International Journal of Core Engineering & Management
Volume-6, Issue-12, 2021 ISSN No: 2348-9510

THE ROLE OF BGP IN MULTI-HOMED NETWORKS BEST PRACTICES FOR INTERNET REDUNDANCY AND TRAFFIC OPTIMIZATION

Nikhil Bhagat, Principal Network Engineer
Independent Scholar, Network Engineering
Aurora, Colorado, USA
nikhil.bhagat90@gmail.com

Abstract

Border Gateway Protocol (BGP) is the foundation for internet routing and has evolved to be used in multi-homed networks where internet redundancy and traffic optimization is important. With businesses and organizations becoming more and more dependent on the internet to operate, high uptime levels and optimal traffic flow are becoming a must. The paper describes how BGP facilitates internet redundancy and traffic management in multi-homed networks. It presents a timeline of BGP's history as both an Interior Gateway Protocol (IGP) and an Exterior Gateway Protocol (EGP), shows how BGP made the internet into the global interconnected network it is today, and explains ways enterprises and service providers can configure multihomed and redundant networks with BGP. It also describes various route-manipulation techniques using BGP to perform traffic engineering. Using these route-manipulation techniques, enterprises and service providers can optimize the traffic flow across their network. The paper concludes with best practices for improving BGP performance over the internet, providing a stable and efficient network.

Index Terms – BGP, EGP, Multi-homed networks, Route manipulation, best practices, Internet.

I. INTRODUCTION

BGP, which is the core routing protocol of the present-day internet, provides data access over large networks with high efficiency [1]. BGP relays routing information across autonomous systems (AS) – large, distinct networks operated by one or more service providers [2]. This protocol supports redundancy and load balancing, which is an essential feature for multi-homed networks, where different internet service providers (ISPs) provide you with seamless connectivity and optimized traffic.

In the current business world where a single day of downtime is likely to cost the business a large financial loss, the necessity of robust network redundancy and traffic planning is greater than ever. Multi-homed networks accomplish this through BGP, which entails the ability to seamlessly route traffic across multiple connections in case one is interrupted [3]. BGP also enables network administrators to control the flow of traffic by selecting policies-enabled routes to ensure that data is transported in the network.

This paper dives deep into the use of BGP in multi-homed networks. It discusses BGP evolution as IGP and EGP, establishing multi-homed and redundant BGP networks, techniques for route-manipulation and best practices for BGP internet performance.

II. EVOLUTION OF BGP AS AN IGP AND EGP

BGP's existence was the fundamental design element of the internet. Originally created in the late 1980s, BGP has passed through various iterations to be the robust protocol on which the internet relies today [4]. To follow the evolution of BGP, it is vital to understand the difference between IGPs and EGPs.

A. IGP (Interior Gateway Protocol)

IGPs are routing protocols implemented within one Autonomous System (AS) to perform routing decisions. Three of the most popular IGPs are Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP) and Routing Information Protocol (RIP) [5]. These protocols are used by routers within an AS to exchange routing information and dynamically select optimal routes to route traffic from source to destination.

BGP, while commonly regarded as an EGP, is sometimes also an IGP. BGP can be implemented as an IGP (iBGP) to facilitate routing between routers in a single AS. But typically, iBGP is deployed with other IGPs to offer fast internal routing and routing loop protection.

B. EGP (Exterior Gateway Protocol)

EGPs are routing protocols that handle routing information exchange between multiple autonomous systems. BGP is the most widely used EGP on the web. Being an EGP (eBGP), BGP enables networks to advertise their routes to other autonomous systems and connect networks.

BGP-1.1, the earliest version released in 1989, was an upgrade to the less scalable Exterior Gateway Protocol (EGP) [6]. As the internet grew larger and more complex, new versions of BGP emerged including the current BGP-4, which supports Classless Inter-Domain Routing (CIDR) [7]. This feature supports efficient allocation of IP addresses and avoid exhaustion of IPv4 addresses.

C. Key Features of BGP:

- Path Vector Protocol: BGP employs path vector mechanism to define the path to different networks [4]. It advertises the complete routes (AS paths) instead of distance-based metrics, to make it easier to manage and enforce route policies.
- Scalability: BGP was created to manage the internet at scale, enabling thousands of routes and networks.
- Policy-Based Routing: In contrast to distance-vector or link-state routing protocols, BGP provides routing policies for controlling the advertised routes and preference.

III. HOW BGP TRANSFORMED THE INTERNET

The impact of BGP on the history of the internet cannot be overstated. When the internet exponentially evolved beyond research institutes and government networks to go into the enterprise, a scalable, dynamic routing protocol was necessary. With BGP, the internet was transformed into a free, decentralized worldwide network of reliable independent systems.

Below are some of the ways BGP transformed the internet:

A. Decentralization of Network Control

Routing protocols such as RIP and EGP, which had lower control on scalability, struggled with managing larger and complex networks [8]. BGP introduced the concept of autonomous systems,

in which specific area of the networks were independent managed by an entity [1]. BGP allowed intern Autonomous Systems communication for different entities to communicate with each other. This decentralization proved essential as ISPs and businesses created their own networks with their own policies, operating requirements and topologies.

B. Supporting Internet Growth

When the amount of traffic was growing exponentially during the 1990s and 2000s, BGP's large-scale routing table implementation and support for CIDR prevented routing table from growing exponentially [8]. As the number of IP prefixes grew, BGP allowed service providers and enterprises to aggregate multiple routes in one, making BGP very agile on the internet.

C. Traffic Engineering and Optimization

With BGP, network operators had policy-based routing to control the course of traffic [9]. This ability allowed for traffic engineering where administrators could route according to business needs, bandwidth costs, or performance expectations. BGP made it possible for Internet Service Providers (ISPs) and enterprises to manage their traffic and transport data across multiple ASes [4].

D. Internet Redundancy

BGP allowed for multi-homing, where the organizations could connect their BGP autonomous systems via multiple ISPs. This redundancy ensured that if one ISP's line was lost, BGP would route traffic through another ISP, maintaining continuous connectivity. The support for multiple connections and routing on AS path revolutionized the internet reliability.

IV. BUILD MULTI-HOMED NETWORKS USING BGP

A multi-homed network is a network that links to two or more ISPs to provide redundancy and better traffic distribution [3]. Multi-homing is commonly implemented using BGP because it governs and moves traffic to and from multiple connections according to policies.
Strategies to build a Multi-Homed Network Using BGP

A. Connect to Multiple ISPs

The key to a multi-homed network is to have physical connections to two or more ISPs. It can be accomplished via direct peering contracts with the service provider or buying internet transit services.

B. Obtaining an Autonomous System Number (ASN):

To implement BGP, the network administrator needs to possess an ASN (Anonymous system Number) [1]. ASN is required to setup External Border Gateway Protocol (eBGP) sessions with upstream ISPs.

C. Configure BGP Sessions

Once the physical connections are established between the ISP and the organization, and ASN is purchased, the network administrator needs to configure BGP sessions with each ISP. These eBGP sessions permit the network to share routing data with its upstream providers.

D. Advertise IP Prefixes

The network operator will need to advertise the organizations IP prefixes to the ISPs. This advertisement informs the ISPs what IP addresses are associated with the business and which routes must be connected to the multi-homed network [10].

E. Implement BGP Policies

BGP has full policy enforcement and gives the network administrator complete control over routing decisions. Policies like Local Preference, AS Path Prepending, and Multi-Exit Discriminators (MEDs) can be used to manage traffic routing from and to several ISPs.

F. Monitor and Optimize

A multi-homed network will need constant monitoring to ensure traffic is being distributed efficiently [9]. The monitoring prevents one of the ISP link from becoming overwhelmed. BGP policies can be configured and optimized as needed to maintain performance and prevent link congestion.

V. BUILD REDUNDANT NETWORKS USING BGP

Redundancy is critical for any network that requires high availability and reliability. BGP is the protocol of choice for building redundant networks because it supports multiple connections and can dynamically reroute traffic in case of a failure.
Key Redundancy Features in BGP:

A. Failover Mechanisms

BGP's most significant redundancy capability is the ability to detect link dropouts and redirect traffic [6]. When one ISP's connection fails, BGP will re-route traffic to another ISP and will keep the network up.

B. Load Balancing

BGP allows load balancing in terms of Equal-Cost Multi-Path (ECMP) in which traffic is split up on different links at equal cost. This allows an organization to prevent throttling a single connection, thus ensuring optimum bandwidth utilization.

C. Route Aggregation

BGP is capable of aggregating routes to make routing tables smaller and the network more manageable [11]. This is particularly crucial for redundant networks where multiple ISPs can advertise the same route.

D. Policy-Based Redundancy

Administrators can enable AS Path Prepending to have some routes prioritized over others [9]. For instance, an organization might want to have traffic sent via a higher bandwidth link until the link goes down, in which case BGP would forward traffic via a redundant link.

International Journal of Core Engineering & Management
Volume-6, Issue-12, 2021 ISSN No: 2348-9510

E. Monitoring and Failover Testing

Redundant BGP environments need continuous monitoring and regular failover testing to make sure the failovers work as intended. The default failover capabilities in BGP are powerful, but it needs to be tested frequently to keep the network secure.

VI. ROUTE MANIPULATION STRATEGIES USING BGP

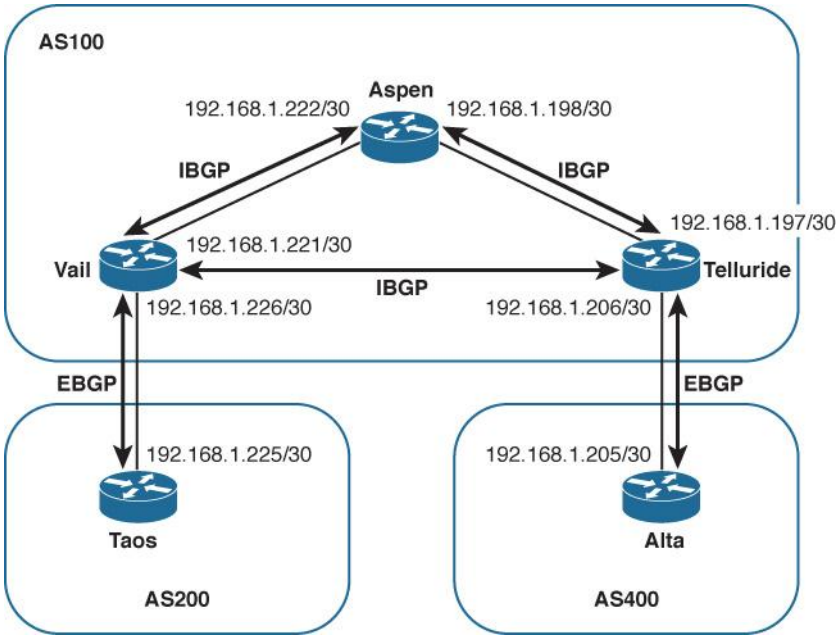

Fig 1. Network topology using BGP [12]

A. Outbound Route Manipulation

Outbound route manipulation involves manipulating the routing of traffic when it leaves your network (how you influence and transfer traffic towards other networks) [13]. The following techniques and features are widely deployed:

1. **Local Preference (LOCAL_PREF)**

- How it Works: The higher LOCAL_PREF value is always preferable [14].
- Usage: Most often implemented in iBGP networks. For instance, if a business wants to prefer one ISP over another for outbound traffic, they can assign a higher LOCAL_PREF to the preferred ISP routes.

```
route-map CONFIGURE_LOCAL_PREF permit 10
 set local-preference 150
```

- Scenario: If there are multiple routes to the same destination, a business could increment the LOCAL_PREF of the preferred path, causing all routers in the AS to route outbound traffic through it.

International Journal of Core Engineering & Management
Volume-6, Issue-12, 2021 ISSN No: 2348-9510

2. **Weight (Cisco-specific)**

- How it Works: The weight attribute has its significance locally within a router, and the most weighted path is selected by BGP [1].
- Usage: Used when a business wants to control the route-selection of a single router without effecting other routers on the network.

```
route-map SET_WEIGHT permit 10
 set weight 2000
```

- Scenario: A business can set the weight of the learned path from a specific neighbor on the router, and the router will always choose that path when sending outbound traffic.

B. Inbound Route Manipulation

Inbound route manipulation is a technique used by businesses to signal external networks their preferred path of entry within the business's network [4]. Since the business is not in control of the external ASes, the businesses would typically use BGP attributes such as AS_PATH and MED to signal their desired paths out to the external networks.

1. **AS_PATH Prepending**

- How it Works: Businesses can add their own AS number to a route multiple times to appear longer and more unfavorable [15].
- Usage: Used to influence inbound traffic by making certain routes less desirable for outside ASes.

```
route-map CONFIGURE_PREPEND_AS_PATH permit 10
 set as-path prepend 65001 65001 65001
```

- Scenario: If the business is using multiple ISPs, they can append their AS number multiple times on one route to make it unattractive, thus attracting traffic from the other ISP.

2. **Multi-Exit Discriminator (MED)**

- How it Works: The lower the MED value, the more preferred route. However, it applies only to routes received from the same neighboring AS [4].
- Usage: It can be used to affect inbound traffic by signaling your AS's preferred entry point when routing to other ASes.

```
route-map CONFIGURE_MED permit 10
 set metric 50
```

- Scenario: If the business's AS has multiple connections to a neighboring AS, then they may set a lowest MED value on the preferred connection and tell the neighboring AS to send inbound traffic through it.

3. **Communities**

- How it Works Communities are route tags. They can be applied to routing policies, such as "no-export" (don't allow a route to be advertised outside the AS), local-preference or other

special behaviors [16].

- Usage: A business can use BGP communities to communicate routing preferences to upstream or downstream ASes, or group of routes for special service (traffic engineering).

```
route-map CONFIGURE_COMMUNITY permit 10
 set community 65002:100
```

- Scenario: You could associate certain routes with a community value to enable your upstream provider to apply a routing policy (e.g., low local preference or no export of the route to other peers).

C. Route Filtering and Policy Control

BGP supports filtering of routes as well, which means an enterprise or service provider can control which routes are advertised or allowed [17]. They can be used to filter on various attributes.

1. Prefix Lists

- How it Works: Prefix lists filters routes based on their destination prefix and subnet mask [18].
- Usage: A business can use prefix lists to control a group of routes being advertised to and received from a BGP peer.

```
ip prefix-list CONFIGURE_FILTER_ROUTES permit 192.168.1.0/24
ip prefix-list CONFIGURE_FILTER_ROUTES deny any
```

- Scenario: A business can create a prefix list to allow in or advertised out a certain number of routes towards a specific BGP peer.

2. Route Maps

- How it Works: Match and set statements, filter routes or adjust attributes of route maps [1]. A business can use them either inbound (routes learned) or outbound (routes advertised).
- Usage: A business can use route maps to manipulate or filter attributes like LOCAL_PREF, AS_PATH, MED, etc. on a specified match criteria (prefix, AS_PATH, communities).

```
route-map CONFIGURE_FILTER_ROUTES deny 10
 match ip address prefix-list CONFIGURE_FILTER_ROUTES
```

- Scenario: A business can use a route map to block specific prefixes from being advertised to or received from a BGP peer and have a full policy control.

3. BGP Peer Groups

- How it Works: Peer group allows several BGP peers to share the same routing policies, simplifying the configuration and enabling behavior consistency for multiple peers [7].
- Usage: Peer groups can be used to apply route maps, prefix list, or other policy to multiple BGP peers at the same time.

International Journal of Core Engineering & Management
Volume-6, Issue-12, 2021 ISSN No: 2348-9510

```
neighbor CONFIGURE_PEER_GROUP_NAME peer-group
neighbor CONFIGURE_PEER_GROUP_NAME route-map SET_LOCAL_PREF out
```

- Scenario: If a business has more than one peer that needs the policy (e.g., creating a LOCAL_PREF), they can create a group of peers and apply the policy to all peers.

D. Route Aggregation and Summarization

BGP allows for the aggregation of multiple routes into a summary route, saving routers from having to advertise multiple routes to peers while scaling up large networks [19].

1. **Aggregate Routes**

- How it Works: When aggregating routes, BGP can display one larger aggregated route as opposed to multiple subnetted routes, minimizing routing table size.
- Usage: Using route aggregation, businesses can reduce the number of routes they share with other ASes, improving scalability.

```
aggregate-address 192.168.0.0 255.255.0.0 summary-only
```

- Scenario: If a business has more than one subnet in the 192.168.0.0/16 network, they can create a summary route (192.168.0.0/16) instead of dozens of unique prefixes (192.168.1.0./24, 192.168.2.0/24).

VII. BEST PRACTICES FOR USING BGP OVER THE INTERNET

While BGP can offer great solutions for internet redundancy and traffic management, a business should be cautious to avoid routing instabilities, misconfiguration, or security risks. A business can achieve better performance with BGP networks if they adhere to these best practices.

A. Use Route Filtering

BGP lacks internal defences against route hijacking or advertisement of incorrect routes [20]. This makes route filtering mandatory, where the network administrators need to decide on prefixes they want to be advertised and accepted from peers. This can help establish protection against malicious or false routes.

B. Implement Secure BGP Sessions

In order to protect BGP sessions, administrators should make use of MD5 authentication for BGP peering. It means that only authorized routers can open BGP sessions and share routing information. Securing BGP sessions helps in avoiding the unauthorized entry and even attack on routing system [21].

C. Monitor BGP Sessions

Monitoring of BGP sessions is required to make sure that the network stays operational [1]. Monitoring tools can show routing failures, a drop in connection, or excessive flapping of the route. Network operators can act quickly if the issues are detected.

D. Use BGP Communities for Traffic Engineering

BGP communities enable traffic engineering because the administrators can assign routes to attributes. Communities enable network operators to influence traffic flow decisions (e.g. favoring certain flows of traffic versus other ISPs for specific traffic).

E. Minimize AS Path Length

Most popular BGP optimization strategy is to keep the AS path length small. BGP likes to use shorter AS paths so that traffic uses the fastest and shortest route to optimize performance. Administrators should constantly monitor and optimize AS routes to reduce latency and avoid redundant hops [22].

F. Ensure Proper Prefix Aggregation

Optimal prefix aggregation reduces number of routes offered to BGP, which makes routing tables smaller. Eliminating redundant route advertisements speeds up the network and reduces the probability of unnecessary routing table expansion.

G. Conduct Regular Failover Testing

Multi-homed and redundant networks should undergo routine failover testing to ensure that BGP enables rerouting on demand [23]. By testing for failures and watching BGP response, it's always optimal to verify that failovers are in place and that the configuration does not contain any issues.

VIII. CONCLUSION

- BGP has been central to the early evolution and expansion of the modern internet, and revolutionized the way networks communicate, share routes, and make redundancy possible.
- In multi-homed networks, BGP allows enterprises and service providers to have uninterrupted connectivity, high-speed traffic, and faster convergence.
- Organizations can create multi-homed, redundant networks through BGP to ensure high availability, scalability, and performance.
- However, BGP should be monitored for misconfigurations and vulnerabilities to maintain route stability.
- Route filtering, session security, traffic engineering within communities and routine failover tests are recommended for an effective and reliable BGP deployment.
- If the implementation best practices are followed, BGP is a highly effective protocol that helps keep the internet secure and maximizes traffic efficiency in multihomed networks.

REFERENCES

1. "RFC 1654: A Border Gateway Protocol 4 (BGP-4)," RFC Editor, 1994.
2. Bergman, R. Hendricks, and G. Belson, "Pied Piper: Rethinking Internet Data Delivery," in Proc. 2020 ACM SIGCOMM Conf., 2020.
3. N. Kushman, S. Kandula, D. Katabi, and B. M. Maggs, "R-BGP: Staying connected in a connected world," in Proc. 2007 ACM SIGCOMM Conf., 2007, pp. 231-242.
4. Y. Rekhter and P. Gross, "Application of the Border Gateway Protocol in the Internet," RFC

Editor, 1994.
5. C. Wijaya, "Performance analysis of dynamic routing protocol EIGRP and OSPF in IPv4 and IPv6 network," Int. J. Adv. Comput. Sci. Appl., vol. 7, no. 12, pp. 135-142, 2016.
6. "RFC 1267: A Border Gateway Protocol 3 (BGP-3)," RFC Editor, 1991.
7. Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC Editor, 2006.
8. A. Abhashkumar et al., "Running BGP in data centers at scale," in Proc. 2015 ACM SIGCOMM Conf., 2015, pp. 435-448.
9. Y. Wang, M. Schapira, and J. Rexford, "Neighbor-specific BGP: More flexible routing policies while improving global stability," in Proc. 2014 ACM SIGCOMM Conf., 2014, pp. 15-28.
10. "Enterprise multihoming using provider-assigned addresses without network prefix translation: Requirements and solution," RFC Editor, 2012.
11. "Link aggregation and load balancing," IEEE, 2020.
12. C. Solis, "The OSI Model: Understanding the Seven Layers of Computer Networks," Cisco Press, Aug. 11, 2023. [Online]. Available: https://www.ciscopress.com/articles/article.asp?p=2738462&seqNum=3.
13. "BGP persistence," Cisco, 2018.
14. "Border Gateway Protocol," Cisco, 2017.
15. J. Tantsura and P. Lapukhov, "Equal-cost multipath considerations for BGP," RFC Editor, 2020.
16. H. Liang, G. Teng, J. Wang, D. Wang, and Y. Gao, "BGP-based inter-domain traffic engineering in BGP/MPLS VPNs," Int. J. Netw. Manag., vol. 16, no. 3, pp. 160-174, 2010.
17. "RFC 4271: A Border Gateway Protocol 4 (BGP-4)," RFC Editor, 2006.
18. S. Hares and K. Patel, "AS path based outbound route filter for BGP-4," RFC Editor, 2020.
19. "BGP fundamentals," Cisco, 2019.
20. G. Carl and G. Kesidis, "Large-scale testing of the Internet's Border Gateway Protocol (BGP) via topological scale-down," IEEE J. Sel. Areas Commun., vol. 21, no. 6, pp. 987-998, 2003.
21. S. Bellovin, "Security problems in the TCP/IP protocol suite," ACM SIGCOMM Comput. Commun. Rev., vol. 19, no. 2, pp. 32-48, 1989.
22. "IP routing - Q&A," Cisco, 2018.
23. "Rapid detection of BGP anomalies," Cisco, 2016.