BGP hijacking

A brief guide on protecting BGP from bad actors

HACKTHEBOX

Anastasios-Theodosios Meletlidis - lean@hackthebox.eu

2024

What is BGP hijacking?

Summary

Unauthorized takeover of BGP routes.

BGP hijacking occurs intentionally or unintentionally when an AS is announcing a route to IP prefixes it doesn't control. If unchecked, this announcement could spread and be included in BGP routing tables across the Internet

This can be abused to redirect and leak traffic through malicious routers in a man-in-the middle attack, or to amplify DDoS attacks.

Why does it happen?

BGP inherently believes the information provided by its neighbor ASes. It is from this inherent belief that security issues arise.

Main causes are the decentralized and trust-based nature of the protocol and lack of universal adoption of security measures like RPKI and BGPsec.

What is BGP hijacking?

IP hijacking can occur one of several ways

- ➔ An AS announces that it originates a prefix that it does not actually originate.
- ➔ An AS announces a more specific prefix than what may be announced by the true originating AS.
- ➔ An AS announces that it can route traffic to the hijacked AS through a shorter route than is already available, regardless of whether the route exists.

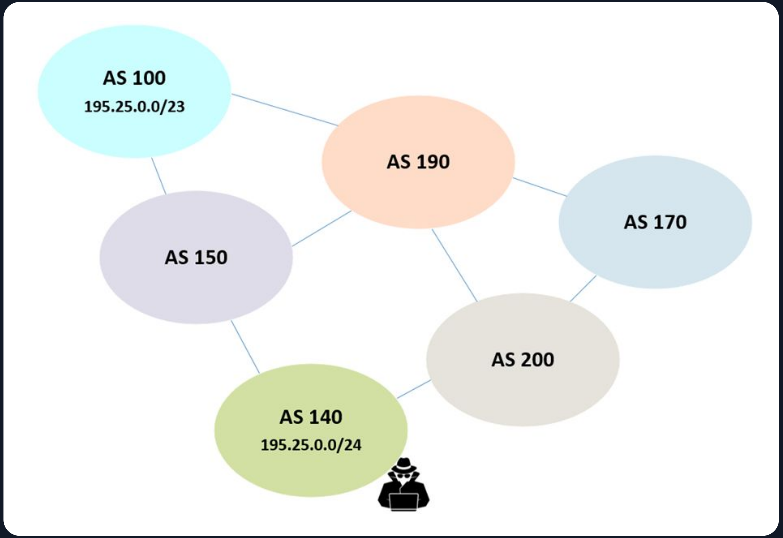

In this example an attacker controlled peer AS140 advertises a more specific prefix (CIDR subnet), thus causing all traffic going to AS100 to be re-directed to AS140.

Other types of BGP attacks

TCP/IP Protocol attacks: Since BGP runs on top of TCP, common TCP attacks can also affect it with examples including, spoofing and TCP reset, session hijacking or SYN flooding attacks.

BGP route manipulation attacks: These attacks abuse the way BGP works by advertising bogus prefixes or paths, they compose BGP hijacking which we will cover today.

Protocol manipulation attacks: These relatively new types of attacks attempt to exploit BGP by causing unintentional behaviour of the protocol, usually caused by multi-exit-discriminator (MED) modification or exploiting the RFD/MRAI timer.

Denial of service attacks via resource exhaustion: Attackers can flood the BGP speaker with too many BGP messages so that it can not even process legitimate BGP packets.

Real world examples

2014

BGP hijacks were used to intercept communication between Bitcoin miners and mining pools. This allowed an adversary to obtain bitcoin that should have been allocated to the mining pool.

2018

Attackers executed a BGP hijack, diverting traffic to Amazon's DNS service. They manipulated DNS queries for "myetherwallet.com" leading users to a fake site. Some victims unknowingly entered login details, resulting in the theft of cryptocurrency wallets.

2022

In August 2022, a BGP hijack on Celer Bridge used fake AltDB entries and forged announcements, tricking a transit provider to redirect cryptocurrency funds to the attacker's account.

2023

Balancer.fi experienced an intricate BGP Hijacking attack leading users to a compromised link, altering HTTPS certificates and deploying malicious JavaScript code from resulting in phishing attacks on users' connected wallets.

Methodology of protecting BGP

Information that must be validated to secure BGP speakers and sessions

- ➔ Does this BGP message come from an authorised BGP peer? (BGP speaker authentication).
- ➔ Is the legitimate holder originating the prefix, or is the AS authorised to originate it? (Origin validation).
- ➔ Does the AS path reflect the sequence of ASes that the BGP UPDATE message has traversed? (Path validation).
- ➔ Are the attributes in the BGP UPDATE message correct, and has anyone tampered with them? (Attribute verification).

https://www.rfc-editor.org/rfc/rfc4271

Protecting the BGP session

Methods for protecting sessions

- ➔ TCP-MD5: is used to protect BGP sessions against spoofed TCP segments and TCP resets, both sides use a pre-shared key to generate and verify a checksum using MD5 on TCP segments, ensuring message integrity and authenticity by discarding messages where the checksums do not match.
- ➔ TCP-AO: enhances the security and authenticity of TCP segments in BGP and LDP sessions by using message authentication codes (MACs) that rely on shared keys, thereby addressing the limitations of TCP-MD5.
- ➔ GTSM: protects directly connected eBGP sessions from attacks through TTL detection by checking the TTL (IPv4) or hop-limit (IPv6) of incoming IP packets to make sure that they have not been spoofed.

TCP-MD5 and GTSM are already widely implemented by default on most vendor software without the need to install additional packages.

Route filtering

Routes exchanged between BGP peers are controlled by using BGP filtering policies. They are rules used to decide what prefixes to announce to your peers and what prefixes you would like to accept from them. Filtering is usually done on ip prefixes or AS paths.

Filtering can be done on:

- Inbound traffic
- Outbound traffic

Approaches:

- Explicit Permit (permit all then deny any)
- Explicit Deny (deny all then permit any)

Rules of filtering:

- Filter as close to the edge as possible:

  You should implement filtering as close to the edge as possible and discard the undesired routes at the edge of your network before they enter your routing domain.

- Filter as precisely as possible:

  You should specify exactly what routes you want to receive or discard to or from your BGP peers.

Route filtering

Prefix List

Configuring a prefix list allows you to set a range of routes you wish to accept or announce. They can be set either by manual bgpd commands or generated by automatic tools.

Tools

- **level3 filtergen**
- **bgpq4**
- **peval**

These tools can connect to the RIPE database (and other sources) to look for the **route(6)** objects associated with an AS and then get the prefix within the **route(6)** object and creates the BGP prefix-list configuration.

Important prefix filters

- Special purpose prefixes (ipv4/ipv6)
- Unallocated prefixes
- Longest accepted prefixes
- Routes that contain own prefixes
- IXP LAN prefixes
- The default route (0.0.0.0/0, ::/0)

RPKI

Resource Public Key Infrastructure (RPKI) is a trust chain system implemented by Regional Internet Registries (RIRs) that utilizes X.509 certificates to implement Route Origin Authorisation (ROA).

In short it provides a robust way to know which AS'es should announce which prefixes in a centralized way using RIRs. This source of authority can be delegated from RIR's to LIR's by signing certificates.

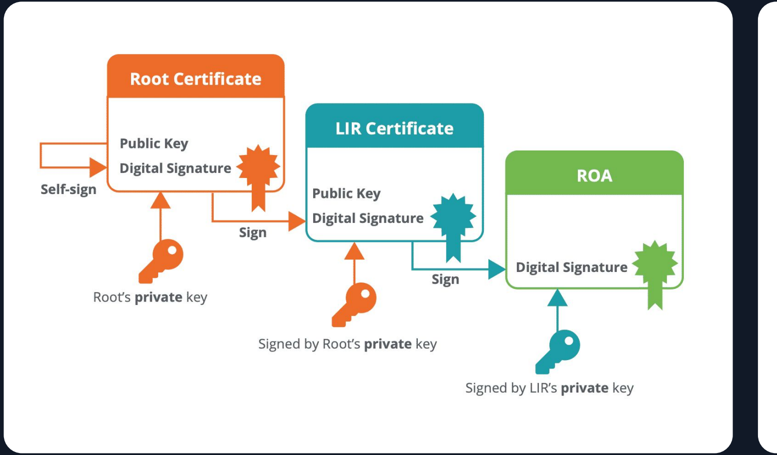

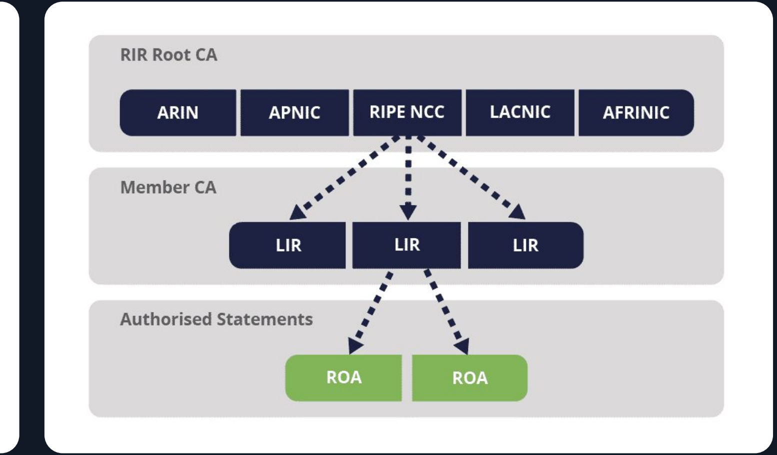

HACKTHEBOX

Thank you for listening!

We hope this presentation gave you a better feel of how to protect BGP from bad actors.

Anastasios-Theodosios Meletlidis - lean@hackthebox.eu