

Pretty Secure BGP (psBGP)*

Tao Wan[†] Evangelos Kranakis[†] P.C. van Oorschot[†]

Abstract

The Border Gateway Protocol (BGP) is the de-facto standard inter-domain routing protocol on the Internet. However, it is well known that BGP is vulnerable to a variety of types of attacks, and that a single misconfigured or malicious BGP speaker could result in large scale service disruption. We first summarize a set of security goals for BGP, and then propose Pretty Secure BGP (psBGP) as a new security protocol achieving these goals. psBGP makes use of a centralized trust model for authenticating Autonomous System (AS) numbers, and a decentralized trust model for verifying the propriety of IP prefix origination. We compare psBGP with S-BGP and soBGP, the two leading security proposals for BGP. Our analysis suggests that psBGP provides a better balance between security and practicality than either S-BGP or soBGP: it significantly reduces the complexity of prefix ownership verification in S-BGP and soBGP, although in theory offering somewhat less security; and psBGP offers more security than soBGP in terms of AS number authentication and AS_PATH verification, albeit requiring expensive digital signature operations. Our performance analysis using real world BGP data suggests that psBGP is practical with respect to the number of certificates to be stored and to be updated per AS. We also raise a number of issues of independent interest about the design of S-BGP and soBGP.

1 Introduction and Motivation

The Internet consists of a number of Autonomous Systems (ASes), each of which consists of a number of routers under a single technical administration (e.g., sharing the same routing policy). The Border Gateway Protocol (BGP) [40] is the de facto standard inter-domain routing protocol for exchanging routing information between ASes on the Internet. It is well-known that BGP has many security vulnerabilities [29, 35], for example: AS numbers and BGP speakers (routers running BGP) can be spoofed; BGP update messages can be tampered with; and false BGP update messages can be spread. One serious problem is that a single misconfigured or malicious BGP speaker may poison the routing tables of many other well-behaved BGP speakers by advertising false routing information (e.g., see [9]). Examples of consequences include denial of service (i.e., legitimate user traffic cannot get to its ultimate destinations) and man-in-the-middle attacks (i.e., legitimate user traffic is forwarded through a router under the control of an adversary).

Many solutions [43, 29, 31, 19, 47, 2, 24] have been proposed for securing BGP. S-BGP [28, 29] is one of the earliest security proposals, and probably the most concrete one. S-BGP makes use of strict hierarchical public key infrastructures (PKIs) for both AS number authentication and IP prefix ownership verification (i.e., verifying which blocks of IP addresses are assigned or delegated to an AS). Besides computational costs, many people consider S-BGP to be impractical because of the viewpoint that requiring strict hierarchical PKIs makes it difficult to deploy across the Internet (e.g., [3]). Our viewpoint is slightly different and we consider that the two PKIs used in S-BGP have different practicalities, as explained below.

Agreeing in part with an important design decision made in S-BGP, we suggest that it is practical to build a centralized PKI for AS number authentication because: 1) the roots of the PKI are the natural trusted authorities for AS numbers, i.e., the Internet Assigned Number Authority (IANA) or the Internet Corporation of Assigned Numbers and Names (ICANN) and the Regional Internet Registries (RIRs), hereinafter IANA;

*Version: September 02, 2004.

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

and 2) the number of ASes on the Internet and its growth rate are relatively manageable, making PKI certificate management feasible. For example, based on the BGP data collected by the RouteViews project [34], there are in total about 17 884 ASes on the Internet as of August 1, 2004. This number has grown by an average of 190 (157 removed and 347 added) per month since January 1, 2004.

However, it would appear to be extremely difficult to build a centralized PKI for verifying IP prefix ownership given the complexity, if not impossibility, of tracing how existing IP address space is allocated, delegated, and tracing all changes of IP address ownership in part due to the large number of prefixes in use and frequent organization changes (e.g., corporations splitting, merging, bankruptcy, etc.). As pointed by Aiello et al. [2], it is exceptionally difficult to even approximate an IP address delegation graph for the Internet. Therefore, it may well be impossible to build a centralized PKI mirroring such a complex and unknown delegation structure. To quote from a study by Atkinson and Floyd [3] on behalf of the Internet Architecture Board (IAB): “*a recurring challenge with any form of inter-domain routing authentication is that there is no single completely accurate source of truth about which organizations have the authority to advertise which address blocks*”.

In contrast, soBGP [47] proposes use of a web-of-trust model for authenticating AS public keys and a hierarchical structure for verifying IP prefix ownership. While a web-of-trust model has strong proponents for authenticating user public keys within the technical PGP community, it is not clear if it is suitable for authenticating public keys of ASes which are identified by AS numbers strictly controlled by IANA; thus it is questionable if any entity other than IANA should be trusted for signing AS public key certificates. With respect to IP prefix ownership verification, soBGP makes use of a strictly hierarchical structure similar to that of S-BGP. Prefix delegation structures might be simplified in soBGP by using ASes instead of organizations as entities, however, it is not clear if there are difficulties in practice to do so since IP addresses are usually delegated to organizations not to ASes [2]. We suggest that soBGP, like S-BGP, faces extreme difficulty in tracing changes of IP address ownership in a strict hierarchical way. Thus, both S-BGP and soBGP have made architectural design choices which arguably lead to practical difficulties.

1.1 Our Contributions

In this paper, we present a new proposal for securing BGP, namely Pretty Secure BGP (psBGP), based on our analysis of the security and practicality of S-BGP and soBGP, and in essence, combining their best features. Our objective is to provide a reasonable balance between security and practicality. psBGP makes use of a centralized trust model for authenticating Autonomous System (AS) numbers, and a decentralized trust model for verifying the propriety of IP prefix origination which is in line with the recommendation of IAB [3]. Our analysis suggests that psBGP provides a better balance between security and practicality than either S-BGP or soBGP: it significantly reduces the complexity of S-BGP and soBGP in prefix ownership verification, although in theory offering somewhat less security; and it offers more security than soBGP in AS number authentication and AS_PATH (see §2.2) verification, albeit requiring expensive digital signature operations. Another advantage of psBGP is that it can successfully defend against threats from uncoordinated, misconfigured or malicious BGP speakers in a practical way. The major architectural highlights of psBGP are as follows (see §3 for other details and Table 4.4 for a summary comparison).

1) psBGP makes use of a *centralized trust model* for AS number authentication. Each AS obtains a public key certificate from one of a number of the trusted certificate authorities, e.g., RIRs, binding an AS number to a public key. We suggest that such a trust model provides perfect authorization of AS number allocation and best possible authenticity of AS public keys. Without such a guarantee, an attacker may be able to impersonate another AS to cause service disruption.

2) psBGP makes use of a *decentralized trust model* for verifying the propriety of IP prefix ownership.

Each AS creates a *prefix assertion list* consisting of a number of bindings of an AS number and prefixes, one for itself and one for each of its peering ASes. An assertion is *proper* if it is consistent among the prefix assertion lists of peering ASes. We consider this approach to be practical because it reflects existing AS peering relationships and common practices (e.g., ingress filtering [15]). In this way, we distribute the extremely difficult task of tracing IP address ownership across all ASes on the Internet, albeit introducing some additional security risk. Assuming reasonable due diligence in tracking IP address ownership of direct peer ASes, and assuming no two ASes in collusion, a single misbehaving AS originating improper prefixes will be detected because they will cause inconsistency with the prefix assertions made by its peering ASes.

The rest of the paper is organized as follows. Section 2 defines notation, discusses BGP threats, and summarizes BGP security goals. psBGP is presented in Section 3, and compared with S-BGP and soBGP in Section 4. Security and performance of psBGP are analyzed in Section 5 and 6 respectively. A brief review of related work is given in Section 7. We conclude in Section 8.

2 BGP Security Threats and Goals

Here we define notation, discuss BGP security threats, and summarize a number of security goals for BGP.

2.1 Notation

A and B denote entities (e.g., an organization, 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 governing the related entity making that assertion. We use the following notation:

\mathbb{S}, s_i	\mathbb{S} is the complete AS number space; currently $\mathbb{S} = \{1, \dots, 2^{16}\}$. s_i is an AS number; $s_i \in \mathbb{S}$.
\mathbb{P}, f_i	\mathbb{P} is a set of all possible IP address prefixes. f_i is an IP prefix; $f_i \subset \mathbb{P}$.
T	an authority of \mathbb{S} and \mathbb{P} , e.g., $T \in RIRs$.
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).
$peer(s_i)$	a set of ASes with which s_i establishes a BGP session on a regular basis. More specifically, a given AS s_i may have many BGP speakers, each of which may establish BGP sessions with speakers from many other ASes. $peer(s_i)$ is the set of all other such ASes.
$k_A, \overline{k_A}$	one of A's public and private key pairs.
$\{m\}_A$	digital signature on message m generated with A's private key $\overline{k_A}$.
$(k_A, A)_{k_B}$	a public key certificate binding k_A to A, signed by B using $\overline{k_B}$.
$(k_A, A)_B$	equivalent to $(k_A, A)_{k_B}$ when the signing key is not the main focus.
$(f_i, s_i)_A$	a prefix assertion made by A that s_i owns f_i .
f_i^A, f_i^B	possible different prefixes asserted by A and B related to a given AS.

2.2 BGP Security Threats

BGP faces threats 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 threats against BGP control messages without considering those against data traffic (e.g., malicious packet dropping). 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 (hereinafter *falsification* [4]) of BGP control messages; deletion, exposure and replaying are beyond the scope of this paper. Deletion appears indistinguishable from legitimate route filtering. Exposure might compromise confidentiality of BGP control messages, which may or may not be a major concern [4]. Replaying is a serious threat and can be handled by setting expiration time for a message, however it seems challenging to find an appropriate value for an expiration time.

There are four types of BGP control messages: OPEN, KEEPALIVE, NOTIFICATION, and UPDATE. The first three are used for establishing and maintaining BGP sessions with peers, and falsification of them will very likely result in session disruption. As mentioned by Hu et al. [24], they can be protected by a point-to-point authentication protocol, e.g., IPsec [26]. We concentrate on falsification of BGP UPDATE messages (hereinafter, often referred to simply as BGP messages) which carry inter-domain routing information and are used for building up routing tables.

A BGP UPDATE message consists of three parts: withdrawn routes, network layer reachability information (NLRI), and path attributes (e.g., AS_PATH, LOCAL_PREF, etc.). Due to space limitations, we omit discussion of how to protect withdrawn routes. NLRI consists of a set of IP prefixes sharing the same characteristics as described by the path attributes. NLRI is falsified if an AS originates a prefix not owned by that AS, or aggregated improperly from other routes. Examples of consequences include denial of service and man-in-the-middle attacks. There are two types of AS_PATH: AS_SEQUENCE or AS_SET. An AS_PATH of type AS_SEQUENCE consists of an ordered list of ASes traversed by this route. An AS_PATH of type AS_SET consists of an unordered list of ASes, sometimes created when multiple routes are aggregated. Due to space limitations, we focus on the security of AS_SEQUENCE in this paper. (Note AS_SET is less widely used on the Internet. For example, as of August 1, 2004, only 23 of 17884 ASes originated 47 of 161796 prefixes with AS_SET.) An AS_PATH is falsified if an AS or any other entity illegally operates on an AS_PATH, e.g., inserting a wrong AS number, deleting or modifying an AS number on the path, etc. Since AS_PATH is used for detecting routing loops and used by route selection processes, falsification of AS_PATH can result in routing loops or selecting routes not selected otherwise. We are interested in countering falsification of NLRI and AS_PATH. We assume there are multiple non-colluding misbehaving ASes and BGP speakers in the network, which may have legitimate cryptographic keying materials. This non-colluding assumption is also made by S-BGP and soBGP, explicitly or implicitly.

2.3 BGP Security Goals

We seek to design secure protocol extensions to BGP which can resist the threats as discussed above. 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. We summarize five security goals for BGP (cf. [28, 29]). G1 and G2 relate to data origin authentication, G3 to data integrity, and G4 and G5 to the propriety of BGP messages.

- G1. (*AS Number Authentication*) It must be verifiable that an entity that uses an AS number s_i as its own is in fact an authorized representative of the AS to which a recognized AS number authority assigned s_i .
- G2. (*BGP Speaker Authentication*) It must be verifiable that a BGP speaker, which asserts an association with an AS number s_i , has been authorized by the AS to which s_i was assigned by a recognized AS number authority.
- G3. (*Data Integrity*) It must be verifiable that a BGP message has not been illegally modified en route.

- G4. (*Prefix Origination Verification*) It must be verifiable that it is proper for an AS to originate an IP prefix. More specifically, it is proper for AS s_1 to originate prefix f_1 if 1) f_1 is owned by s_1 ; or 2) f_1 is aggregated from a set F of prefixes such that $f_1 \subseteq F$, i.e., $\forall f_x \subseteq f_1, f_x \subseteq F^1$.
- G5. (*AS Path Verification*) It must be verifiable that an AS_PATH ($p_k = [s_1, s_2, \dots, s_k]$) of a BGP route m consists of a sequence of ASes actually traversed by m in the specified order, i.e., m originates from s_1 , and has traversed through s_2, \dots, s_k in order.

3 Pretty Secure BGP (psBGP)

psBGP makes use of a centralized trust model for authenticating AS numbers and AS public keys. RIRs are the root trusted certificate authorities. Each AS s is issued a public key certificate (ASNumCert), signed by one of the RIRs, denoted by $(k_s, s)_T$. An AS with an ASNumCert $(k_s, s)_T$ creates and signs two data structures: a SpeakerCert $(k'_s, s)_{k_s}$ binding a public key k'_s to s ; and a prefix assertion list (PAL), listing prefix assertions made by s about the prefix ownership of s and s 's peers. PAL_s is an ordered list: the first assertion is for s itself and the rest are for each of s 's peers ordered by AS number. Figure 1 illustrates the certificate structure used in psBGP. We next describe psBGP with respect to five security goals, corresponding to G1-G5 above.

3.1 AS Number Authentication in psBGP

Following S-BGP and soBGP, we make use of a centralized PKI [42] for AS number authentication, with four root Certificate Authorities (CAs), corresponding to the four existing RIRs. When an organization B applies for an AS number, besides supplying documents currently required (e.g., routing policy, peering ASes, etc.), B additionally supplies a public key, and should be required to prove the possession of the corresponding private key [42, 1]. When an AS number is granted to B by an RIR, a 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 $(k_s, s)_T$, binding s to a public key k_s signed by one of the RIRs. The proposed PKI for authenticating 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 probably 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, though admittedly non-trivial (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 §6 - Table 2). Considering there are four RIRs, the overhead of managing ASNumCerts should certainly be feasible as large PKIs are currently commercially operational [21].

To verify the authenticity of an ASNumCert, an AS must have the trusted public key (or certificate) of the signing RIR. These few root trusted public key certificates can be distributed using *out-of-band* mechanisms. ASNumCerts can be distributed with BGP UPDATE messages. An ASNumCert is revoked when the corresponding AS number is not used or reassigned to another organization. Issues of revocation, though

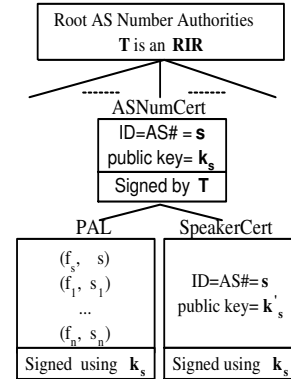


Figure 1: psBGP Certificate Structure

¹If s_1 does not own f_1 and $\exists f_x \subseteq f_1$ such that $f_x \not\subseteq F$, then s_1 *overclaims* IP prefixes, which is considered to be a type of falsification.

extremely important, are beyond the scope of the present paper; we restrict comment to the observation that revocation is a well-studied issue, if albeit still challenging (e.g., see [1]). So far, 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.

There is much debate on the architecture for authenticating the public keys of ASes in the BGP security community, particularly on the pros and cons of using a strict hierarchical trust model vs. a distributed trust model, e.g., a web-of-trust model. We make the use of a strict hierarchical trust model with depth of one for authenticating AS numbers and their public keys since it is not clear if a web-of-trust model is suitable here. Some of the issues with a web-of-trust model are discussed in Appendix 1, e.g., *trust bootstrapping*, *trust transitivity*, *vulnerability to a single misbehaving party* [33, 41].

3.2 BGP Speaker Authentication in psBGP

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 peer relationship with another AS. In psBGP, an AS with a certified ASNumCert issues an operational 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: 1) each BGP speaker is issued a unique public key certificate; 2) group signatures (e.g., see [7]) are used, i.e., each BGP speaker has a unique private key but shares a common public key certificate with other speakers in the same AS; or 3) all BGP speakers in a given AS share a common public-private key pair. We propose the latter for its simplicity and practicality. Choice 1) provides stronger security but requires more certificates, and discloses BGP speaker identities. Such disclosure may introduce new competitive security concerns [46]. Choice 2) provides stronger security, requires the same number of certificates, and does not disclose BGP identities, but involves a more complex system.

The private key corresponding to the public key of a SpeakerCert is used for establishing secure connections with peers (§3.3), and for signing BGP messages. Therefore, it must 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 and a PAL, it need not be stored in a BGP speaker. Thus, compromising a BGP speaker only discloses the private key of a SpeakerCert, requiring revocation and reissuing of a SpeakerCert, without impact on an ASNumCert. This separation of ASNumCerts from SpeakerCerts provides a more conservative design (from a security viewpoint), and distributes from RIRs to ASes the workload of certificate revocation and reissuing resulting from BGP speaker compromises. In summary, an ASNumCert must be revoked if the corresponding AS number is re-assigned or the corresponding key is compromised. A SpeakerCert must be revoked if a BGP speaker in that AS is compromised, or if that AS decides for other reasons to reissue it (e.g., if the private key is lost).

3.3 Data Integrity in psBGP

To protect data integrity, BGP sessions between peers must be protected. Following S-BGP, psBGP uses IPsec Encapsulating Security Payload (ESP) [27] with null encryption for protecting BGP sessions. Since many existing BGP speakers implement TCP MD5 [22] with manual key configurations for protecting BGP sessions, it must be supported by psBGP as well. In psBGP, automatic key management techniques can be implemented to improve the security of TCP MD5 as each BGP speaker has a public-private key pair

(common to all speakers in that AS).

3.4 Verification of Prefix Origination in psBGP

When an AS s_i originates a BGP UPDATE message $m = (f, [s_i, \dots])$, another AS needs to verify if it is proper for s_i to originate a route for a prefix f . As stated in §2.3 (G4), it is proper for s_i to originate a route for prefix f if: 1) s_i owns f ; or 2) s_i aggregates f properly from a set F of prefixes carried by a set of routes s_i has received.

3.4.1 Verification of Prefix Ownership in psBGP

Facing the extreme difficulty of building an IP address delegation infrastructure (recall §1), we propose a *decentralized* approach for verifying the propriety of IP address ownership, and more specifically by using *consistency checks*. Our approach is inspired by the way humans acquire their trust in the absence of a trusted authority: by corroborating information from multiple sources.

In psBGP, each AS s_i creates and signs a *prefix assertion list* (PAL_{s_i}), consisting of a number of tuples of the form (IP prefix list, AS number), i.e., $PAL_{s_i} = [(f_i^{s_i}, s_i), (f_1^{s_i}, s_1), \dots, (f_n^{s_i}, s_n)]$, where $\forall 1 \leq j \neq i \leq n, s_j \in peer(s_i)$ and $s_j < s_{j+1}$. The first tuple $(f_i^{s_i}, s_i)$ asserts that s_i owns $f_i^{s_i}$; the rest are sorted by AS number, and assert the prefix ownership of s_i 's peers. $(f_j^{s_i}, s_j)$ ($s_j \neq s_i$) asserts by s_i that s_j is a peer of s_i and s_j owns prefix $f_j^{s_i}$ if $f_j^{s_i} \neq \emptyset$. Otherwise, it simply asserts that s_j is a peer of s_i .

As a new requirement in psBGP, each AS is responsible for carrying out some level of due diligence offline: for the safety of that AS and of the whole Internet, to determine what IP prefixes are delegated to each of its peers. We suggest the effort required for this is both justifiable and practical, since two peering ASes usually have a business relationship (e.g., a traffic agreement) with each other, allowing offline direct interactions. For example, s_i may ask each of its peer s_j to show the proof that f_j is in fact owned by s_j . Similar due diligence might have been taken by service providers for implementing ingress filtering [15] on the Internet. Publicly available information about IP address delegation may also be helpful.

Two assertions $(f_i, s_i), (f'_i, s'_i)$ made by two ASes are *comparable* if they assert the prefix ownership of a given AS, i.e., $s_i = s'_i$ and the asserted prefixes are non-empty, i.e., $f_i, f'_i \neq \emptyset$; and are *incomparable* otherwise, i.e., they assert the prefix ownership of different ASes or one of the asserted prefixes is an empty set. Two comparable assertions (f_i, s_i) and (f'_i, s_i) are *consistent* if $f_i = f'_i$; and are *inconsistent* if $f_i \neq f'_i$.

Let n be the number of s_i 's peers. (f_i, s_i) is *k-proper* if there exist some fixed number k ($2 \leq k \leq n+1$) of consistent assertions of (f_i, s_i) made by s_i or s_i 's peers. Requiring $k = n+1$ means that the assertion (f_i, s_i) made by s_i and all of its peers must be consistent for (f_i, s_i) to be proper; this provides maximum confidence in the correctness of (f_i, s_i) if the condition is met. However, it is subject to attacks by a single misbehaving AS. For example, if $\exists s_j \in peer(s_i)$, and s_j makes a false assertion $(f_i^{s_j}, s_i)$ inconsistent with $(f_i^{s_i}, s_i)$, then $(f_i^{s_i}, s_i)$ will not be verified as proper, or will be verified as *improper*, although it might indeed be proper. From the perspective of assertion list management, the greater k is, the larger prefix assertion lists will grow, and the more updates of prefix assertion lists will be required since a change to an AS number s_i or a prefix f_i requires the update of all PALs making an assertion about s_i or f_i . Moreover, there are a large number of ASes which have only one peer. For example, as of August 1, 2004, there were 6619 ASes which have only one peer based on one BGP routing table collected from the RouteViews project [34]. Requiring $k \geq 3$ will prevent these ASes from originating authorized prefixes.

To begin with, we suggest $k = 2$ in psBGP, i.e., $(f_i^{s_i}, s_i)$ is *proper* if there exists any single $s_j \in peer(s_i)$ such that s_j make an assertion $(f_i^{s_j}, s_i)$ which is consistent with $(f_i^{s_i}, s_i)$. When verifying $(f_i^{s_i}, s_i)$, an AS checks its consistency with the prefix assertion related to s_i made by each of s_i 's peers until a consistent one is found, or no consistent assertion is found after all relevant assertions made by s_i 's peers have been

checked. In the former case, $(f_i^{s_i}, s_i)$ is verified as proper; in the latter case, it is verified as improper. For simplicity, the consistency among the prefix assertions related to s_i made by s_i 's peers amongst themselves is not checked. A non-aggregated route $(f, [s_i, \dots])$ originated by s_i is verified as proper if $(f_i^{s_i}, s_i)$ is proper and $f \subseteq f_i^{s_i}$.

psBGP assumes that no two ASes are in collusion. AS s_i and s_j are said in collusion if they make factually false but consistent assertions related to s_i 's prefix ownership. Note that a false prefix assertion made by s_j about a remote AS s_k , i.e., $s_j \notin \text{peer}(s_k)$, will not be checked when the self prefix ownership assertion by s_k is verified. Thus, a misbehaving AS s_j is only able to cause inconsistency with the own prefix ownership assertion by one of s_j 's peers. If $\forall s_j \in \text{peer}(s_i)$, s_j issues $(f_i^{s_j}, s_i)$ inconsistent with $(f_i^{s_i}, s_i)$, $(f_i^{s_i}, s_i)$ will be verified as *improper* by other ASes, even if it might be actually proper. This is the case when misbehaving ASes form a network cut from s_i to any part of the network. It appears impossible to counter such an attack, and many other techniques can also be used to deny the routing service of s_i , e.g., link-cuts [6], filtering, or packet dropping. Note that an attacker in control of a BGP speaker in AS s_j is unable to issue valid false prefix assertions if the private key of s_j 's ASNumCert is not compromised.

3.4.2 Verification of Aggregated Prefixes

Suppose s_i owns IP prefix f_i . When receiving a set of routes with a set of prefixes $F = \{f_j\}$, the BGP specification [40] allows s_i to aggregate F into a prefix f_g to reduce routing information to be stored and transmitted. We call f_j a prefix to be aggregated, and f_g an aggregated prefix. s_i can aggregate F into f_g if one of the following conditions holds: 1) $\forall f_j \subseteq f_g, f_j \subseteq f_i$; or 2) $\forall f_j \subseteq f_g, f_j \subseteq F$.

In case 1), s_i must own f_i which is a superset of the aggregated prefix f_g . Most likely, f_i will be the aggregated prefix, i.e., $f_g = f_i$. This type of aggregation is sometimes referred to as prefix *re-origination*. From a routing perspective, prefix re-origination does not have any effect since traffic destined to a more specific prefix will be forwarded to the re-originating AS and then be forwarded to the ultimate destination from there. From a policy enforcement perspective, prefix re-origination 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 re-origination is no different than normal prefix origination since the aggregated prefix is either the same as, or a subset of, the prefix owned by the aggregating AS. Therefore, the aggregated route f_g can be verified by cross-checking the consistency of s_i 's prefix assertion list with those of its peers (§3.4.1).

In case 2), s_i does not own the aggregated prefix f_g . Therefore, f_g cannot be verified in the same way as for prefix re-origination. To facilitate verification of the propriety of route aggregation by a receiving AS, psBGP requires that the routes to be aggregated 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.

3.5 Verification of AS_PATH in psBGP

There is no consensus on the definition of ‘‘AS_PATH security’’, and 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, AS_PATH security 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).

Since AS_PATH is used by the BGP route selection process, great assurance of the integrity of an AS_PATH increases the probability that routes are selected based on proper information. While the BGP specification [40] does not explicitly state that AS_PATH is used for route selection, it commonly is in practice (e.g., by Cisco IOS). Without the guarantee of AS_PATH integrity, an attacker may be able to modify an AS_PATH in a such way that it is 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 the falsified route over any correct routes. As a result, traffic flow might be influenced. Thus, we believe that it is insufficient to verify only the existence/non-existence of an AS_PATH, and it is desirable to obtain greater assurance of the integrity of an AS_PATH; we acknowledge that the cost of any solution should be taken into account as well. In what follows, we define AS_PATH security according to the original definition of AS_PATH [40], as “an ordered set of ASes a route in the UPDATE message has traversed”.

We choose the S-BGP approach with the improvement of the bit-vector method by Nicol et al. [37] (see next paragraph) for securing AS_PATH in psBGP, since it fits into the design of psBGP and provides greater assurance of AS_PATH integrity with reasonable overhead. Hu et al. [24] propose a secure path vector protocol (SPV) for protecting AS_PATH using authentication hash trees with less overhead than S-BGP. psBGP does not use the SPV approach since it has different assumptions than psBGP. For example, SPV uses different public key certificates than psBGP.

Let $n_i = |\text{peers}(s_i)|$ be the number of peers of s_i . Given $m_k = (f_1, [s_1, s_2, \dots, s_k])$, a psBGP speaker s_i ($1 \leq i \leq k - 1$) generates a digital signature $\{f_1, [s_1, \dots, s_i], v_i[n_i]\}_{s_i}$ where $v_i[n_i]$ is a bit vector of bit-length n_i , with one bit corresponding to each peer in s_i 's prefix assertion list (§3.4.1). If s_i intends to send a routing update to a peer s_j , it sets the bit in $v_i[\]$ corresponding to s_j . In this way, a message sent to multiple peers by a BGP speaker need be signed only once. For s_{k+1} to accept m_k , s_{k+1} must receive the following digital signatures: $\{f_1, [s_1], v_1[n_1]\}_{s_1}, \{f_1, [s_1, s_2], v_2[n_2]\}_{s_2}, \dots, \{f_1, [s_1, s_2, \dots, s_k], v_k[n_k]\}_{s_k}$.

4 Comparison of S-BGP, soBGP, and psBGP

We compare the different approaches taken by S-BGP, soBGP, and psBGP for achieving the BGP security goals listed in §2.3. Table 4.4 provides a summary. We see that psBGP falls somewhere between S-BGP and soBGP in several of the security approaches and architectural design decisions, but makes distinct design choices in several others.

4.1 AS Number Authentication

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

certificates $(k_X, X)_T, (k_s, s)_{k_X}$ might be revoked, and new certificates $(k_Y, Y)_T, (k'_s, s)_{k_Y}$ might be issued.

4.2 BGP Speaker Authentication

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

Data Integrity

S-BGP use IPsec for protecting BGP session and data integrity. Both soBGP and psBGP adopt this approach. TCP MD5 [22] is supported by all three proposals for backward compatibility. In addition, automatic key management mechanisms can be implemented for improving the security of TCP MD5.

4.3 Prefix Origination Verification

Both S-BGP and soBGP propose a hierarchical structure for authorization of the IP address space; although S-BGP traces how IP addresses are delegated among organizations, while soBGP only verifies IP address delegation among ASes. It appears that soBGP simplifies the delegation structure and requires fewer certificates for verification; however, it is not clear if it is feasible to do so in practice since IP addresses are usually delegated between organizations, not ASes. In psBGP, consistency checks of PALs of direct peers are performed to verify if it is proper for an AS to originate an IP prefix. Therefore, psBGP does not involve verification of chains of certificates (instead relying on offline due diligence). We note that while psBGP does not guarantee perfect security of the authorization of IP address allocation or delegation, as intended by S-BGP and soBGP, as discussed in (§1), it is not clear if the design intent in the latter two can actually be met in practice.

4.4 AS_PATH Verification

Both S-BGP and psBGP verify the integrity of AS_PATH based on its definition in the BGP specification [40]. In contrast, soBGP verifies the plausibility of an AS_PATH. Thus, S-BGP and psBGP provide stronger security of AS_PATH than soBGP, at the cost of digital signature operations which might slow down network convergence.

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 BGP speaker	one certificate per AS
G3: Data Integrity	IPsec or TCP MD5	IPsec or TCP MD5	IPsec or TCP MD5
G4: Prefix Origination Verification	centralized (multiple levels)	centralized (multiple levels)	decentralized (no trust transitivity)
G5: AS_PATH Verification	integrity	plausibility	integrity

Table 1: Comparison of S-BGP, soBGP, and psBGP approaches for achieving BGP security goals.

5 Security Analysis of psBGP

We analyze psBGP against the listed security goals from §2.3. The analysis below clarifies how our proposed 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 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 encourage such formalized reasoning for lack of better alternatives.

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

Proof Outline: For an AS number s to be certified, psBGP requires an ASNumCert $(k_s, s)_T$. Since T controls s , and is the trusted guardian of AS numbers (by assumption), any assertion made by T about s is proper. Thus $(k_s, s)_T$ 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 formally², $(T \text{ controls } s) \wedge (k_s, s)_T \Rightarrow (k_s, s)$ is proper.

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

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

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

Proof Outline: psBGP uses the IPsec Encapsulating Security Payload (ESP) [26, 27] with null encryption for protecting BGP sessions, and relies upon IPsec ESP for data integrity.

Before presenting Proposition 4, we establish two Lemmas.

Lemma 1 *Assume that $\forall s_i \in \mathbb{S}, \exists s_j \in \text{peer}(s_i)$ such that s_j carries out reasonable due diligence to create a proper prefix assertion $(f_i^{s_j}, s_i)$ (A1); and that no two ASes are in collusion (A2), then psBGP provides reasonable assurance of prefix ownership verification, i.e., a prefix assertion $(f_i^{s_i}, s_i)$ that is actually proper will be verified as such; otherwise not.*

Proof Outline: Suppose $(f_i^{s_i}, s_i)$ is proper. Since $\exists s_j \in \text{peer}(s_i)$ which makes a proper assertion $(f_i^{s_j}, s_i)$ (by assumption A1), then $(f_i^{s_i}, s_i)$ is consistent with $(f_i^{s_j}, s_i)$ since two proper assertions must be consistent. Thus, $(f_i^{s_i}, s_i)$ will be verified as proper because there exists a prefix assertion from s_i 's peer s_j , $(f_i^{s_j}, s_i)$, which is consistent with $(f_i^{s_i}, s_i)$.

Suppose $(f_i^{s_i}, s_i)$ is improper. To show that $(f_i^{s_i}, s_i)$ will not be verified as proper, we need to show that there does not exist $(f_i^{s_j}, s_i)$, $s_j \in \text{peer}(s_i)$, such that $(f_i^{s_j}, s_i)$ is consistent with $(f_i^{s_i}, s_i)$. $\forall (f_i^{s_j}, s_i), s_j \in \text{peer}(s_i)$, if s_j carries out due diligence successfully, then $(f_i^{s_j}, s_i)$ is proper and will be inconsistent with the improper $(f_i^{s_i}, s_i)$. If s_j misbehaves or its due diligence fails to reflect actual IP ownership, then $(f_i^{s_j}, s_i)$ is

²Here we adapt BAN-like notation, modified for our purpose (cf. [8, 16, 18]).

improper. We consider it to be a collusion of s_j and s_i if $(f_i^{s_j}, s_i)$ and $(f_i^{s_i}, s_i)$ are improper but consistent. This case is ruled out by assumption A2. Thus, an improper prefix assertion $(f_i^{s_i}, s_i)$ will be verified as improper since there does not exist an improper assertion which is consistent with $(f_i^{s_i}, s_i)$ without collusion. This establishes Lemma 1.

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 5). If $f_g \subseteq \cup\{f_i\}$, then s_x aggregates f_g properly. If s_x cannot provide required evidence, s_x 's aggregation of f_g is verified as improper. This establishes Lemma 2.

Proposition 4 *psBGP provides reasonable assurance of IP prefix origination verification, i.e., an AS s_i 's origination of a prefix f is verified as proper if f is owned by s_i or is aggregated properly by s_i from a set of routes received by s_i . Otherwise, s_i 's origination of f is verified as improper.*

Proof Outline: Lemma 1 allows verification of the propriety of prefix ownership. Suppose $(f_i^{s_i}, s_i)$ is verified as proper, i.e., $f_i^{s_i}$ is verified to be owned by s_i . If s_i owns f , then $f \subseteq f_i^{s_i}$. In psBGP, s_i 's origination of f is verified as proper if $f \subseteq f_i^{s_i}$. If $f \not\subseteq f_i^{s_i}$, psBGP requires that s_i provide proof that f is aggregated properly from a set of received routes (see Lemma 2). If s_i does not own f and s_i does not provide proof of the propriety of prefix aggregation, psBGP verifies s_i 's origination of f as improper. This establishes Proposition 4.

Proposition 5 *psBGP provides assurance of AS_PATH verification (G5).*

Proof Outline: Let $m_k = (f_1, p_k)$ be a BGP route, where $p_k = [s_1, s_2, \dots, s_k]$. Let r_i ($1 \leq i \leq k-1$) be a BGP speaker in s_i which has originated ($i = 1$) or forwarded ($2 \leq i \leq k-1$) m_i to s_{i+1} . In psBGP, the integrity of p_k implies that m_k has traversed the exact sequence of s_1, s_2, \dots, s_n . In other words, there does not exist i ($2 \leq i \leq n-1$) such that s_{i-1} didn't send $(f_1, [s_1, \dots, s_{i-1}])$ to s_i .

By way of contradiction, assume that it is possible in psBGP that $(f_1, [s_1, \dots, s_k])$ is accepted by a BGP speaker r_{k+1} and there exists i ($2 \leq i < k$) such that s_{i-1} didn't send $(f_1, [s_1, \dots, s_{i-1}])$ to s_i . psBGP requires that for $[s_1, s_2, \dots, s_k]$ to be accepted by r_{k+1} , $\forall i$ ($1 \leq i < k$), r_{i+1} has received a valid digital signature $\{p_1, [s_1, \dots, s_i], v_i[\]\}_{s_i}$ where the bit in $v_i[\]$ corresponding to s_{i+1} is set. $\{p_1, [s_1, \dots, s_i], v_i[\]\}_{s_i}$ serves as a signed assertion that s_i does send that routing update to s_{i+1} . This contradicts the above assumption. Thus, Proposition 5 is established.

The above results establish the desired psBGP security properties, and are summarized by Theorem 1.

Theorem 1 (psBGP Security Property) *psBGP achieves the following five security goals: AS number authentication (G1), BGP speaker authentication (G2), data integrity (G3), IP prefix origination verification (G4), and AS_PATH verification (G5).*

6 Performance Analysis of psBGP

We use BGP data collected by the RouteViews project [34] to estimate the number of ASNumCerts, SpeakerCerts, and PALs that are required by psBGP on the Internet, and their monthly changes. We retrieved one

BGP routing table of the first day of each month from January to August 2004. Despite likely incompleteness of the RouteViews data set, it is one of most complete data repositories publicly available, and has been widely used in the BGP community. We present our preliminary results of performance analysis for psBGP, specifically, the stability of certificate structures used in psBGP.

6.1 ASNumCerts and SpeakerCerts

We observed in total 17 884³ ASes 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. The same holds for SpeakerCerts. However, more efficient certificate distribution mechanisms (e.g., see [1, 30]) may be used; further discussion is beyond the scope of present paper.

The monthly number of ASes has grown an average of 190 since January 1, 2004, with an average of 347 ASes added and 157 ASes removed (see Table 2). When an AS number is added or removed, the corresponding ASNumCert must be issued or revoked by an RIR. Thus, four RIRs between them must issue an average of 347 new ASNumCerts and revoke an average of 157 existing ASNumCerts per month. This would certainly appear to be manageable in light of substantially larger PKIs existing in practice (e.g., see [21]). Note the issuing and revocation of a SpeakerCert is performed by an AS, not an RIR.

	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

Table 2: AS Number Dynamics from January 1 to August 1, 2004

6.2 Prefix Assertion Lists (PALs)

Each AS issues a PAL, which might be large if the number of peers or the number of prefixes assigned to a peer is large. To be distributed with BGP UPDATE messages whose size is limited to 4096 bytes, a large PAL must be split into multiple smaller ones. For simplicity, we consider one PAL per AS in our analysis. Thus, in the worst case, each AS needs to store 17 884 PALs.

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug
Start of Month	148 903	148 014	151 174	156 019	157 925	160 818	155 118	161 796
Stable During Month	143 200	144 422	146 139	151 481	153 171	148 280	151 436	N/A
Stable During Jan-Aug	119 968	119 968	119 968	119 968	119 968	119 968	119 968	N/A
Removed During Month	5 703	3 592	5 035	4 538	4 754	12 538	3 682	N/A
Added During Month	4 814	6 752	9 880	6 444	7 647	6 838	10 360	N/A

Table 3: IP Prefix Dynamics from January 1st to August 1st, 2004

A prefix assertion list PAL_{s_i} must be changed (removed, added, or updated) if: 1) the AS number s_i changes (i.e., removed or added); 2) an IP prefix owned by s_i changes; 3) s_i 's peer relationship changes, i.e., a peer is removed or added; or 4) an IP prefix changes which is owned by one of s_i 's asserted peers (i.e., a peer whose prefix ownership is asserted by s_i). Table 3 depicts the dynamics of prefixes, Figures 2-(a) and (b) illustrate AS peer relationships and AS prefix delegation, respectively, based on July 2004 data.

³AS numbers used by IANA itself are not counted.




Figure 2: AS Peer Relationships and Prefix Delegation

We study the number of changes of prefix assertions (PAs) 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.

In §6.2.1, we present the projected PA additions, removals, and the combined PA changes for ASes as a result of the changes of their own AS numbers or IP prefixes. In §6.2.2, we present the projected PA additions, removals, and the combined PA changes for ASes as a result of the changes of their peers' AS numbers or prefixes. We separate PA additions from removals because we consider that PA additions should be performed with high priority to minimize service outage, while PA removals can be performed with low priority without impact. In §6.2.3, we present the projected PA changes of all ASes.


6.2.1 ASes Changing their Own AS Numbers or Prefixes

We count the number of prefix assertion changes for each AS as a result of addition/removal of its own AS number or prefixes, as described above. We then count the number of ASes with a given number of PA additions, removals, and the combined changes respectively. We plot the number of PA additions versus the number of ASes with that number of PA additions in Figure 3(a), and the same for PA removals in Figure 3(b). The combined PA changes versus the number of ASes with those specified PA changes are illustrated in Figure 4. Note that in Figure 4, there is one AS which needs 4 936 PA changes. This AS (701) added 4 924 prefixes and removed 12 prefix during the month.


6.2.2 ASes whose Peers Changing their AS Numbers or Prefixes

Here we project the number of PA changes for ASes as a result of their asserted peers changing their AS numbers or prefixes. Let $t \geq 1$ be the exact number of peers for a given AS s_i , let n be our desired number of peers asserting prefix ownerships for s_i , and let m be the actual number of asserting peers of s_i which we will choose in the AS topology graph for our analysis. If $t \geq n$, we set $m = n$ since we desire n asserting peers and this number is possible. Otherwise, set $m = t$ since only t peers are available to make assertions.

We study four scenarios based on a given AS topology derived from the July 2004 dataset, and a desired value of n . ($n = 1$): for each AS, there is exactly $m = 1$ peer asserting prefix ownerships for that AS; ($n = 2$): for each AS, there are $m = 2$ peers asserting prefix ownerships for that AS if it has two or more peers, otherwise, set $m = 1$; ($n = 3$): for each AS, there are $m = 3$ peers asserting prefix ownerships for that AS if it has three or more peers; otherwise, set $m = t$; ($n = all$): for each AS, all $m = t$ of its peers



(a)



(b)

Figure 3: Projected PA additions and removals for ASes as a result of changing their own AS numbers or prefixes, based on July 2004 data. In (a), point $(x=8, y=28)$ indicates that 28 ASes need 8 PA additions.







Figure 4: Projected PA changes for ASes as a result of changing their own AS numbers or prefixes, based on July 2004 data




(a) $n=1$



(b) $n=2$




(c) $n=3$




(d) $n=all$


Figure 5: Projected PA additions for ASes as a result of their asserted peers adding prefixes, or newly appearing, based on July 2004 data.




(a) $n=1$



(b) $n=2$




(c) $n=3$




(d) $n=all$


Figure 6: Projected PA removals of ASes as a result of their asserted peers removing their AS numbers or prefixes, based on July 2004 data.




(a) $n=1$



(b) $n=2$



(c) $n=3$



(d) $n=all$

Figure 7: Projected PA removals of ASes as a result of their asserted peers changing their AS numbers or prefixes, based on July 2004 data.

assert prefix ownerships for that AS. According to these scenarios, Figures 5, 6, 7 illustrate PA additions, removals and the combined changes for ASes whose asserted peers change their AS numbers or prefixes during the month. We can see that more ASes require more prefix assertion changes as n increases (i.e., more asserting peers are desired).

6.2.3 Prefix Assertion List Stability

Table 4 depicts the projected PAL dynamics based on the data set of July 2004. The total number of ASes observed during July 2004 is 18048, including 17884 ASes observed on August 1, 2004 and 164 removed during July 2004. We can see that the more peers asserting the prefix ownership of other ASes, the more PA changes required. We recommend the scenario $n = 2$, where an AS has only $m = 2$ of its peers asserting its prefix ownerships even if it has more than two peers. that only peer will assert its prefix ownership. For $m = 2$, it provides a level of redundancy in the case that one of the two asserting peers fails to carry out its due diligence.

We see from Table 4 that in the recommended scenario $n = 2$, 20.8% of the ASes need to update their PALs during the month. 9.8% of the ASes need to only one PA change in the month, 5.8% need 2 to 4 PA changes, 2.8% need 5 to 10 PA changes. However, a small number of ASes need more than 100 changes, and AS 701 (UUNET) needs 5 465 changes. For a large organization like UUNET (in this case), we believe that this worst case³ of 5 465 updates is feasible. Table 5 in Appendix 2 lists the organizations which need more than 100 PA changes in the month. We can see that those requiring many PA changes are large ISPs. Exceptions are ASes 23311 and 26224 which are not large ISPs, but need 4924 PA changes. This is because they peer with AS 701 and are randomly selected in our analysis to assert prefix ownerships for AS701. We suggest that ASes should choose large ISPs (e.g., their upstream service providers) to assert prefix ownerships for them since large ISPs usually have more resources and capabilities to respond to changes more quickly. For example, AS 701 also peers with large ISPs, e.g., AS 209 (Qwest), 3356 (Level3) and 1239 (Sprint); for $n = 2$ it could choose any two of them as its asserting peers.

# of PA Changes		1	2-4	5-10	11-30	31-60	61-100	101-200	201-300	301-1000	1001-5000	over 5000	Total
n=1	# of ASes (percentage)	1 650 (9.1%)	824 (4.6%)	392 (2.2%)	215 (1.2%)	56 (0.3%)	19 (0.1%)	20 (0.1%)	1 (0%)	0 (0%)	1 (0%)	1 (0%)	3 179 (17.6%)
n=2	# of ASes (percentage)	1 767 (9.8%)	1 052 (5.8%)	513 (2.8%)	267 (1.5%)	86 (0.5%)	30 (0.2%)	31 (0.2%)	3 (0%)	3 (0%)	2 (0%)	1 (0%)	3 755 (20.8%)
n=3	# of ASes (percentage)	1 864 (10.3%)	1 217 (6.7%)	602 (3.3%)	365 (2.0%)	106 (0.6%)	28 (0.2%)	37 (0.2%)	5 (0%)	3 (0%)	3 (0%)	1 (0%)	4 231 (23.4%)
n=all	# of ASes (percentage)	1 270 (7.0%)	1 865 (10.3%)	1 930 (10.7%)	2 748 (15.2%)	2 476 (13.7%)	1 819 (10.1%)	744 (4.1%)	56 (0.3%)	20 (0.1%)	1 951 (10.8%)	429 (2.4%)	15 308 (84.8%)

Table 4: Projected number of ASes absolute number, and as percentage of all ASes requiring the specified prefix assertion changes based on July 2004 Data. We recommend row $n = 2$.

6.3 Discussions

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 a peer of a given AS updates its asserted prefix ownerships for that AS in a timely manner, a service agreement between them would likely be required, e.g., an extension to their existing agreements. Since there is usually some window before newly delegated prefixes are actually used on the Internet, an

³This worst case assumes a separate update for each PA change. In practice, the actual number of updates might be considerably less, e.g., if one update reflects a large number of PA changes.

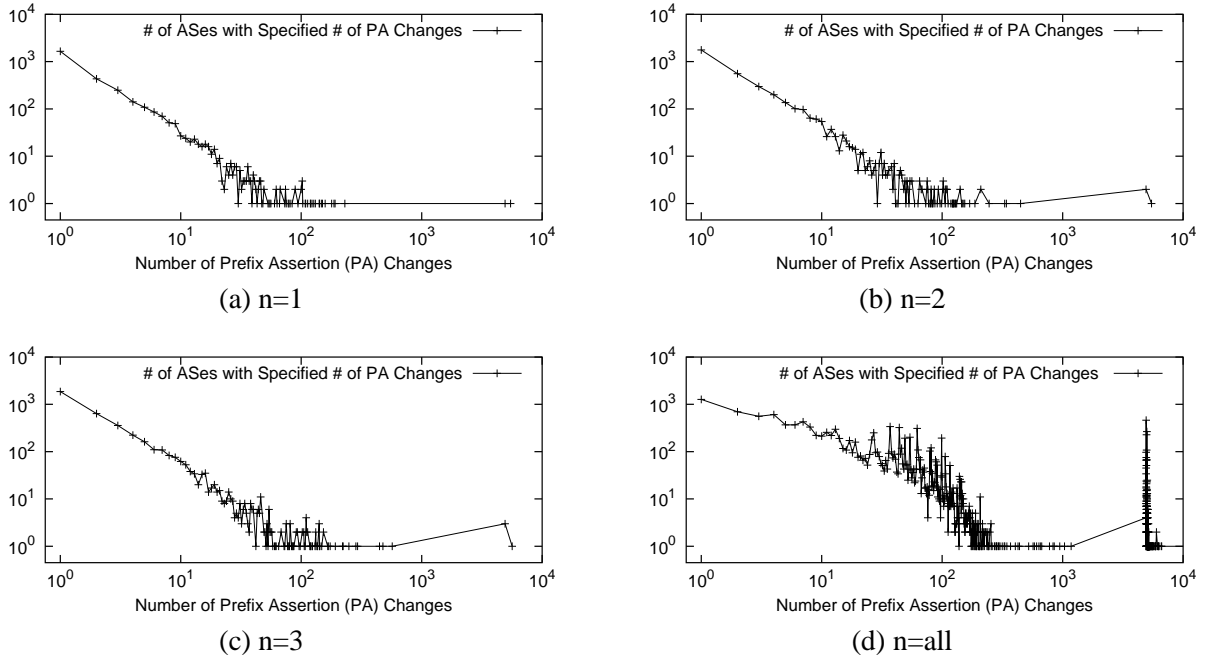


Figure 8: Projected Prefix Assertion Changes based on July 2004 Data (based on Figures 4, 7).

asserting peer should be required to update its PAL to include newly delegated prefixes of the asserted peer 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 can be distributed with BGP update messages in newly defined path attributes [30], thus, they can be distributed as fast as announcements of prefixes. PALs might also be stored in centralized directories [30]. However, a “pull” model might make it challenging to decide how often centralized directories should be checked.

To the best of our knowledge, there is no similar study projecting the number certificate updates per AS required by S-BGP and soBGP. S-BGP has been evaluated for the requirements of storage, CPU, and memory [28], but not certificate updates. We are not aware of any performance study for soBGP. Thus, it is difficult to directly compare psBGP with S-BGP and soBGP on their requirements of prefix certificate updates at this time.

7 Related Work

Significant research has been published on securing routing protocols. Perlman [39] was among the first to recognize and study the problem of securing routing infrastructures. Bellovin [5] discussed security vulnerabilities of Internet routing protocols as early as 1989. More recently, Bellovin and Gansner [6] discussed potential link-cutting attacks against internet routing. Kumar [32] proposed the use of digital signatures and sequence numbers for protecting the integrity and freshness of routing updates. Smith et al. [43] proposed the use of digital signatures, sequence numbers, and a loop-free path finding algorithm for securing distance vector routing protocols including BGP. Thorough analysis of BGP vulnerabilities was performed by Murphy [35].

The most concrete security proposal to date for addressing BGP vulnerabilities is S-BGP [28, 29, 42],

which proposes the use of centralized PKIs for authenticating AS numbers and IP prefix ownership. The S-BGP PKIs are rooted at RIRs, and parallel to the existing system of AS number assignment and IP address allocation. AS_PATH is protected using nested digital signatures, and the integrity of an AS_PATH is guaranteed.

soBGP [47] 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 peering ASes. A global AS graph can be constructed from those certificates. Thus, the existence (only) of an AS_PATH can be verified.

Goodell et al. [19] proposed a protocol, namely Interdomain Routing Validator (IRV), for improving the security and accuracy of BGP. Each AS builds an IRV server which is authoritative of the inter-domain routing information of that AS. An IRV can query another IRV to verify BGP UPDATE messages received by its hosting AS. Improper prefix origination and AS_PATH might be detected by uncovering the inconsistency among responses from other IRVs. One advantage of IRV is that it supports incremental deployment since it does not require changes to the existing routing infrastructure.

Kruegel et al. [31] 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 appears easy to deploy.

In a rigorous study of prefix origination authentication, Aviello et al. [2] formalize the IP prefix delegation system, present a proof system, and propose efficient constructions for authenticating prefix origination. Real routing information is analyzed for restoring 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 [45] are proposed 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 in [45]) 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. [24] propose a Secure Path Vector (SPV) protocol for securing BGP. SPV makes use of efficient cryptographic primitives, e.g., authentication trees, one-way hash chains for protecting AS_PATH. It is shown that SPV is more efficient than S-BGP.

8 Concluding Remarks

Different approaches have been taken by S-BGP and soBGP for addressing security in BGP. In essence, psBGP combines their best features, while differing fundamentally in the approach taken to verify IP prefix ownership. As no centralized infrastructure for tracing changes in IP prefix ownership currently exists, and it would appear to be quite difficult to build such an infrastructure, we suggest that the decentralized approach taken by psBGP provides a more feasible means of increasing confidence in correct prefix origination. We also suggest that the certificate structure and trust model in psBGP has practical advantages. We hope that our comparison of S-BGP, soBGP and psBGP will help focus discussion of securing BGP on the technical merits of the various proposals. We also hope this paper will serve to stimulate discussion in the Internet community about alternate design choices for securing BGP.

Acknowledgments

The first author is supported in part by Alcatel Canada, MITACS (Mathematics of Information Technology and Complex Systems), and the NCIT (National Capital Institute of Telecommunications). The second author is supported in part by MITACS and NSERC (Natural Sciences and Engineering Research Council of Canada). The third author is Canada Research Chair in Network and Software Security, and is supported in part by NCIT, an NSERC Discovery Grant, and the Canada Research Chairs Program.

References

- [1] C. Adams and S. Lloyd. *Understanding Public-Key Infrastructure*, 2nd edition. Addison Wesley Professional, 2003.
- [2] W. Aiello, J. Ioannidis, and P. McDaniel. Origin Authentication in Interdomain Routing. In *Proc. of the 10th ACM Conferences on Computer and Communication Security (CCS'03)*, Washington, D.C., USA. October 2003.
- [3] R. Atkinson and S. Floyd. IAB Concerns & Recommendations Regarding Internet Research & Evolution. RFC 3869, August 2004.
- [4] A. Barbir, S. Murphy, and Y. Yang. Generic Threats to Routing Protocols. Internet Draft (work in progress), April 13, 2004.
- [5] S.M. Bellovin. Security Problems in the TCP/IP Protocol Suite. *ACM Computer Communications Review*, 19(2): 32-48, April 1989.
- [6] S.M Bellovin and E.R. Gansner. Using Link Cuts to Attack Internet Routing. May 2003. <http://www.research.att.com/smb/papers/>
- [7] D. Boneh, X. Boyen, and H. Shacham. Short Group Signatures. In *Proceedings of Crypto 2004*, LNCS vol 3152, pp. 41-55. Santa Barbara, USA. August 15-19, 2004.
- [8] M. Burrows, M. Abadi, and R. Needham. A Logic of Authentication. *Research Report 39*, Digital Systems Research Center, February 1989.
- [9] V.J. Bono. 7007 Explanation and Apology. <http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html>.
- [10] R. Chandra, P. Traina, and T. Li. BGP Communities Attribute. RFC 1997, August 1996.
- [11] D. Estrin, J. Postel, and Y. Rekhter. Routing Arbiter Architecture. <http://www.isi.edu/div7/ra/Publications/>, June 1994.
- [12] S. Convey, D. Cook, and M. Franz. An Attack Tree for the Border Gateway Protocol. IETF Internet Draft, October 2002.
- [13] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-Law Relationships of the Internet Topology. In *Proceedings of ACM SIGCOMM*, 1999.
- [14] J. Farrar. Cable and Wireless Routing Instability. <http://www.merit.edu/mail.archives/nanog/2001-04/msg00209.html>.
- [15] P. Ferguson and D. Senie. Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing. RFC 2267 (Informational), January 1998.
- [16] K. Gaarder and E. Sneekenes. Applying a Formal Analysis Technique to the CCIT X.509 Strong Two-Way Authentication Protocol. In *Journal of Cryptology*, 3: 81-98, 1991.
- [17] L. Gao. On Inferring Autonomous System Relationships in the Internet. In *Proceedings of IEEE Global Internet*, November 2000.
- [18] V.D. Gligor, R. Kailar, S. Stubblebine, and L. Gong. Logics for Cryptographic Protocols - Virtues and Limitations. In *Proceedings of the Computer Security Foundations Workshop IV*, pp. 219-226. IEEE Computer Society Press, Los Alamitos, California, USA. 1991.
- [19] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin. Working around BGP: An Incremental Approach to Improving Security and Accuracy in Interdomain Routing. In *Proceedings of 2003 Internet Society Symposium on Network and Distributed System Security (NDSS'03)*, San Diego, California, USA. February 2003.
- [20] R. Govindan and A. Reddy. An Analysis of Internet Inter-Domain Topology and Route Stability. In *IEEE InfoCom*, 1997.

- [21] R. Guida, R. Stahl, T. Bunt, G. Secrest and J. Moorcones. Deploying and Using Public Key Technology: Lessons Learned in Real Life. *IEEE Security and Privacy*, July/August 2004. pp. 67-71.
- [22] A. Heffernan. Protecting of BGP Sessions via the TCP MD5 signature option. RFC 2385 (Standards Track), August 1998.
- [23] Y.C. Hu, A. Perrig, and D.B. Johnson. Efficient Security Mechanisms for Routing Protocols. In *Proc. NDSS'03*, San Diego, USA. Feb 2003.
- [24] Y.C. Hu, A. Perrig, and M. Sirbu. SPV: Secure Path Vector Routing for Securing BGP. In *Proc. of SIGCOMM'04*, Portland, Oregon, USA. Aug.30 - Sep.3, 2004.
- [25] J. Ioannidis and S.M. Bellovin. Implementing Pushback: Router-Based Defense Against DDoS Attacks. In *Proceedings of 2002 Internet Society Symposium on Network and Distributed System Security (NDSS'02)*, San Diego, California, USA. February 2002.
- [26] S. Kent and R. Atkinson. Security Architecture for the Internet Protocol. RFC 2401 (Standards Track), November 1998.
- [27] S. Kent and R. Atkinson. IP Encapsulating Security Payload (ESP). RFC 2406 (Standards Track), November 1998.
- [28] S. Kent and C. Lynn, J. Mikkelsen, and K. Seo. Secure Border Gateway Protocol (Secure-BGP) - Real World Performance and Deployment Issues. In *Proceedings of 2000 Internet Society Symposium on Network and Distributed System Security (NDSS'00)*, San Diego, California, USA. February 2000.
- [29] S. Kent and C. Lynn and K. Seo. Secure Border Gateway Protocol (Secure-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4): 582-592, April 2000.
- [30] S. Kent. Secure Border Gateway Protocol: A Status Update. In *Proceedings of the 7th IFIP TC-6 TC-11 Conference on Communications and Multimedia Security*, Italy, October 2-3, 2003.
- [31] C. Kruegel, D. Mutz, W. Robertson, and F. Valeur. Topology-based Detection of Anomalous BGP Messages. In *Proceedings of the 6th Symposium on Recent Advances in Intrusion Detection (RAID'03)*, September 2003.
- [32] B. Kumar. Integration of Security in Network Routing Protocols. In *ACM SIGSAC Review*, 11(2): 18-25, Spring 1993.
- [33] U. Maurer. Modelling a Public-Key Infrastructure. In *Proceedings of the Fourth European Symposium on Network and Distributed System Security (ESORICS'96)*, pp. 324-350, 1996.
- [34] D. Meyer. The RouteViews Project (<http://www.routeviews.org/>). August 2004.
- [35] S. Murphy. Border Gateway Protocol Security Analysis. IETF Internet Draft, draft-murphy-bgp-vuln-00.txt. November 2001.
- [36] S. Murphy. BGP Security Protection. IETF Internet Draft, draft-murphy-bgp-protect-02.txt. February 2002.
- [37] D.M. Nicol, S.W. Smith, and M.Y. Zhao. Evaluation of efficient security for BGP route announcements using parallel simulation. *Simulation Practice and Theory Journal*, special issue on Modeling and Simulation of Distributed Systems and Networks. June 2004.
- [38] University of Oregon - Looking Glass. <http://antc.uoregon.edu/route-views/>
- [39] R. Perlman. *Network Layer Protocols with Byzantine Robustness*. PhD thesis, Massachusetts Institute of Technology, August 1988.
- [40] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4), RFC 1771, March 1995.
- [41] M. Reiter and S. Stubblebine. Toward Acceptable Metrics of Authentication. In *IEEE Symposium on Security and Privacy*, pp. 10-20, 1997.
- [42] K. Seo, C. Lynn, and S. Kent. Public-Key Infrastructure for the Secure Border Gateway Protocol (S-BGP). *IEEE DARPA Information Survivability Conference and Exposition II*, 2001.
- [43] B.R. Smith and J.J. Garcia-Luna-Aceves. Securing the Border Gateway Routing Protocol. In *Proceedings of Global Internet 1996*. London, UK. November 1996.
- [44] L. Subramanian, S. Agarwal, J. Rexford, and R.H. Katz. Characterizing the Internet Hierarchy From Multiple Vantage Points. In *IEEE INFOCOM*, 2002.
- [45] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. Katz. Listen and Whisper: Security Mechanisms for BGP. In *Proc. of the First Symposium on Networked Systems Design and Implementation (NSDI'04)*, San Francisco, CA, USA. March 2004.
- [46] R.White, D. McPherson, and S. Sangli. *Practical BGP*. Addison-Wesley. June 2004.

- [47] R. White. Securing BGP Through Secure Origin BGP (soBGP). In *The Internet Protocol Journal*, 6(3): 15-22, September 2003.
- [48] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S.F. Wu, and L. Zhang. An Analysis of BGP Multiple Origin AS (MOAS) Conflict. In *ACM SIGCOMM Internet Measurement Workshop*, San Francisco, USA. Nov. 2001.
- [49] X. Zhao, D. Pei, L. Wang, L. Zhang, D. Massey, A. Mankin, and S.F. Wu. Detection of Invalid Route Announcement in the Internet. In *International Conference on Dependable Systems and Networks*, 2002.

Appendix 1: Issues with a Web-of-Trust Model

There is much debate on the architecture for authenticating AS public keys in the BGP community, and in particular the pros and cons of using a strict hierarchical trust model vs. a distributed trust model (e.g., a web-of-trust model). While a web-of-trust model is widely used within the technical PGP community for authenticating user public keys, it is not clear if it is suitable for authenticating AS public keys in practice due to a number of issues. Some of these are discussed below.

- *Issue of Bootstrapping Trust.* To bootstrap trust, some entities must be trusted for signing a certificate binding an AS number to a public key. Top ISPs have been proposed for functioning as such trusted certificate authorities [47]. However, their authority for signing AS public key certificates is at best questionable, since only IANA/ICANN and RIRs have authority over AS numbers. Top ISPs may be trusted for forwarding subscriber traffic because of their large scale networks, but probably not for authenticating AS numbers because that is beyond their jurisdiction. In addition, a top ISP may be trusted by people within its geographic area, but may not be trusted by outside entities especially those who might have conflicts of interest with them.
- *Issue of Trust Transitivity.* A web-of-trust model relies upon trust transitivity for expanding trust. Given a chain of public key certificates, an entity must trust the signature on the first certificate and every intermediate certificate on the chain to trust the authenticity of the last public key. Given trust in the first entity, it is not clear why one should trust the downstream entities of the chain. For example, in real life, it is well accepted that trust is not transitive.
- *Vulnerable to a single bad party.* A web-of-trust model is vulnerable to a single misbehaving party involved in a certificate chain, essentially requiring the assumption that there is no single misbehaving entity on a certificate chain. This seems to contradict to the threat model of many security proposals (e.g., S-BGP, soBGP) which allow and try to resist uncoordinated misbehaving entities. Requiring multiple signatures may be of little value in a web-of-trust model since a single misbehaving entity may be able to obtain multiple trusted public key certificates [33, 41]. This is possible due to the fact that in a web-of-trust model, no one has authority over (“owns”) the name space involved.

Appendix 2: ASes with Top 40 Number of PA Changes

# of PA Changes	AS Number	Organization Name
102	17633	ASN for Shandong Provincial Net of CT
	17773	CNNIC, China Network Information Center
	6347	SAVVIS Communications Corporation
103	21578	Universidad autonoma de Bucaramanga
	8054	Ticsa
106	19832	20twenty Financial Services [Pty] Ltd
107	11744	Investec Bank
	6467	E.Spire Communications, Inc.
108	2905	The Internetworking Company of Southern Africa
	5400	Concert European Core Network
110	25653	Pegasus Web Technologies
	29791	Voxel.net, Inc.
112	19429	E.T.B.
118	174	PSINet Inc.
121	4323	Time Warner Communications, Inc.
122	30893	Glassbilen Networks
123	16150	Port80 AB, Sweden
124	8473	Bahnhof Autonomous System
125	2914	Verio, Inc.
126	4755	Videsh Sanchar Nigam Ltd. Autonomous System
129	724	DLA Systems Automation Center
130	9600	SONY CORPORATION
132	7303	Telecom Argentina Stet-France Telecom S.A.
141	22597	Syngy, Inc
	30544	People First Federal Credit Union
145	286	KPNQwest Backbone AS
147	2497	IJJ
151	1785	AppliedTheory Corporation
154	7474	Optus Communications Pty Ltd
169	10026	IXNet Hong Kong Limited
187	3549	Global Crossing
209	4134	Data Communications Bureau
	721	DLA Systems Automation Center
245	209	Qwest
330	3356	Level 3 Communications, LLC
341	7018	AT&T
447	1239	Sprint
4924	23311	Hinda Incentives
	26224	PRE Solutions, Inc.
5465	701	UUNET Technologies, Inc.

Table 5: ASes with Top 40 Numbers of PA Changes based on July 2004 Data