SECURITY IN INTERDOMAIN ROUTING

Research and Development work

By

Hemant Kumar IT032028
D.R.Esesve IT032022
K.Sriram IT032033

Network Technology
Symbiosis Center for Information Technology

Abstract

The Internet wasn't built with security in mind; it was built with communication in mind, security is an Internet afterthought.

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 research comprehensively examines about BGP security. The limitations and advantages of the currently proposed solutions are analyzed in this paper.

1. Introduction

BGP is the interdomain routing protocol of the Internet. Its primary purpose is to route Internet traffic, not to ensure the secure delivery of that traffic. Accidental misconfigurations of BGP can interrupt Internet connectivity and create havoc. It's conceivable that BGP could become the target of attacks that could disrupt Internet services on a large scale. Think globally in terms of bank databases, telephone networks, defense systems and the like.

1.1 Scope:

- Once we understand how BGP works, we begin to understand the security issues, what can be done and what experts are proposing as possible solutions.

2. Analysis of Problem in Research

BGP is a path-vector routing protocol that runs between autonomous systems on the Internet. Instead of keeping track of the Internet's entire topology, BGP routers receive information (known as reachability information) from neighbor routers and then choose the route with the shortest path for inclusion in the routing table. Each router will then announce the path to other neighbors if its routing policy permits that.

The path is a list of every autonomous system (AS) between the router and the destination. BGP groups networks together within autonomous systems so they may be seen as single entities. It doesn't matter whether there is a single BGP-speaking router within the AS or hundreds of BGP-speaking and non-BGP-speaking routers within the AS. An AS is sometimes described as a "single administrative domain," which makes one think that each business entity on the Internet has its own administrative domain. That isn't correct. An AS can and does include more than one organization.

As a business connected to the Internet, you pay for the services of an Internet service provider and therefore become part of the ISP's administrative domain. Basically, your organization's Internet-bound traffic routes through its ISP, which in turn routes traffic through its upstream ISP and so on.

Routing policies take precedence over reachability information for the simple reason that transit services aren't free. An ISP will receive routes from its upstream ISPs and announce all routes to its customers. Announcing a route is an invitation to route through the announcer, so the basic rule of thumb is *"send routes only to paying customers"*. If a customer has two or more ISPs, it gets a little trickier -- but that's another conversation altogether. The point is that systems grouped together can be targeted together.

2.1 **Security issues**:
Here's the basic issue with the current BGP setup: BGP routers trust one another. There aren't any true authentication mechanisms built in, and there's no such thing as a BGP digital signature. Cryptographic authentication isn't mandated.

There are basically two ways someone can harm a BGP session:

1. 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 unless filters are strict or, conversely, the attacker can garner information. The attacker might even route some of your address space to himself and appear to the world as you. That's a little scary.
2. The other form of attack is to force a reset of your BGP session, which is more than annoyingly disruptive. BGP is subject to the same kinds of attacks at TCP/IP: *IP spoofing, session stealing, denial of service* and the like. When we talk about routing, we are talking about the path that data takes. An attacker can reroute traffic down a path that will enable him to view the data along that path, or he can send the data into a black hole. In any case, there is too much risk with the current setup to not pay attention to what needs to happen next to solve the problems.

2.2 **Security problems in BGP** [3]
Some of the key problem areas under BGP are enumerated below.

Hop integrity [3.1]
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.

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

- Data integrity: 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: verification performed at each hop to assure that the sender of a message is who it claims to be and not a pretender.
- 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 ASs 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?"
- Path Validation: Path validation is process of validating:
    1. All the digital certificates in a certification path
    2. 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.

3. Current Solutions & Implementations

The scientific Internet community has long been working on solutions for securing BGP. Many ideas and possible solutions have been offered. Everyone agrees that the idea of trusted systems has to go. This implies the use of cryptographic authentication of some kind. As we know, encryption carries with it high demands in terms of hardware and bandwidth. However there are a few good trends in securing BGP. Their methodologies and architectures are discussed below in brief.

3.1 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.

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 single rooted hierarchy that parallels the existing IP address allocation system. ICANN is the root in the certificate hierarchy. Next tier consists of Internet registeries (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. BGP speakers to authenticate one another use these two types of certificates, 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 is assigned to second tier, and third type certificates are assigned to third tier.

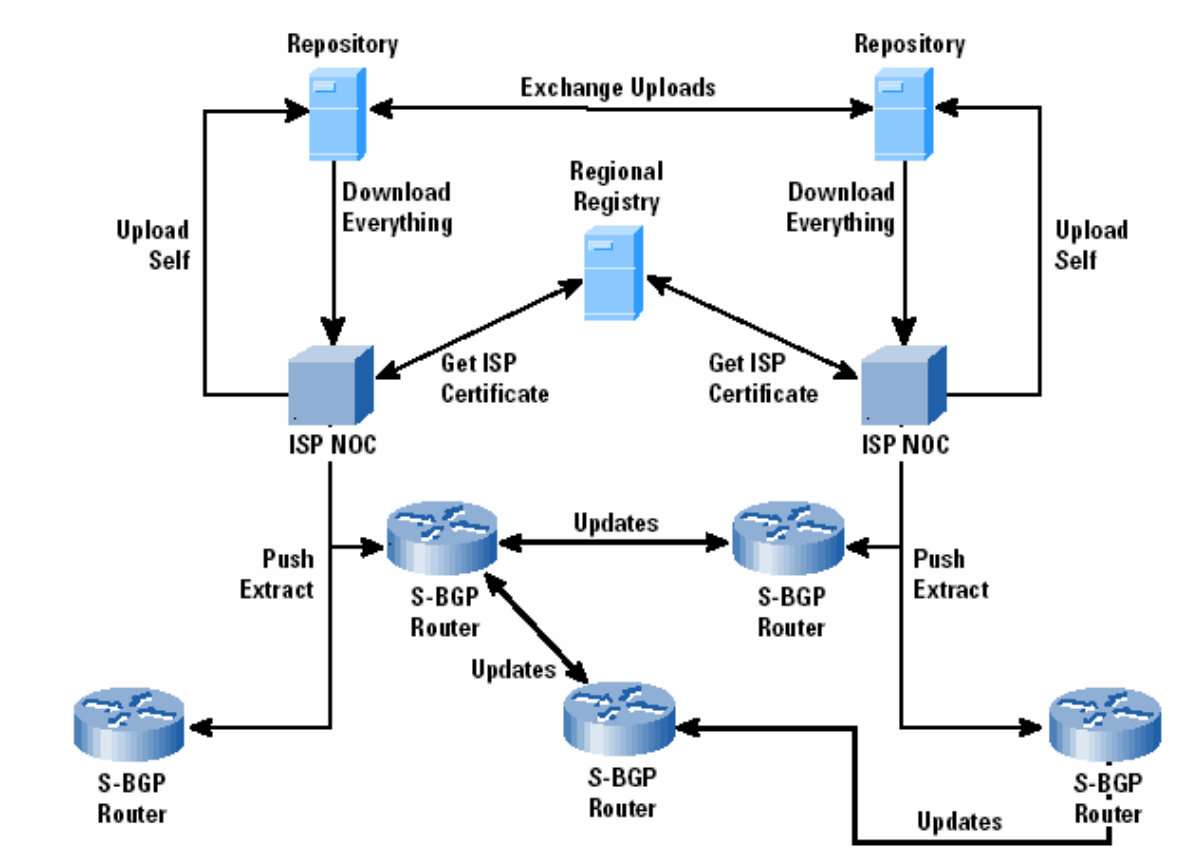

Figure 1 [8]

Attestations form the most important part of S-BGP. Digital signatures protect attestations. 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 preceding AS along the path to advertise the route has authorized each AS along the route. Attestations are also used to verify that the owner of each IP prefix contained in the UPDATE message to advertise these prefixes has authorized the originating AS. 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. An AS issues route attestations 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 ASs that are authorized to these advertise prefixes.

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].

Using PKI and address attestations does origin authentication in S-BGP. First types 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.

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 [3].

Residual Vulnerabilities in S-BGP

Despite the extensive security offered by S-BGP, architectural vulnerabilities exist that are not eliminated by its use. For example, an S-BGP router may reassert a route that was withdrawn earlier, even if the route has not been readvertised. The router also may suppress UPDATEs, including ones that withdraw routes. These vulnerabilities exist because BGP UPDATEs do not carry sequence numbers or time stamps that could be used to determine their timeliness. However, RAs do carry an expiration date and time, so there is a limit on how long an attestation can be misused this way. S-BGP restricts malicious behavior to the set of actions for which a router or AS is authorized, based on externally verifiable, authoritative constraints [4,5,6].

3.2 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.

Designers of soBGP addressed four goals when designing it. [5]

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. Does the originator, or owner, of the destination, authorize the peer advertising the route 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. So soBGP targets to achieve only the first two goals.

soBGP adds a new message type SECURITY to current BGP protocol. BGP speakers to share three different types of certificates use SECURITY type message. 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.

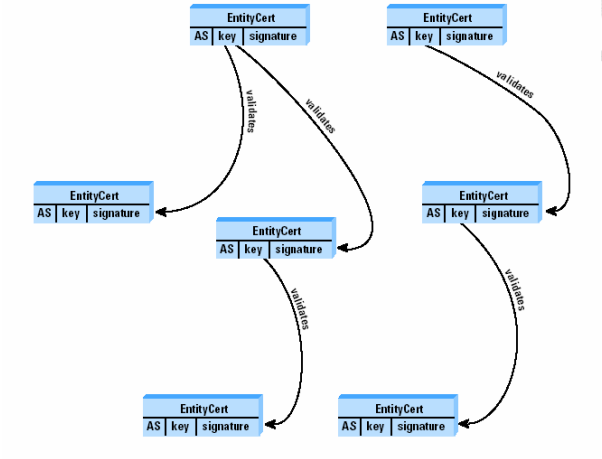

There are three types of certificates in soBGP.

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. A third party, validating who an entity is within the routing system signs EntityCerts. 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. Originator of policies signs this certificate 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) that 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 knows 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.

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.

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 a 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:

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.

3.3 IRV (Interdomain Route Validation)

IRV defines a service that protects against completely ruined or misconfigured ASs, 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.

4. Proposed solutions

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:

1. 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 behavior. Transitive trust in parties causes mistakes to propagate.
2. Solutions must demonstrate similar efficiency, performance and reliability as the other parts of BGP.
3. The requirements of a solution should scale well within BGP.
4. Integrity and authenticity of BGP messages should be guaranteed at the traffic (Classified as hop integrity).
5. A BGP router should be able to verify the owner of each prefix that authorized the origin AS (Classified as origin authentication).
6. 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).

Improvements suggested based on the above facts

Literature shows a number of ways to secure BGP. The current secure version uses MD5 Signature Option. This is discussed more in RFC2385 and is currently mandatory in all implementations of BGP [7]. Implementation of MD5 signature option ensures integrity of the message transmitted and peer entity authentication as long as we can assume that the key has not been compromised. This requires the MD5 algorithm to be secure and that the key used to secure the communication is well protected and difficult to guess. However the MD5 signature option is not enough for all situations.

Research shows that there should be the below 5 changes which could be crucial.

These changes are

1. Encryption of all BGP messages between peers using keys exchanged at BGP link establishment time.
2. Addition of message sequence number

3. Addition of UPDATE sequence number of timestamp
4. Addition of PREDECESSOR path attribute indicating the AS prior to the destination AS for the current route
5. Digitally signing all unchanging UPDATE fields at the point of origin.

Note: BGP uses four different types of messages. These types are:

1. OPEN
2. UPDATE
3. NOTIFICATION
4. KEEPALIVE

If these changes were to be implemented the security of BGP would increase considerably. Use of key exchange during the link establishment time would guarantee that good keys are used and that they are changed often enough. Also encryption of all messages would provide confidentiality (even though it is not always needed). Message and UPDATE sequence numbers would protect against replay attacks. The PREDECESSOR path attribute would allow verification of the path information and digital signatures of unchanging UPDATE fields would not only provide authentication and integrity between BGP peers but also of the full AS_PATH.

However, the uptake of these changes would require all of the BGP speakers to be updated. Therefore, at least possibility to negotiate the options used should be possible between BGP peers to get over the transition time.

The fact that BGP is run on top of ordinary TCP/IP allows the use of any security methods available on TCP/IP. Mainly this would allow the use of IPsec, which could be used to authenticate and secure BGP sessions. IPSec could be run in ESP (Encapsulated Secure Payload) mode, which provides both authentication and integrity and could also be used to encrypt the entire payload. However, again it would take time to implement IPsec in all BGP speaking hosts.

The method also includes two other means of securing BGP. These are a new path attribute, attestations, which establishes that the subject of the attestation is authorized by the issuer to advertise a path to the specified blocks of address space and the use of PKI Public Key Infrastructure) certificates. They describe in detail how a PKI certification tree could be build and how the attestation attribute and PKI certificates could be used to validate routes. IPsec would be used to prevent an active wire-taper from spoofing route withdrawals or replaying intercepted UPDATE messages.

Conclusion

Introduction of these changes would address a number of BGP's vulnerabilities. However, a malfunctioning BGP speaker could still disturb Internet traffic and the lack of UPDATE sequence numbers could still cause BGP peers to reassert routes that have been withdrawn earlier.

Again also the implementation of this method will require changes.

- First of all a method to distribute PKI certificates to BGP peers must be established as well as the entire certification tree should be implemented.
- In addition to this IPsec should be implemented on all peers and the attestations path attributes should be added.

Therefore this cannot be implemented overnight but work must be done to make this possible. This technology is developed and is available as open source. Anyhow a large-scale deployment will require coordination of all parties to implement required changes and to build the PKI certificate tree.

5. References

1. IETF Draft on BGP vulnerabilities - draft-ietf-idr-bgp-vuln-00

2. RFC 2725 - Routing Policy System Security

3. Security In Inter-domain Routing - Tuna Vardar - Helsinki University of Technology - T-110.551 Seminar on Internetworking 2004 - 26-27.04.2004

3.1 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.

4. What's BGP got to do with Internet security? by Marcia J. Wilson, CCSP Staff Writer, March 30, 2004

5. Interdomain routing with BGP - Issues and challenges - Olivier Bonaventure - Department of Computing Science and Engineering - Université catholique de Louvain (UCL)

6. Origin Authentication in Interdomain Routing - John Ioannidis, Patrick McDaniel, William Aiello - AT&T Labs Research

7. Security of Inter-Autonomous Systems Routing - Iikka Väkiparta - Helsinki University of Technology - Iikka.Vakiparta@iki.fi

8. The Internet Protocol Journal – September 2003 – Volume 6, number 3.