PEC 1 Schedule Securing the Border Gateway Protocol BGPv4

Máster Interuniversitario de Seguridad de las TIC

Advisor: Joan Borrell Viader

Student: José María Foces Vivancos

June 2018

Contents

1. Introduction

This document defines my final master project about BGPv4 Security.

The section below, state of art, summarizes the status of security measures implemented to secure route exchange through protocol BGPv4. Considering the current status, personal and technical motivations are explained and after that, it provides a detailed description about the problem to solve and the main goals of the project.

Once objectives are clear, I explain the tasks to be executed to accomplish and the work methodology.

2. State of Art

Inter-domain routing is managed by BGPv4. *"This protocol provides a set of mechanisms for supporting Classless Inter-Domain Routing (CIDR)"* (RFC4271 - A Border Gateway Protocol 4 (BGP-4)).

This protocol, now on version 4, has evolved since the first version released (RFC1105) 1989, through several versions: (RFC1105 - A Border Gateway Protocol (BGP)), (RFC1163 - A Border Gateway Protocol (BGP)), (RFC1267 - A Border Gateway Protocol 3 (BGP-3)), (RFC1654 - A Border Gateway Protocol 4 (BGP-4)), (RFC1771 - A Border Gateway Protocol 4 (BGP-4)).

The initial draft for this version ((RFC1654 - A Border Gateway Protocol 4 (BGP-4))) dates from 1994, and it was built on the experience gained with EGP and its usage in the NSFNET Backbone and obviously previous versions of BGP.

The protocol has evolved, and the latest specification can be found at (RFC4271 - A Border Gateway Protocol 4 (BGP-4)).

It's very important to remember that the protocol does not provide any security measures by itself and currently, it makes internet work.

To analyze the current trends to secure BGPv4 the following information sources have been considered:

1. ACM's article (Rethinking Security for Internet Routing). At high level it can be extracted that there are some attacks that may succeed despite of the security measures implemented currently (the article dates from October of 2016).
2. Out of the specification, there is a set of documents important to keep in mind in order to establish the current context of route exchange. The most important documents out of the standard are:
    a) (RFC3013 - Recommended ISP Security, 2000)

b) (RFC4272 - BGP Vulnerability Analysis, 2006)

c) (RFC7454 - BGP Operations and Security, 2015)(February)

3. MANRS ((Mutually Agreed Norms for Routing Security, s.f.)): this is a commitment by network operators around the world to clean their part of the street and improve the security of the global routing system. MANRS was born on early 2014.

It would be hard to state with no doubt what are most popular security measures applied to the route exchange, since it would require information about what security measures are enforced by ISPs and in some cases, it may disclose vulnerabilities. Therefore, that is not likely to be available. It would be possible if all ISPs and network operators have already joined MANRS, but they do not.

To provide a not fully precise but enough explanation about current trends concerning BGPv4 security I will use the sources mentioned above. They give some clues about that.

2.1. Best practices extracted from RFCs

RFC7454 BGP OPsec advises to implement the following actions in order to secure the route exchange:

1. **Protection of the BGP speaker:** mainly by ACLs to permit connections to port 179 from a reduced set of IP addresses. Moreover RFC 6192 provides more information about protection of the control plain.

2. **Protection of BGP Sessions:** the BGP session is the connection between two speakers. Attacks on TCP connections are possible and a MITM attack can succeed for example. These connections can be secured by authenticating the packets received. The most popular way to achieve this was TCP MD5 Signature option [(RFC2385 - Protection of BGP Sessions via the TCP MD5 Signature Option)] in 2015 and is obsoleted by TCP-AO [(RFC5925 - The TCP Authentication Option)]. The drawbacks are the additional maintenance and configuration overheads.

3. **BGP TTL Security (GTSM):** as stated at section 2 of (RFC5082 - The Generalized TTL Security Mechanism (GTSM)), the vast majority of protocol peerings are between adjacent routers. This protection is based on the impossibility for an attacker to make packets arrive with TTL 255 to the victim router. Unless he is directly connected to the victim. It reduces the attack area in a very simple and robust way.

4. **Prefix Filtering:** this is a simple white list technology currently implemented by almost all manufacturers. It controls the IP prefixes that can be received and announced on BGP peerings. It's important to remember that IANA allocates IPv4 and IPv6 spaces.

   a) There are pretty obvious prefix filters that must never be considered by BGP peers: special-purpose and unallocated prefixes. Moreover, as detailed on

(RFC7454 - BGP Operations and Security, 2015) there is no guarantee that the list of allocated prefixes is updated regularly and therefore it's better not to configure any filters based on allocated prefixes.

b) IANA relays prefix allocations on RIRs (Regional Internet Registries). RIRs manage the property of prefixes. These prefixes are owned by LIRs (Local Internet Registries). LIRs are for example: ISPs, enterprises or academic institutions. Therefore, a more precise check can be performed to make sure that the prefixes received are originated or transited by ASes in charge of them. From the previous experience, it's been observed that an AS can easily advertise a prefix that does not belong to it and create black holes or security threats. There are two options to achieve this:

   i. A short-term solution is to use prefix filters created from IRRs (Internet Routing Registries). IRRs are databases containing Internet routing information in language described at (RFC4012 - Routing Policy Specification Language next generation (RPSLng)). At high level, it's possible to build a white list of originated or transited prefixes that one should accept. This list should be kept up to date regularly.

   ii. A long-term solution called SIDR (Secure Inter-Domain Routing). It's an infrastructure, described in (RFC6810 - The Resource Public Key Infrastructure (RPKI) to Router Protocol), (RFC6811 - BGP Prefix Origin Validation), (RFC6480 - An Infrastructure to Support Secure Internet Routing), (RFC6483 - Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs))and (RFC7715 - Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI)). It was designed to secure internet advertisements. It's mainly divided in two services:

      1. Origin Validation: the main purpose is to validate the route originating a given route. It's now operational and some entities implement this security measure.

      2. Path validation via BGPsec. On 2015 BGPsec was being designed and a threat model could be found at RFC7132.

c) Prefixes that are too specific. For example, RIPE community documented that at 2015 IPv4 prefixes longer than 24 are rarely announced or accepted in the Internet.

d) Filtering prefixes of the Local AS and Downstreams. Small prefixes should not be sent between peers. This prevents intra-AS traffic from leaking over external peerings. This is only applicable to single-homed scenarios since it could break desired redundancy on multi-homed ones. An operator with a multi-homed

customer should keep accepting customer prefix from the upstream peers. This will make possible to reach customer prefix even if the customer-provider link is down, through the other ISP/s.

e) IXP (Internet Exchange Points) LAN Prefixes

   i. Network Security: IXP members using a LAN prefix should not accept more specific prefixes from any external BGP peers. Doing so would create black holes to the IXP LAN.

   ii. PMTUD and Loose uRPF problem

f) The default route: obviously the prefix 0.0.0.0/0 should never be accepted or advertised except in specific customer/provider configurations. Typically, it's recommended to filter this prefix.

5. **BGP Route Flap Dampening**: it's a mechanism that makes possible to fine routes each time they change in the BGP routing table. RIPE community has recommended to not to implement this, as some studies revealed that it could harm more than help the Internet.

6. **Limit prefixes on a peering**: there are two kinds of limits:

   a) From peers: the limit should be lower than the number of routes in the internet.

   b) From upstreams providing full routing: the limit should be higher than the number of routes in the Internet. This is mainly oriented to protect router's memory.

7. **AS Path Filtering:** the idea is to filter received routes by ASes path. The following policies are recommended:

   a) Accept short routes from customers, 2-4 bytes AS paths.

   b) Do not accept/advertise prefixes with private AS numbers in the path.

   c) Do not accept prefixes when the first AS in the path is not a peer. Unless the peering is done toward a BGP Route Server with transparent path handling.

   d) Do not advertise prefixes with nonempty AS path unless you are providing transit for these prefixes.

   e) Do not advertise prefixes with upstream AS numbers unless you are providing transit to them.

8. **Next-Hop Filtering:** sometimes it's desired not to filter this kind of route advertisements, for example for BGP Route-Server setups where the route server just relays routing information, but it will never route traffic by itself. On other cases like:

a) Direct Peering ISP-ISP' this is not desired since it may create blackholes or may trick the other into sending traffic to an unsuspected third party. Therefore, a policy to filter this case should be applied in order to just accept next hop for accepted prefixes to the BGP peer address.

Since this document was published in February 2015 the technology has evolved and there are some RFCs that should be mentioned since they have or may have impact on system structure, operation and evolution of the Internet:

(RFC7947 - Internet Exchange BGP Route Server), purposes a standard to centralize routing information updates inside IXPs. As mentioned the Route Server is not a router but provides necessary information to the routers so they can route traffic. From September 2016.

(RFC8205 - BGPsec Protocol Specification): BGPsec specification, released on September 2017.

2.2. ACM's article: Rethinking Security for Internet Routing

It makes a very good analysis and syntheses about what defenses protect against given attacks in certain contexts. A brief summary:

"BGP is insecure because any AS can announce any path it wants to any subset of its neighbors."

With that in mind and after explaining protocol operations with examples. The document is mainly focused on vulnerabilities derived from the protocol itself. No threats or defenses derived from the use of other protocols such as TCP are explained. It details the following defenses that counter the following threats:

1. **Defense: Origin Validation:** BGPv4 does not provide a way to validate the prefix allocation. This defense is implemented by RPKI. As mentioned above, [4.b.ii] SIDR provides a way to authenticate advertised routes by providing a trusted database binding ASes to prefixes. Route advertisements are therefore authenticated by using key pairs. Therefore, it protects against the following attacks:

    a) Sub prefix hijack: remember that longest prefix matching is used to get the best route for a network. When an AS advertise to all its neighbors that it has the shortest path to a given CIDR that does not belong to it and this domain is sub-prefix of a prefix in charge of another AS. If it's successful it will be attracting all traffic from its neighbors to this domain.

    b) Prefix hijack: the same technique employed before but, in this case, it won't be able to attract the 100% of the traffic destined to this domain.

    c) Route leak: in a route leak the attacker violates the agreed export policies by advertising legitimate route to many of its neighbors, some of them should know

that path but others do not. The impact produced is that the ASes involved in the path will attract more traffic to the prefix that they are expected to.

2. **Defense: Topology Validation:** BGPv4 does not provide a way to validate the network topology and therefore, discard advertisements that does not match with it. This defense is implemented by soBGP, that uses a similar PKI infrastructure as RPKI does. The PKI usage is different since the purpose of this defense is to validate the presence of links (direct connections between ASes).

    a) One-hop hijack: the attacker can advertise a route meaning that traffic can go through it to the destination prefix (in charge of the legitimate AS). Therefore, the origin validation cannot protect against that. But this route is bogus since the attacker has no direct connectivity with the legitimate AS. The experience shows that this attack can trick attacker's neighbors to deliver traffic to the legitimate prefix to it.

3. **Defense: Path Validation:** BGPv4 does not provide a way to validate paths and discard advertised routes that do not match with export policies of ASes involved. BGPsec is one of the most popular way to implement this validation. It requires RPKI and each AS on the route append its signature to BGPv4 messages. The main drawback of this security measure is that the computational overhead involved would require routers with dedicated cryptographic processors. Without this security measure implemented the following attack is possible:

    a) Announce an unavailable path: even when the path to the destination prefix is valid (in terms of connectivity, there are links between all routers the full path), there may be another restriction, for example, one of the intermediate AS in the path does not want to forward traffic to the destination prefix. Therefore, in this case it does not announce this route. The attacker may want to announce this route to all his neighbors. Therefore, in certain circumstances, for example when the advertised path is shorter than the legitimate one. The attacker achieves to get a portion of the traffic destined to the given prefix.

4. **Prefix filtering:** this is just a white list. As stated before, it controls the IP prefixes that can be received and announced by ASes on BGP peerings. It would be hard to maintain this list for all customers of ISPs or by using information in IRRs, so in practice a conservative policy is implemented. ISPs can filter route advertisements of customer ASes when the given ISP is the only provider for the customer AS. This is not applicable when the customer has several providers, since that brakes the redundancy that the customer is buying when it has two providers. As an example, consider that any other customer of the given ISP can reach the multi-homed customer even if the link through the ISP is down, by the other ISP. It won't be possible if the ISP filters this.

a) Announce an available path: this is also a route leak. In some contexts, an AS may don't want to tell others that it can reach the given prefix through a peer. It's produced when the attacker violates agreed export policies and announces a shorter path than the agreed one.

2.3. MANRS (Mutually Agreed Norms for Routing Security)

As stated above, it is a commitment by network operators. It was born on February 2014. The main objective is to contribute to a better global routing system. This objective is accomplished by implementing at least the following actions:

1. Filtering to ensure the correctness of announced prefixes and received announcements from customers or adjacent networks.
2. Anti-Spoofing by enabling source address validation for at least single-homed stub customer networks, end users and infrastructure.
3. Coordination to maintain an accessible up-to-date contact information.
4. Global validation by publishing routing information so other members can validate routes at global scale.

Now it has been adopted by almost 60 entities including RIRs (RIPE NCC, LACNIC), ISPs (Level 3, Swisscom) and enterprises (Verisign). The most of them implement the previous the full set of security measures listed above.

They also have an active news section. Concretely, the article "*14000 Incidents: a 2017 Routing Security Year in Review*" (Mutually Agreed Norms for Routing Security, s.f.) dating from February's 26 of 2018. It summarizes the incidents in the Internet routing system. Highlights (as stated on the article):

1. 13935 total incidents of any types (outages, route leaks or hijacks) had place.
2. Over 10% of all autonomous systems were affected.
3. 3106 ASes were a victim of at least one routing incident.
4. 1546 networks caused at least one incident.

2.4. Conclusions

Internet routing system is currently vulnerable to known types of attacks. Security measures, to mitigate the impact of these attacks, have been explained but not all of them are implemented today.

Each security measure solves a part of the problem. The most popular ones are what improve the security of the AS/es implementing them, for example BGP session protections or prefix filtering white-listing. There are others that will make the AS

implementing them to be a bit safer than before, while improving security of neighbor ASes. Concretely, implementing certain solutions is expensive (BGPsec or soBGP) and do not provide a security improvement worth the effort if not all parties involved implement them, if they are not fully deployed. But others, like origin validation (RPKI) are capable to provide tangible benefits in partial deployments. Therefore, the trend it's to evolve to an RPKI secured routing system and nowadays we may find some ASes implementing them. This allows Internet's routing system to evolve, in the future, to be safer. It's important to consider that more robust security measures require RPKI to be deployed before they can operate, such as BGPsec.

Moreover, it's hard to make the routing system to evolve to be safer since there are a lot of people involved and some security measures won't provide a benefit unless all these parties involve decide to implement them. Therefore, enterprises or ISPs, do not deploy fully secured ASes if the others are not committed to do the same. This is the most important reason by what MANRS arises.

As a benefit in general, for all members of the Internet: a robust routing system may help to defeat some types of attacks out of the scope of the routing system itself. For example, Source Address Spoofing based DDoS attacks would be mitigated or even impossible if all entities or at least the biggest internet carriers enforce security measures outlined and enforce them customers to implement them.

3. Motivation

From a personal point of view, I selected this project since I want to learn about these technologies. It provides me an extension of the knowledge acquired during the master and my professional career about computer networks while delving into protocols that I currently know at higher level.

It's also interesting for me since management and operation of this kind of systems is usually out of the standard network user and therefore, it implies a technical evolution for me. In addition, as stated by the project description, BGPv4 packets usually lack security measures. If I remember correctly, there are some TCP extensions that provide message authentication such as TCP-AO that provides message authentication by adding a HMAC to the packet, as detailed at RFC-5925. No doubt it's an exciting project.

From a technical point of view: considering facts explained above, that define the current context of BGPv4 security, there are some attacks that are still possible against Internet's routing system. Therefore, it's interesting to perform an analysis over them.

Results may help other people understand them and generate awareness, so we can build a better Internet.

4. Problem definition

This section is broken down into two parts. The first one quotes from project description form alternatives provided by the university, and the second one gives more detail about the concrete problem and how I will solve it.

4.1. Project description

Resource Public Key Infrastructure (RPKI) [RFC 6810, 2013] is a specialized public key infrastructure (PKI) framework designed to secure the Border Gateway Protocol (BGPv4) of the TCP/IP protocol suite.

The basic blocks of the Internet's architecture are called Autonomous Systems (AS), each of them composed by a set of networks with a common administrative policy. Border Gateways are the routers of an AS in charge of announcing the network prefixes inside this AS to the other AS, by using BGPv4. Despite its central role in Internet routing, the original BGPv4, as proposed in [RFC 4271 (2006)], lacks any security measure. Therefore, a number of attacks against BGPv4 have been deployed in the past, being "Pakistan Telecom vs. YouTube" (2008) the most popular one (https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study). Among all the proposals to secure BGPv4, BGP Prefix Origin Validation (POV) [RFC 6811, 2013], based on the use of RPKI, seems to be the most promising one. POV has been adopted by a number of commercial routers, and recently an open source version of it (known as BGP Secure Routing Extension (BGP-SRx) has also been made available by the USA's National Institute of Standards and Technology (NIST). BGP-SRx runs over Quagga, the open source version of BGPv4, on CentOS boxes. It is available at http://www-x.antd.nist.gov/bgpsrx/ The goal of this project is double. First, to provide an up-to-date analysis of the security of POV and RPKI. Second, to construct and evaluate a working scenario of several simulated secured AS, communicating through BGP-SRx, over some virtualization environment. The simulated scenario should show its correct behavior in front of some attacks. The simulated scenario should show its correct behavior in front of some attacks.

4.2. Detail and Objectives

Project description states two main objectives:

1. to provide an up-to-date analysis of the security of Prefix Origin Validation (POV) and Resource Public Key Infrastructure (RPKI).
2. to construct and evaluate a working scenario of several simulated secured AS, communicating through Border Gateway Protocol - Secure Routing Extensions (BGP-SRx), over some virtualization environment.

Considering that Internet's routing system security evolves slower than protocols to secure it, my will is to extend the current objectives as follows:

The main objective is to provide an up-to-date analysis of the security of BGPv4 and other AS-to-AS, route-communication protocols. This objective can be broken down into the following:

1. **To provide an analysis of the current guarantees provided to Internet's routing system.**
    a) As explained before, the most popular configuration contains the following security measures (mainly from (RFC7454 - BGP Operations and Security, 2015)):
        i. BGP Speakers protections
        ii. BGP Session protections
        iii. GTSM
        iv. Prefix Filtering: by white lists
2. **To provide an analysis of the benefits provided by SIDR**:
    a) As explained at the state of art, SIDR was designed to secure route advertisements. It offers two services:
        i. Origin Validation, that seeks to check the correctness of attributes associated with routes. ROA (Route Origin Authorization), validated against RPKI (Resource Public Key Infrastructure).
        ii. Path Validation: that attempts to ensure that no one advertise bogus routes. This is currently implemented by BGPsec. In addition, and if it fits in time, my will is to extend the project and analyze it.
3. **Summarize and show quantitative results of the security improvements.**

To achieve these objectives several working scenarios implementing mentioned security measures may be deployed in order to perform analyses over them when it's needed.

5. Methodology

To summarize the current status of BGPv4 security I've used several public information sources. After this status is clear, the method to analyze each alternative to secure BGPv4 route exchange will be analyzed in the following way (the Scientific Method):

1. Ask question

2. Do background research
3. Construct hypothesis
4. Test with PoCs
5. Analyze results
    a) If true, then 6
    b) Otherwise back to 3
6. Report results

6. Tasks and Schedule

Beforehand, it's important to say that I will write a log book to have trace of every task, results and decisions taken during the process.

This log book with the results of each phase is the base to write project's memory and slide show.

1. Define the State of Art of BGPv4 Security.
2. Prepare the host for the virtualization environment on the top of Linux-KVM.
3. For each alternative to secure route exchanges:
    a) Define the simulation, composed of several inter-connected virtual machines.
    b) Build and start the simulation
        i. Decide BGPv4 implementation. This may change between alternatives.
        ii. Configure the routers
        iii. Validate simulation operation.
    c) Establish attacks over the infrastructure.
    d) Perform attacks over the infrastructure
    e) Analyze results
    f) Build conclusions about the alternative.
4. Once the loop of alternatives and security analyses have been carried out over all of them.
    a) Write the memory
    b) Pack simulations
    c) Create the slide show, record the video.

7. Bibliography

Mutually Agreed Norms for Routing Security. (n.d.). Retrieved from https://www.manrs.org

Rethinking Security for Internet Routing. (n.d.).

(n.d.). *RFC1105 - A Border Gateway Protocol (BGP).*

(n.d.). *RFC1163 - A Border Gateway Protocol (BGP).*

(n.d.). *RFC1267 - A Border Gateway Protocol 3 (BGP-3).*

(n.d.). *RFC1654 - A Border Gateway Protocol 4 (BGP-4).*

(n.d.). *RFC1771 - A Border Gateway Protocol 4 (BGP-4).*

(n.d.). *RFC2385 - Protection of BGP Sessions via the TCP MD5 Signature Option.*

(2000, 11). *RFC3013 - Recommended ISP Security.*

(n.d.). *RFC4012 - Routing Policy Specification Language next generation (RPSLng).*

(n.d.). *RFC4271 - A Border Gateway Protocol 4 (BGP-4).*

(2006). *RFC4272 - BGP Vulnerability Analysis.*

(n.d.). *RFC5082 - The Generalized TTL Security Mechanism (GTSM).*

(n.d.). *RFC5925 - The TCP Authentication Option.*

(n.d.). *RFC6480 - An Infrastructure to Support Secure Internet Routing.*

(n.d.). *RFC6483 - Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs).*

(n.d.). *RFC6810 - The Resource Public Key Infrastructure (RPKI) to Router Protocol.*

(n.d.). *RFC6811 - BGP Prefix Origin Validation.*

(2015). *RFC7454 - BGP Operations and Security.*

(n.d.). *RFC7715 - Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI).*

(n.d.). *RFC7947 - Internet Exchange BGP Route Server.*

(n.d.). *RFC8205 - BGPsec Protocol Specification.*