An In Depth Discussion of BGP Hijacking Attacks and How Systems Detect and Mitigate Such Attacks

Parth Diwane[1] and Rucha Vaidya[#]

[1]Amador Valley High School, USA
[#]Advisor

ABSTRACT

While internet architecture and software have made drastic developments since the birth of the internet, the internet is still an unsecure place where attacks are common. Despite the security of the internet, hundreds of BGP hijacking attacks occur each year. In 2021 alone, there were 775 possible instances of BGP hijacking attacks. These attacks redirect information to an alternative network where it could be compromised to misused, but they are also a major inconvenience for the user as connection to a particular website could be completely dropped (known as black holing). The purpose of this review is to synthesize and link two discrete topics, being cybersecurity and network architecture, and how those topics work together. To find evidence for this review, databases such as Google Scholar and IEEE Xplore were queried. Additional information on how evidence was gathered, how quality of information was ensured, and strategies used to synthesize information is discussed in the Methods section. This review provides an in depth discussion on BGP hijacking attacks, the different types of BGP hijacking attacks, how the BGP protocol (and other routing protocols) function, the structure of a BGP message, how systems detect and mitigate BGP hijacking attacks, and a summary of internet architecture and how devices on the internet communicate.

Background Information

The Internet is defined as a web of many interconnected computer networks. Devices, also known as end systems, communicate over the internet by sending pieces of information, known as packets, from source to destination address. The device that sends the information is known as the sending end system, conversely, the device that receives the message is known as the receiving end system. However, certain measures need to be implemented for the packet to be sent and received. To account for this, network engineers implemented a layered architecture for the Internet.

There are three main layers to the internet, the application layer, the transport layer, and the network/IP layer. Additionally, each layer oversees a different function that allows the packet to go from the sending end system to the receiving end system. To accomplish these functions, each layer has specific protocols set into place. Protocols allow packet information to be added or manipulated to the packet so it can travel from the sending to the receiving end system. The layers of the internet can be seen as linked nodes (devices that relay information) that send information in a forward fashion with respect to the sending end system but in a reverse fashion with respect to the receiving end system.

The application layer is where all of today's websites and apps reside. The application layer is home to the famous HTTP, FTP, and SMTP protocols. These protocols allow internet users to gain access to websites, send information from one device to another, and allow for the sending and receiving of emails respectively.

The transport layer transports messages (information such as website files, for example, HTML files), to the network layer. It does so by using either the TCP or UDP protocol. The TCP protocol is a wired protocol that allows for guaranteed delivery of the application message, congestion control, and flow control (matched bit rates for the sending and receiving end system). In contrast, the UDP protocol is a wireless protocol that does not guarantee the delivery of a message and does not have congestion or flow control.

The network layer does the action of delivering the information (which it obtained from the transport layer) to the receiving end system's transport layer. To accomplish this, the network layer uses the IP protocol. The IP protocol inserts information into the application message that informs routers where to send the application message and what route to take to get to the destination source.

Routing is when a packet must be taken from the source to the destination address. To accomplish this, a router must use something known as forwarding. Forwarding is the action of sending the packet from one router to another router over the predetermined path.

Computer networks connect nodes, such as routers or other computers, via coaxial cable or wireless signal. These nodes can be placed into groups, known as autonomous systems (ASs), based upon the administrative control (the person or thing controls these nodes). As seen in Figure 1 below, routers 2a through 2d all fall under AS2 as they all have the same administrative control. In many cases, packet information may want to be sent from one node to another node that is on a similar AS. For example, router 1a might send packet information to router 1b. The process of sending packet information within the same AS is known as intra-domain routing. In other cases, we might want to send a packet of information across ASs, this is known as interdomain routing. An example of interdomain routing would be the sending of packet information from router 1d to 3d. This packet information's route can either be AS2 AS3 or it can AS1 AS3. The route the packet information takes to get from a node on one AS to a node on another AS is determined by routing algorithms such as Dijkstra's algorithm. However, to send packet information from router 1a to router 3d, router 3d must first announce its location. This is done by using the Border Gateway Protocol (BGP). The BGP protocol allows for the sending of prefix advertisements between ASs. This is done by sending BGP connection messages over port 179. These connections can either be sent across ASs (known as eBGP connections) or within ASs (known as iBGP connections). Gateway routers (routers such as 1c, 2c, and 3a) sit at the very edge of an AS and relay the BGP connection messages to other ASs as well as routers inside their respective AS. For example, gateway router 2c would relay the BGP connection message to AS1 as well as inform the routers within its own AS (2a, 2b, 2c, and 2d) about the location of router 3d. While BGP is highly effective, it unfortunately lacks the assurance of security, thus presenting a significant problem.

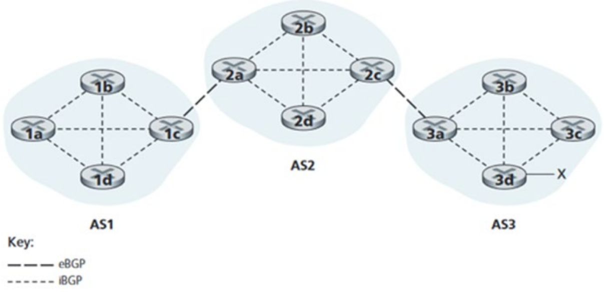

Figure 1[7]. Example of multiple ASes and routers in such ASes communicate with one another using eBGP and iBGP links.

Structure of BGP Message and Types of BGP Hijacking Attacks

There are various different types of routing attacks present in our world, some of which include DDOS (Denial of Service), TCP RST (TCP reset), Route flapping, and BGP hijacking. One example of a DDOS attack is when the time to live (TTL) for a packet is exceeded resulting in the router no longer forwarding the packet and subsequently dropping it. TCP RST occurs when a remote hacker resets the TCP connection between two routers. Route flapping occurs when a router continuously goes offline and online, resulting in the router's advertisements being continuously removed and advertised to its peers. While this is happening, the router is also adding advertised information from its peers into its routing table, as when the router goes offline its routing table resets. Lastly, BGP hijacking, the main focus of this paper, is when malicious ASes falsely claim ownership or advertise false routes to neighboring ASes.

BGP hijacking relies on the announcement of IP prefixes to other routers and other ASes. There are thousands of ASes present on the internet, and hundreds of routers that belong to such ASes. However, those routers do not know the existence, or path, of other ASes. To display the existence and route of an AS to other ASes, routers must use BGP advertising. BGP advertising is the process of conveying reachability and existence information to other ASes and the routers that comprise those ASes. To convey reachability and existence information, routers use iBGP and eBGP connections. An example of conveying reachability information could be taken from the image below. Consider advertising the reachability and existence information of the prefix “x” to all routers in AS2. To do this, router 3d sends an iBGP message to gateway router 3a. Then, gateway router 3a sends an eBGP message to gateway router 2c. Router 2c would then send this eBGP message to all routers (including router 2c) in AS2. This eBGP message would look something like “AS2 AS3 x”, indicating that to get to the prefix “x”, data must first travel through AS2 and then through AS3.

BGP messages are formatted in a specific way containing specific fields. The common size for a BGP message can be anywhere from 19 Octets to 4096 Octets. The format and size (in Octets) for the message fields can be seen in Figure 2 below.

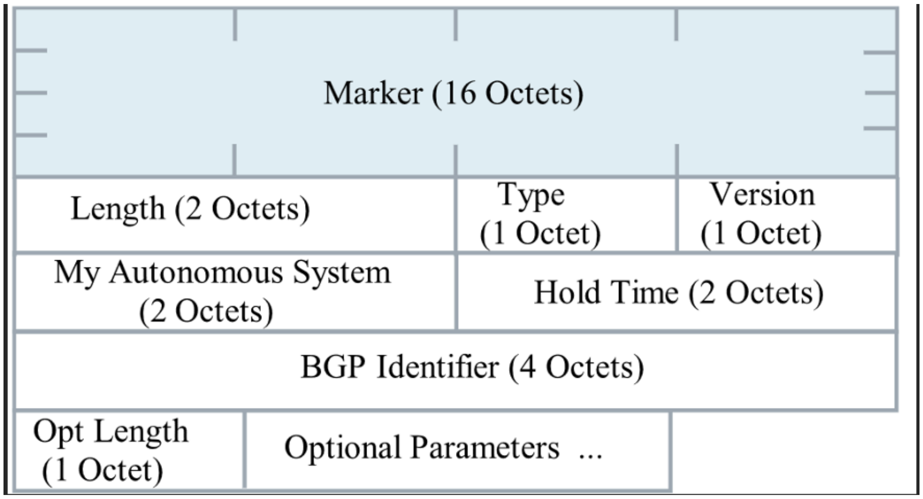

Marker (16 Octets)		
Length (2 Octets)	Type (1 Octet)	Version (1 Octet)
My Autonomous System (2 Octets)	Hold Time (2 Octets)	
BGP Identifier (4 Octets)		
Opt Length (1 Octet)	Optional Parameters ...	

Figure 2[1]. An example of the structure of a BGP message and the fields that make up such a message.

The first sixteen Octets dedicated to the “Marker” field mark the start of the message. The length field denotes the total message length. The type field can be one out of four possibilities: OPEN, UPDATE, NOTIFICATION, and KEEPALIVE. The “OPEN” message type is the first one sent after a TCP connection has been established between two routers. After this, the KEEPALIVE messages are sent periodically to ensure the connection between the two routers. A NOTIFICATION message states information regarding the termination of a session (during which information is exchanged) between two routers. Lastly, an UPDATE message is used to announce a new route, take back a previously announced route, or update the path to a current route.

The types of BGP attacks include *attacks against confidentiality*, *attacks against message integrity*, *fraudulent origin attacks*, and BGP hijacking. *Attacks against confidentiality* occur when two routers are sending information to one another and a third-party source chimes into the message stream and listens to the information being sent. The attacker does this to learn route policy and route information.

Attacks against message integrity occur when two routers are sending information and a third party intercepts and tampers with the data being sent. Examples of such tampering include message insertion (the sending of forged

BGP messages next to the message stream, resulting in the connection between the two routers shutting down), message deletion (when the third party removes the data being sent), and message modification (the third party intercepts the data being sent and changes it, then forwards the altered message to its desired location).

Fraudulent origin attacks (also known as prefix hijacking) occur when a malicious AS advertises a prefix that it does not own. As a result, neighboring ASes will believe that the malicious AS is the owner of such prefix, and will route necessary packets to the malicious AS. The true owner of the prefix will not receive the packets it was intended to receive. The malicious AS can choose to drop all the packets that are sent to it, this is known as a *black hole*. When a black hole occurs, the user would not be able to access a website or an application as the packet is never received by its designated location (for example, a Google server).

The types of BGP hijacking include prefix hijacking, sub-prefix hijacking, prefix and its AS hijack, sub-prefix and its AS hijack, and hijack of a legitimate path. The goal of prefix hijacking is to obtain information being sent to another AS by advertising a general element of that AS, resulting in information being sent to the malicious AS rather than the actual AS. Specifically, prefix hijacking occurs when an attacker configures a router to announce a prefix that belongs to another AS. For example, say there is some prefix “x” belonging to AS1 and a malicious AS called AS99. The attacker can accomplish prefix hijacking by announcing to neighboring ASes that it owns the prefix “x”. Neighboring ASes will send an UPDATE message to all of their routers, and their routers will compare the path length of this message to the current path length in the table. If the path length is shorter than the current in the table, only those neighboring ASes will update the route. Other neighboring ASes will disregard this UPDATE message and will continue to use the path that is currently in the routing table. Neighboring ASes who have updated the path for prefix “x” will route all information pertaining to prefix “x” to AS99 rather than AS 1. This attack is effective in that it allows the attacker to gain access to information that should have been sent to another AS. However, the downside of this type of BGP hijack is that the attacker is sometimes unable to get information from all ASes. As sometimes the advertised route has a longer path length than the one in the routing table. This issue is fixed in the “sub-prefix hijacking” attack

Another type of BGP hijacking is known as sub-prefix hijacking. A sub-prefix is a prefix that is contained inside a larger prefix. For example, prefix 10.10.0.0/16 holds prefix 10.10.0.0/24, thus the /24 prefix is a sub-prefix of 10.10.0.0/16. The goal of sub-prefix hijacking is to receive information belonging to another AS by advertising a specific element of that AS. More specifically, in this attack, the attacker announces a sub-prefix it does not own, resulting in information being sent to a malicious AS rather than the actual AS which the sub-prefix originates from. An example of sub-prefix hijacking can be taken from figure 3.

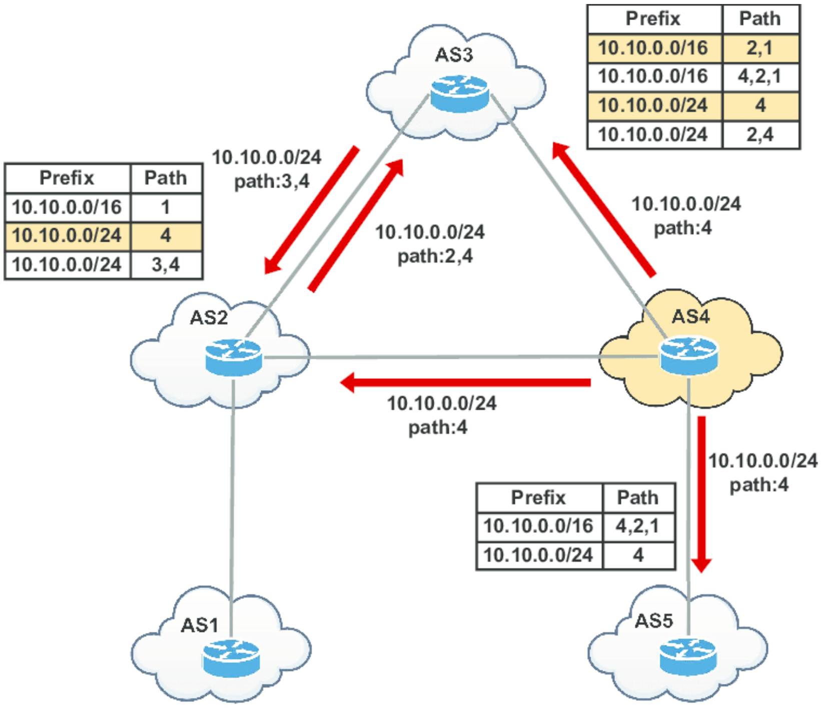

Prefix	Path
10.10.0.0/16	2,1
10.10.0.0/16	4,2,1
10.10.0.0/24	4
10.10.0.0/24	2,4

Prefix	Path
10.10.0.0/16	1
10.10.0.0/24	4
10.10.0.0/24	3,4

Prefix	Path
10.10.0.0/16	4,2,1
10.10.0.0/24	4

Figure 3[1]**.** An example of sub-prefix hijacking.

AS4 announces the sub-prefix 10.10.0.0/24, which belongs to AS1 and prefix 10.10.0.0/16, to its neighbors. All surrounding ASes will read this advertisement and inform their routers to update their routing table, stating that prefix 10.10.0.0/24 belongs to AS4 rather than AS1. Routers would then route information that needs to be sent to prefix 10.10.0.0/24 to AS4 rather than AS1. Even routers belonging to ASes that are directly connected to AS1, for example, AS2, will route information belonging to 10.10.0.0/24 to AS4 rather than AS1. This is because the path advertised by AS4 is more specific than the path advertised by AS1. This is because AS1 states that in order for information to be sent to 10.10.0.0/24 neighboring ASes need to send it over a more general prefix (10.10.0.0/16). Neighboring ASes would much rather send information to the prefix itself rather than send it over a general prefix. Thus, all information that needs to be sent to 10.10.0.0/24 is sent to AS4 rather than AS1.

The prefix and its AS hijack attack is one where a malicious AS announces to neighboring ASes that it has a direct connection to the victim AS. The main goal of this type of attack is to convince neighboring ASes that the malicious AS has a direct connection to the victim AS. The malicious AS will announce that it has a direct connection to the victim AS. Like the prefix and sub-prefix hijack, neighboring ASes will send an UPDATE message to their routers. The routers for such ASes will then compare the path length of such an UPDATE message to the path length currently in the routing table. If the UPDATE message from the malicious AS has the same or greater path length, then the routers on the AS will choose not to use the newly announced route. However, if the path length for the newly announced route is smaller than the one currently in the routing table, the router will accept this route and add it to the routing table. Once the router has added the route to its routing table, it would send all information that needs to be sent to the victim AS to the malicious AS first. The malicious AS can choose to black hole, alter, or view the information being sent. The malicious AS in this type of hijack can be seen as a “middle-man” between the AS that is sending information and the “victim AS”. An example of a prefix and its AS hijack can be further explained using figure 4.

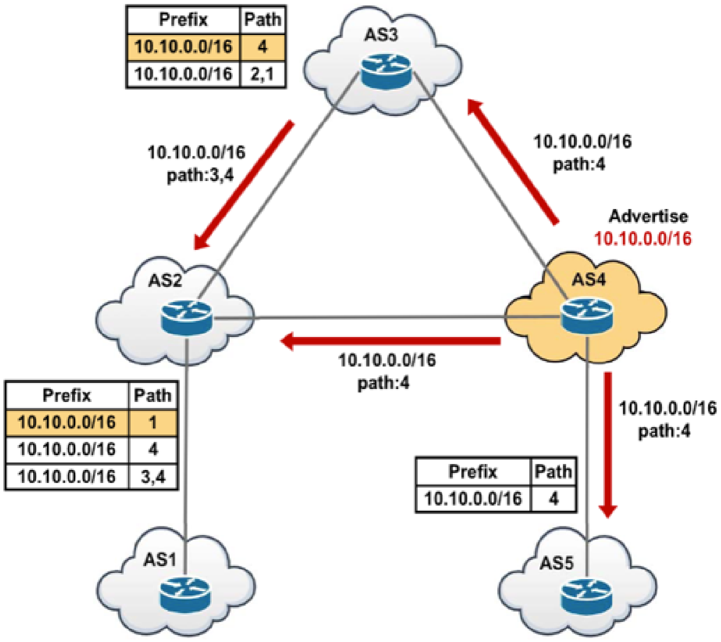

Prefix	Path
10.10.0.0/16	4
10.10.0.0/16	2,1

Prefix	Path
10.10.0.0/16	1
10.10.0.0/16	4
10.10.0.0/16	3,4

Prefix	Path
10.10.0.0/16	4

Figure 4[1]. Example of a prefix and its AS hijack attack.

AS4 would announce to neighboring ASes that it has a direct connection to AS1. The neighboring ASes will receive this UPDATE message and compare it to the route currently in the routing table. AS5 will accept this as a new route as it has a shorter path length than the previous route in the table. Thus, when AS5 wants to send information to AS1, it would be sent to AS4 first and then to AS1, resulting in the path (5,4,1). However, AS2 and AS3 will not use the newly advertised route as it has the same (if not longer) path length as the one currently in the table.

The last type of BGP hijacking attack is known as sub-prefix and its AS hijack. The goal of this attack is to announce a fake path to a victim prefix. Like the prefix and AS hijack attack, a malicious AS sends out an UPDATE message to its peers (surrounding routers) and informs them that it has a direct connection to a sub-prefix on another AS. Surrounding ASes will send this UPDATE message to its routers and they will immediately replace it for the current route in the routing table. Even if this new path is longer compared to the old path, routers will replace it as the new path is more specific than the old path. An example of a sub-prefix and its AS hijack could be taken from figure 5.

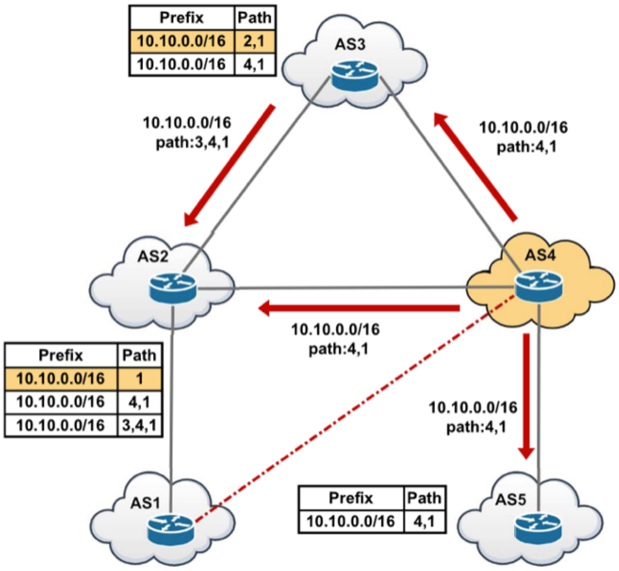

Figure 5[1]. Example of a sub-prefix and its AS hijack attack where the prefix and paths for a particular route is given in the form of a table.

AS4 announces that it has a path to 10.10.0.0/24 (a sub-prefix which belongs to AS1). AS2, 3, 4, and 5 will take the UPDATE message sent out by AS4 and distribute it across all routers contained in that respective AS. The routers would then compare this "new" route to the one currently in the table and replace the one currently in the table with the new route (due to specificity). All ASes will replace the old route for the "new route", even if the "new route" has a longer path length than the "old route". Once packets are sent to the malicious AS from surrounding ASes, the attacker can choose to alter or even blackhole the packets. This attack is beneficial in that it targets and gets information from all ASes.

Current Detection and Mitigation Techniques and How They Function

Instances of BGP hijacking are more common today than ever before. In 2020 alone, there were 2,477 BGP hijacking events recorded, with a peak in April. In order to reduce and protect the system from BGP hijacking attacks, engineers developed techniques to help with the detection and mitigation of BGP hijacking attacks. An example of such a technique could be the Route Origin Validation (ROV). ROV relies on the Resource Public Keys Infrastructure (RPKI) framework and is used on border routers to reject any advertisements that are marked as RPKI-invalid. If something is marked as RPKI-invalid it means that the advertised prefix does not contain origin AS (the AS that is advertising the prefix) or the AS number for origin AS does not match the AS number for the prefix. In order to determine if an advertised route is RPKI-invalid or not, the router queries a database of validated prefix-to-AS mappings, downloaded from a cache server, and checks whether the advertised prefix belongs to the announcing AS. Should this be returned false, the route would register as RPKI-invalid and the router would not accept the advertised route. However, in many cases the settings for ROV can be misconfigured, resulting in a router possibly invaliding a valid route or validating an invalid route. If an attacker hijacks a system and advertises a false route, ROV could possibly validate such a route, which could potentially bring down the entire system. In such cases, engineers need to be able to recognize that their system is under attack. To do this, engineers use a commercial tool provided by ISPs known as a BGP monitoring service. These services alert engineers via forms of communication such as Slack or Microsoft Teams that their system is under attack. Such services are also tied to a JSON file which contains information about the affected

ASes and prefixes. The alert via popular forms of communication paired with a JSON file identifying which ASes and prefixes are under attack allows engineers to solve the issue and prevent further damage to the system. However, human intervention, while sometimes beneficial, is slow and highly prone to error. Systems should rely on reactive detection-assisted mitigation software rather than human intervention to stop and reduce BGP hijacking attacks. To do so, systems use a recently developed detection and mitigation technique known as Bogus route purging. This technique relies on the use of pre-selected life-saver ASes (typically large ISPs that act as pivotal nodes through which data traffic traverses en route to a designated destination) which delete the "bogus route" from its routing table. Due to the large size of such ASes, the spreading of the "bogus route" is throttled. If the life-saver AS acts fast enough, the spreading of such a "bogus route" can be stopped almost immediately.

Methods

In order to search for relevant literature I queried the Google Scholar database and searched for keywords such as "BGP hijacking detection" and "BGP hijacking". The only filter I applied while searching for relevant literature is the "Articles" filter. The publication date of the literature I was viewing did not matter for my topic as BGP hijacking detection technology has not been updated since the early 2000s, and all the literature I viewed had been published in that time.

To select studies and literature, I opted to look for survey and review papers rather than research papers. Research papers in this field simply discuss the effects of such detection techniques and not how such detection techniques function at a high level. Other factors considered while selecting a study were the relevance of the study relative to my topic as well as the reliability of such as study.

To ensure the reliability of a study, all literature was either drawn from Google Scholar or university-level books. Synthesis of evidence for this paper was done by gathering multiple pieces of evidence and drawing a connection and conclusion from one piece of evidence to the focus of the paper. This allowed me not only to gain a deeper understanding of the material but also showed me how one piece of evidence connected to another. Within the paper itself, evidence was organized based on the section of the paper and how that evidence relates to the discussion of the section. For example, all evidence relating to the background section of the paper was placed in the background section and was checked to ensure the evidence was relevant to the background section. The quality of evidence was assessed by seeing if the ideas and conclusions present in the evidence agree with other pieces of evidence relative to the same topic. If the results of one study say one arise to one conclusion and the results of a similar study arise to the same conclusion, then the original study is credible. However, if the findings two studies differ, the original study is not credible. This ensures the reliability of the evidence used.

Quality was also assessed by the institution from which the evidence came from. The institution had to be creditable, for example, universities and RFCs.

Conclusion

Techniques in order to reduce and hopefully stop BGP hijacking are still being developed, refined, and used today (for example the RPKI framework). We see multi-billion dollar companies, such as Cloudflare, invest millions of dollars in order to make a network more secure. However, just because time has passed and security has become better does not mean that BGP hijacking, and any network attack for that matter, has ceased to exist. In 2020 alone there were 2,477 instances of BGP hijacking attacks. Additionally, in 2019 we see Allegheny Technologies leak many false routes to the internet, resulting in lots of information being misdirected to Allegheny themselves. Even if BGP hijacking attacks have occurred in the past couple years, it does not mean that BGP security had not improved since the development of the border gateway protocol. Developments such as ROV and the RPKI framework have left a drastic impact on BGP security and whether or not a route is accepted.

However, just because BGP security has come a long way does not mean that it can still improve. In recent years we have seen a massive development in the field of artificial intelligence (AI). We see AI models such as ChatGPT and Gemini, developed by OpenAI and Google respectively, make life increasingly efficient. However, we also see AI models creating new ways to solve problems and taking old ideas and making them even more efficient. In the past couple years, Deep Mind's AI model has developed a new, and efficient way, to multiply two matrices. I believe as time goes on we will see AI and machine learning models being implemented into BGP and overall network security. These models would be extremely efficient as they would have the capability to learn from their mistakes, thus making a network even more secure than it is today.

Acknowledgments

I would like to thank my advisor Rucha Vaidya for guiding me through this project and helping me understand my topic of research. I would also like to thank the Empowerly research team for guiding me through the publishing process and reviewing my writing.

References

[1] Al-Musawi, B., Branch, P., & Armitage, G. (2017). BGP Anomaly Detection Techniques: A Survey. IEEE Communications Surveys & Tutorials, 19(1), 377–396. *https://doi.org/10.1109/comst.2016.2622240*
[2] BGP Hijacking: Understanding Threats to Internet Routing. (2023, July 19). Kentipedia. *https://www.kentik.com/kentipedia/bgp-hijacking/*
[3] Bush, R., & Austein, R. (2013, January 1). The Resource Public Key Infrastructure (RPKI) to Router Protocol. IETF. *https://datatracker.ietf.org/doc/html/rfc6810*
[4] Butler, K., Farley, T. R., McDaniel, P., & Rexford, J. (2010). A Survey of BGP Security Issues and Solutions. Proceedings of the IEEE, 98(1), 100–122. *https://doi.org/10.1109/jproc.2009.2034031*
[5] IP Routing: BGP Configuration Guide, Cisco IOS XE Release 3S - Configuring Multiprotocol BGP (MP-BGP) Support for CLNS [Cisco IOS XE 3S]. (n.d.). Cisco. Retrieved November, 2023, from *https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/configuring_multiprotocol_bgp__mp-bgp__support_for_clns.html*
[6] Kent, S., Lynn, C., & Seo, K. (2000). Secure Border Gateway Protocol (S-BGP). IEEE Journal on Selected Areas in Communications, 18(4), 582–592. *https://doi.org/10.1109/49.839934*
[7] Kurose, J. F., & Ross, K. W. (2021). Computer networking: a top-down approach (8th ed.). Pearson.
[8] Siddiqui, A. (2021, February 5). BGP, RPKI, and MANRS: 2020 in review. MANRS. *https://manrs.org/2021/02/bgp-rpki-and-manrs-2020-in-review/*
[9] Siddiqui, A. (2022, February 21). BGP Security in 2021. MANRS. *https://manrs.org/2022/02/bgp-security-in-2021/*
[10] Shapira, T., & Shavitt, Y. (2022, April 11). AP2Vec: An Unsupervised Approach for BGP Hijacking Detection [Review of AP2Vec: An Unsupervised Approach for BGP Hijacking Detection]. IEEE Xplore; IEEE. *https://ieeexplore.ieee.org/abstract/document/9754706*. DOI: 10.1109/TNSM.2022.3166450
[11] What Is a Network Node? - IT Glossary | SolarWinds. (n.d.). *Www.solarwinds.com*. *https://www.solarwinds.com/resources/it-glossary/network-node*