International Journal of Networks and Communications 2026, 15(1): 13-19
DOI: 10.5923/j.ijnc.20261501.02

Network Resilience & Fault Tolerance: Enhancing Hybrid Network Robustness Against Failures and Cyber Attacks

Mwaba Besa*, Simon Tembo

University of Zambia, School of Engineering, Department of Electrical and Electronic Engineering, Lusaka, Zambia

Abstract Modern networks have become an indispensable part of our daily lives, facilitating communication, industrial processes, and essential services. However, these networks are increasingly vulnerable to a range of disruptions, including hardware failures, software malfunctions and cyber-attacks. Such vulnerabilities can lead to significant downtime, financial losses, and compromised critical services. Current solutions for network resilience often fall short due to their lack of adaptability to dynamic network conditions and limited capabilities in handling unexpected large-scale failures. This research proposes a comprehensive approach to enhance network resilience and fault tolerance by focusing on three key areas; investigating fault-tolerant routing protocols i.e., ensuring seamless connectivity during link failures, developing automated recovery mechanisms to minimize downtime through real-time and efficient recovery strategies, strengthening hybrid network security from cyber-attacks. The study uses an experimental design with simulation-based modelling to replicate disaster scenarios that can't be tested in live networks. It enables controlled fault injections, protocol behaviour analysis, and cybersecurity testing. Both quantitative metrics (e.g., convergence time, jitter, packet loss) and qualitative insights (e.g., protocol behaviour, security risks) are evaluated. Results showed that automated recovery mechanisms, and robust firewall defenses can significantly reduce downtime and improve reliability. Effective recovery strategies were validated, with results confirming that networks can recover quickly and continue operating with minimal disruption.

Keywords Fault-Tolerant, Resilience, Hybrid, Jitter, Packet-loss

1. Introduction

The digital infrastructure that powers today's interconnected world is built on complex networks, ranging from the internet and telecommunication systems to power grids and transportation networks. In an era where digital communication underpins nearly every facet of human activity, from healthcare systems and emergency services to global finance and critical infrastructure, network availability and reliability have become indispensable. As the backbone of modern civilization, computer networks are expected to remain operational around the clock, even amidst faults, failures, or cyber threats. However, with the rapid proliferation of distributed systems, cloud computing, and the Internet of Things (IoT), today's networks face unprecedented challenges in maintaining consistent performance, especially under threats? This has made network resilience and fault tolerance not just a technical necessity, but a strategic imperative for organizations worldwide. Network robustness is a critical factor determining the efficiency and reliability of systems that underpin communication, data transmission, and infrastructure. The resilience of these networks is deeply influenced by their underlying graph topology, which governs how nodes and connections are structured.

Network resilience refers to a system's ability to maintain acceptable service levels in the face of faults and disruptions. Meanwhile, fault tolerance encompasses the techniques and mechanisms that enable a network to continue functioning even when some of its components fail. Together, these concepts form the bed-rock of dependable network design, enabling infrastructures to withstand and recover from diverse threats such as hardware failures, software bugs, misconfigurations, natural disasters, and sophisticated cyberattacks.

Traditional approaches to network design often prioritize performance and scalability over resilience, leaving systems exposed to cascading failures and prolonged downtimes. Furthermore, the increasing complexity and interdependence of networked systems mean that a single point of failure can ripple through entire ecosystems. This has sparked a growing research interest in the development of intelligent, fault-tolerant mechanisms, robust routing proto-cols, and automated

* Corresponding author:
mwabasuwilanji1996@gmail.com (Mwaba Besa)
Received: Feb. 18, 2026; Accepted: Mar. 9, 2026; Published: Mar. 13, 2026
Published online at http://journal.sapub.org/ijnc

recovery strategies that can pre-emptively address or swiftly mitigate failures with-out human intervention. This paper explores techniques to enhance network resilience and fault tolerance, focusing on three core pillars: tolerant routing protocols, automated network recovery mechanisms, and disaster recovery planning. It investigates how modern simulation tools can be leveraged to model fault scenarios and evaluate the effectiveness of various mitigation strategies under realistic conditions. Through this study the goal is to identify practical, scalable, and adaptive solutions that can fortify modern networks against an increasingly hostile and uncertain operational landscape. Ultimately, this research aims to contribute to the body of knowledge by not only analysing current resilience strategies but also proposing novel frameworks that push the boundaries of what is achievable in self-healing, failure-aware, and resilient network architectures. To the best of our knowledge, no prior study has jointly evaluated routing fault tolerance and cybersecurity resilience in a single hybrid topology using pfSense and VyOS under controlled simulation conditions.

2. Literature Review

The rapid evolution of digital communication systems has made the reliability, robustness, and availability of network infrastructures more critical than ever. Modern networks are expected to maintain continuous connectivity and service delivery even in the face of diverse failures ranging from hardware malfunctions and software bugs to cyber-attacks and natural disasters. As networks grow in scale and complexity, so too does their vulnerability to a wide range of fault conditions [16]. This has driven significant academic and industry interest in the fields of network resilience and fault tolerance, which aim to ensure that networks can sustain operations, recover swiftly from failures, and adapt to changing conditions without compromising overall functionality [1].

Network resilience refers to a network's ability to withstand failures and attacks while maintaining essential operations. It encompasses several dimensions, including robustness, adaptability, self-healing, and survivability [11]. Fault tolerance, on the other hand, is a closely related concept that focuses on the ability of a system to continue operating correctly even when some of its components fail. Together, these concepts form the cornerstone of designing dependable communication systems that are crucial to supporting critical infrastructure, cloud computing, IoT deployments, and other mission-critical services [5].

The literature on resilient and fault-tolerant networking has grown significantly in the past two decades. Research efforts have explored a wide range of solutions, including fault-tolerant routing protocols, automated failure recovery mechanisms, and improve network behaviour under adverse conditions. Simulation tools such as GNS3 have become essential in enabling researchers to experiment with routing configurations and evaluate failure recovery strategies in controlled, scalable environments. These tools provide cost effective and flexible platforms for emulating real-world network scenarios without the need for physical hardware.

The purpose of this literature review is to critically examine and synthesize current research on techniques that enhance network resilience and fault tolerance, with a specific focus on technologies and protocols that can be tested through GNS3 simulations. By analysing foundational theories, established protocols, recovery frameworks, and experimental models, this Section aims to build a comprehensive understanding of the current state of the field, identify gaps in the literature, and highlight the limitations of existing approaches.

2.1. Impact of Natural Disasters on Network Stability

Natural disasters such as floods, earthquakes, hurricanes, wildfires, and lightning storms pose significant challenges to the stability and availability of communication networks. As society increasingly relies on inter-connected systems for essential services such as emergency response, health care, banking, and utilities, the failure of network infrastructures during disasters can lead to catastrophic consequences. The effects of natural disasters are not confined to a single layer of the net-work stack. They propagate across multiple layers of the OSI model, degrading both the performance and survivability of communication systems. This section reviews how specific types of disasters impact different network layers, drawing from real-world events, simulation based studies, and theoretical research.

2.1.1. Physical Layer Disruptions

The physical layer is highly vulnerable during disasters because it depends on tangible infrastructure like fibre cables, towers, and power lines. Earthquakes, floods, and hurricanes can sever cables, submerge equipment, or collapse towers, causing complete service outages. Mitigation includes geo-redundant links, reinforced underground conduits, flood -resistant designs, and backup power systems [2].

2.1.2. Data Link Layer Disruptions

Disasters indirectly affect the data link layer by degrading wireless signals and damaging switches or access points. This leads to frame collisions, dropped connections, and higher latency or jitter, especially when users overload surviving access points [11]. Resilience can be improved through link aggregation, self-healing mesh networks, and redundant switching fabrics.

2.1.3. Network Layer Disruptions

At the network layer, disasters destabilize routing and connectivity. The loss of routers or gateways creates black holes, route flapping, and long convergence delays. Studies show protocols like BGP adapt slowly to large-scale failures [3]. Fault-tolerant strategies such as ECMP, MPLS, and software-defined routing help maintain connectivity by providing alternate paths and programmable recovery.

2.1.4. Transport and Application Layer Disruptions

Failures cascade upward, disrupting sessions, ports, and service quality. Applications like VoIP, banking, and emergency systems are sensitive to jitter, latency, and packet loss [15]. Congestion can trigger TCP timeouts and DNS failures, while data centres in disaster zones may go offline. Resilience at these layers relies on regional load balancing, failover mechanisms, and replication or caching [2].

2.1.5. Cascading Failures across Layers

Disasters often trigger cascading failures across multiple layers. A fibre cut at the physical layer can destabilize links, break routing, disrupt sessions, and crash applications. These chain reactions affect critical sectors such as healthcare and finance [2]. Addressing this fragility requires holistic, cross-layer resilience strategies that integrate adaptive fault detection and mitigation to ensure continuity.

2.2. Fault-Tolerant Routing Protocols in Traditional Networks

Fault-tolerant routing ensures communication continues during failures by maintaining or quickly restoring paths. Traditional protocols differ in resilience, scalability and convergence speed [4].

2.2.1. Fault Recovery and Convergence

Protocols vary in recovery speed and redundancy. RIP is slow (up to 180 seconds), OSPF recalculates quickly using shortest-path algorithms, BGP offers policy control but converges slowly without enhancements, and EIGRP achieves rapid failover through backup paths and load balancing [14].

2.2.2. Redundant and Multipath Techniques

Multipath routing improves resilience by providing alternate routes. OSPF and EIGRP support equal-cost multipath, while MPLS enables predefined backup paths and fast reroute [14]. Tools like Bidirectional Forwarding Detection (BFD) allow failures to be detected in under 50 ms minimizing disruption.

2.2.3. GNS3-Based Studies

Simulation studies using GNS3 show protocol differences under failures. OSPF adapts within 10–30 seconds but depends on topology; EIGRP converges faster (4–8 seconds) due to proactive backup paths; RIP is slow (>90 seconds) and prone to loops. Comparative tests confirm EIGRP as fastest, OSPF moderate, and RIP un-suitable for fault-tolerant networks.

2.3. Cyber Security and Disaster-Induced Vulnerabilities

Natural disasters not only disrupt physical networks but also create cybersecurity risks. Damaged or misconfigured systems, reduced monitoring, and emergency workarounds give attackers opportunities to exploit weaknesses, steal data, or disrupt services. Integrating resilience and security is therefore essential in disaster-tolerant network design [8].

2.3.1. Cyber-Attack Vectors in Disaster Scenarios

Disasters increase exposure to attacks such as denial of service, man-in-the-middle interception, phishing, routing manipulation, and data theft. These threats are especially dangerous because human and technical response capacity is reduced during emergencies [7].

2.3.2. Vulnerabilities across Network Layers

At the physical layer, damaged monitoring systems and unsecured backup stations reduce visibility. At the data link and network layers, unencrypted Wi-Fi and weak routing authentication expose networks to spoofing and manipulation. At the transport and application layers, insecure VPNs, DNS hijacking, and VoIP abuse further compromise service reliability.

2.3.3. Case Studies of Cyberattacks during Disasters

Past events highlight these risks: after Hurricane Katrina, unsecured Wi-Fi enabled phishing and call interception; during the Nepal earthquake, poor key management led to VPN breaches and phishing scams; and during COVID-19, rushed remote access deployments were exploited for ransomware attacks on hospitals and logistics networks.

3. Methodology

The methodology integrates experimental simulations, comparative evaluations, and structured analyses. By combining emulation platforms with structured evaluation metrics, the study provides a reproducible and rigorous approach to understanding how networks behave under stress and what solutions enhance resilience.

3.1. Research Design

The study uses an experimental design with simulation-based modelling to replicate disaster scenarios that can't be tested in live networks. It enables controlled fault injections, protocol behaviour analysis, and cybersecurity testing. Both quantitative metrics (e.g., convergence time, jitter, packet loss) and qualitative insights (e.g., protocol behaviour, security risks) are evaluated.

3.2. Simulation Environment and Tools

3.2.1. GNS3 for Tradition Networks

GNS3 is used to emulate traditional IP networks with Cisco IOS, QEMU appliances, and Linux hosts. It sup-ports testing of routing protocols (OSPF, RIP, and BGP), simulating failures like link outages and router crashes, and deploying Linux VMs for traffic generation.

Table 1. IP Addressing Scheme

Component	Interface / Port	IP Address	Role/Function
Vyos router 1	Eth6- LAN Switch	192.168.1.1	LAN gateway
VyOS Router 1	Eth5- pfSense em0	10.0.0.2	WAN link to pfsense
VyOS Router 1	Eth7- VyOS router 2	10.0.2.1	Redundant link
VyOS Router 2	Eth5-LAN Switch	192.168.1.2	LAN gateway
VyOS Router 2	Eht7–VyOS router 1	10.0.2.2	Redundant link
PfSense firewall	Em0-VyOS router 1	10.0.0.1	WAN interface
pfSense firewall	Em2 - Cloud node	DHCP/WIFI	Internet uplink
pfSense firewall	Em1-LAN Switch	192.168.1.254	Default gateway
VPC 1	Eth0-LAN Switch	192.168.1.10	Client endpoint
VPC 2	Eth0-LAN Switch	192.168.1.11	Client endpoint
VPC 3	Eth0-LAN Switch	192.168.1.12	Client endpoint
Kali Linux host	Eth1- LAN Switch	192.168.1.50	Attack/Link failure simulation node
Cloud node	Wifi uplink	Public IP	External connectivity

3.2.2. Supplementary Tools

To support the simulation, the following tools are used for traffic analysis, fault injection, and security testing:

- Wireshark – for packet capture and analysis
- iPerf3 – for measuring throughput and latency
- Netem – for injecting jitter, latency, and packet loss
- Kali Linux – for simulating cyberattacks like DDoS, Port Scanning and routing poisoning

3.3. Experimental Setup

3.3.1. Network Topology Design

The experiment uses a hybrid topology built in GNS3, as shown in figure 1, integrating virtual routers, switches, pfSense firewall, and VMs. A cloud node represents upstream connectivity. Redundant paths and Layer 2/3 switches ensure realistic distribution and access. All components run on virtualized hosts using GNS3VM.

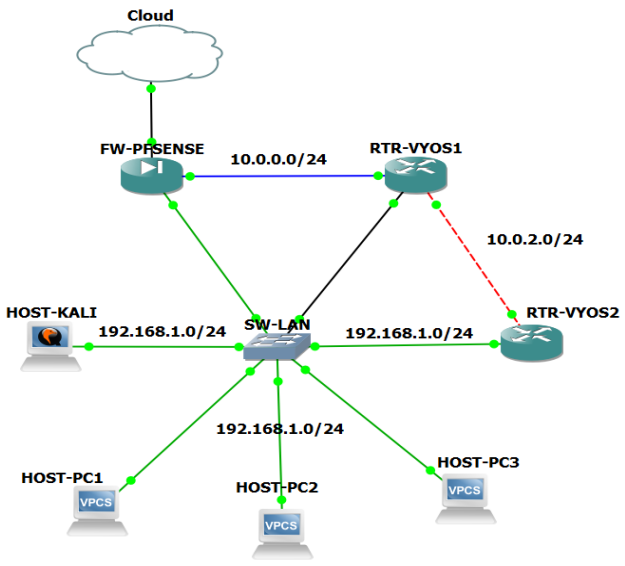

Figure 1. Network topology design

3.4. IP Addressing Scheme

The addressing scheme was designed to support both routing functionality and attack simulations. VyOS Router 1 and Router 2 were configured with redundant links (10.0.2.1/ 10.0.2.2) to test failover and convergence times. pfSense served as the default gateway (192.168.1.254) for all LAN devices, with its WAN interface (10.0.0.1) connected upstream to VyOS Router 1. VPC nodes (192.168.1.10 -12) acted as traffic sources and sinks, while the Kali Linux host (192.168.1.50) generated attack traffic. The cloud node provided external connectivity, enabling WAN-LAN interactions.

3.5. Data Collection and Analysis

To evaluate network resilience under fault conditions, data was collected through controlled simulations using GNS3 as shown in Figure 2. Faults such as link failures and node outages were injected into hybrid topologies running pfSense and VyOS routers. Performance metrics including latency, jitter, packet loss, and recovery time were captured Using Wireshark and iperf3.

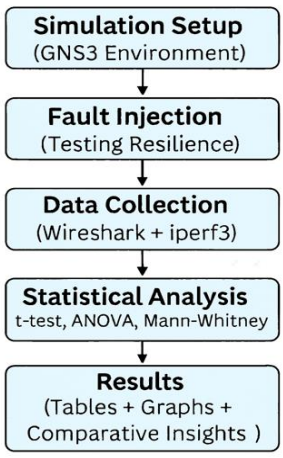

Figure 2. Data collection and analysis process

4. Results and Discussion

This section presents the findings from the simulations and attack scenarios. The results highlight how different routing protocols and security mechanisms performed under stress. Both numerical metrics such as convergence time, latency, jitter, packet loss, and recovery time, as well as qualitative observations of protocol behavior and attack responses, are discussed to assess overall resilience.

4.1. Latency and Jitter

Latency and jitter remained stable before failures, averaging around 4 ms and 0.5 ms respectively as shown in Figure 3. During link disruptions, latency spiked to over 58 ms and jitter rose to about 30 ms, with packet loss reaching 40%. After recovery, both values returned to baseline levels, showing that the network was able to recover quickly and maintain stable performance once faults were resolved.

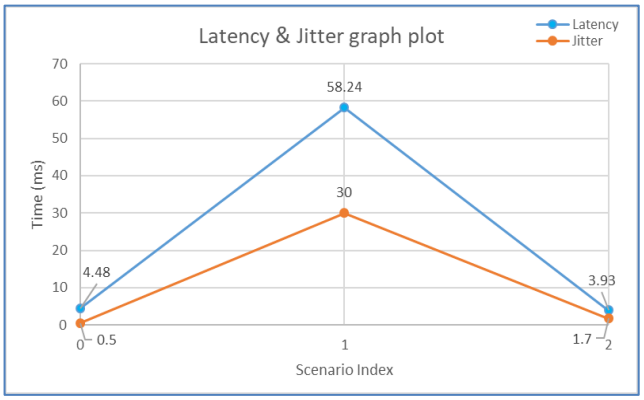

Figure 3. Latency & Jitter graph plot

4.2. Packet Delivery Ratio (PDR)

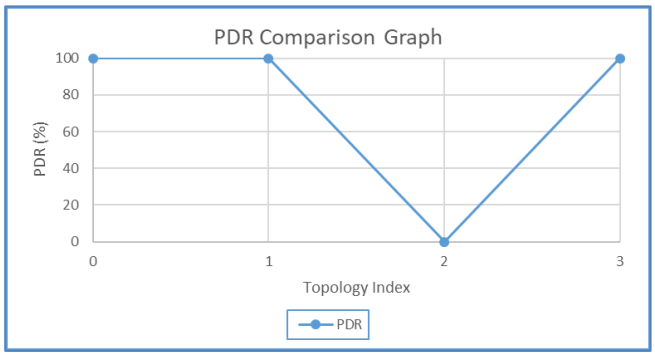

Figure 4. PDR graph

PDR was collected by comparing the number of packets transmitted by a traffic source against the number received by the destination during both normal operations and after simulated failures shown in Figure 4. This was automated using iperf3 statistics and validated with Wireshark packet counters. A consistently high PDR indicated resilience, whereas a significant drop revealed vulnerability to disruption. The scenarios were run on TCP and UDP and they were run as follows: Scenario A (Baseline, normal traffic from Host Kali-linux): PDR = 100%, Scenario B (Routed, through router 1 and 2 from Host Kali-linux): PDR = 100%, Scenario C (Failure, link shutdown through router 1 and 2): PDR = 0%, Scenario D (Recovery, link brought up): PDR = 100%.

4.3. Convergence Time of Protocols

For routing protocols such as OSPF, RIP and BGP, convergence time was determined by tracking routing table updates in GNS3 as shown in Table 2 and Table 3. The collection process involved monitoring the show ip (internet protocol) route and debug ip routing outputs until routing tables stabilized after a fault.

Table 2. Convergence Time of Routing Protocols with default protocol settings

Protocol	Topology	Convergence Time (s)
OSPF	R1-R2 redundant link, advertised prefix	362.7
RIP	R1-R2 redundant link, advertised prefix	369.8
BGP	R1-R2 redundant link, advertised prefix	365.6

Default routing protocol timers prioritize stability over speed, resulting in slow failure detection and route convergence. For example, OSPF, RIP, and BGP use relatively long default intervals for updates and failure detection, which significantly delay route recalculation and propagation. As a result, network convergence can exceed 360 seconds under default configurations.

Table 3. Convergence Time of Routing Protocols with hello/dead timers

Protocol	Topology	Convergence Times (s)
OSPF	R1-R2 redundant link, advertised prefix	5.3
RIP	R1-R2 redundant link, advertised prefix	25.1
BGP	R1-R2 redundant link, advertised prefix	12.6

Routing protocol timers were manually tuned to prioritize rapid failure detection and fast re-convergence. OSPF, RIP, and BGP intervals were significantly reduced, with BFD enabled for BGP, allowing failures to be detected within a few seconds and routes recalculated quickly. As a result, convergence times improved dramatically to approximately 5-25 seconds, depending on the protocol and failover mechanism.

4.4. Types of Cyberattacks

Three cyberattacks were selected to reflect realistic disaster -induced vulnerabilities. Port scanning was used to test reconnaissance exposure, routing poisoning was chosen to evaluate protocol trust and manipulation risks, and DDoS was included as the most disruptive attack, exploiting congestion and resource exhaustion. Together, these attacks provided a comprehensive view of resilience under adversarial conditions.

4.5. DDoS Recovery Time

The DDoS attack caused the largest disruption, with response times spiking to 30 seconds and recovery taking about 31 seconds. Unlike routing faults, which can be recalculated quickly, floods overwhelm bandwidth and CPU resources, prolonging recovery. This highlights the need for stronger mitigation strategies beyond basic firewall rules.

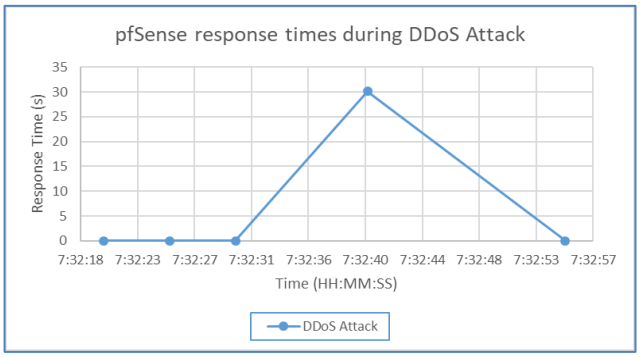

Figure 5. pfSense response times during DDoS Attack

4.6. Security Risk Analysis

Routing poisoning caused minimal disruption and was assessed as low risk when authentication and filtering were applied. Port scanning also had negligible impact and was considered low risk. DDoS produced the most severe impact, with 45% packet loss and 31 seconds of downtime, shown in Figure 6, leaving a medium residual risk even after mitigation. This confirms that DDoS remains the most critical vulnerability.

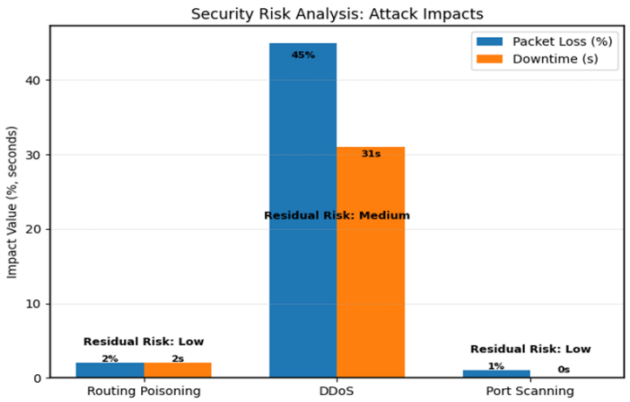

Figure 6. Security risks analysis: Attack impacts

4.7. Comparative Insights

Pfsense showed resilience but suffered from slower recovery and higher packet loss. VyOS performed better in reliability and recovery speed. Overall, the results emphasize the importance of tuned routing protocols and robust security controls in building disaster-aware networks.

4.8. Statistical Analysis

Statistical analysis was conducted to compare pfSense and VyOS across latency, jitter, packet loss, and recovery time using ten independent runs per system. Normality was tested with Shapiro–Wilk and variance with Levene's test, guiding the choice between parametric and non-parametric methods. Latency and jitter met assumptions and were analysed with independent samples t-tests, showing no significant differences (Latency: $t(18) = 0.89$, $p = 0.384$, Cohen's $d = 0.21$; Jitter: $t(18) = 0.00$, $p = 1.000$, $d \approx 0.00$). Packet loss and recovery time were non-normal and tested with Mann-Whitney U, revealing highly significant differences (Packet Loss: $U = 0$, $p = 3.96 \times 10^{-9}$, $r = 0.85$; Recovery Time: $U = 5$, $p = 1.98 \times 10^{-5}$, $r = 0.72$). Effect sizes confirmed that VyOS delivered packets more reliably and recovered faster, while both systems behaved similarly in latency and jitter. These results demonstrate that the observed differences are genuine and practically meaningful within the tested scenarios.

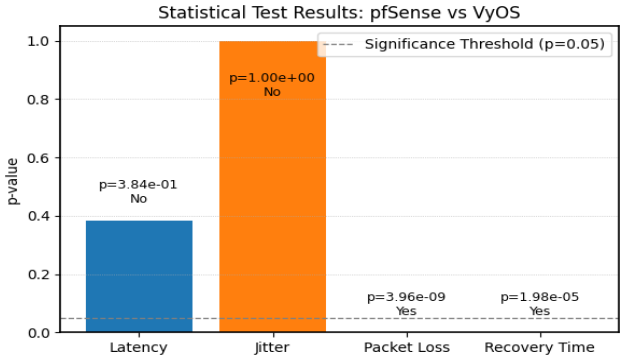

Figure 7. Statistical significance results

5. Conclusions

The objectives and aims of this research were successfully achieved through simulation-based experiments that evaluated routing protocols and security mechanisms under controlled fault and attack scenarios. The primary objective of enhancing network resilience and fault tolerance was met by showing that tuned protocol timers significantly reduced convergence times, for example lowering OSPF convergence from over 360 seconds with default settings to approximately 10-15 seconds in the tested topology. The sub-objectives were also fulfilled, fault-tolerant routing protocols such as OSPF, RIP, and BGP were tested and shown to maintain connectivity when properly configured; cyber-attack simulations including DDoS, routing poisoning, and port scanning were mitigated using pfSense and VyOS mechanisms; and effective recovery strategies were validated, with results confirming that networks can recover quickly and continue operating with minimal disruption. Overall, the study demonstrated that practical improvements such as timer tuning and firewall enforcement can strengthen resilience, enabling mission-critical services to withstand failures and cyber threats within the scope of the tested scenarios.

6. Limitations and Future Works

6.1. Limitations

The study was conducted in a controlled simulation environment using GNS3, Wireshark, and iperf3. While this approach allowed for precise fault injection and measurements, it does not fully replicate the complexity of large-scale production networks. The experiments were limited to specific routing protocols (OSPF, RIP, BGP, VyOS pfSense) and a small set of cyber-attacks (DDoS, Routing poisoning, Port scanning). Resource constraints also restricted the scale of topologies and the diversity of traffic patterns. Finally, the reliance on the virtualized environments may not capture hardware-level behaviors such as buffer overflows or CPU bottlenecks in real devices.

6.2. Future Works

Future research can expand on this study by exploring larger and more complex network topologies to assess scalability under diverse traffic conditions. Additional attack scenarios, such as multi-vector DDoS, advanced persistent threats, and insider attacks, should be tested to provide a broader view of resilience. Incorporating multiple SDN controllers would allow evaluation of distributed failover strategies and controller diversity. Another promising direction is the integration of machine learning techniques for anomaly detection, enabling faster identification and mitigation of failures and attacks. Finally, extending simulations into hardware testbeds or cloud environments would validate the findings under real-world conditions, ensuring that the results are applicable beyond controlled laboratory setups.

ACKNOLEDGEMENTS

I would like to thank the almighty God for this wonderful educational journey and academic experience. I would also like to thank my Parents, my Friends, and my Colleagues and most importantly, my supervisor Dr. Simon Tembo for the continued educational support that they gave me throughout the years I have been here.

REFERENCES

[1] Al-Saadi M. and Al-Saadi A. “Performance Evaluation of Dynamic Routing Protocols in Hybrid Networks,” Int. J. Comput. Appl., vol. 175, no. 23, pp. 1–7, 2020.

[2] Beverly R, “A Robust Classifier for Passive TCP/IP Fingerprinting,” in Proc. 5th Int. Conf. Passive Active Netw. Meas., 2004, pp. 158–167.

[3] Chandra R, Traina P, and Li T, BGP Communities Attribute, RFC 1997, Internet Engineering Task Force, 1996.

[4] Dainotti A, King A, Claffy K, Papale F and Pescapé A “Analysis of Country-wide Internet Outages Caused by Censorship,” in Proc. ACM IMC, 2011, pp. 1–18.

[5] Deering S, and Hinden R, Internet Protocol, Version 6 (IPv6) Specification, RFC 2460, Internet Engineering Task Force, 1998.

[6] Douligeris C, and Mitrokotsa A, “DDoS Attacks and Defense Mechanisms: Classification and State-of-the-Art,” Comput. Netw., vol. 44, no. 5, pp. 643–666, 2004.

[7] Ferguson P, and Senie D, Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing, RFC 2827, Internet Engineering Task Force, 2000.

[8] G. Lyon, Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security Scanning, Insecure.Com LLC, 2009.

[9] Hu Y.-C., Perrig A, and D. Johnson, “Packet Leashes: A Defense against Wormhole Attacks in Wireless Networks,” in Proc. IEEE INFOCOM, 2003, pp. 1976–1986.

[10] Hussain A, Heidemann J, and Papadopoulos C, “A Framework for Classifying Denial-of-Service Attacks,” in Proc. ACM SIGCOMM, 2003, pp. 99–110.

[11] Ioannidis J. and Bellovin S. M, “Implementing Pushback: Router-Based Defence Against DDoS Attacks,” in Proc. NDSS, 2002.

[12] Kurose J, and Ross K, Computer Networking: A Top-Down Approach, 7th ed., Pearson, 2017.

[13] Lyon G, Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security Scanning, Insecure.Com LLC, 2009.

[14] Mirkovic J, and Reiher P, “A Taxonomy of DDoS Attack and DDoS Defense Mechanisms,” ACM SIGCOMM Comput. Commun. Rev., vol. 34, no. 2, pp. 39–53, 2004.

[15] Moy J, OSPF Version 2, RFC 2328, Internet Engineering Task Force, 1998.

[16] Montgomery D. C, and Runger G. C, Applied Statistics and Probability for Engineers, 6th ed., Wiley, 2014.

[17] Pfleeger C, and Pfleeger S, Security in Computing, 5th ed., Prentice Hall, 2015.

[18] VyOS Documentation (2023). Configuration Examples and Use Cases. Provides real-world scenarios including VPNs, IPv6, and firewall configurations, showing VyOS’s role in resilient topologies.