ANALYSIS OF BORDER GATEWAY PROTOCOL

By

Muhammad Nasim

A thesis submitted to the faculty of Information Security Department Military College of Signals, National University of Sciences and Technology, Rawalpindi in partial fulfillment of the requirements for the degree of MS in Information Security

July 2014

SUPERVISOR CERTIFICATE

It is to certify that Final Copy of Thesis has been evaluated by me, found as per specified format and error free.

Dated_____ _____

(Mr. Ali Imran)

ABSTRACT

The Border Gateway Protocol is the inter-domain routing protocol for internet. Despite being a central part of the most critical network infrastructure i.e., the internet, it lacks in security features and is vulnerable to variety of attacks. BGP was originally designed with minimal security features to work in trusted networks. However, the evolution in threat environment for the internet presents the challenge to secure its backbone infrastructure, governed by BGP routing protocol. The configuration mistakes on routers also have serious and devastating effects on BGP routing.

Security provisioning for BGP has been researched extensively. This research work analyzes the threats posed to BGP routing protocol. A Threat Model for BGP is presented, inclusive of all important and latest threats on the protocol. The existing solutions to secure BGP including S-BGP, SoBGP, IRV, PsBGP, BGPSec are studied and evaluated based on security features. An Evaluation Model for BGP security architectures is established based on two approaches of evaluation, security evaluation and deployment evaluation. BGP security architectures and extensions are evaluated in thorough detail against the defined evaluation model for security features and performance attributes. A deployment strategy for BGP is also presented. The security of the proposed deployment strategy is based on BGPSec, while the deployment cost is greatly reduced; by off boarding computational and storage overhead from existing BGP border routers. The need for high performance specifications at existing border routers has remained the main challenge for successful deployment of BGP. The presented deployment strategy presents an efficient solution to this challenge and provides a practical deployment solution for BGP.

DEDICATION

In the name of Allah, the most beneficent, the most merciful.

This research work is dedicated to my dear homeland Pakistan, my teachers and my family.

ACKNOWLEDGEMENT

Immense gratitude to Almighty Allah, Who endowed me with strength and intellect to accomplish what I have today.

First of all, I would like to thank my supervisor Mr. Ali Imran. His vision, knowledge and intellect remained a constant source of inspiration for me. Despite his busy schedule and commitments, he was always available for meetings and discussions. He guided me throughout the course of my research work. His valuable suggestions and constant encouragement helped me during the most difficult phases of my research. It is only with his supervision that I could successfully complete this challenging work.

I am also very thankful to my co-supervisor Dr. Asif Masood, guidance committee members: Lecturer Waseem Iqbal, Lecturer Adeela Waqar and Lecturer Waleed Bin Shahid for their support, help and valuable suggestions. I am especially thankful to my Head of Department Dr. Baber Aslam and his team for administrative help and support.

I would also like to thank all my family members for the help and support they provided during the entire course of my studies.

Table of Contents

LIST OF FIGURES

LIST OF TABLES

Table Number

Page

KEY TO ABBREVIATIONS

BGP	Border Gateway Protocol
AS	Autonomous System
IGP	Interior Gateway Protocol
EGP	Exterior Gateway Protocol
OSPF	Open Shortest Path First
IANA	Internet Assigned Number Authority
ICANN	Internet corporation for Assigned Name and Numbers
RIR	Regional Internet Registries
ISP	Internet Service Provider
NLRI	Network Layer Reach-ability Information
RIB	Routing Information Basis
SLA	Service Level Agreement
DoS	Denial of Service
TCP	Transport Control Protocol
TTL	Time to Live
GTSM	Generalized TTL Security Mechanism
ASN	Autonomous System Number
DSUA	Documenting Special Used Prefixes
ACL	Access Control List
IRR	Internet Routing Registry
S-BGP	Secure Border Gateway Protocol
PKI	Public Key Infrastructure

CRL	Certificate Revocation List
SoBGP	Secure Origin Border Gateway Protocol
IRV	Inter-domain Route Validation
NME	Network Management Element
TLS	Transport Layer Security
PsBGP	Pretty Secure Border Gateway Protocol
PAL	Prefix Assertion List
IETF	Internet Engineering Task Force
SIDR-WG	Secure Inter-domain Routing Working Group
RPKI	Resource Public Key Infrastructure
ROA	Route Origin Authorization
SIA	Subject Information Access
AIA	Authoritative Information Access
MOAS	Multi Origin Autonomous System
SPV	Secure Path Vector

Chapter 1

Introduction

1.1 Background

As communication systems in today's world continue to grow and evolve, the utility of networked systems is now more than ever. Information can now reach to widely placed networks even when they are far and wide from one another. The world's largest network is the internet, which provides connectivity to many smaller networks. Network traffic travels across computing devices through the internet, and is managed by routing protocols, that help reach network traffic from source to destination.

Border Gateway Protocol (BGP) is the internet's *de facto* Internet Engineering Task Force's standard inter domain routing protocol. It provides routing infrastructure for the internet and connectivity for exchange of network routing information between different independently managed smaller networks. BGP was originally designed to be used with a set of trusted networks, lacking in security features. As the threats and attacks on the internet continue to grow and evolve, insufficient security provisioning for BGP remains a challenging problem. Because of the increased importance of and growing reliance on internet, there is a need to secure the underlying routing infrastructure of the internet for which, BGP is the most dominant inter-domain routing protocol.

This research work is based on exploring the security vulnerabilities in BGP protocol and the security extensions to mitigate attacks on BGP. The security analysis and deployment evaluation of BGP security extensions and related techniques is a major contribution of this research. A deployment strategy for BGPv4 is also presented. This

chapter presents an overview of the problem area, and the research objectives and methodology. In the end, overall organization of this thesis is described.

1.2 Overview

Internet is globally the world's biggest network comprising of many inter-connected smaller networks. In usually topology, smaller networks are connected with larger or countrywide networks which lead network traffic to densely connected backbone providers networks. Network traffic thus travels on one of many possible paths from the source to the destination in different networks. Border Gateway Protocol (BGP) is the most dominant routing protocol used today on the internet, for governing traffic flow between independent networks, making them accessible to each other.

BGP was originally designed for a set of trusted networks so it lacks in security features but with the evolution and growing complexity of internet, this trust can no longer be assumed. The need for security features is driven by the growing significance and dependence on the internet. Attacks on BGP pose a serious threat to internet connectivity that can cause devastating damage to the today's connected world. Thus, there is a need to secure the underlying routing infrastructure of the internet for which, BGP is the most dominant inter-domain routing protocol.

The aim of this research work is to examine the threats posed to BGP routing protocol and the security extensions devised to mitigate such threats. The proposed security architectures and extensions to BGP include S-BGP, SoBGP, IRV, psBGP, BGPSec, Hop Integrity Protocols, Invalid MOAS conflicts detection schemes, Path Authentication and SPV. The security extensions are assessed on two criteria; Security Evaluation and the Deployment Evaluation.

1.3 Problem Statement

Several problems related to BGP security were identified with reference to the background presented above. The first problem was to understand the basic working of BGP routing protocol. Next problem was to formulate the threat model for BGP so that the required security attributes can be established. An Evaluation Model was defined including all necessary security and deployment requirements. BGP security extensions were evaluated thoroughly based on the evaluation model. Last problem area was to present a deployment strategy using BGPSec.

1.4 Research Methodology and Achieved Goals

The research work carried out during the course of this thesis was divided into many phases. The first phase was to explore the working of BGP in densely networked environments like the internet; which is more prone to security threats and attacks. The next phase was to develop a threat model for BGP. The threat model includes BGP session attacks, origin falsification attacks, path subversion attacks, denial of service attacks and attacks due to wrong configuration of BGP. The third phase was to conduct analysis of BGP extensions including S-BGP, SoBGP, IRV, psBGP and BGPSec.

The next phase was to establish an evaluation model for BGP security requirements and performance requirements. BGP security extensions were evaluated based on the devised evaluation model. The last phase of this research was to propose a strategy for BGP deployment with optimum computational complexity and security provisioning.

1.5 Thesis Organization

This thesis is organized into 8 chapters. Chapter 2 is based on the literature review carried out for this research. It explains the working of BGP routing protocol. Chapter 3 presents a threat model for BGP, highlighting security vulnerabilities in the protocol, followed by the protocol's limitations that cause security weakness. Chapter 4 presents analysis of the existing BGP security architectures, including S-BGP, SoBGP, IRV, psBGP and BGPSec. An evaluation model for BGP is presented in Chapter 5, based on requirements for BGP security provisioning and performance criteria. This evaluation model is used to evaluate proposed BGP security extensions, a detailed account of which is made in Chapter 6. A deployment strategy for BGP is proposed in Chapter 7. The security of the proposed scheme is the same as that of the latest and most comprehensive extension BGPSec. The proposed strategy achieves performance optimization and incremental deployment for BGPSec implementation. Chapter 8 concludes this report and describes the directions to extend this work in future.

Literature Review

2.1 Introduction

The Border Gateway Protocol (BGP) is the existing inter-domain routing protocol on the Internet. It provides routing mechanism for the internet providing connectivity between many smaller networks. The main function of BGP speaking system is to exchange network reachability information with other BGP speakers. This reachability information is the IP prefixes which are reachable through neighboring Autonomous Systems (AS). The BGP speaker, also known as BGP peer, is the device which implements BGP. Latest description of the protocol is Border Gateway Protocol version 4 which is defined in IETF RFC [1, 2] along with its applications in RFC [3]. In this chapter, an overview of inter-domain routing is provided and working of BGPv4 is explained. The chapter also explains BGP routing information base and routing policy, followed by single hop and multi-hop sessions in BGP. In the end, the BGP route aggregation behavior is explained.

2.2 Inter-Domain Routing

The Internet, being a collection of many large and small networks linked together is envisioned as interconnected Autonomous Systems (AS). An AS itself is a group of peers with IP prefixes which are linked together and work on the same and well defined routing policy [4, 5]. In the usual setting, an AS works under a single administrative domain but if the routing policies of two interconnected network operators are the same, they can also share a single autonomous number. The autonomous systems can be

classified into three types, Stub Autonomous System, Multi-homed Autonomous System and Transit Autonomous System. Stub AS has connection to only one other AS. They are usually customers connected with transit AS in order to connect to the Internet. Multi-homed AS is same as Stub AS but has more than one connection to different AS, for sharing load and redundancy purposes. It does not carry traffic of one AS to another. Transit AS provides backbone services and carries traffic of one AS to another. It is densely connected with a mesh of other AS.

The addressing scheme used on the internet is based on the Internet protocol (IP) [6]. The IP address is used to indicate the network address and for uniquely identifying specific host on the network. Routers use routing protocols to exchange this network reachability information with other routers, telling the other routers which networks are accessible through them. Upon receiving advertisements from other routers, the router selects the best path and maintains a routing table for the available paths.

2.2.1 Interior and Exterior Gateway Protocols

Interior Gateway Protocols (IGPs) are Intra-Autonomous routing protocols that exchange routing information and maintain routing information within the autonomous system. Common IGPs includes Routing Information protocol (RIP), Open shortest path first (OSPF), Intermediate System to Intermediate System (IS-IS), Interior Gateway Routing protocol (IGRP) [7, 8, 9, 10, 11, 12]. Many IGPs can run simultaneously within an AS. Inter-Autonomous routing protocols are called Exterior Gateway Protocols (EGPs). EGPs maintain routing information between different AS. The De-facto standard EGP used on the Internet is Border Gateway Protocol.

Since an AS is under a single administrative domain, there is complete control over the routing behavior within AS. This makes it easier to implement protocol specific

changes in intra-autonomous routing protocol and to replace the routing protocol to meet changing security and performance requirements. But the same is not true for inter-autonomous routing protocols. Without control over the routing behavior of AS in other networks/domains, other AS have to be trusted for the routing information provided by them and to propagate IP prefixes correctly.

2.2.2 IP Address Space Delegation

Originally IP Address space was assigned to regional registries and organizations by the Internet Assigned Number Authority (IANA) according to the guidelines prescribed in [13]. However in 1993, US department of commerce selected Internet Corporation for Assigned Names and Numbers (ICANN) for this task. Guidelines setup hierarchy of address space delegation process, with ICANN being the top level authority delegates IP address space (IP prefixes) to the Regional Internet Registries (RIRs). Only a few RIRs operate in the world, each operating in large geopolitical region like the continents. RIRs further assign and allocate address space to Local Internet Registries (Internet Service providers (ISPs)) and similar organizations. ISPs further delegate IP prefixes to the customers or other ISPs when required. IP address space or IP prefixes are assigned to entities in the form of block of IP addresses. Address delegation authorities keep updated records of assigned allocation of IP prefixes. Figure 2.1 shows the IP Address space hierarchy model.

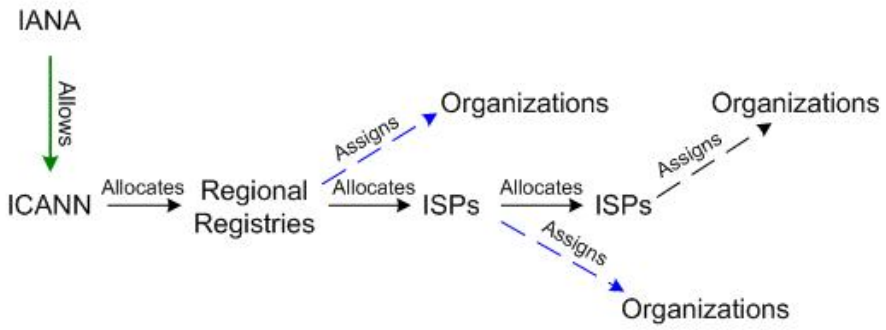

Figure 2.1: IP Address space delegation hierarchy.

2.2.3 AS Number Assignment Process

Autonomous System Number (ASN) uniquely identifies a network on the Internet. IANA (or ICANN) is a top authority for delegating AS numbers to regional authorities, which further assigns ASN to ISPs and similar organizations. ASN is only assigned to identify a network as an AS that implements its own routing policy. ASN can be public or private. Public ASN are visible on the Internet but private ASN are assigned by ISPs to customers who want to benefit from BGPv4 features like load balancing without unique identification on the internet. ISP drops the private ASN and attaches its own ASN before propagating received routing information from private peer. Figure 2.2 show the AS number assignment model.

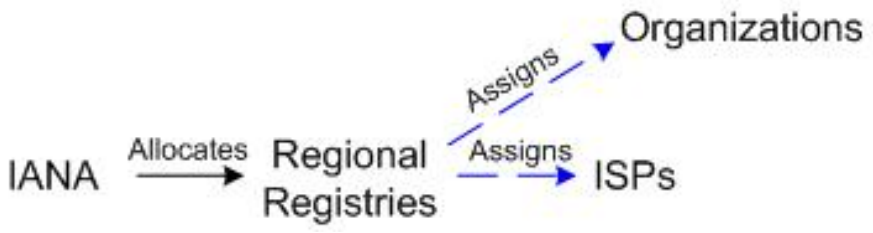

Figure 2.2: AS Number assignment hierarchy

2.3 BGPv4 Protocol Working

BGPv4 is the most dominant inter-autonomous/inter-domain routing protocol for the internet. It governs the exchange of network reachability information between BGP speakers/BGP peers (devices that implement BGP). The network reachability information is the IP prefixes which are reachable through neighboring AS [14, 15, 16]. Figure 2.3 explains the difference between IGP and EGP. Inter-autonomous routing information is exchanged via EGP and within AS via IGP [17, 18, 19].

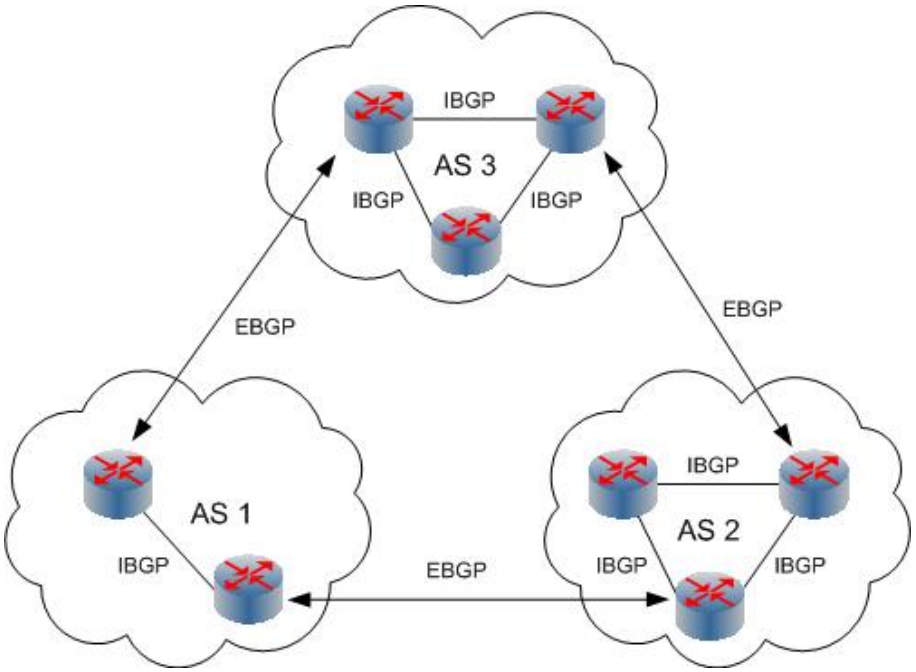

Figure 2.3: The difference between IBGP and EBGP.

Each AS may have one or more BGP speakers. BGP can be used both to exchange reachability information within AS and between different AS. Inter-autonomous routing between external peers is accomplished by using external Border Gateway Protocol (eBGP) and the link in between is called external link. The AS also maintains consistent view of global routing within the AS by interconnecting all the BGP speakers within the AS and exchange reachability information using internal Border Gateway Protocol (iBGP). BGP runs over reliable TCP protocol so it doesn't have its own transport level

error correction and detection mechanism. TCP port no. 179 is used to listen for incoming BGP connections. Peers in different AS are usually directly connected with each other using any OSI Layer-2 media but BGP also supports multi-hop configuration.

2.3.1 BGPv4 Messages

There are four types of BGP messages; OPEN, KEEPALIVE, UPDATE, and NOTIFICATION. Initially TCP connection is established between the peers. It has been made mandatory for BGP implementation to support signatures like MD5 for TCP segments to protect BGP peering sessions [14, 20].

2.3.1.1 OPEN Message

Once the TCP connection is established, any of the BGP peers initiate BGP protocol connection establishment and send the OPEN message. Upon receiving OPEN message, the other BGP peer confirms the connection by sending KEEPALIVE message. Then both BGP peers start sending their entire BGP routing table using UPDATE messages (if allowed by their export policy). BGP router sends only routes, which it uses itself. After receiving updates from neighbor BGP speakers, BGP speaker runs its decision process and selects the best route to the destinations as per its routing policy and further propagates this information to other connected peers. Each BGP peer executes the same process, making the routing information spread throughout the Internet.

2.3.1.2 KEEPALIVE Message

The KEEPALIVE message is sent periodically to check aliveness of BGP peer. If BGP speaker does not receive KEEPAILVE or update message within the specified time period (Hold Down time), a notification message is sent and BGP connection is closed.

2.3.1.3 UPDATE Message

UPDATE message is used both to withdraw previously advertised routes and to advertise new routes. Each time the BGP speaker receives update message, it runs its decision process to select the best available route to the destination and propagates the updates to other BGP peers. BGP speakers constantly exchange this Network Layer Reachability Information (NLRI) with its peers by sending UPDATE messages. NLRI contains the IP prefixes and associated path attributes. The most noticeable are AS_Path and Next_HOP attributes. AS_Path contains the list of AS numbers in sequence (if AS_path type is AS_Sequence) through which this information is traversed (AS_path type can also be AS_SET which is the result of route aggregation by some AS in the transit, as explained later in this chapter). The AS number on the right most side in AS_path is the originator of IP prefix. The receiving BGP speaker modifies the AS_Path attribute and attaches its own AS number on the left most side of this field. Next_HOP attribute contains the IP address of the next hop router (border router) through which particular prefixes mentioned in the NLRI field of update message can be reached. Every BGP speaker contains information about Next Hop router (direction) and the complete path (AS_path) to the destinations.

2.3.1.4 NOTIFICATION Message

NOTIFICATION message is sent if there is some error condition or if some special condition is true. After NOTIFICATION is sent, all BGP resources are released and TCP connection is closed. BGP runs its decision process and selects the alternative routes to unreachable destinations and propagates this information to other peers. This ensures delivery of outstanding before the connection is closed [14].

2.3.2 BGP Routing Information Base

BGP speakers maintain three different kind of routing information bases.

2.3.2.1 Adj-RIBs-In

This Routing Information Bases (RIB) stores routes learned from external BGP peers. BGP speakers logically maintain separate Adj-RIB-In for each external peer. These routes are used in the decision process to select the best route.

2.3.2.2 Loc-RIB

Loc-RIB stores routes that are selected after applying internal policies to the routes stored in Adj-RIBs-In. These routes are used by local BGP speaker. All BGP speakers within the autonomous system maintain consistent view of these routes.

2.3.2.3 Adj-RIBs-Out

This RIB stores routes selected by local BGP speaker to advertise to external BGP peers. The decision is made based on export policy of local BGP speaker. Figure 2.4 illustrates the use of RIBs and processing of UPDATEs received from external peers.

2.3.3 Routing Policy

Routing Policy is used in AS to decide which routes are to be imported into and which to be exported from the local routing table to external peers. Organizations running different AS are bound by business contracts called Service Level Agreements (SLAs). The SLAs include quality of service and legal liability issues, which effectively influence the routing policy in an AS to make routing decisions. Therefore, routing on the Internet is influenced by operational, economical and political factors [18, 21]. Routing policy

specifies conditions on the basis of IP prefixes, ASN and BGP community attribute which can trigger the specified decision.

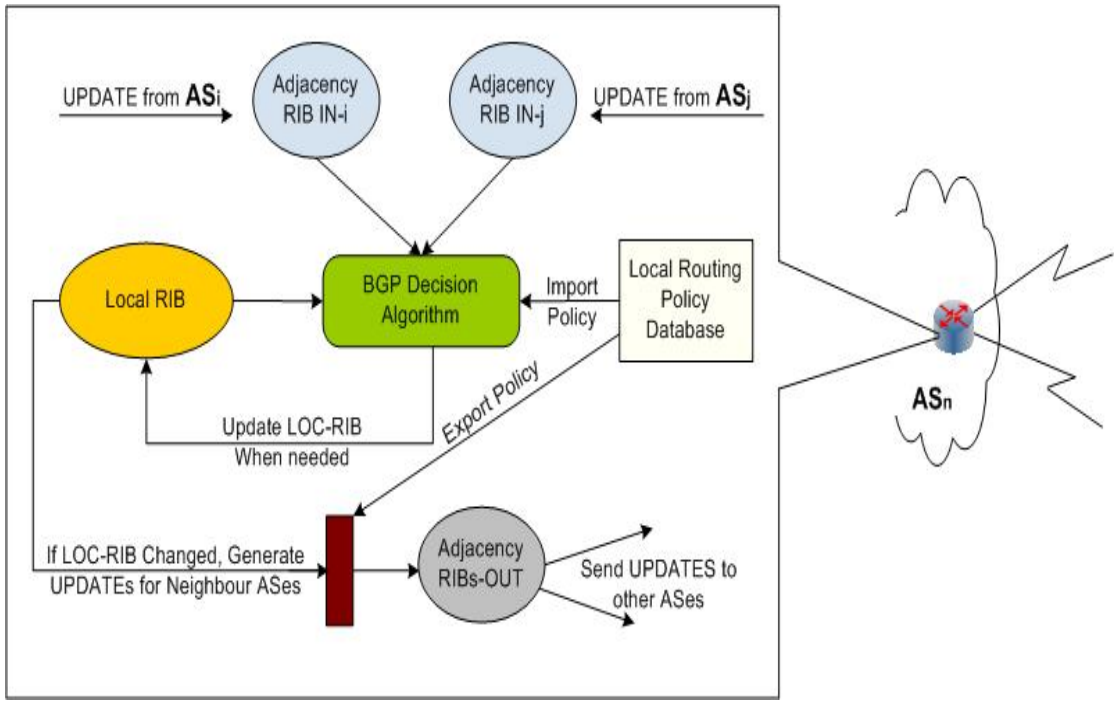

Figure 2.4: Processing of UPDATEs received from external peers.

2.3.4 Community Attribute

Initially, BGP speakers made decision about distribution of routing information on the basis of IP prefixes and on the value of AS_Path attribute. BGP community attribute was introduced to tag the destinations so that routing decisions can be made on the basis of TAG value [22]. Current BGP implementations support BGP community attribute and it is widely in use to make TAG based routing decisions.

2.3.5 Route Flap Damping

BGP connection can sometimes be lost due to excessive errors on the communication media and device failure. In such a case, BGP speaker release all resources and runs its decision process, selects alternative routes and propagates the new

information to other peers. Upon receiving updated information, other peers also run their decision process to update their routing table. Meanwhile if the lost BGP connection with the peer is re-established, BGP routing table is exchanged and decision process is run once again. Thus significant CPU utilization is wasted throughout the internet and BGP speakers bear extra processing load.

To counter such a situation, BGP Route Flap Damping was introduced [23, 24]. If the customers' BGP session is flapping then the ISPs configure routers to set timer that has to be expired before the lost BGP connection can be re-established. This gives some time for the communication link to be stable.

2.3.6 Single Hop and Multi-Hop Sessions

In the present internet environment, there is no router or hop in between two BGP peers. But because BGP runs over TCP so Multi-Hop configuration (i.e., with more than one routers present in between two BGP speakers) is also supported. In the current internet environment, multi-hop BGP sessions are only found between satellite downlink service providers and Multi-homed AS or Stub AS. In many parts of the world satellite downlink (only) services are cheaper than symmetric links. Because of asymmetric nature of internet traffic, Multi-homed and Stub AS customers prefer to use such services. Multi-hop BGP sessions are used to exchange routing information as satellite service providers' network is far away from the customer networks.

2.3.7 Route Aggregation Behavior

The decision making process of BGP speaker can aggregate a group of IP prefixes to make one single short prefix that represents all aggregated IP Prefixes, provided that IP prefixes can be represented as a single IP prefix according to the CIDR rules [25, 26].

Route Aggregation process also include aggregation of AS_Path attribute by listing AS numbers only once in AS_Path attribute, regardless of how many time it appeared in Aggregated routes. AS_Path type is set to AS_SET after route aggregation [14]. Aggregation of routes reduces the number of entries, which any BGP speaker needs to store in its routing table thus saving memory space and CPU utilization on the BGP speakers in the internet.

BGP speaker's decision-making process always prefers longest prefix (that represent more specific route) over the relatively short prefix (that represent a less specific route). This feature is useful to support redundancy configurations in multi-homed ASes and to avoid manual route aggregation mistakes [14, 26].

2.4 Summary

This chapter briefly describes the working of BGP. An overview of the protocol's functionality is given with details on its message types, routing information bases and routing policy. In addition to this, an account of BGP operational features like the community attribute, route flag damping and session types (single and multi hop) is made. This chapter also describes the route aggregation behavior in BGP.

A Threat Model for BGP

3.1 Introduction

The internet was originally designed for a group of trusted networks to communicate with each other. Over the period of time, scientific advancements in communication technology and business interests encouraged new networks to become a part of the internet. Today it forms a mesh of different interconnected networks which do not hold trust relationship. Likewise, BGP was designed to exchange reachability information within and between trusted AS. Although BGP has been refined over the period of time to incorporate improvements identified with its operational experience, its open design remains vulnerable to a lot of attacks. This chapter presents an overview of the threats and attacks on BGPv4 protocol [14]. The consequences of the security threats are also discussed, followed by the limitations of BGP protocol that create vulnerabilities. In the end, the existing mechanisms to secure BGP are discussed.

3.2 BGP Session Attacks

As described in Chapter 2, BGP routers are either directly connected with each other or through a series of routers in case of multi hop configuration. BGP session attack is explained through a scenario where two BGP routers, router A and router B maintain BGP connection and exchange routing updates. An adversary C intends to maliciously manipulate the information in transit between the two routers. As BGP design is based on trust relationship, routers A and B trust each other for correctness of received information. This trust relation can be exploited by C. In case C has full control over the

communication link between A and B, also the MD5 signature option for protecting TCP segments is not configured on the link between A and B (as even today most ISPs do not configure this option on peer links). Figure 3.1 illustrates this scenario.

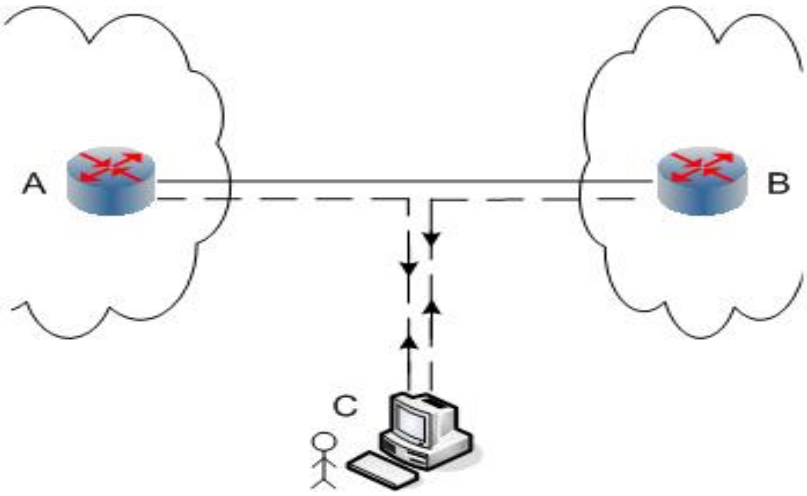

Figure 3.1: Attacks scenario for simple BGP session.

3.2.1 Confidentiality Attacks

All the information between A and B is sent without any confidentiality protection mechanism, so adversary C has access to all the information in transit over the communication link. Routing data confidentiality is not a strict requirement over BGP. Organizations and customers are normally bound for confidentiality protection by business agreements. The routing policy of any ISP clearly reflects the business agreements with their customers. By eavesdropping on the routing information which A and B send to each other, C can try to determine their routing policy and get valuable information about their business relationship. This is an example of industrial espionage.

3.2.2 Integrity Attacks

Instead of just passively listening to updates, adversary C can play an active role and can launch a number of attacks to violate integrity over BGP.

3.2.2.1 Insertion Attacks

Once the connection is fully established between A and B, C can insert spoofed messages into the message stream. C manipulates and sends Update message to B with spoofed source as A. This enables C to inject false routing information into the B's routing table (withdrawal of A's previously advertised routes are advertisement of victim's network with more specific prefix). This information is propagated from B to other peers, causing A's withdrawn networks to disappear from the internet and victim's traffic will start coming towards A and finally be dropped, creating the black hole effect. Adversary C can also terminate the BGP session between A and B by sending spoofed malformed BGP message to either of A or B. Thus with the ability to spoof messages, C can use BGP messages in a number of ways to disrupt routing services.

3.2.2.2 Modification Attacks

Adversary C can act as a Man-in-the-middle to temper with messages exchanged between A and B and modify the messages to insert false routing information in A and B's routing table or to terminate BGP connection.

3.2.2.3 Deletion Attacks

Adversary C can delete messages coming from A and B, selectively removing messages from the message stream. For example if C deletes certain UPDATE messages, it directly results in routing misbehavior on both sides. BGP peers periodically send KEEPALIVE messages to each other to check the other peer is still alive. If A doesn't receive periodic KEEPALIVE or UPDATE message within Hold_Down_Time (its value is negotiated through OPEN message at the time of BGP connection establishment) from

B, it sends a NOTIFICATION message to B and terminates the BGP connection. This will release all BGP resources allocated for connection with router B.

3.2.2.4 Replay Attack

The adversary C can save older messages exchanged between A and B and replay them after some time with the intention to cause routing misbehavior. In case A had previously withdrawn some prefixes through update message and later re-advertised them, C can send the saved UPDATE message to B through which A previously withdrew routes. This would withdraw A's prefixes once again, resulting in disruption of services.

3.2.3 Session Termination Attacks

If C can successfully insert spoofed messages in the message stream between A and B, then C can easily terminate BGP session between A and B. The BGP state machine enables C to accomplish this task in a number of ways due to a number of related vulnerabilities in BGP state machine model [18, 27]. For instance, C can simply send malformed UPDATE message to terminate the connection or C can directly send spoofed NOTIFICATION message indicating some error condition. The receiving BGP speaker thinks that the other BGP speaker has sent the message due to some error condition, following immediate closing of BGP session after which, the BGP connection will have to re-establish. If C successfully terminates the BGP connection between A and B a few times, the configured BGP Route Flap Damping Feature [23] on either router A or B would cause a longer delay for the connection to be established again. Thus by sending just a few spoofed messages, each a small interval apart, C can cause two AS to be disconnected successfully for sufficiently long time even when they are physically directly connected.

3.2.4 MD5 Signature for TCP Segments

Implementing MD5 signatures for TCP segments for securing BGP peer sessions is now a mandatory requirement in the current version of BGP [14, 20]. This option is not an integral part of BGP protocol itself [27]. Although using MD5 signatures on TCP segments (which is an underlying transport protocol for BGP) may reduce the risk of insertion, modification attacks and replay attacks but MD5 signature relies on manual keying. Managing manual keys for every BGP peer within a large network is resource intensive. Additionally, the keys need to be changed regularly otherwise MD5 is vulnerable to cryptanalytic attacks. Key management procedures also consume resources and can cause service disruption as TCP connection is dropped between the BGP peers.

3.3 Origin Falsification Attacks

A BGP speaker in one AS can advertise false routing information to its peers in other AS through update messages; neighboring BGP speakers accept this false information because of the weak design of BGP based on trust relationship. Such information further propagates to other BGP speakers on the internet. Thus a single mis-configured, malicious or compromised BGP speaker can poison the routing table of many other BGP speakers. Consider a case where a prefix 210.57.10.0/24 is actually owned and advertised by AS7500. A malicious BGP speaker in AS7070 (on some other side of the internet) can advertise this prefix 210.57.10.0/24 to its peers and can falsely claim that it is originating this prefix. The neighboring peers would accept this false information as it came from a trusted peer. A large number of routers accept this information and believe that 210.57.10.0/24 is accessible through AS7070. The routers would start sending data traffic for 210.57.10.0/24 to AS7070. The actual owner of the prefix AS7500 will not

receive data traffic which should have been sent to it. This attack is known as prefix-hijacking. A malicious AS7070 may wish to drop all of the traffic destined for 210.57.10.0/24, the effect is called black holing [18, 28]. If AS7070 uses the IP addresses of victim's network (i.e., 210.57.10.0/24) and assigns the addresses to its hosts, it can also launch more dangerous attacks. Consider an example of a web server for online banking facilities that runs on IP address 210.57.10.20 on TCP port 80. A malicious attacker in AS7070 can run the same kind of web server on the hijacked IP address 210.57.10.20 with identical web interface as of the original server, and can launch Phishing attack to capture sensitive customer information. The part of the internet which accepted the malicious routing information would direct the request for IP address 210.57.10.20 towards AS7070, which is impersonating as online bank service for customers [technical trends in phishing]. Thus the attacker doesn't have to spoof the IP Address to launch an attack (which allows only one way communication); it can actually hijack the IP address, enabling the adversary to launch devastating attacks.

Another kind of fraudulent origin attack can be launched by exploiting the BGP Route Aggregation feature. In BGP decision making process, more specific prefixes (longest prefixes) are always preferred over less specific prefixes (comparatively short prefixes) as already Chapter 2. Also BGP aggregates routes in order to save the space used to store BGP routing table. This phenomenon can be exploited by malicious AS7070 by falsely originating de-aggregated prefixes (many more specific routes) for which any aggregated prefix (single less specific route) is already present in the global BGP routing table. When the neighboring BGP speakers receive more specific routes to the destinations, they further propagate this information, causing all BGP speakers on the internet to update their routing table with the false information and send data traffic for

these prefixes towards malicious AS7070. The malicious AS can also drop all the packets. The real owners of IP prefixes will not receive any data. Historically, mis-configured routers are the source of such type of incidents, which caused wide scale communications disruption. This kind of de-aggregation attacks can even represent Denial of Service for the BGP routing protocol itself as advertising large numbers of longer prefixes can cause BGP traffic and router table size to increase exponentially [27].

3.4 Path Subversion Attacks

A malicious BGP speaker can mischievously alter the AS_Path attribute of the received UPDATE message before forwarding routing information to other peers. Tempering with the path attributes in the transit can misguide other BGP speakers to believe that the UPDATE information actually traversed through the path mentioned in AS_Path attribute when actually it did not. BGP, being a path vector protocol, believes that the path in terms of series of AS is actually the proper path to the destination. Traffic can thus be directed towards sub-optimal path or the path which does not use expected routing policy or towards an AS with no route to the destination [27]. Malicious path manipulation can shorten the valid path (referred to as truncating attacks) or elongate the AS path to cause delay in data traffic. An attacker can also choose a victim AS to launch a bandwidth consumption Denial of Service attack (DOS) against it. For this purpose, the attacker simply uses a compromised BGP speaker to put victim's ASN in AS_Path attribute of the tempered UPDATE message and spread this information on the internet. Many BGP speakers take this information for granted as it comes from a trusted peer. The victim AS is flooded with unexpected data packets creating a bottleneck in traffic flow to the victims AS. An attacker can also divert the traffic towards a malicious AS who can eavesdrop on the data traffic to capture sensitive information.

In the context of Path Subversion Attacks on BGP, the adversaries can be classified into two types; isolated adversaries and colluding adversaries. Isolated adversary can launch an attack independently (e.g., a malicious AS) while the colluding adversaries are a group of malicious attackers which can have access to many compromised routers and help each other in launching an attack. Colluding adversaries can also establish tunneled links with each other to transport routing advertisements and data traffic [29]. Due to the help of all attack participants, colluding adversaries are able launch more sophisticated attacks in comparison to isolated adversaries. If a malicious AS A, wants to eavesdrop on the data traffic of some other AS but also wants that data traffic to be delivered back to its legitimate destination so that the victim AS does not suspect eavesdropping. A can conspire with two other malicious AS, B and C, in order to successfully launch an attack, with B being nearer to the victim AS. C simply alters the path in all received UPDATE messages for victim's prefixes and include the AS number of A in the path. The data traffic for victim AS is routed through A, which after eavesdropping on data, sends the data to B through tunneling. B can forward the data towards actual destination as it has a valid route to the victim AS.

3.5 DoS Attack Through Underlying TCP Protocol

Many of the above-discussed attacks can deny or degrade different kind of routing services for BGP. Modification, insertion, deletion and session termination attacks between two BGP peers can disconnect two AS logically. IP prefix hijacking can deny services for legitimate prefix owner. Prefix de-aggregation attack can deny services to many different prefix owners and also cause routing table to be exploded with too many routes. Malicious manipulation with path attributes can misdirect data packets to

networks that either cannot deliver the data to the destination or cannot handle such a large volume of data and collapse.

In addition to these DoS attacks, BGP is indirectly vulnerable to TCP attacks as BGP uses TCP as its transport protocol. There are a number of attacks which can directly target the TCP implementation of the router on which BGP is running. Few such attacks are discussed below.

3.5.1 TCP SYN Flooding Attack

TCP is a three-way handshake protocol. Initiator of the connection sends SYN message to open the connection, the receiver acknowledges the SYN by sending SYN ACK message. Initiator of the connection finally send ACK message and connection is established. An attacker can send spoofed TCP SYN packets on the BGP router's port 179 impersonating that the packets are coming from legitimate peer. Upon receiving each SYN message BGP router reserves some memory resources and acknowledges the SYN, which in turn is never acknowledged back as legitimate peer never initiated connection attempt. An attacker can flood victim's router with spoofed TCP SYN packets and the router will ultimately run out of resources.

3.5.2 Packet Flooding on TCP port 179

Another similar kind of attack can be lunched on TCP port 179 by just sending the flood of packets. The packets directed to TCP port 179 are passed to the BGP process, flooding a router with TCP port 179 packets is an effective DOS attack on the router [27].

3.5.3 TCP RST Attack

TCP RST message is used to reset the TCP connection. An attacker can send the spoofed TCP RST message to BGP router to reset the TCP connection between the peers, causing the BGP connection between the peers to terminate. This requires the knowledge of TCP sequence number which must be within the receive window [30, 31].

All TCP attacks discussed above cause the BGP router to crash or terminate the TCP session between two BGP peers. Either way BGP connection is lost. In such event, the BGP session is terminated and re-established repeatedly, indicating that the route is flapping which triggers the BGP Route Flap Damping feature on neighbor BGP speaker. This stops the BGP connection to be established for even longer time, results in denial of service for elongated time. For peer/connection aliveness test, BGP relies on periodical KEEPALIVE messages. In case a KEEPALIVE is not received within the specified time; BGP connection is terminated following the NOTIFICATION message to other peer. Data flooding attacks consume significant bandwidth of the link, and contribute greatly in termination of BGP sessions, as communication link may not be able to pass BGP KEEPALIVE messages.

3.6 BGP Mis-Configuration

Although the attacks described above present critical threats to the BGP routing on the internet but surveys show wide scale internet outage (due to internet routing infrastructure) was caused by mis-configured BGP routers. Mis-configuration is a situation where operator of the network makes a mistake in configuring BGP on the routers. Configuring BGP on the routers requires the knowledge and understanding of how BGP actually works and how BGP routes can be distributed to internal routers

through IGP. Operators make mistakes, which directly affects BGP routing like any other BGP attack. Studies show that 75% of all new prefix advertisements are the result of configuration mistakes on BGP [32, 33]. The first category of mis-configurations includes a BGP speaker that can accidentally inject routes into the global routing table, including address space hijacks (i.e., accidental hijacking). In the second category, a BGP speaker can advertise routes which it should have filtered according to the AS export policy [34], [Route Views].

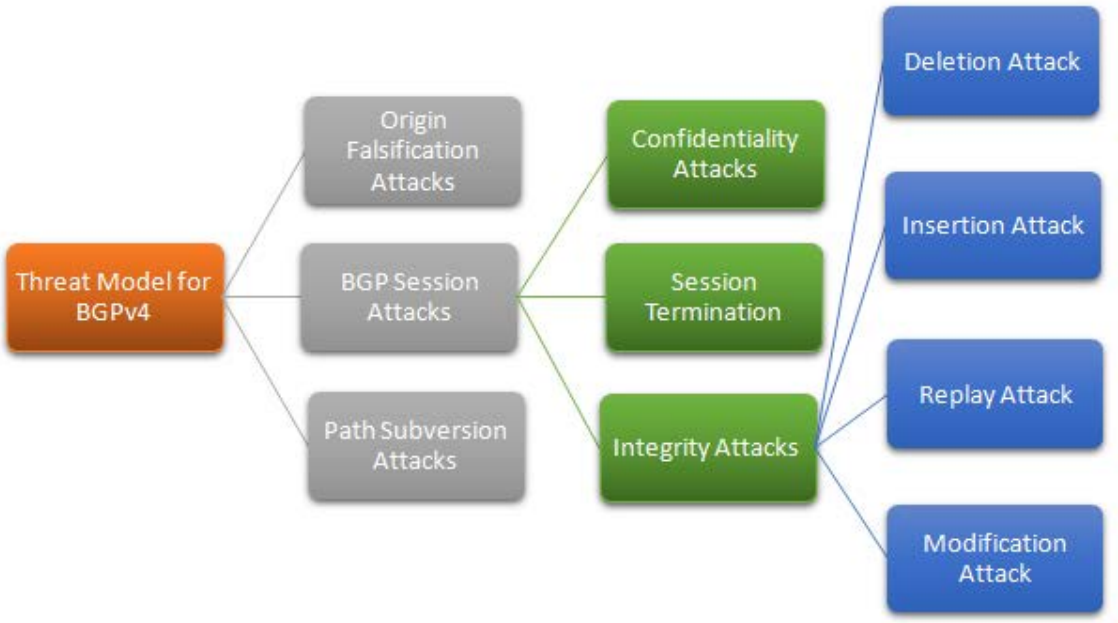

Figure 3.2 Threat Model for BGPv4 Routing Protocol.

3.7 Consequences of Attacks

The consequences of BGP attacks are devastating in the connected world and current information age. BGP attacks are unlike other attacks that target single host or single web server, causing limited damage. Attacks on BGP routing protocol affect a large number of hosts that are rendered unable to communicate with rest of the internet. Complete series of networks can disappear from internet and their traffic can be dropped.

Malicious path manipulation can result in network congestion, black hole attack, delayed network traffic and routing loops [27]. Session terminations can disconnect two AS and force other BGP speakers to re-calculate routes increasing computational complexity. Sensitive data traffic can be eavesdropped by path manipulation. De-aggregation attacks affect the internet as a whole, network and router resources can be exhausted. Prefix hijacking not only disrupt services for original owner but also provide the foundation for eavesdropping and phishing attacks. Spammers can hijack the IP address space and use it for spamming [18, 28, 25, 36]. These attacks affect all services, which are dependent on internet infrastructure.

3.8 Limitations of BGP

Analysis shows there are three fundamental vulnerabilities, which are the main cause for most of the risks associated with BGP. The first limitation in BGP is absence of internal mechanism to provide strong protection of the integrity, freshness, and peer entity authenticity of the messages. Integrity of the messages ensures that the message is not modified in the transit; freshness ensures that the received message is new one and not the replay of old message, peer entity authentication provides the assurance that the received messages came from authentic source. The second limitation in BGP is deficiency of validating the authority of an AS to announce NLRI information. This is relevant to AS's ability of originating IP prefixes, which it doesn't own. Any AS can hijack prefixes owned by other organizations. Third limitation in BGP is that it does not specify any mechanism to ensure the authenticity of the path attributes announced by an AS. Any BGP speaker with malicious intention is capable of tempering with the path attributes of the BGP updates, which can change the normal routing behavior, routing the traffic to malicious networks [27].

Because of the first vulnerability, a mandatory support of MD5 to protect TCP segments has been added into the current version of BGP [14, 20]. If BGP sessions are protected by using MD5 signatures for TCP segments, message integrity and peer-to-peer authenticity can be achieved. This is subject to the security of MD5 (with reference to hash collisions; two messages with the same MD5 hash value) [37, 38].

3.9 Existing Security Mechanisms

The above mentioned attacks and incidents highlights security weaknesses and risks associated with existing BGP routing protocol. Researchers around the world have contributed to design secure and practical solutions for BGP. The currently available techniques to add security to BGP are discussed below in view of above mentioned limitations of BGP.

3.9.1 BGP Session Protection

The currently deployed solutions for protecting BGP sessions between peers are discussed below.

3.9.1.1 MD5 based TCP Segment Protection

Researchers suggested the use of MD5 signature for TCP segments that carry BGP routing information [20]. The security of MD5 output is based on the assumption that it is non-invertible and collision resistant [39, 40]. Non-invertible means that MD5 uses one-way function and collision resistant means that it is computationally impossible to produce same MD5 output with two different messages. The security also relies on the secret key used to create messages digest. Because of these properties of MD5, researchers suggested using it to protect TCP traffic that carries BGP information

between two peers [41]. For use within BGP, every TCP segment carries 16 byte MD5 messages digest, computed using 4 factors; TCP pseudo-header (i.e., source IP address, destination IP address), TCP header (excluding options fields), TCP data and shared secret value [20, 42, 43].

Upon receiving messages, the receiver computes the same value using MD5 and compares the computed value with the received one. If both values match, the TCP segment is accepted and rejected otherwise. This mitigates the risks of modification and insertions attacks which are based on sending spoofed messages since it would require the knowledge of secret key for generating digest. There are few limitations of using this technique. Firstly, configuration and management of keys in a large network is difficult. An AS with having ‘n’ number of BGP peers needs to keep track of ‘n’ different keys, one for each BGP peer. Secondly, key management should ensure key synchronization, failure of which can terminate BGP session since TCP segments will not be authenticated without the updated key. The time required for changing the keys may be more than hold down time set at both BGP speakers, which terminates BGP session.

3.9.1.2 IPSec

Due to the limitations of using MD5 for BGP session protection, many researchers have proposed to use IPSec protocol suite to secure BGP sessions. IPSec is a protocol suite which provides protection at network layer of OSI Reference model [44 – 48]. It was specially designed to provide integrity, authentication and/or confidentiality services to Internet Protocol (IP) [6]. IPSec can operate in two basic modes, Transport and tunnel mode. The transport mode provides host-to-host security by providing protection to IP datagram payload and few of the IP header fields while the Tunnel mode provides protection from intermediate IPSec gateway to Intermediate gateway (end hosts need not

to be IPSec aware) and protects the entire IP datagram (including source and destination IP address). The keying material for IPSec is provided by IKE protocol which provides the flexibility of negotiating the cryptographic algorithms to be used between participating parties [49, 50]. Due to flexibility and reasonable level of security provided by IPSec, it is suitable for protecting peer-to-peer BGP sessions.

3.9.1.3 Generalized TTL Security Mechanism

Time-to-live (TTL) is an 8-bit value used in IP protocol in order to discard packets lost on the internet [6]. When a packet is passed through a router, the TTL value is decremented by one. Sometimes due to the temporary routing loops in global routing table and network failures, some packets do not find their destination and are lost. Routers discard them when the packet's TTL value is zero. Generalized TTL Security Mechanism is a simple technique based on TTL field of an IP datagram which adds an extra layer of security against attackers to mount remote spoofed attacks [51]. GTSM suggests that every BGP peer should set the TTL value to its maximum 255 before sending data to its peer and other BGP peer should only accept data packets with TTL 255. Remote attacker can send spoofed BGP messages but attacker cannot ensure TTL value to be 255 as it is decremented by one when the packet passes through any router. GTSM also suggest TTL value for the Multi-hop BGP sessions to be 255 (configured-range-of-acceptable hops). This is a very simple way to mitigate remote spoofing attacks and is widely adopted in industrial applications.

3.9.2 Protection against Prefix-Hijacking

Studies show that router mis-configurations and compromised BGP router can cause wide-scale service disruption. Mechanisms are devised to minimize such kind of

incidents. The current best practices mechanisms to mitigate prefix-hijacking and maliciously originating IP prefixes are discussed below.

3.9.2.1 Ingress and Egress Filtering

Filtering the incoming BGP updates received from external peers before deciding to add these routes into local routing table is called ingress filtering and filtering routes before advertising them to other peers is called egress filtering [18, 28, 52]. Figure 3.3 illustrates the use of ingress and egress filters on border routers.

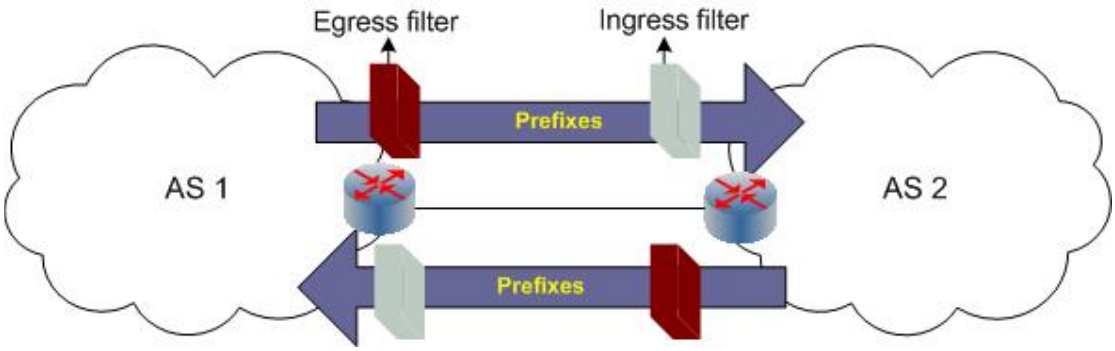

Figure 3.3: Ingress and egress prefix filtering at border routers.

Defensive ingress filtering can counter variety of neighboring mis-configurations and intentional attacks. BGP best practices suggests that an AS operating as an ISP should accept only those prefixes from their stub or multi-homed customers which they declared to own. ISPs can validate the ownership of IP prefixes and ASN through Regional Internet Registries (RIRs) or customer's peers. If the ISP has assigned the address space to the customer then ISP must allow that customer to advertise it back. ISPs must never allow Martian address to be advertised [63]. Martian addresses also known as Documenting Special Use prefixes (DSUA) have been reserved for special use and therefore they must not be announced on the public internet [52]. Martian addresses don't change so once the ingress Martian filter is applied on the routers then there is no

need to update it. ISPs must also filter unallocated address space (that IANA has not allocated to regional registries) as any malicious AS can advertise these addresses and may use it for other attacks. But the filtering rules must be updated regularly to cater for latest changes [53]. It is also recommended that AS filter and reject prefixes longer than 28 Bits (i.e., > /28) and less than 6 bits (i.e., < /6). An ISP should receive only those routes which the other ISP indicated at the time of contract and an ISP must export only those routes which it indicated at the time of contract. Careful ingress and egress filtering at ISP-customer peering edge minimizes the risk of prefix-hijacking.

3.9.2.2 BGP Packet Filtering Using ACLs

The packets received on router's TCP port 179 are filtered through Access control lists (ACLs) before passing them to CPU. This provides defense against SYN-Flooding attacks since BGP routers only accept SYN requests from their neighbors' IP address.

3.9.2.3 Prefix Limiting

Prefix hijacking attacks on BGP motivated the addition of prefix limiting feature in the current version of BGP. Prefix-limiting sets the upper bound on the number of address prefixes which a BGP speaker accepts from its neighbor. When this upper bound is reached, a BGP speaker may discard the new prefix advertisements from peer or may terminate that particular BGP session [14]. The BGP speaker's decision of either discarding new prefix updates or terminating BGP session depends on the configured settings. But reaching upper bound of maximum prefix length indicates something being wrong with neighboring BGP peer. This security feature effectively minimizes the effect of prefix de-aggregation attacks.

Prefix filtering along with defensive route filtering defends BGP against unauthorized route injections into global routing table. If filters and prefix-limiting is not implemented at any point in the internet, any mis-configured or compromised BGP router in the internet can be used to launch attacks and create wide-scale internet outage.

3.9.3 Internet Routing Registry

Internet Routing Registry uses the concept of distributed database development to contain latest records of all routes originated by every AS, every AS's routing policy and the connected peers. This information is stored in routing registries in a standard format using Routing Policy Specification language (RPSL) and can be accessed by anyone to troubleshoot routing problems [54, 55]. There are number of routing registries (RR) actively operating on the internet, well known being APNIC, RIPE, RADB [56, 57, 58]. Every Routing Registry keeps the mirror of all the data of every other routing registry.

Routing registries simplifies filtering process. Two AS may agree to filter incoming routes from each other on the basis of other AS's registered routes in the RR. ISPs can make it compulsory for their customers to register their AS and prefix ownership records in internet routing registry (IRR). ISPs then deploy ingress filters to check the records in IRR and allow the IP prefixes and AS path on the basis of information retrieved from the registry. Figure 3.4 illustrates how ISP's ingress can be updated automatically using the information stored in routing registry [63, 59, 60].

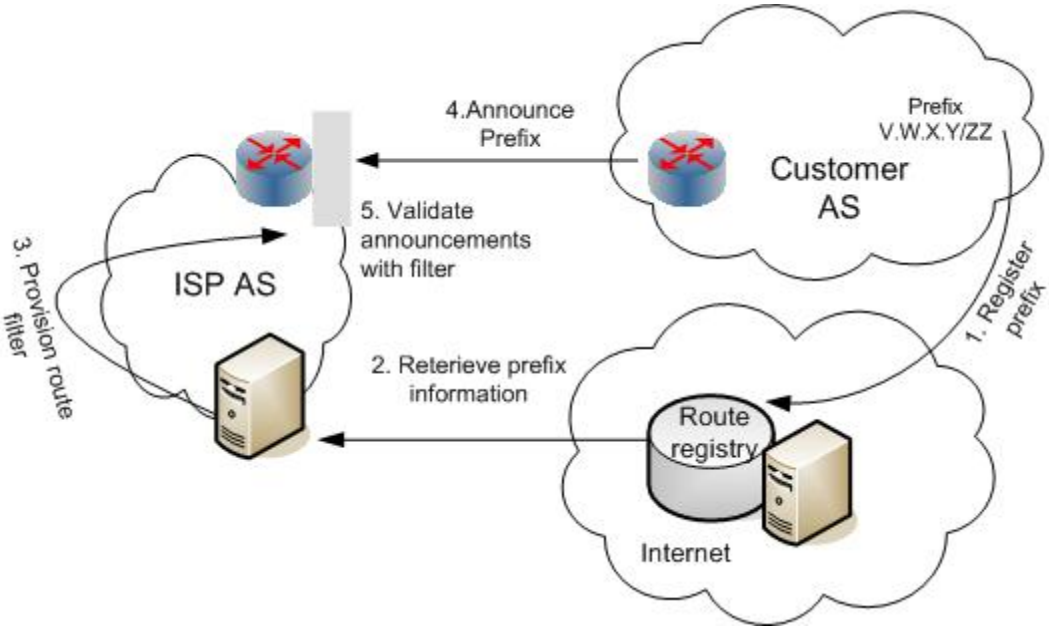

Figure 3.4: Automatically updating prefix filters through information stored in route registry.

3.9.3.1 Strong Security Policies

Other than available solutions to secure BGP routing on the internet, security policies also need to be enforced to minimize the incidents of mis-configuration and compromised routers. Procedures should be defined to configure and change routers settings and ensure accountability. Internal security practices include SSH access to routers based in RADIUS server authentication, logging the access to routers, and keeping record of router's configuration changes [61]. Security solutions and policies are necessary to build a secure solution for BGP in general.

3.9.4 Gradus

Gradus is a commercial system to detect prefix hijacking in real time. It monitors BGP traffic in real time and maintains the global view of BGP routing. Gradus customers install the client application on their work station, which connects with Gradus server. If some AS on the internet originates the customer's prefixes registered with Gradus server, the system will generate alarm on clients indicating the ASN of the autonomous system

that is trying to hijack the customer's prefix [62, 63]. This system is very suitable for critical online services to prevent IP prefix hijacking.

3.10 Summary

This chapter briefly describes the threats and possible attacks on BGP. The consequences of these attacks are discussed that can cause wide scale outage of networks supported by BGP infrastructure. The limitations of BGP protocol that enables adversaries to launch the attacks are also explained. This is followed by a description of the existing mechanisms and practices to secure BGP against routing attacks.

Analysis of Existing BGP Security Architectures

4.1 Introduction

The latest advancements in technology have facilitated the growth and use of internet. This has also changed the threat environment for the internet and systems and applications dependant on it for providing numerous services. Threats on BGP have also increased exploiting the open design and insufficient security provisioning for the most dominant global routing protocol. There are currently practiced techniques and mechanisms that exist to secure different aspects of BGP routing but there is no comprehensive solution to address all risks associated with BGP. This chapter presents detailed analysis of BGP security extensions and some other defined solutions to mitigate threats to BGP. Operational working, network arrangement and security aspects for each extension is explained and evaluated in detail.

4.2 Secure-Border Gateway Protocol

Secure-Border Gateway Protocol (S-BGP) was the first complete and concrete solution that addressed BGP security aspects. It was proposed by BBN technologies and tested the proof-of-concept prototype of S-BGP on DARPA's test bed network. S-BGP provides protection against Origin falsification attacks by using address attestations, Path subversion attacks by using route attestations and BGP session attacks by using IPSec ESP with NULL encryptions [64].

S-BGP is designed to be backward compatible with existing BGP routing protocol so that it can be deployed incrementally within the existing internet infrastructure. The

main elements of this approach are two Public key infrastructures (PKIs) based on X.509v3 certificates, a new optional transitive BGP attribute containing attestations and using IPSec between BGP peering sessions.

4.2.1 A PKI for Address Allocation

The sole purpose of this PKI is to issue a certificate to grant the ownership of the portion of the IP address space. The proposed infrastructure for this PKI is based on same hierarchy of the current address allocation system. ICANN being the root authority allocates IP address space to regional registries by issuing a X.509v3 certificate for their identity, public key and the allocated address space. ICANN signs the certificate with its private key. The regional registry then further allocates sub-blocks of its address space to ISPs/DSPs by issuing a certificate to ISP/DSP in the same way.

4.2.2 A PKI for Assignment of AS and Router Associations

This PKI is responsible for issuing three kinds of certificates to authenticate AS, BGP speakers and the relation between AS and the BGP speakers. The root in the hierarchy is ICANN which allocates AS numbers to regional registries by issuing the certificate with allocated range of AS numbers and registry's public key. The regional registry then assigns one or more AS numbers to an organization (e.g., ISPs/DSPs) by issuing certificates to organization. An organization is then authoritative of issuing two different kinds of certificates. The first certificate which an organization issues will bind each AS number (owned by itself) with that AS number's public key. The second certificate represents the relationship between an AS and the BGP speaker. One certificate will be issued for every BGP speaker in an AS. This certificate binds router's DNS name, router ID, AS number and router's public key.

4.2.3 Attestations

Attestations are considered as special kind of certificates in which subject is an AS and the issuer is the organization which owns the address space. These certificates prove that an AS is authorized by an organization to advertise the path to IP prefixes owned by it. There are two kinds of attestations, Address attestations and Route attestations. Address attestations are certificates issued by the owner of the address space to allow an AS to originate the prefixes owned by an issuer. These certificates bind the AS number with the address space that it can originate. An owner of address space signs the certificate with its private key.

Route attestations are used to check the integrity of the AS_Path attribute. The S-BGP introduces new optional transitive BGP attribute to carry route attestations. The route attestation uses compact encoding scheme to fit into UPDATE message. The first AS which originate the NLRI information add its AS number in path attribute, signs path attribute and other sensitive attributes of the UPDATE message with its private key and put the signed information in optional transitive attribute called ‘attestations’ of that UPDATE message. Upon receiving this message, the next S-BGP speaker checks that the NLRI information is originated by authorized autonomous system and validates the signed information in attestation attribute to ensure that message traversed through the path mentioned in AS_Path attribute. S-BGP speaker adds it’s AS number in AS_Path attribute on the left-most side of the list and signs AS_Path attribute along with other sensitive attribute (e.g., BGP community) with its private key. This information is added in attestation attribute of that Update message. All S-BGP speakers that receive this information, repeats the same validation before propagating this information any further.

This allows every AS to validate the complete path. Figure 4.1 illustrates this concept graphically.

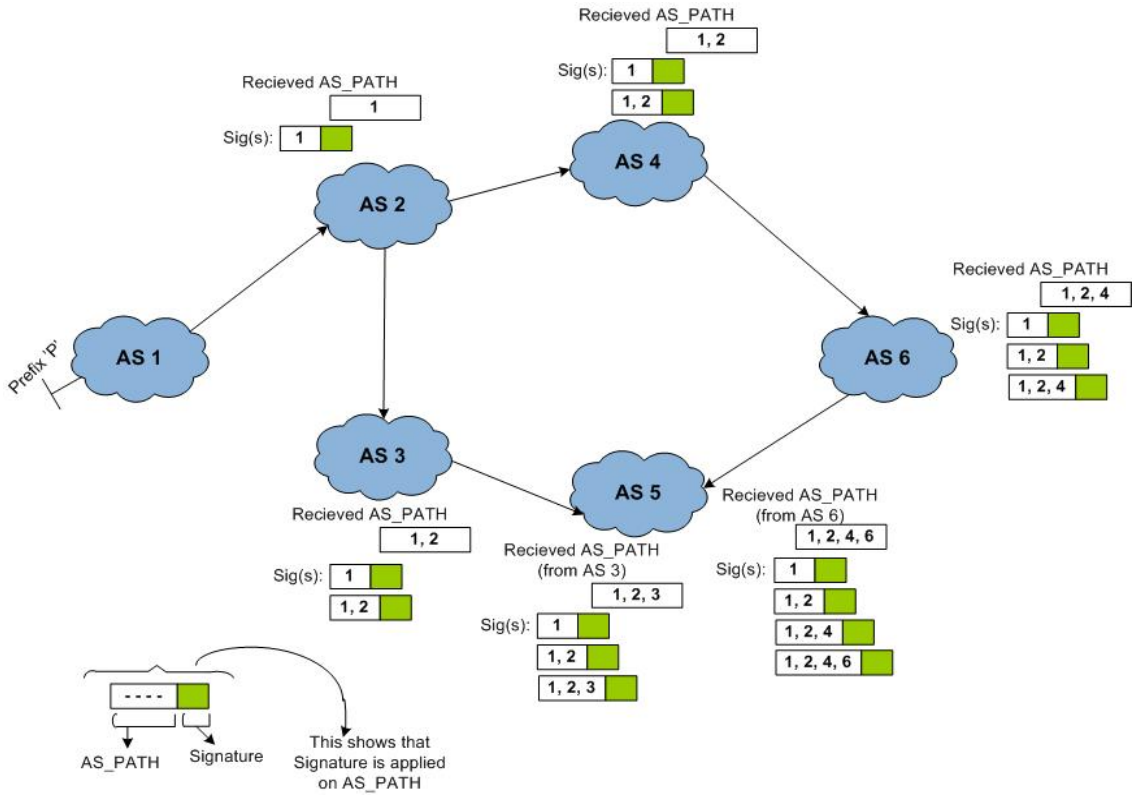

Figure 4.1: Path authentication in S-BGP.

4.2.4 Distribution of Certificates, Address Attestations and CRLs

In order to validate the address and route attestations, each non-leaf S-BGP speaker needs the valid public keys of all organizations, all AS and all BGP speakers on the internet. This information is stored on the device running S-BGP. The memory requirement to store all these certificates and attestations is significantly high for S-BGP speaker. S-BGP proposed to use out-of-band mechanism for distributing certificates, attestations and Certificate Revocation Lists (CRLs) which make use of two tiers of repositories. The top tier consists of several replicated sites to store all information. The second tier repositories are AS level repositories. Each AS maintains a separate local

repository which downloads/requests the complete database and other information of top tier repository using FTP or TFTP. Second tier repository validates all received certificates, CRLs and address attestation to produce more compact information for storage on S-BGP speakers. This pre-processing by AS's local repository significantly reduces the memory requirement and processing time to validate attestations. The certificates revocation issues are also handled by AS's local repository. Figure 4.2 illustrates the concept of 1st tier and 2nd tier repositories for distribution of certificates out of band.

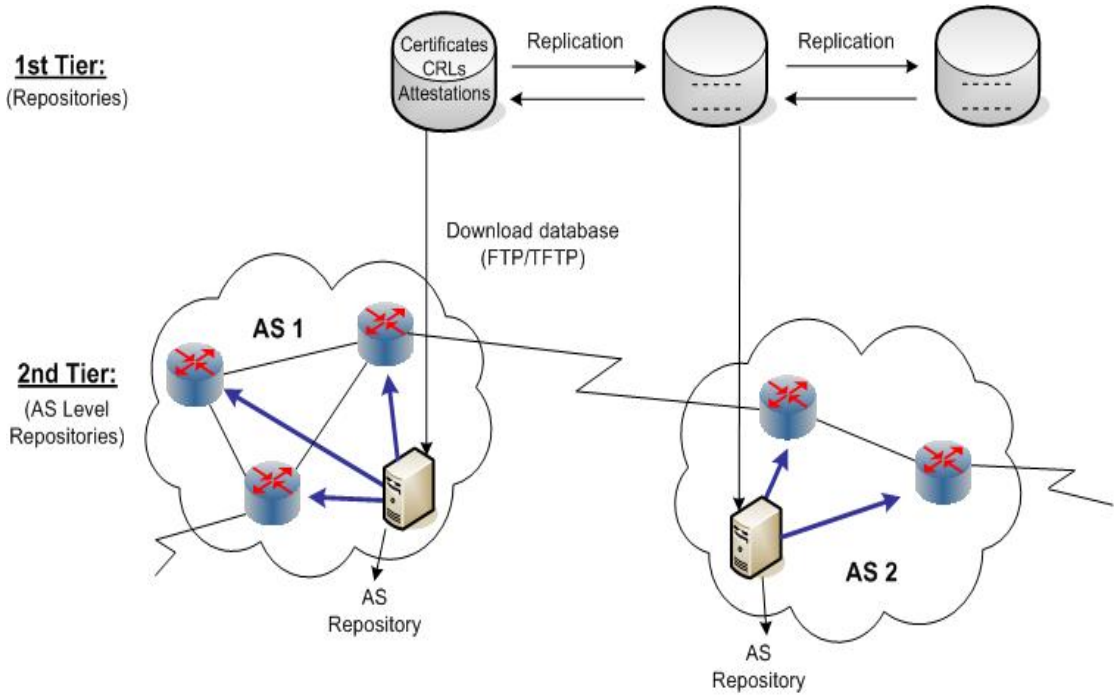

Figure 4.2: Distribution of S-BGP certificates, CRLs and attestations.

4.2.5 IPSec for BGP Sessions Protection

S-BGP proposes the use of IPSec Encapsulating Security Payload (ESP) with NULL encryption for authenticating BGP peers to each other and to prevent insertion, modification and replay attacks [46, 48]. The use of NULL Encryption algorithm is

specified in [65]. It essentially doesn't encrypt the data stream but is used because it is imperative that an ESP association must specify at least one cryptographic algorithm [46].

4.3 Secure Origin Border Gateway Protocol

Secure Origin BGP (SoBGP) is another security architecture. It is comparatively light-weight in comparison to S-BGP which makes use of PKI for authenticating and authorizing entities. SoBGP doesn't rely on any central authority and uses web of trust model for issuing certificates. SoBGP mainly protects against origin falsification attacks (by using authorization certificate called AuthCert), path subversion attacks (by using logical topology map driven by information in ASPolicyCert) and BGP session attacks (using IPSec ESP with NULL encryption). This section discusses main aspects of SoBGP.

4.3.1 Public Key Infrastructure

SoBGP makes use of four different types of certificates; Entity Certificate, Authentication Certificate, PrefixPolicyCert and ASPolicyCert.

4.3.1.1 Entity Certificate

SoBGP assumes there are trusted entities, like top level internet back-bone providers and key-authentication service providers. Their public keys act as root-keys which can be distributed on the internet out of band. These well-known entities issue certificates to AS which bind AS number with its public key. This kind of certificate is called EntityCert. An AS with EntityCert can issue other entity certificates to other AS to make a web of trust based on few well-known entities. Once an AS's public key is verified through EntityCert, it can generate other certificates using its private key corresponding to AS's public key mentioned in EntityCert.

4.3.1.2 Authorization Certificate

AuthCert certificates binds AS with the address block(s) it can originate. The concept is similar to address attestation in S-BGP but in SoBGP, an AS authorizes other AS to originate the address blocks by issuing a certificate called AuthCert. This binds AS number of authorized AS, the address block and AS number of authorizing AS. The authorized AS can further authorize other AS in similar way. When SoBGP speaker receives the route Update, it checks the origin AS authorization using the AuthCerts.

4.3.1.3 Prefix Policy Certificate

AuthCerts are not advertised independently. They are wrapped in another certificate called Prefix Policy Certificate (PrefixPolicyCert). PrefixPolicyCert is actually issued by the AS with authorization to originate a block of address (through AuthCert). It binds the IP address block in authorization certificate with the assigned policies for the block of address mentioned in AuthCert. It contains AuthCert and the set of policies which an originator may want to apply to the prefixes in that block of address. PrefixPolicyCert can define per-prefix policies and this certificate is used to securely distribute the prefix policy information.

4.3.2 Route Validation using ASPolicyCert

SoBGP introduces the concept of internet topology map to verify the advertiser of route has the path to the destination. In order to create the topology map of the paths of entire internet, SoBGP proposed that every AS builds a AS Policy Certificate (ASPolicyCert) which contains the list of its connected peers and the AS's local policy information. ASPolicyCert is distributed on the internet forming a topology map of entire Internet. Figure 4.3 illustrates topology map creation using ASPolicyCert.

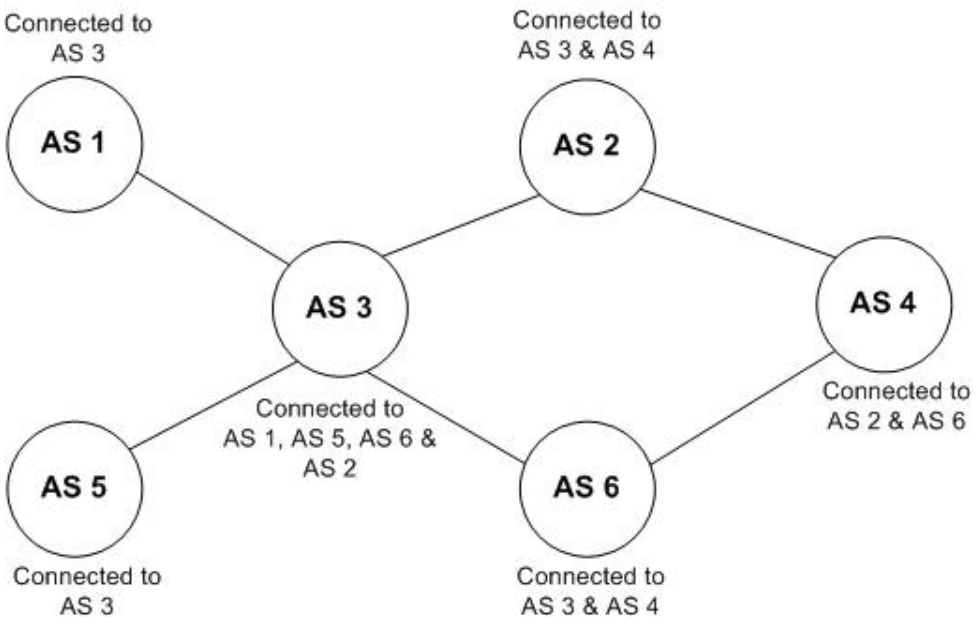

Figure 4.3: SoBGP build Virtual Topology Map using the information distributed through ASPolicyCert.

4.3.3 Distribution of Certificates

SoBGP suggested distributing all the certificates in the same way as routing information is distributed on the internet. A new BGP SECURITY message is added to distribute the EntityCert, PrefixPolicyCert and ASPolicyCert certificates. Existing BGP protocol needs to be modified to support this new BGP message.

4.3.4 Proposed Options to Reduce Processing Cost

SoBGP proposed different options to reduce the processing overhead for a SoBGP speaker. First option is to distribute the processing cost. The receiving border router processes the certificate and sends validated information to other routers in the same AS over an encrypted link and is stored by them. Using this approach, not every SoBGP router has to process all the certificates. Second proposed option is to use a server to process the certificates. When any border router receives a certificate, it sends the

certificate to AS's internal server which process and stores all certificates. BGP routers query the server about the validity of information in the update message when needed.

4.3.5 IPSec for Securing BGP Sessions

Like S-BGP, SoBGP also proposed the use of IPSec for securing sessions between two BGP peers as IPSec prevents session attacks.

4.4 Inter-domain Route Validation

The need for PKIs, high storage requirement, message overhead and high computational cost associated with S-BGP and SoBGP motivated the development of more efficient solutions to secure BGP. The Inter-domain Route Validation (IRV) is one such effort. It is an independent protocol which works as a companion to BGP protocol to validate route information [18, 66]. IRV mainly protects against origin falsification attacks (naively by directly querying AS whether it originated the route UPDATE) and path subversion attacks (by directly querying every AS on the AS_PATH).

IRV is primarily designed to protect against mis-configured AS, providing weak notion of security against sophisticated attacks. There are two main components in IRV architecture. The Inter-domain routing validator, called the IRV server, responds to the queries of other AS and the Network Management Element (NME) also called IRV client, queries IRV servers in other AS on behalf of its AS. Each AS maintains IRV server and NME. When new BGP update is received, NME of the AS try to verify the correctness of received information from every AS mentioned in AS_Path of received UPDATE message. NME does this by directly querying the corresponding AS's IRV server.

4.4.1 IRV Server

IRV Server is a dedicated server which stores data and responds to the queries generated by IRV clients in other AS on the basis of stored information. In order to get updated information of routes advertised by BGP speakers to other AS, border routers are configured in such a way that for each eBGP session, border routers establish a shadow eBGP connection with IRV. This connection is used to send the same data as the corresponding session with remote AS. Thus making IRV server act as a eBGP listener.

4.4.2 IRV Client

In order to validate the route information, an AS needs to form systematic queries upon reception of UPDATE messages. For this purpose IRV client (NME) listens to I-BGP messages from E-BGP listeners (border routers) and use this information to form queries on be-half of an AS to validate the received routes. NME determines the correctness of received routes by directly querying IRV server in other AS. NME receives the response to queries. The path is used only if all AS mentioned in the path validates the route. Figure 4.4 illustrates the working of IRV server and client.

4.4.3 Finding an IRV Server

It is proposed that single well known registry can store IRV server's IP address for each AS. This location information must be authenticated, for this purpose one or more public key certificate is used.

4.4.4 Prefix Origin Verification

IRV provides very weak notion of validating prefix origin. When an AS receives new UPDATE, NME queries the IRV server in the AS. This approach can only defend

against mis-configurations or unintentional mistakes. Any malicious prefix-hijacking can't be mitigated using this protocol

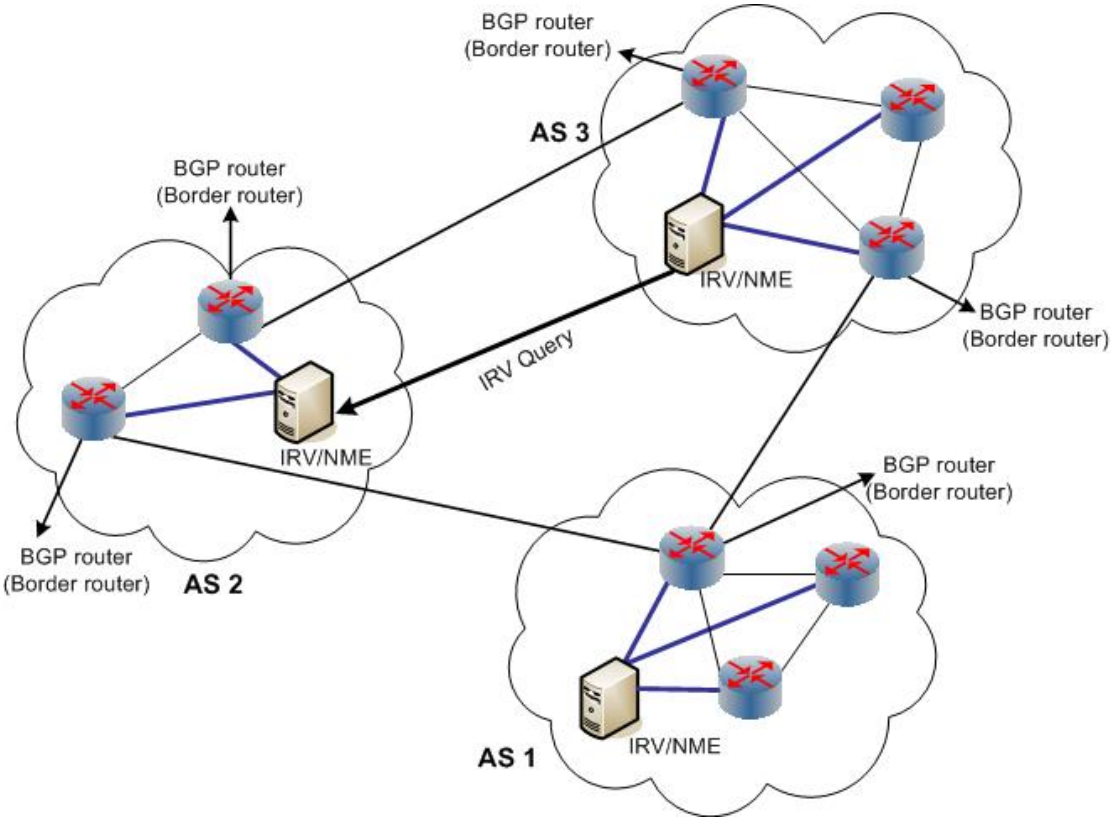

Figure 4.4: Working of IRV Server and Client.

4.4.5 Path Validation

IRV is quite effective in validating path information. All AS in the Path attribute of UPDATE message are queried one by one to check whether they advertised the route and to whom they sent the advertisement.

4.4.6 Other Issues

It is proposed that wherever possible, IPSec or Transport layer Security (TLS) should be used between IRV client and IRV server to ensure the integrity, authentication and timeliness of the queries and responses. Encryption is also facilitated using IPSec and

TLS [67]. The IRV server need not to reply all the queries uniformly, it can verify and authenticate the source of the queries and can respond to queries accordingly. The server may restrict access to sensitive data (e.g., Routing policy, internal routing info etc) so that only authorized sources can query it.

IRV proposed that AS can choose algorithm to define when and how often to check the validation of routes. The real cost of route validation depends on this algorithm. The more often routes are validated, the higher the computational cost and better the authenticity of routing information. This cost can be amortized by checking validation of routes at random interval. The cost can be further reduced by caching previously validated information. Furthermore, one single query can be sent to check the route validation for several routes if originating or coming through the same AS.

4.5 Pretty Secure Border Gateway Protocol

Pretty Secure BGP (psBGP) is a comprehensive security architecture based on the analysis of security and practicality of S-BGP and SoBGP. It combines their best feasible features [34, 35] and provides protection against origin falsification attacks (using Prefix Assertion Lists (PALs)), path subversion attacks (using S-BGP route attestation with bit-vector improvement) and BGP session attacks (IPSec ESP with NULL encryption and also supports MD5 signatures for TCP segments mechanism).

Like S-BGP, psBGP uses centralized PKI for validating AS numbers but makes use of decentralized trust model for verifying IP prefix ownership. There are few authorities for delegating AS and it is easier to manage the growth of AS numbers on the internet., making centralized PKI model inherently suitable. On the other hand, the number of IP prefixes in use is relatively high and it is difficult to trace back and maintain updated

record of IP prefix assignment through a centralized PKI. Thus, centralized PKI is theoretically attractive but practically infeasible [34, 35].

4.5.1 Centralized PKI for AS Number Allocation

In psBGP, four major Regional Internet Registries (RIRs) are the root authorities. Their public key certificates are distributed out of band to all AS. Any regional registry can allocate AS number by issuing public key certificates (X.509v3) which binds AS number with its public key. This certificate is called ASNumCert and is distributed on the internet. An organization with more than one AS must obtain a special certificate called MultiASCert from the RIR. MultiASCert binds list of AS to the name of the organization.

4.5.2 Decentralized PKI for Authentication

An AS with ASNumCert creates two data structures; SpeakerCert and Prefix Assertion List (PAL) and signs them with its private key. SpeakerCert binds AS number and a second different public key which is used to authenticate BGP speakers of this AS. Its corresponding private key is shared among all BGP speakers to sign BGP messages and to establish secure connection with external peers. If a BGP speaker is compromised, SpeakerCert is revoked without having any impact on ASNumCert.

Prefix Assertion List contains the prefixes owned by an AS itself and prefixes owned by neighboring AS. This facilitates neighboring ASes to share and verify information. The structure of PAL appears as follows; AS number and list of prefixes owned by it, the 1st neighbor AS number and list of IP prefixes owned by it, the 2nd neighbor AS and list of IP prefixes owned by it and so on. Every AS distributes the SpeakerCert and PAL on the internet.

4.5.3 Verification of Prefix Origin Information

PsBGP makes use of PALs distributed by all AS in order to verify the ownership of advertised prefix. In psBGP, prefix origin verification is based on the assumption that if one AS erroneously claims the ownership of the address space in its PAL, none of AS's neighbor endorse the same information in their PALs. In order to achieve this, psBGP insists that AS are responsible for carrying some due diligence offline in order to ensure that the neighbor AS's declared prefixes are actually owned by it.

In order to verify that the AS 'S' owns prefix 'f', psBGP compares the S's PAL records with any of S's neighbour PAL records. If S's PAL claims that S owns prefix 'f' and S's neighbour PAL endorses it, it is verified (S can originate 'f'). Otherwise, psBGP checks the PALs of the other neighbours of 'S', if any one of the S's neighbour PAL record endorse that S owns prefix 'f' then it is verified as proper, otherwise it is verified as improper and route is not accepted by psBGP speaker. Moreover, two assertions made by two different AS but owned by single organization (i.e., appearing in MultiASCert) is note accepted even if they are consistent with each other.

4.5.4 Path Validation

PsBGP suggests using the same method for path validation as proposed by S-BGP but with bit-vector improvement [68]. The Bit-vector improvement effectively amortizes the cost of signing the path when S-BGP speaker sends the same advertisement to multiple peers. BGP advertisement sent to the multiple peers is identical but in S-BGP, recipient's name is added as part of signed message. This leads S-BGP speaker to sign every update message separately for all peers even if the update message is identical. PsBGP suggests that same security can be achieved by using a simple bit-vector (or bit

string) instead of recipient's identity in Update message. Each bit in the vector represents a different peer. The S-BGP speaker can set the bit position of the intended recipients to 1 and sign the message including this bit-vector. In this way, just a single message can be used to send all peers.

4.5.5 Protection for BGP Sessions

Like previous solutions, psBGP also propose to use IPSec Encapsulating Security Payload protocol with NULL encryption. psBGP also supports protection of BGP session via TCP MD5 described by [20] but suggests that dynamic session keys are generated based on public keys of psBGP speakers to improve security.

4.5.6 Distribution of Certificates and PALs

In psBGP, SpeakerCert, ASNumCert and others certificates are distributed with BGP Update message but the method is not specified. Similarly distribution of PALs is not explained but it is suggested that any efficient method for distribution of certificates and PALs can be used [34, 35].

4.6 BGPSec

BGPSec is a 'work in Progress' effort by Secure Inter-domain Routing Working Group (SIDR WG) of Internet Engineering Task Force (IETF) to enhance the security of BGP protocol by enabling full BGP path validation using cryptographic principles. According to latest internet draft on BGPSec design choices, there are four key high-level design goals for BGPSec. First is rigorous path validation for all announced prefixes; not merely showing that a path is not impossible. Second goal is incremental deployment capability; no flag-day requirement for global deployment. Third is protection of AS

paths only in inter-domain routing (eBGP); not applicable to iBGP (or to IGPs). Fourth aim for no increase in provider's data exposure (e.g., require no disclosure of peering relations, etc). Two main vulnerabilities of BGP are addressed by SIDR; to check if an AS is authorized to originate a prefix and to validate if the AS-Path represented in the route is the same as the path through which NRLI travelled.

Therefore, the BGPSec focuses on providing these countermeasures against the above mentioned vulnerabilities; cryptographic validation of Prefix Origin Authorization using RPKI and ROA, cryptographic validation of AS-Path Integrity using extension of RPKI and introduction of new optional (non-transitive) attribute called BGPSec_Path_Signatures.

4.6.1 Resource Public Key Infrastructure

The SIDR WG proposed an infrastructure for secure inter-domain routing [http://tools.ietf.org/html/rfc6480]. The foundation of the proposed standard is 'Resource Public Key Infrastructure (RPKI); PKI for allocation of IP Address space and AS numbers to resources, and the distributed repository for storing and distributing PKI certificates and other signed objects necessary of improved security of internet routing.

The purpose of this infrastructure is to have cryptographic mechanism to enable a legitimate owner of IP prefix to verifiably authorize one or more ASs to originate that prefix through BGP. Such kind of verifiable authorizations are used by BGP peers to filter routes. BGPSec uses RPKI for validation of prefix origin authorization.

4.6.2 Architecture of RPKI

There are three main components of RPKI Architecture proposed by SIDR WG. First is X.509 based PKI for allocation of IP Address Space and AS to entities to assert

attestation of holding of IP address space and AS numbers. Second is Non-Certificates signed Objects (e.g., Route Origination Authorizations (ROAs and manifests). Third is distributed repository for storing and making available all signed objects including X.509 certificates to ISPs for routing decisions.

4.6.3 PKI for Allocation of Internet Number Resources to Entities

The proposed X.509 based RPKI makes use of existing hierarchical structure for certificate issuance and mirrors the way in which the internet number resources are already distributed. The existing structure for IP address space and AS number allocation is described in previous sections. The certificate issued through this PKI is called Resource Certificate. Any internet number resource holder authorized to sub-allocate these resources can issue the resource certificates corresponding to sub-allocation process. Therefore, all resource holders authorized to sub-allocate resources have their CA certificates for issuance of certificates for sub-allocations. CA certificate also enables resource holder to issue Route Origin Authorizations (ROAs). In some cases, CA certificate will be issued to resource holder who is not authorized to sub-allocate resource because it would require for issuing ROA for its provided-independent allocation e.g., multi-homed customer with provider independent allocation.

RPKI uses X.509 certificates with addition of extension for IP address space and AS number. The resource certificates holders have verifiable proof of holder ship (i.e., resource certificate) as each certificate attests allocation of resources to resource holder. However, these certificates do not need to contain descriptive subject name field as these certificates are to be issued just for the proof of holder ship and not for authentication or identification. Moreover, resource holders capable of sub-allocating resources and have

multiple resources from multiple providers will have multiple CA Certificates in order to sub-allocate resources allocated by different higher level resource allocators.

In addition to CA Certificate, RPKI makes use of another certificate named “End-Entity” Certificate. This end-entity certificate will be signed using the corresponding private key of the resource holder and will be used to sign ‘Route Origination Authorization (ROA)’ and manifests. One end-entity certificate is used to sign one ROA / manifest only. There is one-to-one relationship between ROA and end-entity certificate. Therefore, it is also not necessary for resource holder to keep private key corresponding to public key mentioned in end-entity certificate once ROA or manifest is signed.

4.6.4 Route Origination Authorization

The above PKI is not sufficient for routing decisions like whether AS is authorize to advertise routes for certain prefixes allocated to any resource holder. Therefore, the signed object named ROA has been introduced. ROA is an attestation that IP address space holder has explicitly authorized any particular AS to originate routes its prefixes. One ROA will be used to authorize one AS only. If the resource holder of IP address space needs to authorize another AS to originate routes against particular prefixes then it has to issue another ROA for that matter. Authorizing more than one AS to originate could be requirement in case of multi-homed (provided-independent) customer.

When holder of IP address space needs to generate ROA to authorize any particular AS to originate its prefixes, it creates new public/private key pair for signing ROA, signs the ROA using private key, creates an end-entity certificate by binding the corresponding public key and authorized IP prefixes using private key corresponding to its CA Certificate. The ROA is distributed to relying parties using distributed registries and

relying parties (i.e., ISPs) could validate the authority of AS that originated the route against particular prefix by the holder of that IP address space.

The detailed syntax of ROA is mentioned in [http://tools.ietf.org/html/rfc6482]. The high level contents of ROA are AS Number, the list of IP Prefixes and optionally maximum length of more specific prefixes for each IP prefix. It is easy to revoke ROA as it can be revoked by revoking the corresponding end-entity certificate. The end-entity certificate is usually embedded in the ROA.

4.6.5 Distributed Repositories

To create a list of table of all prefixes authorized by prefix holder to AS to originate routes against, ISPs need to acquire and validate all ROAs. To validate all ROAs, the ISPs need to have access to all certificates involved in the chain and Certificate Revocation Lists (CRLs). For this purpose, a repository system is proposed. The primary function of this repository system is to store and maintain databases of all signed objects (including CA certificates, CRLs, ROAs and manifests) so that this information can be pulled by replying parties at whatever frequency they deem appropriate.

4.6.6 Structure of Repository System

Although there is a single repository system comprising of multiple databases but the structure of the repository system is distributed. The databases of repository system will be distributed among registries (i.e., RIRs, NIRs, LIRs/ISPs) and each repository will maintain all signed objects including CA and EE certificates, CRLs, manifest signed by CA(s) associated with that registry. Moreover, each repository will maintain all signed objects of its customers and customers of its customers. Therefore, ideally the RIRs will contains all PKI based signed objects from all entities within its geopolitical scope.

The repository contains file system directory against each certificate. This directory is the authoritative publication point for all signed objects verifiable through that certificate including certificates, CRLs, ROAs and manifests. RPKI makes use of Subject Information Access (SIA) and Authoritative Information Access (AIA) extension defined in [69, 70]. The certificate's SIA points to the certificate publication directory containing all verifiable signed object through that certificates. The AIA extension of the certificate contains the URI authoritative location of the parents certificate (a certificate whose corresponding private was used to sign this certificate). This structure is illustrated in figure 4.5.

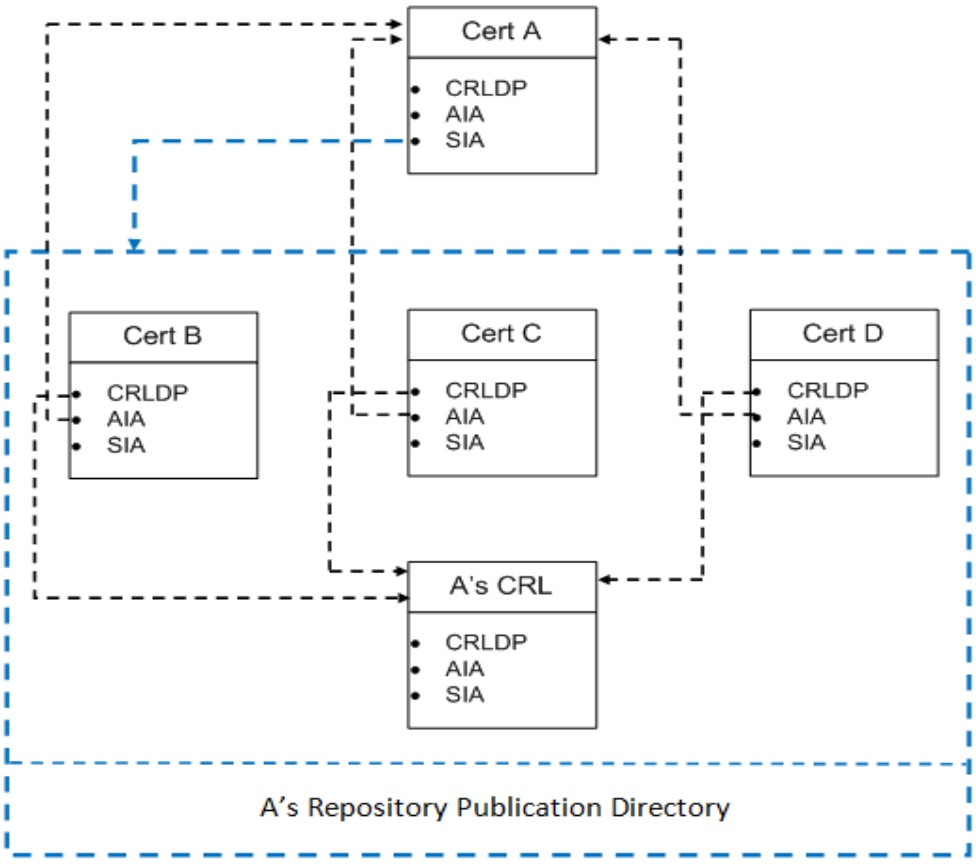

Figure 4.5: Use of file systems CERT's public directory in RPKI

In order to prevent unauthorized manipulation and deletion of signed objects in the repositories, RPKI pushes repositories to use an access control mechanism to allow only

authorized entities to make modifications to objects (upload, change and delete) that comes under their authoritative scope. However, no access controls mechanism is enforced. For downloading of contents from repositories, RPKI enforces to support rsync [71]. Although, other download protocol could also be used.

In order to prevent the threat of deletion or replacement of older version of signed object in the repository, RPKI introduces the concept of 'Manifest'. Manifest is the signed object that lists all the files and their hashes. It is signed by private key corresponding to EE certificate. Like ROA, there is one-to-one relationship between EE certificate and Manifest. One EE certificate will be issued for exactly one manifest.

4.6.7 Cryptographic Validation of AS-Path Integrity

ROAs introduced in RPKI just provide protection against unauthorized origination of IP prefixes. However, any sophisticated attacker could do the route hijacking attack by appending authorized origin AS to otherwise illegitimate AS-path. BGPSec therefore, extends RPKI and introduces another certificate referred to as BGPSec router certificate. This certificate binds an AS number to a public signature verification key. The corresponding private key is held by one or more border BGP speakers within that AS.

Moreover, BGPSec introduces an optional (non-transitive) attribute named as "BGPSec_Path_Signatures". The BGP border speakers in AS using the private key corresponding to BGPSec router certificate will sign the AS Path attribute before advertising the route to peer and append the signature in BGPSec_Path_Signatures attribute. This will be done by every BGPSec enabled router before advertising the route to its external peers. The relying parties (i.e., peer BGP routers) can then verify that the AS-Path in BGP Update messages is in fact the path the route traversed through. The concept is similar to that introduced in S-BGP. The BGPSec_Path_Signatures attribute

will contains the series of signatures as the route traversed through from one BGP Speaker to other, one for each AS in the AS-Path attribute of BGPSec update message.

4.6.8 Signing and Signature of BGPSec AS-Path

Whenever BGPSec speaker originates a route, it signs the NLRI, AS number of originating AS, AS number of peer AS to whom the update is being sent and some other fields necessary for security guarantee using the private key corresponding to BGPSec router certificate and appends the signature in BGPSec_Path_Signatures. The NRLI in BGPSec is restricted to contain only single prefix. When the neighboring BGPSec speaker receives the route and wishes to advertise the route to its external peers, it will then add another signature to the BGPSec_Path_Attribute using private key corresponding to its BGPSec router certificate. This signature will be calculated based everything protected by previous signature plus AS number of route advertiser and AS number of the external peer to whom the route is advertised.

Before advertising the route to its peer, each BGPSec speaker will also reference to its BGPSec router certificate, known as Subject Key Identifier (SKI). The SKI against each signature will be used by relying parties to identify and acquire the certificate used for verification of each of the signature in BGPSec_Path_Signatures attribute. Figure 4.6 explains how BGPSec speaker in AS 1 originates and advertises the prefix 115.186.0.1/24 to its peer BGPSec speaker in AS 2, the AS2 advertises that route to AS3 that is connected to AS2 and AS3 advertises the same route to AS4. Thus any BGPSec speaker in the AS_Path in BGPSec update message is able to validate the integrity of path that the route traversed through.

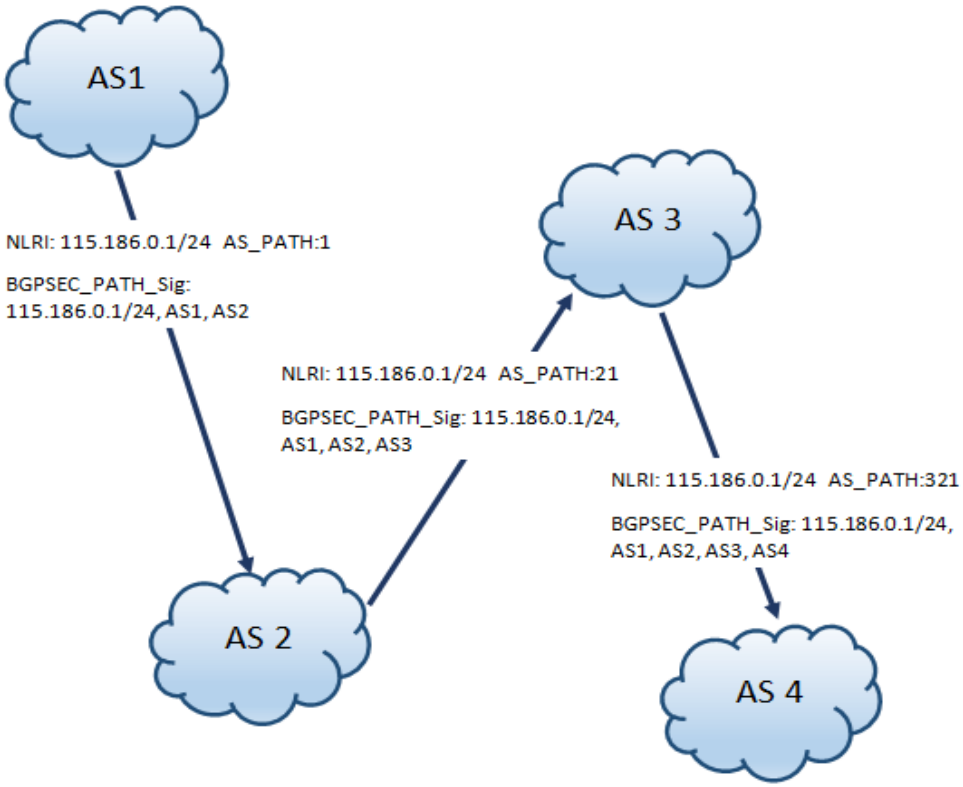

Figure 4.6: BGP_Sec_Path_Signature when route advertised from AS1 to AS2, AS2 to AS3, AS3 to AS4.

4.7 Hop Integrity Protocols

The IETF has also chartered RPSec Working Group to establish the security requirements for routing protocols. Many solutions have been proposed which address different aspects of BGP security, including Hop Integrity Protocols. Hop integrity ensures the authenticity, integrity and freshness of the messages sent between any two adjacent routers. Gouda et al. proposed a protocol suite which can be added in routers to provide hop integrity at network (IP) layer [72]. All the proposed protocols are stateless, require small overhead and based on public keys. This Protocol suite defines three protocols; Secret Exchange Protocol, Weak Integrity Protocol and Strong Integrity Protocol. The Secret Exchange protocol is used to dynamically generate and exchange

new secrets between two routers. These secrets are utilized by weak integrity and strong integrity protocols to compute and verify integrity of the messages exchanged between two routers. The secret exchange protocol assumes that both the participating routers have valid public keys of each other and they share some secret before time of using this protocol (initially secrets can be exchange out of band).

To understand working of these protocols, let's assume that two routers A and B want to exchange secrets periodically using secret exchange protocol. According to the protocol, Router A generates a new secret and encrypt, encrypts the old (which is already shared among routes A and B) and new secret with the public key of router B and send the message to router B. Router B upon receiving this message, decrypt the message using its private key and compare its currently used secret with the old secret mentioned in the received message. If they both match, router B installs the new secret and encrypt the new secret with router's A public key and send it to router A. Router A receives and decrypts the message using its private key and compares the received key with the new secret. If both match, Router A gets acknowledgement that the router B has received and installed the new secret. If Router A doesn't receive the encrypted new secret from router B, router A sends the same message to router B again after waiting for random time.

The Weak Integrity protocol only provides the integrity of the messages (i.e., defends only against modification attacks) between two routers. Assuming both routers A and B shared a fresh secret using Secret key exchange. The weak integrity protocol specifies that each data packet 't' and shared secret will be put into the Message digest (MD) function which will generate the fixed digest 'd' based on the input. Then the message (t, d) will be sent to other router. When the other router receives this message, it

will compute the Message digest (MD) on the data packet ‘t’ and compare the output with the‘d’. If they both match, the packet will be accepted otherwise it is rejected.

The strong Integrity protocol provides integrity as well the freshness of the messages (i.e., defend against modification and replay attacks) between two routers. It’s same like weak integrity protocol but makes use of sequence numbers of order to provide freshness to the messages. The sequence number mechanism used in this protocol is also fault-tolerant in case two routers lose synchronization.

This protocol suite describes very lightweight approach, which can be used to protect the BGP sessions between two peers. It can be used as an alternate to IPSec.

4.8 Invalid MOAS conflicts Detection

Hawkinson and Bates [4] recommend that a prefix should generally originate from only one single AS. Multi Origin AS conflict occurs when a prefix is appearing to be originated from more than one ASes simultaneously. Although at first, MOAS conflicts appear to be the result of an intentional attack or misconfiguration but MOAS conflicts do occur because of the operational needs, for example to support different types of multi-homing. The analysis of MOAS conflict [73] showed that there can be many different causes for MOAS conflicts. Exchange point prefixes (the IP addresses configured on the link that connects two ASes) may cause MOAS conflict because each exchange point AS may advertise these prefixes as it can directly reach these prefixes (this doesn’t create a problem for routing). Multi-homing when used without BGP or with using private AS number may cause MOAS conflicts when Multi-home network transition from one provider to other provider. Faulty or malicious configuration is also the major reason for such MOAS conflicts. Further, they described that prefix aggregation and any-cast addresses can also be the cause of MOAS conflicts.

MOAS conflicts can be categorized as Valid MOAS and Invalid MOAS. MOAS conflict is valid if the originating AS can directly reach to the destination and MOAS is invalid when the originating AS cannot reach to the destination prefix [18, 28, 73, 74]. It is very difficult to differentiate between valid and invalid MOAS conflicts in the current BGP routing environment. The authors proposed a solution to detect and mitigate invalid MOAS by creating MOAS list [73]. This list will contain all the ASes who are entitled to originate any prefix ‘p’. This list then can be attached with route advertisements by all those originating ASes. BGP community attribute can be used to carry MOAS list. When any BGP speaker receives the conflicting advertisements for prefix ‘p’, it can check to see that whether MOAS list attached with conflicting routes is consistent with each other (i.e., the same ASes are listed in all MOAS lists). For BGP speaker to verify MOAS list, it is needed to change current BGP implementation so authors proposed that an off-line monitoring system can be used to verify MOAS conflicts, which can periodically check routing table for invalid MOAS conflicts. This approach although can check against mis-configurations but it is limited in a sense that BGP community is an optional transitive attribute which can dropped by transit BGP speaker or MOAS list can be modified in transit, which can either deny service to legitimate AS or router can be forced to accept false routes.

Other proposals to detect MOAS conflicts include the proposed solution presented in [75]. The authors proposed that a DNS server can be used to store AS and prefix pairs and BGP speaker can check the prefix’s origin AS in the DNS and compare it with the received update information. The roué must only be accepted if they both match. The similar kind of approach is used these days by ISPs who implement filters which automatically be updated by checking the records in Routing registries.

There is another model proposed by Kruegel et al. [76] which detects the MOAS conflicts. This model uses the approach similar to intrusion detection in computer networks and although it doesn't differentiate between valid or invalid MOAS but avoids some kind of MOAS conflicts which can occur due to some valid reasons. Their proposed model stores the mapping between IP address blocks and ASes (who can originate these address blocks) in order to detect the address violations. This model has two phases. Model building phase and the Detection phase. During the model building phase, data is collected over the period of weeks to build the mapping between IP address block and ASes who are originating this address space. This information is extracted from BGP update messages. (i.e., IP Prefix and the last AS number in AS_Path attribute). Mapping phase assumes that network was running normally and there weren't any serious attacks or faults condition present on the Internet when the data was collected to create a mapping table. While model building phase, all aggregated routes from core AS are ignored. Once the IP address block and AS mapping is done the system can detect if MOAS conflict occurs. A major issue for consideration is that if the prefix ownership is changed, the system would generate false alarm. Authors argued that the prefix ownership is usually stable but if the address ownership is changed it can be solved by having an automated system to detect whether the previous owner still advertising the address block or not, if not then the bindings will be changed automatically. This kind of system however needs to collect data from different vantage points on the Internet in order to truly map IP prefixes with ASes and detect false announcements [76, 77].

4.9 Path Authentication

Hu et al. proposed a symmetric key cryptography based path authentication solution for path-vector protocols [78]. They presented the cumulative authentication mechanism which assumes that each node in the path shares a private key (symmetric key) with the authenticating node. Each packet maintains path authenticator and the Address list. Address list can be thought of as AS_Path attribute. Authenticator value has some well known initial value (e.g., 0). When the packet traverses a node, node append its address in the address list and compute the MAC (using the shared key between itself and the authenticating node) on the all immutable fields of the packet and the authenticator value and overwrite the authenticator value with the computed MAC. The node then sends this data to next node in the path. The next node receives the packet which also contains the address list (AS_Path in case of BGP) and the authenticator value. It appends its address in the address list and computes the MAC (using shared secret between itself and authenticating node) on all immutable fields of the packet and the authenticator value received from previous node. This node then replaces the authenticator value with computed MAC and sends packet to the next node. All the nodes in the path repeat the same process until the packet reaches the authenticator node. The authenticator node then computes the recursive MAC on the packet using information from the address list. If the computed value matches with the received authenticator value, the authenticating node is ensured that the packet actually traversed through that path.

This cumulative mechanism has its limitations that this doesn't protect the second node from removing the first node. First of all , Any node in the path can remove all previous node from address list and initialize the authenticator value to zero again and then can send the message forward to authenticating node pretending that the message

actually originated by it. Secondly any node in the path can remove any node only from the address list. This way, authenticator will not be able to computer recursive MAC properly. Thirdly this mechanism is not suitable for BGP as path authentication is not performed at each node. The main problem with this mechanism is the large number of AS on the internet. It is not feasible that each AS shares a key with every other AS, i.e., number of keys equivalent to the total number of AS on the internet.

Authors suggested that instead of using shared private keys, this cumulative authentication mechanism can be used along with TESLA broadcast authentication protocol [79] to perform authentication at each recipient. TESLA protocol is specifically designed to provide the origin authentication based on symmetric key in broadcast applications. It provides the public-key semantics using symmetric key cryptography. The main idea of TESLA is that the sender and receivers are loosely time synchronized before time. The sender uses a private key known only to him to generate a message MAC on the message. The Sender then broadcast the message to receivers and later releases the key (broadcast the key) after some specified time. Receiver when receive the message, check the time to make sure that the key for this particular message is not released yet (as there is time already specified by sender that when it will release the key for this particular message). If the key is not released then the receiver buffers the message and waits for the key. After receiving key it can compute the MAC and make sure that the message came from authentic source. The whole point is that the receiver receives the message before the key is released. If the receiver gets the message after the specified time of key release then receiver has no guarantee of authenticity of the message as anyone then can compute the MAC on the message. As the key is released

after sending the message so the TESLA requires one key for one message. The keys are generated using One-way Hash key chain mechanism.

By using TESLA authentication protocol, proposed cumulative authentication mechanism will be little bit modified. Each node when receives the packet from previous node, it will append its address in address list, compute the MAC on the Packet and the Authenticator value using its current TESLA key, buffers the packet for verification. Later it gets previous node's TESLA key it verifies the authenticator value sent by previous node using that node's TESLA key. If Packet is verified then node sends the packet forward. Every node in the path does the same processing. This approach if implemented for BGP path validation would require all BGP routers to be loosely synchronized which is another problem in itself.

4.10 Secure Path Vector

Secure Path Vector solution (SPV) was proposed to secure AS path, which is based on purely symmetric cryptographic primitives [80]. SPV efficiently removes the need for routers to bear computational load of public-key cryptography. It protects the path from the origin so also defends against prefix-hijacking attacks. Although signature generation and verification process is purely based on symmetric cryptography but SPV makes use of two public-key certificates in order to validate the public keys. The authors proposed similar kind of PKI as was proposed by S-BGP for address attestations, in order to ensure that the prefix belongs to the originating AS. The PKI proposed by SPV is discussed below.

4.10.1 PKI for Prefix Allocation

Same like S-BGP the root authority is ICANN, which delegates address space to regional registries (RRs), RR then further delegates address space to service providers and service providers further delegate the address space to their customers or other organizations. But unlike S-BGP, at each stage when higher authority in address delegation hierarchy assigns address space to lower organization, it issues a certificate which binds address space or prefix with the prefix-public key. The addresses are of course delegated to organizations but the certificate binds the public key to prefixes. The corresponding private key is named prefix-private key. Now any AS can originate the prefix if and only if it has got the knowledge of prefix-private key. This prefix-private key will be associated with the block (UPDATE message) which is used to originate the corresponding prefix. The prefix public key certificate is then distributed on the Internet.

SPV uses the combination of one-way hash chains, hash trees and one-time signatures to provide path authenticity using symmetric cryptography. To prevent the replay of old BGP UPDATE message SPV assumes that the time is divided into fixed slots called epoch and BGP UPDATE messages are valid for this epoch. They need to be re-advertised after the epoch is expired [81].

4.10.2 ASPATH Protector

SPV proposes a cryptographic mechanism called ASPATH protector. The ASPATH protector in every update message makes sure that the attacker cannot shorten the path and the AS path cannot be modified.

4.10.2.1 Origin AS Computation

The prefix originator AS first constructs the ASPATH protector offline. The maximum number of expected ASes that may be in the AS path (at any place) are estimated first. The same number of private keys is generated by randomly choosing one key and generating one way hash chain on it. This kind of key is called single ASN private key. Each of these keys is used to generate one-time signature at each AS. The originating AS will compute all the values before sending the update message. Each one-time signature is generated by expanding ASN private key using peudo-random number Function (PRF), hashing all the generated values and constructing the hash tree over them. The root value of this tree is called single ASN public key which serves as a public key for a single signature generated with corresponding single private key. Then the hash tree is computed over all single public keys. The root node of this tree is called epoch public key. This key is used to verify all one-time signatures within this epoch. But epoch expires so there is need of one key which must be used to authenticate all epochs so AS construct a hash tree over all epoch public keys. The root value is called multi-epoch public key. All the ASPATH protectors for all epoch can be verified with this key (for a specified prefix). An AS then issues the certificate which bind multi-epoch public key and prefix and sign this certificate with prefix-private key. This certificate is then distributed on the internet using some efficient mechanism.

4.10.2.2 ASPATH Protector Use and Verification

SPV proposes to use optional transitive attribute of BGP to carry signatures and the key. This attribute can be named ASPATH Protector. Just to illustrates, consider, A→ B → C is the path which update message will traverse. The originating AS ‘A’ add its AS

number and the number of the AS to which it is sending the update (i.e. A, B)in the AS_Path attribute of the BGP message and generates the one time signature on AS_Path (which includes the AS number of next AS in the path) || value of epoch, using first key in the hash chain (i.e, single ASN private key). Then AS A applies hash function to its single ASN private key and place the hashed key (which will be used as single ASN public key for next AS) and the signature in ASPATH protector. A then forwards the update to B. An AS B when receives the Update with ASPATH protector it must be able to verity all the one time signatures on the AS_Path using multi-epoch public key. SPV assume that every AS has a multi-epoch public key for the prefix which is needed to be verified. So first of all B verifies the A's one-time signature on the AS_path. If there are more than one signatures on the path, B can verify all of them because it knows the root of all signatures. The signatures in ASPATH protector provide enough information to construct the single ASN public key for their corresponding single ASN private key. An AS B can compute all the following one time signatures by repeatedly hashing its single ASN private key and each key then can used to construct single ASN public key for that single private public key. The hash tree on these single ASN public keys can provide epoch keys which can then be used to get multi-epoch public key value. This computed value will be compared with the Multi-epoch public key. If the computed value matches with multi-epoch public key then all the signatures are verified (that's the cryptographic property of hash tree). Precisely in this scenario, B will verify A's one-time signature on the AS path (by comparing calculated root value with Multi-epoch public key) and it will then compute a one time signature on AS_Path including next AS number C (i.e., A, B, C) and epoch value, using the key sent by A. B then takes the hash of its single private key, put the signature and the hashed key in ASPATH protector field and forward the

Update to the next AS. All the nodes on the path will use the same process and they can verify signatures of all previous nodes.

Using simple hashing function in one time signature makes this process much faster then public key cryptography based functions. Hash tree based one time signatures provide the semantics of public-key cryptography while providing the efficiency of symmetric cryptography. Although verification process involves more computations but authors experiment and discovered that its 20 fold faster than digital signatures. Moreover SPV reduces the no. of certificates each BGP router needs to store.

4.11 Summary

This chapter presents detailed analysis of security extensions and other solutions proposed to secure BGP routing protocol against numerous threats. These include S-BGP, SoBGP, IRV, psBGP and BGPSec security architecture extensions. Some other solutions for securing BGP are also presented.

An Evaluation Model for BGP

5.1 Introduction

The inter domain routing infrastructure for the internet is supported by BGP protocol. BGP provides and processes routing information to enable network traffic to reach its destination. The threat model for BGP, based on earlier and recent attacks and exploitations of the protocol are presented in Chapter 3. Analysis of the security extensions for BGP has also been given in detail in Chapter 4. In this chapter, an evaluation model for BGP, based on the practical assessment of BGP attacks is proposed. The evaluation model is established in two parts; security evaluation model and performance evaluation model. It encompasses all necessary security requirements to prevent attacks on BGP and performance requirements for optimum utilization of resources in BGP deployment. A balance is achieved between security and performance features in the proposed evaluation model.

5.2 Practical Assessment of BGP Attacks

Most of the major historical service disruptions resulted through BGP misbehavior was due to mis-configured routers. Despite the open design of BGP that makes it vulnerable to attacks as explained in Chapter 3, any major specific attack on BGP is not observed. This is due to inherent difficulties in launching attacks on BGP, which are discussed in this section.

5.2.1 Spoofing

To launch a spoofing attack, an attacker sends masqueraded packets to the target, pretending that it is coming from a legitimate peer. BGP spoofing attacks are of two types; TCP based or spoofed BGP messages. Launching a spoofed attack against any BGP peer is complicated. An attacker has to obtain IP addresses of the peers and determine the exact port number. Any one of the two BGP peers initiates TCP connection and selects a random source port and sends TCP SYN message to the neighbor with destination TCP port 179. It is not easy for an attacker to determine which of the peer is using port 179 and which one is using randomly selected port (and initiated the connection). So the attacker has to either capture packets or guess which peer is using TCP port no. 179.

TCP protocol makes use of TCP sequence numbers in order to assemble the packets correctly at receiver end and for end-to-end error detection and correction. The TCP sequence numbers specified in the TCP segments must be within or above the TCP sliding window otherwise the receiver simply drops the packets. The mature TCP implementation uses the random sequence numbers and it is not easy for any outsider to guess the acceptance value for TCP sliding window. Since BGP peers are mostly directly connected with each other, it is difficult to determine the port number and acceptable TCP sequence number in this scenario. The use of MD5 signature option for TCP along with GTSM makes remote attacks (insertion, deletion, modification and session hijacking etc) even more difficult to materialize.

5.2.2 Prefix Hijacking

Most of the ISPs implement route filters and accepts only the prefixes declared by customers. ISPs make sure that the customer declared prefixes are actually owned by them. ISPs use Routing Registry to update their route filters automatically. In case customer gets new prefix, it registers it with the relevant AS in the RIR and updates the records. Routing registries follow procedures to make sure that the records updated are actually proper. Prefix hijacking in this environment is much more difficult to execute.

5.2.3 Countermeasures at ISPs

ISPs are usually well aware of security threats and have resources to implement best security practices in their network. This makes it difficult for an attacker to compromise an ISP's BGP router. On the other hand, multi-homed customers and stub customers are more vulnerable to attacks because of lack of expertise in securing their infrastructure. This is also why the main focus of current practices to secure BGP in on the ISP and customer peering. As a general practice, ISPs prevent malicious route advertisement from their customers. This is one of the contributing factors in avoiding major BGP attacks.

Additionally, routers are special purpose devices and don't run any extra services that can be exploited. Historically, there have been much less security vulnerabilities discovered in routing software in comparison to other operating systems and web servers. This is also the reason which makes it difficult to compromise a BGP router.

5.3 Evaluation Model for BGP

Security and performance requirements defined so far for BGP are incoherent and not comprehensive. The balance between security provisioning and deployment requirements is also not maintained. The difficulties in BGP deployment are increased

manifolds due to the lack of such a comprehensive evaluation model. In this section, we define an evaluation model for BGP. The proposed model is inclusive of all security and performance requirements for secure and efficient deployment of BGP. Equilibrium is achieved between the needs for security provisioning and performance optimization, for operational feasibility of the protocol. This makes the presented evaluation model suitable for meeting practical needs.

5.3.1 Security Evaluation Model

The security evaluation criteria are defined for BGP based on the following security requirements, necessary for protection of BGP against attacks.

5.3.1.1 Protection against BGP Session Attacks

Provide assurance to BGP speaker that the routing updates it received from neighboring AS were in fact came from authorized BGP speaker of that AS and not from some spoofed source. Also, provides assurance to BGP speaker that the BGP updates it receives is authentic and nothing has been added, modified and deleted (from AS_Path) from it by any malicious entity while in transit from origin to the current AS.

5.3.1.2 Protection against Origin Falsification Attacks

Provide assurance against each BGP update message that the Autonomous System (AS) which originated the route for any particular prefix was explicitly authorized by the holder of that prefix.

5.3.1.3 Protection against Path Subversion Attacks

Provides assurance to every BGP speaker which receives the BGP update message that the autonomous system numbers mentioned in the AS_Path are the actual ASes through which this update traversed through.

5.3.2 Deployment Evaluation Model

The performance requirements necessary for operational deployment of BGP are listed in the performance evaluation model defined below.

5.3.2.1 Scalability

The proposed BGP secure extension should be able to accommodate existing BGP routes on the internet and also be scalable enough to accommodate future internet growth resultantly growth in BGP routing table and increasing rate of BGP updates.

5.3.2.2 Convergence Speed

The convergence time or speed is the measure of how fast the set of BGP routers get in the state to have same topological view of the routing information of prefixes. The goal of any proposed BGP solution is to at least meet the existing convergence time / speed of the BGP.

5.3.2.3 Backward Compatibility

One of the most important factors to consider while proposing any new security extension to BGP is to have its backward compatibility with existing BGP protocol. Because it is not possible to transition from already widely deployed BGP protocol in one go. The incremental deployment is the only workable solution.

5.3.2.4 Computational Overhead

The computational overhead on the BGP speaking router must be considered when evaluating or proposing any cryptographic secure extension to BGP. It is not easy and feasible to upgrade processing power on all existing BGP speakers. Any solution that proposes to change the routing platform on all BGP speakers will not be categorized as practical or workable solution.

5.3.2.5 Memory / Storage Overhead

The storage or memory overhead of any proposed secure solution need to be considered while evaluating any secure extension to BGP because existing BGP speaker may not be able to handle this memory or storage overhead.

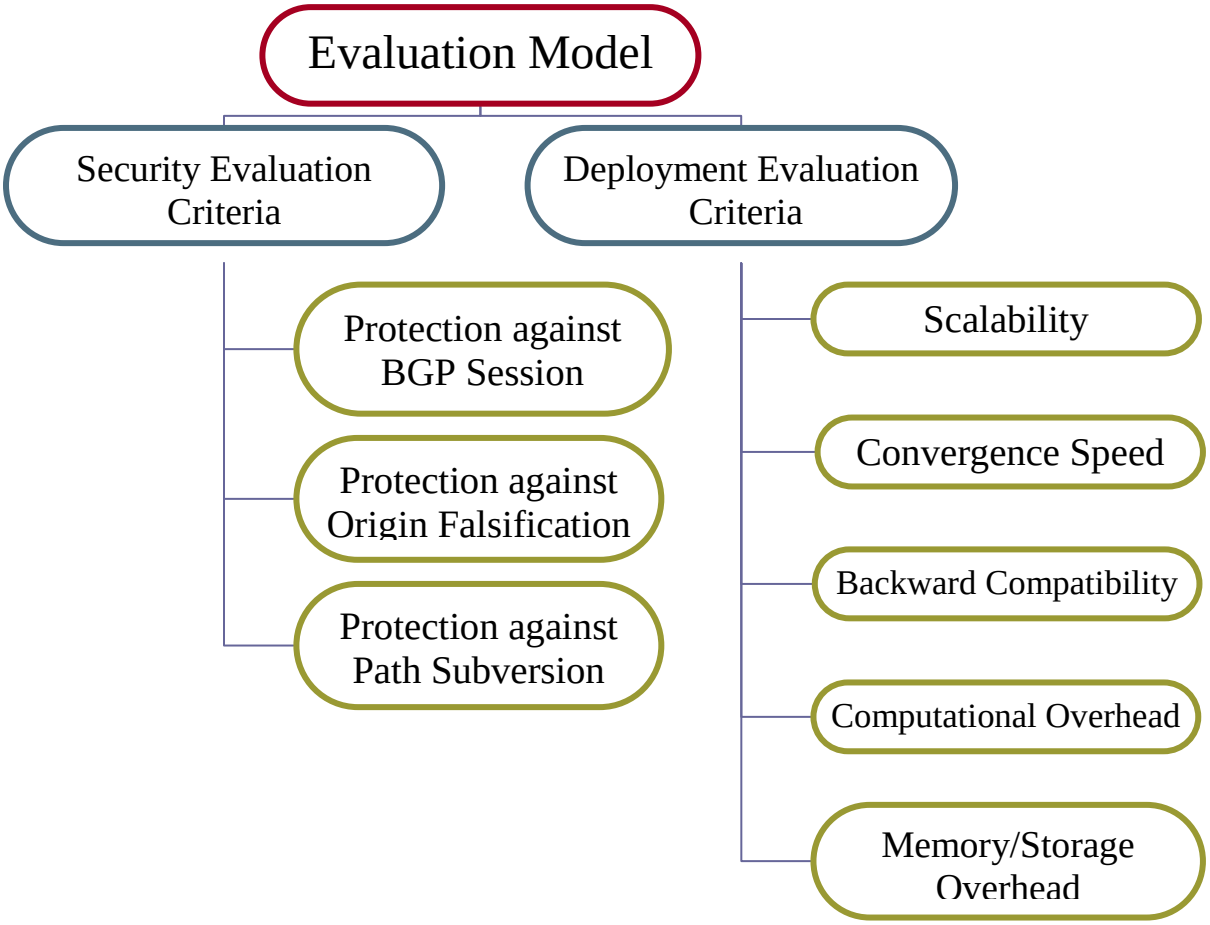

Figure 5.1 Evaluation model for BGP.

5.4 Summary

In this chapter, an evaluation model for BGP is proposed. The model is based on security and performance requirements for BGP evaluation. The model achieves the best tradeoff between security and performance requirements for BGP, hence also being suitable for evaluation for developing deployment strategy for BGP.

Evaluation of BGP Security Architectures

6.1 Introduction

The model and working of security architectures are explained in detail in Chapter 4. This chapter presents evaluation of the discussed security extensions. The evaluation is done by making use of two types of approaches. The first approach is evaluation of security services provided by these solutions to assess the degree to which they mitigate the attacks discussed in Chapter 3. The second approach evaluates the operational feasibility of these solutions with respect to performance requirements. The criteria used for evaluation of security architectures is the Evaluation Model presented in the previous chapter, which covers both security and performance/deployment requirements.

6.2 Security Evaluation

The security solutions for BGP explained in detail in Chapter 4, are evaluated for security features. The evaluation is conducted to check the extent to which security threats and attacks are addressed and prevented. The assessment is based on the Security Evaluation Model presented in Chapter 5.

6.2.1 Protection against BGP Session Attacks

Three kinds of BGP session attacks are described in the threat model (Chapter 3); confidentiality, integrity and session termination attacks. [20]. Security solutions including S-BGP, SoBGP and psBGP make use of IPSec ESP protocol (with NULL

encryption) to protect BGP sessions, which is accepted as the most suitable solution for this purpose.

BGPSec is still in the state of ‘work in progress’ and it has not yet identified exactly what security mechanism would be used for protecting BGP Session between BGP peers. According to latest draft published on ‘Security Requirements for BGP Path Validation’, BGPSec design must resist attacks against BGP session including protection against message insertion, deletion, modification, or replay [92].

NULL encryption algorithm doesn’t provide confidentiality service it is actually used to enable using IPSEC ESP protocol without confidentiality. IPSec ESP is chosen because if confidentiality on peer connections is needed, any other block cipher can be used instead of NULL algorithm. IPSec authenticates both peers to each other and protects against insertion, modification and replay attacks. It also protects against the DOS attacks on the protocol and provides flexibility of negotiation of cipher suite at the time of session establishment. Replay attacks are prevented using authenticated sequence numbers which provides freshness mechanism. Deletion attacks cannot be prevented using any secure protocol; it requires physical security measures of the link between peers.

Hop Integrity protocol suite [72] assumes that each peer has an authentic public key of the other peer. The strong authentication protocol is used to provide integrity and freshness to the messages between BGP peers, preventing insertion, modification and replay attacks. GTSM [51] itself is not any concrete security mechanism but effectively adds another layer of protection against outsider attacks. It offers security between directly connected peers. In multi-hop BGP session, it adds another layer of protection only if used along with other protection mechanisms.

Thus, the use of IPSec, MD5 and hop integrity protocols provide message integrity and origin authentication. This prevents all attacks based on TCP implementation (e.g., TCP SYN flooding, TCP RST attacks). IPSec in addition also provides confidentiality. The above discussed solutions only provide protection against malicious spoofed insertion, modification and replay attacks. There is no check for correctness of information exchanged between both the peers. In case a peer is compromised, other mechanisms have to be employed to check the correctness of the information.

6.2.2 Protection against Origin Falsification Attacks

Malicious or mis-configured AS can originate prefixes which are not owned by it. This can lead to service disruption if neighboring AS accept this information without validation. S-BGP introduced the concept of route attestations, through which the owner of the prefix gives permission to an AS to advertise its prefixes and sign the address attestation with its private key. The receiving BGP speaker verifies the origin of prefix, achieving origin authentication. If any transit AS forwards the update message without verification, the next AS will not accept this information if address attestation is not present in its local database or address attestation doesn't testify that the authority of AS for originating the prefix. Thus, S-BGP's address attestations provide strong protection against prefix-hijacking, de-aggregating attacks and mis-configurations. SoBGP uses a similar kind of concept as introduced by S-BGP for verifying address ownership. The difference is the address delegation process. In SoBGP, address space is delegated directly to ASes by other AS (which may be regional registry, ISP etc) instead of organizations. This different delegation system actually provides one AS a facility to authorize other AS to originate prefixes owned by it. This is done by using the Authorization Certificate (AuthCert). It verifies that the authorizing AS has any authority

to gives permission for the particular prefix. Thus the same kind of security guarantee about the ownership of prefix is achieved, as in case of S-BGP's address attestations.

BGPSec makes use of same concept to protect against origin falsification attacks that was introduced by S-BGP. However, instead of using the name 'address attestation' it uses the name 'Route Origination Authorization" for signed object. ROA is a digitally signed object using the private key of the holder of IP prefix and through ROA, the holder of the prefix explicitly authorizes one or more Autonomous Systems (ASs) to originate route for the prefix. For each AS, separate ROA will be generated. Any BGPSec speaker, wants to verify that whether AS that originated the prefix has authorization from the prefix holder, it will make use of certificates and ROA downloaded from RPKI repository and can verify the ROA issued in this regard. Like S-BGP, BGPSec's ROA provide strong cryptographic protection against prefix hijacking and BGP mis-configuration scenarios.

IRV provides very weak notion of security to validate the origin of the prefix. Any AS which receives an UPDATE for any prefix directly sends the query to the AS which originated this update to check that the prefix was actually originated by it. This protects against mis-configurations up to some extent but doesn't provide any protection against the AS which maliciously originated the prefix. psBGP makes use of Prefix Assertion lists (PALs) in order to verify the prefix ownership. The PAL is kind of endorsement from neighbor AS for the prefixes owned by its neighbor. The psBGP approach assumes that no two neighboring AS can be malicious. This assumption stands weak as majority of attacks are launched by insiders [82]. Such AS can simply assert prefix to be hijacked in one AS's PAL and other AS's PAL only have to endorse it. Also, this approach puts all the responsibility on the AS to verify which addresses are owned by their neighboring

AS which can render this security mechanism for authenticating prefix origin not helpful. Any malicious AS can insert wrong prefix assertion about its neighbor AS in its PAL to force BGP routers to not to accept and forward the prefixes originated by the neighbor. These threats and limitations hinder the adoption of such security mechanism [34, 35]. The MOAS propose the use of BGP community attribute to carry MOAS list which is an optional transitive attribute. The main limitation of this approach is that optional transitive attribute can be dropped by any AS who receives this update. Secondly, the authenticity of the MOAS list cannot be guaranteed without cryptographic protection. The information in the MOAS list can be changed by any malicious AS in the transit and the service may be denied to the legitimate owner of the prefix [73].

SPV also provides strong prefix origin authentication. The multi-epoch public key which is used to authenticate all the signatures on the path is the value computed by the originator of the prefix using one way hash on all one-time signatures (generated using single ASN private keys). The multi-epoch public key is signed by the prefix private key which is only owned by the owner of the prefix. Any UPDATE is only accepted if the BGP speaker verifies all one-time signatures on the path. SPV has a complex design but is very efficient in terms of speed.

S-BGP and SoBGP and SPV provide strong cryptographic verification of prefix origin. These techniques mitigate all kind of origin falsification attacks and also provide protection against router configuration mistakes. IRV can only provide protection against mis-configurations if and only if every AS's IRV server maintains up to date record of the configuration changes and provides reliable view of AS's routing policy. MOAS list can naively detect and mitigates MOAS conflicts but it is vulnerable to isolated adversary attacks. The assumption of psBGP that no two neighbouring AS can be malicious, is not

practical. Moreover, many AS have only one neighbor and any malicious neighbor can deny services to its neighboring prefix, thus making the protocol unsecure.

6.2.3 Protection against Path Subversion Attacks

Malicious manipulation of path information in transit can result in black holing, eavesdropping of data, denial of service and degradation of service etc. The proposed solutions present different schemes for providing path authentication. S-BGP introduce the concept of route attestations, which are carried in an optional transitive attribute called attestations. Every BGP speaker when receives update message, verifies the signatures on the path information in that update message. This recursive signature verification provides the authenticity of the path. Any malicious ASN can not shorten the path or add any AS in the path. The information is integrity protected by the signature of the preceding AS. Thus strong path authentication mechanism is provided. A major drawback is that attestations are optional transitive attribute. It is very likely that during the incremental deployment stage some of BGP speakers which don't participate in S-BGP route authentication may drop this attribute, which goes unnoticed at the receiving S-BGP speaker. S-BGP also provides no protection against the attack launched by colluding AS. A malicious AS can forward the Update to other malicious AS through tunneling link which when forwards the update to next AS. This way colluding AS can force BGP speaker to accept this shorter route to prefix.

Although work still in progress, BGPSec makes use of same concept that was introduced by S-BGP for protection against BGP Path Subversion attacks. However, instead of using the optional transitive attribute named 'attestation' to carry the digital signatures, BGPSec uses optional non-transitive attribute named "BGPSec_Path_Signatures". The difference between optional transitive and non-

transitive BGP attributes is that optional transitive attributes if not recognized by BGP implementation on BGP router then will be marked as ‘partial’ and propagated to other neighbors but optional non-transitive attributes would be discarded by BGP router, if not recognized. BGPSec provides the same level of Path Integrity as discussed for S-BGP already above.

SoBGP uses global topology map to validate the path. Every AS distributes a special kind of certificate which contains the information of AS’s peers and AS’s policy. This information is integrity protected (signed with private key). Every BGP speaker stores these certificates and uses them to check the received path in the corresponding AS ASPolicyCerts. SoBGP provides protection against configuration errors and attacks lunched by malicious autonomous systems but it doesn’t provide any guarantee that the update was actually traversed through the path mentioned in AS_Path. SoBGP is thus vulnerable to same kind of colluding AS attack as in case of S-BGP. Any two malicious ASes can insert false information in their ASPolicyCerts to send route advertisements through tunneling link. This can be propagated to their neighboring AS to force them to accept shortest paths to prefixes. This way malicious AS can eavesdrop on data or can drop data packets to cause denial of service attack for the victim’s prefixes.

IRV is a separate protocol which runs in parallel to BGP protocol and makes use of IRV server and IRV client. IRV client queries all the AS along the path to validate received prefix. It also checks the AS they advertised this update. IRV Servers in corresponding AS respond to the queries and provides the relevant information. These responses from servers are protected by IPSec or TLS connection. The received AS_Path is not protected by any cryptographic mechanism. Thus this protocol provides partial protection against mis-configurations. The weak point is that it relies on the un-

authenticated routing to authenticate routes. Any malicious AS can alter the path and claim to be the originator of the prefix which its IRV server can endorse. An AS whose IRV client validates the information from this malicious AS will accept the route update. Moreover, colluding AS can also launch all kinds of sophisticated attacks against IRV route validation mechanism. For example, any malicious AS can add the AS numbers of other malicious AS in the AS_Path attribute. All malicious AS can validate that the information was advertised by them which may result in data traffic routed through chain of malicious AS. The effect would be BGP execution without any security mechanism. The static data stored by IRV server may not be up to date. The proposed solution is the use of front end to automatically transfer generated configurations to routers and IRV server. This causes mis-configurations to be mirrored in IRV server. Since some dynamic data (advertised routes) is obtained from border routers in that AS, the IRV server may give positive reply in response to the queries from other AS [66].

PsBGP doesn't provide any path authentication mechanism; instead it recommends the same mechanism as proposed by S-BGP but with bit-vector improvement [68]. Bit-vector optimization technique is concerned with efficiency of forwarding UPDATEs. The security evaluation for psBGP is the same as for S-BGP. The Cumulative Authentication Mechanism itself is not suitable for BGP as path authentication is not performed at each node [78]. Only authenticator node at the end of the path can compute recursive MAC on the path. The use of TESLA broadcast protocol enables cumulative authentication mechanism to provide path authentication at each node. This mechanism is not suitable for BGP as even using TESLA protocol receiving node can only authenticate that the information received from previous node is authentic but in order to authentic signatures on the path recursively receiving node must be having access to all the TESLA keys

released by previous nodes. Secondly cumulative authentication mechanism does stop second node from removing first node from the address list. If this mechanism is used in BGP without any separate cryptographic proof of address ownership then any malicious AS can remove the previous nodes from the list and signature from the authenticator field and can claim to be the owner of the prefix. Moreover, any malicious AS can add as many AS in the address list and generate their signature from its own. The BGP speaker which receive such an update message, if it is able to compute recursive MAC on the authenticator value then it will be verified as there was no proof that which TESLA belonged to which AS or BGP speaker. The colluding adversaries' attacks are even easier. Furthermore, the use of TESLA authentication protocol demands that all the participating nodes must be loosely time synchronized with each other which is rather very difficult in practical situation.

SPV uses a symmetric cryptography based path validation mechanism. In SPV, the update message includes onetime signatures on the path and the single ASN private key (which preceding BGP speaker generated applying hash function to its own single ASN private key). The signature is actually generated on the epoch value and the AS_PATH. This concept is similar to S-BGP route attestations. In addition, the next AS in the path will validate that the received path from last AS is consistent with the previous signatures on the path. This is to prevent any malicious AS to alter the path as any AS can generate next AS's 'single ASN private key' and thus the signatures as well. SPV protects against path alteration attacks mounted by isolated adversaries. One drawback of using one time signature is that if the same parameters are used to sign multiple messages then the signatures can be forged; a malicious BGP speaker receives multiple routes for same prefix different peers, the probability of falsifying the AS_Path is increased by this

speaker. Secondly, SPV uses optional transitive attribute to carry signatures which can dropped by any malicious AS and next SPV capable router in the path will not be able to verify the authenticity of the route. Furthermore, malicious AS can send route advertisement through tunneled interface to other malicious AS somewhere else in the world, creating routing instability for any prefix or denial of service and degraded service.

Using public key cryptography based signatures (as used in S-BGP) provides better level of security in comparison to one-time signature used in SPV and other consistency checks (used by SoBGP and IRV). S-BGP, SoBGP and SPV protects against configuration errors and most of the path alteration attacks launched by an isolated adversary but none protects against sophisticated attacks launched by colluding adversaries. S-BGP and SPV uses optional transitive attribute to carry route attestations (the use of this attribute is to back ward compatibility) but optional transitive attribute can be dropped by any transit AS. IRV is vulnerable to both attacks launched by individual AS and colluding ASes. The psBGP doesn't provide any path authentication mechanism and uses the path authentication proposed by S-BGP with efficiency improvement. The cumulative authentication mechanism is not suitable for BGP and also vulnerable to isolated and colluding adversaries attacks. The invalid routes which can result in colluding adversaries' attacks cannot be detected unless the complete topology is known and enforced [29]. SoBGP tries to build topology map but this topology is built on the basis information provided by individual ASes. Any two malicious colluding AS can provide fake information in order to maliciously manipulate a path.

Table 6.1 Overview of security evaluation of BGP security solutions.

Security Solutions	Protection Against Origin Falsification Attacks	Protection Against Path Subversion Attacks	Protection Against BGP Session Attacks
S-BGP	Yes	Yes	Yes
SoBGP	Yes	Yes	Yes
IRV	Yes	Yes	No
PsBGP	Yes	Yes	Yes
BGPSec	Yes	Yes	Work in Progress (Planned)
Hop Integrity	N/A	N/A	N/A
MOAS Scheme	Yes	N/A	N/A
Path Auth Scheme	N/A	Yes	N/A
SPV	Yes	Yes	No

6.3 Deployment Evaluation

This section evaluates performance of security architectures for BGP. The assessment is made to check how well the security solutions perform with respect to the operational requirements defined in Evaluation Model. The evaluation also explains why the security mechanisms are not utilized when provide protection against a good number of threats to BGP.

6.3.1 Deployment Evaluation of SBPG

S-BGP makes heavy use of digital certificates and assumes that two hierarchical PKIs are in place. If it is assumed that such infrastructure can be established, it is still difficult to find how existing IP address space is allocated and delegated [34, 35] and re-issue IP address space and AS numbers to organizations by providing them digital certificates. Every S-BGP speaker needs to store public keys and address attestations to validate prefix origin and route attestations. Each non-leaf S-BGP requires a number of public keys; one per internet registry, one per organization that has been assigned an address prefix, one per organization that has been assigned an AS number, one per AS (on assumption that each AS has atleast one BGP speaker) and one per AS/BGP speaker [83, 84]. Each S-BGP speaker also need to store one Address Attestation (AA) per organization which has been assigned an address prefix [83]. The BGP speaker doesn't store certificates; instead 2nd tier repository validates all the certificates and produces more compact information for each certificate and address attestation. Still significant storage space is required for S-BGP speaker which also stores countless routes to different destinations in the world.

Every S-BGP speaker needs to verify the route attestations on the path and generates signatures on updates which it advertises to other AS. This creates computational overhead as public key based computations are very expensive. Every BGP speaker needs to validate many updates per second, the actual processing time depends upon the CPU speed of the router. During peak traffic load and because of route flapping, a BGP speaker may receive enormous amount of Updates from its peers which it cannot handle if it needs to verify all attestations. Moreover when a new S-BGP speaker is added, it is flooded by routing updates from its peers which needs verification for all the route

attestations, increasing computational requirements. The bandwidth consuming DDOS attacks on the internet can also affect S-BGP routing on the internet and BGP speaker receives more updates than they normally expect to receive.

Later enhancements in S-BGP proposed that a route will only be verified if it is selected as the best route to the destination and certificates signatures when verified can be cached for some time. This reduces the processing load on the routers. When S-BGP speaker send updates to other peers it needs to create separate UPDATE message for each peer because of security reasons. The Bit-Vector improvement reduces this processing issue. Furthermore, signature amortization scheme (S-A) enhances processing speed of S-BGP [68]. Every BGP speaker stores UPDATE message received from peer in Routing information bases (RIBs) until withdrawn by corresponding peers. Routing attestations add mega bytes of data at peer. So routers with many peers need Giga-bytes of RAM in order store this data, which incurs very high storage cost.

The convergence speed of any update will be affected on the internet by processing delays at S-BGP routers. Under normal circumstances the convergence speed may be tolerable but route flapping, link cuts and other factors contribute in slow convergence of BGP updates in case S-BGP is used. S-BGP is backward compatible with BGP, so is suitable for incremental deployment. It uses optional transitive attribute to carry route attestations which existing BGP speakers can transparently pass on to other routers. If PKI is established then few of the S-BGP speakers can benefit from address attestations which are distributed out of band and can verify prefix ownership.

S-BGP cannot be deployed without having special processing power at routers and Giga-bytes of additional RAM so it is less likely that it can tolerate the growth of internet.

The traffic on the internet is expected to increase manifolds which make it difficult for S-BGP speakers to handle all the data.

6.3.2 Deployment evaluation of SoBGP

The web-of-trust model for authenticating AS proposed by SoBGP is questionable as AS numbers are controlled by IANA. Further the IP addresses are currently allocated to organizations not to AS directly, so this distributed web-of-trust model needs major changes in the current address delegation system. Like S-BGP, SoBGP also uses many digital certificates but require much lesser storage space. Each SoBGP speaker needs to store one public key certificate for each AS, one prefix policy certificate (PrefixPolicyCert) per AS which has been assigned address prefix and one AS policy Certificate (ASPolicyCert) for each AS. SoBGP doesn't use any cryptographic signature generation or verification process. In order to validate the certificate chains, routers need to compute cryptographic signatures which consume significant processing power. SoBGP provides variety of ways to reduce processing. If AS's internal repository is used to validate the certificate then routers do not have to perform expensive cryptographic operations and the processing cost would be justifiable.

The main problem with SoBGP is that it doesn't give a single solution to distribute certificates but provides multiple approaches for that and doesn't mandate any choices, which is a problem in itself. Any good solution must address each and every component clearly so that others can understand and properly estimate the processing cost. At first SoBGP proposed a new BGP security message to carry all certificates but now it says that certificates can be also be distributed out of band. Without any mandatory choices in the proposed solution could create interoperability issues if it is implemented [84].

SoBGP router doesn't compute expensive cryptographic operations (if it was provided already validated information from local repository) so the convergence speed is likely to be near to normal BGP convergence speed. Moreover, SoBGP also propose that it is possible to install the routes first in local routing table and later authenticate the route and prefix origin (which is obviously a trade-off between security and speed). SoBGP proposed a new BGP security message (although not mandated) to carry certificates which doubts its backward compatibility with existing BGP implementation. The existing BGP implementation cannot understand this BGP messages so its backward compatibility can only be achieved if existing BGP implementation is modified in order to recognize and forward this new SECURITY message and the new version is updated on all the routers on the Internet, which is not called backward compatibility in a true sense. If out-band-band mechanism is used then benefits can be achieved in incremental deployment phase. SoBGP speakers on the Internet can authenticate the AS's authority to originate prefix. If all the AS's distribute their ASPolicyCerts then path can also be validated by SoBGP speakers.

6.3.3 Deployment Evaluation of IRV

IRV is an out-of-band protocol which doesn't interact with BGP routers; instead it is a separate system/protocol (i.e., IRV Server/Client) which automatically checks for validity of routes according to the AS local security policy (i.e., check the routes as soon as update receives or first install the routes and periodically check the authenticity of routes). This system easy to be deployed but the main problem is that it doesn't guarantee strong security. IRV requires IRV server and client in every ISP but not every AS can be convinced to install an extra server which will provide security to the internet routing.

Also, it cannot be guaranteed that the IRV server represents the exact view of AS's policy and it is properly updated.

IRV requires a central database to store the addresses (IP addresses) of all the IRV servers in each AS. The problem is management of this central repository (because it may need secure access to upload and download data) and authentication of AS to upload their server address. IRV is a separate protocol so it is totally compatible with existing BGP infrastructure. But the benefits during incremental deployment can only be achieved if majority of the autonomous system implement IRV Server/Client. Convergence speed will only be an issue if the AS wishes to authenticate the route as soon as the update is received by an AS.

6.3.4 Deployment Evaluation of psBGP

PsBGP requires the same kind of centralized PKI for AS number assignment as proposed by S-BGP. Current authorities do not provide this kind of service so there is need to make changes to system which allocates AS number to organizations. PsBGP makes use of public keys like S-BGP and SoBGP. Each psBGP speaker needs to store one public key certificate per AS, one public key certificate per AS for BGP speakers (SpeakerCert), one prefix assertion list (PAL) per AS and one certificate (MultiASCert) per organization that has more than one AS. PsBGP significantly reduces the storage requirement for each psBGP speaker in comparison to S-BGP. Its storage requirement is equivalent to SoBGP in no. of certificates but may differ in size [34, 35].

PsBGP doesn't propose any new mechanism to path authentication and relies on S-BGP mechanism for path authentication so the above description for processing overhead, memory overhead and convergence speed holds true for psBGP. The only difference is

that psBGP recommends bit-vector improvement [68] which can reduce processing time to generate signatures when psBGP speaker sends same update message to multiple peers.

6.3.5 Deployment Evaluation of SPV

Like S-BGP, SPV also requires centralized public key infrastructure for address delegation. Each SPV speaker needs to store one public key certificate per prefix (address space) per organization which has been assigned an address space and one multi-epoch key per prefix. The authors of [34] didn't specify the exact size of the multi-epoch key but it is likely to be significantly high. The one time signatures used by SPV are based on symmetric cryptographic primitives which are comparatively much faster than public key based signatures, especially signature generation is much faster than signature verification. But processing is still a significant overhead for existing routers. SPV also requires that the processing accelerators must be added to existing routers but authors of [34, 35] argued that accelerators for SPV are cheaper than accelerators for S-BGP.

The most fundamental problem with using one time signature is that its size is much bigger than the public key signature used in S-BGP. So it requires Giga bytes of RAM in order to store the routes received from neighboring AS in router's Routing Information Bases (RIBs). This requirement may hinder the adoption of such kind of mechanism. SPV route convergence speed is likely to be better than S-BGP as the signature and verification process is faster than S-BGP. SPV is fully backward compatible with BGP. It uses optional transitive attribute to carry route signatures. During Incremental deployment phase it does provide some benefits as if the originating AS uses SPV then the rest of the SPV speakers can atleast validate that the prefix was originated by authorized AS but authenticity of the whole path cannot be guaranteed until all the BGP speakers implement SPV.

6.3.6 Deployment Evaluation of BGPSec

Like S-BGP, BGPSec also makes heavy use of digital certificates and the hierarchical PKI named RPKI. RPKI infrastructure and repositories has been implemented by all five RIRs using self-signed trust anchor AFRINIC RPKI [93], APNIC RPKI [56], RIPE RPKI [57], ARIN RPKI [94] and LACNIC RPKI [95]. At the moment RIRs offering hosted and delegated models for generation of resource certificates and ROAs. However, it will take significant amount of time for resource holders to ask for generation of certificates and ROAs. But it is not an impossible effort and will certainly take time to deploy and issue resource certificates and generate corresponding ROAs.

Every BGPSec speaker needs to store resource certificates, BGPSec router certificates and ROAs to validate prefix origin authorizations and signatures in BGPSec_Path_Signatures attribute. Like S-BGP, each BGPSec router requires a number of public keys; one per IANA, one per each RIR, one per each NIR, one per each LIR, at least one per organization that has been assigned an address prefix, one per organization that has been assigned an AS number, at least one per AS for its each BGPSec border router. Each BGPSec speaker also need to store one ROA per each IP Prefix. According to the study conducted by NIST [96], if on average 04 AS numbers are present in an update message then using RSA-2048 would increase the current BGP update size to 15 times of the current size. Moreover, the same study concludes that in BGP-4, there is an average of 3.83 prefixes per update but BGPSec is restricted to include just one prefix per update. So there is significant amount of storage / memory requirement on each BGPSec router that would require replacement of most of the internet BGP routers.

Every BGPSec speaker needs to verify the multiple signatures on the path kept in BGPSec_Path_Signatures attribute as well as validate ROA against each BGP update message. BGPSec also needs to generate signatures on updates which it advertises to other AS. This creates computational overhead as public key based computations are very expensive. Every BGPSec speaker needs to validate many updates per second, the actual processing time depends upon the CPU speed of the router. During peak traffic load and because of route flapping, a BGPSec speaker may receive enormous amount of Updates from its peers which it cannot handle if it needs to verify all attestations. Moreover when a new BGPSec speaker is added or in BGPSec speaker failure-recovery scenarios, it is flooded by routing updates from its peers which needs verification for all the route attestations, increasing computational requirements. The bandwidth consuming DDOS attacks on the internet can also affect BGPSec routing on the internet and BGP speaker receives more updates than they normally expect to receive.

Like S-BGP, the convergence speed of any update will be affected on the internet by processing delays at BGPSec routers. Under normal circumstances the convergence speed may be tolerable but route flapping, link cuts and other factors contribute in slow convergence of BGP updates in case BGPSec is used. BGPSec is backward compatible with BGP, so is suitable for incremental deployment. BGPSec cannot be deployed without having processing power to handle cryptographic computations and verification at BGP routers and multi-giga bytes of additional RAM so it is less likely that it can tolerate the growth of internet if any workable deployment strategy is not proposed. The traffic on the internet is expected to increase manifolds which make it difficult for BGPSec speakers to handle all the data.

6.4 Summary

In this chapter, detailed analysis of the security attributes and performance features of BGP security architectures is presented. The evaluation is conducted on the basis of the evaluation model defined in Chapter 5, which includes all necessary security and performance requirements for BGP deployment.

A Deployment Strategy for BGP

7.1 Introduction

The latest security extension for BGPv4 is BGPSec. The thorough analysis of security and performance of BGPSec is presented in the previous chapters. The work on BGPSec is still under progress by SIDR WG of IETF. BGPSec is the prime effort at international scale to secure current BGP protocol so no new security extensions to BGP are proposed after it. The only active official work in progress at the moment is BGPSec [85, 86]. This chapter presents a deployment strategy for BGPSec, that overcomes the difficulties in its operational utility.

7.2 Constraints in BGPSec

BGPSec assumes addition capability in the BGPSec routers which includes additional memory to store routes with added signatures in BGPSec_Path_Signatures attribute. Moreover, border BGPSec router is required to cryptographically verify ROA and multiple signatures in BGPSec_Path_Signatures attribute which requires additional processing power on the routers [85]. The processing requirement in terms of memory and computation on border BGP routers is the major hurdle in operational deployment of BGPSec.

One of the goals of this research work is to propose a workable strategy for incremental deployment of proposed secure extension to BGP. Thorough analysis and evaluation shows that BGPSec provides security guarantees required against the major vulnerabilities identified in the current BGP Protocol and is based on best features of all

previously proposed security extensions to BGP. Efforts are in progress for developing optimized and workable solutions for BGPSec implementation. NIST is heavily involved in the design of BGPSec protocol and has also developed an open source prototype software reference implementation for BGPSec with the name 'BGP Security Routing Extension (BGP-SRx) [87]. Current research is focused on developing efficient router platform capable of handling cryptographic operations at nearly line speed even under router failure-recovery scenarios to meet the speed of BGP convergence [88, 89].

This chapter presents a workable deployment solution to implement BGPSec using the existing BGP routers and off boarding all cryptographic computations from router platform.

7.3 Deployment Strategy

The solution proposed for successful and incremental deployment of BGPSec requires a dedicated server named as 'Cryptographic Verification & Computation Server' (CV&CS) within each ISP. This addition of this server off boards all cryptographic verification and computation workload from border BGP speakers. Moreover, an optional (non-transitive) attribute named as "Route_ID" is introduced for iBGP, having its significance and visibility only internally within the Autonomous System (AS).

7.3.1 Role of CV&CS

The CV&CS acts as a RPKI repository for the ISP and stores and maintains up-to-date copies of all signed objects that exist within the repository system of RPKI. This can be accomplished by establishing Rsync connections with RPKI repositories of all five Regional Internet Registries (RIRs) and IANA.

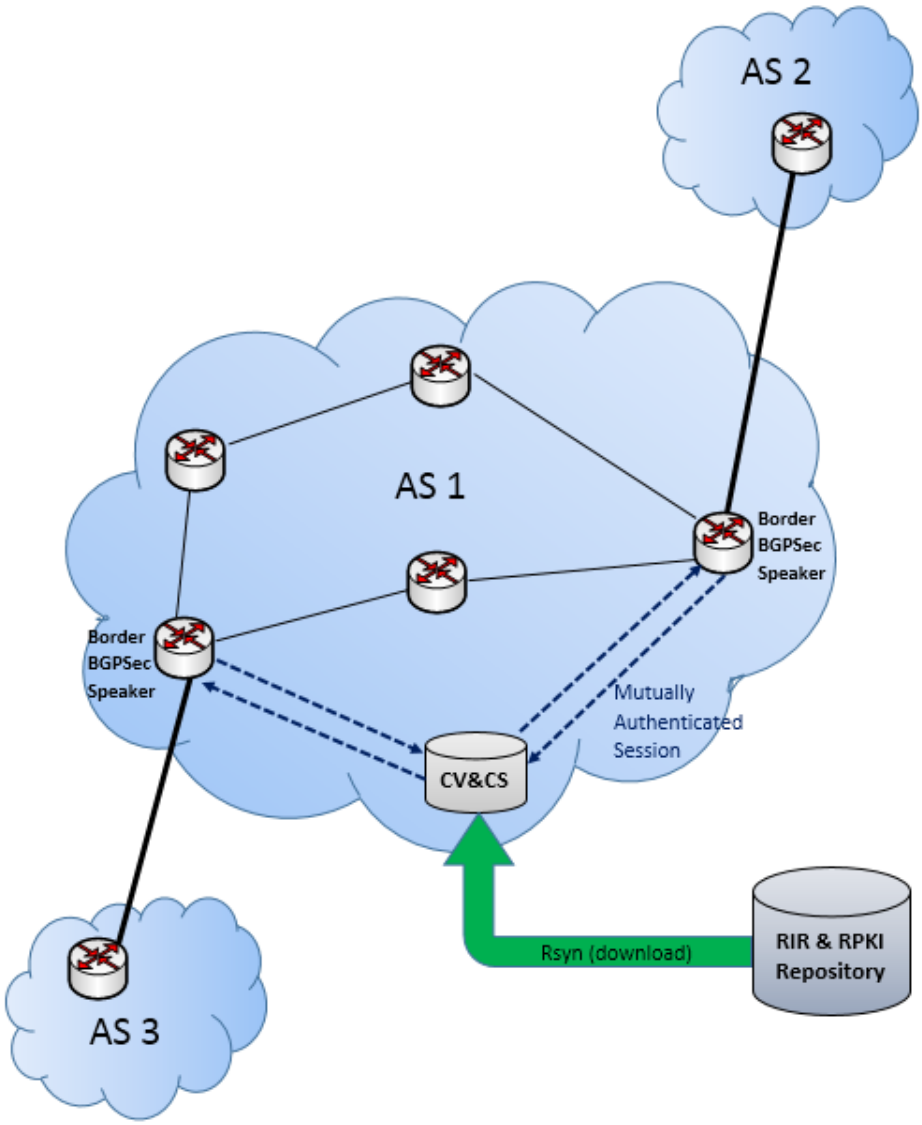

Figure 7.1 Proposed deployment solution for BGPSec.

The BGPSec router certificate can also be distributed through the distributed repository system. Therefore, CV&CS would contain all BGPSec router certificates in the same manner as it downloaded other signed objects including certificates, CRLs, ROAs and manifests. To improve performance, CV&CS would maintain up-to-date indexed data of all signed objects and will have access to updated Routing Policy information for all peering relationship of its AS. Only CV&CS will have access to private key corresponding to BGPSec router certificate issued for signing and adding

signatures in BGPSec_Path_Signatures attribute on behalf of its Autonomous System. The use of CV&CS completely off boards all cryptographic computational loads of the border BGP routers including Path verifications and ROA validations.

7.3.2 Receiving Routing Updates from BGPSec Peers

All border BGP routers are proposed to have mutually authenticated secure connections established with CV&CS within their Autonomous System (AS). Upon receiving BGPSec enabled updates from other AS's BGP router, the border BGP speaker will forward the update message to the CV&CS for verification of path and ROA. CV&CS is to perform all cryptographic computations to validate ROA for the update received as well as verify all digital signatures in BGPSec_Path_Signatures attribute of BGP update message. If the result of verification process is 'OK', CV&CS removes BGPSec_Path_Signatures attribute from BGPSec Update message. It generates unique route ID and insert it in Route_ID attribute. It then append Route_ID attribute to the subject BGP Update message and sends the update to the same BGP router again.

CV&CS will also store indexed Route_ID and corresponding complete BGPSec update message in the local database. It checks the received update against routing policy. If routing policy allows to further propagate the routing update to peers, CV&CS pre-computes the BGPSec signature on the AS_Path for each BGPSec neighbor and stores the updated BGPSec update message in the local export database.

In case the verification process fails, CV&CS would discard the BGPSec update. Border BGP speaker that forwarded the BGPSec update to CV&CS will only accept verified BGP updates from CV&CS. These are normal verified BGP updates with single addition of Route_ID attribute. After receiving BGP update from CV&CS, the border BGP router would add the update to its Adj-RIB-In for particular peer BGP in order to

process it for decision making. The protocol resumes conventional BGP execution from this point onwards. The working is explained in Figure 7.2.

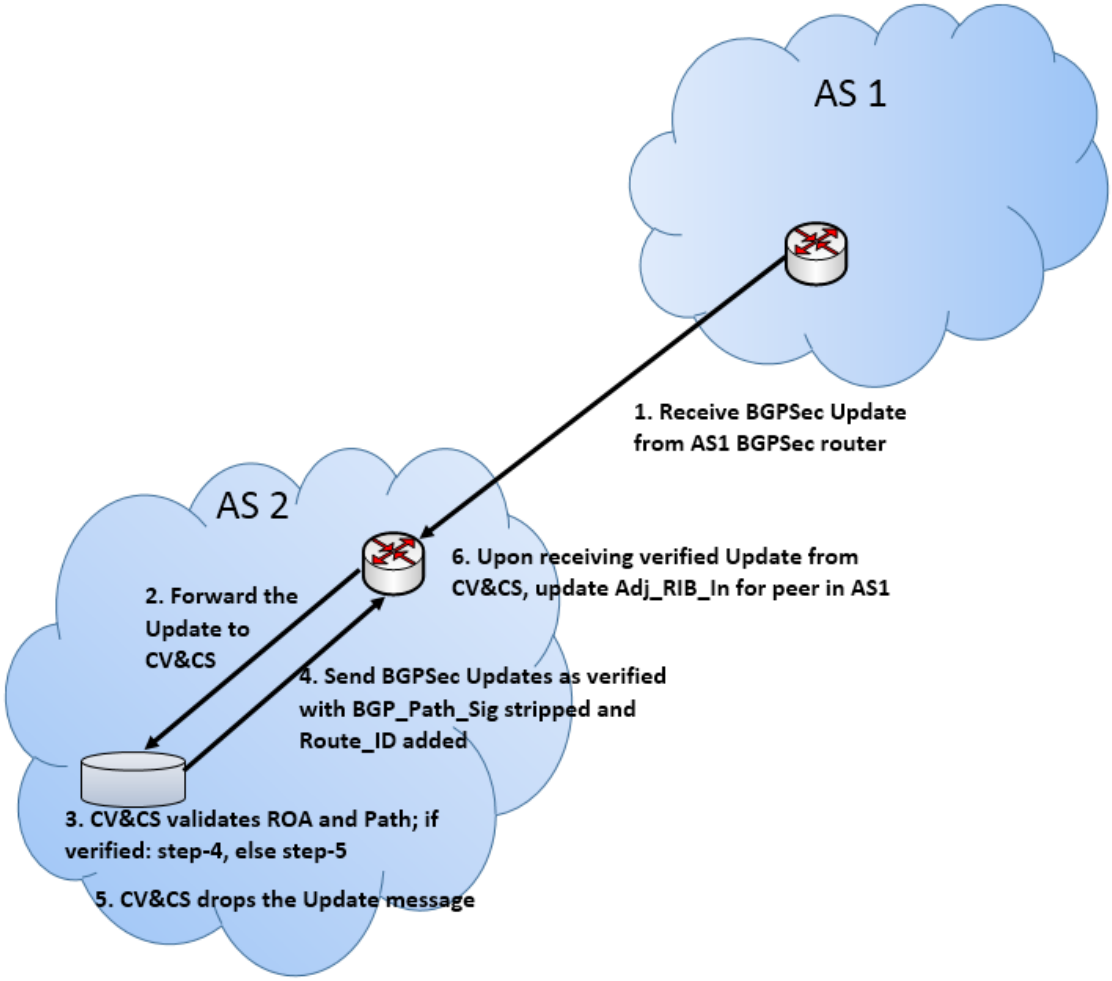

Figure 7.2: CV&CS role while receiving BGPSec Updates.

7.3.3 Propagating Routes to Peers

Before propagating received routes to BGPSec peers in other AS, border router will send BGP update to CV&CS. The CV&CS checks its export DB using the value in Route_ID for BGP Update message with pre-computed BGPSec_Path_Signatures for the particular peer to whom the route is intended to be advertised. The CV&CS immediately returns pre-computed BGPSec update message back to the same BGP border router (without Route_ID attribute). The BGPSec router immediately sends BGPSec enabled

update to its BGPSec peer in other AS. This mechanism is used for route propagation to peers. CV&SC Role for advertising routes is explained in Figure 7.3.

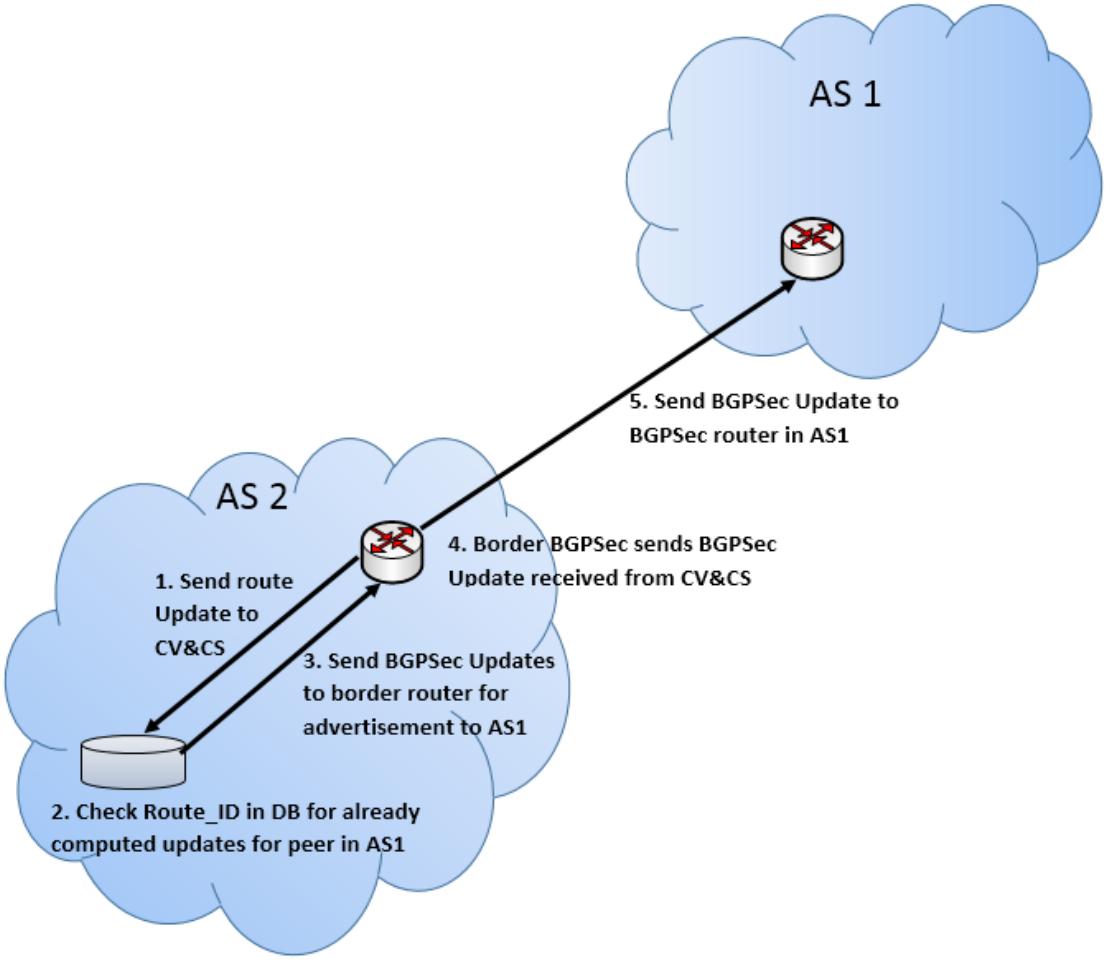

Figure 7.3: CV&CS role while advertising routes.

7.4 Security Analysis

The proposed deployment strategy is based on BGPSec, the latest and most comprehensive security extension for securing BGP. Thus, the security offered by this strategy is essentially the same as its underlying protocol; BGPSec. The offered security services include protection against BGP Session Attacks, Origin Falsification Attacks and Path Subversion Attacks, in the same way as offered by BGPSec.

The proposed solution doesn't introduce any change in BGPSec itself so it provides same security guarantees as that of BGPSec. The proposed solution addresses the deployment issue for BGPSec that requires significant additional computational and memory resources for validating cryptographic signatures and signed objects computed by BGP routers as part of the process. This has remained a major hurdle in BGPSec deployment so far.

The proposed strategy inherently provides security benefit over currently proposed BGPSec extension. This solution requires the private key corresponding to BGPSec router certificate to be kept with just CV&CS server instead of being shared among multiple border BGPSec routers of an Autonomous System (AS). The private key corresponding to BGPSec router certificate is used to compute signatures on AS-Path and other attributes on behalf of AS that is to be put in BGPSec_Path_Signatures attribute of BGPSec updates advertised by the AS. This minimizes the risk of compromise of private key corresponding to BGPSec router certificate.

However, there could be a concern that the introduction of additional server (i.e., CV&CS) and communication protocol between BGP routers and CV&CS may introduce additional security threats and vulnerabilities that must be identified. A few security concerns relating to CV&CS may arise. In case CV&CS server of an AS is compromised, the attacker can inject false BGP updates in the Local Routing Information Base (RIB) of BGP routers having established trust relationship with CV&CS server. In such an event, the attacker can also compromise private key corresponding to BGPSec router certificate. The CV&CS server can be compromised in a manner that the attacker is able to stop providing service to the BGP routers. Additionally, the communication protocol can be compromised so that the attacker is able to replay old messages and launch integrity

attack on the communication session or to masquerade as CV&CS server to router or BGP router to CV&CS server.

To counter these threats, the CV&CS server has local significance and shouldn't be accessible from outside the AS. The best practices to secure critical servers must be followed to mitigate most of the risks posed by threats scenarios and reduce the threat level to an acceptable level. Such practices include implementation of two factor authentication to access the management or database administration interface, four-eye principle and effective patch management and restricting access management to only authorized internal machines only.

For communication protocol, in high level proposed design it is recommended to use mutual authenticated session. IPSec with mutual authentication and ESP can be used. The public keys of Border BGP routers can easily be exchanged with CV&CS server manually and vice versa. The communication protocol can thus use IPSec Tunnel Mode for secure communication between BGP router and CV&CS. For protecting private key corresponding to BGP router certificate, special cryptographic hardware cards could be used to prevent compromise of private key even if the server itself gets compromised. Therefore, the established best security practices should be followed to reduce the risks posed by proposed deployment solution to the acceptable level.

7.5 Performance Analysis

7.5.1 Network Throughput Analysis

The proposed deployment solution introduces additional network traffic and communication protocols within the AS. This is attributed to the secure mutually authenticated tunnel session between border BGP routers and CV&CS server and the

introduction of 4 bytes Route_ID attribute that would impact the BGP updates for BGPSec enabled updates within the AS.

Keeping in the view the AS being the single administrative domain and availability of high network bandwidth (usually exceeding 10 GBE) within transit AS, this additional throughput on network bandwidth is considered negligible and would not be of any major concern within the autonomous system.

7.5.2 Incremental Deployment Analysis

The proposed deployment strategy strongly supports incremental deployment requirement of BGPSec. Only BGPSec updates will be forwarded to CV&CS server and the rest of normal BGP updates will be dealt as they are currently being processed by BGP routers.

7.5.3 Memory Overhead Analysis

The proposed BGPSec deployment technique significantly reduces memory overhead in BGPSec. It doesn't require BGP routers to store BGP updates with BGPSec_Path_Signatures attribute comprising of more than 20 bytes of public key based cryptographic signatures. Moreover, BGP router doesn't need to load thousands of public key certificates and other signed objects into the memory to compute and validate the cryptographic signatures which significantly reduces memory requirement of BGPSec speakers to work. In fact, the additional memory requirement to store BGPSec updates has been completely offboarded to CV&C server.

The only memory overhead that has been added to current BGP routers is through the introduction of optional non-transitive attribute named Route_ID. Route_ID has local significance only and we propose this attribute size to be of 4 Bytes as this attribute only

stores unique number of the route for which CV&CS has performed validation and computation. This attribute is only used to refer back to the pre-computed data that has been stored in CV&CS server for further propagating BGPSec updates to peers.

7.5.4 Computational Requirement Analysis

The proposed solution completely offboards the processing requirement introduced by BGPSec on BGP routers. No cryptographic computation and validation will be performed on border BGP routers. Keeping in view cryptographic processing introduced by BGPSec, it may require special hardware based cryptographic processors on all BGP routers. In the proposed solution, this requirement is off-boarded to one specialized server so the overall computational cost is minimized as well.

7.5.5 Single Point of Failure Analysis

Another availability concern of CV&CS is being a single point of failure. This concern can also be addressed using already established best practices for high availability of critical services. CV&CS server can also be configured in high availability hardware clusters with dual and backup power source. This can further be supplemented by OS Cluster, Database Cluster and service level cluster technologies. The redundancy is available in N+1 and N+N configuration at hardware and software level.

7.6 Summary

This chapter presents the proposed deployment strategy for BGPv4, making use of BGPSec. The operational working of the protocol as per the proposed strategy is explained. The addition of a single server per AS removes computational overhead from the border routers, presenting a workable solution for BGPSec deployment.

Chapter 8

Conclusion and Future Work

8.1 Introduction

The ever growing threat environment has made it a challenge to ensure security of the largest network infrastructure, i.e., the internet, for which BGP is the most dominant inter-domain routing protocol. A number of security solutions have been proposed for securing BGP that provide necessary security for the inherently vulnerable protocol. However, the high performance cost incurred by the proposed security solutions renders them unfeasible for practical deployment. Thus there is a need to identify the essential security and performance requirements for BGP, and devise a comprehensive deployment strategy for BGP, so that the benefits of a secure routing infrastructure can be fully achieved. This chapter briefly describes the conducted research work on BGP security solutions as explained in detail in previous chapters. Objectives achieved during this research are explained and directions to continue this work in future are given.

8.2 Overview of Research Work

BGP is widely deployed inter-domain routing protocol on the internet but its design assumes that only trusted networks are operating on the internet. This assumption does not hold true today because internet today has taken a shape similar to a mesh of networks, where threats are posed by both insiders and outsiders. This research work is based on the threats to BGP and the security solutions proposed to secure the protocol.

The first phase of this research work was to study the BGP protocol working and formulate the threat model for BGP. Internet infrastructure is constantly under threat of

attacks because of the weak design of BGP. Malicious attacks on BGP motivated research community to focus their attention on its security. Some countermeasures to protect against BGP mis-configurations and attacks have already been proposed but currently deployed BGP security practices do not provide effective solution to mitigate threats to BGP. This is explained in Chapter 2 & 3.

The next phase was to conduct security and performance analysis of the proposed BGP security architectures. This included the analysis of all major BGP security extensions including S-BGP, SoBGP, IRV, PsBGP, BGPSec and some other solutions. The detailed examination is discussed in Chapter 4. This analysis led to formulation of the evaluation model for BGP. As BGP's huge base is installed on the internet so an obvious requirement from new solution is to be backward compatible with existing BGP in order to facilitate incremental deployment. Every presented solution represents some limitations either in terms of security or usability. Thus a comprehensive set of security and performance requirements has been established in Chapter 5.

The BGP security architectures were evaluated in the next phase. The solutions were reviewed based on the formulated evaluation model for BGP. The detailed analysis and reasoning is presented in Chapter 6. Analysis showed that there is a need for the security solutions to be deployed, because of the wide utility and growing dependence on the internet infrastructure. Thus a deployment strategy for BGP was devised in next phase that provides the security of BGPSec and practical feasibility for deployment. The detailed account of the deployment strategy is made in Chapter 7.

8.3 Achievements

This aim of this research work was to study the Border Gateway Protocol, the existing solutions to secure the protocol and the associated security and performance

requirements. The latest security extension BGPSec was also explored in detail, along with previously presented solutions including S-BGP, SoBGP, IRV, PsBGP, Hop integrity protocols, Invalid MOAS conflicts detection, Path authentication solution and SPV. Security weaknesses in these protocols were identified and the performance analysis of each solution was made.

Based on the findings, an evaluation model was established, which is complete in all aspects of the desired security and performance requirements for BGP. This evaluation model defines comprehensive evaluation criteria for assessing security and performance features of BGP security solutions. All existing solutions were evaluated against this model, highlighting major security weaknesses and performance drawbacks. A new deployment strategy for BGP is proposed, that achieves the best of the security and performance features for BGP security solutions. The proposed deployment strategy for BGPSec achieves many advantages over other security solutions for operational working of BGP. It provides all security services offered by BGPSec. The main advantage of the proposed strategy is that the computational and storage overhead is off-boarded from the existing BGP border routers. This is a major achievement since changing the hardware specification of existing border BGP routers has remained the main challenge for successful deployment of BGPSec. The inclusion of a single dedicated CV&CS server with adequate computational and storage capabilities within a single Autonomous System is more feasible than replacing each border BGP routers to handle all this load. Moreover, this server also provides benefits of pre-computing BGP_Path_Signatures for the peers. This greatly helps in minimizing the BGP convergence speed especially in case BGP router failure-recovery scenarios.

8.4 Future Work

This research work can be extended in many directions in future. For the proposed deployment strategy, hardware platform for proposed “Cryptographic Validation and Computation Server” can be designed. This work should focus on defining minimum hardware resources required for the said purpose keeping in the view the internet growth tolerance and adoptability of IPv6 address space. In this future work, consideration should be focused on proposing and defining high availability design for CV&CS server. The research work presented in this thesis can also be effectively utilized to design a detailed protocol for secure communication between border BGPSec routers and CV&CS, design database architecture to work with CV&CS and the structure of data indexing for fast processing.

8.5 Conclusion

This chapter presents an absolute overview of the work carried out during the course of this research work. The major findings and goals achieved have been highlighted. The security and efficiency parameters accomplished by the proposed deployment strategy are discussed. In the end, suggestions for continuing this work in future are presented. This concludes the research and thesis work.

BIBLIOGRAPHY

[1] Rekhter, Y., Li, Ti. A Border Gateway Protocol 4 (BGP-4).Internet Engineering Task Force RFC 1771. (1995)

[2] Rekhter, Y., Li, Ti., Hares, S. A Border Gateway Protocol 4 (BGP-4). Internet Engineering Task Force RFC 4271. (2006)

[3] Rekhter, Y., Gross, P. Application of the Border Gateway Protocol in the Internet. Internet Engineering Task Force RFC 1772. (1995)

[4] Hawkinson, J., Bates, T. Guidelines for Creation, Selection, and Registration of an Autonomous System (AS). Internet Engineering Task Force RFC 1930. (1996)

[5] Mitchell, J. Autonomous System (AS) Reservation for Private Use. Internet Engineering Task Force RFC 6996. (2013)

[6] Postel, J. Internet Protocol. Information Sciences Institute, University of Southern California. Internet Engineerign Task Force RFC 791. (1981)

[7] Malkin, G. RIP Version 2. Internet Engineering Task Force RFC 1723. (1994)

[8] Malkin, G. RIP Version 2. Internet Engineering Task Force RFC 2453. (1998)

[9] Moy, J. OSPF Version 2. Internet Engineering Task Force RFC 2178. (1998)

[10] Moy, J. OSPF Version 2. Internet Engineering Task Force RFC 2328. (1998)

[11] Callon, R. Use of OSI IS-IS for Routing in TCP/IP and Dual Environments. Internet Engineering Task Force RFC 1195. (1990)

[12] Fu, Y., Guo, G., Huang, T. S., Interior Gateway Routing Protocol. Journal of Computers and Operations Research, Elsevier. (2012)

[13] Hubbard, K., Kosters, M., Conrad, D., Karrenberg, D., Postel, J. Internet Registry IP Allocation Guidelines. Internet Engineering Task Force RFC 2050. (1996)

[14] Rekhter, Y., Li, T., Hares, S. A Border Gateway Protocol 4 (BGP-4). Internet Engineering Task Force RFC 4271. (2006)

[15] Vohra,Q., Chen, E.BGP Support for Four-Octet Autonomous System (AS) Number Space. Internet Engineering Task Force RFC 6793. (2012)

[16] Chen, E., Yuan, J. Autonomous-System-Wide Unique BGP Identifier for BGP-4. Internet engineering Task Force RFC 6286. (2011)

[17] Butler, k., Farley, T., Mcdaniel, P., Rexford, J. A Survey of BGP Security. AT&T Labs Research, Vol. V, No. N. (2005)

[18] Butler, k., Farley, T., Mcdaniel, P., Rexford, J. A Survey of BGP Security Issues and Solutions. Proceedings of the IEEE|Vol.98,No.1. (2010)

[19] Gill, P., Schapira, M., Goldberg, S. A Survey of Interdomain Routing Policies. Journal of SIGCOMM Comput. Commun. Rev. Vol. 44, No. 1, pg. 28-34. (2014)

[20] Heffernan, A. Protection of BGP Sessions via The TCP MD5 Signature Option. Internet Engineering Task Force RFC 2385. (1998)

[21] Caesar, M., Rexford, J. BGP Routing Policies in ISP Networks. Network, IEEE , Vol.19, Issue. 6, pg. 5-11. (2005)

[22] Chandra, R., Traina, P., and Li, T. BGP Community Attribute. Internet Engineering Task Force RFC 1997. (1996)

[23] Villamizar, C., Chandra, R., Govindan, R. BGP Route Flap Damping. Internet Engineering Task Force RFC 2349. (1998)

[24] Handley,M., Rescorla, E. Internet Denial-of-Service Considerations. Internet engineering Task Force RFC 4732. (2006)

[25] Fuller, V., Li, T., Yu, J. Varadhan, K. Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy. Internet Engineering Task Force RFC 1519. (1993)

[26] Fuller, V., Li, T. Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan. Internet Engineering Task Force RFC 4632. (2006)

[27] Murphy, S. BGP Security Vulnerabilities Analysis. Internet Engineering Task Force RFC 4272. (2006)

[28] Butler, K., Farley, T., McDaniel, P., Rexford, J. A Survey of BGP Security Issues and Solutions. Citeseer Draft. (2008)

[29] Subramanian, L. Decentralized Security Mechanisms for Routing Protocols. PHD Thesis, University of California, Berkeley. (2005)

[30] Watson, P. Slipping in the Window: TCP Reset Attacks. Technical White Paper. www.terrorist.net. (2003)

[31] Huston, G., Rossi, M., Armitage, G. Securing BGP — A Literature Survey. Communications Surveys & Tutorials, IEEE, Vol: 13, Issue: 2. (2011)

[32] Mahajan, R., Wetherall, D., and Anderson, T. Understanding BGP Misconfiguration. ACM SIGCOMM, Pittsburgh, PA, USA. (2002)

[33] Cetinkaya, E.K., Sterbenz, J.P.G.A Taxonomy of Network Challenges. Proceedings of 9th International Conference on the Design of Reliable Communication Networks (DRCN), IEEE. (2013)

[34] Wan, T., Paul, C., Oorschot,V. Analysis of BGP Prefix Origins During Google's May 2005 Outage. 20th International Parallel and Distributed Processing Symposium, IPDPS. (2006)

[35] Wan, T., Kranakis, E., Oorschot, V. Pretty Secure BGP (psBGP). In Proceedings of Network and Istributed Systems Security, Internet Society (ISOC), San Diego, CA. (2005)

[36] Mujtaba, M. Analysis of Intrusion Detection System (IDS) in Border Gateway Protocol. MSc Thesis, Faculty of Engineering and Information Technology. University of Technology, Sydney. (2012)

[37] Wang, X., Feng, D., Lai, X., Yu, H. Collisions for Hash Functions MD4, MD5, HAVAL-128 and RIPEMD. Cryptology ePrint Archive, Report 2004/199, IACR. (2004)

[38] Hao, F., Kodialam, M.S., Song, H. Hash Functions for Applications Such As Network Address Lookup. Google Patents, US Patent 7,990,973. (2011)

[39] Rivest, R. The MD5 Message-Digest Algorithm. Internet Engineering Task Force RFC 1321. (1992)

[40] Turner, S., Chen, L. Updated Security Considerations for the MD5 Message-Digest and the HMAC-MD5 Algorithms. Internet Engineering Task Force RFC 6151. (2011)

[41] Hethmon, P., Elz, R. Feature Negotiation Mechanism for the File Transfer Protocol. Internet Engineering Task Force RFC 2389. (1998)

[42] Touch, J., Mankin, A., Bonica, R.The TCP Authentication Option. Internet Engineering Task Force RFC 5925. (2010)

[43] Borman, D. TCP Options and Maximum Segment Size (MSS). Internet Engineering Task Force RFC 6691. (2012)

[44] Kent, S., Atkinson, R. IP Authentication Header. Internet Engineering Task Force RFC 2402. (1998)

[45] Housley, R. Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH). Internet Engineering Task Force RFC 4835. (2007)

[46] Kent, S., Atkinson, R. IP Encapsulating Security Payload. Internet engineering Task Force RFC 2406. (1998)

[47] Kent, S. and Atkinson, R. Security architecture for the Internet Protocol. Internet Engineering Task Force RFC 2401. (1998)

[48] Kent, S., Seo, K. Security Architecture for the Internet Protocol. Internet Engineering Task Force RFC 4301. (2005)

[49] Hoffman, P. Algorithms for Internet Key Exchange Version 1 (IKEv1). Internet Engineering Task Force RFC 4109. (2005)

[50] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P.Internet Key Exchange Protocol Version 2 (IKEv2). Internet Engineering Task Force RFC 5996. (2010)

[51] Gill, V., Heasley, J., Meyer, D. The Generalized TTL Security Mechanism (GTSM). Internet Engineering Task Force RFC 3682. (2004)

[52] Greene, B. BGPv4 Security Essentials version 0.5. (2004)

[53] Team Cymru. Lake Mary, FL 32746, US. https://www.team-cymru.com/

[54] Alaettinoglu, C., Villamizar, C., Gerich, E., Kessens, D., Meyer, D., Bates, T., Karrenberg, D., Terpstra, M. Routing Policy Specification Language (RPSL). Internet Engineering Task Force RFC 2622. (1999)

[55] Blunk, L., Damas, J., Parent, F., Robachevsky, A. Routing Policy Specification Language next generation (RPSLng). Internet engineering Task Force RFC 4012. (2005)

[56] APNIC. Asica Pacific Network Information Centre. https://www.apnic.net, http://www.apnic.net/services/services-apnic-provides/resource-certification

[57] RIPE. European IP Networks. http://www.ripe.net/, http://www.ripe.net/lir-services/resource-management/certification

[58] RADB. Routing Assess Database, Merit Networks INC. http://www.ra.net/

[59] ISC, Internet Systems Consortium. http://www.isc.org/

[60] ISC IRRToolSet. Internet Systems Consortium, Internet Routing Registry Tool set. http://irrtoolset.isc.org/

[61] Ylonen, T., Lonvick, C. The Secure Shell (SSH) Protocol Architecture. Internet Engineering Task Force RFC 4251. (2006)

[62] Renesys, The Internet Intelligence Authority. http://www.renesys.com.

[63] Benkis, L. Practical BGP Security: Architecture, Techniques and Tools, White Paper. Renesys. (2005)

[64] Kent, S., Lynn, C., and Seo, K. Secure Border Gateway Protocol (S-BGP). IEEE Journal on Selected Areas in Communications Vol. 18, No. 4. (2000)

[65] Glenn, R., Kent, S. The NULL Encryption Algorithm and Its Use With IPsec. Internet Engineering Task Force RFC 2410. (1998)

[66] Goodell, G., Aiello,W., Griffin, T., Ioannidis, J., McDaniel, P., and Rubin, A. Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing. ISOC NDSS'03, San Diego, CA, USA, pg. 75–85. (2003)

[67] Dierks, T., Rescorla, E.,The Transport Layer Security (TLS) Protocol Version 1.2. Internet engineering Task Force RFC 5246. (2008)

[68] Nicol, D., Smith, S., and Zhao, M. Efficient Security for BGP Route Announcements. Dartmouth Computer Science Technical Report TR-2003-440. (2002)

[69] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., Polk, W.Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. Internet Engineering Task Force RFC 5280. (2008)

[70] Yee, P. Updates to the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. Internet Engineering Task Force RFC 6818. (2013)

[71] Weiler, S., Ward, D., Housley, R. The rsync URI Scheme. Internet Engineering Task Force RFC 5781. (2010)

[72] Gouda, M. G., Elnozahy, E. N., Huang, C.-T., and McGuire, T. M. Hop Integrity in Computer Networks. 8th International Conference on Network Protocols, Osaka, Japan. (2000)

[73] Zhao, X., Pei, D., Wang, L., Massey, D., Mankin, A., Wu, S. F., and Zhang, L. An analysis of BGP Multiple Origin AS (MOAS) Conflicts. ACM SIGCOMM Internet Measurement Workshop, San Francisco, CA, USA. (2001)

[74] Farley, T., Mcdaniel, P., Butler, K. A Survey of BGP Security Issues and Solutions. AT&T Labs - Research, Florham Park, NJ. (2004)

[75] Bates,T., Bush, R., Li, T., Rekhter, Y. DNS-based NLRI Origin AS Verification in BGP. Internet Engineering Task Force Internet Draft. (1998)

[76] Kruegel, C., Mutz, D., Robertson, W., Valeur, F. Topology-Based Detection of Anomalous BGP Messages. Proceedings of 6th International Symposium, RAID Pittsburgh, PA, USA. (2003)

[77] Hong, S. C., Ju, H., Hong, J. W. K. Network Reachability-based IP Prefix Hijacking Detection. International Journal of Network Management, Vol. 23, Issue 1, pg. 1–15. (2013)

[78] Hu, Y., Perrig, A., Johnson, D. Efficient Security Mechanisms for Routing Protocols. Internet Society Network and Distributed Systems Security, San Diego, CA, USA. (2003)

[79] Perrig, A., Canetti, R., Tygar, JD., Song, D. The TESLA Broadcast Authentication Protocol. RSA Laboratories CryptoBytes Newsletter, vol. 5. (2005)

[80] Hu, Y.-C., Perrig, A., Sirbu, M. SPV: Secure Path Vector Routing for Securing BGP. In ACM SIGCOMM. (2004)

[81] Wong, C., and Lam, S. Digital Signatures for Flows and Multicasts. IEEE/ACM Transactions on Networking Vol.7, Issue: 4, Pg. 502–513. (1999)

[82] Shaw, E., Ruby, K., Post, J. The Insider Threat to Information Systems, The Psychology of the Dangerous Insider. (2007)

[83] Kent, S., Lynn, C., Mikkelson, J., Seo, K.. Secure Border Gateway Protocol (S-BGP) Real World Performance and Deployment Issues. ISOC Symposium on Network and Distributed System Security. (2000)

[84] Kent, S. Securing the Border Gateway Protocol: A Status Update. 7th IFIP TC-6 TC-11 Conference on Communications and Multimedia Security, Torino, Italy. (2003)

[85] Lepinski, M., Turner, S. An Overview of BGPSEC. Internet Engineering Task Force Internet Draft. (2014)

[86] Lepinski, M. BGPSEC Protocol Specification. Internet Engineering Task Force Internet Draft. (2013)

[87] BGP Secure Routing Extension (BGP-SRx). National Institute of Standards and Technology. http://www-x.antd.nist.gov/bgpsrx/

[88] Cryptographic Acceleration for Border Gateway Protocol Security (BGPSEC). DOC-NIST – 2014-NIST-SBIR-01. (2014)

[89] Sriram, K. BGPSEC Design Choices and Summary of Supporting Discussions. Internet Engineering Task Force Internet Draft. (2014)

[90] Deering, S., Hinden, R. Internet Protocol, Version 6 (IPv6) Specification. Internet Engineering Task Force RFC 2460. (1998)

[91] Nordstrom, O., Dovrolis, C. Beware of BGP Attacks. ACM SIGCOMM Computer Communications Review, Vol. 34, No. 2. (2004)

[92] Bellovin, S., Bush, R., Ward., D. Security Requirements for BGP Path Validation. Internet Engineering Task Force, Internet Draft. (2014)

[93] AFRINIC RPKI. African Network Information Centre. http://www.afrinic.net/en/initiatives/resource-certification

[94] ARIN RPKI. American Registry for Internet Numbers. https://www.arin.net/resources/rpki/index.html

[95] LACNIC RPKI. Latin America and Caribbean Network Information Centre. http://www.lacnic.net/en/web/lacnic/certificacion-de-recursos-rpki

[96] Sriram, K. RIB Size Estimation for BGPSEC, Information Technology Laboratory, National Institute of Standards and Technology. (2011)