10th International Conference on Control Engineering &Information Technology (CEIT-2025)
Proceedings Book Series –PBS- Vol 23, pp.71-79

Cost and Availability Improvements for Fault-Tolerant Systems Through Bidirectional Forwarding Detection (BFD) over the Gateway Load Balancing Protocol (GLBP)

Mahmud Mansour[#1], Mohammad Mehdi Berrish[#2], Ahmed Ben Hassan[#3]

[#]*Dapartment of Network, University of Tripoli*
Tripoli, Libya

[1]Mah.mansour@uot.edu.ly
[2]M.berrish@uot.edu.ly
[3]A.benhassan@uot.edu.ly

Abstract— **The Internet has experienced explosive growth since its emergence. As a result, enterprises are increasingly prioritizing network availability and minimizing downtime due to the growing demand for online applications and services. Maintaining high availability can be costly, but a lack can damage an organization's reputation and cause significant financial losses. To enhance availability, the Gateway Load Balancing Protocol (GLBP) is one of the necessary protocols that is used to achieve this goal. GLBP is a redundancy protocol that is used to manage network default gateway routers by using one or more redundant routers that will take over in case of default router failure. However, late failure detections and slow responses can lead to packet loss during failure. Bidirectional Forwarding Detection (BFD) is an effective solution to increase availability by rapidly detecting link failure and monitoring IP connectivity. This paper discusses how the integration of BFD with GLBP will enhance availability and reduce downtime. The evaluation focuses on convergence time, packet loss, CPU utilization, and bandwidth consumption. Following the implementation, testing, optimization, and simulation using PNETLAB. We have verified that GLBP with BFD shows very fast failure detection and recovery with reduced downtime and packet loss, thus improving network reliability and stability.**

Keywords— ***FHRP, BFD, GLBP, Cost, Availability.***

I. INTRODUCTION

The exponential growth of the internet and the ubiquitous integration of computing systems into various facets of our lives have propelled us into an era where dependence on reliable networks is more critical than ever. In the early 1990s, the concept of electronic mail connectivity marked a novel leap in communication. Fast forward a decade, and we witnessed a seismic shift where reputable companies transformed their websites into platforms for direct consumer transactions. Then quality of service has emerged as a critical requirement in today's IP networks that carry a multitude of real-time multimedia applications such as voice over IP, video conferencing, live streaming, and online gaming [1]. This transformative journey reflects the staggering progress of the internet and networking technologies, becoming integral to our daily lives and professional landscapes [2].

Modern society requires certain systems, such as air traffic control or life support systems, to be continuously available. Because of this, availability has become a significant concern for enterprises and businesses in today's network. Every minute of service interruption has the potential to result in significant financial losses for a firm, amounting to hundreds or even thousands of dollars. To avoid outages, we aim to enhance the network's uptime by implementing redundant lines and nodes and fault isolation, fault detection and notification, and online repair. Nevertheless, such systems may fail if several redundant units fail simultaneously or if single points of failure exist. While redundancy might be beneficial, it also comes with a high cost. Achieving optimal network availability is dependent on the client's specific business objectives and their tolerance for network downtime [3].

ISSN: 2961-6611

Availability refers to the length of time a network is available for users and is generally a crucial goal for network design clients. Availability can be defined as a percent uptime per year, month, week, day, or hour, relative to the entire time in that period. For example, in a network that delivers 24-hour, 7-day-a-week service, if the network is up 165 hours in the 168-hour week, availability is 98.21 percent [4].

Normally, availability is represented as the percentage of time the network is functional. It was here that the phrase “five nine” came into usage. Five-nines refer to the percentage of 99.999%, which is a generality that has for long been used for marketing and has been seen as the desirable target for availability in many networks, at least at the core level [5].

As we delve deeper into the realm of network reliability, it becomes increasingly crucial to shed light on a critical yet often underestimated aspect – the financial repercussions of network downtime. Many firms may not fully understand the impact of downtime on their business. Calculating the impact's cost may be tough, as it demands a thorough understanding of both physical and intangible losses. Actual losses are physical expenditures; they include lost income, the cost to retrieve lost information, catastrophe recovery, and business continuity costs. Intangible costs include damage to your company's reputation, lost customers, and staff productivity losses. In many circumstances, the damage associated with intangible costs may have a greater long-term effect on an organization than that of actual expenses. Downtime cost is defined as any profit that a corporation loses when its equipment or network stops working.

ITIC's 2022 Global Server Hardware and Server OS Reliability Survey found that 91% of respondents now estimate that one hour of downtime costs enterprises $301,000 or more; this is an increase of two (2) percentage points in less than two years. Of that number, 44% of those polled indicated that hourly downtime costs now exceed $1 million. Since 2021, only one (1%) percent of respondents said a single hour of downtime costs them $100,000 or less. Nine percent (9%) of respondents valued hourly downtime at $101,000 to $300,000 [6].

II. Related Work

In a study by Mehdi Berrish et al. (2024) [7], “Performance Analysis of Bidirectional Forwarding Detection (BFD) over the Hot Standby Router Protocol (HSRP)," they compared the performance of HSRP with and without BFD in terms of packet loss, convergence time, CPU utilization, and bandwidth consumption.

In Ben Hassan (2024) [8], “Performance Evaluation of Bidirectional Forwarding Detection (BFD) over Virtual Router Redundancy Protocol (VRRP)," they compared the performance of VRRP with and without BFD in terms of convergence time, CPU utilization, bandwidth consumption and packet loss.

In Niu and Li (2023) [9], “Design and Implementation of VRRP and BFD Linkage Technology in Campus Information Service Platform Network,” the study examined how to integrate BFD and VRRP technologies into existing networks. However, how BFD affects convergence time, packet loss, bandwidth usage, and CPU usage has not been evaluated.

In Ben Saud (2023) [10], “Performance Evaluation of First Hop Redundancy Protocols in IPv4 and IPv6 Networks," they compared the FHRPv4 and FHRPv6 performance in terms of packet loss, convergence time, and CPU utilization without evaluating bandwidth consumption and used IP SLA as a method for detecting ISP failures, which has a failure detection time of at least 1 second.

In Mansour (2022) [11], “Performance analysis and functionality comparison of first hop redundancy protocol IPV6” focused on the FHRPv6 performance in terms of packet loss and convergence time and used IP SLA as a method for detecting ISP failures.

In study Kim et al. (2019) [12], “FDVRRP: Router implementation for fast detection and high availability in network failure cases," they studied the implementation of fast detection BFD with VRRP to improve failure detection and a failover. But it was limited to an on-premise failure scenario involving the failure of a master router.

III. Bidirectional Forwarding Detection Protocol

Bidirectional Forwarding Detection (BFD) is a fast millisecond failure detection mechanism that rapidly detects link failure and monitors IP connectivity on the entire network independent of media and routing protocols while maintaining low overhead. It also provides a single, standardized method of link/device/protocol failure detection at any protocol layer and over any media.

The BFD protocol is designed to provide low overhead and fast detection of link failures on any type of path, including direct physical links, virtual circuits, tunnels, MPLS, label switched paths (LSPs), and

ISSN: 2961-6611

multihop routed paths. Furthermore, it operates independently on the transmission media, data protocol, and routing protocol, without any need to modify the existing protocols [12].

A. BFD Echo mode

BFD Echo is a rapid failure detection mechanism in which the local system sends BFD Echo packets and the remote system loops back the packets. BFD echo mode is enabled by default, but you can disable it so that it can run independently in each direction. BFD echo mode works with asynchronous BFD. Echo packets are sent by the forwarding engine and forwarded back along the same path in order to perform detection. The BFD session at the other end does not participate in the actual forwarding of the echo packets. The echo function and the forwarding engine are responsible for the detection process, while BFD control packets maintain the BFD session as shown in Fig 1; therefore, the number of BFD control packets that are sent out between two BFD neighbors is reduced. In addition, because the forwarding engine is testing the forwarding path on the remote (neighbor) system without involving the remote system, there is an opportunity to improve the interpacket delay variance, thereby achieving quicker failure detection times [13].

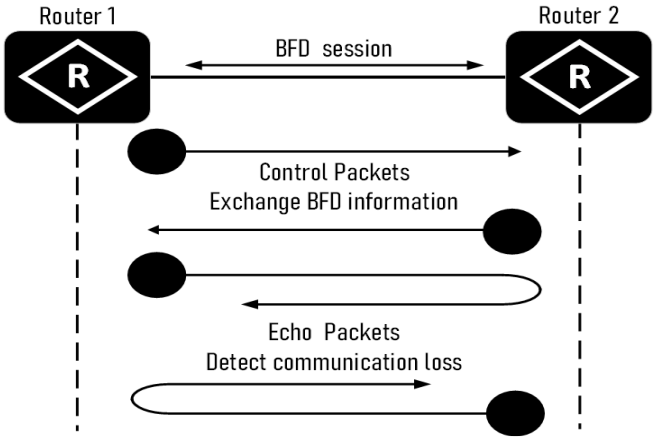

Fig. 1. Asynchronous mode with Echo mode

IV. GATEWAY LOAD BALANCING PROTOCOL

Gateway Load Balancing Protocol (GLBP) is one of the First Hop Redundancy Protocols (FHRP), which provides redundancy, like other First Hop Redundancy Protocols, as well as load balancing. It is a Cisco proprietary protocol that can perform both functions. It provides load balancing over multiple routers using a single virtual IP address and multiple virtual Mac addresses. Each host is configured with the same virtual IP address, and all routers in the virtual router group participate in forwarding packets.

GLBP works by making use of a single virtual IP address, which is configured as the default gateway on the hosts. When the routers are set to a GLBP group, they first choose one gateway to be the Active Virtual Gateway (AVG) for that group. The election is based on the priority of each gateway (the gateway with the highest priority wins). If all of them have the same priority, then the gateway with the highest actual IP address becomes the AVG [14].

The group members provide backup for the AVG in the event that the AVG becomes unavailable. The AVG assigns a virtual MAC address for each member of the GLBP group. Each gateway assumes responsibility for forwarding packets sent to the virtual MAC address assigned to it by the AVG . These gateways are active virtual forwarders (AVFs) for their virtual MAC addresses. The AVG is responsible for answering Address Resolution Protocol (ARP) requests for the virtual IP address. Load sharing is achieved by the AVG replying to the ARP requests with different virtual MAC addresses [14].

Each gateway that is issued a virtual MAC address is termed an Active Virtual Forwarder (AVF). A GLBP group only has a maximum of four AVFs. If there are more than 4 gateways in a GLBP group, then the remainder will become Standby Virtual Forwarders (SVF), which will take the place of an AVF in the case of failure [14].

Priority determines if a GLBP device functions as a standby virtual gateway and the order of ascendancy to becoming an AVG if the current AVG fails. By default, the GLBP virtual gateway preemptive scheme is disabled. A backup virtual gateway can become the AVG only if the current AVG fails, regardless of the priorities assigned to the virtual gateways.

ISSN: 2961-6611

GLBP also uses a weighting scheme to determine the forwarding capacity of each device in the GLBP group. Thresholds can be set to disable forwarding when the weighting for a GLBP group falls below a certain value. By default, the GLBP virtual forwarder preemptive scheme is enabled with a delay of 30 seconds. A standby virtual forwarder can become the AVF if the current AVF weighting falls below the low weighting threshold for 30 seconds.

The Gateway Load Balancing Protocol (GLBP) has four timers:

- Hello time is the estimated time that routers transmit in a hello message to communicate that the peer router is active, with a default value of 3 seconds.
- Hold time is the projected duration during which the standby router will report that the peer is down and becomes active, with a default value of 10 seconds.

V. Design And Configuration

This paper focuses on implementing the gateway load balancing protocol with bidirectional forwarding detection and evaluating the performance in comparison to GLBP without BFD.

Fig.2 shows the network is designed hierarchically to have two default gateway routers, each connected to a different ISP. In this work, we implemented network scenarios using PNELAB network emulator software. In the network design topology, routers on the left R1 have been configured with higher priority than the routers on the right R2.

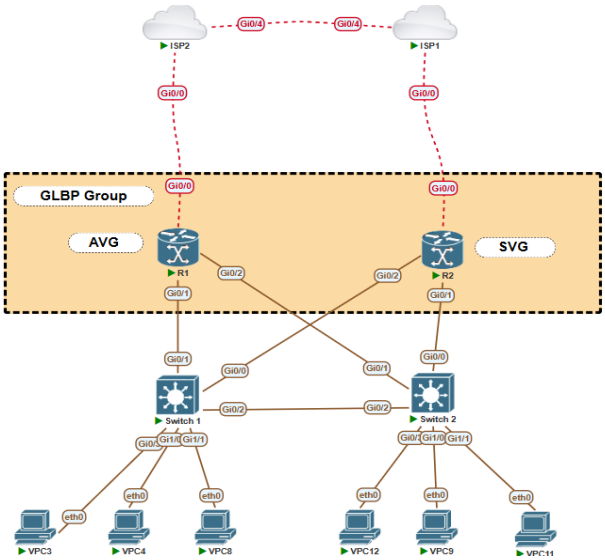

Fig. 2. Network Topology

Initially, the GLBP will be implemented without BFD. In this configuration, an IP Service Level Agreement (IP SLA) will be utilized instead of a BFD to monitor the reachability of ISPs. The IP SLA will be configured on the enterprise routers to check the reachability of ISPs. If an ISP becomes unreachable, the IP SLA will detect this loss of connectivity and report it to the GLBP installed on the router. The route object associated with the IP SLA will detect that the ISP is down and reduce the weight value of the router, allowing an router with a higher weight to take over routing the traffic instead of the router less weight. We set the detection time to 1 second, the fastest available for IP SLA, and configured the AVG router with 200 weight while SVG with 100 weight.

When implementing GLBP with BFD, BFD is operated to monitor the connectivity to ISPs, detecting access issues within milliseconds. We set the detection time to 50 milliseconds, the fastest available for BFD. The configured BFD on AVG router.

ISSN: 2961-6611

TABLE I. SIMULATION PARAMETER

Simulation	
parameter	**Value**
Number of ISP	2 ISP (ISP1- ISP2), Cisco Version 15.7(3)M3
Number of Router	2 Routers, Cisco Version 15.7(3)M3
PCs Numbers	6 PCs
Test Duration	1 Hour
GLBP- Hello –Hold time	Default : (3 – 10) s With Optimization: (1 - 3)s
Forwarder preemptive delay	Default 30 second With Optimization 0 second
BFD Detection Time	50 milliseconds
IP SLA Detection Time	1 second

VI. RESULTS

This section will present and discuss the measurements conducted in order to test the performance of GLBP, both with and without BFD, and then present and evaluate the results of GLBP without BFD compared to GLBP with BFD. The testing process comprised transmitting 3600 ICMP packets over a range of 1 hour; this duration was enough to capture critical metrics like CPU utilization and bandwidth consumption. Following this, an intentional ISP failure was generated to observe and study the network's response. To ensure the reliability and consistency of the findings, all tests were repeated multiple times. This repetition allowed us to verify the stability of the results.

A. GLBP without BFD Results

1) Convergence Time

- By using default timers for hello and hold, the convergence time between R1 and R2 is equal to 36.07 seconds, and the convergence process between ISP1 and ISP2 takes 37.02 seconds. During the convergence process, 18 ICMP packets were lost. This is caused by the forwarder preemption delay.
- By using optimized timers for hello, hold, and forwarder preemption delay, the convergence time between R1 and R2 is equal to 2.86 seconds, and the convergence process between ISP1 and ISP2 takes 3 seconds. During the convergence process, 1 ICMP packet was lost.

2) CPU Utilization

- Without timers' optimization, GLBP consumed 0.05% CPU usage on R1 and 0.04% on R2, while both routers had CPU utilization at 3%.
- With timers' optimization, GLBP consumed 0.12% CPU usage on R1 and 0.09% on R2, while both routers had CPU utilization at 3%.

3) Bandwidth Consumption

- During the testing period without timers' optimization, the traffic generated by GLBP packets accounted for approximately 270 KB. This estimate is based on the default configuration, where hello packets are sent every 3 seconds.
- During the testing period with timers' optimization, the traffic generated by GLBP packets accounted for approximately 800 KB. This estimate is based on the optimized hello packet interval, where packets are sent every 1 second. GLBP hello packet size is 102 bytes.

B. GLBP with BFD Results

1) Convergence Time

- By using default timers for hello and hold, the convergence process between ISP1 and ISP2 takes is equal to 1 second. The convergence time between R1 and R2 is equal to 35.2 seconds. During the fast convergence process, no ICMP packets were lost.

ISSN: 2961-6611

- By using optimized timers for hello and hold, the convergence process between ISP1 and ISP2 takes is equal to 1 second. The convergence time between R1 and R2 is equal to 0.59 seconds. During the fast convergence process, no ICMP packets were lost.

2) CPU Utilization

- During the testing period, BFD consumed 2.55% CPU usage on R1 and 2.62% on R2, while both routers had CPU utilization at 7%.

3) Bandwidth Consumption

- During the testing period, traffic generated by BFD packets accounted for approximately 16.54MB. This estimate is based on a failure detection duration of 50 milliseconds, during which BFD echo packets are sent every 50 milliseconds and BFD control packets are sent every second. As mentioned previously, BFD Echo packets are responsible for detecting failures, while BFD Control packets maintain the BFD session between R1 and ISP1. BFD echo packet size is 54 bytes, while BFD control packet size is 66 bytes.

Note: The network topology used in this study is relatively small; some results, such as CPU utilization and convergence time, may differ as the network scales or becomes more complex.

VII. COMPARISON AND EVALUATION

This section compares and evaluates the performance of GLBP before and after BFD implementation. The comparison parameters are convergence time, packet loss, CPU utilization, and bandwidth consumption.

A. Convergence Time

We can see from Fig. 3 that GLBP with BFD has the best convergence time result at 1 second in both default and optimized mode, thanks to a BFD failure detection time of 50 milliseconds, compared to GLBP without BFD, which has an IP SLA failure detection time of 1 second. Meanwhile, GLBP with BFD-Optimize has the best convergence time to switch between active and standby mode at 0.59 seconds; this is because of the optimized hello packet sent every 1 second combined with the BFD failure detection time.

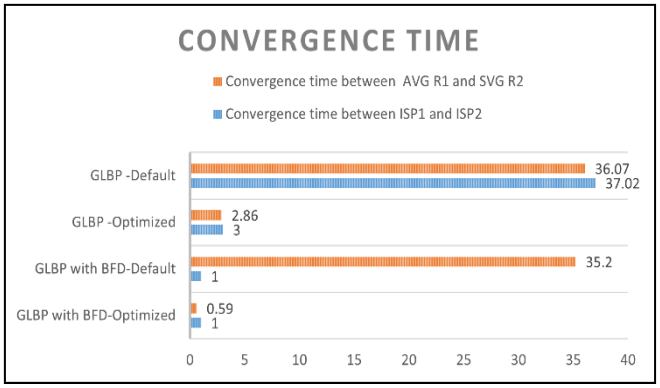

Fig. 3. Convergence Time Comparison

B. Packet Loss Comparison

Fig. 4 shows that during convergence, for GLBP without BFD, 18 packets were lost before optimization "default" due to an IP SLA failure detection time of 1 second with default hello packets sent every 3 seconds. With optimization, only 1 packet was lost thanks to the optimized hello packet sent every 1 second and reduced forwarder preemptive delay to 0 seconds. while for GLBP with BFD, no packets were lost either before optimization or with optimization, thanks to a BFD failure detection time of 50 milliseconds.

ISSN: 2961-6611

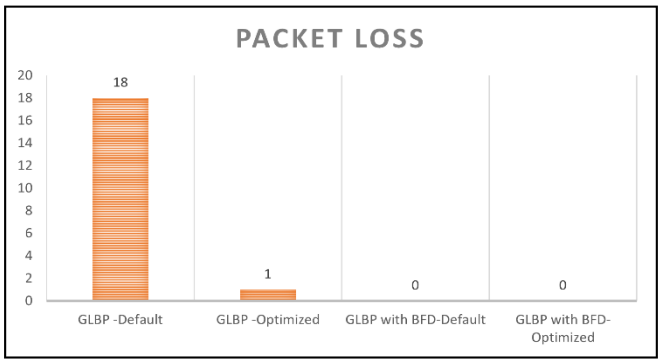

Fig. 4. Packet Loss Comparison

C. CPU Utilization

Fig. 5 shows the increase in CPU usage observed when using GLBP with BFD due to the high load resulting from sending BFD echo packets every 50 milliseconds and BFD control packets every second, so it can be concluded that GLBP with BFD has the worst CPU usage compared to GLBP without BFD.

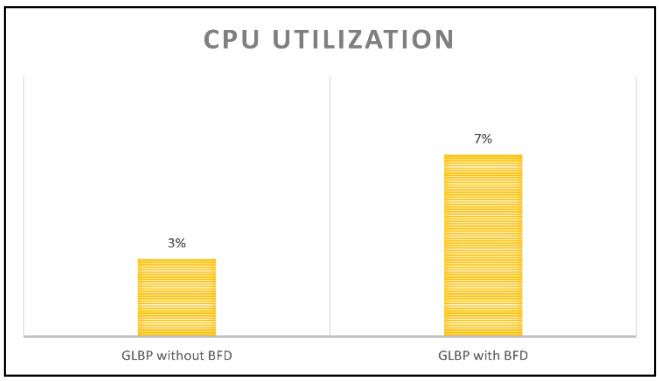

Fig. 5. CPU Utilization Comparison

D. Bandwidth Consumption

Table II shows that BFD consumes very high bandwidth, about 16.54 MB compared to IP SLA, which was 562 KB, due to the result of sending BFD echo packets every 50 milliseconds and BFD control packets every second, which we mentioned earlier. Also, we notice a significant increase in bandwidth consumption in GLBP-Optimized compared to GLBP-Default due to sending Hello packets every 1 second.

TABLE II. BANDWIDTH CONSUMPTION COMPARISON

Bandwidth Consumption	
Protocols	**Value**
GLBP-Default	270KB
GLBP-Optimized	800KB
IP SLA	562KB
BFD	16.54MB

VIII. CONCLUSION

Following the implementation and testing of GLBP both with and without BFD, and after examining its outputs concerning four critical factors, convergence time, packet loss, CPU use, and bandwidth consumption, it is evident that the integration of BFD with GLBP markedly improves network performance. The convergence time decreased from 37 seconds to 1 second, with no packet loss observed during failure circumstances. Nonetheless, this resulted in elevated CPU utilization and significant bandwidth consumption. Consequently, organizations must judiciously weigh the advantages of less packet loss and expedited failover times against the resource requirements of BFD. The implementation of BFD is contingent upon the

ISSN: 2961-6611

significance of business downtime, assuming adequate resources are accessible to fulfill CPU and bandwidth demands.

REFERENCES

[1] M. Mansour, A. Samood, and N. B. Saud, “Assessing queue management strategies to enhance quality of service in MPLS VPN networks," *Libyan Journal of Informatics*, vol 1, pp. 29-48.
[2] Chris Oggerino, “High Availability Network Fundamentals”, CiscoPress, 1st edition, 2001.
[3] L. Felsberger, B. Todd, and D. Kranzlmüller, "Cost and availability improvements for fault-tolerant systems through optimal load-sharing policies," *Procedia Computer Science*, vol. 151, pp. 592–599, 2019.
[4] Priscilla Oppenheimer, “Top-Down Network Design”, Cisco Press, 3ed edition, 2010.
[5] Mattias Thulin, “Measuring Availability in Telecommunications Networks”, “Master’s thesis report at Song Networks AB”, 2004, pp 13.
[6] L. DiDio, “Server and application reliability by the numbers: Understanding ‘The Nines,’” Nov. 30, 2022. https://itic-corp.com/server-and-application-by-the-numbers-understanding-the-nines/.
[7] M. M. Berrish, M. Mansour, and A. B. Hassan, "Performance Analysis of Bidirectional Forwarding Detection (BFD) over the Hot Standby Router Protocol (HSRP)," *International Journal of Computer Science & Security (IJCSS)*, vol. 18, no. 4, pp. 48 - 64, 2024.
[8] A.B.Hassan and M.Mansour, "Performance Evaluation of Bidirectional Forwarding Detection (BFD) over the Virtual Router Redundancy Protocol (VRRP)," *Procedia Computer Science*, vol. 251, pp. 256–264, 2024.
[9] Y. Niu and X. Li, “Design and Implementation of VRRP and BFD Linkage Technology in Campus Information Service Platform Network,” *In Proceedings of the 2023 4th International Conference on Machine Learning and Computer Application (ICMLCA ’23)*, pp. 197–201, Oct. 2023, doi: 10.1145/3650215.3650250.
[10] N. B. Saud and M. Mansour, "Performance Evaluation of First Hop Redundancy Protocols in IPv4 and IPv6 Networks," 2023 IEEE 3rd International Maghreb Meeting of the Conference on Sciences and Techniques of Automatic Control and Computer Engineering (MI-STA), Benghazi, Libya, 2023, pp. 440-445, doi: 10.1109/MI-STA57575.2023.10169462.
[11] M. Mansour, M. Agomati, M. Alsaid, M. Berrish, and R. Alasem, “Performance analysis and functionality comparison of First Hop Redundancy Protocol IPV6,” *Procedia Computer Science*, vol. 210, pp. 19–27, Jan. 2022, doi: 10.1016/j.procs.2022.10.115.
[12] C. Lee, S. Kim, and H. Ryu, “FDVRRP: Router implementation for fast detection and high availability in network failure cases,” *ETRI Journal*, vol. 41, no. 4, pp. 473–482, May 2019, doi: 10.4218/etrij.2018-0309.
[13] “Routing Configuration Guide, Cisco IOS XE Everest 16.6.X (Catalyst 9500 Switches) - Configuring Bidirectional Forwarding Detection [Support],” *Cisco*, Mar. 28, 2024. https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/16-6/configuration_guide/b_166_rtng_9500_cg/b_166_rtng_9500_cg_chapter_00.html
[14] Cisco Systems, "GLBP - Gateway Load Balancing Protocol," [Online]. Available: https://community.cisco.com/kxiwq67737/attachments/kxiwq67737/5991-discussions-wan-routing-switching/29508/1/17533-GLBP.pdf

ISSN: 2961-6611