

USING DIFFERENT ATTRIBUTES IN BGP MULTIHOMING ROUTING

Milan Pavlović¹ Marija Zajeganović² Dragan Ristić³

Abstract: The Border Gateway Protocol (BGP) serves as a fundamental routing protocol for exchanging route information between autonomous systems on the Internet. Its application is essential in enabling multihoming configurations, allowing corporate networks to maintain connections with multiple Internet Service Providers (ISPs) simultaneously. This ensures improved reliability and load balancing for network operations. The introductory section provides an overview of BGP's functionalities and highlights its key characteristics, followed by an exploration of techniques for establishing multihoming setups. In the practical segment, a simulated network scenario demonstrates the use of BGP in connecting a private network to two different ISPs, utilizing the GNS3 emulator. The focus is placed on advanced configuration capabilities, including various attributes such as route optimization, confederation setup, and route mapping. Finally, particular attention is given to network administration, including identifying and resolving potential issues to maintain a robust and efficient network infrastructure.

Key words: BGP, multihoming, BGP attributes, route maps, route optimization, confederation

1. INTRODUCTION

Dynamic routing protocols are implemented on routers to enable them to exchange routing information (routing updates) about networks. Through these protocols, routers establish neighbor relationships, communicate, and share specific parameters. Based on the received information, each router maintains its own routing table [1]. This table contains routes to remote networks that the router has learned. The primary advantage of dynamic routing protocols is their ability to automatically update and adapt routing tables in response to network changes, reducing the need for manual intervention by administrators. The global computer network consists of numerous independent networks, referred to as Autonomous Systems (AS). These systems are composed of groups of routers managed by a single network administration team. Each AS establishes its own routing policy, manages interconnections with other systems, and regulates the flow of traffic through its network. Within an AS, decisions are made to optimize data routing across the complex infrastructure of computer networks.

Routing protocols are classified based on whether routing information is exchanged within an AS or between different ASs. Interior Gateway Protocols (IGPs) manage routing within a single AS, utilizing specific metrics to determine the best paths within the network. Regardless of how many IGPs are deployed, the AS presents itself as a cohesive entity in routing interactions with other AS. Exterior Gateway Protocols (EGPs), such as the Border Gateway Protocol (BGP), facilitate the exchange of routing information between ASs [2]. Every AS is assigned a unique Autonomous System Number (ASN) by IANA (Internet Assigned Numbers Authority) or an appropriate ISP, enabling its identification in the global network. BGP serves as the backbone for global routing on the Internet. It propagates network advertisements, Network Layer Reachability Information (NLRI) across ASs. BGP operates based on policy rules and follows the hop-by-hop paradigm. This means that a BGP router constructs its own routing table with next-hop addresses based on its policies but does not dictate how the next-hop router will forward the traffic toward the destination [2]. Establishing a BGP session between two routers is indispensable in scenarios where corporate networks use multihoming—simultaneous connections to multiple ISPs—to achieve higher reliability and fault tolerance [3]. Beyond its core functionality, BGP supports advanced configuration options, such as route optimization through policy-based routing, the use of confederations for managing large

¹ Professor of applied studies, The Academy of Technical and Art Applied Studies Belgrade, Department School of Information and Communication Technologies, Zdravka Čelara 16, Belgrade, 11000, Republic of Serbia, email: milan.pavlovic@ict.edu.rs

² Professor of applied studies, The Academy of Technical and Art Applied Studies Belgrade, Department School of Information and Communication Technologies, Zdravka Čelara 16, Belgrade, 11000, Republic of Serbia, email: milan.pavlovic@ict.edu.rs

³ Head of IT security planning section, Telekom Srbija, ristic.dragan@gmail.com

networks, and the implementation of route maps for precise traffic control [4]. These features make BGP a scalable and robust solution for managing over a million IPv4 routes currently present on the Internet.

This study is structured to provide an overview of BGP's fundamental characteristics and functionalities in the introductory chapters, while subsequent chapters address the configuration of multi-homed networks. The practical section includes a simulation of a private network connected to two ISPs using the GNS3 emulator. Special attention is given to advanced BGP configurations, the use of attributes. The conclusion offers guidelines for further research and refinement of the presented solutions.

2. BGP ATTRIBUTES

BGP attributes are a set of parameters that contain specific information about individual routes. These attributes include path attributes, route preferences, next-hop, and aggregation parameters. These attributes are used in the BGP decision-making process to determine which route is valid and will be selected for routing. Once selected, the route is placed in the Loc-RIB for advertisement and then installed into the FIB for actual packet forwarding. The attribute format includes three fields: attribute type, attribute length, and attribute value. BGP attributes are categorized into two main groups: well-known and optional attributes [2].

Well-known attributes must be supported by every BGP implementation. Among them are well-known mandatory attributes, which must be included in every BGP update, and well-known discretionary attributes, which may or may not be included in BGP updates. Optional attributes, on the other hand, are not required to be supported by BGP. Optional attributes are further divided into optional transitive attributes, which must be passed to neighboring routers, and optional non-transitive attributes, which are not propagated to adjacent BGP routers.

There are eight basic BGP attributes. Origin Attribute (Type Code: 1) is a well-known mandatory attribute used to indicate how the route was introduced into the BGP process. It can carry one of three values: IGP - The route was redistributed from an Interior Gateway Protocol; EGP - The route was learned via an Exterior Gateway Protocol; Incomplete - The route was manually injected into the BGP process and advertised further. These values prioritize routes as follows: IGP > EGP > Incomplete. The Origin attribute ensures that the advertised route reflects its true visibility within the network. AS_Path Attribute (Type Code: 2) is another well-known mandatory attribute that contains a list of AS numbers that the route advertisement has traversed. Each BGP router in an AS adds its AS number to the AS_Path attribute when advertising a route to a router in a different AS.

Next_Hop Attribute (Type Code: 3) is a well-known mandatory attribute that specifies the IP address of the next-hop router responsible for forwarding packets toward the destination. Multi-Exit Discriminator (MED) Attribute (Type Code: 4) is an optional non-transitive attribute used to influence how neighboring ASs should route traffic when multiple links exist between two ASs. By setting the MED value, an AS can suggest a preferred entry point for inbound traffic. Each router makes a decision for itself. The Local Preference (Local_Pref) Attribute (Type Code: 5) is a well-known discretionary attribute used within an AS to prioritize preferred routes. It is advertised to all routers within the same AS but not shared with external neighbors. Routes with higher Local_Pref values are preferred. Local_Pref is generally set on inbound advertisements, and its value influences outbound traffic decisions.

Furthermore, Atomic_Aggr (Atomic Aggregate) Attribute (Type Code: 6) is a well-known discretionary attribute used for route aggregation. Aggregator Attribute (Type Code: 7) is an optional transitive attribute that records the AS number and router ID of the device that performed the aggregation. Finally, Community Attribute (Type Code: 8) is an optional transitive attribute that assigns a tag (community value) to a route. Community attributes are widely used for advanced routing decisions and policy enforcement.

3. BGP DECISION PROCESS

For each NLRI route received from a BGP neighbor, the route undergoes the BGP Decision Process [4]. The process starts by verifying the validity of the received route, primarily through the examination of the Next-hop and AS_Path attributes. Next, it checks whether there is an existing route to the same destination already learned and recorded in the routing table, i.e., the Loc-RIB. If no such route exists, the received route is directly written into the routing table. However, if there is an existing route to the same destination, the received route is compared with the existing one based on their attributes in a specific order. The decision is then made on which route will be preferred.

Initially, the validity of the received route is determined by inspecting the Next-hop attribute, followed by the AS_Path attribute. The process verifies whether the next-hop in the attribute is reachable, i.e., whether the next-hop router is accessible for advertising the route. To check the next-hop address, the routing table is referenced, and it may indicate that accessing the next-hop address requires forwarding packets to another third-party next-hop address. This demonstrates that the next-hop attribute does not necessarily have to refer to the directly visible neighboring router. The AS_Path attribute is the next to be verified. This involves checking whether the AS_Path attribute contains the local AS number of the BGP router that received the route. If the local AS number is present, it indicates an AS loop, and the route is ignored due to the detected routing loop. Once the route is confirmed as valid, the next step is to check whether there is another route to the same destination already in the local RIB. If no such route exists, the received route is directly written into the routing table, i.e., the local RIB. However, if another route to the same destination exists in the local RIB, the routes are compared based on their attributes. The BGP Decision Process defines a strict order of attribute evaluation, determining which attribute has higher priority in deciding the best route [5].

The first attribute analyzed is the WEIGHT attribute, which is optional and specific to Cisco routers. This attribute is valid only locally and is not advertised to other routers. It enables the administrator to prioritize one route over another on the local router. Different WEIGHT attributes can be defined on different routers, allowing distinct decision-making policies. If two routes have the same WEIGHT attribute or do not have it defined, the process moves to the next rule: the Local Preference (Local Pref) attribute. Routes with higher Local Pref values are preferred. Local Pref is manually configured by the administrator. If the routes have the same Local Pref value, the process considers the AS_Path attribute. The route with the shorter AS_Path (fewer AS numbers in the sequence) is preferred. If the AS_Path lengths are equal, the route with the lower ORIGIN attribute (IGP/EGP/Incomplete) is preferred. ORIGIN specifies how the NLRI route was introduced into the BGP process, with IGP being the most preferred, followed by EGP, and finally Incomplete, which indicates manually configured routes. If the ORIGIN attribute is also equal, the MED attribute is considered. Routes with lower MED values are preferred. MED represents a metric where smaller values indicate better routes.

If the MED values are the same, the route received via an external BGP (eBGP) connection is preferred over one received via an internal BGP (iBGP) connection. eBGP connects links between different ASs, whereas iBGP connections exist within the same local AS. If all these attributes are identical, the process considers the Next-hop attribute, specifically which next-hop is closer. A route with a directly reachable next-hop address is preferred over one requiring an additional hop through another router. If the Next-hop attributes are also identical, the router ID is used as the tiebreaker, with the route having the lower router ID being preferred.

In conclusion, the BGP Decision Process emphasizes that BGP routing policies are hop-by-hop and independent. The routing policy implemented on one BGP router is entirely independent of those implemented on other BGP routers, even in different ASs. These routing policies do not need to be identical across BGP routers.

4. NETWORK REALIZATION

In this paper, the topology of a computer network is implemented using the GNS3 emulator and is illustrated in Fig. 1. The presented topology is based on the "Multi-homed" concept, where a single AS is connected to an Internet Service Provider (ISP) via a primary and a backup link. The user is part of

AS3, while the ISP operates within AS1. The user is connected to the ISP through two border routers, R3 and R4, using two separate physical links, X1 and X2.




Figure 1 - One multi-homed computer network realized in GNS3 emulator

The ISP establishes connectivity with the user in AS3 via its two border routers, R1 and R2. As illustrated in Fig. 1, link X1 connects router R4 (the user's router in AS3) with router R2 (the ISP's router in AS1), while link X2 connects router R3 (the user's router in AS3) with router R1 (the ISP's router in AS1). Within the user's AS3, an internal iBGP session is established between routers R3 and R4. The user's routers are connected to local networks, specifically the 172.16.220.0/24 network for router R3 and the 172.16.65.0/24 and 172.16.10.0/24 networks for router R4. Router R3 is also configured with a Loopback interface, L0, assigned the IP address 172.16.2.254. On the ISP side, within AS1, an iBGP session is also established between border routers R1 and R2. Through router R1, the ISP maintains connectivity with another user's AS. Specifically, router R1 in AS1 is connected to router R in AS6. On the other hand, through router R2, the ISP connects to a Network Access Point (NAP). From there, the ISP in AS1 accesses the 193.78.0.0/16 network, which provides connectivity to the rest of the global network. Between routers belonging to different ASes, external eBGP sessions are established.

The routing policy in this scenario is designed to ensure that the user's AS3 accepts only those BGP advertisements related to the local prefixes advertised by its ISP in AS1, as well as prefixes advertised by other customers of the same ISP, such as the prefixes from AS6. Additionally, AS3 will accept a single route from the Internet, which will serve as the default route pointing to the ISP in AS1. Regarding the routing policy for outbound traffic from AS3, all traffic destined for networks within the ISP's AS1 and networks in the other customer's AS6 must be forwarded exclusively through link X2. In the event of a failure on link X2, this traffic will be redirected to link X1. Other outbound traffic from AS3 destined for the global Internet should primarily be forwarded via link X1 using the default route. If link X1 fails, the default route for the remaining outbound traffic will be redirected through link X2. The inbound traffic policy must meet the requirement that the ISP routes traffic destined for the 172.16.220.0/24 network in AS3 via the primary link X2, while all other traffic destined for AS3 is routed via link X1, which serves as the primary link in this case. This approach achieves partial advertisement of the BGP table and ensures a balanced utilization of available resources.

Figure 2 illustrates the configurations on routers R3 and R4 that fulfill the specified routing policies. First, the BGP process was enabled on both routers. Synchronization was disabled, and the networks from AS3 that R3 and R4 will advertise through the BGP process were announced. An internal iBGP session was then configured between routers R3 and R4. On router R3, the next-hop address for the iBGP neighbor was set to 172.16.1.2, which is the IP address of Ethernet interface E0 on router R4. Similarly, on router R4, the next-hop address for the iBGP neighbor was set to 172.16.2.254, which is the IP address of Loopback interface L0 on router R3, used for the iBGP session between R3 and R4. Additionally, commands were added to the iBGP configuration to ensure that routers R3 and R4 advertise all prefixes exchanged between them with their own IP addresses in the NEXT_HOP attribute. External eBGP sessions were also configured with routers R1 and R2 in AS1, specifying the next-hop addresses of neighbors on the opposite ends of links X1 and X2. Two route maps were applied to these eBGP sessions. The first, named SET_OUTBOUND_TRAFFIC, was

applied to inbound advertisements (In direction) to control outbound traffic from AS3. The second, named SET_INBOUND_TRAFFIC, was applied to outbound advertisements (Out direction) to control inbound traffic to AS3. In addition to these route maps, filter list 10 was applied to outbound advertisements from routers R3 and R4 to the neighboring AS1.

<pre> R3 router bgp 3 no sunchronization network 172.16.1.0 mask 255.255.255.0 network 172.16.10.0 mask 255.255.255.0 network 172.16.65.0 mask 255.255.255.0 network 172.16.220.0 mask 255.255.255.0 neighbor 172.16.1.2 remote-as 3 neighbor 172.16.1.2 update-source loopback0 neighbor 172.16.1.2 next-hop-self neighbor 172.16.20.1 remote-as 1 neighbor 172.16.20.1 route-map SET_OUTBOUND_TRAFFIC in neighbor 172.16.20.1 route-map SET_INBOUND_TRAFFIC out neighbor 172.16.20.1 filter-list 10 out no auto-summary ! ip route 0.0.0.0 0.0.0.0 193.78.0.0 ip as-path access-list 10 permit ^\$ ip as-path access-list 4 permit ^1\$ ip as-path access-list 4 permit ^1 6\$ ip as-path access-list 4 permit ^1 ?[0-9]*\$ access-list 2 permit 172.16.220.0 0.0.0.255 access-list 101 permit ip 193.78.0.0 0.0.255.255 255.255.0.0 0.0.0.0 ! route-map SET_OUTBOUND_TRAFFIC permit 10 match ip address 101 set local-preference 200 route-map SET_OUTBOUND_TRAFFIC permit 20 match as-path 4 set local-preference 300 ! route-map SET_INBOUND_TRAFFIC permit 10 match ip address 2 set local-preference 200 route-map SET_INBOUND_TRAFFIC permit 20 set metric 300 </pre>	<pre> R4 router bgp 3 no sunchronization network 172.16.1.0 mask 255.255.255.0 network 172.16.10.0 mask 255.255.255.0 network 172.16.65.0 mask 255.255.255.192 network 172.16.220.0 mask 255.255.255.0 neighbor 172.16.2.254 remote-as 3 neighbor 172.16.2.254 next-hop-self neighbor 192.68.5.2 remote-as 1 neighbor 192.68.5.2 route-map SET_OUTBOUND_TRAFFIC in neighbor 192.68.5.2 route-map SET_INBOUND_TRAFFIC out neighbor 192.68.5.2 filter-list 10 out no auto-summary ! ip route 0.0.0.0 0.0.0.0 193.78.0.0 ! ip as-path access-list 10 permit ^\$ ip as-path access-list 4 permit ^1\$ ip as-path access-list 4 permit ^1 6\$ ip as-path access-list 4 permit ^1 ?[0-9]*\$! access-list 101 permit ip 193.78.0.0 0.0.255.255 255.255.0.0 0.0.0.0 ! route-map SET_OUTBOUND_TRAFFIC permit 10 match ip address 101 set local-preference 250 route-map SET_OUTBOUND_TRAFFIC permit 20 match as-path 4 set local-preference 250 ! route-map SET_INBOUND_TRAFFIC permit 10 set metric 250 </pre>
--	--

Figure 2 - Realized BGP configuration in routers R3 and R4

Beyond the BGP configuration, a static default route was configured on both R3 and R4, specifying the next-hop address for the 193.78.0.0 network prefix. This binds the default route to the 193.78.0.0 prefix received by router R3 via the BGP protocol. This setup directs all outbound traffic not destined for networks in AS1 or AS6 toward the Network Access Point (NAP) or Internet Exchange Point (IXP) network hosting the 193.78.0.0 prefix. If the 193.78.0.0 prefix is available via a BGP session, it is assumed that the entire Internet is accessible. This configuration does not specify a particular link for default outbound traffic; instead, the direction is determined based on the route from which the advertisement for the 193.78.0.0 prefix was received. Two AS_Path access lists were configured on R3 and R4. The first, AS_Path access list 10, permits all prefixes with an empty AS_Path attribute "^\$". An empty AS_Path indicates that the prefix originated from AS3 and was initially advertised by router R3 itself. This AS_Path access list is used by filter list 10, applied to filter outbound advertisements through the configured eBGP session from R3. This ensures that AS3 does not inadvertently become a transit AS and that only prefixes belonging to AS3 are advertised externally. The second AS_Path access list is AS_Path access list 4, which consists of three rules. The first rule permits all prefixes with an AS_Path attribute having the sequence AS_Path:1 6. The second rule permits all prefixes with an AS_Path attribute having the sequence AS_Path:1. In other words, the AS_Path access list allows all prefixes originally advertised in AS6 that passed through AS1 to reach router R3 in AS3, as well as prefixes originating in AS1 and reaching R3 in AS3. The third rule of this AS_Path access list permits all AS_Path attributes that include a sequence of two AS numbers where the first is AS1 and the second can be any number, expressed as "^1[0-9]*". This satisfies the routing policy requirement to accept advertised prefixes originating from any AS directly neighboring AS1. This AS_Path access list is utilized by the route map SET_OUTBOUND_TRAFFIC, applied to inbound advertisements via the eBGP session on router R3.

Additionally, prefix-related access lists are configured. Standard access list 2 permits the network 172.16.220.0 /24 and is configured only on router R3. Extended access list 101 is configured on both

routers and permits only the prefix 193.78.0.0 /16. Access list 2 is used by the route map SET_INBOUND_TRAFFIC, applied to outbound advertisements via the eBGP session from router R3, while access list 101 is used by the route map SET_OUTBOUND_TRAFFIC, applied to inbound advertisements via the eBGP sessions on both routers R3 and R4.

The configuration of the route map SET_OUTBOUND_TRAFFIC, applied to inbound advertisements, consists of two rules: permit 10 and permit 20. The first rule references access list 101, which permits the prefix 193.78.0.0 /16 (treated as the Internet prefix). If this access list is matched and the prefix 193.78.0.0 /16 is accepted, router R3 assigns the Local_Preference attribute a value of 200 via this route map, while router R4 assigns the Local_Preference attribute a value of 250. The second rule references AS_Path access list 4, indicating that if this AS_Path access list is matched, i.e., prefixes with AS_Path attributes of 1 6 or 6 are received this route map assigns those prefixes a Local_Preference attribute of 300 on router R3 and 250 on router R4.

The configuration of the route map SET_INBOUND_TRAFFIC, applied to outbound advertisements, includes two rules on router R3 and one rule on router R4. On router R3, the first rule references access list 2, permitting the prefix 172.16.220.0 /24, a local network associated with router R3. If this access list is matched and the prefix 172.16.220.0 /24 is accepted, router R3 assigns this prefix a MED attribute value of 200 before advertising it further via the X2 link. For all other prefixes not matching 172.16.220.0 /24, this route map assigns a MED attribute value of 300 before advertising those prefixes via the X2 link. On router R4, it is defined that all prefixes are assigned a MED attribute value of 250 before being advertised further via the X1 link.

5. CONCLUSION

Based on the implemented BGP configuration and the defined route maps SET_INBOUND_TRAFFIC and SET_OUTBOUND_TRAFFIC for inbound and outbound advertisements on both routers R3 and R4, the routing policy is fully satisfied. The advertised prefix 193.78.0.0 /16, received from AS1 in AS3 via the X1 and X2 links, is assigned a Local_Preference attribute value of 200 by router R3 and 250 by router R4. Consequently, it is concluded that outbound traffic to the network 193.78.0.0 /16, i.e., to the Internet, will primarily traverse the X1 link. It is observed that for the advertised prefixes that were initially advertised in AS1, in AS6, or some other AS next to the AS1 and then reached AS3 through AS1 via links X1 and X2, router R3 assigned a Local_Preference attribute of 300, while router R4 assigned a Local_Preference attribute of 250. Based on this, it is concluded that outbound traffic to networks in AS1 and AS6 will primarily flow over link X2. Regarding outbound advertisements and the SET_INBOUND_TRAFFIC route map, router R3 assigned a MED attribute of 200 to the prefix 172.16.220.0 /24. For all other prefixes, router R3 assigned a MED attribute of 300. On the other hand, router R4 assigns a MED attribute of 250 to all the prefixes it advertises. Based on this, it is concluded that incoming traffic to the network 172.16.220.0 /24 in AS3 will primarily arrive from AS1 over link X2. Incoming traffic for all other networks in AS3 will arrive from AS1 primarily via link X1. Using BGP for LAN-WAN routing can pose challenges and complexities for a multi-homed network. Configuring and maintaining BGP involves establishing equitable relationships with ISPs and other ASs, along with configuring routers with appropriate attributes, access filters, and route maps.

6. REFERENCES

- [1] Lacoste, R., Edgeworth, B. *CCNP Enterprise Advanced Routing ENARSI 300-410 Official Cert Guide*, Cisco Systems, Inc. (2020), ISBN-13: 978-1-58714-525-4, ISBN-10: 1-58714-525-1
- [2] Zhang, R., Bartell, M. *BGP Design and Implementation*, Cisco Press (2004), ISBN: 1-58705-109-5
- [3] Jain, V., Edgeworth, B. *Troubleshooting BGP: A Practical Guide to Understanding and Troubleshooting BGP (Networking Technology)*, 1st Edition, Cisco Press (2016), ISBN-13: 978-1587144646
- [4] Caesar, M. and Rexford, J. (2005). BGP routing policies in ISP networks. *IEEE Network*. 19(6), 5-11. doi: 10.1109/MNET.2005.1541715.

- [5] Flavel, A., McMahon, J., Shaikh, A., Roughan, M., & Bean, N. (2010). BGP route prediction within ISPs. *Computer Communications*, 33(10), 1180-1190.