Bachelor's Thesis (UAS)

Degree Program Information Technology

Specialization Telecommunication and Networking

2012.

KALU OKPO UME

BGP SIMULATION

TURUN AMMATTIKORKEAKOULU
TURKU UNIVERSITY OF APPLIED SCIENCES

BACHELOR'S THESIS | ABSTRACT

TURKU UNIVERSITY OF APPLIED SCIENCES

Degree Programme | Telecommunication and Networking

August 2012 | 51

Instructor(s): Riikka Kumala

KALU OKPO UME

BGP PROTOCOL SIMULATION

Abstract

An overview of BGP Simulation

The increase in network traffic due to over dependence on the Internet has made the Border Gates Routers (BGR) vulnerable to attacks. This over the years has been a major issue of discussion within the network engineers and the user communities.
These draw backs have led to the proposing of the use of BGP network Simulators to allow for design, proper analysis and evaluation of network behaviors under different scenarios. The use of a centralized Remote Router Validator (RRV) for verifying router attributes from different Autonomous System (AS) is implemented as a solution to solve these drawbacks (issue of false identity and announcement) by adversaries pretending to be who they are not. The RRV checks all parameters sent to it by comparing it with the attributes cached locally in its database and validate them by appending a signature on the packets and sending them back to the sending AS if all attributes are matched.

FOREWORD

This thesis was borne out all the courses I took in Communication and backed by my knowledge in Cisco Certified Network Professional (CCNP).Route This work wouldn't have been possible without all the staff members of the Cisco Lab and, my supervisor who was always there for me.

I would like to thank all the members of my family for their support, my course mates who were the reason that we are here today and to a special friend in the person of Arto Toppinen who believes in me.

August 2012 | Turku.

TABLE OF CONTENTS

FIGURES

TABLES

Notations

AH…… Authentication Header protocol
AS……Autonomous System
ASN……Autonomous System Number
BBN…… Bolt, Beranek and Newman

BGP……..Boarder Gate Protocol
BGR..... Boarder Gate Router
ESP……..Encapsulating Security Payload
GTSM…….....Generalized TTL Security Mechanism
GUI……..Graphical User Interface
ICMP……Internet Control Message Protocol
IANA.……Internet Assigned Numbers Authority
ISP……Internet Service Provider
IP……...Internet Protocol
IPsec.....Internet Protocol Security
MD5 ……Message-Digest algorithm 5
NRI…….Network Reachability Information
PKI……Public Key Infrastructure
RIR……..Regional Internet Registry
ROA……..….Route origin validation
RPKI……Resource Public Key Infrastructure
RRV……..Remote Router Validator
sBGP……..Secure Boarder Gate Protocol
soBGP……Secure Origin Boarder Gate Protocol
SHA-1……..Secure Hash Algorithm
TCP…….Transmission Control Protocol
TTL…….Time –To-Live
VPN……Virtual Private Networks
VNE……..... Virtual Network Environment
VS……...Visual Studio
XML……Extensible Markup Language

CHAPTER 1

INTRODUCTION

When the boarders of a country are weakened, infiltrated, become porous and left unprotected, the security of that country becomes a serious issue or compromised.
The same applies to the world of Boarder Gate Routers (BGR) where there are over 40,000 route announcements on the Internet today [BGP Origin Validation].

Boarder Gate Protocol (BGP) routers act as a power or water mains distribution line directing generated traffics. The inter domain distribution of generated traffic streams, whose source and destination are within domains is called AS.
Since all generated internet traffic from any domain goes through the BGP, edge routers act as the main transit points for each AS or network. This poses serious security issues to network engineers and the Internet at large, hence the need for proper security measures.

Regardless of the large deployment, and importance of BGP on the overall performance of the Internet, it is still vulnerable to many security attacks [Rekhter et al, 2006 and Murphy, 2006]. One of these problems is the fact that any BGP speaker can inject unintentionally or maliciously invalid information about reachability of prefixes that will be propagated to the rest of the system. These errors may have several implications regarding security, reachability or stability of the network. This brings us to the saying that every network is as strong as its weakest point.

1.1 Why BGP Simulators?

The inherent drawbacks of theoretical measurements and analytical methods have led to a considerable increase in the use of simulations for BGP routers networks analysis. Despite requiring extensive computations, simulations generally allow for simpler evaluation of more realistic configurations, most of which are simply beyond the power of theoretical approaches.

In a simulation environment, there are total controls of the system to achieve desired goals, for example:

- ✓ Predict the course and results of the actions.
- ✓ Understand why observed events occur.
- ✓ Identify problem areas before implementation.
- ✓ Explore the effects of modifications.
- ✓ Confirm that all variables are known.
- ✓ Evaluate ideas and identify inefficiencies.
- ✓ Gain insight and stimulate creative thinking.
- ✓ Communicate the integrity and feasibility of the simulation.

Also, it is still much more cost-effective to test the proposed and new extensions of BGP using simulators than to deploy them in a real system.

The process of prototypical development by way of simulators requires a thorough study and proper understanding of the system in question. The process of frequently modeling with simulators always discovers inherent problems that were unknown, not well-understood or over-looked over the past few years.

1.2 Contribution

The aim of this thesis will be to achieve the following:

- look at the existing BGP securities and its drawbacks,
- design a BGP simulator using C#,
- test the implementation of the network simulator,
- and address BGP trust security issues using Router Origin Authentication and the RPKI management system using a centralized Remote Router Validator.

1.3 How Does BGP Work?

BGP uses TCP as the transport protocol, on port 179. Two BGP routers form a TCP neighbor connection between one another. These routers are called peer routers. The peer routers initiate the exchange of messages to open and confirm the connection parameters.

BGP routers exchange network reachability information and it is mainly an indication of the establishment of the full paths (the paths are BGP AS numbers) that a route must take in order to reach the destination network. BGP is a distance vector protocol, which chooses the shortest path to the destination based on the number of AS it has to traverse. This does not depend on the physical hops, because we do not know the internals of the AS and how they will actually transmit data.

Any two BGP routers that form a TCP connection for the purpose of exchanging BGP routing information are called "peers" or "neighbors". BGP peers initially exchange the full BGP routing tables. After this exchange, the peers send incremental UPDATES as the routing table changes.

BGP keeps a version number of the BGP table and this is the same for all the BGP peers and changes whenever BGP UPDATES the table with routing information changes. BGP routers send out keep-alive packets to ensure that the connection between the BGP peers is still alive and notification packets in response to errors or special conditions.

1.4 When and When Not to Use BGP

BGP is a very complex routing protocol and does not always need to be implemented in all scenarios in order to route to different autonomous systems. The use of static and default routes is an alternative to use in place of BGP. The questions that need to be answered are, when should one use BGP and when should one use default or static routes? The answer to these questions depends on the scenario [Cisco-BGP Case Studies].

1.4.1 When to use default and static routes

Memory and processing power: when the routers in the network do not have enough

memory and/or processing power and the number of routes contained in the Internet is huge, it can introduce unnecessary delays in the network.

Single – homed AS: if an AS has only one exit point to the outside network and there is no requirement to enforce any policies (how the routes are redistributed between routing domains) [Cisco-CCNP Route].

Bandwidth connecting two AS: BGP depends only on the stability of the Internet. If the Internet is stable, then the link bandwidth and router CPU cycles consumed are due to the exchange of the BGP KEEPALIVE messages (every 30seconds) exchanged only between peers [Kent et al, 2000]. If the Internet is unstable, only the changes to the reachability information (that are caused by the instabilities) are shared between routers via UPDATES [Rekhter et al, 1991]. The greatest overhead per UPDATE message occurs when each UPDATE message contains only a single network and this can be reduced by grouping multiple networks into a single UPDATE message, thus significantly reducing the amount of bandwidth required. Most of the bandwidth is consumed by the exchange of the Network Reachability Information (NRI).

Table1. Illustrates a typical amount of bandwidth consumed during the initial exchange between a pair of BGP speakers based on the above assumptions (ignoring bandwidth consumed by the BGP Header) [Kent et al, 2000].

Table1. Bandwidth Consumed.

# Network	Mean AS Distance	#AS's	Bandwidth
2,100	5	59	9,000 bytes
4,000	10	100	18,000 bytes
10,000	15	300	49,000 bytes
100,000	20	3,000	520,000 bytes

1.4.2 When to use BGP

Policy Requirement: when there is need to enforce inbound and/or outbound policies on information entering or leaving the network [Cisco CCNP Route].

.

Multi-homed AS

When the network has multiple connections to different autonomous systems, and there is need to exchange traffic from one autonomous system to another autonomous system (transit AS) [Cisco-CCNP Route].

ISP connections

When connecting to different internet service providers (ISP) to one another.
In general, when there are different policy requirements than the ISP, it is necessary to use BGP to connect to an ISP.

Organization of the Thesis
This thesis will be structured into six chapters and each of the chapters is as follows:

- ✓ Chapter 2 presents some potential problems and inherent vulnerabilities of BGP and its security concerns.

- ✓ Chapter 3 presents a description of solutions implemented within current protocols.

- ✓ Chapter 4 shows results from the simulations of BGP, obtains estimation of the overheads and compares them with the gains arising from the Secure BGP modification.

- ✓ Chapter 5 shows a summary, discussion and future work in this area.

CHAPTER 2

SECURITY AND THREAT MODEL FOR BGP

BGP was designed to enable inter-domain routing within and between trusted Autonomous Networks. However, commercial interests and increase in user communities has caused a drastic growth of the Internet, have changed the initial trend of the network, hence the assumptions of trust present in the Internet's original design now present a security and complex problem.Trust should do more than monitor BGP security threats it detect diagnoses and mitigates them 24/7. This ensures that communication channels between AS are legitimately secured.The loose collaborations that BGP was designed for are fundamentally flawed and different from interactions in the current prevailing environment. With this development the changing models of trust have led to series and forms of problems or attacks on the Internet as will be discussed in this chapter.

2.1 Inherent Vulnerability

The vulnerability issue exists because BGP's architecture is based on trust. For different AS to communicate, the BGP routers identify the quickest and most efficient route for the data to reach its destination. But BGP assumes that when a router announces a best path, it is telling the truth. This gullibility makes it easy for eavesdroppers to manipulate routers into sending malicious traffic [Murphy 2006].

In order to have a clear view and understanding of the BGP's vulnerabilities, it is pertinent to consider a threat model. This model should be able to provide an outline of all the sort of attacks that characterizes the ability of adversaries to attack the protocol. The outline should specify the name of the attack; identify the malicious entities, consequences of the attack, the intent of the attacks and a complete documentation of the finding [Nordstrom and Dovrolis, 2004].

2.2 Attack Models

BGP is a distributed protocol run by hundreds of thousands of routers. Hence, there are many points at which an adversary can mount an attack. Each Autonomous System is indirectly linked to other AS on the Internet. Adversaries can affect routers and networks far removed from their peers by exploiting this scale and interconnectedness [Kelvin et al. 2010].

2.2.1 Hijacking Prefixes

Each AS has one or more routers on the edge of its network which routes traffic to its entire peer ASs. ASs then communicate routing information and establish peering relationships using the Border Gateway Protocol (BGP).This is all done in an effort to allow each AS to make announcements about the IP address space it controls .

IP space is allocated and announced in blocks, so if an AS controls all IP addresses between 10.0.0.0 and 10.255.255.255, then it could announce the block 10.0.0.0/8. The numbers before the slash indicate the IP address mask, and the number after the slash is how many bits of the mask should be considered important.

Allocating addresses in blocks leads to smaller routing tables and fewer route advertisements, as most routers need only know how to direct traffic toward the block of addresses, rather than storing separate routing information for every IP address [Bellovin, 2003]. Since prefixes have variable length, one IP prefix may be completely contained within another (sub-netted). A router may have routing information for two prefixes 10.2.0.0/12 and 10.2.132.0/22, where the first prefix completely covers the second one. To decide how to forward a data packet, an IP router identifies the longest prefix that matches the destination IP address. For example, a packet with destination IP address 10.2.132.37 would match 10.2.132.0/22, since this prefix is more specific than 10.2.0.0/12.

The AS that introduces a destination prefix into the global routing system by announcing the prefix to neighboring ASs is called the originating AS [Karlin, 2006]. In Figure 2.1 AS 0

advertises a BGP route for 10.2.7.0/16 with an AS path of 0 to its upstream provider AS 1, which prepends its own AS number to the front of the AS path before sending the BGP advertisement to other neighbors like AS 2 and 7. However, there is no BGP security mechanism to ensure that a BGP-speaking router uses the AS number it has been assigned, or that the AS has the prefixes it originates. A router can be configured to advertise routes into BGP with any AS number, as long as the neighboring router accepts them. Similarly, a router can originate routes for any destination prefix, including very small address blocks (e.g., 10.2.132.4/30) and address blocks it does not hold. The neighboring router will accept these advertisements unless configured to do otherwise, based on prior knowledge of the acceptable prefixes or prefix lengths. This makes the routing system extremely vulnerable to misconfiguration or malicious attack [Karlin, 2006]. An AS can advertise a prefix from an address space unassigned by or belonging to another AS, an action known as prefix hijacking. Neighboring ASs receiving this announcement may select this route and direct traffic toward the wrong AS; these ASs may, in turn, advertise the BGP route to their own neighbors [Zheng et al. 2009]. In the example in Figure 2.2, if an adversary AS 4 announces 10.2.7.0/16 and all ASs select shortest-path routes, then ASs 5 and 6 mistakenly choose routes through AS 4 rather than AS 0.

Prefix hijacking can happen in one of three ways - a block containing:

- ✓ unallocated space can be announced,
- ✓ a sub-block of an existing allocation can be announced,
- ✓ or a complete announcement for exactly the same space as an existing allocation can be announced.

Sub-block hijacking is the easiest, most steady attack and one of the biggest concern to network operators. No matter what the style of attack, the announcements will probably be short-lived, relative to legitimate announcements. This is because attackers, wishing to hide their tracks, will withdraw their announcement once they are done, as opposed to legitimate network operators who generally strive for as much uptime as possible.

2.2.2 De-Aggregation

De-Aggregation or sub-prefix hijacking is one of the most virulent methods of spreading false information. When used as an attack, it breaks up an address block into a number of more specific (i.e., longer) prefixes. This is similar to prefix hijacking and, occurs when the announcement of a large block is fragmented into a collection of announcements for other blocks (mostly by misconfiguration) [Mahajan et al 2002]. Since BGP performs longest prefix matching, whereby the longest mask associated with a prefix will be announced for routing purposes.

De-aggregation negatively impacts on the performance of BGP and indirectly the network by increasing the size of BGP tables and flooding the network with redundant and sometimes incorrect bogus UPDATES. If an AS falsely claims to be the origin of a prefix and the update has a longer prefix than others currently in the global routing table, it will have fully hijacked that prefix. Not only will neighboring routers believe this update, but they will flood the false update to their peers. This flooding eventually propagates the attack throughout the Internet.

2.2.3 Subversion

Subversion is a special case of redirection in which the attacker forces the traffic to pass through a desired link; unlike redirection, the intention is not to change the destination but to have access to the path of the traffic so as to eavesdropping or modify the data [Nordstrom and Dovrolis, 2004]. In subversion attacks, the traffic is still forwarded to the intended destination, making it more difficult to detect.

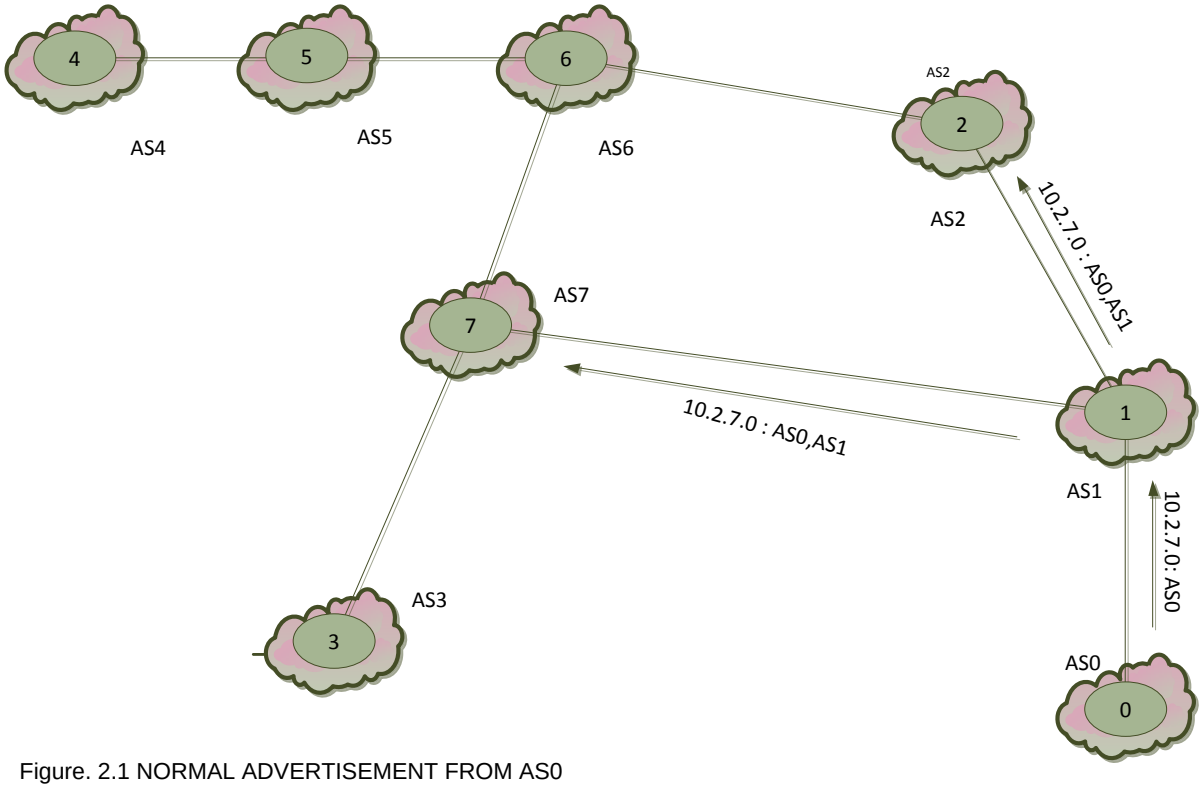

Figure. 2.1 NORMAL ADVERTISEMENT FROM AS0

2.2.4. Redirection

Redirection occurs when traffic going to a particular network is forced to take a different path and to be delivered to illegal destination [Ballani, 2007]. One objective of redirection attacks is that the illegal destination impersonates the original destination to receive confidential information. Another objective may be to redirect excessive amounts of traffic to a certain link or network and cause congestion collapse.

The Pakistan Telecom attack on YouTube similarly involved announcing a smaller address block that effectively misdirected all packets meant for the YouTube site to the wrong place, where they were dropped [BBC News]. These till date is believed not be necessarily malicious attacks, but simply innocent configuration mistakes by the network operators [Nordstrom and Dovrolis, 2004].

2.2.5 Black-Holing

If the malicious AS4 in Figure 2.2 simply drops all packets destined to the hijacked addresses (AS 0), the effect is called a black hole and the destination seems unreachable to some part of the AS that believe the bogus BGP routes. It can be used to intentionally block reserved or unallocated blocks, or it can be maliciously done by an attacker in order to disrupt service [Horn, 2009, Nordstrom and Dovrolis, 2004].

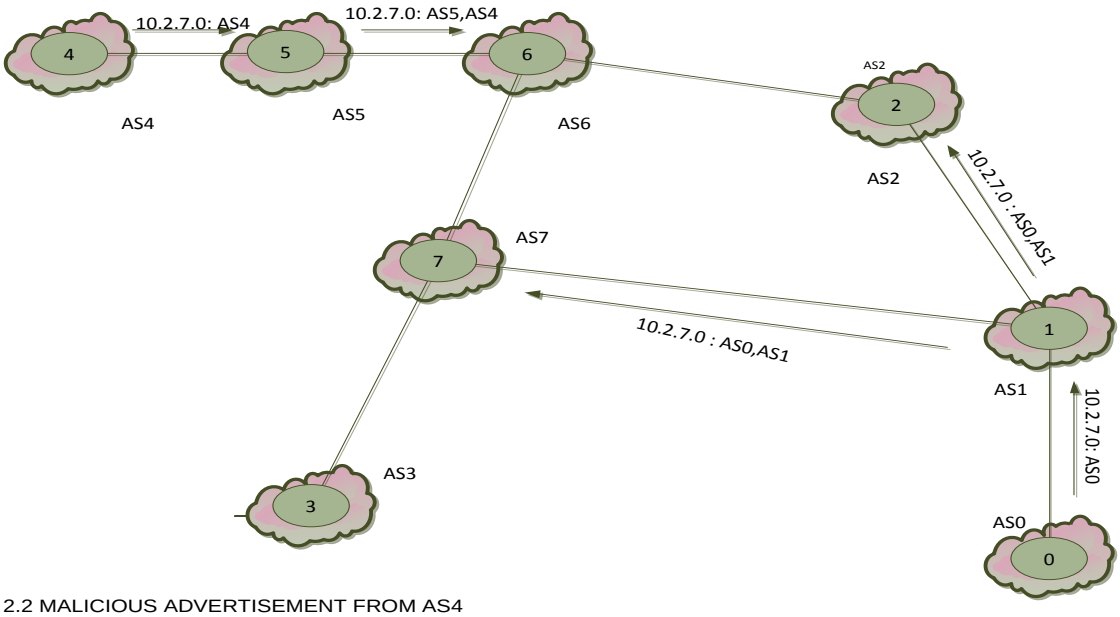

Figure. 2.2 MALICIOUS ADVERTISEMENT FROM AS4

2.2.6 Contradictory Advertisements

A contradictory advertisement is the propagation of different routing announcements sent by the same AS to different BGP peers. It is a legitimate technique for inter-domain traffic engineering and can be used by an attacker for modifying the AS Path so as to control the flow of traffic [Nordstrom and Dovrolis, 2004].

BGP offers a number of attributes that can be used in the AS path selection process to choose the most preferred path to a certain destination. For instance, in a multi-homed network, AS can send UPDATES with a padded AS-PATH to one of its providers so that the link to that provider is only used if the primary link to another provider fails.

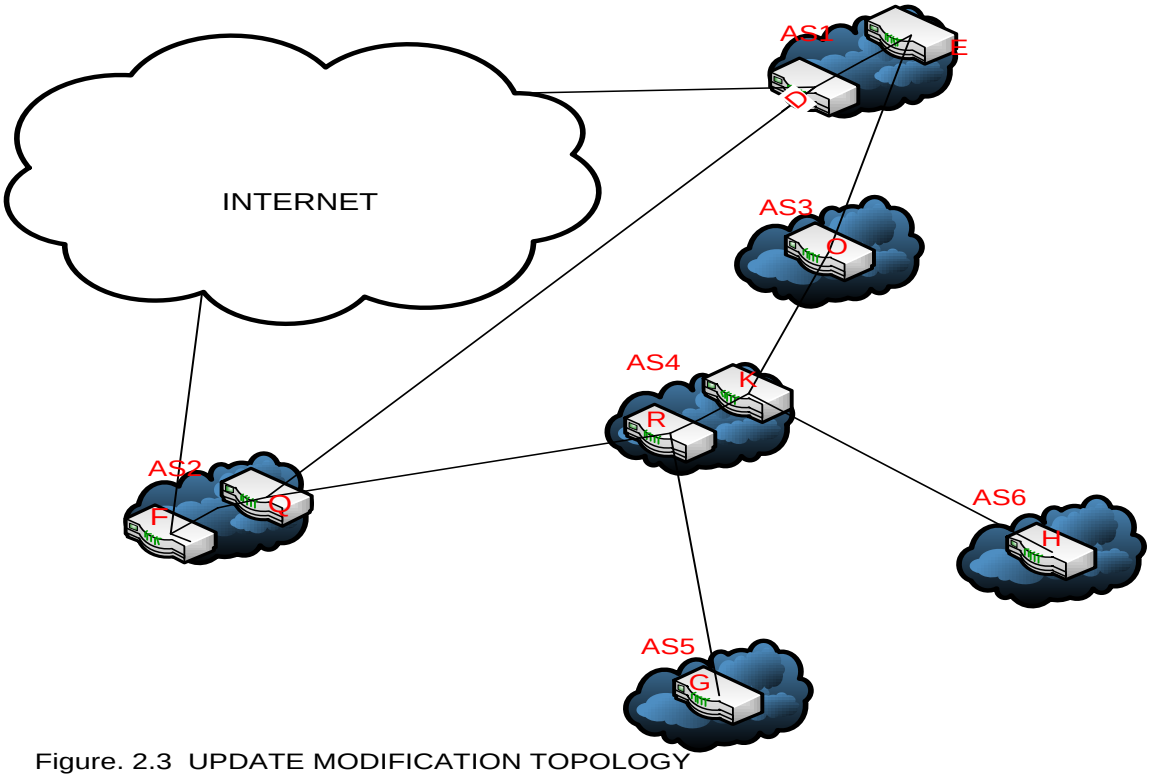

Figure. 2.3 UPDATE MODIFICATION TOPOLOGY

Suppose that AS4 uses link Q-R as is its primary connection to the global Internet, and link O-K as backup. To accomplish this policy, the AS4 border router K can pad or extend the AS-PATH of the UPDATES going to AS3 with several repetitions of its own AS number.

The AS-PATH for AS5 and AS6 to AS4 will then be AS5, AS4 and gAS6, AS4g respectively. On the other hand, the AS-PATH of the UPDATES sent to AS3 can be artificially padded as in gAS6, AS4, AS4, AS4f and fAS5, AS4, gAS4, AS4f. This will make the path through AS3 longer and less attractive for other ASs. Contradictory advertisements can be used by a malicious router to redirect traffic to itself or to another AS. To illustrate, the compromised router Q should normally announce the AS5 route that goes through gAS5, AS4, AS2f only. Instead, Q can propagate that route only to F indicating that it should not be announced any further, and announce the padded route that goes through AS3 to D. This means that part of the Internet (excluding AS2) will be able to reach AS5 only through AS3. The attacker may want to do so in order to create congestion in AS3, or to redirect traffic destined to AS5 through a suboptimal backup path [Nordstrom and Dovrolis, 2004].

2.2.7 Update Modifications

Update modifications can be used by a compromised router to redirect traffic in a way that affects the origin AS by modifying the path and padding [Bonaventure, 2002]. In this case, the attacker wants to modify the AS Path in a way that it hurts the victim indirectly. An attacker that is aware that its victim uses padding can avoid using a more expensive route; the attacker may remove the padding from UPDATES so that traffic is still sent over the more expensive link.

Example: if AS4 uses the link O-K in Figure 2.3 only for backup purposes because it is cheaper to use link Q-R instead as the primary link, AS2 will not advertise its AS4 route to AS3, because doing so would enable AS3 to use AS2 to reach AS4 instead of using its own link O-K. To prevent other ASs from using the link O-K, router K can pad the UPDATES going to O, making the corresponding AS-PATH longer. Assume now that router D is compromised and it redirects traffic to AS4 through the more expensive link O-K. D can drop the padding in the route that includes the AS3, AS4 link, and instead pad the route that includes the AS2, AS4 link (or simply not announce it). This would force traffic for AS4 to take the more costly O-K route. As long as connectivity is preserved, this can be very difficult to detect.
This is a result of the fact that business relationships and policies between providers are largely kept secret, and makes the ability to detect illegitimate routing in terms of policy constraints very difficult for a third party [Nordstrom and Dovrolis, 2004].

2.3. Consequences of Attacks

The consequences of these attacks are as diverse as their approach. BGP sessions can be prematurely severed, networks and ASs can be made unreachable, the address space can become fragmented, and other undesirable outcomes can result from an attack. Attacks can be used in concert to amplify their effect or to enable further malicious activity. The generic consequences of routing threats are further discussed in - Barbir et al. [2003]. Examples of these consequences include the disclosure of confidential information, deceptive or incorrect information introduced into the network through message modification, the disruption of network activity through denial of service attacks, and the usurping of router services and functions. Consider the ramifications of a dysfunctional

routing system under attack. An individual router is subject to being overloaded with information, knocked offline or taken over by an attacker. An autonomous system can have its traffic black-holed or otherwise misrouted, and packets to or from it can be grossly delayed or dropped altogether. Malfunctioning ASs harm their peers by forcing them to recalculate routes and alter their routing tables. As the misconfiguration examples have shown, these events can disrupt international backbone networks and have the potential to bring a large part of the Internet to a standstill. From the individual level of an organization's traffic being stolen to the worldwide scale of IP traffic being globally subverted, the threats against BGP are a matter of grave concern to anybody reliant on the Internet [Smith and Garcia 1998].

Route Flapping

When a route is repeatedly advertised and withdrawn, it is considered to be 'flapping'. It is different from form route oscillation. Oscillation is periodic but flaps are not. Route flaps are the leading contributor to instability on the Internet and on any internetwork. It occurs when a valid route is declared invalid and then declared valid again. This problem is evident, in that it causes the router to change its state continuously and the change is advertised throughout the internetwork and this forces the router to make appropriate recalculations.

Most often we quickly name unstable physical links or failing router interfaces as leading to causes of route flapping, and we are right. But another common cause of route flaps, possibly the most common of all, is humans; Technicians tinkering in the Telco central office or in the wiring closet can certainly cause outages leading to flaps, but also inexperienced network administrators innocently configuring or troubleshooting his router may repeatedly delete a route, changing the state of the of an interface, or clearing a BGP session or an attacker triggers route dampening for a victim's route.

How bad can the effects of instability be? Consider a single somewhat overloaded or underpowered BGP router. An upstream connection becomes unstable, causing many routes to flap simultaneously. The router cannot handle the changes, and it fails. Now downstream routers have to process not only the original flapping routes, but also all the now-unreachable routes originated from the failed router. The effects can snowball cascading throughout the internetwork, possibly causing more routers to fail [Mao et al, 2002].

CHAPTER 3

SECURITY SOLUTIONS TO BGP

The current proposed or implemented BGP security solutions are still limited in their effectiveness to address all the inherent vulnerabilities so far discovered and finding complete solutions to BGP security problem is a wide aspect of research. We will examine few of the numerous proposals because of the variety of the issues involved, the different methodologies employed; the great number of new proposals being proposed, has made a complete categorization of solutions a difficult task to achieve (Kelvin et al 2000) because these issues are relatively new and no solutions have been universally deployed to address it.

The question that comes to mind is what actually do we need to protect?

- ✓ Original data:
  Address blocks
  AS numbers
- ✓ Path data:
  (Transitive) path attributes
- ✓ Originating addresses:
  Who owns the original address block?
  Has it been allocated by registries according to IANA/RIR hierarchy
- ✓ Route announcements:
  Uses transitive trust - even though you trust your neighbor, do you trust your neighbors' neighbor?

In this section we present the current implemented and proposed solutions, and their level of protection which include cryptographic techniques, architectural solutions and formal properties of routing.

3.1 BGP Cryptographic Techniques

Numerous researches have proposed a series of ways to address some of the security challenges inherent in BGP routing security. Some lay emphasis on formal properties of routing protocols while others focus on the application of the novel cryptographic

structures that provide better security guarantees. Cryptography is the tool that has been applied most often in works on BGP security. In addition to the ability to protect information from observation using encryption and secure digital signatures, the purpose has been to provide authentication information. For example, where did this information come from, and who authorized the information? Cryptography can be applied at many levels to secure BGP, including link, routing update, and routing database level.

3.1.1 Pairwise Keying

Many of the cryptographic mechanisms protecting a pair of parties rely on the existence of a shared secret key, often as input for a message authentication code. The first step is to obtain appropriate domain parameters that are generated, either the entity itself generates the domain parameters, or the entity obtains domain parameters that another entity has generated [Butler et al 2010]. The owner of a key pair is the entity that is authorized to use the private key of that key pair. Having obtained the domain parameters, the entity obtains assurance of the validity of those domain parameters. The two parties agree, often in an offline manner, on how the key is to be shared between them, and this key is then configured manually at each end point.

This approach is limited in that implementing and maintaining shared secret keys between many peer routers concurrently can be difficult as a result of the complexity of pairwise key management [NIST, 2007]. Moreover, such secrets, if not replaced frequently, are subject to exposure by cryptanalysis.

3.1.2 Cryptographic Hash Functions

Cryptographic hash functions, also known as digest algorithms, compute a fixed length of hash value from an input text and form the basis for message authentication codes and digital signature. The most common hash functions currently in use are the Message-Digest algorithm 5 (MD5) and the Secure Hash Algorithm family, particularly SHA-1 [Rivest, 1992]. A hash function is cryptographically sound if it is computationally infeasible to find a pre-image of a hash parameters (it is non-invertible) [Butler et al. 2005] and it is computationally infeasible to find two inputs with same output hash parameters (collision resistant). For MD5, the output is 128 bits in length. To demonstrate infeasibility, consider an attempt to find a message that will map to a particular MD5 digest with a 128-bit digest,

one would require an average of 2127 messages to find the particular message that mapped to the digest value, or 264 messages to find a message that created a collision, a different message that maps to the same digest value. The MD5 digest mechanism requires that a shared secret key be configured manually at each session end-point [Heffernan, 1998].

3.1.3 Public Key Infrastructure

The cryptographic techniques rely on a shared key between two parties. Because announcements can originate from any of the over 35,000 ASs on the Internet, being able to establish the integrity of these messages through mechanisms such as message authentication codes and digital signatures is necessary, but these rely on the establishment of keys between AS peers. Managing these pairwise keys between over 35,000 ASs will be problematic. Key management on a global scale requires public key cryptography. As applied to BGP, every AS has a public key, distributed freely to any other AS in the Internet, and a private key, which is never divulged. Two ASs without a priori knowledge of each other can negotiate a key for secure communication with each other (through a Diffie-Hellman key exchange) if they can find the public key for the AS they wish to communicate with. Public key infrastructure, or PKI, provides a framework for assignment and delegation of public keys. The PKI handles requests for public keys originating from other ASs. Keys are distributed in a hierarchical manner. Currently, such an infrastructure does not exist, but there has been considerable research in the field [Seo et al. 2001].

3.2 Architectural Security

Current efforts to provide more comprehensive BGP security architectures have been attempted within the standard bodies and the research communities. The architecture provides a suite of security services and an explicit threat model [Bellovin, 2003]. In this section we are going to look at the most comprehensive proposals to make certain that the data in BGP announcements are correct. The truth is nothing is in place presently.
: Butler et al [2005] propose:

- ✓ Secure BGP (S-BGP)

- ✓ Secure Origin BGP (soBGP)
- ✓ Interdomain Routing Validation(IRV)

3.2.1 S-BGP

Secure BGP (S-BGP) was designed by researchers at BBN (Bolt, Beranek and Newman) as an extension to BGP and was the first comprehensive routing security solutions targeted specifically to BGP with the objective to protect BGP from erroneous or malicious UPDATES [Kent et al. 2000]. S-BGP implements security by validating path attributes in BGP UPDATE messages passed between ASs through the use of digital signatures and associated public key cryptography.

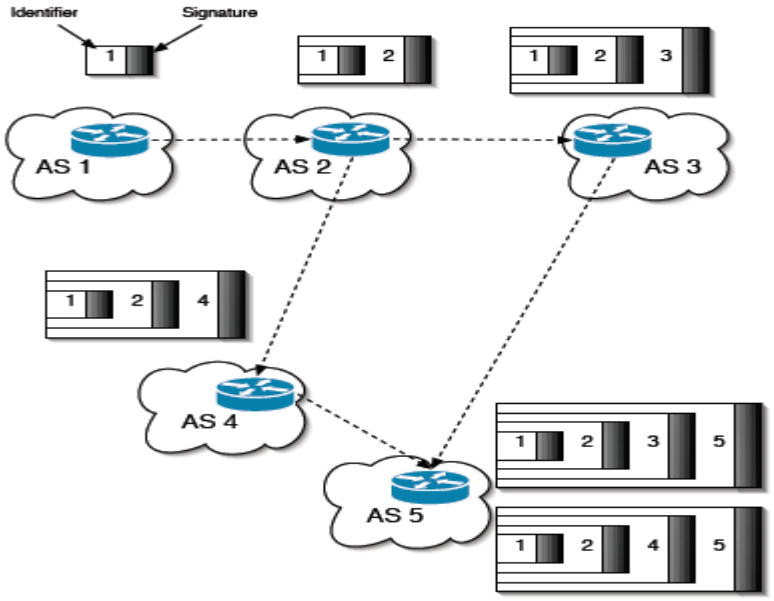

Figure 3.1 Route attestations in S-BGP.

As UPDATE messages are passed between peers, the receiving peer signs the received message before passing it to another neighbor. The result is an “onion-style” attestation that contains signatures from all routers along the path as shown in Figure 3.1[Butler et al 2010].

S-BGP implemented three major additions to BGP:

- ✓ It introduces a Public Key Infrastructure (PKI) in the interdomain routing infrastructure to authorize prefix ownership and validation of routes,

- ✓ A new transitive attribute is introduced to BGP UPDATES that ensure the authorization of routing UPDATES, and prevents route modifications from intermediate S-BGP speakers.
- ✓ All routing messages can be secured using IPSec, if routing confidentiality is a requirement.

All information exchanged in S-BGP is validated using the certificates in the PKI and Statements made by the AS are signed using the associated private key. An entity receiving the signed data verifies this by using the two key features of S-BGP, Address Attestations (AA) and Route Attestations (RA).

Address Attestations are digitally signed statements generated by the originator of a prefix and used to assert the authenticity of prefix ownership and advertised routes. Address attestations claim the right to originate a prefix, to sign and distributed it out-of-band. An out of-band mechanism does not directly use the BGP protocol to transmit information, instead, it uses some external interface or service to communicate relevant data.

Route attestations are distributed within S-BGP in a modified BGP UPDATE message as a new attribute authorizing a neighboring AS to propagate the route contained in that UPDATE. A route attestation is signed by each AS as it traverses the network and all ASs on the path sign previously attached signatures (nested signatures). Hence, the validator can validate not only the path that the ASs were traversed in the order indicated by the path, but also that no intermediate ASs were added or removed by an adversary. Figure 3.1 shows a simplified use of route attestations as they propagate between routers [Butler et al. 2010]. As UPDATE messages are passed between peers, the receiving peer signs the received message before passing it to another neighbor. The result is an ‘‘onion-style’’ attestation that contains signatures from all routers along the path.

However, while S-BGP proposes the most comprehensive security guarantees of all proposals by providing full authentication of origins and the paths to destinations, there are significant barriers that hamper its adoption. A deployment obstacle, however, is that it requires the presence of a hierarchical PKI infrastructure and distribution system, trusted by all participating ISPs.

Another obstacle is that S-BGP is quite cryptographically intensive, requiring each UPDATE to be verified and signed by each S-BGP router (or by each participating AS) it passes through. This performance overhead is unacceptable upon initialization (or reboot) of a BGP peering session due to the large number of routes that would be generated in a short time interval.

Aggregation is an additional problem for S-BGP. Route aggregation provides a means to coalesce several prefixes into a larger address block, thus reducing the number of UPDATES generated by a BGP speaker. S-BGP however, requires that all UPDATES be signed by the prefix originator. An upstream router performing aggregation will generally not be the owner of all the constituent prefixes.

Also, S-BGP cannot prevent collusion attacks. Such attacks are possible when two compromised routers fake the presence of a direct link between them.

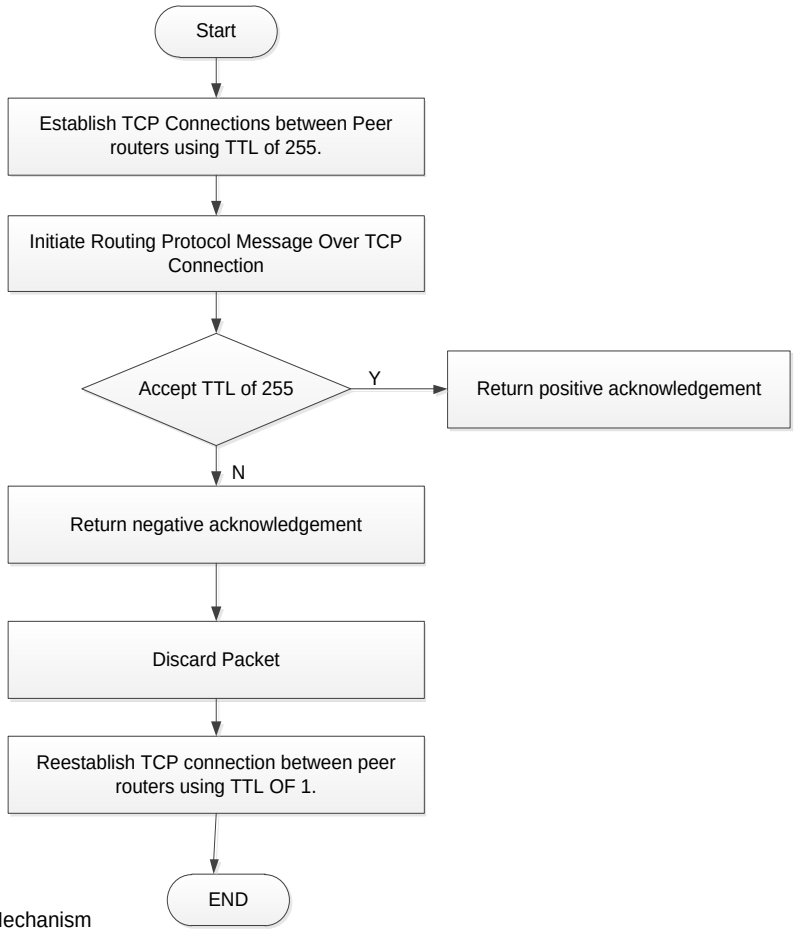

Figure.3.2 Generalized TTL Security Mechanism

Furthermore, the implementation issue is that routers need a large memory space (equivalent to the CPU and memory provided by a desktop PC [Kent et al. 2000] to store the public keys needed for route attestations. The space requirement can be significant for a speaker with tens of peers. However, because of the amount of data and number of possible signers, validation can be costly. These and similar results have raised concerns about the feasibility of S-BGP on the Internet, and led many to seek alternative solutions.

3.3 Formal Properties of Routing

The efficient operation of the network relies on the configuration of the individual routers and securing the connections between two BGP speaking routers. This also depends on both the TCP and implementing a good protection for the BGP session. Below is a description of some methods of protecting peer communication between two BGP speaking routers.

3.3.1 Generalized TTL Security Mechanism

Originally called the BBGP TTL Security Hack, the Generalized TTL Security Mechanism (GTSM) provides a method for protecting peers from remote attacks [Gill et al. 2004]. The time-to-live, or TTL, attribute in an IP packet is set to a value that is decremented at every hop. By default, IOS sends BGP messages to neighbors with a TTL of 1 and this requires that the peer be directly connected or the packets will expire in transit.
However, there is an inherent vulnerability to this approach; it is trivial for remote attackers to adjust the TTL of sent packets to appear as if it is originating from a directly connected peer. The solution as discussed in RFC 5082 is to avert the direction in which the TTL is counted. The maximum value of the 8-bit TTL field in an IP packet is 255.

For example, if a packet traverses n hops from source to destination, the TTL decrements by n – 1 at each hop till it gets to the destination with a TTL value of n. Routers using GTSM set the TTL of an IP packet to its maximum value of 255 as shown in Figure 3.2. When a BGP peer receives a packet, it checks the TTL and if this value is less than 254 (decremented by one i.e. n- 1), the packet is flagged or discarded. This prevents remote attacks which come from more than a hop away, as those packets will have TTLs less than the threshold value of 254.

However, GTSM weakly guarantees against attacks that are more than one hop away. It does not prevent against subverted peers sending malicious information or other similar insider attacks, and it is less useful in multi-hop scenarios where BGP peers are farther than one hop away from each other. The TTL threshold can be lowered to account for how many hops away the peer is, but there will consequently be no protections against attackers the same number of hops away. Additionally, if an attacker tunnels an IP packet by encapsulating it within another IP packet to a peer a hop away from the victim, the de-capsulated packet, with a TTL set to the maximum value, will automatically evade GTSM. GTSM is simple, low cost, and generally effective against unsophisticated attackers. However, the effectiveness of the solution to mitigate motivated attackers is limited.

Table 2. BGP Peer Session Security Solutions Comparisons [Butler et al 2010].

	Integrity	Confidentiality	Replay Prevention	DOS Prevention
MDS Integrity	yes	no	yes	no
Countermeasures	yes	yes	yes	no
HOP Protocol	yes	no	yes	no
GTSM	no	no	no	no
IPsec (AH)	yes	no	yes	yes
IPsec (ESP)	yes	yes	yes	yes

3.3.2 IPsec

Internet Protocol security (IPSec) is a framework of open standards which has been proposed as a mechanism for securing BGP session through the use of cryptographic security services [Thayer et al. 1998].

Many recent proposals have suggested the use of IPsec as a mechanism for securing the BGP session. IPsec is not specific to BGP, but is a suite of protocols that provide security at the network layer. These protocols define methods for encrypting and authenticating IP headers and payload, and provide key management services for the maintenance of long term sessions [Kent (RFC4303) 2005].

The IPsec Authentication Header protocol (AH) and Encapsulating Security Payload (ESP) protocol implement packet-level security with differing guarantees [Kent (RFC 4302), 2005]. All of these services work in concert to establish and maintain the secret keys used to guarantee the confidentiality and authenticity of data. IPsec is often used as the security mechanism for implementing Virtual Private Networks (VPNs) [Gleeson et al. 2004] and it provides the desirable security guarantees for authenticity of data, integrity, message replay prevention, data theft and data confidentiality.

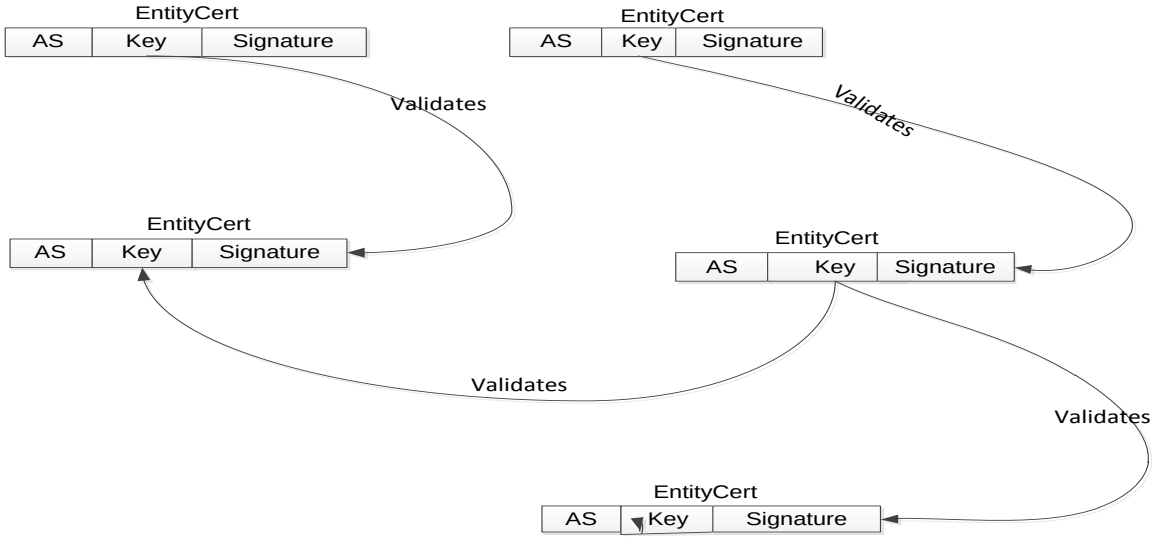

Figure.3.3 EntityCert Trust Validation

IPsec sessions only implement security between peers and address many issues relating to session-local vulnerabilities, but they do little to address widespread attacks.

As shown in Table 2, out of the existing solutions, IPsec provides the most comprehensive protection.

3.4 Secure Origin BGP

Secure origin BGP (soBGP) is a 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. soBGP is a lightweight alternative to S-BGP, mostly proposed by researchers at Cisco Systems [White, 2004]. All information pertaining to security in soBGP is transmitted between peers via a SECURITY message, a new message type in BGP introduced by soBGP.

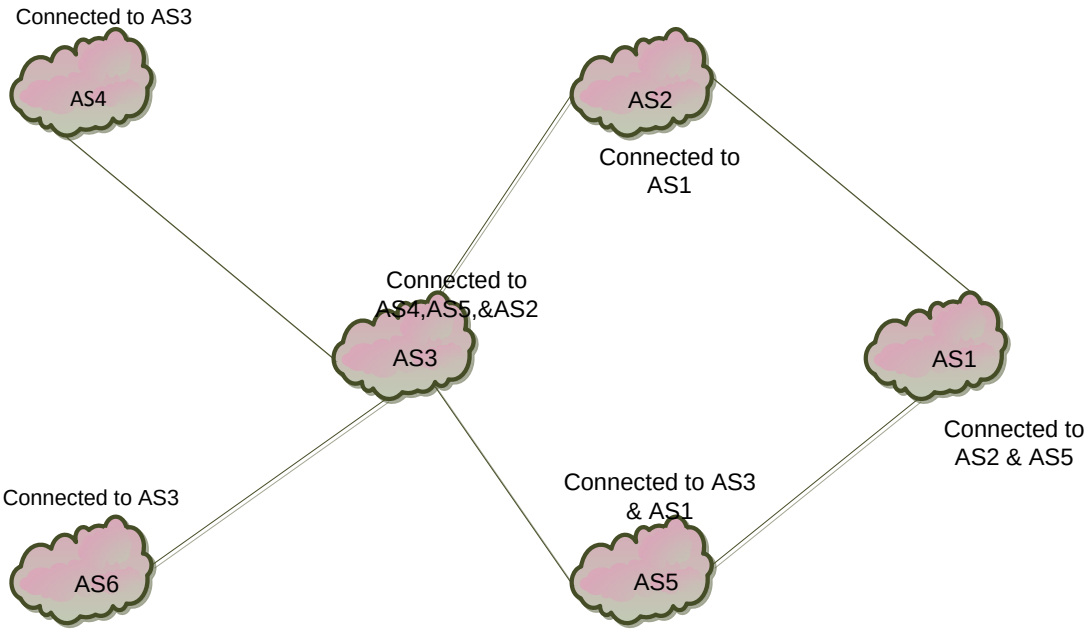

Figure. 3.4 MAP OF INTERNET TOPOLOGY

soBGP aims to authenticate two aspects of routing information. First, soBGP validates that an AS is authorized to originate a given prefix. Second, soBGP attempts to verify that an AS advertising a prefix has at least one valid (in terms of policy and topology) path to that destination. soBGP is based on the use of three certificate types.

The Entity Certificate (EntityCert) is used to establish the identity and public key of an AS. The EntityCert ties an AS number, an entity within the routing system, to a public key which belongs to that AS. The issuing Organizations that authorize the use of AS numbers and blocks of address space are not required to validate the public key of an entity, nor should organizations which validate an organization's public key be required to validate the advertisement of certain address spaces, rather a third party should be required. Thus, the EntityCert stands as a separate certificate type, validating who an entity is within the routing system as shown in Figure 3.3.

The Authorization Certificate *(*AuthCert) ties a particular AS to a particular block of addresses. The organization which authorizes an AS to advertise that a block of addresses signs this certificate, which is a very narrow piece of information, not including any policy about these prefixes, or attempting to prove that an entity is who they claim to be, nor that their public key is what they claim it to be. This authorization is provided through an Authorization Certificate*,* or AuthCert. An AuthCert ties an AS (AS in the topology) to a block of addresses that the AS may advertise [White, 2002].

Although the Cisco system researchers are currently working to develop prototypes of soBGP on several platforms to show how the technology will be deployed on a wide range of devices but it is believed that the implementation stage is not within reach yet and this will still pose a lot of security problems.

CHAPTER 4

SECURING sBGP ROUTING

Route origin validation ROA and RPKI is the underlying technology used in this project for securing sBGP routing (this are inter boundary issues (prefixes). This chapter attempts to describe the work done in this thesis to achieve its objectives. The main building blocks in the project are trust, ROAs and validators. Before proceeding and to put the records straight, and to avoid some misunderstanding, it is worth stressing that ROA is always stated to be the certificate but, this is not entirely true.

Rather, ROA are the documents used to link prefixes with an origin ASN. ROAs are created by ISPs, signed by the resource holder with a private key and the signing creates a chain of trust which allows the sBGP routers to validate its announcements. The ROAs can contain many prefixes and only have a single origin AS as shown below in Table 4.

Table 3 Multiple Prefix with a single As

Origin ASN	17771
Prefixes: 1.	172.128.200.0/24
2.	172.128.200.0/24

To achieve the verification of announced route or ASN, the sBGP routers rely only on the RRV database and believe that information provided by it is valid. Routes/AS_Paths are added to the database with an origin AS.

4.1 VNE Library

VNE is an acronym from Virtual Network Environment and it is a proprietary library that can be used for several purposes. VNE library contains the most common network functions and components, for Ethernet switch, routers, TCP/IP stacks, fragmentation and ARP. IPv4 and IPv6 are both supported. Library is actively developed and future releases will support basic firewall and DNS server resolving. Library is written in “pure (although sometimes 'unsafe') C#

code”. The VNE components can be used both in Windows and Linux without any special privileges. Different components can be used together to achieve different training and testing setups. VNE library also be used for BGP daemon, VPN and as a real time BGP monitor to investigate loops and convergence.

4.2 Network configuration

Every network design must be simple, scalable, easily configured and understood. The network configuration was done in an xml-file. Several xml-files can be produced and selected with a command line attribute. The engines read the configuration from the selected file and generate the network. In the xml-file the autonomous systems are defined and their corresponding network prefix and mask assign with an AS-number.

```
 <AS asn="11">
                    <Prefix network="172.16.12.0" mask="24"/>
</AS>
```

BGP-routers are mainly defined by their AS-number, unique ids, and addresses of the network there belong.

```
<BGPRouter asn="1" id="1">

                                <address source="10.1.2.2" destination="10.1.2.1"/>
                                <address source="10.1.5.5" destination="10.1.5.1"/>
                                <address source="10.1.6.6" destination="10.1.6.1"/>
                                <address source="10.1.10.10" destination="10.1.10.1"/>
            </interface>

</BGPRouter>
```

4.3. Main form

The main form is initialized in the code with the Visual Studio generated code portion of InitializeComponent(). The code is always run on form start-up and defines parameters like the name of the form, textual topic of it, client size and its reactions to resizing. Since the method content is generated by Visual Studio, changing it can stop the VS Designer from functioning, and thus the code should only be changed by using the Designer itself located in the Form1.Designer.cs file.

4.3.1 Graphical User Interface

The GUI for the project program consists of three separate portions (main form, controls and the canvas) within the code. The main form holds all the items displayed on the GUI. On top of the form are two panels that contain a graphical representation of the given router configuration, and other controls for operating the network. The Random_ping button is found on the top-left corner of the form as shown in Figure 4.3 is for testing purposes and retains the functionality of sending a packet from the source to the destination router in the network.

4.3.2 Canvas and Graph-Panel

The graph-Panel contains a visual representation of the router network, given to the program in the form of an xml file. When the amount of routers is parsed from the file, a set of coordinates is calculated so that the routers can be placed on the circumference of a circle, thus providing a way for having non-router-overlapping connections between them.

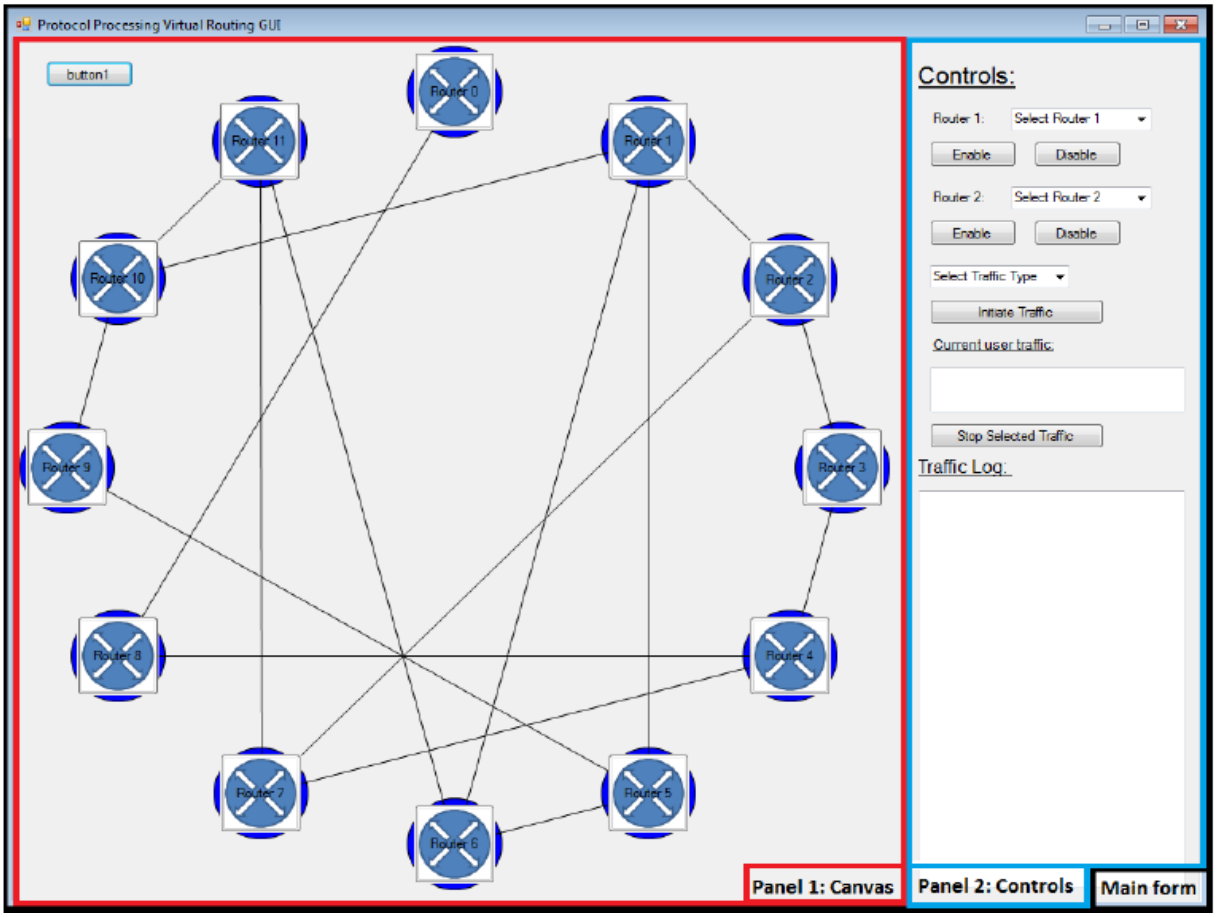

Figure 4.1 Canvas and Graph-Panel.

Each router is drawn in the panel as a button, making it easy to create events and handlers.

The connection lines in the panel are parsed from the configuration xml file, and all added connections are drawn as Line-Shape objects using router coordinates on top of a Shape-Container named canvas. The router backgrounds (blue circles behind the router images, see Figure 4.1.) are drawn on top of the canvas as Oval-Shapes. The reason for using different shapes on top of the Shape-Container is for the flexibility of changing the colors of the objects as desired.

The shapes are added as controls on top of the canvas the canvas added as a control on top of the graph-Panel. The Graph-Panel is a control on top of the main form and the router buttons are added as controls on top of the graph-Panel. In order to have the router buttons on top, the Bring-To-Front() method is invoked and Send-To-Back() method in the canvas. All the drawing-related codes can be found in the Form1.Designer.cs file:

4.4 Network Topology Drawing and Discovery

The topology was drawn as to reduce the amount of intersecting connection lines, thus providing a clear image of the network structures and links by simply drawing the routers to the edge of a circle.

The connection listing was done in a recursive process, so that when an arbitrary router is started and checks are done on its unknown connections. Since the list of found connections are done by sorting through the router numbers by this process we eventually ended up with a list containing all the unique connections in the network.

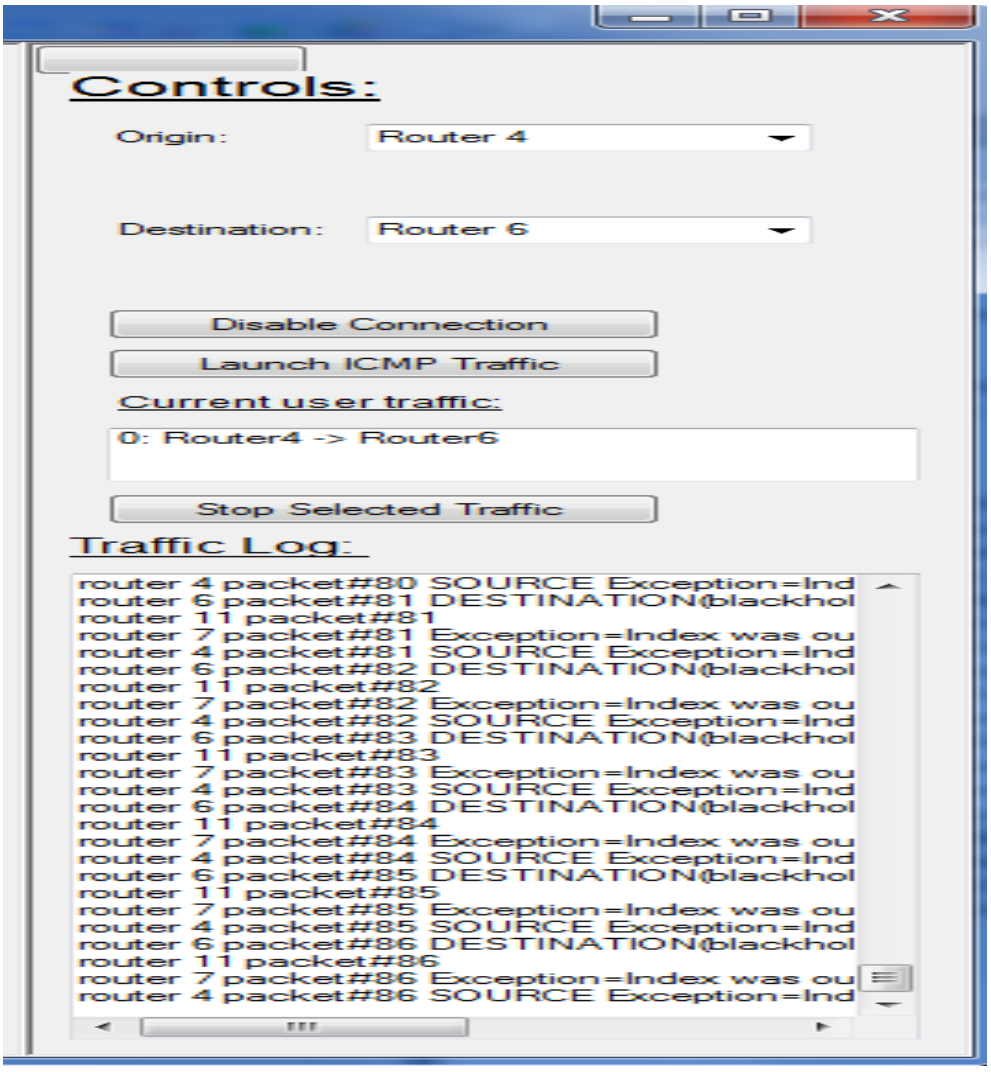

Figure 4.2 Control Panel.

4.5 Control Panel and the components

The Control Panel contains the control functionality of the form as shown in Figure 4.2. It is placed on the right side of the form and consists of buttons (enable, disable, launch ICMP traffic, stop selected traffic etc.), labels (all the routers, example R1 etc.), combo-boxes (the drop-down lists) and list-boxes (the empty white boxes). The combo-boxes work as selectors for the routers in the simulation network configuration. By selecting a router, the user can either decide to enable it by clicking on the buttons below the combo-box and disabled by clicking on the router diagram. When the source and destination routers are selected and enabled, the user initiates traffic into the network by pressing the “Initiate Traffic” button.
Initiating traffic produces a listing to the list-box in the form of “*index of traffic+ Router X.>Router Y”. The traffic creates a thread inside the program which sends the data packets from the first router to the next and highlighted in the graph-Panel router backgrounds and connection lines. Logging data about the packets are produced in the list-box under the “Traffic Log” label. Any traffic can be stopped by selecting the list and clicking the “Stop selected traffic” button. A slight delay is noticed when the traffic is stopped because it takes a fraction of a time for the list to be cleared from the list table.

4.6 Engine Code

Engine is the main code for the program, IP4 Router Extended class is used for storing BGP4 router Extended reference attached to the IP4 Router because there is no direct access to know from which IP4 router the BGP4 router is attached . “My id” property is used for the router and IP4 level as to have the “asn” number and “router id”. The BGP4 router knows about the IP4 router and can initiate communication with it.

Every router has one-to-one correspondence to some single router or AS. Figure 4.4 describes the structure of the program. When a router is switched off, the connections or routes are broken and when it is switched on again, it creates new routes and new lists of connections are established. The routing table for the IP4 router is saved into a temporary variable or directory. Routes are stored and deleted based on traffic requirements. The solution here is to copy the whole list of connections, make a new list that is equivalent to the

previous one. The route prefix to interfaces is contained in the route and traffic will correspond to the prefix that is already stored in route table. The invoke function of delete connections for a router is achieved by comparing the internal routes to the external routes of another router and a Public class net packet extension as enhancement only for storing the packet id. A packet path class was created to trace and debug the looping behavior in the routing network. For saving the history of every packet flowing through the routers, the packet id is stored as to have an array list of packet path and also an array of array list.

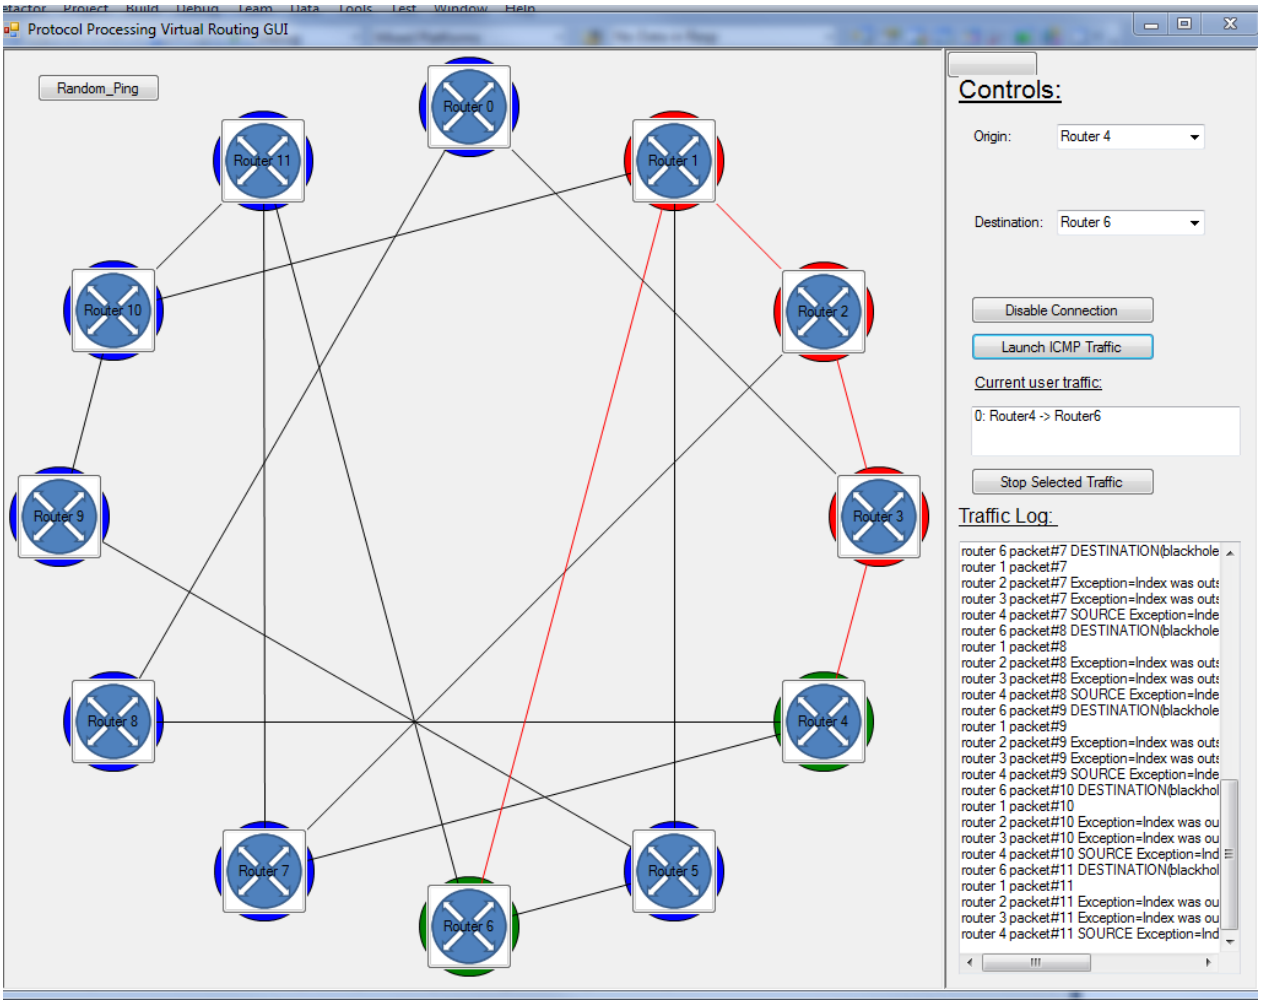

Figure 4.3 Initiating traffic

The array corresponds to array of packets and the index in the array to the “packet id”, and array list will be an array of dynamically changed array of paths the packet traverses. The

array is resized and added as a new element by entering it in array list. The array list actually holds the actual path and id of router from path. The serialized array gets the parameter file name and opens the string to the file name which will either create or overwrite the parameter. The sender component allows the serialization of the whole object. The Soap parameter is used as a converter for serializing in XML. The defined XML route packet path and its directives to anyone who need to serialize the class are included in the public class.

4.6.1 Init network function

Init Network is the main function which is being invoked in the library form externally. It reads the file externally, loads the file to Extensible Markup Language (XML) and sends the XML node list to the routers. By doing this, the numbers of routers are known as well as and the array of routers initialized and their connections. For each router the id is compared with the index i. The initialized routers and the created router function are put to the collection. The IP router4 are put in the packet hook event handler for processing of paths through the routers. The function will process the created stack attached to the IP4 router and also a backward link and reference from IP router to stack is established. The BGP router with the IP stack CREATED provides us the socket constructor that is able to listen to any connections.

Therefore, the socket will be a listener socket at port 179 (BGP Port and the BGP router communicates with the IP Router and listens to all incoming TCP packets directly. After creating the router, the BGP router comprising of Stack and other routers, the XML (Path_request, query, all the AS numbers and prefix children numbers) is advertised. All prefixes are taken and combined with the function “add local routes” to BGP routes and this will announce all local routes not only to its peers but also to all routers. These routes will be created as black hole routes and the prefixes emerge with interface nodes and the black hole routes are treated as local routes.

The interface parameters of XML are actually physical connections from source to destination. For the IP Stack part, the IP stack reference is added with a source address and the defined attributes. The BGP socket will independently communicate to IP stack through its own socket listener.

This connection is not physical; it is just a logical connection as only the IP4 router is communicating with everyone in network but the BGP router does not know where its neighbors are physically. All the prefixes it has are either black hole routes to node interfaces or IP stack and for link communicating with neighbors and the links have to be defined, so that it runs through all the routers.

This was achieved using a function of feedback route that returns the remote address. When it returns the remote address, then the neighbor's router is identified. In the stack it has the remote destination addresses that communicate to the router. Actually, the routers will pass its local addresses to this function and then there will be a comparison between the addresses to check if it is assigned to it by another as source address.

4.6.2 Initializing Network

The IP Packet hook function is called every time the packet arrives to an IP4 router. The packet hook event contains the packet reference of packet which is being passed. The IP class handles the packet content and obtains the source or destination address. With the source address it will be able to check if the current router is the originator of the packet or if it is just traversing the packet. The packet length is set at a default length 24 Bytes for ICMP packet. If the length of the packet is 24 Bytes, it means that the packet is not modified. Otherwise it is modified.

The function for highlighting the (path) wire actually colors the path on which packets traverses. Wire is defined by the current router id and the next point. Also, there is a function to highlight the router that are involved in (path routers) passing the packet. The wire or path is always red, and the transit routers are green while the source and destination router is red when it is passing traffic.

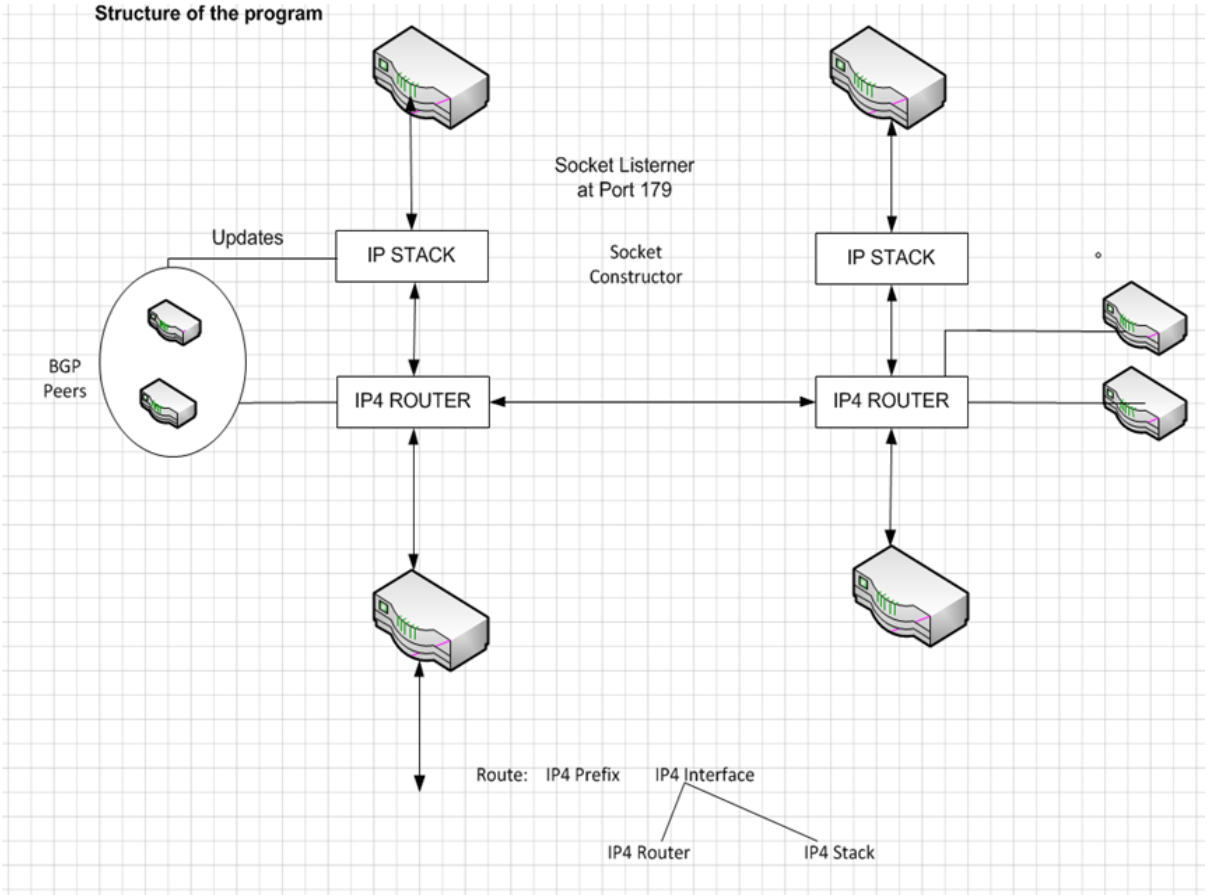

Figure 4.4 Network structure.

4.7 Packet Transmission

Below are five different stages or steps that are involved in the traversing of packets between routers in the network Simulation as shown in Figure 4.6.

(1) Packet transmissions are initiated by the user through the GUI or by some timers in the networking protocols used (like BGP).

(2) After the traffic is initiated, the packets are created at the local IP Stack without any payload and the packets are created and passed to the local IP stack through a TCP socket connection.

(3) When the IP Stack has proper network packet formed, it invokes the send packet function of its local IP router, the calling of which automatically invokes the route packet function. By this, it checks out whether the target interface is an IP Stack (local address) or an IP Router (non-local address).

(4) If the packets destination address is local, the packet is passed to the local IP Stack for delivery. The IP stack recognizes the TCP frame within the packet and passes it to the relevant TCP socket for delivery to the BGP router.

(5) If the destination address is an IP router, the target's send packet function is called, leading again to launching Route-Packet on the target Router.

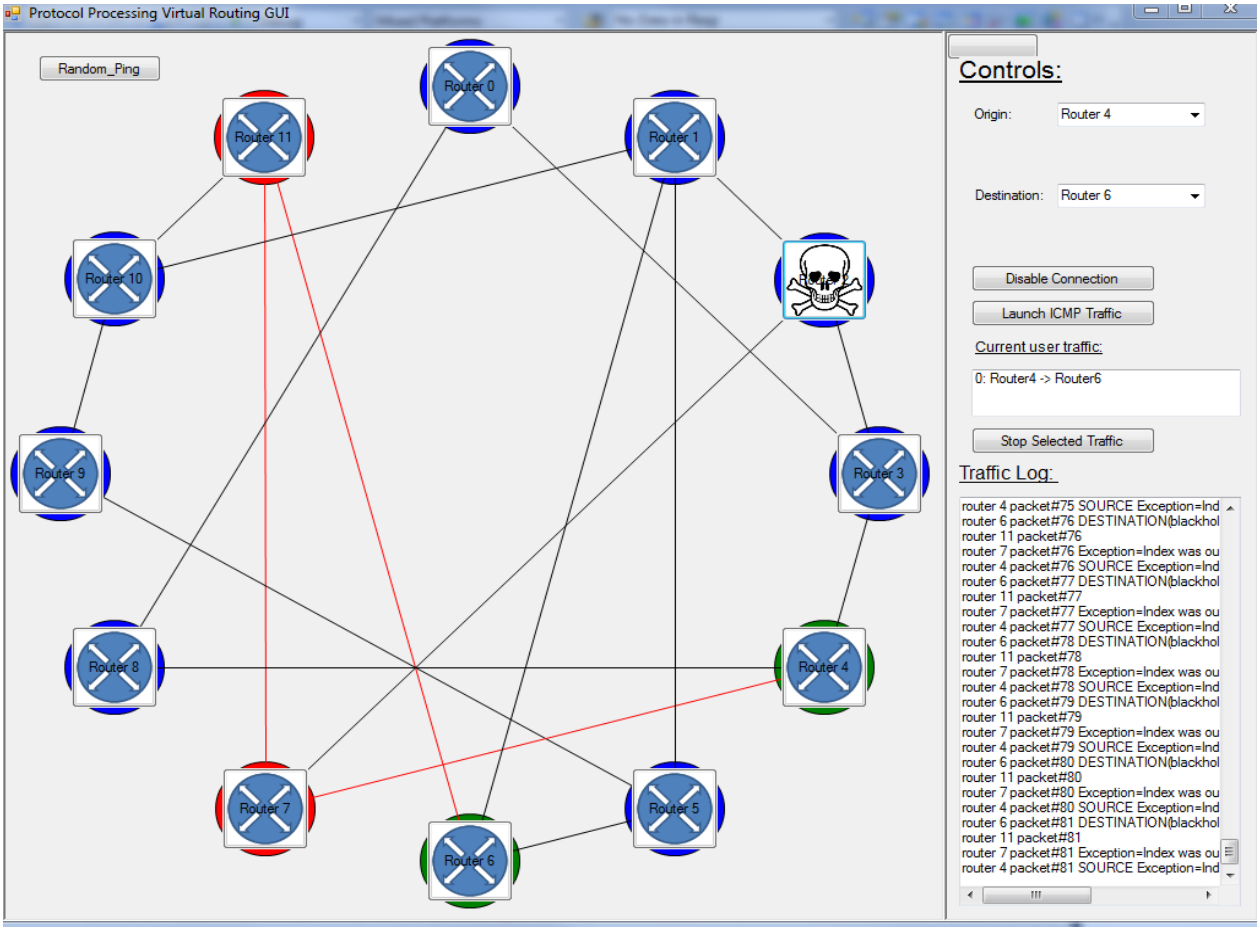

Figure 4.5 Enabling and disabling the routers.

4.8 Enabling and disabling of routers from GUI

The network is controlled through the GUI which, can be used to enable or disable routers at will. This was done by creating methods that either move a router's routing table into a temporary storage, nullify the actual list (disable), or restores the list from the temporary storage (enable). There are two problems associated with the implementation of these methods. Firstly, the state of the router is not checked before calling of these methods. The solution to this problem was to always check the current state of the selected router before invoking either of the methods, and allowing a switch of the state.

Secondly, having a null route table in the case of disabling a router allows the methods to be invoked more than once. Null routing table can easily cause unhandled exceptions as the route list can be nullified or an empty list copied. The solution to this is that a disabled router should be left as null list rather with the routing table list set to a length of 0.

In the Panel-Graph as shown in Figure 4.5, the router changes to the sign of the skull showing that it has been disabled. Any number of routers or paths can be disabled and the network readjusts and finds alternate paths to its destination if any.

Remote Router Validator

When a sBGP router announces an asn/ prefix, the peer sBGP router accesses the Remote Router Validator (RRV) that stores key attributes before responding to the announcement. The illustration below in Figure 4.7 shows the function of the key architecture.

Prior to authentication (signing and validation processes) access to private keys goes through the key directories and files in the remote validator router, which sets a comprehensive set of criteria for managing and using private keys.

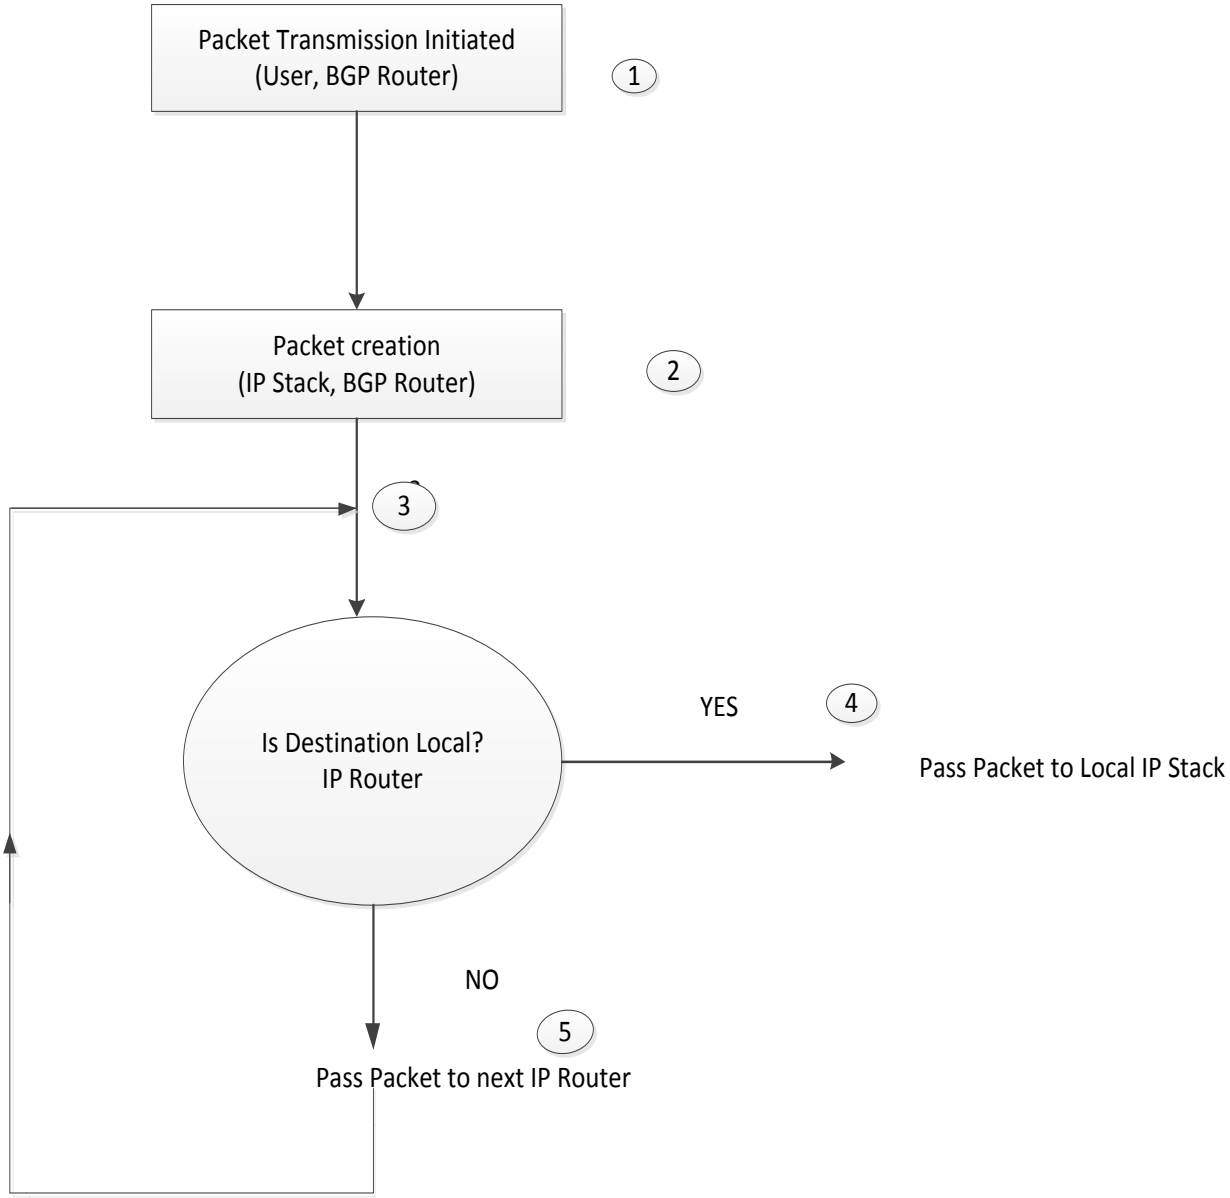

Figure 4.6 Packet transmissions.

Using a single Remote Router Validator removes the difficulties of searching from one domain to another without losing private keys. Users create the keys and these are stored manually in RRV.

Implementation of the validation

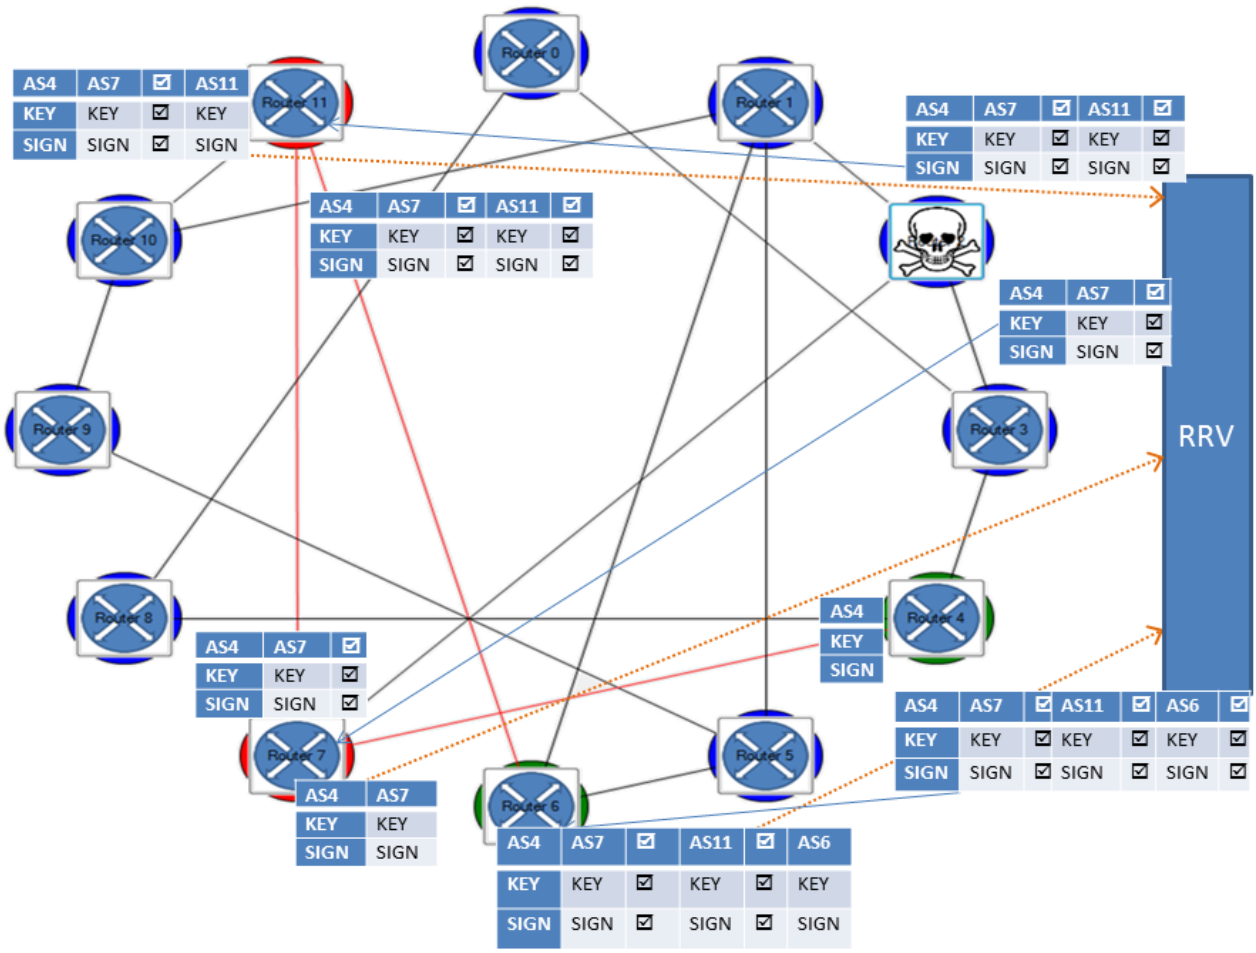

Figure 4.7 ROA validations and signing

The full implementation of this project depends on the RRV. The RRV does a lookup of all ROAs from the trust anchors and caches them locally. For each request the corresponding attribute will be checked and examined for ROA before a response. The response to request should show the results of validation. When a failure occurs the RRV notifies the sBGP routers why it failed (false positive from the speaker). The ROA`s major role is to link the set of prefixes to an ASN. The signing is done using the private key (from a certificate) of the resource holder and this works as an onion or chain of trust.

Chapter 5

Summary

The major issue is security, and the network might be compromised in terms of security so all of the traffic would need to be encrypted and the trust client forming a TCP connection to the IPStack by having all of its traffic routed automatically within the network. We figured out that this can be achieved by introducing RPKI in such a way that when registering an AS, an organization would be granted its own key-pair, which it would use to communicate.

The RPKI would be such that all communications between the trust clients would be signed with the organizations' (AS') private key. The public keys would then be available at a centralized server, so that when a new connection to a previously unknown AS is made, its messages could be verified and encrypted by requesting the associated public key from the central server. The communications towards the main server would obviously need to be signed by both sides with their respective secret keys. In this way, all communications inside the trust client network would be secure, and the communication could easily be identified. What would still remain a problem to be solved is the dilemma of having a RRV (centralized authoritative server).

Problem

Non-safe threading

After implementing threads for changing the colors of router backgrounds and connection lines when there was traffic, we soon ran into exceptions concerning non-safe threading. The problem was that we were accessing the OvalShapes and LineShapes from a separate thread of that in which they had been created.

As a solution to this, we introduced lock-objects, with which we could ensure that only one thread is accessing a given oval or line object at a time. Essentially a thread willing to change the color value of a line or oval would first try to get a hold of the lock object associated with the target object. If the lock was available, it took hold of it, did the changes it was meant to and released the lock after the operations. If the lock was already in use the thread would

wait until it was released. The same method was later used when it was made possible for the user to create and remove traffic-threads.

Further Work

This section is to mainly bring up the topic of continuation, even though the project is completed, there is still much that could be achieved:

- Implementation of trust and the RRV
- Convergence time graph
- Relationship between the rates of propagation to the number of hops.
- Behavior of the simulation (how it reacts to changes?)

REFERENCES

Gleeson, B., Lin, A. J. H, Armitage, G. and Malis, A. February 2000. "A Framework for IP Based Virtual Private Networks", RFC 2764.

BBC News: Pakistan blocks YouTube website. http://news.bbc.co.uk/2/hi/south asia/7261727.stm .Accessed: Jan. 20.2012.

Bonaventure, O. 2002. Interdomain routing with BGP: Issues and challenges. IEEE SCVT2002, Louvain-la-Neuve, Belgium.

Bellovin, S., Bush,R., Griffin, T., and Rexford, J. 2001. Slowing routing table growth by filtering based on address allocation policies. http://www.research.att.com/jrex/. Accessed: Feb. 23.2012.

Barbir, A. Murphy, S., and Yang, Y. October 2006. Generic threats to routing protocols (Request for Comments: 4593). http://www.ietf.org/rfc/rfc4593.tx. Accessed: Feb. 12.2012.

BGP Case Studies: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800c95bb.shtml#intro. Accessed: February 24.2012.

BGP Origin Validation. 07 Dec. 2011. RIPE Network Coordination Centre. http://www.ripe.net/lir-services/resource-management/certification/bgp-origin-validation. Accessed: Dec. 6.2012.

BGP Origin validation: http://www.ripe.net/lir-services/resource-management/certification/bgp-origin-validation. Accessed: February 24.2012.

Cisco - BGP Case Studies - 4shared.comdc304.4shared.com/doc/UDMObwkB/preview.html

Cisco CCNP Route (Cisco Academy).

Cranor, L. and La, M. B. 1998. Spam! Communications of the ACM 41, 8 (Aug.), pp. 74–83.

Ballani, H.,Francis, P. and Zhang, X. 2007. Study of prefix hijacking and interception in the Internet. In Proc. of SIGCOMM ' Aug. 2007. New York, NY, USA: ACM, pp. 265–276.

Butler, K., Farley, T., McDaniel, P. and Rexford, J. April 2005. A Survey of BGP Security. Technical Report TD-5UGJ33, AT&T Labs-Research, Florham Park, NJ.

Butler, K., Farley, T., McDaniel, P. and Rexford, J. January 2010. A Survey of BGP Security Issues and Solutions. Vol. 98. No. 1, pp. 100 – 122.

Gill, V., Heasley, J. and Meyer, D. Feb. 2004. The Generalized TTL Security Mechanism (GTSM), RFC 3682. http://tools.ietf.org/html/rfc3682. Accessed: May 8.2012.

Horn, C. June 8, 2009. Understanding IP Prefix Hijacking and its Detection. Technische University of Berlin. http://www.net.t-labs.tu-berlin.de/teaching/ss09/IR_seminar/talks/prefix_hijacking_horn.handout.pdf. Accessed: July, 4.2012.

Karlin, J. June, 2006. A fun hijack: 1/8, 2/8, 3/8, 4/8, 5/8, 7/8, 8/8, 12/8 briefly announced by as 23520 (today). NANOG Archive. http://seclists.org/nanog/2006/Jun/81.

Kent, S. and Seo, K. Dec. 2005. Security Architecture for the Internet Protocol, RFC 4301.

Kent, S. Dec.2005a. IP Authentication Header, RFC 4302.

Kent, S. Dec.2005b. IP Encapsulating Security Payload (ESP), - RFC 4303.

Kent, S., Lynn C., Mikkelson, J. and Seo, K. Secure Border Gateway Protocol (S-BGP) - Real World Performance and Deployment Issues. Proc. of Network and Distributed System Security Symposium, (San Diego, California), Feb. 2000.

M. Zhao, S. W. Smith, and D. M. Nicol. Nov. - Dec. 2005.The performance impact of BGP security, [IEEE Network]. Vol. 19, No. 6, pp. 42–48. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1541720&tag=1. Accessed: Nov. 5.2012.

Mahajan, R., Wetherall, D. and Anderson, T. Aug. 2002. Understanding BGP Misconfiguration. In Proceedings of ACM Sigcomm. Volume 32, No.4, pp. 3 – 16.

Murphy, S. June 2003. BGP Security Vulnerabilities Analysis, IETF Internet Draft (draft-ietf-idr-bgp-vuln-00.txt).

National Bureau of Standards. March, 2007.Secure Hash Standard, FIPS PUB 180-2, Aug. NIST SP 800-56A, Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography.

NIST SP 800-56A. March, 2007. Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography.

Nordstrom, O. and Dovrolis, C. April 2004. Beware of BGP attacks. In ACM SIGCOMM Computer Communication Review. Vol. 34. No. 2, pp. 1 - 8.

Premore, B. May 2003. An Analysis of Convergence Properties of the Border Gateway Protocol Using Discrete Event Simulation, Ph.D. thesis, Dartmouth College.

Rekhter, Y., Li, T., and Hares, S. 1996. A border gateway protocol 4 (bgp-4) - RFC 4271.

Rekhter, Y. and Watson, T. J. 1991. BGP protocol Analysis - RFC 1265.

Russ, W. October 2002. Deployment considerations for secure origin BGP (soBGP). Internet – Draft, Cisco Systems.

Russ, W. September 2003. Securing BGP through Secure Origin BGP. Technical report, Internet Protocol Journal, Cisco Systems. Volume 6, No. 3, pp 1- 6.

Seo, K., Lynn, C., and Kent, S. Jun. 2001. BGP Public-key infrastructure for the secure border gateway protocol (S-BGP), in IEEE DARPA Information Survivability Conference and Exposition II, Anaheim, CA.

Smith, B. and Garcia, L.A.J.1998. Efficient Security Mechanisms for the Border Gateway Routing Protocol. Computer Communications (Elsevier). Vol.21. No. 3. pp. 200–210.

Thayer, R., Doraswamy, N., and Glenn, R. 1998. IP Security Document Roadmap. RFC 2411.

Varghese, G. and Randy. H. K. Oct. 2002. Route Flap Damping Exacerbates Internet Routing Convergence," in Proceedings of ACM SIGCOMM, pp. 221-233. New York, NY, USA. http://dl.acm.org/citation.cfm?id=633047. Accessed: March 5.2012.

Zheng, Z., Ying, Z., Charlie, Y. H., and Morley, Z. M. Practical Defenses Against BGP Prefix Hijacking. http://web.eecs.umich.edu/~zmao/Papers/conextDefendHijack07.pdf. Accessed: April 24.2012.