Performance Analysis of Bidirectional Forwarding Detection (BFD) over the Hot Standby Router Protocol (HSRP)

Mohammad Mehdi Berrish *m.berrish@uot.edu.ly*
University of Tripoli
Tripoli, Libya

Mahmud Mansour *mah.mansour@uot.edu.ly*
Department of Network
University of Tripoli
Tripoli, Libya

Ahmed Ben Hassan *a.benhassan@uot.edu.ly*
Department of Network
University of Tripoli
Tripoli, Libya

Abstract

Organizations 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 IP network availability, the Hot Standby Router Protocol (HSRP) is a necessary protocol that is used to achieve this goal. HSRP 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 HSRP will enhance IP network 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 HSRP with BFD shows very fast failure detection and recovery with reduced downtime and packet loss, thus improving network reliability and stability.

Keywords: FHRP, HSRP, BFD.

1. INTRODUCTION

The Internet has experienced explosive growth since its emergence. Internet traffic keeps growing at an exponential rate of almost doubling itself each year, and this trend is expected to continue. Various and vast amounts of Internet-based applications and services emerge with this growth, and surveys reveal trends towards them. More and more people come to depend on them, and all kinds of business processes are built around them(Mansour &Hmeed, 2019).

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 (Felsberger et al., 2019).

2. AVAILABILITY

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 (Oppenheimer, 2010).

According to Oppenheimer (2010), availability means how long the network is operational. Availability is linked to reliability, but it has a more specific meaning (percent uptime) than reliability. Reliability refers to a variety of issues, including accuracy, error rates, stability, and the amount of time between failures. Availability is closely linked to resilience, a concept that is becoming more common in the networking field. Resiliency is how much stress a network can bear and how rapidly it can bounce back from problems, including security breaches, natural and unnatural disasters, human error, and catastrophic software or hardware failures(Cisco Systems, 2004).

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. Five Nine translates to five minutes of downtime a year (Thulin, 2004). Figure 1 shows a comparison of weekly downtime according to different availability percentages.

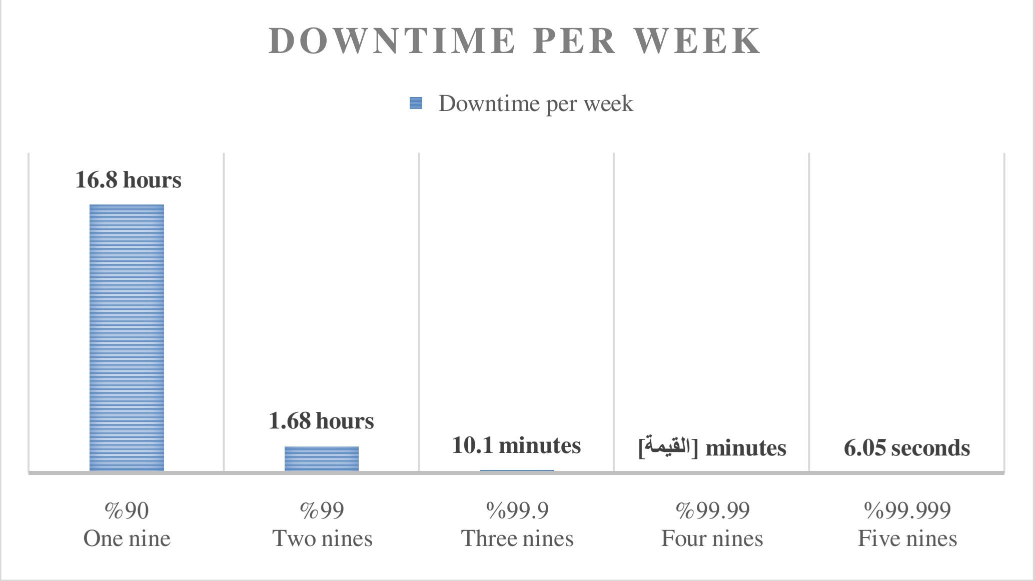

FIGURE 1: Downtime According to Different Availability Percentages.

To determine theoretical availability, the network is separated into each dependent item, such as hardware, software, physical connections, power supply, etc. For most equipment, the manufacturer will offer information on availability expectations, generally characterized as the mean time between failures (MTBF). For those elements of the network that do not have this data, such as a power source, statistical data and guesses must be employed. The projected time to repair each portion of the network has to be calculated. This is generally referred to as the mean time to repair (MTTR) (Mansour et al., 2022). Each unit's availability is determined by:

$$Availability = \frac{MTBF}{MTBF + MTTR} \cdots\cdots\cdots\cdots\cdots\cdots\cdots(1)$$

Mean time between failure (MTBF) is a common unit to measure reliability. MTBF is the average expected interval between failures of a product in its steady state (Cisco Systems, 2004). To compute the overall availability of the network, the availability of all units has to be totaled together.

Logically, a chain can never be stronger than its weakest link. Adding redundancy will result in greater availability. However, adding redundancy does not necessarily boost availability in a linear sense. A switchover from one route to another takes time, and during this moment, the connection will be offline (Thulin, 2004).

3. COST 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.

According to Research, the losses associated with network downtime include:

- Reputation
- Productivity losses
- Opportunities
- Data

In a July 2009 white paper titled “Navigating Network Infrastructure Expenditures during Business Transformations,” authored by Lippis Consulting, the cost of network downtime for a financial firm’s brokerage business was determined to be $7.8 million per hour. A one-hour interruption for a financial firm’s credit card processing can cost upwards of $3.1 million. A media organization might lose money on pay-per-view revenues, an airline company in ticket sales, and a retail company in catalog sales (Sholomon & Kunath, 2011).Figure 2 offers a few additional instances of losses by industry sector.

According to Opsworks (2022), in 2020, ITIC research found that since 2016, the average cost of a one-hour layover had grown by 30%. The bottom line is that of the 1,000 firms that participated in the study, more than 30% reported spending between $1 and $5 million on one hour of downtime. Meanwhile, over $300,000 is worth 1 hour of downtime for approximately 80% of enterprises. Finally, 98% claimed that one hour of downtime costs them roughly $100,000.

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(Didio, 2022).

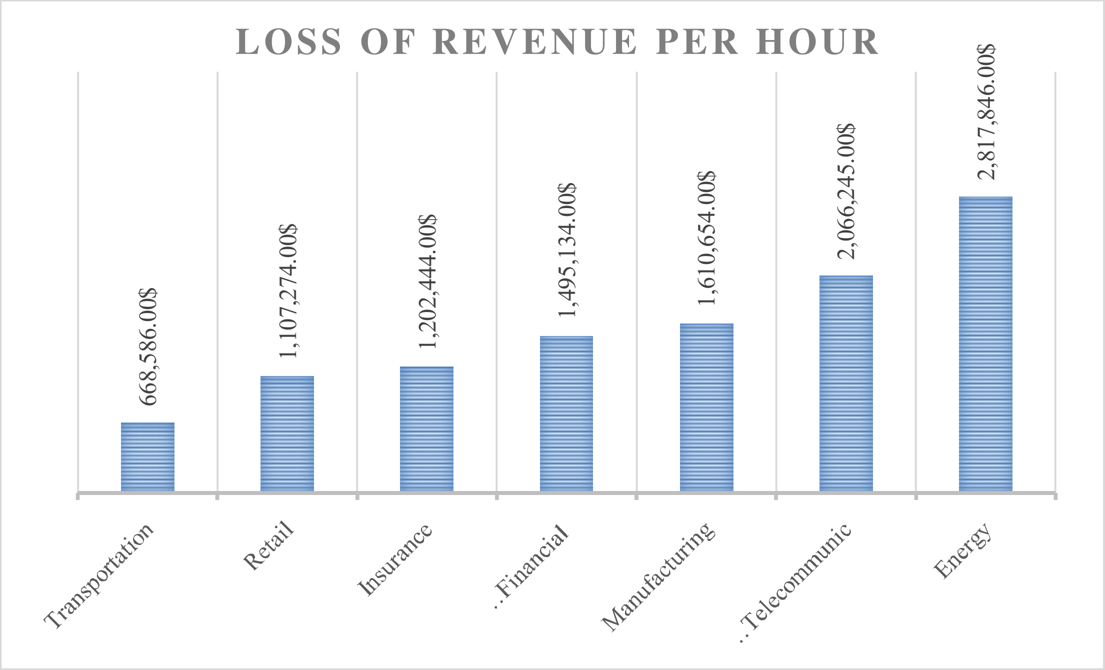

FIGURE 2: Loss of Revenue for Network Downtime by Several Industries.

4. RELATED WORK

In Ben Hassan (2024), “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 packet loss, convergence time, CPU utilization, and bandwidth consumption.

In Niu and Li (2023), “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), “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), “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 a previous study by Imelda et al. (2020), titled "Performance Analysis of VRRP, HSRP, and GLBP with the EIGRP Routing Protocol," a comparison in performance between VRRP, HSRP, and GLBP was introduced, and the EIGRP routing protocol was applied.

In Mansour (2021), “Performance Analysis and Functionality Comparison of First Hop Redundancy Protocols," the effect of FHRPv4 in terms of different parameters, mainly bandwidth consumption, traffic flow, convergence time, and CPU utilization, is analyzed.

In a previous study, Mansour (2020), under the title “Performance Evaluation of First Hop Redundancy Protocols,” investigated the impact of several factors such as CPU utilization, bandwidth consumption, and traffic flow without clarifying the difference in convergence time.

In study Kim et al. (2019), “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.

5. FIRST HOP REDUNDANCY PROTOCOLS

First Hop Redundancy Protocol (FHRP) is a suite of protocols that enables a router on a network to instantly take over if the main default gateway router fails. The devices in a shared network segment are set with a single default gateway address, which relates to the router that connects to the rest of the network. The trouble emerges when this main router fails, and there is a second router on the segment that is also capable of becoming the default gateway, but end devices don’t know about it. Hence, if the initial default gateway router fails, the network will stop (Dubey et al., 2013). First hop redundancy protocols are one of the solutions to this problem. The three primary First Hop Redundancy Protocols are:

- Hot Standby Router Protocol (HSRP; Owned by cisco).
- Virtual Router Redundancy Protocol (VRRP; Open Standard).
- Gateway Load Balancing Protocol (GLBP; Owned by cisco).

First hop redundancy techniques like as HSRP and VRRP offer default gateway redundancy with one router functioning as the active gateway router with one or more additional routers retained in standby mode. While others like GLBP allows all available gateway routers to load share and be operational at the same time (Dubey et al., 2013).In this paper we use only Hot Standby Router Protocol.

5.1 Host Standby Routing Protocol

HSRP is a redundancy protocol that allows failover to the default gateway. The active-standby model supports end-user traffic with one device at a time and one on standby to take over if the active device fails.

HSRP routes IP traffic without relying on the availability of any single router. It enables a set of router interfaces to work together to present the appearance of a single virtual router or default gateway to the hosts on a LAN.

HSRP allows a set of routers to function in harmony, giving the hosts on the LAN the illusion of a single virtual router. This set is known as an HSRP group or a standby group. A single router chosen from the group is responsible for forwarding the packets that hosts provide to the virtual router. This router is known as the active router, as shown in Figure 3; another router is designated as the backup router. In the case that the active router fails, the standby will take over the active router's packet forwarding tasks. This procedure is transparent to users. Although an arbitrary number of routers may run HSRP, only the active router transmits the packets sent to the virtual router. Devices in an HSRP group chose the active router based on device priority. To lower network traffic, only the active and backup routers send periodic HSRP messages once the protocol has finished the election process (Li et al., 1998).

Each standby group emulates a single virtual router. For each standby group, a single well-known virtual MAC and IP address are given to the group. The IP address should belong to the principal subnet in use on the LAN but must vary from the addresses allocated as interface addresses on all routers and hosts on the LAN, including virtual IP addresses issued to other HSRP groups.

Multiple hot standby groups could be configured; each group operates independently of other groups (Liet al., 1998).

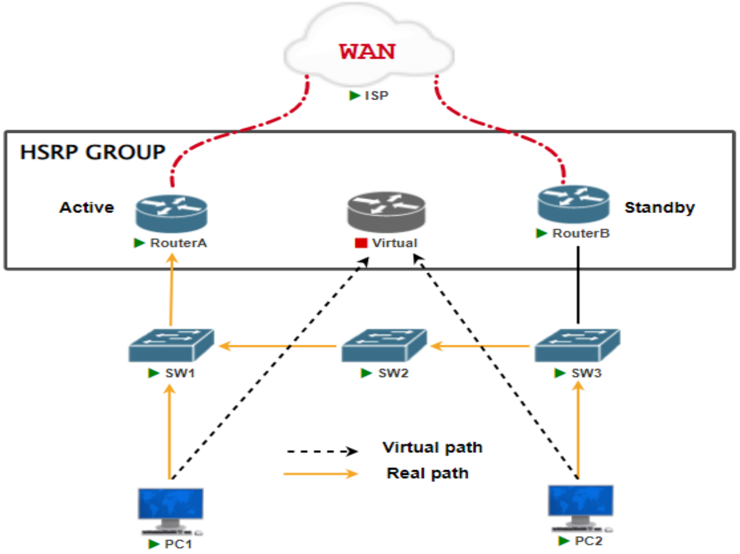

FIGURE 3: HSRP Operation.

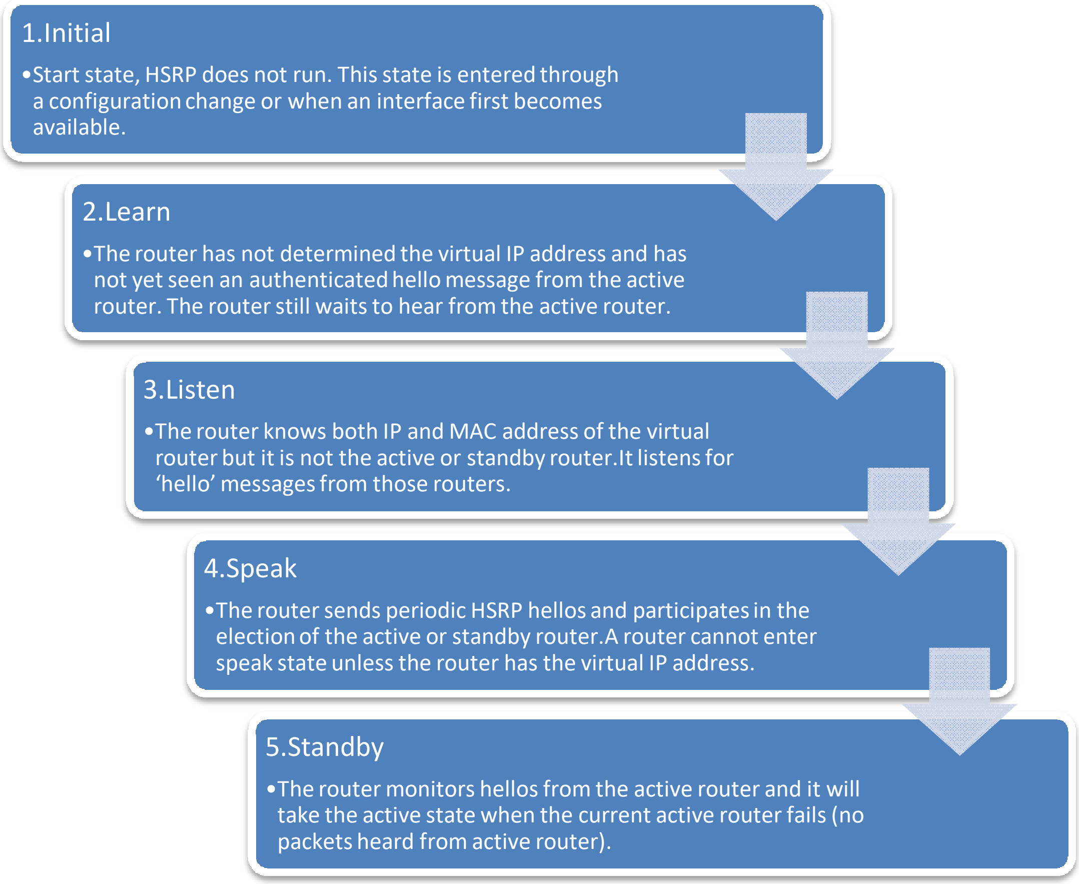

1.Initial

- Start state, HSRP does not run. This state is entered through a configuration change or when an interface first becomes available.

2.Learn

- The router has not determined the virtual IP address and has not yet seen an authenticated hello message from the active router. The router still waits to hear from the active router.

3.Listen

- The router knows both IP and MAC address of the virtual router but it is not the active or standby router.It listens for ‘hello’ messages from those routers.

4.Speak

- The router sends periodic HSRP hellos and participates in the election of the active or standby router.A router cannot enter speak state unless the router has the virtual IP address.

5.Standby

- The router monitors hellos from the active router and it will take the active state when the current active router fails (no packets heard from active router).

FIGURE 4: HSRP Status.

The Hot Standby Redundancy Protocol (HSRP) has two timers:(1)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.(2)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.These timers can be tuned and tweaked to achieve the lowest convergence, making a network highly accessible.

6. 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 multihop routed paths. Furthermore, it operates independently on the transmission media, data protocol, and routing protocol, without any need to modify the existing protocols(Kim & Ryu, 2019).

BFD does not have a discovery mechanism; sessions must be explicitly configured between endpoints. BFD may be used on many different underlying transport mechanisms and layers and operates independently of all of these. Therefore, it needs to be encapsulated by whatever transport it uses. For example, protocols that support some form of adjacency setup, such as OSPF, IS-IS, BGP, or RIP, may be used to bootstrap a BFD session. These protocols may then use BFD to receive faster notification of failing links than would normally be possible using the protocol's own keepalive mechanism.

In BFD, there is a set of parameters used to determine the failure detection time:

- Detect Multi:Detection timeout multiplier is the number of packets that have to be missed in a row to declare the session to be down.
- Required Min Rx Interval (RMRI): minimum interval for receiving BFD control packets.
- Desired Min Tx Interval (DMTI): minimum interval for sending BFD control packets.
- Required Min Echo RX Interval (RMERI): minimum interval for receiving Echo packets.

These parameters are within the BFD control packet that will be sent.

6.1 BFD Detection Modes

There are two operating modes to BFD, asynchronous mode and demand mode.The asynchronous mode is similar to the hello and hold-down timers. The system periodically sends BFD control packets. The system considers that the session is down if it does not receive any BFD control packets within a specific intervalas shown in Figure 5. In demand mode, a system sends several BFD control packets that have the Poll (P) bit set at the negotiated transmit interval as shown in Figure 6. If no response is received within the detection interval, the session is considered down. If the connectivity is found to be up, no more BFD control packets are sent until the next command is issued(Hewlett Packard Enterprise, 2017).

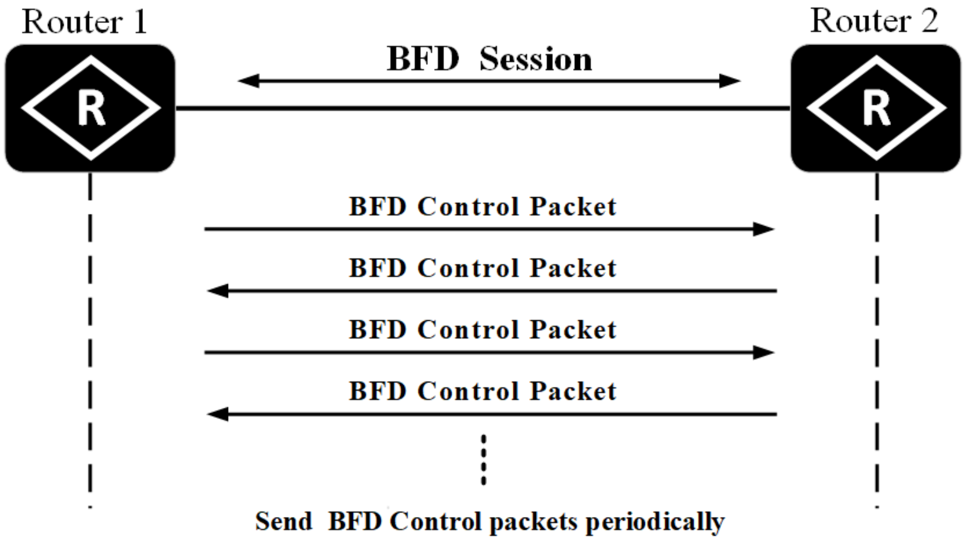

FIGURE 5: Asynchronous Mode.

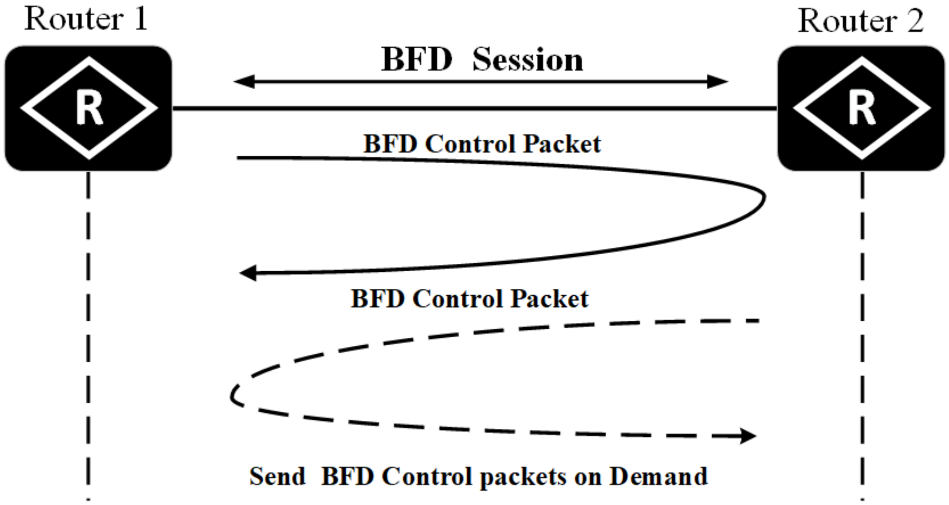

FIGURE 6: Demand Mode.

6.2 BFD Detection Time

The detection time (the period of time without receiving BFD packets after which the session is determined to have failed) is not carried out explicitly in the protocol. Rather, it is calculated independently in each direction by the receiving system based on the negotiated transmit interval and the detection multiplier. There may be different detection times in each direction (Katz & Ward, 2010).

In asynchronous mode, the detection time calculated in the local system is equal to the value of Detect Mult received from the remote system, multiplied by the agreed transmit interval of the remote system (the greater of the required min Rx interval and the last received desired min TX interval). Detection time in asynchronous mode = received Detect Multi of the remote system x max (local RMRI/received DMTI).

In demand mode, the detection time calculated in the local system is equal to the value of Detect Multi of the local system, multiplied by the agreed transmit interval of the remote system (the greater of required min Rx interval and the last received desired min TX interval).Detection time in demand mode = Detect Multi of the local system x max (local RMRI/received DMTI).

6.3 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 Figure 7; 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(Cisco, 2017).

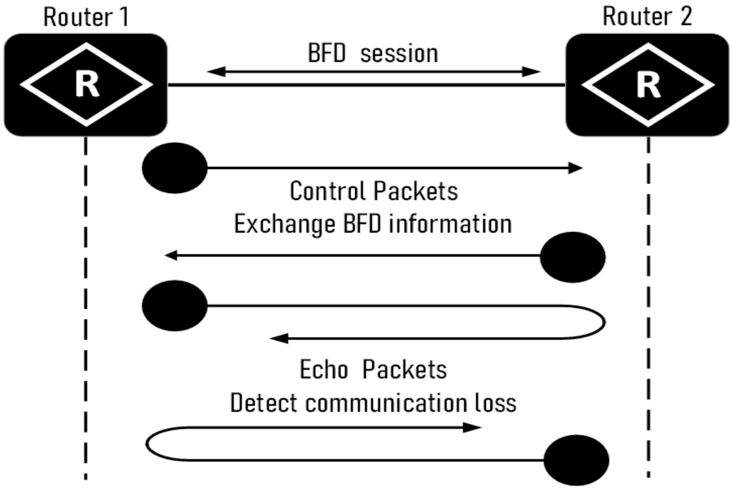

FIGURE 7: BFD Asynchronous Mode with Echo Function.

7. DESIGN AND SIMULATION

This paper focuses on implementing the Host Standby Router protocol with Bidirectional Forwarding Detection and evaluating the performance in comparison to HSRP without BFD. The enterprise site is connected to two different ISPs to ensure high availability. In the event of a connection failure between an ISP and a gateway, or if the ISP experiences a period of unavailability, the gateway will promptly identify the breakdown. This will enable the backup gateway, which is connected to the other ISP, to take over and assume control. This approach effectively minimizes network downtime, a crucial objective for enterprises operating in contemporary network environments.

7.1 Simulation Tools

In this work, PNELAB network emulator software was used to implement network scenarios. captured data such as convergence times, packet loss, and bandwidth consumption during failover scenarios using the Wireshark tool.

7.2 Network Design

The network is designed hierarchically to have two default gateway routers, each connected to a different ISP. In this work, PNELAB network emulator software was used to implement network scenarios. On the LAN side, there are two access switches connecting to end devices.

Access switches are connected to gateway routers in a partial mesh network design in order to eliminate single points of failure in the enterprise network. The EIGRP routing protocol is applied to provide routing between nodes in an enterprise. The topology is shown in Figure 8.

Router has a track object that is used to verify the connection to the ISP. In case the connection goes down, the track object decrements a value for the priority of the active router, which will make it have less priority than the standby, and it will result in making the standby become the active router. In the network design topology, routers on the left, R1, have been configured with higher priority than the routers on the right, R2.

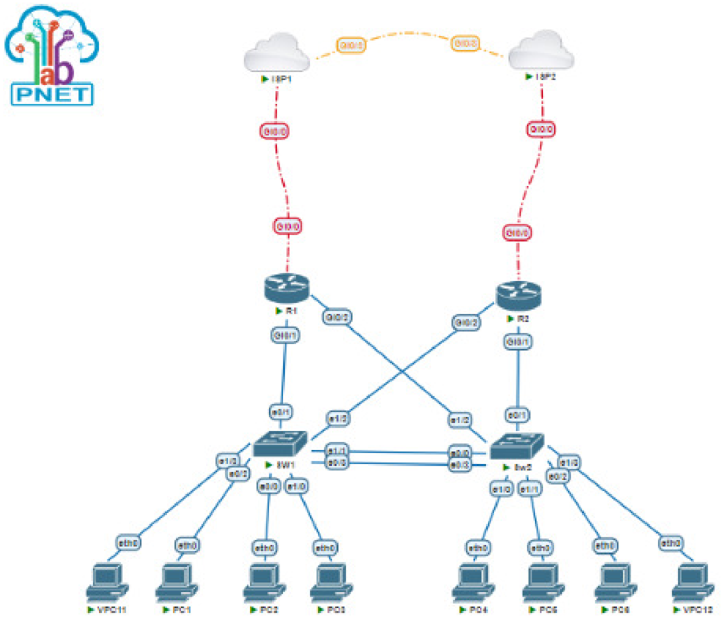

FIGURE 8: Network Topology.

8. CONFIGURATION

Initially, the HSRP 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. IP SLAs are network performance measurement and diagnostic tools that use active monitoring. One of its purposes is to verify whether a given IP address is reachable and report the status. The IP SLA will be configured on the enterprise routers to check the reachability of the ISPs. If an ISP becomes unreachable, the IP SLA will detect this loss of connectivity and report it to the HSRP installed on the router. A track object tied to the IP SLA will detect an ISP is down and reduce the router's priority value, allowing a higher-priority router to become the active router. We set the detection time to 1 second, the fastest available for IP SLA.

When implementing HSRP with BFD, BFD is operated to monitor the connectivity to ISPs, detecting access issues within milliseconds. BFD sessions are configured on the enterprise router to check ISP reachability. We set the detection time to 50 milliseconds, the fastest available for BFD. This minimizes the time it takes to detect a link failure compared to IP SLA. The configured BFD on R1.

Simulation parameter	Value
Number of ISP	2 ISP (ISP1- ISP2), Cisco Version 15.7(3)M3
Number of Routers	2 Routers, Cisco Version 15.7(3)M3
Number of Switches	2 Switches, Cisco Version 15.1
PCs Numbers	8 PCs
Traffic Generator	Iperf
Packet Rate	3600 Packet
Test Duration	1 Hour
HSRP- Hello –Hold time	Default 3 10 With Optimization 1 3
BFD Detection Time	50 milliseconds
IP SLA Detection Time	1 second

TABLE 1: Simulation Parameters.

9. RESULTS

This section will present and discuss the measurements conducted in order to test the performance of HSRP, both with and without BFD, and then present and evaluate the results of HSRP without BFD compared to HSRP 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. The obtained data will be utilized to examine the influence of BFD on HSRP performance and its effectiveness in minimizing downtime. The measurements are taken in terms of convergence time, CPU utilization, and bandwidth consumption. 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.

9.1 HSRP without BFD Results

Convergence Time:

- By using default timers for hello and hold, the convergence time between R1 and R2, where R2 transitions to the active state, is equal to 7.3 seconds, from the last packet received by ISP1 before the failure at 14:25:51.62 to the moment when R2 sends an advertisement as the active state at 14:25:58.92. The convergence process between ISP1 and ISP2 takes is equal to 9 seconds. From the last packet received by ISP1 before the failure and the first packet received by ISP2 after the failure at 14:26:00.63. During the convergence process, 4 ICMP packets were lost.

- By using optimized timers for hello and hold, the convergence time between R1 and R2, where R2 transitions to the active state, is equal to 3.78seconds. From the last packet received by ISP1 before the failure at 15:01:11.99 to the moment when R2 sends an advertisement as the active state at 15:01:15.77. The convergence process between ISP1 and ISP2 takes is equal to 5 seconds. From the last packet received by ISP1 before the failure and the first packet received by ISP2 after the failure at 15:01:16.99. During the convergence process, 2 ICMP packets were lost.

CPU Utilization:

- Without timers' optimization, HSRP consumed an average of 0.05% of the CPU usage on routers R1 and R2, while the CPU usage was 2% on R1 and 1% on R2.

- With timers' optimization, HSRP consumed an average of 0.09% of the CPU usage on R1 and 0.07% on R2, while the CPU usage was 2% on R1 and 1% on R2.

Bandwidth Consumption:

- During the testing period without timers' optimization, the traffic generated by HSRP packets accounted for approximately 172 KB, with a total of 2792 hello packets exchanged between routers R1 and R2. 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 HSRP packets accounted for approximately 499 KB, with a total of 8052 hello packets exchanged between routers R1 and R2. This estimate is based on the optimized hello packet interval, where packets are sent every 1 second.
- During the testing period, the traffic generated by IP SLA packets accounted for approximately 562 KB, a count of 7200 ICMP packets exchanged between routers R1 and ISP1 every 1 second.

Figure 9. (a) shows the bandwidth consumption of HSRP packets without optimization sent by R1 and R2, where HSRP packets consist of a hello packet of fixed size of 62 bytes sent every 3 seconds and an advertise packet of 60 bytes sent from time to time. In most seconds, only one 62-byte packet is sent per second from R1 or R2. However, in some seconds the bandwidth will exceed 122 bytes because a packet from R1 and a packet from R2 are sent in the same second. In rare cases, the bandwidth will exceed 182 bytes because three packets are transmitted simultaneously: a hello packet from both R1 and R2, along with an HSRP advertise packet from either R1 or R2.

Figure 9. (b) illustrates the bandwidth consumption of HSRP packets after optimization, where R1 and R2 each send hello packets every second. In most instances, the bandwidth will exceed 182 bytes because three packets are transmitted simultaneously: a hello packet from both R1 and R2, along with an HSRP advertisement packet from either R1 or R2. In rare cases, the bandwidth will exceed 244 bytes, as four packets are sent within the same second: a hello packet from both R1 and R2, along with an HSRP advertise packet from both routers in the same second.

Figure. 9. (c) shows the bandwidth consumption of IP SLA packets sent between R1 and ISP1, with each packet having a fixed size of 78 bytes. In every second, 156 bytes of data are transferred fixed due to the consistent exchange of two ICMP packets between R1 and ISP1, one a request from R1 and one a response from ISP1, and this happens evenly over time.

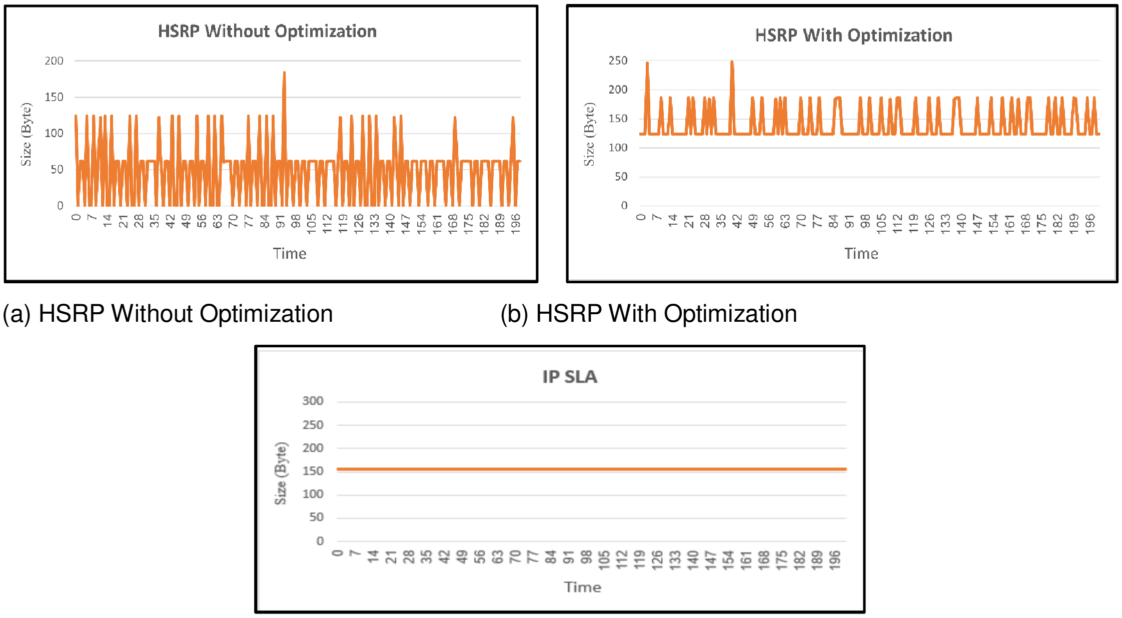

(a) HSRP Without Optimization

(b) HSRP With Optimization

(c) IP SLA

FIGURE 9: Bandwidth Consumption of IP SLA and HSRPpacketsinBytes/Sec.

9.2 HSRP with BFD Results

Convergence Time:

- By using default timers for hello and hold, the convergence process between ISP1 and ISP2 takes is equal to 1 second. From the last packet received by ISP1 before the failure at 16:27:58.34 and the first packet received by ISP2 after the failure at 16:27:59.34. The convergence time between R1 and R2 is equal to 6.38 seconds. From the last packet received by ISP1 before the failure to the moment when R2 sends an advertisement in the active state at 16:28:04.72. During the fast convergence process, no ICMP packets were lost.

- By using optimized timers for hello and hold, the convergence process between ISP1 and ISP2 takes is equal to 1 second. From the last packet received by ISP1 before the failure at 16:28:11.34 and the first packet received by ISP2 after the failure at 16:28:12.34. The convergence time between R1 and R2 is equal to 3.33 seconds. From the last packet received by ISP1 before the failure to the moment when R2 sends an advertisement in the active state at 16:28:14.67. During the fast convergence process, no ICMP packets were lost.

CPU Utilization:

- During the testing period, BFD consumed an average of 2.57% of the CPU usage on routers R1, while the CPU usage was 5% on R1 and 1% on R2.

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. A total of 321228 BFD packets were exchanged between R1 and ISP1, of which 313036 were echo packets and 8192 were control BFD packets. As mentioned previously, BFD Echo packets are responsible for detecting failures, while BFD Control packets maintain the BFD session between R1 and ISP1.

Figure10shows the bandwidth consumption of BFD packets sent between R1 and ISP1, where each BFD echo packet has a fixed size of 54 bytes, while the BFD control packet has a size of 66 bytes. Every second, the amount of data transmitted varies between 4600 and 5200 bytes, thus the pattern is unstable. The number of packets transmitted every second is different from the other due to the rapid transmission of BFD echo packets every 50 milliseconds and BFD control packets every second.

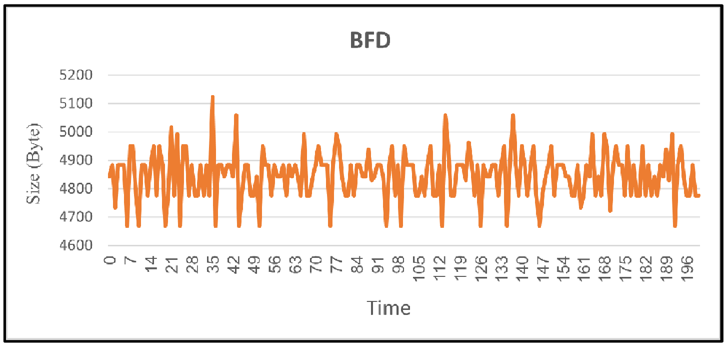

FIGURE 10: Bandwidth Consumption of BFD packetsinBytes/Sec.

The results indicate that the integration of BFD with HSRP markedly enhances network performance by diminishing convergence time and packet loss. The implementation of BFD reduced the convergence time from 7.3 seconds in the HSRP without BFD to just 1 second, which can be attributed to its rapid failure detection capacity. This mitigated network downtime, thereby improving network availability. Compared to Mansour (2022) and Ben Saud (2023), who evaluated FHRP without BFD, our research demonstrates shorter convergence times and fewer packet losses. InMansour (2022) reported a minimum convergence time for HSRP of 3.271 seconds with one lost packet after timers optimization, while Ben Saud (2023) observed a minimum convergence time for HSRPv2 of 2.94 seconds with two lost packets after timers optimization.However, our BFD implementation attained a convergence time of 1 second without any packet loss, signifying a substantial enhancement. Furthermore, Kim and Ryu (2019) concentrated on VRRP but neglected to include ISP-level failure scenarios, which our study thoroughly addresses using HSRP. Nevertheless, these benefits, BFD has a trade-off characterized by heightened resource consumption. CPU utilization climbed from 2% to 5%, while bandwidth consumption escalated from 562 KB with IP SLA to 16.54 MB with BFD.

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.

10. COMPARISON AND EVALUATION

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

10.1 Convergence Time

We can see from Figure 11 that HSRP 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 HSRP without BFD, which has an IP SLA failure detection time of 1 second. Meanwhile, HSRP with BFD-Optimize has the best convergence time to switch between active and standby mode at 3.33 seconds; this is because of the optimized hello packet sent every 1 second combined with the BFD failure detection time.

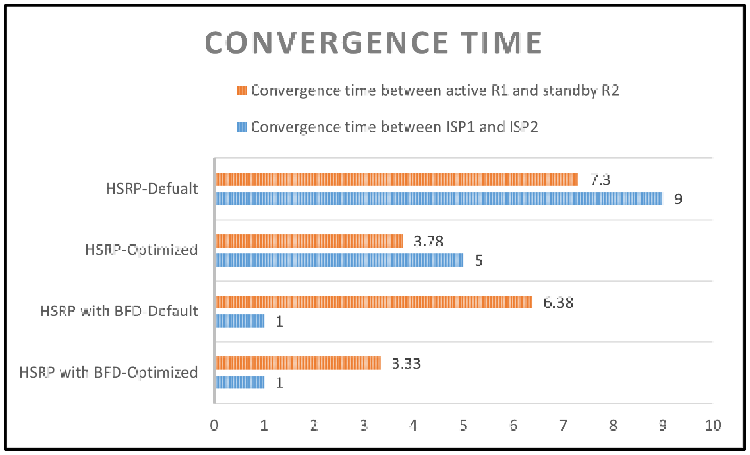

FIGURE 11: Convergence Time Comparison.

10.2 Packet Loss Comparison

Figure12 shows that during convergence, for HSRP without BFD, 4 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 2 packets were lost thanks to the optimized hello packet sent every 1 second. while for HSRP with BFD, no packets were lost either before optimization or with optimization, thanks to a BFD failure detection time of 50 milliseconds.

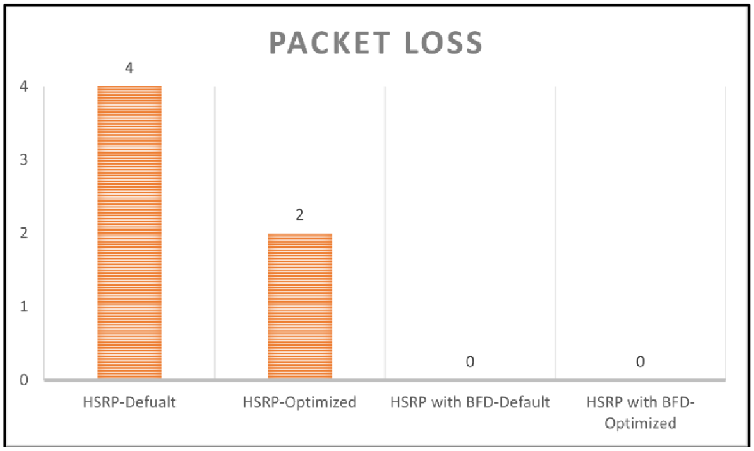

FIGURE 12: Packet Loss Comparison.

10.3 CPU Utilization Comparison

Figure 13 shows the increase in CPU usage observed when using HSRP 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 HSRP with BFD has the worst CPU usage compared to HSRP without BFD.

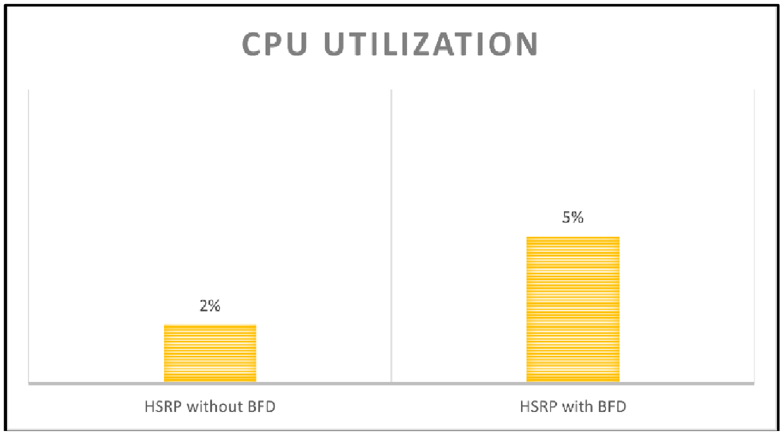

FIGURE 13: CPU Utilization Comparison.

10.4 Bandwidth Consumption

Table 2 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.

Protocols	Bandwidth Consumption
HSRP-Default	172KB
HSRP-Optimized	499KB
IP SLA	562KB
BFD	16.54MB

TABLE 2: Bandwidth Consumption Comparison.

11. CONCLUSION

After implementing and testing HSRP without and with BFD and studying and analyzing its outputs for four important factors, namely convergence time, packet loss, CPU utilization, and bandwidth consumption, it is clear that using BFD with HSRP significantly enhances network performance. The convergence time improved from 7.3 seconds to 1 second, and no packet loss occurred during failure scenarios. However, the cost was increased CPU utilization and high bandwidth consumption. Therefore, organizations must carefully balance the benefits of reduced packet loss and faster failover times with the resource demands of BFD. The use of BFD depends on the importance of business downtime, provided that sufficient resources are available to meet CPU and bandwidth requirements.

This paper is valuable for network engineers, IT administrators, and decision-makers trying to strengthen the resilience of their networks. These studies illustrate practical ramifications for businesses such as telecom corporations, airports, healthcare, and others, where even modest downtime can result in huge financial losses. By using BFD with several FHRPs, enterprises can ensure continuous service availability and reduce operational interruptions. This paper also gives significant information for academics and researchers focused on network replication strategies.

12. REFERENCES

Ben Hassan, A., & Mansour, M. (2024). Performance Evaluation of Bidirectional Forwarding Detection (BFD) over The Virtual Router Redundancy Protocol (VRRP).*In The 15th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN 2024),251,*256-264.

Ben Saud, N., & Mansour, M. (2023). Performance Evaluation of First Hop Redundancy Protocols in IPv4 and IPv6 Networks.*In 2023 IEEE 3rd International Maghreb Meeting of the Conference on Sciences and Techniques of Automatic Control and Computer Engineering (MI-STA)*. https://doi.org/10.1109/MI-STA57575.2023.10169462.

Cisco. (2017, November 3). *Routing configuration guide, Cisco IOS XE Everest 16.6.X (Catalyst 9500 switches)*. Cisco.
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.

Cisco Systems, Inc. (2004). *Network availability: How much do you need? How do you get it?*. Cisco.https://www.cisco.com/web/IT/unified_channels/area_partner/cisco_powered_network/net_availability.pdf.

Didio, L. (2022, November 30). *Server and application reliability by the numbers: understanding "The Nines",* ITIC. https://itic-corp.com/server-and-application-by-the-numbers-understanding-the-nines/.

Dubey, P., Sharma, S., & Sachdev, A. (2013). Review of First Hop Redundancy Protocol and Their Functionalities. *International Journal of Engineering Trends and Technology,4*(5),1085-1088.

Felsberger, L., Todd, B., &Kranzlmüller, D. (2019). Cost and Availability Improvements for Fault-Tolerant Systems Through Optimal Load-Sharing Policies. *Procedia Computer Science*, *151*, 592–599. https://doi.org/10.1016/j.procs.2019.04.080.

Hewlett Packard Enterprise. (2017, December). *High availability configuration guide: Bidirectional forwarding detection*. HewlettPackardEnterprise.
https://support.hpe.com/techhub/eginfolib/networking/docs/switches/5930/5200-4571_hi-avail_cg/content/491355164.htm.

Julia, I. R., Suseno, H. B., Wardhani, L. K., Khairani, D., Hulliyah, K., & Muharram, A. T. (2020, October). Performance Evaluation of First Hop Redundancy Protocol (FHRP) on VRRP, HSRP, GLBP with Routing Protocol BGP and EIGRP. In *2020 8th International Conference on Cyber and IT Service Management (CITSM)*, 1-5.https://doi.org/10.1109/CITSM50537.2020.9268799.

Katz, D., & Ward, D. (2010). Bidirectional Forwarding Detection (BFD). *Internet Engineering Task Force (IETF)*, RFC 5880.https://datatracker.ietf.org/doc/html/rfc5880.

Kim, S., & Ryu, H. (2019). FDVRRP: Router Implementation for Fast Detection and High Availability in Network Failure Cases. *ICT R&D program of MSIP/IITP,41*(4),473-482.https://doi.org/10.4218/etrij.2018-0309.

Li, T., Cole, B., Morton, P., & Li, D. (1998). *Cisco Hot Standby Router Protocol*, RFC 2281.https://www.rfc-editor.org/rfc/rfc2281.

Mansour, M. (2020). Performance Evaluation of First Hop Redundancy Protocols. In *The 11th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN 2020),177*,330–337. https://doi.org/10.1016/j.procs.2020.10.044.

Mansour, M., &Hmeed, A. (2019). A Comparison of Queueing Algorithms Over TCP Protocol. *International Journal of Computer Science and Security (IJCSS)*, *13*(6), 275-294.

Mansour, M., Agomati, M., Alsaid, M., Berrish, M., &Alasem, R. (2022). Performance Analysis and Functionality Comparison of First Hop Redundancy Protocol IPV6. In *The 13th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN 2022)*,*210*,19-27. https://doi.org/10.1016/j.procs.2022.10.115.

Mansour, M., Ghneimat, A., Alasem, R., &Jarray, F. (2021). Performance Analysis and Functionality Comparison of First hop Redundancy Protocols. *Journal of Ubiquitous System & Pervasive Networks, 15*(1), 49-58.

Niu, Y., &Li, X. (2023, October). 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*, 197–201. https://doi.org/10.1145/3650215.3650250.

Oppenheimer, P. (2010). *Top-down network design* (3rd ed.). Cisco Press.

OpsWorks. (2022, September 12). *The cost of downtime: The truth and facts of IT downtime.*OpsWorks.https://opsworks.co/cost-of-downtime-truth-and-facts-of-it-downtime.

Sholomon, A., & Kunath, T. (2011). *Enterprise network testing* (1st ed.). Cisco Press.

Thulin, M. (2004). *Measuring availability in telecommunications networks*. KTH.https://people.kth.se/~e98_thu/thesis/Availability_Thesis_final.pdf.