy psBGP first, then more and more ASes will follow. It is desirable that those ASes deploying psBGP first can achieve some immediate benefits to justify their investment before psBGP is widely deployed. Here we analyze benefits and constraints of psBGP deployment ($\beta = 1$), assuming all certificates and *PALs* required for verifying a prefix origin are available to a psBGP-enabled AS (cf. Section 4.1.2).

The first AS adopting psBGP does not gain any immediate benefit, since none of the other ASes speaks psBGP. The second AS adopting psBGP will have some benefit collectively with the first psBGP-enabled AS if they are direct neighbors. In this case, one psBGP-enabled AS (s_i) will likely prefer the route originated by the other (s_j) over routes originated by a non-psBGP-enabled AS regarding a prefix assigned to s_j (see Section 4.5). Since s_i and s_j are also directly connected, traffic originated from s_i and destined to s_j will likely arrive at s_j and not be attracted to another AS if everything else besides BGP also works correctly. In the case that s_i and s_j are not directly connected, i.e., connected by one or more non-psBGP-enabled ASes, s_i will still likely prefer the route originated by s_j over an erroneous one by a non-psBGP-enabled AS (see Section 4.5), resulting in containment of any erroneous announcements. However, there is no assurance that traffic destined to s_j can reach their ultimate destinations from s_i . This is because such traffic must traverse through non-psBGP-enabled ASes (or unsecured zones), some of which could have poisoned routing tables and direct traffic over incorrect paths. Thus, security that can be achieved by two remote psBGP-enabled ASes is less than that achieved by two psBGP-enabled neighbors.

We say that one or more psBGP-enabled ASes with direct links among themselves form a *secure zone* and one or more non-psBGP-enabled ASes with direct links among themselves form a *nonsecure zone*. Assume at one point, a number of ASes on the Internet have deployed psBGP. The Internet can then be

viewed to consist of a number of secure and nonsecure zones. Since two directly connected secure or nonsecure zones can always form a larger secure or nonsecure zone, a secure zone will always directly connect with nonsecure zones and a nonsecure zone can have only secure zones as its direct-zone neighbors. This implies that secure zones can always form a network cut for a nonsecure one. To this end, we can draw two conclusions: (1) an AS improperly originating a route for a prefix assigned to a psBGP-enabled AS will be contained once it reaches a secure zone. In other words, if a misbehaving AS is within a secure zone, the erroneous route will be immediately contained. If it is within a nonsecure zone, it will propagate within the nonsecure zone and be contained once it reaches a secure zone. (2) An improper origination of a prefix assigned to a non-psBGP-enabled AS will be propagated (without detection by psBGP) through all nonsecure and secure zones, i.e., over the entire Internet.

It is clear from the above conclusions that prefixes assigned to a psBGP-enabled-AS are protected to a certain degree from being hijacked, while there is no such protection for non-psBGP-enabled ASes. While a psBGP-enabled AS might find limited protection when the number of other psBGP-enabled ASes is small, the protection increases as this number grows. As a starting point, it might be beneficial for an organization, which owns multiple ASes (such as a large or even medium-sized government), to deploy psBGP so that a secure zone can be formed within that organization.

6.2 Complexity Analysis of psBGP

Here we consider the computational complexity resulting from AS_PATH verification and AS prefix graph related operations. The former involves computationally expensive operations, such as digital signature generation and verification, while the latter involves much simpler (less costly but potential numerous) operations, such as data structure insertion, deletion, comparison, and query. We do not attempt to provide a detailed, mathematically rigorous running-time analysis for psBGP operations, but rather to provide enough insight to allow ballpark estimates sufficient to provide confidence that computational costs of psBGP are reasonable and will not be a reason to avoid deploying psBGP.

6.2.1 Complexity of AS_PATH Verification. Let a be the average number of external ASes with which a BGP speaker establishes BGP sessions and b the average number of ASes on an AS_PATH. A psBGP-enabled BGP speaker needs to generate on average a unique digital signatures (one per AS neighbor) for each BGP update message it sends to a neighbors and to verify, on average, b unique digital signatures (for maximal security, i.e., $\theta = 1$) for each BGP update message received (see Algorithm 3). Signature verifications related to certificate revocation and certificate chains are ignored here.

6.2.2 Complexity of AS Prefix Graph Operations. Let n be the total number of ASes on the Internet, d the average number of AS neighbors, and h the average number of prefixes assigned to an AS. Let $x \leq d$ be the average number of neighboring ASes whose prefix assertions are endorsed by an AS and y

the average number of prefixes endorsed by an AS for each such neighbor. Accordingly, each AS on average has x endorsing neighbors.

Thus, each *PAL* (cf. Section 4.1) on average consists of: (1) h prefix assertions, one for each assigned prefix; (2) y prefix endorsements for each endorsed neighbor (x of them), resulting in xy prefix endorsements in total; and (3) $d - x$ null prefix endorsements, one for each non-endorsed neighbor. Assume there are z MultiASCerts. We next estimate the computational costs of the construction, update, and query of an AS prefix graph in psBGP. Note all operations mentioned here are simple database operations (e.g., comparison), not computationally expensive operations, such as digital signature generation or verification.

1. *Complexity of AS Prefix Graph Construction* (Algorithm 4). For the first pal_i received from each AS on the Internet, an AS needs to update the APAS $H(s_i)$ for s_i (lines 6–13), resulting in $h\{1 + d[2 + xy(1 + z + 1 + 1)]\}$ operations. In addition, an AS also needs to update the APAS $H(s_j)$ for each of s_i 's endorsed neighbors s_j (lines 14–20), resulting in $d\{1 + h[xy(1 + z + 1 + 1) + 1]\}$ operations. Thus, in total $2hdxyz + 6hdxy + 3hd + h + d$ operations are required for processing each pal_i , resulting in $n(2hdxyz + 6hdxy + 3hd + h + d)$ operations for constructing a complete AS prefix graph from n *PALs*.
2. *Complexity of AS Prefix Graph Update* (Algorithm 5). Consider the worst case that an AS s_i issues a new pal'_i that is completely different from the existing pal_i , i.e., all of its prefix assertions and endorsements have changed. In Algorithm 5, lines 6–7 result in h operations, lines 8–11 result in $5xy$ operations, lines 12–18 result in $5d$ operations, lines 19–25 result in $h\{1 + d[xy(1 + z + 1 + 1)] + 1\}$ operations, and lines 26–31 result in $d\{xy[1 + h(1 + z + 1 + 1)]\}$ operations. Thus, one update might require, in total, $2hdxyz + 6hdxy + dxy + 5xy + 3h + 5d$ operations.
3. *Complexity of AS Prefix Graph Query* (Algorithm 6) When an AS receives a BGP update message, it verifies that the origin AS is allowed to announce the prefix by comparing the announced prefix with the h prefixes asserted by the origin AS, resulting in up to h operations for one prefix origin verification.

6.3 Performance Analysis of psBGP

Here we present our preliminary estimation of memory, bandwidth, and CPU overhead, and the analysis of certificate dynamics in psBGP. While rigorous study has been performed by Aiello et al. [2003] on the prefix delegation stability on the Internet as a whole, and by Nicol et al. [2004] and Zhao et al. [2005b] on PKI impact on BGP security using simulation, it is desirable to study certificate dynamics of a secure system and to project certificate management overhead on a per-AS level. We use BGP data collected by the RouteViews project [RouteViews 2005]. We retrieved one BGP routing table the first day of each month from January to August 2004. Despite known shortcomings including incompleteness of the RouteViews data set, it is one of the most complete data repositories publicly available, and has been widely used in the BGP community.

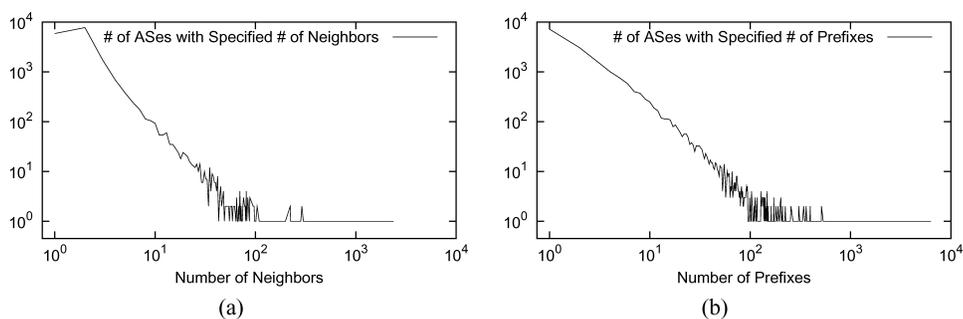


Fig. 4. Distribution of AS Neighbors and prefix assignments.

6.3.1 Memory Overhead. Four types of certificates, one AS prefix graph, and digitally signed BGP update messages require memory storage for a BGP speaker to support psBGP. We estimate the memory overhead for each type and then estimate the total. We omit the memory requirement for storing SessionCerts, since a BGP speaker only needs to store them for a small number of direct neighbors (e.g., fewer than tens).

6.3.1.1 ASNumCerts and SpeakerCerts. We observed in total 17,884 ASes⁶ as of August 1, 2004. One ASNumCert is required per AS. In the worst case, an AS may need to store the ASNumCert of every AS on the Internet; in this case, 17,844 ASNumCerts would be stored. As with S-BGP and soBGP, psBGP recommends use of the X.509v3 certificate structure because of wide industrial support. Assuming the average size of a certificate is 600 bytes [Kent 2003] based on 1024-bit RSA keys, 10.5 MB of memory would be required for storing 17,844 ASNumCerts. The same holds for SpeakerCerts.

6.3.1.2 PALs and MultiASCerts. The size of pal_i , issued by each AS s_i , is primarily determined by the number of prefixes assigned to s_i , the number of s_i 's neighbors, and the number of prefixes assigned to each of s_i 's neighbors that are endorsed by s_i . Figures 4a and b, respectively, illustrate the distribution of AS neighbors and AS prefix assignments, based on the RouteView dataset in July 2004. We can see that while some ASes have many neighbors, and some are delegated many prefixes, many ASes have only a small number of neighbors and are delegated a small number of prefixes. On average, each AS has 4.2 neighbors and is delegated 9.1 prefixes. Assuming the average size of a PAL is 1024 bytes (600 bytes for an X.509v3 certificate plus 424 bytes for about 60 prefix assertions and endorsements), 17.8 MB of memory would be required to store 17,844 PALs, one for each AS. For MultiASCerts, a BGP speaker needs to store one certificate for each organization that owns multiple ASes. Based on the data from Aiello et al. [2003], there are 385 multi-AS organizations, which, in total, own 1259 ASes. On average, each multi-AS organization owns 3.3 ASes. Assuming the average size of a MultiASCert is 600 bytes, 0.226 MB of memory are required by each AS for storing all MultiASCerts.

⁶AS numbers used by IANA itself for experimental purposes are not counted.

Table III. psBGP Memory Requirements per AS

ASNumCerts	10.5 MB
SpeakerCerts	10.5 MB
PALs	17.8 MB
MultiASCerts	0.2 MB
AS Prefix Graph	3.8 MB
AS_PATH Digital Signatures	147.0 MB
Total	189.8 MB

6.3.1.3 AS Prefix Graph. Each AS must construct an AS prefix graph for prefix origin verification. Memory required for storing an AS prefix graph depends on the data structures being used. For simplicity, we use an adjacency list currently consisting of 17,844 entries, one entry per AS. Each entry consists of a 16-bit AS number and two 32-bit pointers, one pointing to a linked list of prefixes assigned to this AS and the other pointing to a linked list of neighboring ASes. On average, each prefix linked list has 10 elements with each of 17 bytes and each neighboring AS linked list has 5 elements with each of 6 bytes. Thus, each entry in the fixed array on average consumes 210 bytes. In total, an AS prefix graph requires 3.7 MB memory (MB = 10^6 bytes), using these (nonoptimized) data structures.

6.3.1.4 AS_PATH Digital Signatures. Each BGP speaker must store digital signatures for protecting the AS_PATH in a BGP update message received from a direct neighbor, which are estimated at 35 MB per neighbor [Kent 2003]. While many BGP speakers have a few neighbors, some at the Internet exchanges may have tens. Assuming an average of 4.2 neighbors, each BGP speaker would need 147 MB memory for storing AS_PATH digital signatures.

In summary, on average, a total of 189.8 MB of memory are required for storing all certificates, an AS prefix graph, and digitally signed BGP update messages to support psBGP (see Table III). While many BGP speakers may require less memory to support psBGP, some would require significantly more. We expect that routers will be equipped with more memory over time, thus mitigating the hurdle of memory overhead.

6.3.2 Bandwidth Overhead. Except for a small number of public key certificates of trusted CAs, which may be distributed using out-of-band mechanisms, all other certificates in psBGP can be distributed with BGP update messages. The latter consumes extra network bandwidth. However, such overhead is not persistent, since those certificates only need to be distributed periodically or upon changes. We expect that such overhead is of little significance and do not discuss it further.

The primary bandwidth overhead is introduced by digitally signed data and signatures carried by each BGP update message for protecting the message. For a fully protected BGP route where every AS on the route digitally signs the update message, the overhead is mainly determined by the number of such ASes and could result in as much as 600% overhead, according to Kent [2003]. We expect no significant difference between the bandwidth overhead of psBGP and

S-BGP. While increased bandwidth overhead resulting from psBGP (or e.g., S-BGP) is significant in terms of percentage, as pointed out by Kent [2003], BGP control messages only account for a small fraction of network bandwidth versus subscriber traffic. Thus, from our preliminary analysis, we expect that bandwidth overhead of psBGP will not create difficulty in the deployment of psBGP.

6.3.3 CPU Overhead. We expect that CPU overhead of psBGP will mainly result from AS_PATH verification, not AS prefix graph operations (cf. Section 6.2). A psBGP-enabled BGP speaker needs to digitally sign each BGP update message sent to each neighbor and to verify some digital signatures carried by each BGP update message it receives and chooses to use. As shown by Kent et al. [2000] in their study of S-BGP performance, such CPU overhead is significant. Especially in the case of reboots, a BGP speaker will receive full routing tables from each of its neighbors, and thus must verify a large number of digital signatures if psBGP is implemented. Note an AS prefix graph need not be rebuilt, since it can be stored in persistent storage and reloaded upon reboot. psBGP provides the flexibility for reducing the CPU overhead resulting from digital signature verification by using a lower confidence threshold, which trades off security for efficiency. In other words, psBGP provides a mechanism, which allows protection to be proportionally achieved in accordance to the CPU power which a router has available to spend on signature verification. However, to achieve higher level of assurance of AS_PATH integrity, significant CPU overhead will be generated by psBGP. To mitigate the problem, various approaches might be helpful, including caching [Kent et al. 2000], delay of signature verification [Kent et al. 2000], using a digital signature algorithm with a faster verification operation (e.g., RSA) [Nicol et al. 2004], and aggregated path authentication [Zhao et al. 2005a].

6.3.4 Certificate Dynamics. Here we discuss dynamics of certificates used in psBGP, including ASNumCerts, SpeakerCerts, and PALs.

6.3.4.1 ASNumCerts and SpeakerCerts. The monthly number of ASes on the Internet has grown by an average of 190 since January 1, 2004, with an average of 347 ASes added and 157 ASes removed (see Table II). When an AS number is added or removed in psBGP, the corresponding ASNumCert must be issued or revoked by an RIR. Thus, five RIRs between them must issue an average of 347 new ASNumCerts and revoke an average of 157 existing ASNumCerts per month. This appears quite manageable in light of substantially larger PKIs existing in practice (e.g., see Guida et al. [2004]). Note the issuing and revocation of a SpeakerCert is performed by an AS, not an RIR.

6.3.4.2 Prefix Assertion Lists (PALs). A prefix assertion list pal_i must be changed (removed, added, or updated) if: (a) the AS number s_i changes (i.e., is removed or added); (b) an IP prefix assigned to s_i changes; (c) s_i 's neighbor relationship changes, i.e., a neighbor is removed or added; or (d) an IP prefix changes which is endorsed by s_i for one of its neighbors. Table IV depicts the dynamics of prefix assignments.

Table IV. IP Prefix Dynamics from January 1 to August 1, 2004

	Jan	Feb	Mar	Apr	May	Jun	Jul
Start of month	148,903	148,014	151,174	156,019	157,925	160,818	155,118
Stable during month	143,200	144,422	146,139	151,481	153,171	148,280	151,436
Stable during Jan–Jul	119,968	119,968	119,968	119,968	119,968	119,968	119,968
Removed during month	5,703	3,592	5,035	4,538	4,754	12,538	3,682
Added during month	4,814	6,752	9,880	6,444	7,647	6,838	10,360

Table V. Projected Number of ASes in Absolute Number, and as Percentage of All ASes, Requiring the Specified Number of Monthly Prefix Assertion (PA) Changes in psBGP^a

# of PA Changes		1	2–4	5–10	11–30	31–100	101–1000	Over 1001	Total
$n = 1$	# of ASes (percentage)	1497 (8.3%)	677 (3.8%)	319 (1.8%)	152 (0.8%)	69 (0.3%)	26 (0.1%)	2 (0%)	2742 (15.2%)
$n = 2$	# of ASes (percentage)	1508 (8.4%)	713 (4.0%)	346 (1.9%)	187 (1.0%)	87 (0.5%)	48 (0.2%)	3 (0%)	2892 (16.0%)
$n = 3$	# of ASes (percentage)	1516 (8.4%)	725 (4.0%)	355 (2.0%)	205 (1.1%)	93 (0.5%)	54 (0.3%)	4 (0%)	2952 (16.4%)
$n = \text{all}$	# of ASes (percentage)	1424 (7.9%)	784 (4.3%)	387 (2.1%)	233 (1.3%)	112 (0.6%)	53 (0.3%)	30 (0.2%)	3023 (16.7%)

^aBased on July 2004 data. We recommend row $n = 2$ (n is the number of endorsing neighbors).

We study the number of prefix assertion (PA) changes required for each AS based on the two routing tables of July 1 and August 1, 2004. Each prefix addition or removal is counted once (i.e., resulting in one PA addition or removal) if the AS number of the AS owning that prefix does not change. If an AS number is newly added (or removed) during the month, all additions (or removals) of the prefixes owned by that AS are counted once, as a whole. One PA change usually represents one update to a *PAL* if such update is done in a timely manner. However, an AS can choose to do multiple PA changes in one *PAL* update (see Section 6.4 for more discussions).

Table V depicts the projected PA dynamics based on the data set of July 2004. The total number of ASes observed during July 2004 is 18,048, including 17,884 observed on August 1, 2004 and 164 removed during July 2004. We can see, the more ASes endorsing an AS's prefix assertions, the more PA changes required. We recommend the scenario $n = 2$, where each AS has at most two endorsing neighbors even if it has more than two neighbors. This provides a level of redundancy in the case that one of the two endorsing neighbors fails to carry out adequate due diligence.

From Table V, in the recommended scenario $n = 2$, 16% of the ASes need to update their *PALs* during the month. More specifically, 8.4% of ASes need only one PA change in the month, 4% need two to four PA changes, and 1.9% need five to ten PA changes. However, a small number of ASes need more than 100 changes and AS 701 (UUNET) and its two endorsing neighbors need around 5000 changes. We expect the number of PA changes will be lower, in practice, since some of prefix changes observed through the RouteViews dataset might only be temporary and result in no changes to *PALs*.

6.4 Discussion

The timeliness of *PAL* updates is important to ensure service availability. *PALs* need to be updated and distributed in a timely manner so that prefix ownerships can be verified using currently correct information. To ensure that an endorsing neighbor of a given AS updates its *PALs* for that AS in a timely manner, a service agreement between them would likely be required, e.g., as an extension to their existing agreements. Since there is usually some time delay window before newly delegated prefixes are actually used on the Internet, an endorsing AS should be required to update its *PAL* to include newly delegated prefixes of an endorsed neighbor within that delay window. Updates of prefix removals can be done with lower priority, since they would appear to have only relatively small security implications. *PALs* along with other certificates (e.g., ASNumCerts and SpeakerCerts) can be distributed with BGP update messages in the previously discussed new optional and transitive path attribute (see Section 4.1.2); thus, they can be distributed as fast as announcements of prefixes and are accessible without any dependence on BGP routes. Since those certificates are not route specific, new rules are required to determine how often a certificate will be included in an update message, e.g., on a daily basis or when a certificate is newly issued. Another approach is to store those certificates in centralized directories [Kent 2003], and to have each AS download them periodically, e.g., on a daily basis.

7. RELATED WORK

Considerable research has been published on securing routing protocols. Perlman [1988] was among the first to recognize and study the problem of securing routing infrastructures. Bellovin [1989] discussed security vulnerabilities of Internet routing protocols as early as 1989 (see also [Bellovin 2004]). More recently, Bellovin and Gansner [2003] discussed potential link-cutting attacks against Internet routing. Kumar and Crowcroft [1993] proposed the use of digital signatures and sequence numbers for protecting the integrity and freshness of routing updates. For a thorough analysis of BGP vulnerabilities and protections, see Murphy [2002a, 2002b].

The most complete and concrete security proposal to date for addressing BGP vulnerabilities is S-BGP [Kent et al. 2000; Seo et al. 2001]. It proposes the use of centralized PKIs for authenticating AS numbers and IP prefix ownership. S-BGP PKIs are rooted at RIRs and parallel to the existing system of AS number assignment and IP address allocation. An AS_PATH is protected using nested digital signatures; its integrity is guaranteed.

soBGP [White 2003] proposes the use of a web-of-trust model for AS public key authentication and a centralized hierarchical model for IP prefix ownership verification. AS_PATH is verified for plausibility by checking against an AS topology graph. Each AS issues certificates listing all neighboring ASes. A global AS graph can be constructed from those certificates. Thus, the existence of an AS_PATH can be verified if all ASes on the path have deployed soBGP. Table VI compares S-BGP, soBGP, and psBGP (recall Section 2.3 re: goals; also see Section 3.5 and Wan et al. [2007] for further background information).

Table VI. Comparison of S-BGP, soBGP, and psBGP—Achieving Security Goals of Section 2.3

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	Full integrity	Plausibility	Stepwise integrity

Goodell et al. [2003] proposed a protocol and architecture interdomain routing validator (IRV), for improving the security and accuracy of BGP. Each AS builds an IRV server, which has the interdomain routing information of that AS. One IRV can query another IRV for nonauthoritative routing information to verify BGP update messages received by its hosting AS. Improper prefix origination and AS_PATH might be detected by uncovering inconsistencies among responses from other IRVs. One advantage of IRV is that it supports incremental deployment requiring no changes to the existing routing infrastructure.

Kruegel et al. [2003] propose a model of AS topology augmented with physical Internet connectivity to detect and stop anomalous route announcements. Their approach passively monitors BGP control traffic and does not require modification to the existing routing infrastructure. Therefore, it would appear to be easy to deploy.

In a rigorous study of prefix origination authentication, Aiello et al. [2003] formalize the IP prefix delegation system, present a proof system, and propose efficient constructions for authenticating prefix origination. Real routing information is analyzed and used to reconstruct the IP delegation relationship over the Internet. They discover that the current prefix delegation on the Internet is relatively static and dense, however, they also note that it is extremely difficult, if not impossible, to determine this delegation structure.

Listen and Whisper [Subramanian et al. 2004] are proposed mechanisms for protecting the BGP data plane and control plane, respectively; they are best used together. The first approach (Listen) detects invalid data forwarding by detecting “incomplete” (as defined by Subramanian et al. [2004]) TCP connections. Whisper uncovers invalid routing announcements by detecting inconsistency among *path signatures* of multiple update messages, originating from a common AS but traversing different paths.

Hu et al. [2004] propose a Secure Path Vector (SPV) protocol for securing BGP. SPV makes use of efficient cryptographic primitives, e.g., authentication trees and one-way hash chains for protecting AS_PATH and is argued to be more efficient than S-BGP; however, it suffers a significant cost in increased memory (perhaps as much as three-to five fold).

8. CONCLUDING REMARKS

Beyond AS_PATH verification in Section 3.5, it is desirable to verify if an AS_PATH conforms to the route-exporting policies of each AS on the path.

Since BGP is a policy-driven routing protocol, each AS can individually decide whether or not a received route advertisement should be further propagated to a neighboring AS. Such route-exporting policies are mainly defined based on the business relationship with a neighboring AS. Without such verification, a misbehaving BGP speaker (e.g., misconfigured) may be able to readvertise routes which are prohibited by its route-exporting policies. For example, a multihomed AS may readvertise routes received from one provider AS to the other, thus functioning as a transit AS for its two providers. Such misbehavior may allow for eavesdropping and may also result in service disruption. New mechanisms for AS_PATH verification appear necessary.

Different approaches have been taken by S-BGP, soBGP, and IRV, among other proposals, for addressing security in BGP. We believe that psBGP adopts their best features, while differing fundamentally with a novel approach taken to verify IP prefix assignments and AS_PATH integrity. As no centralized infrastructure for tracing changes in IP prefix assignments currently exists and it would appear to be quite difficult to build such an infrastructure, we believe that the decentralized approach taken by psBGP provides a more feasible means of increasing confidence in correct prefix origin. We hope this work stimulates consideration of alternate design choices and trust models for securing BGP.

ACKNOWLEDGMENTS

We thank Stephen Chou, John Ioannidis, Steve Kent, Angelos Keromytis, Stefan Mink, and Peiter (Mudge) Zatkó for helpful discussion, as well as Steve Bellovin who also pointed out the collusion problem of multi-AS organizations and for motivating the proposal as described in Section 4.2. We also thank anonymous reviewers, whose comments considerably improved positioning and corrected numerous technical points. The first author is Canada Research Chair in Network and Software Security and is supported in part by MITACS, an NSERC Discovery Grant, and the Canada Research Chairs Program. The second author is supported in part by Alcatel Canada and MITACS. The third author is supported in part by MITACS and NSERC.

REFERENCES

- ADAMS, C. AND LLOYD, S. 2003. *Understanding Public-Key Infrastructure*, 2nd Ed. Addison-Wesley, Reading, MA.
- AIELLO, W., IOANNIDIS, J., AND MCDANIEL, P. 2003. Origin authentication in interdomain routing. In *Proceedings of the 10th ACM Conference on Computer and Communications Security*. Washington, D.C., 165–178.
- BARBIR, A., MURPHY, S., AND YANG, Y. 2004. Generic threats to routing protocols. Internet Draft.
- BELLOVIN, S. 1989. Security problems in the TCP/IP protocol suite. *Computer Communications Review*. 19, 32–48.
- BELLOVIN, S. 2004. A look back at “security problems in the TCP/IP protocol suite”. In *The 20th Annual Computer Security Applications Conference (ACSAC’04)*. Tucson, Arizona.
- BELLOVIN, S. AND GANSNER, E. 2003. Using link cuts to attack internet routing. Unpublished manuscript.
- BELLOVIN, S., IOANNIDIS, J., AND BUSH, R. 2005. Position paper: Operational requirements for secured BGP. In *DHS Secure Routing Workshop*.

- BONEH, D., BOYEN, X., AND SHACHAM, H. 2004. Short group signatures. In *Proceedings of Crypto 2004*. Vol. 3152. 41–55.
- BURROWS, M., ABADI, M., AND NEEDHAM, R. 1990. A logic of authentication. In *Research Report 39*. Digital Systems Research Center of Digital Equipment Corporation, Palo Alto, CA, 8, 1(Feb.), 18–36.
- DEMPSTER, A. 1967. Upper and lower probabilities induced by a multivalued mapping. *The Annals of Statistics* 28, 325–339.
- DHS. 2005. DHS secure routing workshop. Department of Homeland Security, Washington, D.C.
- GAARDER, K. AND SNEKKENES, E. 1991. Applying a formal analysis technique to the CCIT X.509 strong two-way authentication protocol. *Journal of Cryptology* 3, 81–98.
- GAO, L. 2000. Inferring autonomous system relationships in the Internet. In *IEEE Global Internet*.
- GLIGOR, V., KAILAR, R., STUBBLEBINE, S., AND GONG, L. 1991. Logics for cryptographic protocols—virtues and limitations. In *Proceedings of the Computer Security Foundations Workshop IV*. Los Alamitos, CA. 219–226.
- GOODELL, G., AIELLO, W., GRIFFIN, T., IOANNIDIS, J., MCDANIEL, P., AND RUBIN, A. 2003. Working around BGP: An incremental approach to improving security and accuracy of interdomain routing. In *Proceedings of the 2003 ISOC Symposium on Network and Distributed Systems Security (NDSS'03)*. San Diego, CA. 75–85.
- GUIDA, R., STAHL, R., BUNT, T., SECREST, G., AND MOORCONES, J. 2004. Deploying and using public key technology: Lessons learned in real life. *IEEE Security and Privacy (July/Aug.)*. 67–71.
- HEDRICK, C. 1988. Routing information protocol. IETF RFC 1058.
- HEFFERNAN, A. 1998. Protection of BGP sessions via the TCP MD5 signature option. IETF RFC 2385.
- HOUSLEY, R., FORD, W., POLK, W., AND SOLO, D. 1999. Internet X.509 public key infrastructure—certificate and CRL profile. IETF RFC 2459.
- HU, Y., PERRIG, A., AND SIRBU, M. 2004. SPV: Secure path vector routing for securing BGP. In *Proceedings of ACM 2004 SIGCOMM*. Portland, OR.
- IRR. 2005. Internet routing registry. <http://www.irr.net>.
- JUST, M., KRANAKIS, E., AND WAN, T. 2003. Resisting malicious packet dropping in wireless ad hoc networks. In *Proceedings of the 2nd Annual Conference on Adhoc Networks and Wireless (ADHOCNOW'03)*.
- KAUFMAN, C. 2005. The internet key exchange (IKEv2) protocol. IETF RFC 4306.
- KENT, S. 2003. Securing the border gateway protocol: A status update. In *Seventh IFIP TC-6 TC-11 Conference on Communications and Multimedia Security*.
- KENT, S. 2005. IP encapsulating security payload (ESP). IETF RFC 4303.
- KENT, S. 2006. An infrastructure supporting secure internet routing. In *Third European PKI Workshop*.
- KENT, S. AND ATKINSON, P. 1998a. Security architecture for the Internet protocol. IETF RFC 2401.
- KENT, S. AND ATKINSON, P. 1998b. IP encapsulating security payload (ESP). IETF RFC 2406.
- KENT, S., LYNN, C., MIKKELSON, J., AND SEO, K. 2000. Secure border gateway protocol (S-BGP) real world performance and deployment issues. In *Proceedings of the 2000 ISOC Symposium on Network and Distributed Systems Security (NDSS'00)*.
- KENT, S., LYNN, C., AND SEO, K. 2000. Secure border gateway protocol (S-BGP). *IEEE Journal on Selected Areas in Communications* 18, 4 (Apr.), 582–592.
- KOBLITZ, N. AND MENEZES, A. 2004. Another look at “provable security.” *Cryptology ePrint Archive*, Report 2004/152. To Appear in *Journal of Cryptology*. <http://eprint.iacr.org/2004/152/>.
- KRUEGEL, C., MUTZ, D., ROBERTSON, W., AND VALEUR, F. 2003. Topology-based detection of anomalous BGP messages. In *Proceedings of the 6th Symposium on Recent Advances in Intrusion Detection (RAID'03)*.
- KUMAR, B. AND CROWCROFT, J. 1993. Integrating security in interdomain routing protocols. *ACM SIGCOMM Computer Communication Review* 23, 5 (Oct.), 36–51.
- LYNN, C., KENT, S., AND SEO, K. 2003. X.509 Extensions for IP Addresses and AS Identifiers. draft-ietf-pkix-x509-ipaddr-as-extn-02.txt.

- MA, C., HU, N., AND LI, Y. 2006. On the release of CRLs in public key infrastructure. In *Proceeding of 15th USENIX Security Symposium*.
- MAURER, U. 1996. Modelling a public-key infrastructure. In *Proceedings of the 4th European Symposium on Research in Computer Security (ESORICS'96)*. 324–350.
- MURPHY S. 2002a. BGP Security Vulnerabilities Analysis. draft-murphy-bgp-vuln-00.txt.
- MURPHY S. 2002b. BGP Security Protections. draft-murphy-bgp-protect-00.txt.
- NICOL, D., SMITH, S., AND ZHAO, M. 2004. Evaluation of efficient security for BGP route announcements using parallel simulation. *Simulation Practice and Theory Journal, Special Issue on Modeling*, 187–216.
- PERLMAN, R. 1988. Network layer protocols with byzantine robustness. Tech. Rep. MIT/LCS/TR-429.
- REITER, M. AND STUBBLEBINE, S. 1997. Toward acceptable metrics of authentication. In *Proceedings of 1997 IEEE Symposium on Security and Privacy*. 10–20.
- REKHTER, Y. AND LI, T. 1995. A border gateway protocol 4 (BGP 4). IETF RFC 1771.
- RETANA, A. AND WHITE, R. 2002. BGP Custom decision process. draft-retana-bgp-custom-decision-00.txt.
- ROUTEVIEWS. 2005. Route views project. <http://www.routeviews.org>.
- SEO, K., LYNN, C., AND KENT, S. 2001. Public-key infrastructure for the secure border gateway protocol (S-BGP). In *IEEE DARPA Information Survivability Conference and Exposition II*.
- SHAFFER, G. 1976. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, NJ.
- SUBRAMANIAN, L., ROTH, V., STOICA, I., SHENKER, S., AND KATZ, R. 2004. Listen and whisper: Security mechanisms for BGP. In *Proceedings of the First Symposium on Networked Systems Design and Implementation (NSDI'04)*, San Francisco, CA.
- VILLAMIZAR, C., ALAETTINOGLU, C., MEYER, D., AND MURPHY, S. 1999. Routing policy system security. IETF RFC 2725.
- WAN, T. 2006. Securing routing protocols through information corroboration. Ph.D. thesis, Carleton University, Ottawa, Canada.
- WAN, T., KRANAKIS, E., AND VAN OORSCHOT, P. 2004. S-RIP: A secure distance vector routing protocol. In *Proceedings of the Applied Cryptography and Network Security (ACNS'04)*. Vol. 3089. 103–119.
- WAN, T., KRANAKIS, E., AND VAN OORSCHOT, P. 2005. Pretty secure BGP (psBGP). In *Proceedings of the 2005 ISOC Symposium on Network and Distributed Systems Security (NDSS'05)*. San Diego, CA.
- WAN, T., VAN OORSCHOT, P., AND KRANAKIS, E. 2007. A selective introduction to border gateway protocol (BGP) security issues. In *Proceedings of the NATO Advanced Studies Institute on Network Security and Intrusion Detection*. Nork, Yerevan, Armenia. IOS Press (to appear, 2007).
- WHITE, R. 2003. Securing BGP through secure origin BGP. *The Internet Protocol Journal* 6, 3, 15–22.
- WHITE, R., MCPHERSON, D., AND SANGLI, S. 2004. *Practical BGP*. Addison-Wesley, Reading, MA.
- ZHAO, M., SMITH, S., AND NICOL, D. 2005a. Aggregated path authentication for efficient BGP security. In *Proceedings of 12th ACM Conference on Computer and Communications Security*. Alexandria, VA.
- ZHAO, M., SMITH, S., AND NICOL, D. 2005b. Evaluating the performance impact of PKI on BGP security. In *Proceedings of 4th Annual PKI Research Workshop (PKI'05)*. Gaithersburg, MD.
- ZIMMERMANN, P. 1995. *The Official PGP User's Guide (second printing)*. MIT Press, Cambridge, MA.
- ZSAKO, J. 1999. PGP authentication for RIPE database updates. IETF RFC 2726.

Received September 2005; revised September 2006 and January 2007; accepted January 2007