

SECURITY IN INTERDOMAIN ROUTING

Tuna Vardar
Helsinki University of Technology
T-110.551 Seminar on Internetworking 2004
26-27.04.2004
vardart@cc.hut.fi

Abstract

Current interdomain routing protocols are limited in implementations of universal security. Because of this, the Internet is vulnerable to many attacks at the AS to AS routing infrastructure. Such attacks can result in Internet outages, manipulation or exposure of Internet traffic, or the loss of control over Internet address space. BGP is the protocol that enables interdomain routing in the Internet. Although BGP has proven to be generally stable, there are serious concerns about its capability to match the requirements of the rapidly evolving Internet. An important limitation of BGP is its failure to satisfy requirements of security. The design of BGP has complicated definitions at securing interdomain routing. This paper comprehensively examines works about BGP security. The limitations and advantages of three proposed solutions are analysed in this paper.

1 Introduction

The Internet is a network of networks, owned and operated by different companies, including Internet backbone providers. These networks are called Autonomous Systems(AS), which are managed independently. Those ASes connect with each other to provide connectivity to end-users. In order to provide connectivity across Internet, each AS must manage its own routing information as well as exchange routing information with other ASes. The interdomain routing is the type of routing across ASes. The de facto interdomain routing protocol in use to route Internet traffic among backbone is version 4 of the Border Gateway Protocol (BGP-4). However, The Internet was not built with security in mind; it was built with communication in mind. Depending on this fact, BGP has intrinsic security failures. BGP's design goals include ability to spread reachability information in a way that would recover some well known problems of routing but little attention was paid to security of it. BGP could become the target of attacks that could disorder Internet services.

2 Security aspects

2.1 Fundamental aspects

When an entity is connected to the Internet, it pays for the services of an Internet service provider (ISP) and therefore

becomes part of the ISP's administrative domain (AS). Entity routes its own Internet traffic to other entities through its ISP, which in turn routes traffic through its upstream ISP and so on. ISPs receive routes from their upstream ISPs and announce all routes to their customers. From this point of view, it is obvious that systems grouped together can be targeted together by a malicious attacker. Furthermore, BGP routers mutually trust each other. BGP router traffic is not encrypted. There are not any true authentication mechanisms built in, and there's no such thing as a digital signature between BGP router traffic. Cryptographic authentication is not mandatory in BGP and it is not widely used.

There are basically two common ways someone can harm a BGP session. The first is to masquerade as a peer router, taking over the IP address of that peer. The attacker can then propagate bad information into the routing tables or the attacker can acquire routing information. The attacker might even route some of router's address space to himself and appear to the world as the router. The other form of attack is to reset a BGP session of a router. Since BGP operates in TCP layer of Internet, it is subject to the same kinds of attacks that might exist in TCP: Session stealing, denial of service, etc. An attacker can reroute traffic down a path that will enable the attacker to view the data along that path, or attacker can send the data into a black hole. Incidentally, several backbone operators do not listen to the BGP traffic from the lower level ISPs, thus they are protected from many attacks.

2.2 Threats, vulnerabilities, attacks

BGP messages are subject to modification, deletion, forgery, and replay [1]. These exploits can be caused by malicious intent as well as faulty or misconfigured BGP routers. Moreover, bogus messages can originate from malicious sources or accidentally misconfigured peers. The effects of misconfiguring a BGP router can be similar to those of an attack [10].

There are two areas of globally visible misconfigurations in BGP [2]:

- An AS accidentally injects a prefix into the global BGP tables (origin misconfiguration).
- A router exports a route it should have filtered (export misconfiguration).

Malicious BGP packet manipulation can lead to erroneous information in the compromised router's routing table. There are three main vulnerabilities in BGP [1]:

- BGP does not protect the integrity, freshness and source authentication of messages.
- BGP does not validate an AS's authority to announce reachability information.
- BGP does not ensure the authenticity of the path attributes announced by an AS.

BGP is subject to following attacks [1]:

Eavesdropping: Attacker intercepts and reads BGP messages. The routing data carried in BGP is carried in cleartext, so eavesdropping is a possible attack against routing data confidentiality.

Replay: Attacker records messages and resends them. BGP does not provide for replay protection of its messages. This can be a side effect for denial of service (DoS) attacks.

Message insertion: An attacker inserts bogus messages into a BGP session. BGP does not provide protection against insertion of messages. However, TCP can discover this type of attacks by sequence numbers in the packets.

Message deletion: An attacker intercepts and deletes a message passed between BGP peers. BGP does not provide protection against deletion of messages. This kind of attack is quite hard to perform in TCP.

Message modification: An attacker removes messages from a BGP session, modifies them, and reinserts them. BGP does not provide protection against modification of messages. If this type of attack does not change the length of TCP payload, it would not be detectable.

Man-in-the-middle: An attacker completely corrupts the communication stream between two peers and poses as the sender to the receiver and vice versa. BGP does not provide protection against man-in-the-middle attacks. As BGP does not perform peer entity authentication, a man-in-the-middle attack is easy to do.

Denial of service (DoS): An attacker floods a resource to the point of exhaustion. This type of attack can be performed by using bogus routing data. For example, advertising large numbers of more specific routes (longer prefixes) can cause BGP traffic and router table size to increase.

3 Security in BGP

3.1 Current status of security in BGP

BGP security mechanisms protect the transmission of routed message across local networks (within the AS); however BGP does not provide integrity or authentication of the routing information itself as it traverses the nodes and links that make up the Internet[10]. Lack of strong security implementation in BGP creates big risks for attacks defined in previous section. ISPs are using their own methods of protection rather than a proposed security mechanism. First method used by ISPs is filtering the Internet traffic at border routers. BGP traffic and misuse of their internal IP addresses is controlled by border routers. This is a good defense method. Second method, in BGP router configuration, ISPs usually hardcode IP address of other BGP peers, which BGP router is linked to. Even though, wrong configurations may lead to

insecure situation, still this is a good way of defense as soon as router configuration is done correctly.

3.2 BGP security requirements

In order to minimize vulnerabilities of BGP, security must be implemented in BGP. After considering possible attacks and their vulnerabilities, there are defined requirements for BGP security [4]:

- Security architectures for BGP should not rely on mutual trust among ISPs: Some ISPs will never be trusted. Also, trusted parties can make mistakes or they can change behaviour. Transitive trust in parties causes mistakes to propagate.
- Solutions must demonstrate similar efficiency, performance and reliability as the other parts of BGP.
- The requirements of a solution should scale well within BGP.
- Integrity and authenticity of BGP messages should be guaranteed at the traffic (Classified as hop integrity).
- A BGP router should be able to verify the owner of each prefix that authorized the origin AS (Classified as origin authentication).
- A BGP router should be able to verify that each subsequent AS in the path has been authorized by its predecessor AS (Classified as path validation).

If a security approach fails to achieve these requirements, BGP routers will be vulnerable to attacks. If a security approach achieves these requirements, BGP routers will be able to detect attacks and reject unauthorized routes.

4 Security problems in BGP

Current efforts try to find solutions for these problems of BGP security:

4.1 Hop integrity

A computer network is said to provide hop integrity if and only if the following condition holds for every pair of adjacent routers p and q in the network. When q receives a message m supposedly from p , then q can check that m was not modified after it was sent by p , and that m was not a replay of an old message sent long ago by p [6].

In BGP, hop integrity is not provided. BGP should provide following to satisfy hop integrity [10]:

- **Data integrity:** A verification performed at each hop to assure that the data in a message has not been modified, destroyed, lost, or replayed in an unauthorized or accidental manner.
- **Source authentication:** A verification performed at each hop to assure that the sender of a message is who it claims to be and not a pretender.

4.2 Origin authentication

Origin authentication is a corroboration (such as by using a digital signature) that the origin of a message or data is as claimed. Origin authentication is validation of AS claims of address ownership. After it has been determined that a BGP router is authenticated, the next logical step is to determine if that BGP router is authorized to advertise the information it had sent. Addresses on the Internet are matched to ASes through a hierarchical network of issuing authorities and organizations. Origin authentication should ask questions such as "Is AS1024 authorized to advertise the prefix 120.40.0.0/16?" [7].

4.3 Path Validation

Path validation is process of validating [8]:

- all the digital certificates in a certification path
- the required relationships between those certificates, thus validating the contents of the last certificate on the path.

Inside a BGP UPDATE message sent by BGP router, each announced prefix has an associated AS path to that prefix. Path validation ensures that the path is valid (each BGP router in the path is accessible from the previous BGP router), and each AS on the path is authenticated [10].

5 Solutions for secure routing in BGP

5.1 S-BGP (Secure BGP)

S-BGP addresses vulnerabilities of BGP by defining scalable methods of verifying the authenticity and authorization of BGP control traffic. The S-BGP architecture uses three security mechanisms to satisfy BGP security requirements: PKIs, attestations, and IPSec [9].

Public Key Infrastructure (PKI) is based on the use of X.509v3 certificates. It is used to support the authentication of IP address block ownership, AS Number ownership, AS identification, and BGP router identification and authorization to represent an AS. To achieve all of these aims, there is need for three kinds of certificates. First type of certificate assigns a public key to an organization and to a set of IP address prefixes. These certificates are used to verify if an originating AS owns a specified portion of IP address space or to specify if the owner has authorized AS to advertise the address space. The certificates are arranged into a singly-rooted hierarchy that parallels the existing IP address allocation system. ICANN is the root in the certificate hierarchy. Next tier consists of Internet registries (e.g. RIPE, ARIN). Third tier consists of major ISPs. And the other tiers consists of others ISPs and subscribers. The second type of certificate assigns a public key to an organization and a set of AS numbers and the third type of certificate assigns a public key to an AS number and to a BGP router ID. These two types of certificates are used by BGP speakers to authenticate one another, and to verify that a given speaker is authorized to represent specified AS. The second and third type certificates are arranged into a singly-rooted hierarchy as well. ICANN

is the root in the certificate hierarchy. Next tier consists of Internet registries, and third tier consists of ISPs and subscribers. Second type of certificates are assigned to second tier, and third type certificates are assigned to third tier [4,9].

Attestations form the most important part of S-BGP. Attestations are protected by digital signatures. Their primary purpose of use is to encounter Byzantine attacks (where the attacker's aim is to see everybody lose). Attestations are signed and validated using the keys from PKI. Each BGP speaker that receives a route advertisement uses attestations to verify that each AS along the route has been authorized by preceding AS along the path to advertise the route. And attestations are also used to verify that the originating AS has been authorized by the owner of each IP prefix contained in the UPDATE message to advertise these prefixes. Attestations are carried in a new and optional BGP transitive path attribute that contains digital signatures covering the route information. There are two types of attestations: Route attestations and address attestations. Route attestations are issued by an AS and they subject a transit AS or another AS providing third party advertisements for an AS that is not running BGP. Address attestations are issued by the organization that owns the address prefixes contained in the attestation and they subject one or more ASes that are authorized to these advertise prefixes [ibid].

IPSec is used to provide data and partial sequence integrity, and peer entity authentication for BGP control traffic. ESP (Encapsulating Security Payload) of IPSec is used in BGP to achieve these goals. IPSec protects the integrity of TCP connections used between BGP speakers, because IPSec works in IP layer. Its anti-replay mechanisms detect and reject replayed packets more quickly than TCP, which helps to overcome DoS attacks. Also IPSec may be used in the future, if needed, to provide confidentiality for BGP control traffic [ibid].

Origin authentication in S-BGP is done by using PKI and address attestations. First type of certificates are used to authenticate organization's ownership of IP addresses. Address attestations are signed by owner's private key. This private key corresponds to a public key in the first type of certificate in PKI used by S-BGP. Hop integrity is done by the use of IPSec. IPSec provides both per-hop data integrity and per-hop source authentication. Path validation is done by use of route attestations and PKI. These attestations are used by a transit AS for verification of path information. When combined with certificates from the PKIs, a BGP speaker is able to validate the authenticity and integrity of every AS on a path from source to nearest neighbor by comparing the attestations with a certificate database [4,9,10].

S-BGP had been experimented at a testbed by DARPA's CAIRN (Collaborative Advanced Interagency Research Network). They tested security achievements of S-BGP, overall S-BGP performance and interoperation capabilities with BGP-4. S-BGP provided desired security improvement in the tests. S-BGP detected and rejected manipulated malicious BGP messages. Tests about performance showed that S-BGP has significant overhead especially in CPU utilization and storage/memory. Basically this is due to cryptography, that is used in many stages of S-BGP. There was little overhead in bandwidth because of increasing size of BGP

UPDATE message. But it is reported that increasing size of BGP UPDATE message does not have too much effect on performance. Interoperation capabilities were satisfactory by using a BGP router that records router traffic and sends to a S-BGP router. Then S-BGP router within the same AS could distribute attestations to other S-BGP speakers. No problems on interoperability were observed within an AS [4].

Although S-BGP provides good solutions to problems of BGP, it is not widely deployed yet. Because, S-BGP is a new protocol that targets to replace BGP. It is said that ISPs cannot afford to replace or upgrade current BGP routers, registries cannot afford to offer certificate authority services, and router manufacturers do not afford to implement S-BGP router software or produce S-BGP router hardware if ISPs will not buy them. Beyond these economical reasons, there are some technical reasons. Router requirements of S-BGP in terms of memory and CPU utilization is very demanding. If S-BGP is wanted to be used efficiently, all ASes in the path between the announcer of a BGP UPDATE message and recipient of the message, must have an implementation of S-BGP running. BGP UPDATE messages are getting more complex if route attestations are wanted to be used. Designers of BGP kept route attestations optional because of the interoperability problems that may be faced after deployment. But it is still possible that S-BGP can be deployed incrementally to overcome these technical problems [4,9,10].

5.2 IRV (Interdomain Route Validation)

IRV defines a service that protects against completely ruined, or misconfigured ASes, and is used to identify and diagnose routing configuration problems. IRV relies on out-of-band communication with a route originator to verify the correctness of a route [11].

IRV has following goals to achieve in its definition [ibid]:

1. Allow ASes to acquire and validate both static and dynamic interdomain routing information.
2. Be incrementally deployable.
3. Allow ASes to securely differentiate the requesters of routing information, in order that responses be tailored to the recipient.
4. Not be tightly coupled with BGP; the protocol must operate independently of the reception of BGP messages, and ASes must be free to validate and acquire routing information whenever they desire.
5. Allow ASes to passively receive routing-relevant information from remote entities; this will permit collections of participating ASes to cooperatively monitor and debug the routing infrastructure.

IRV architecture is a decentralized query system. IRV provides access to dynamic routing data (e.g. BGP route announcements, current routing tables) and static routing data (e.g. routing policy) through a query interface. IRV allows ASes to confirm that they have announced and propagated particular routes. These confirmations are done by an Interdomain Routing Validator (IRV) server, which is located

within the AS that it represents. It is a dedicated machine or a set of machines. IRV server is responsible for answering queries from other ASes. So it should provide an interface by which external entities query routing data. During the BGP UPDATE message exchange, each AS designates an IRV server. IRV server is responsible for answering queries from other ASes. Other ASes query IRV server to validate received BGP data or to acquire additional and relevant route information. IPsec or TLS is used to ensure the integrity, authenticity and timeliness of the queries and responses. The procedure in IRV is very simple. Whenever an AS finds BGP data suspicious, it can check all ASes along the path by querying IRV servers. IRV servers are assigned to ASes uniquely. This is done by using a well-known registry which stores AS unique data (e.g. ASN) and unique IRV server contact information for each AS (e.g. IP address). Implementation of IRV is targeted to be easily deployable, robust and simple. Designers have proposed to place IRV system running on top of HTTP as a web-based service which uses TLS or if needed IPsec. Query requests are done by HTTP POSTs. Responses from IRV servers are XML documents. IRV servers use XQuery to fetch data from their database [ibid].

In IRV system, path validation is done by querying each AS in the path given in an UPDATE message at an IRV server. Hop integrity is provided by the IRV system because IRV servers can talk to all elements in the network. Origin authentication is done by querying home and outer IRV servers, and comparing received results together with BGP routing data [10].

Unlike S-BGP, IRV does not target to replace BGP or put additions to BGP. IRV architecture require a centric registry like S-BGP's PKI certificate authorities. Each AS is responsible for defining its own IRV server, so they can define protection methods of the server and its access with other nodes in the network [10,11].

5.3 soBGP (Secure Origin BGP)

soBGP is a proposed specification for adding security to BGP and it is proposed as an alternative to S-BGP. Under soBGP, ISPs can authenticate route advertisements and can implement policy on them. Designers of soBGP aimed it to be a deployable mechanism for validating the correctness and authorization of the data carried within BGP, and also for preventing the sorts of attacks resulting from misconfiguration or intentional insertion of bad data into the Internet routing system [12].

Designers of soBGP addressed four goals when designing it [12, 13].

1. Is the AS originating the destination (prefix) authorized to advertise it? If a router receives an advertisement for the 10.1.1.0/24 network originating in AS65500, is there any way to verify that AS65500 is supposed to be advertising 10.1.1.0/24? (Origin authentication)
2. Does the AS advertising the destination actually have a path to the destination? In other words, if a router is receiving an advertisement from a BGP peer in AS65501 that it can reach 10.1.1.0/24, is there any way to verify

that AS65501 actually has a path to the AS origination 10.1.1.0/24? (Path verification)

3. Is the peer advertising the route authorized by the originator, or owner, of the destination, to advertise a path to the destination?
4. Does the path advertised by a peer AS fall within the policies the local network administrators have set forward? The most obvious issue is whether or not the AS path advertised by the peer is an acceptable path to send the traffic along.

Although designers wanted soBGP to achieve these four goals, they concluded that reaching goals 3 and 4 is not quite possible in the operation of Internet because of many reasons described at [15]. So soBGP targets to achieve only first two goals.

soBGP adds a new message type SECURITY to current BGP protocol. SECURITY type message are used by BGP speakers to share three different types of certificates. These certificates contain public keys. Certificates are signed by private key of the sender. Receiver validates public and private key pair. And so, receiver is able to validate all BGP traffic messages [12,13].

There are three types of certificates in soBGP. [12,13]

1. EntityCert is used to verify, through a trust model, the existence of an entity within the routing system, and the value of that entity's public key for use in the routing system. Each entity within the routing system must generate a public/private key pair. The public key portion of this pair is then signed, verifying that anyone using this public key is actually the entity in question. EntityCerts are signed by a third party, validating who an entity is within the routing system. So after signed by a third party, EntityCerts can form a web of trust. Web of trust can be built on the public keys of a small number of well-known entities, such as top-level backbone service providers, key authentication service providers (e.g. Verisign), and others. These "root keys" can be distributed out of band and could be used to validate a set of advertised EntityCerts. These are used in turn to build up the database of known good AS/key pairs in the system, allowing even more EntityCerts to be validated.
2. PolicyCert provides information about policies. Policies are requested by an AS, which originates routes. There is only one valid PolicyCert for each AS which originates routes at any given time. This certificate is signed by originator of policies because it is not necessary for any entity outside AS to validate or verify these policies.
3. AuthCert ties an AS to a block of addresses that the AS may advertise. The organization (e.g. ISPs) which authorizes an AS to advertise a block of addresses signs this certificate.

Source authentication in soBGP is done using an EntityCert. EntityCert ties an AS number to a public key (or a set of public keys) corresponding to a private key that the AS

will be using to sign various other certificates. An EntityCert is defined in soBGP to be an X.509v3 certificate, similar to those used by TLS (Transport Layer Security) and IPSec. The main problem when accepting an EntityCert is knowing whether or not the key carried within the certificate is actually the key of the advertising AS. soBGP resolves this by requiring the EntityCert to be signed by a third party, validating that this AS actually belongs with this key. The key each AS distributes in its EntityCert is actually the public half of a private/public key pair. An AS would keep its private key entirely private, holding it on one highly secure device in its network and generating signatures for other certificates as needed [12].

First goal of soBGP is achieved by using of certificates. Any device receiving AuthCerts can check them by looking up the public key of the authorizer, and verifying the signature on the AuthCert, as well as by making certain the authorizer is permitted to advertise the address space it has suballocated this block of address space from. The device then builds a local table of address blocks and corresponding ASes authorized to advertise prefixes within those address blocks. Received updates can be checked against this database to verify authorization of the originating AS to advertise a prefix [ibid].

Second goal of soBGP is achieved by building a topology map of the paths of the entire internetwork. Each AS attached to the internetwork builds an PolicyCert, which contains, primarily, a list of its peers, and signed using the originator's private key. Using this list of transit peers, a map of the internetwork topology may be built. Topology map is a database. And this database stores paths. Using the PolicyCerts announced by each AS, BGP speakers can build the path database of all possible paths to a prefix. As each prefix is processed, path databases can be queried to confirm that the questioned path is valid or not [ibid].

Deployment of soBGP provides a wide variety of options, because it is not transport-dependent, nor dependent on a yet-to-be constructed centralized set of servers. Deployment involves primarily with distribution of certificates. Designers of soBGP propose three different options about deployment [13]:

1. Direct certificate exchange and processing between border routers. With this option, routers that are capable of the cryptographic processing required to validate received certificates exchange certificates with their peers in other ASes (just as they exchange routing information today), process those certificates, and build local databases from which they perform security checks on received updates. This spreads the processing along all the edges in the AS.
2. The edge routers exchange the certificates, but not process them. Instead, each edge router would relay the not-yet-validated certificates to internal servers, thereby validating the certificates by performing the necessary cryptographic operations. As the border routers receive updates, they can query the server about the validity of each update, and take action based on the reply received.

3. It is possible for the internal servers within an AS to exchange certificates directly, over a multihop session, without relays of border routers or processing at border routers. So internal servers would then process the certificates, and the border routers would query these servers to determine whether received updates are valid or invalid.

soBGP is a lightweight solution compared to S-BGP. It has strong security mechanisms. One missing thing with soBGP is how hop integrity would be provided. Source authentication is done by EntityCerts but still there is need for ensuring data integrity of BGP messages [10,12,13].

6 Conclusion

Interdomain routing is very stable and reliable in the transmission of routing data. But security is not provided within interdomain routing protocols. For instance, BGP routing traffic is easily readable if an access to BGP links is achieved, because BGP routing traffic is not encrypted. There is need for interdomain routing protocols that has concerns about security. Three of the proposed solutions are mentioned in this paper. S-BGP targets to replace BGP-4 incrementally, soBGP proposes enhancements to the protocol and IRV proposes a new infrastructure that can be used by BGP routers. These proposals have strong capabilities but still deployment of them is an issue. At the deployment side, biggest concerns are about performance of overall routing procedure. If security is implemented, performance is decreasing mainly because of cryptography usage. Despite this fact, hardware upgrades on BGP routers and bandwidth increases in BGP links may probably increase performance. Also a single decision about right security solution is very important. Nobody wants to spend huge amounts of money for a slow, irrelevant, unsatisfactory solution. And everybody wants a feasible solution that is easily adoptable to currently running BGP. A replacement of BGP with a new secure and fast interdomain routing domain protocol seems quite imaginary nowadays. A solution that adds enhancements to BGP (which are easily deployable) is more preferred. But still there is not any secure solution that is agreed to be an easily deployable enhancement to BGP by ISPs. It seems like there will be discussions going on this subject. And it seems like there will not be any change soon in interdomain routing security mechanisms.

References

- [1] Murphy, S. *Bgp security vulnerabilities analysis*. IETF Draft, 2003.
- [2] Mahajan, R., Wetherall, D., and Anderson, T. *Understanding bgp misconfiguration*. ACM SIGCOMM, 2002.
- [3] Barbir, A., Murphy, S., and Yang, Y. *Generic threats to routing protocols*. IETF Draft, 2003.
- [4] Kent, S., Lynn, C., Mikkelsen, J., and Seo, K. *Secure border gateway protocol (s-bgp) real world performance and deployment issues*. ISOC Symposium on Network and Distributed System Security, 2000.
- [5] Rekhter, Y. and Li, T. *A border gateway protocol 4 (BGP-4)*. IETF RFC 1771, 1995.
- [6] Gouda, M.G., Elnozahy E. N., Huang C.T., McGuire T.M. *Hop Integrity in Computer Networks*. Proceedings of the IEEE International Conference on Network Protocols, 2000.
- [7] Aeillo, W., Ioannidis, J., andMcDaniel, P. *Origin authentication in interdomain routing*. ACM CCS, 2003
- [8] R. Shirey *Internet Security Glossary*. The Internet Society RFC 2828, 2000.
- [9] Stephen Kent, Charles Lynn, and Karen Seo *Secure Border Gateway Protocol (Secure-BGP)*. IEEE Journal on Selected Areas in Communications Vol. 18, No. 4, April 2000, pp. 582-592
- [10] Tony Farley, Patrick McDaniel *A Survey Of BGP Security Issues and Solutions*. AT&T Labs Research, 2003.
- [11] Goodell, G., Aiello, W., Griffin, T., Ioannidis, J., McDaniel, P., and Rubin, A. *Working around BGP: An incremental approach to improving security and accuracy of interdomain routing*. Internet Society Network and Distributed Systems Security, 2003.
- [12] White, R. *Securing BGP Through Secure Origin BGP*. The Internet Protocol Journal, 2003.
- [13] White, R. *Deployment Considerations for Secure Origin BGP (soBGP)*. Network Working Group, Cisco Systems, 2003.
- [14] James, Ng. *Extensions to BGP to Support Secure Origin BGP (soBGP)*. Network Working Group, Cisco Systems, 2002.
- [15] White, R. *Considerations in Validating the Path in Routing Protocols*. Network Working Group, Cisco Systems, 2003.