

# Addressing Security Issues in the Border Gateway Protocol

Nicholas Barton  
UC Riverside  
bartonn@cs.ucr.edu

Jorge Mena  
UC Riverside  
jmena@cs.ucr.edu

March 20, 2005

## Abstract

The Internet is a critical component in today's society. This makes it ever so important to have a secure routing protocol. In the current Internet routing protocol, BGP, security was only an afterthought as it was designed on the basis of trust. With trust, misconfigurations and attacks can disrupt the Internet easily. Therefore, many proposals have been made to make BGP more secure. In this paper, we explain what attacks can occur, then perform a survey of many types of proposals that have been presented in the past. We lastly present our ideas to secure the BGP protocol. We believe that the key of success for the BGP protocol is to provide one with the characteristics of backwards compatibility of the existing version, scalability, enhancement of security, and easy deployment. These features are the goals of the protocol we present in this paper. Our ideas guard the AS Path from alterations using cumulative authentication and then make sure that the origin AS is able to advertise the prefixes.

## 1 Introduction

The Internet today has become a primary way of communication for both the general public and businesses. For this reason, it is clearly important to provide a reliable medium of communication, especially if it involves the transmission of sensible or private data for any of the parties; to obtain such reliability, it is necessary first to understand how the Internet works, identify its current problems, then address them. The Internet is a global network composed of interconnected networks. As such, it is necessary to use protocols to provide routing of packets between these networks to establish communication. In order to simplify this task, Autonomous Systems (AS) of networks have been established to manage a group of networks by the same system administrator and have common characteristics such as owned by the same entity, i.e. Internet Service Provider, etc. These protocols are classified as intra-domain and inter-domain routing protocols. Intra-domain protocols are in charge to provide routing within a particular AS, while inter-domain protocols aim to route among autonomous systems. The Border Gateway Protocol (BGP) is an inter-domain, path-vector routing protocol

and its description is given in [12]; currently it is at version 4. The goal of this paper is to address the inherited security issues that exist in BGP as well as to review current proposed protocols and describe our own.

Now we provide some definitions to clarify what we mean, based on RFC 3552 [9]. When we say security, we mean communication security just as it is defined in [9]. It is composed of Confidentiality, Data Integrity, and Authentication. Confidentiality means that when communication is established, no unauthorized entity is learning the information being communicated. Data Integrity means that the receiver in a communication received data that the sender intended to send. Authentication means that both ends of the communication can be assured that they are communicating with whom they wanted to. Routing is the process by which nodes exchange topological information to build correct forwarding tables, as defined in [8]. Therefore, by routing security we mean routing that has addressed the communication security issues described above.

Policy and scalability are two features that are desirable in BGP, but often appear in opposite sides of a trade-off. Policies represent the business relationships among ASes, which must cooperate to achieve global reachability [1]. These policies help control the outgoing traffic of an AS and to a certain extent, the incoming traffic. Scalability refers to the ability to provide routing as the network size increases. In BGP, scalability is achieved by the use of aggregation of routing information [1], of the summarization of information to avoid the transmission of large routing tables. The trade-off between the efficient use of policies and scalability is evident. ASes have better understanding of the network topology when they transmit more specific prefixes, thus, they can design better routing policies. However, this means that larger tables (less aggregation) are needed, which hurts scalability.

Finally, it is necessary to describe the Internet Threat Model. Just as RFC 3552, we assume that both ends of communication have not been compromised when communication is attempted. However, the communication channels are assumed to be almost

under control of an attacker, in other words, the attacker can read any packets as well as change and insert forged ones into the wire [9]. For the purposes of this paper, this definition is enough, but we follow the same guidelines of [9].

**Outline:** Our paper is organized in the following way. Section 2 explains the security flaws and vulnerabilities of BGP. Section 3 introduces various counter-measures that have been proposed. This section is divided into Section 3.1 Active detection techniques, Section 3.2 Passive detection techniques and Section 3.3 Other detection techniques. Section 4 describes our ideas to make BGP secure. In Section 5, we provide some discussion. Lastly, in Section 6 we give concluding remarks.

## 2 BGP Attacks and Vulnerabilities

In [7], Nordstrom and Dovrolis provide a study of BGP attacks and its countermeasures in an attempt to warn the routing community about the need to address the security vulnerabilities of the interdomain protocol. They identify four attack objectives: prefix blackholing, traffic redirection, traffic subversion, and routing instability.

Prefix blackholing occurs when a prefix is unreachable from a large portion of the Internet. When traffic is forced to take a distinct path to arrive to a distinct destination in order to cause damage to the true destination is called traffic redirection. A special case of redirection happens when the traffic is forced to take a distinct path with the purpose of learning or modifying the data in the traffic, even if this traffic is forwarded to the correct destination; this is called traffic subversion. Route instability refers to the constant updates and withdraws advertisements sent for the network, with the objective to cause route dampening or cause large convergence delays. Although attacks can be made with an objective like the ones described above or unintentionally due to policy misconfiguration, it is in the nature of the BGP protocol where the problem arises.

Due to the nature of the BGP, it is relatively easy for anyone to obtain routing messages. It is precisely for this reason that there exists public banks of BGP messages such as RIPE. Smith and Garcia-Luna-Aceves provide one of the first reviews of BGP in terms of security since the introduction of version 4 in [10]. It is easy to disclose routing messages that transit in the Internet due to the lack of confidentiality of peer links and the level of trust among BGP speakers [7]. In later attempts to remedy this, some cryptographic mechanisms have added security to the protocol, such as the MD5 digest, but none have become part of the standard protocol yet. The reason of this is that there is no point on securing the BGP messages in the Application layer if the underlying layers (TCP) are not protected. Ultimately, the disruption of routing messages are caused by the exploits of the

vulnerabilities existing in the protocol; these vulnerabilities are the lack of access control, authentication, and integrity of the BGP contents [7], as well as the defenseless underlying layers; in BGP, this is the TCP protocol.

Besides the natural problems of BGP, there exists problems that are specific to the management of policies and those that affect the scalability of the protocol. [1] Finds policy oscillations and weak security as part of policy-induced problems. Policy disputes cause oscillations that can lead to an unstable network that may never converge. For example, an AS could be misconfigured to send updates to reach a location in the Internet that is better than having direct access to the network. Traffic will always take this path, rather than the direct access and will never arrive to the destination. This could be a misconfiguration of the system administrator or an attack with the objective to create a blackhole to a particular IP prefix. Other policy problems are Multiple Origin AS (MOAS), studied by Zhao in [17], where multiple ASes claim to be the origin of a particular IP prefix. De-aggregation could be part of a policy in an AS if it is using multiple connections to connect to the Internet (multihoming). If misconfigured, de-aggregation can represent a problem because traffic to a specific prefix could use a less efficient path to reach its destination simply because the longer prefix match is the primary way to select a route in the protocol.

Weak security problems are natural to BGP since it lacks of it at all. This means that there is no way for a BGP router to guarantee with to whom it is talking, and even less to know if its peer has authority over the IP prefix space it includes in its updates. Other security issues refer to the integrity of the data in the messages transmitted. For example, even if an AS has been authenticated, if some method some method has been implemented for such thing, then the data that this AS sends to its peer could be subject of modification (subversion intent), by an attacker. As a result, it is necessary to address the issue of data integrity in order to verify the validity of the path advertised.

Scalability is also desirable in BGP; after all, it is the protocol that connects the network of networks. The main problems with scalability caused by BGP are due to the lack of topology details. In [1], there are three main scaling techniques: ASes are represented as single nodes in BGP, route reflection inside an AS, and aggregation of IP prefixes. These techniques abstract the characteristics of a particular AS in order for the protocol to scale better. However, such techniques make it more difficult to find and analyze the origin of problems when a misconfigured or compromised BGP router introduces bad routing updates. For example, when de-aggregation is used as a policy to improve traffic flow, such action goes against the scalability goal of BGP. It is because of this that there is a trade-off between scalability goals and policy goals.

These vulnerabilities contribute to the lack of guarantee of correctness as information travels among ASes. Also, BGP has no way to detect the source of instabilities of messages. For this reason, it is easy to introduce misconfigurations in the routing advertisements without any mechanism to detect its source and correct it. To these problems, we add the convergence delays caused by routing instabilities mentioned above, which may provoke that the protocol does not converge after all; thus, more complex to analyze and correct. Finally, we add to this the fact that the Internet is dynamic, so this all the above problems multiply in terms of complexity.

## 3 Countermeasure Mechanisms

### 3.1 Active Detection Methods

The methods described below physically alter the BGP Update packet so they can send information to subsequent ASes. The information that is sent to subsequent ASes is in the form of cryptographic methods that create signatures that attempt to secure some or all aspects of the BGP Update packet. This information could be located in the path attributes field which is allowed in the current BGP protocol or in changes to the BGP packet layout which will break the protocol. Any protocol that breaks the current protocol will ensure little to no acceptance by the Internet community.

#### 3.1.1 Secure BGP (S-BGP)

Kent et al proposed Secure BGP [5] which utilizes public key cryptography to add robust authorization and authentication to BGP. The goals of S-BGP was to protect unauthorized prefix advertisements and protect the AS Path from truncation and modification in the BGP Update packet. To achieve this, several modifications are done to the protocol, but backwards compatibility is achieved. This approach involves three additions: (1) two parallel Public Key Infrastructures (PKI) which are certificate hierarchies, (2) a new transitive path attribute and (3) the use of the IP Security protocol suite (IPSec).

S-BGP utilizes two parallel certificate hierarchies. The first is the AS ownership and router PKI which has large entities signing (using their private key) AS numbers over to smaller entities where they are ultimately designated to specific routers inside a particular network. This is for validating ownership over ASes. The second is the address space PKI which binds address block(s) to a public key belonging to the entity to which the addresses are being assigned to [5]. These hierarchies will be used to authorize prefix ownership and to validate routes. Both hierarchies have ICANN as the root and these certificate structures are based on existing delegation systems, so deployment and trustworthiness is achievable.

The key additions to BGP are attestations. Attestations are used to ensure that the issuer has allowed the AS in question to advertise the blocks of IP space in question. There are two types of attestations, route and address. Address attestations are used to ensure that the originator AS is allowed to advertise the specified prefixes. This is done by a delegation message being signed by the issuer of the IP space. To give a further AS the right to propagate prefixes in the BGP Update, route attestations are used. A route attestation is where the AS Path thus far is signed. The attestations are stored in the BGP Update as a new transitive path attribute which will prevent S-BGP speakers to modify the AS route path.

When a BGP Update is received, an address attestation is performed to ensure that the originating AS is allowed to advertise the prefixes. Next, each AS signs a route attestation to allow the next AS to propagate the prefixes. This method ensures that the AS Path cannot be modified since doing so would require the private key of the AS right before the one that the malicious AS wants removed. Also, an attacker could not introduce its AS unless a delegation is given to it.

IPSec is used by S-BGP to get around the vulnerabilities of TCP. The problems with TCP is that TCP does not provide data integrity or authentication. Therefore, this protocol utilizes the Encapsulated Payload (ESP) with NULL encryption from IPSec to achieve these on a point-to-point basis. Certificates are needed for ESP and the AS ownership and router PKI provides these.

There are a few problems to note about this protocol. This protocol has a high computational cost. This is because there are many signatures used to certify the BGP Update. These signatures in the form of route attestations are also stored in the BGP Update, thus bloating it heavily. Because of this, the AS owners have been reluctant to use this protocol. Though a definite positive is that this protocol can be used when the origin AS does not use this protocol.

#### 3.1.2 Secure Path Vector (SPV)

Hu et al [4], describes the SPV protocol which is similar to S-BGP but uses symmetric key cryptography to secure the AS Path. SPV uses a combination of one-time signatures, hash trees, and one-way chains to achieve this. This removes the need for long term asymmetric private keys which have long life times and could become insecure if hijacked, and for the routers to perform expensive computations. The routers still need to manage short-lived private key for the one-time signatures.

To authenticate BGP messages from a given prefix, a certificate hierarchy is used. This PKI structure is the same structure as the address space certificate hierarchy utilized in S-BGP. This enforces that an attacker cannot advertise a prefix it doesn't own.

There are three cryptographic mechanisms used in this protocol: (1) one-way signatures, (2) one-way hash chains and (3) hash trees. The AS Path protector is constructed using these mechanisms. One-way signatures are hash functions that are supposed to be computationally infeasible to invert, but if they are used multiple times, the security degrades. These are used to achieve AS Path integrity. One-way hash chains are created by selecting the final value at random, then deriving previous values by repeatedly applying a one-way hash function. These are used to enable authentication of the AS Path. Hash trees reduce the complexity of authenticating a sequence of values to merely authenticating a single value. This is used to reduce the size of the AS Path protector [4].

To secure an AS Path, the prefix owner first creates a sequence of one-time signatures. Each signature is used to secure one AS in the AS Path. The address owner then passes the one-time signature list in the BGP Update onto the next AS. At each forwarding of the BGP Update message, one of the private keys is used up to sign it into the AS Path and removed from the list. Authenticating the AS Path is done by verifying the one-time-signatures. Authenticate the one-time signatures public key using address attestations [4].

### 3.1.3 Cumulative Authentication

As described in Hu et al [3], cumulative authentication provides a framework to assure that the middle of the AS Path is not modified. Though, if a malicious AS that uses this method, a truncation attack could be used. A PATH\_AUTHENTICATOR transitive attribute is added and is used to hold the current ASes message authentication code (MAC). The way they sign this packet is to create a new MAC computed over the previous ASes MAC and the packets immutable fields. This is done until the packet reaches the destination AS. This AS can verify the path depending on how the MAC was initialized. Ie, a MAC of 0 to start the chain. If the reconstructed path authenticator value matches the attributes value, then the packet is authenticated. Assymmetric or symmetric keys can be used on this framework.

There is a problem with this framework. In order to verify the path, a well-known value must be the first MAC in the chain. This means that a malicious node could delete all hops in the AS Path and act like its node is the originator of the packet by recreating the MAC in the PATH\_AUTHENTICATOR attribute.

### 3.1.4 Secure Transport Services

As proposed in Smith, Garcia-Luna-Aceves [10], this proposes many changes that are not backward compatible with the current BGP standard. The changes are as follows: (1) Upon establishing a connection with another AS, a session key is sent so that the path can be encrypted for the rest of the session. This is to

provide confidentiality and integrity of the BGP update. (2) A message sequence number is placed in the header of the BGP Update. This is set to zero at the start of each session, so this is for local purposes only. This will prevent missing update information and replay of updates for the session. (3) To protect against replay attacks, a per BGP speaker Update sequence number or time-stamp is used. This is placed right after the header in the BGP packet. To be a valid packet, the number must be greater than a previously received sequence number. A database on each BGP speaker is needed to keep track of the sequence numbers on a per BGP speaker <domain name, public key> pair basis. (4) A PREDECESSOR attribute is added to the path attributes section of the BGP Update which identifies the AS prior to the destination AS. This allows the integrity of the complete AS path given that a digital signature calculated by the originating router is provided. (5) Lastly, in the marker field of the header, an UPDATE message digital signature is added which signs the update message using the originating BGP speakers key to maintain the integrity of the message. [10]

These proposed changes do look like they will achieve their security goals but at a high cost. This protocol cannot be released incrementally as there are too many changes to the base protocol. A protocol that can be implemented incrementally is key to its survival.

### 3.1.5 Secure Origin BGP (SoBGP)

Secure Origin BGP [15] is a proprietary protocol developed at Cisco. This protocol addresses the most important issues in routing: validates the correctness of routes, the authority of the issuing advertisements, and prevents attacks that come from misconfiguration and intentional insertion of incorrect data. They also argue that the protocol is deployable in the current environment where BGP is being used.

To provide authentication of the peers, they proposes the use of certificates of authenticity called EntityCert. These certificates carry a public key, from a private/public pair of keys to validate the ownership of the certificate; this certificate also includes the AS number of the owner of the certificate, thus tying the certificate to the AS. The problem of this is that the receiving entity must assure the procendence of this certificate; SoBGP addresses this by crating a web of trust, where there is a third party entity that every one trusts and signs the certificates thus providing authenticity to the transmitted certificate. Consequently, the new entities can also sign other certificates, thus creating a web of trust. The problem with this approach is that it needs an existing infrastructure to manage this new kind of certificates, which we believe it is not too complicated; however, in the worst case, a new infrastructure would be needed, which is complicated to deploy.

To provide the correctness of an advertisement, the owner of the IP space issues another certificate to establish the ownership of its IP space it is trying to advertise. These certificates protect blocks of addresses or individual prefixes, and are distributed by AS individually. To check the validity of the certificates, the recipient checks the signature of the certificate against the key in the certificate to be certain that this AS is allowed to issue this certificate.

Finally, SoBGP checks connectivity problems, caused by mis-configurations or intentionally induced, by issuing another kind of certificates, ASPolicyCert. When one AS claims to be connected to another AS, each has one ASPolicyCert that assures the connection with the other AS, and vice versa. So when one AS advertises a path that uses this connection, the recipient can check the validity of the path by checking this certificate at each AS. If at a particular link both certificates assure that both AS are connected, then the path is valid at this point, otherwise, the whole path would be invalid.

## 3.2 Passive Detection Methods

The methods described below use analysis techniques on the BGP Updates to detect anomalies in the internet. These analysis techniques do not alter the BGP protocol in any way. These non-intrusive methods have benefits like fast deployment since the protocol isn't changed, and that they allow the AS owners to keep using the same equipment already being used since no heavy cryptography is being used. There are also negatives as well. These include that none of the proposed methods work in real time and that none propose ways to contact AS owners about anomalies when they arise.

### 3.2.1 Signature-Based Detection

As proposed in Zhang et al [16]; to identify and search routing anomalies, signature-based detection utilizes a set of fixed patterns. To create these patterns, they take consecutively announced routes and assign preference values to them. This value indicates whether the new route is preferred over the previous route, the new route is less preferred, both routes have the same preference or that the route was withdrawn. Using a sequence of these values that are given from a BGP update burst, they define seven patterns. These patterns can indicate transient failure followed by fast fail-over, anomaly in community attributes or aggregation or same length AS path oscillation, or that of normal route changes [16].

To accurately detect anomalies with this method, the patterns have to be updated persistently and well-defined. In their testing, they have noticed that it is difficult to get well-defined patterns. Their system may incorrectly treat consecutive updates as separate events if they are not close enough. Lastly, the patterns described have a minimum pattern length to them, so that

some anomalies could be missed if they don't meet this minimum length.

### 3.2.2 Statistics-Based Detection

As proposed in Zhang et al [16], they use the NIDE/STAT algorithm to perform statistics-based detection. The NIDE/STAT algorithm raises an alarm when the short-term behavior of a subject differs drastically from the long-term profile, which is the expected behavior. A set of detection measures are used to describe a subject's behavior.

There are five measures listed in three categories that classify a subject's behavior. In the intensity category, there are the BGP Updates message arrival frequency measure, and the number of AS paths in a period measure. In the categorical category, there are the BGP Updates type measure and the AS path occurrence frequency measure. Lastly, in the counting category, there is the AS path difference measure [16].

Before these measures are used, they have to be normalized. This is to ensure that each measure is comparable to each other. These five normalized measures are combined to form a score value. When the score value passes the set threshold, an anomaly is flagged.

Statistics-based detection has the benefit that it does not need to know about anomaly patterns in advance as it measures against a historical backdrop. The problem is with the historical backdrop. There has to be a clean data-set for this detection method to be accurate in flagging anomalies and keep false positives/negatives low. This is not an easy task to achieve.

### 3.2.3 Visual and Automated Data Mining

As described in Teoh et al [14], this method utilizes the human visualization system and computers for the automated data mining to achieve near-real-time anomaly detection and analysis. The visualization aspects of this method is that humans can pick out anomalies by recalling related images that are hard for data mining. The data mining aspects are achieved by combining the statistical and signature based detection as described above. There are four modules that describe this system. There is the real-time data collection module which sends its data to the anomaly detection engine which sends its analysis results to the data server and communication module which lastly interacts with the interactive visualization client [14].

The real-time data collection module records raw BGP updates in intervals that are user customizable. The anomaly detection engine takes this data and performs the following analysis: (1) prefix selection to strip out the unneeded prefixes, (2) signature-based detection with parameters that are user customizable and is described above, and (3) statistics-based

detection which is described above and the threshold value is user customizable. The data server collects this analysis data and stores it for use in the visualization client. The client can then pick which IP prefixes, peers and time periods to display. The visualization of anomalous events are done using EventShrubs [14].

To define an anomaly detection event, the statistical method is first used. An anomaly is found if the BGP Update exceeds the threshold value. Then to figure out the length of a single anomaly event, signature based detection is used. This also classifies what type of anomaly the event is. With the help of the visualization aspect, the root cause of an anomaly event may be also described.

### 3.2.4 Topology-based Detection

As proposed in Kruegel et al [6], this method constructs a topology graph of the internet to detect anomalous messages. To construct this graph, they passively monitor BGP Updates and connect ASes together if they can exchange information directly between them. They then utilize the **whois** databases located throughout the world to determine the geographical location of the ASes. Using this connectivity graph, they run an algorithm to classify the ASes as either a *core* or *periphery*.

Nodes that are periphery nodes will have a low degree (two or less) of connectedness. These nodes are pruned using an algorithm. When all the periphery nodes are removed, what is left is the core nodes. Core nodes represent the backbone of the internet or major ISP's and can span across large distances. Periphery nodes represent city networks, universities, or company networks and they span generally small distances.

After the ASes are classified, the periphery nodes are placed into clusters. The algorithm removes the core nodes and the graph now looks like a forest of interconnected periphery nodes and is not fully connected anymore. Each set of periphery nodes that can contact each other without going through a core node is called a cluster. These clusters represent generally small distances because the cost to set up a direct link would be infeasible for small entities.

Using this enhanced connectivity graph, they can now classify what an invalid AS Path announcement looks like. They observe that a valid AS Path will only enter the core infrastructure once, thus an AS Path sequence will look like this: a sequence of ASes located in its local cluster, then a sequence of core nodes, and lastly a sequence of ASes that are located in the destinations cluster. The periphery nodes can also establish a link between two previously unconnected nodes and they must be geographically close [6]. This means that if the AS Path enters the core infrastructure more than once, or that the sequence of periphery nodes contain nodes from more than one cluster, then it is labeled as an anomalous BGP message.

There are problems with this approach. (1) This method is not self updating, so the connectivity graph will get stale over time and need to be reconstructed. This reconstruction can take a day. (2) They also only use the BGP update messages and ignore the BGP withdrawal messages. This makes the connectivity graph slightly erroneous before it is used. (3) The algorithm they use to classify the nodes as either a periphery or core node is only correct most of the time. This is partly because of sparse core networks in growing countries. This could lead to the method flagging BGP messages as anomalous that are in reality correct BGP messages. (4) Lastly, the term geographically close is somewhat ambiguous. People have different perceptions on what is close or not.

## 3.3 Other-Detection Methods

These methods use alternative techniques outside what is meant to be a passive or active based detection method.

### 3.3.1 Interdomain Routing Validation (IRV)

Presented by Goodell in [2], the IRV architecture is used with BGP to validate its data and thus obtain additional routing information relevant to an AS. It uses a querying system that stores both static and dynamic information in a centralized location inside its AS. This location is under the responsibility of the IRV architecture to validate the data in a query and also obtain information from within its AS in order to keep the data up to date. The authors claim that this protocol protects against rouge, subverted, or grossly misconfigured ASes, as it is used to diagnose routing configuration problems [2]. Finally, the protocol proposes a method to find an IRV inside an AS by creating a well-known registry to distribute contact information.

### 3.3.2 Listen and Whisper

Proposed by Subramanian, et al in [13], Listen and Whisper is a protocol that completely abandon the idea of perfect security and instead, it trades it off for a more easily deployable mechanism with some improved security. The protocol is described in two parts: the Listen and the Whisper. The Listen part is a data integrity verification mechanism that detects reachability problems in the advertisements after the session has started. The features it provides are the following, according to [13]: (1) it is a passive mechanism, (2) incrementally deployable and standalone solution with no modifications to BGP, (3) quick detection of reachability problems for popular prefixed, and (4) a low overhead protocol. Listen constantly records and draws conclusions about the state of a route while monitoring TCP flows.

Whisper is a much more security mechanism that aims to detect invalid routes from misconfigured or malicious routers [13].

The properties of the Whisper part of the protocol, according to the authors are: (1) any misconfigured or malicious router propagating an invalid route will always trigger an alarm, (2) a single malicious router advertising more than a few invalid routes will be detected and the effects of these routes will be contained. To achieve such goals, the protocol does Route Consistency Testing. When an update to a route needs to be sent, two update messages are assembled and are sent to two distinct route paths to the same destination. When both arrive, then the validity of these advertisements is verified. If both are valid, then there is no problem, however, if one is not valid, then it raises an alarm flag. The problem with this protocol is that it does not catch the difference when both routes are valid or invalid. Especially the case when both routes are invalid, no flag is raised for two possible compromised routes.

## 4 Our Proposal

### 4.1 Assumptions

BGP vulnerabilities are evident since security was not a primordial issue during the design of the protocol. To address these vulnerabilities, some countermeasures have been proposed, but have failed in acceptance and deployment. The reason is because some of them require the creation of a totally new infrastructure which makes it very unattractive due to costs, or the protocol addresses partially some subset of these problems.

We propose a protocol that aims to be backwards compatible to the current existing infrastructure in order to facilitate its deployment, but at the same time, adds path verification in the advertisements as well as confidentiality. Providing path verification will make sure that there exists at least one path from the current location of the advertisement to the origin of it, thus addressing immediately problems with misconfigurations or malicious advertisements from compromised routers. In the same manner, the protocol provides IP Space authentication with the use of certificates, thus, avoiding problems with IP address hijacking. To distribute these certificates, we propose to use the hierarchy of ASes to hold and distribute the certificates upon request. ICANN would be in charge of generating a certificate per IP Space that exist in the Internet; then delegate the authority to the major ISPs in the Internet. These will recursively delegate the authority to their clients until the last client at the end of the hierarchy. Also we propose the use of certificates of authenticity, such as the ones provided by VeriSign or E-Trust, to authenticate each BGP router that uses our protocol.

In order to address backwards compatibility to our proposal, we want to create an environment that closely resembles the one that exists today with BGP. Our idea is that if there exists a router that does not want to participate with our proposal, then it can use the current version of BGP and interoperate with other

routers that do use our proposal. The main motivation to our protocol is that everyone should defend itself and not worry for no one else but itself. In other words, this is a defensive proposal and by greedily defending oneself, the whole network becomes more secure.

Before continuing with the individual description of each part of the protocol, we will address some assumptions:

1. End nodes are not compromised
2. Anyone can listen the BGP advertisements
3. Business policies are designed in accordance with the Business decisions
4. All routers must own a certificate of authenticity from a well known source

We require end nodes not to be compromised because it would be practically impossible to guarantee confidentiality. If one router gets compromised, then the malicious individual will have access to the most private information of the protocol that is used to identify and authenticate itself to anyone else.

Business policies such as MED and local preference are assumed that are designed according to the business decisions of the AS. It is important to give this flexibility because we do not want to design a protocol that constrains a business administration in the ways they handle their network. For this reason, we intrinsically assume that these policies are well designed and well implemented in the protocol and that they represent what the business demands. If this is not the case, it is not responsibility of the protocol to correct them, but of the system administrators.

Finally, we require that each BGP speaker owns a certificate of authenticity from a well-known source, such as VeriSign or E-Trust. This is enforced in order to provide the proposal some means of verifying the source of the advertisement to later provide path verification. If one router decides not to have a certificate, then the route that it advertises will be treated as suspicious and such route would not be preferred. Therefore, as the number of routers that do not have certificates of authenticity increases, this proposal degrades the current BGP protocol.

### 4.2 Infrastructure

The main goal of our protocol is to be backwards compatible to the existing implementation of BGP; thus, our protocol will be primarily an addition, rather than a dramatic update.

We require no additional hardware to the existing infrastructure. However, we do require ICANN to issue certificates of ownership for the IP Space in the Internet that is currently under

use. This is not a trivial task, but it is necessary because there has to be some way to know who owns what. This is analogous to the system of ownership of land that exists today. A person owns land if it holds a certificate called land scriptures that describe the characteristics of the land; moreover, there is a copy of these scriptures in a central location called City Hall.

In the case of IP Space ownership, we do not propose to have ICANN as a central distributor of certificate of ownership. When ICANN generates the certificates of ownership of the current IP Addresses that are in use, ICANN should generate only one certificate per ISP that has registered as the owner of this space. By issuing this certificate, ICANN delegates the authority to this ISP over the IP Space. In turn each ISP would generate more certificates of ownership per every subportion of IP Space that it delegates. In this model, the owner of the IP Space is finite at the leaves of this hierarchy, and only delegation pointers are found at the intermediate ISPs. The characteristics of these new certificates of ownership are similar to those in the currently existing certificates of authority X.509; the only difference is the purpose of the certificate. We will not discuss X.509 any further, as it is beyond the purpose of this paper.

We leave to the adopter of our protocol the implementation of how these certificates are stored and managed. We will simply assume that each AS will handle these certificates in somehow, but with the condition that they must be publicly available and that each BGP speaker within the AS will know where to find them. Thus, if a foreigner request a certificate for a particular IP Space that the AS owns to any BGP speaker, the router shall be able to find and send the certificate back to the requester. We preferred to provide this flexibility in order to promote the acceptance of the protocol and easy the update. Also, it abstracts unnecessary implementation details.

### 4.3 Path Verification

Many of the BGP misconfigurations and malicious attacks are caused by the lack of a verification mechanism of the routes created by the BGP messages. Some protocols [5, 4] propose the use of nested signatures that must be verified by each node, all the signatures attached by each hop. On top of that, it is necessary to create a new infrastructure to support such method. Others [2, 13] propose a decrease in confidentiality of the messages with the advantage of having a verifying model for the messages. However, they either introduce a central point of failure by having a centralized system for verification or provide not an absolute testing mechanism.

We propose an improved version of Cumulative Authentication [3] and adapt it to our protocol. Cumulative Authentication uses shared private keys or the TELSA broadcast authentication protocol to sign the new BGP message created with the previous

signature and the immutable fields of the original message. By using a well-known value then any receiver can reconstruct an expected final path authenticator value, given the address list [3], or in the BGP case, the AS\_PATH. The problem with this mechanism is that if one of the intermediate nodes is malicious or is misconfigured, it could start a new BGP message using the well-known value and thus hijack the IP Space of the originator of the message.

Our Cumulative Authenticator has the same goals as the original [3], but we aim to protect the originator of the message. In the field of the BGP message, we add a new peace of information: Origin Contact Information. This is a new field of the new BGP attribute that we are introducing to the current BGP v4 protocol existing today. The signatures created by the Cumulative Authenticator include this new field in its signature, as well as the address list (AS\_PATH) and the immutable fields. However, the problem still exists, any malicious node in the AS\_PATH can include itself as the origin of the announced path. This is resolved in our protocol with the use of Certificates of Authentication and Origin Certificates.



Figure : Cumulative authentication of packet  $p$  to a target  $T$

Using the figure above from [3] as a reference, when target  $T$  received a message  $p$  from node  $E$ ,  $T$  will try to validate the message by using either the private shared key or the TELSA protocol. Then the  $T$  continues the validation of the message by checking the validity of the Origin Contact Information.  $T$  will request the Certificate of Authority from this originator and try to establish a secure connection to this node. If the connection can be established and verified,  $T$  knows that there exists a path to this originator (note, however, this path is not necessarily the one advertised in the BGP message) and it is reachable. What it is left to do is to validate the advertised IP Space in the BGP messages. For this we use the Origin Certificates that are issued by ICANN and have been delegated to the IP Space owners.

### 4.4 BGP Attribute: PATH\_VERIFY

The PATH\_VERIFY is an optional transitive attribute that allows BGP Speakers to add the validation features of our protocol. The Attribute Type is left undefined until further review, but this number is not relevant at the moment. The Attribute Length determines the length of the contents of the attribute. The Attribute Value is the field that contains the required information we describe in our protocol. This information is the Origin Contact

Information and the signature of the AS\_PATH Attribute plus the Origin Contact Information.

## 4.5 Putting it all together

Our protocol is designed to introduce fewer amounts of changes to the currently existing infrastructure, thus making it scalable. We add a new BGP attribute to the existing protocol where we plan to hide all of our changes, thus providing backwards compatibility with the current BGP v4. If one message contains this attribute and does not support it, then it can disregard it and assume that the message contents are valid. This means that the protocol degrades to the current implementation of BGP v4.

We start with ICANN. This authority is the most appropriate entity that should issue the Certificates of Ownership. We believe that it is not a difficult task since it keeps a list of all the organization that owns portions of the IP Space; therefore, it is trivial for ICANN to generate and maintain the certificates. ICANN should contain those certificates that have not been assigned to any organization or are not to be assigned at all.

The certificates provide authority over IP Space, but have the property to authority delegation. We use this property to transmit this authority to the organizations that own the space, and so recursively to the clients of these organizations. This hierarchy invalidates the assumption of having a central point of failure because it is not necessary to contact ICANN all the time. When a router needs a certificate, it will use the ASes in the AS\_PATH to find the next hop AS to which it should sent the request; then the process repeats like so until the certificate is found. The careful reader probably has guessed that this imitates the behavior of the Domain Name System.

In order to communicate among routers to do path verification, routers must establish secure connections. Therefore, all BGP speakers must have a Certificate of Authority from a well-known source, such as VeriSign or E-trust, to guarantee authenticity. Since it is assumed that the entities that will handle these certificates, they will cooperate with ICANN closely, then there is no need of a new infrastructure since the existing one will suffice. Therefore, the assumption that all BGP speakers have a Certificate of Authority is not invalid.

After the establishment of secure connections, BGP speakers will guarantee the identity of the peer they are talking to. This characteristic is useful to verify the Origin Contact Information in the PATH\_VERIFY attribute of our protocol. Once we have confirmed the validity of the origin, all it is left is to verify the validity of the remaining portion of the path. For this we use the idea of Cumulative Authentication in [3] with the modifications explained before. This way our protocol verifies the entire path in the AS\_PATH attribute received with the advertisement.

## 4.6 Optional Deployments

Our protocol abstracts implementation details inside ASes. However, this is left to the implementation of the participants using our protocol, thus guaranteeing flexibility in the protocol. Here we present two options of implementation.


If the participants prefer a centralized environment, BGP speakers can act as interfaces for the AS to the rest of the world and a centralized system could be used to handle the added overhead of our protocol. In this environment, the routers can receive requests for certificates and forward them to the centralized system. These, in turn can return the certificate to the router, which replies to the original request. However, the real gain is obtained when the routers forward the advertisements that use our protocol. The centralized system will be in charge of validating the advertised paths and then distribute them to the BGP speakers. Though this implementation can introduce a point of failure in the network, BGP speakers will have more resources devoted to do routing decisions, rather than validating paths.

A second implementation is to let each router handle the received advertisements in a distributed environment. In this case, each router will validate its own received paths and then distribute them among BGP Speakers, using iBGP. Each router will have to handle the added overhead of verifying the paths but there is no notion of central point of failure.

The advantage of both implementations, however, is that not all advertisements need to be validated if they are repeated. When an advertisement is received, iBGP will distribute the new learned routes. If another router has learned an external route internally, then another BGP Speaker must have already verified this route, therefore, it is not necessary to verify it again.

## 5 Discussion

We attempted to use the SSFNet simulation framework [11] to simulate our protocol. We figured out the basics, but we didn't have the time to figure out how to implement parts of our protocol in this. So, we use this framework on a small topology shown in the next figure to measure convergence time as the delay at each router increases. The delay can indicate authentication and processing of our enhanced BGP Packets. The CPU delay will cover the time needed to get the certificate, process its validity, and readvertise to its peers. This topology consists of six ASes, each having a host for IP space, forming a connected graph.



The following table is a list of convergence times that we received from the simulations on SSFNet compared with the ideal time (no time needed) and increments of two seconds.

Delay (sec)	Convergence Time	% Increase
0	57.3821455	0
2	58.4123463	1.01795333
4	60.4121467	5.28039011
6	92.1788149	60.6402376
8	102.284097	78.2507371
10	105.002451	82.9880192

From the table above we can see that if transmission and validation of the messages were instantaneous, then this topology would converge in about 57 seconds, which we consider it as the ideal time. More realistically speaking, we can see that a delay time of less than 4 seconds will add up to a 5% increase of the ideal time, which means about 3 additional seconds to validate all the paths, all the routers. An exaggerated delay of ten seconds will add an 82% increase of the ideal time, which is not even twice as much the ideal time. This shows that processing delays can significantly lengthen the time for the topology to converge. This is only for a small topology, so the numbers will be substantially higher when used on large networks, like the Internet.

We allow either symmetric key cryptography or asymmetric key cryptography to be used to validate the AS Path. From [4] they note that it takes  $401\mu\text{s}$  to verify a 1024-bit RSA signature and about  $2\mu\text{s}$  to verify a Message Authentication Code that uses symmetric key cryptography. This means there are orders of magnitudes of difference in the time delay between the types of cryptography an implementor could use. Without going into a scenario, we note that on large networks, our protocol proposal will have problems with authenticating all of the BGP Updates in a reasonable time if RSA type cryptography is used. If symmetric cryptography is used then we believe our proposal will be able to be used on large networks.

## 6 Conclusion

In this paper we provide a description of the current security issues of the BGP protocol v4 that is currently in use. Then we surveyed a list of the most promising countermeasures available today; our goal was to describe them and point out where they are

strong and where they fail. Finally, we proposed our own ideas; we borrowed some of these ideas and modified them when we felt that they lacked of robustness or confidentiality, and we designed others completely from scratch. The main goal of our protocol is simple: provide an easy transition from the current version of BGP to ours and at the same time keep backwards compatibility. In a sense, our protocol is greedy; it aims at protecting itself rather than the network, and by everyone protecting itself, the whole network can protect itself. Finally, we analyze the convergence time of our protocol by adding some delay time  $t$  to the router processing of messages. We believe that our protocol has the most important characteristics that the solution of the current BGP must have: backwards compatibility, scalability, and easiness of implementation, with the addition of security.

## References

- [1] Nick Feamster, Hari Balakrishnan, and Jennifer Rexford. Some foundational problems in interdomain routing. In *Proc. ACM SIGCOMM HotNets Workshop* (November 2004).
- [2] Geoffrey Goodell, William Aiello, Timothy Griffin, John Ioannidis, Patrick McDaniel, and Aviell Ruben. Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing. In *Proc. of Internet Society Symposium on Network and Distributed System Security (NDSS'03)* (February 2003).
- [3] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Efficient Security Mechanisms for Routing Protocols. In *Proc. of NDSS'03* (February 2003).
- [4] Yih-Chun Hu, Adrian Perrig, and Marvin Sirbu. SPV: Secure Path Vector Routing for Securing BGP. In *Proceedings of ACM SIGCOMM* (August 2004).
- [5] Stephen Kent, Charles Lynn, and Karen Seo. Secure Border Gateway Protocol (Secure-BGP). In *IEEE Journal on Selected Areas of Communications*, 18(4):582-592 (April 2000).
- [6] Christopher Kruegel, Darren Mutz, William Robertson, and Fredrik Valeur. Topology-based detection of anomalous BGP messages. In *RAID* (2003).
- [7] Ola Nordstrom, and Constantinos Dovrolis. Beware of BGP Attacks. In *ACM SIGCOMM Computer Communication Review*, v.34 n.2 (April 2004).
- [8] Larry Patterson and Bruce Davie. *Computer Networks A Systems Approach*. 2nd Edition. (Morgan Kaufmann, 2000).
- [9] Eric Rescorla, and Brian Korver. Guidelines for Writing RFC Text on Security Considerations. *RFC 3552* (July 2003).
- [10] Bradley R. Smith, and J.J Garcia-Luna-Aceves. Securing the Border Gateway Routing Protocol. In *Proc. Global Internet'96*, London, UK, 20-21 (November 1996).
- [11] Scalable Simulation Framework (SSFNet). <http://www.ssfnet.org/>

- [12] John W. Stewart, III. *BGP Book: inter-domain routing in the Internet*. (Reading, Mass: Addison Wesley, 1999).
- [13] Lakshminarayanan Subramanian, Volker Roth, Ion Stoica, Scott Shenker, and Randy H. Katz. Listen and Whisper: Security Mechanisms for BGP. In *First Symposium on Networked Systems Design and Implementation* (2004).
- [14] Soon Tee Teoh, Ke Zhang, Shih-Ming Tseng, Kwan-Liu Ma, and S. Felix Wu. Combining Visual and Automated Data Mining for Near-Real-Time Anomaly Detection and Analysis in BGP. In *Proc. of the ACM Conf. on Computer and Communications Security Workshop on Visualization and Data Mining for Computer Security* (2004).
- [15] Russ White. Securing BGP: soBGP-The Internet Protocol Journal - Cisco Systems. At [http://www.cisco.com/en/US/about/ac123/ac147/ac174/ac236/about\\_cisco\\_ipj\\_archive\\_article09186a00801c5a9b.html](http://www.cisco.com/en/US/about/ac123/ac147/ac174/ac236/about_cisco_ipj_archive_article09186a00801c5a9b.html).
- [16] Ke Zhang, Amy Yen, Xiaoliang Zhao, Dan Massey, S. Felix Wu, and Lixia Zhang. On Detection of Anomalous Routing Dynamics in BGP. In *Proceedings of Networking* (2004).
- [17] Xiaoliang Zhao, Dan Pei, Lan Wang, Dan Massey, Allison Mankin, S. Felix Wu, and Lixia Zhang. An Analysis of BGP Multiple Origin AS (MOAS) Conflicts. In *Proc. ACM SIGCOMM Internet Measurement Workshop* (2001).