

Concordia University College of Alberta
Master of Information Systems Security Management (MISSM) Program
7128 Ada Boulevard, Edmonton, AB
Canada T5B 4E4

Study of BGP Security Issues and Technique for AS Route Validation

by

SIDDIQI, Abid Jamal

A research paper submitted in partial fulfillment of the requirements for the degree of

Master of Information Systems Security Management

Date: December 2008

Research advisors:

Pavol Zavarsky, Director of Research and Associate Professor, MISSM

Dale Lindskog, Assistant Professor, MISSM

Study of BGP Security Issues and Technique for AS Route Validation

by

SIDDIQI, Abid Jamal

Research advisors:

Pavol Zavarsky, Director of Research and Associate Professor, MISSM

Dale Lindskog, Associate Professor, MISSM

Reviews Committee:

Andy Ignor, Assistant Professor, MISSM

Dale Lindskog, Assistant Professor, MISSM

Ron Ruhl, Assistant Professor, MISSM

Pavol Zavarsky, Associate Professor, MISSM

The author reserve all rights to the work unless (a) specifically stated otherwise or (b) refers to referenced material the right to which is reserved by the so referenced authors.

The author acknowledges the significant contributions to the work by Academic Advisors and Review Committee Members and gives the right to Concordia Univeristy College to reproduce the work for the Concordia Library, Concordia Websites and Concordia MISSM classes.

Concordia University College of Alberta

Masters of Information System Security Management (MISSM) Program

7128 Ada Boulevard, Edmonton, AB. T5B 4E4, Canada

Study of BGP Security Issues and Technique for AS Route Validation

A research paper submitted in partial fulfillment of the requirements for the degree of

Master of Information Systems Security Management

November 2008

Prepared By:

Abid Jamal Siddiqi
asiddiqi@csa.concordia.ab.ca

Research Advisors:

Dr. Dale Lindskog, Associate Professor, MISSM

Pavol Zavorsky, Director of Research and Associate Professor, MISSM

Draft No: 4.0, November 2008.

Table of Content:

Abstract	3
1 Introduction	3
2 Definition and Terminology	4
2.1 BGP Attributes Types	4
3 Risks to Border Gateway Protocol	6
3.1 Limitations of BGP	6
3.2 Blackholing	6
3.3 Traffic Redirection / Interception	6
3.4 Instability in inter-domain	7
3.5 Denial of Service to Service Provider and Customer	7
3.6 Prefix Hijacking	7
4 Route Invalidity Due to Configuration Faults	7
4.1 Wrong AS_PATH Announcement	8
4.2 Poor Access Control	9
4.3 Uncertain Community	9
4.4 Miscalculated Summery/Aggregate Address	9
5 Route Validations and Filtering	10
5.1 Protecting the Internal Infrastructure of BGP Network	10
5.2 Discarding Private and Well-Known	11
5.3 Filtering Unallocated Address	13
5.4 Validating the New Entries	13
5.5 Filter Customers	15
5.6 Checking for the Attribute	16
5.7 Controlling Advertisement	16
6 Best Practices	17
7 Related Work	18
8 Conclusion	19
References	19

Abstract

BGP (Border Gateway Routing Protocol) is the only global internet routing protocol is the only global Internet routing protocol. Unfortunately BGP lacks the fundamental authentication or validation, and for this reason false Internet routes can be easily introduced into the routing infrastructure, either accidentally or by malicious attack. This paper identifies the misconfigurations that could cause harm to routing information; it also suggests the policy, configuration best practices and consideration which would protect against bogus routes announced in BGP internal infrastructure and the Internet. Rather than suggesting an entire protocol replacement, or cryptographic solution based on encryption, this paper concentrates on the current implementation of BGP.

1. Introduction

There are many routing protocols available today, and these protocols are divided in to two major categories the intradomain (internal) category and interdomain (external) category. There are many interior routing protocols to connect the internal networks, and there one an exterior routing protocol (BGP). BGP is a successor of the Exterior Gateway Protocol (EGP), which had some serious limitations. EGP was built for a backbone centered tree and was unable to accommodate the expansion of the Internet. The current version of BGP is 4 (also written as BGP-4). BGP-4 was launched in March 1995 for Classless Inter-Domain Routing (CIDR), running over TCP (Transmission Control Protocol) and uses port number 179 [1]. The use of TCP eliminates the need for retransmission and acknowledgement in the BGP protocol.

BGP looks at the entire Internet as a collection of connected Autonomous Systems. An Autonomous System is a collection of networks with the same routing policy a single management perhaps under one ownership and administrative control, that uses a also use single routing protocol. BGP supports two types of routing internal and external routing. External routing (EBGP) refers to exchanges between different ASes

(Autonomous Systems) and internal routing (IBGP) is used by peers within an AS, BGP peers are BGP neighbour connected to each other.

BGP is classified as a path vector protocol; a path vector protocol defines a route to a destination as a pairing between the destination and the attributes of the path to the destination [2].

2. Definitions and Terminology

BGP route contain a list of Autonomous Systems, known AS paths, along with a set of IP address prefixes reachable from that AS path. BGP uses various types of messages to communicate with its peers; it also uses many attributes for tuning and selecting path to the destination. BGP peer discover paths from internal and external speakers it picks up the best path and updates the forwarding table.

2.1 BGP Attributes Types

BGP attributes describe paths and helps to choose the best route among them [3] which are Local preference, Multi-Exit Discriminator (metric), Origin, AS_PATH, NEXT_HOP, and Community.

- *LOCAL PREFERENCE* is used inside an AS to favour a confident exit point from the AS. If there is more than one exit point within the AS, the one with highest local preference will be chosen.
- A *Multi-Exit Discriminator (MED)* is a suggestion to an external AS regarding the entry point to the originating AS. It is a suggestion since the external AS may be using other attributes when choosing the route to the AS.
- The *ORIGIN* attribute specifies where or how the routes were learned from. There are three different values for this attribute: An IGP route indicates that the route is interior to the AS, EGP indicates that the route was learned via EGP from another AS, and Incomplete indicates that the origin is unknown.
- The *AS_PATH* attribute: when a route passes through an AS, that AS adds its own AS number to the list of AS numbers.

- The *NEXT_HOP* attribute is in EBGp, and indicates that the IP address used to reach the advertising router.
- The *COMMUNITY* attribute can be used to group communities, or destinations. These groups can be used to apply routing decisions differently. For example *NO_EXPORT* indicates the route is not advertised to BGP peers outside a confederation [4], *NO_ADVERTISE* indicates that the route is not advertised to any peer, and *NO_EXPORT_SUBCONFED* indicates that the route must not be advertised to any EBGp peer [5].

If no routing policy is in effect, the default BGP routing process consults the attribute list below to decide which route is best to a destination. Unlike other routing protocols, which depend on bandwidth and delay BGP breaks the tie by comparing attributes sequentially, by using the table 1.1 below which should be kept in the routing table, and wherever the tie breaks the routing decision is made.

Route Selection Criteria Table

1. Exclude routes with unreachable next hop
2. Prefer routes with highest weight (CISCO proprietary)
3. Prefer routes with local preference
4. Prefer routes locally originated by the router
5. Prefer routes with the shortest As-path
6. Prefer routes with lowest origin code (IGP<EGP<Incomplete)
7. Prefer routes with lowest MED
8. Prefer EBGp routes over IBGP
9. Prefer nearest IGP neighbour for IBGP path
10. Prefer oldest path in case of EBGp
11. Prefer route from router with lowest BGP router ID

Table 1.1

3 Risks to Border Gateway Protocol

3.1 Limitations of BGP

BGP does not guarantee the validity of the path attributes declared by an AS. On the contrary, path attributes are one way that a malicious AS can damage or disturb the routing infrastructure. Plus, study of BGP activities on Internet demonstrates that the routing can exhibit out of order behaviour. For example, packets originating in Canada and destined for US, may be mistakenly routed through the UK.

BGP does not have the mechanism to validate that the route announced by the AS is actually owned by that AS or not. An AS can announce that it has the shortest path to a destination by forging the path vector, even if it is not part of the destination path at all. Therefore, BGP does not guard the integrity, newness and origin authentication of messages (Integrity guarantees that a message has not been tampered with; newness guarantees that the receiver has really received a new message, or the replayed one, and origin authentication confirms that the originator of the update message is authentic).

3.2 Blackholing

Blackholing occurs when an AS takes responsibility to serve a prefix that they actually don't have a route to, or when a specific destination is unreachable from a large section of the Internet. Legitimate blackhole routing is used to implement private and non-allocated IP ranges. Malicious blackholing refers to bogus route announcement intended to attract specific routes and then drop the traffic.

3.3 Traffic Redirection / Interception

Traffic Redirection/Interception occurs when an organization owning a specific network is difficult to take a deferent path, and to reach a wrong, maybe compromised destination. One purpose of redirection attacks is for the compromised destination to impersonate the true destination, to obtain secret information. In these types of attacks, the organization is still forwarded to the correct destination, making the attack more difficult to distinguish.

3.4 Instability in inter-domain

Instability in inter-domain routing can be caused by consecutive advertisements (using various attributes) and removal of routes. The purpose of such attacks can be to activate route dampening in upstream routers, and thus originate connectivity loss.

3.5 Denial of Service to Service Provider and Customer

As BGP runs over TCP, so it inherits all the weakness of TCP, such as TCP syn-flood against port 179. ACLs are used to protect these type of attacks, by applying controls to incoming traffic, but once the source address is spoofed, ACLs cannot help because ACLs only filters on the source IP address of the BGP peer. Once spoofed, the packet passes through the ACL, and ends up flooding the resources of the hardware, i.e. routers, e.g. the input queuing mechanism, forcing the processor to reach the spot where control plane packets starts dropped.

3.6 Prefix Hijacking

Prefix Hijacking is a type of attack in which an AS announces to its peers a block of addresses that do not belongs to them. It can also allow a malicious entity to access unreachable IP-address space which may be declared for a deliberate attack against a certain organization, or might be a result of misconfiguration. Most recently an incident of this nature was, in Feb 2008, launched against YouTube [9] [10].

4. Route Invalidity due to Configuration Fault

Configuring routing protocols is not a simple job, since it involves numerous parameters and policy settings. A slight mistake can be difficult to identify, and can disturb the whole Internet communication. The filter policies and configurations between peers is very complex and these configurations can only be seen by the border routers. Therefore, configuration faults even from a very small ISP can cause large portions of the Internet to be out of order for a long period. Here are a number of examples:

1. In April 1997, AS 7007, a small ISP announced misconfigured routes which broke down the whole Internet for one day.

2. In December 1999, AT&T traffic was misdirected, due to the announcement of their network by another ISP.
3. In December 2004, an accidental misconfiguration misdirected the traffic for CITY Corp and other companies to Turkey [18].
4. Again in September 2005, AT&T traffic was misdirected to Bolivia.
5. In January 2006, PANIX traffic was misdirected to a New York ISP.
6. Most recently, in Feb 2008 Pakistan Telecommunication blocked YouTube all over the Internet. Their intention was to block it in the country, but YouTube was inaccessible for more than an hour around the world. All these incidents were due to mistakes and errors but, everything probable through human error can also be done by intention. For that reason, to control and check configuration is very important to protect our self and others.

4.1 Wrong AS_PATH Announcement

BGP Operators can tweak BGP AS_PATH attribute to influence the traffic coming in or going out from their network. This could be done with the help of AS prepending, which allows an operator to add AS_PATH to the routes being announced to others. After consulting the path selection criteria (table 1.1) the receiving peers will prefer the routes with the shortest AS_PATH. While prepending the AS_PATH operators should be very careful as any AS number can be entered for prepending (eg.6000). If the added AS exists on the Internet it will cause a conflict with the original AS in the internet, which can effect the internet routing as the receiver will try to direct or reply to original AS, and if the original AS does not know about requesting network it will drop the packets. There are two main categories of AS numbers, public AS numbers (1-64511) and private AS numbers (64512-65535) in addition to the discussion private numbers should also not be propagated on the internet, many well known vendors gives us the option to remove the private AS number from the out going advertisement e.g. a simple Cisco IOS command “#neighbour x.x.x.x. remove-private-AS” can do the job for us.

4.2 Poor Access Control

ACLs and route-maps are used to implement BGP routing policies, permitting and denying traffic with ACLs can be very complex, and mistakes in filtering IP addresses of neighbours are a common problem. Configurations Keywords “in” and “out” are used to prevent or allow incoming and outgoing traffic on internal and external interfaces. A single miscalculation of address could allow unwanted traffic and smallest mistake in applying the mask may block entire allowed network. So there are many things to consider when using ACLs [7]. The order of statements in an access control list should be clearly entered keeping in mind that there is an implicit denial at the end, always be careful when using the mask to permit or deny. According to the requirement the specific internal or external should be identified, with permitting or denying traffic coming in or going out, slight misunderstanding changes every thing a minor mistake could cause opposite effect.

4.3 Uncertain Community

A community is a group of prefixes that share some common property and can be configured with the BGP community attribute [12], it is used by the service providers for various purposes, such as filter out traffic, mark the routes for routing policy such as should the route be advertised out of the AS or should the receiving routes be propagated to another AS. When attaching community with routes administrator should be completely aware of the architecture and requirement, because it becomes very complicated in large infrastructure and a wrong propagation or misconfiguration can cause huge effect, resulting in propagation of prefix beyond the required boundaries, losing substantial information, or it may not propagate where the intended propagation should reach.

4.4 Miscalculated Aggregate Address

BGP is not a secure routing protocol and numerous incidences discussed above have caused problem in the internet backbone. The issue is not just that BGP lacks the mechanisms to verify the validity of an advertisement route but misconfigurations has also played a big part in it weakness, especially regarding the address aggregation. Once

again discussing the YouTube [9] [10] incidence which was mainly due to incorrectly advertisement of a single /24 prefix which belonged to YouTube's /22 network (this /24 prefix held the DNS servers of YouTube). As a matter of fact that BGP depends on longest match, for that reason routers on the Internet redirected packets towards the announced /24 prefix where they were dropped as it became impossible for youtube.com to resolve names, hence was unreachable for 2 hours, for more information visit NANOG mailing list [11].

5. Route Validations and Filtering

Researchers are working to find the solution for the problem faced by BGP such as in 2002, Xiaoling Zhao worked on the Multiple Origin AS (MOAS) which describes a protocol enhancement to detect bogus BGP routes [13], another researcher Geoffrey Goodell proposed a new protocol that works with BGP to detect false and malicious routing information [14], but it will take time for the whole internet to adopt an optimized or a standard solutions.

In BGP the mechanism for validating a prefix, and authenticating the originator of that prefix is not present. Service providers and other operators are using there own methods for filtering internet traffic and validating the internet routes. The following section presents the technique for validating routes and securing advertisements by filtering.

5.1 Protecting the Internal Infrastructure of BGP Network

Protection is the most essential part of any network design. This portion of the paper focuses on the deployment of edge infrastructure protection. Internal infrastructure could include addresses that are not accessible by the public Internet, and internal services which should not be seen or heard by the outsiders, such as SNMP services, IBGP connection, remote access services (telnet, SSH) to other network devices and etc. To completely secure the internal infrastructure, both control plane and forwarding plane should be protected. ACLs can be applied to the boundary clearly blocking all traffic intended for internal address space and services.

Always test the requirement and design in a lab environment before deploying. The ACL should be applied on outgoing and incoming interfaces connected to peers/providers and

costumers. Clearly categorize the infrastructure with all the necessary/authorized protocols that are compulsory including all addresses that are used for the internal network and are seldom accessed by outside sources such as router interfaces, critical services, and point-to-point link address. Assure that the accurate permit/deny statements are in place. Anti-spoof filters should be used i.e. packet originated from another AS should not have your address. Summarization is vital, these addresses should be clustered anywhere best possible into CIDR classless interdomain routing blocks.

5.2 Discarding Private and Well-Known

When considering the security of BGP networks the main option present today is through filtering the network traffic coming in and going out of the network, depending on the type of organization such as a major ISP or a corporation large enough that require the need of BGP implementation. This part of the paper suggests how to validate route and filter the unwanted. Preferably, the peering policy should be specific so that exact filters can be put in place, all address coming in the and going out of the network that are considered private [15] should be filtered.

Discard RFC1918 etc prefixes

Discard your own prefix on the interfaces

Don't accept default unless required

Discard prefixes from ISPs you do not have membership with

10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16 private addresses RFC 1918
127.0.0.0/8 Host loopback
0.0.0.0/8 and 0.0.0.0/32 broadcast/default
192.0.2.0/24 for testing purpose
169.254.0.0/16 DHCP node auto configuration

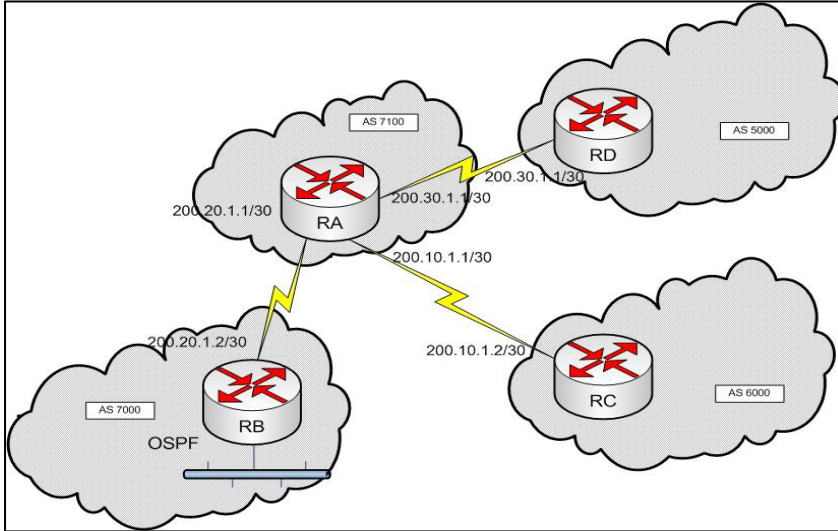


Figure 1.1

In Figure 1.1 RA should use filtering on the interfaces connected to RB, RC and RD on the AS border for all inbound and outbound advertisements to protect our self and the peers.

```

router bgp 7100
no synchronization
no auto-summary
neighbor 200.30.1.1 remote-as 6000
neighbor 200.30.1.1 version 4
neighbor 200.30.1.1 prefix-list DUSA in
neighbor 200.30.1.1 prefix-list DUSA out
neighbor 200.10.1.1 remote-as 5000
neighbor 200.10.1.1 version 4
neighbor 200.10.1.1 prefix-list DUSA in
neighbor 200.10.1.1 prefix-list DUSA out
neighbor 200.20.1.1 remote-as 7000
neighbor 200.20.1.1 version 4
neighbor 200.20.1.1 prefix-list DUSA in
neighbor 200.20.1.1 prefix-list DUSA out
ip prefix-list DUSA deny 0.0.0.0/8 le 32
ip prefix-list DUSA deny 10.0.0.0/8 le 32
ip prefix-list DUSA deny 127.0.0.0/8 le 32
ip prefix-list DUSA deny 169.254.0.0/16 le 32
ip prefix-list DUSA deny 172.16.0.0/12 le 32
ip prefix-list DUSA deny 192.0.2.0/24 le 32
ip prefix-list DUSA deny 192.168.0.0/16 le 32
ip prefix-list DUSA deny 224.0.0.0/3 le 32
ip prefix-list DUSA deny 0.0.0.0/0 ge 25
ip prefix-list DUSA permit 0.0.0.0/0 le 32

```

5.3 Filtering Unallocated Address

The IP address space reserved, and not been allocated yet by the Internet Assigned Numbers Authority (IANA) or any Regional Internet Registry (RIR), also known as Bogon address [24], should also be validated and filter out. This would be a continuous monitoring and updating task as address are being assigned every day and blocking a laminate address would cause problem. The allocations will be of the size requested from the RIRs (AfriNIC, APNIC, ARIN, LACNIC and RIPE) will be according to the network operators' requirements. The most common practice today to announce /24s and some of the /24 announcements are due to traffic engineering efforts for multihoming [16] [8], However the task of aggregation still depends on the operators .

Here is an example configuration and filter using BGP Prefix-List for unallocated address

```
RIPE [18]
ip prefix-list ALLOCATED permit 62.0.0.0/8 ge 9 le 20
ip prefix-list ALLOCATED permit 80.0.0.0/7 ge 9 le 20
ip prefix-list ALLOCATED permit 193.0.0.0/8 ge 9 le 20
ip prefix-list ALLOCATED permit 194.0.0.0/7 ge 9 le 20
ip prefix-list ALLOCATED permit 212.0.0.0/7 ge 9 le 20
ip prefix-list ALLOCATED permit 217.0.0.0/8 ge 9 le 20

APNIC [19]
ip prefix-list ALLOCATED permit 61.0.0.0/8 ge 9 le 20
ip prefix-list ALLOCATED permit 202.0.0.0/7 ge 9 le 20
ip prefix-list ALLOCATED permit 210.0.0.0/7 ge 9 le 20
ip prefix-list ALLOCATED permit 218.0.0.0/8 ge 9 le 20

ARIN [20]
ip prefix-list ALLOCATED permit 24.0.0.0/8 ge 9 le 20
ip prefix-list ALLOCATED permit 63.0.0.0/8 ge 9 le 20
ip prefix-list ALLOCATED permit 64.0.0.0/6 ge 9 le 20
ip prefix-list ALLOCATED permit 199.0.0.0/8 ge 9 le 20
ip prefix-list ALLOCATED permit 200.0.0.0/8 ge 9 le 20
ip prefix-list ALLOCATED permit 204.0.0.0/6 ge 9 le 20
ip prefix-list ALLOCATED permit 208.0.0.0/7 ge 9 le 20
ip prefix-list ALLOCATED permit 216.0.0.0/8 ge 9 le 20
```

5.4 Validating the New Entries

Every BGP network operators maintains there own list of every prefix they announce, and every network they have adjacencies with. All the new routes passed through policy and configured ACLs are trusted and accepted as valid, but a prefix that is spoofed can easily penetrate the filter. So after all the filtering discussed above, it would not guarantee

security, still large possibility of malicious route entry will be there. To work around this problem we should consult the IRR Internet Routing Registry [17] database to validate routing information; its purpose is to ensure stability and consistency of the Internet-wide routing by sharing information between network operators. These public registries contains information for networks of hundreds of organizations including ISPs, universities, and business enterprises who publicly register their routing policy and route announcements to their database to facilitate the operation of the Internet. Registering to these Registries and consulting these databases to generate access lists, troubleshoot routing problems, automatically configure backbone routers, and validate routes would be the answer to a certain extent.

There design Prefix list filters that contains all the prefixes of authentic peer (trusted or believed routes) and the entire transitive downstream peer in the IRR. Stable peers for along time should be assumed valid, discard any new origination by those peers which cannot be confirmed in an IRR. AS based filtering in this case would not give granularity as regular expression would not check each and every prefix therefore filtering should be done prefix based for every peer.




Figure 1.2

In figure 1.2 RA is configured base or IRR information


```

neighbor 200.20.2.2 peer-group
neighbor 200.20.2.2 soft-reconfiguration inbound
neighbor 200.20.2.2 update-source Loopback0
neighbor 200.20.2.2 next-hop-self
neighbor 200.20.2.2 route-map IRR-IN in
neighbor 200.20.2.2 route-map IRR-OUT out
neighbor 200.20.2.2 remote-as 7000
neighbor 200.20.2.2 peer-group PEER
neighbor 200.20.2.2 ebgp-multihop 255
neighbor 200.20.2.2 prefix-list 7000 in

route-map IRR-IN permit 20
set metric (according to requirement)
set community (according to requirement)

route-map IRR-OUT deny 10
match ip address prefix-list UNALLOCATED ! address considered above

route-map IRR-OUT permit 20
match community
set metric(according to requirement)
set community (according to requirement)
ip prefix-list 7000 permit (allowed networks)

```

5.5 Filter Customers

For an ISP any customer could be and always are a potential threat as they are open to configurationally error that might cause big problem. Not filtering your customers puts your network at risk to. ISPs should only accept prefixes which have been assigned or allocated to their downstream peer/customer, i.e. allowing packets with source addresses belonging to the customer's prefixes and denying packets with any other source address.

Considering figure 1.2 and Configuration Upstream on RA

```

router bgp 100
neighbour 200.20.1.2 remote-as 7000
neighbour 200.20.1.2 prefix-list CUSTOMER in
ip prefix-list CUSTOMER permit (SPECIFIC CUSTOMER ANNOUNCED PREFIX)
ip prefix-list CUSTOMER deny (ALL LOCAL ADDRESS)
ip prefix-list CUSTOMER deny (BOGON ADDRESS)
ip prefix-list CUSTOMER deny (UNALLOCATED ADDRESS)

```

5.6 Checking for the Attribute

By looking at the route in the routing table we can easily say that we know exactly how this route originated, has it been redistributed or entered via network command, is it internal or external, but it may not be true because any administrator can manipulate the attributes setting to influence the routing decision for their network, even they could influence the remote autonomous system for traffic engineering or for malicious activity. It can be done when advertising the routes and also when receiving the routes. Assume that AS 200 advertises a prefix A, setting the origin to “e” (External BGP routes), through ISP X and again the same prefix A with another ISP Y setting the origin to “?” (Redistributed), consulting table 1.1 most of the time parsing through the selection criteria the tie will break at the origin as (e>i>?).

When an ISP or peer receives the routes it will prefer the route with origin with ‘e’. No organization on the internet would want any peer to influence their routing decision therefore when ever receiving the routes from peer check and change the attributes to according to your own requirements for routing decision, never use default let the costumers or peer influence their policy on your routing.

5.7 Controlling Advertisement

It is never wise to advertise routes learned from one external peer to another external peer as it will end up turning your organization into a service provider for the service providers with no financial incentive, plus it will cause the network excessive internet routes, making it vulnerable for all sorts of attack. In BGP to control and influence the route advertisement community attribute is one of the most powerful tool, instead of filters the outgoing advertisements based on individual prefixes, community based filtering can be used. As you receive the prefixes from a peer you can define one or more communities to those prefixes and filter advertisement based on those communities.

A simple solution would be to construct a route map matching the set of prefixes you don't want to be advertised to the external peer, and set the community value to NO-Export (Subsection 2.1 discusses some of the well-known community attributes such as NO-Export) that will restrict the advertisement to be propagated to another EBGP peer.

In conjunction with the community based filtering, the bogon filter discussed above should also be used in both in coming and out going direction. This will protect us as well as our neighbours from excessive and unwanted routing information and I will also guard against the configuration mistakes.

6. Best Practices

- A small ISP do not require the entire internet routing table instead configure default route would be better [22].
- The route information received from, an unfrosted neighbor should be carefully, controlled and monitored. The best and the most secure way are to use static routes.
- Always use MD5 authentication with all BGP peers, using strong password, and never use same password with different peers [23].
- Never, ever establish an IGP session with a router that you do not have authority on, do not let your IGP talk to it even if they are your most trusted partners.
- Always make shore that the routes you are advertising are reachable.
- An ISP can protect themselves against excessive advertisement from any external peer by filtering longer prefix e.g. /24 [22].
- Always aggregate address where ever possible, specially at the boundaries.
- Always set a limit on the number of prefixes you are accepting from a peer and set a warning when neighbor reaches a certain percentage of the limit “neighbor x.x.x.x maximum-prefix 200 80” This command used in Cisco IOS sets the limit to 200 prefixes from a peer and will issue warning messages when the neighbour reaches 80% of the limit.
- Prefix limiting filter should be precise or slightly larger than what you anticipate from the external peer, this will protect your network from excessive unwanted prefixes from the peer.
- Prefix advertised by a customer they do not own should always be considered invalid and must be discarded (Consult IRR) [21].

- All the BGP advertisements with private AS numbers (64512 – 65535) should be considered invalid and must be blocked from entering the network [22].
- Similarly private numbers should not be propagated on the internet, remove the entire private AS number block from outgoing advertisement [22].
- AS number of the external peer should always be the first number on AS_PATH list and filter every thing that does not match the criteria, this would protect against spoofed prefixes entering into your domain.
- Only allow packets to and from your external peer with TCP port 179 and IP address of the border router external interface and the peer's external interface.
- Always consult all the RIR's published CIDR blocks, and continuously update your access list filters based on that, every ISP should validate the traffic using these lists.

7. Related Work

Interdomain routing security has been studied for a long time but still an efficient solution has not been defined, researchers are working towards proposing new protocol such as Secure Border Gateway Protocol (S-BGP) to support the authentication of routing, based on Public Key Infrastructure [25]. The SoBGP protocol using a topology database to validate paths being advertised, which does not possess path authentication mechanism but basically provides a mechanism to detect inconsistent routes [26]. Other solution such as IRV which proposes a services that protect against AS misconfiguration by validation both dynamic and static interdomain routes, unlike SBGP IRV, do not requires the replacement of existing BGP [14]. To improve the overall security future work is required to agree upon a single solution which is accepted by the whole internet, either to adopt a new protocol for the internet infrastructure with all required security functionality, or may by attaching new services or protocols to exiting BGP to achieve the required result.

8. Conclusion

In this paper I have highlighted some of the configuration error that can cause a major damage to the internet routing infrastructure, and proposed some of the approaches for BGP routing policy and route filtering with best practices, to filter and validate BGP routes and secure the BGP routing protocol. These proposals will help in achieving the purpose of security but they can not completely solve the problem as day by day new incidences are being recorded and new vulnerabilities are highlighted, some experts argue BGP have served its purpose and now it time for change, I agree with it to a certain level because BGP was designed more than a decade ago and to support such a huge internet architecture was not in its design. So inevitability we have to enhance BGP, or use BGP in conjunction with another protocol that can provide the solution or creating a new protocol with all the required security and scalability features, but the question of how and will the whole internet (all the autonomous systems in the world) agree on a single solution will still remain.

References

- [1] Y. Rekhter, T. Li, A Border Gateway Protocol 4 (BGP-4) RFC 1771, IETF Network Working Group, March 1995.
- [2] D. Estrin, Y. Rekhter, S. Hotz A Unified Approach to Inter-Domain Routing RFC 1322, IETF Network Working Group, May 1992.
- [3] Cisco Systems Inc. Border Gateway Protocol E-Book, Cisco Systems Inc., December 2003.
<http://tinyurl.com/3d5ub>
- [4] P. Traina Autonomous System Confederations for BGP RFC 1965, IETF Network Working Group, June 1996.
- [5] R. Chandra, P. Traina BGP Communities Attribute RFC 1997, IETF Network Working Group, August 1996.
- [6] S. Murphy, BGP Security Vulnerabilities Analysis, RFC 4272, Sparta Inc. January 2006.
<http://www.ietf.org/rfc/rfc4272.txt>
- [7] Mahajan, R., Wetheral, D. and Anderson, T. Understanding BGP misconfiguration. In Proc. ACM SIGCOMM (Pittsburgh, PA, Aug.2002).
- [8] Vince Fuller, Tony Li, Jessica Yu and Kannan Varadhan, RFC1519 - Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy, September 1993.

- [9] Tao Wan & Paul C. Analysis of BGP Prefix Origins During Google's May 2005 Outage, School of Computer Science Carleton University, Ottawa, Canada.
<http://www.scs.carleton.ca/~paulv/papers/ssn06-fine.pdf>
- [10] Anton Tony Kapila, Alex Pilosov, August, 2008, Revealed: The Internet's Biggest Security Hole.
<http://blog.wired.com/27bstroke6/2008/08/revealed-the-in.html>
- [11] Sargun Dhillon, YouTube IP Hijacking, NANOG mailing list, Sun Feb 2008
<http://www.merit.edu/mail.archives/nanog/msg06299.html>
- [12] E.Chen, T. Bates, an Application of the BGP Community Attribute in Multi-home Routing, August 1996.
<http://www.ietf.org/rfc/rfc1998.txt>
- [13] Xiaoliang Zhao, Dan Pei, Lan Wang, Dan Massey, Allison Mankin, S. Felix Wu, Lixia Zhang, Detection of Invalid Routing Announcement in the Internet, 2002.
http://irl.cs.ucla.edu/papers/fniisc_dsn02.ps
- [14] Geoffrey Goodell, William Aiello, Timothy Griffin, John Ioannidis, Patrick McDaniel, Aviel Rubin, Working Around BGP: An Incremental Approach to improving Security and Accuracy of Interdomain Routing.
www.isoc.org/isoc/conferences/ndss/03/proceedings/papers/5.pdf
- [15] IPv4 Global Unicast Address Assignments (last updated 2008-11-11).
<http://www.iana.org/assignments/ipv4-address-space>
- [16] V. Fuller, T. Li, J. Yu and K. Varadhan, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy, RFC1519, September 1993.
<http://www.faqs.org/rfcs/rfc1519.html>
- [17] List of Routing Registries.
<http://www.irr.net/docs/list.html>
- [18] Current RIPE Documents by Number
<http://www.ripe.net/ripe/docs/titletoc.html>
- [19] Allocation sizes within APNIC IPv4 address ranges
<http://www.apnic.net/db/min-alloc.html>
- [20] IP Address Space Allocated to ARIN
http://www.arin.net/reference/ip_blocks.html
- [21] Best Practices for Securing Service Provider Networks, Juniper Networks, Inc, Sept 2008.
http://www.juniper.net/solutions/literature/white_papers/200180.pdf
- [22] Cisco ISP Essentials, Cisco Systems Inc, June 2001.
<http://www.ccxx.net/books/Cisco%20ISP%20Essentials.pdf>
- [23] Rick Kuhn, Kotikalapudi Sriram, Doug Montgomery. Border Gateway Protocol Security, recommendations of the National Institute of Standards and Technology. Special Publication 800-54 July 2007.
<http://www.csrc.nist.gov/publications/nistpubs/800-54/SP800-54.pdf>

[24] RFC 3330 - Special-Use IPv4 Addresses, September 2002.
<http://www.rfc-archive.org/getrfc.php?rfc=3330>

[25] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4), Apr. 2000.

[26] J. Ng. Extensions to BGP to support secure origin BGP (soBGP). Internet Draft, Oct. 2002.