Christian Matlock

ICTN 4040

Dr. Li

April 10, 2017

BGP Hijacking and Mitigation Techniques

Border gateway protocol (BGP) is an exterior gateway protocol which is used to route between two autonomous systems (AS) to make up the global internet. The most current form of BGP, BGP4, has been in use on the internet since 1994 and was initially developed using a system of trust and security was not a main priority.

Due to the nature of how BGP operates and how it was designed, a series of both unintentional mistakes or malicious attacks could take place to take down an entire autonomous system; or a more nefarious act of intercepting traffic and then routing it to the correct destination could also take place. In order to combat the trusting nature of BGP, a series of threat mitigation techniques have been implemented in order to protect BGP against attacks. This includes setting up BGP neighbor authentications, filtering BGP prefixes with AS path access lists, BGP time to live (TTL) security check, and the future resource public key infrastructure (RPKI) have been put in use.

To fully understand how to protect and harden BGP routing tables in your enterprise networks, we must first discuss how BGP operates. The main concept of BGP is the fact that it routes between different AS which are owned by entire organizations such as Google, Amazon, etc and within those AS are networks consisting of interior gateway protocols such as open shortest path first(OSPF), enhanced interior gateway protocol (EIGRP), or even interior BGP itself. For a router to receive BGP prefixes from another router it must first create a BGP peering

and a neighbor relationship with that neighboring router. This is accomplished by creating a TCP session on port 179 and utilizes TCP's reliable nature with the concept of a three-way handshake and the acknowledgement of packets (Beijnum 13).

Once when a neighborship is formed via TCP and the BGP state is in the "established" state, the routers can then share prefixes via BGP update messages. BGP then selects the best path to use and install in the router information base (RIB) table; the default setting BGP uses to select this path is via hop count. However, these metrics can be manipulated and changed to meet an organization need using concepts such as route maps (Beijnum 36).

Attacks on BGP which could cause it to select a non-optimal, unintended, or a route which causes it to be blackholed could occur due to a mistake or malicious intent. This is a consequence due to the system of trust that BGP was inherently designed with. For instance, an enterprise customer edge (CE) router or tier 3 internet service provider will take the word of any prefixes that it learns from its upstream neighbor and the best path on how to get there. If one service provider or enterprise network could advertise a whole range of prefixes out, this could cause a wide outage or man in the middle attack for a large portion of the internet. Butler describes attackers "sometimes introduce false information into BGP to enable them to exchange e-mail with mail servers using unallocated IP addresses that are hard to trace" (Butler 1).

This was the case for Google's YouTube service in February 2008 when the country of Pakistan attempted to block YouTube for its entire nation to dissuade dissidents from sharing information on the platform. Unfortunately, the prefixes it was advertising were not filtered outbound and they were leaked to their upstream service provider and subsequently over two-thirds of the world received a bogus null route to YouTube and knocked out access (Butler 10). This type of mistake or intentional act is commonly called BGP hijacking. Instead a mechanism

must exist for organizations to protect against faulty prefix advertisements or from unintended prefixes from being announced outside of their organization.

Another potential attack vector for criminals to conduct a man in the middle attack, utilizing BGP, to intercept an organization's traffic would be to create a BGP peering session with a target to announce prefixes that it would like to intercept. This type of attack could be accomplished through spoofing a legitimate neighbor or exploiting a misconfigured customer edge router. Also, an attacker could exploit the fact that "messages could introduce incorrect information into the routing system or trigger routers to abort the session" (Nicholes 15) To counter this attack, an organization could implement BGP MD5 authentication in which the neighboring routers exchange a shared key with all routing updates. The receiving router computes that router's update with the shared key to come up with a MD5 hash; if that hash is not an identical value from the one that it has configured, then it will drop that BGP update packet.

To prevent an enterprise's router from receiving BGP updates about prefixes that they don't wish to receive or send out prefixes that are only meant to be internal; a Network Administrator could use prefix filtering techniques to achieve this goal. One way to accomplish this would be to create an IP prefix-list, which through a series of sequence lists, you can match IP ranges and select an action such as deny or permit. Once the prefix list has been created you can apply the prefix-list to the BGP router process in an inbound or outbound direction to achieve your goal. Another method utilizes the route-map concept in which you can create an access control lists (ACL), apply them to a route map, and then apply a series of attributes such as setting the priority or even dropping updates for a certain prefix (Beijnum 18).

With the concept of BGP TTL checks, you can prevent your router from accepting BGP updates from routers that are further way from the router that you are peering with. This could be used to prevent a neighbor spoofing attack.

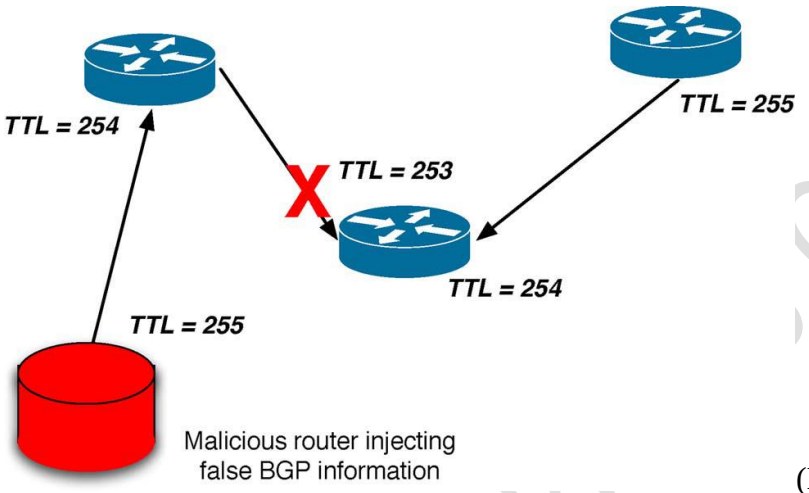

(Hutson 12)

Similar to the concept that regular IP packets utilize, each router or hop in a packets path decrements the TTL field by one. When a router receives a packet and the TTL reaches 0, the packet is immediately discarded. By enabling BGP TTL security, a network administrator can set the exact TTL value that the valid BGP peer is away from their router, thus preventing the BGP spoofing attack. However, as Kevin Butler, describes the pitfalls in this security mechanism that it "weakly defends against attackers who are more than one hop away. It does not defend against subverted peers sending malicious information or other similar insider attacks" (Butler 27). With this limitation in mind, it highlights the need to consider every avenue of attack and plan to utilize multiple security principles to properly secure BGP.

Many leading researchers in the global internet infrastructure speculate that the implementation of public key infrastructure (PKI) in BGP will be the ideal solution moving forward. Similar to how PKI is utilized in the HTTPS protocols and many new secure messaging

applications that have been developed lately. The general idea that "every AS has a public key, distributed freely to any other AS in the Internet, and a private key, which is never divulged" (Butler 32). The two routers that are trying to peer and create a relationship with each other can then use the Diffie-Hellman key exchange process to exchange public keys and then begin communicating via BGP updates (Bruhadeshwar 5). By implementing the PKI in BGP, this would give organizations a clear hierarchy and centralized way to ensure that enterprises that own a certain autonomous system can advertise out to other neighbors.

With society's ever increasing reliance on traffic via the unsecure internet, it is becoming important that organizations hire professionals that fully understand how BGP is implemented in their enterprise networks. Criminals on the internet are starting to utilize BGP hijacking to route traffic to their location to preform man-in-the-middle attacks to steal information or shut down access to an autonomous system. Network administrators will need to fully understand the strengths and weaknesses of each BGP security technique and technology to understand fully how it could be implemented into their own network. Enterprise network administrators have a variety options such as neighbor authentication to ensure that they are exchanging prefixes with a valid neighbor, route filtering to prevent the leaking of unwanted prefixes onto the global internet or prevent bogus updates from valid neighbors, and TTL checks to ensure that neighbors distant on the internet cannot spoor and inject malicious information. Brand new techniques such as the implementation of the public key infrastructure will need to be widely adopted to ensure trust in a routing protocol that was not build with security in mind. With a wide range of threat mitigation tactics, organizations can ensure that their information stays safe on the internet.

Works Cited

Beijnum, Iljitsch Van. BGP. Beijing: O'Reilly, 2002. Print.

Bruhadeshwar, Bezawada, Sandeep S. Kulkarni, and Alex X. Liu. "Symmetric Key Approaches to Securing BGP—A Little Bit Trust Is Enough." IEEE Transactions on Parallel and Distributed Systems 22.9 (2011): 1536-549. Web.*

K. Butler, T. R. Farley, P. McDaniel and J. Rexford, "A Survey of BGP Security Issues and Solutions," in Proceedings of the IEEE, vol. 98, no. 1, pp. 100-122, Jan. 2010.

Huston, Geoff, Mattia Rossi, and Grenville Armitage. "Securing BGP — A Literature Survey." IEEE Communications Surveys & Tutorials 13.2 (2011): 199-222. Web.*

Nicholes, Martin, and Biswanath Mukherjee. "A survey of security techniques for the border gateway protocol (BGP)." IEEE Communications Surveys & Tutorials 11.1 (2009): 52-65. Web.