



Research Article

Design and Implementation of a Secure Segmented Network Architecture for Nigerian SMEs Using pfSense and VLAN Technology

Oluwatosin Gold Adedeji¹ , Theophilus Aniemeka Enem² , Obunike Arinze Ubadike³ ,
Ramat Abba Imam⁴ , Samson Adeyinka⁵ and Enosegbe Daniel Lucky⁶

^{1,2,3,4,5&6}Department of Cyber Security, Faculty of Computing, Air Force Institute of Technology, Kaduna, Nigeria

Article Information

Abstract

Article History

Received: 15 July 2026

Revised: 16 August 2026

Accepted: 3 September 2026

Published online: 5 October 2026

Keywords

Pfsense

VLAN Segmentation

Network Security

SME Cybersecurity

Firewall Policy

Nigeria

Correspondence*

goldadedeji@gmail.com

ORCID

Oluwatosin Gold Adedeji

<https://orcid.org/0009-0006-7397-3489>

Theophilus Aniemeka Enem

<https://orcid.org/0000-0002-5245-8828>

Obunike Arinze Ubadike

<https://orcid.org/0009-0000-2914-6570>

Ramat Abba Imam

<https://orcid.org/0009-0004-6647-7212>

Samson Adeyinka

<https://orcid.org/0009-0005-9361-7337>

Enosegbe Daniel Lucky

<https://orcid.org/0000-0002-8986-5574>

Small and medium-sized enterprises (SMEs) in Nigeria are increasingly dependent on computer networks for communication, data storage, and day-to-day business operations. Yet, a large proportion of these organisations continue to run flat, unsegmented networks with no internal boundary between departments, guests, and critical systems. This arrangement leaves them exposed to unauthorised access, malware propagation, and lateral movement once a single device on the network is compromised. At the same time, the cost and technical demands of commercial security appliances place conventional segmentation solutions out of reach for most SMEs. This study designed and implemented a secure, segmented, multi-tenant network architecture for Nigerian SMEs built around pfSense, an open-source firewall and routing platform, and IEEE 802.1Q Virtual Local Area Network (VLAN) technology. pfSense was configured as the centralised hub through which all routing, Dynamic Host Configuration Protocol (DHCP) allocation, Network Address Translation (NAT), and firewall policy enforcement passed, while three VLANs—Administration, Finance, and Guest—represented logically isolated departmental segments that shared the same physical infrastructure. The system was built and tested in a virtualised laboratory environment using Oracle VirtualBox, with Kali Linux client machines representing departmental and guest endpoints. Evaluation covered DHCP address allocation, gateway reachability, permitted and denied inter-VLAN communication, guest isolation, and outbound internet access. The results showed that each VLAN received correctly scoped DHCP addressing, and the Administration segment was able to reach Finance resources as permitted by policy. At the same time, the reverse path was blocked; the Guest segment was completely denied access to internal departmental resources while retaining internet connectivity via Automatic Outbound NAT, and all gateway and internet-reachability tests recorded zero packet loss. These findings demonstrate that centralised firewall enforcement, combined with VLAN-based logical segmentation, can deliver enterprise-grade network isolation for Nigerian SMEs using only open-source software and standard commodity hardware, without the licensing fees or specialist staffing that conventional segmentation projects typically require.

© 2026 Centre for Research and Innovation (CRI). This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

I. INTRODUCTION

Small and medium-sized enterprises occupy a central place in Nigeria's economy, accounting for a substantial share of employment and contributing directly to innovation and service delivery across sectors [1]. As these businesses move more of their operations onto computer networks for communication, record-keeping, and customer-facing services, the security of the underlying network infrastructure ceases to be a peripheral IT concern and becomes a determinant of business continuity. Despite this shift, many Nigerian SMEs still run their networks flat, with no logical

boundary separating administrative and financial systems from the devices used by visitors or temporary staff [2]. In an environment of that kind, a single compromised laptop or phishing victim gives an attacker an open path to every other system sharing the same broadcast domain [3].

The absence of segmentation is rarely a deliberate choice. Existing literature attributes it to a combination of limited capital, a shortage of in-house security expertise, and an awareness gap regarding what segmentation actually buys an organisation in terms of risk reduction [4]. Enterprise-grade segmentation appliances, when considered at all, tend to be

priced and licensed for organisations several times the size of a typical Nigerian SME, leading many owners to conclude that proper network security is not affordable for businesses of their scale. That conclusion does not hold up, but it persists because few of the published segmentation case studies speak directly to the constraints SMEs actually face.

pfSense, an open-source firewall and routing distribution built on FreeBSD, narrows that gap considerably. It provides VLAN tagging, stateful firewall rules, NAT, and a web-based management console without licensing fees, and it has already been shown to perform competitively against proprietary firewalls in campus and enterprise deployments [5]. Pairing pfSense with VLAN segmentation under the IEEE 802.1Q standard [6] allows several logically isolated departmental or tenant networks to share a single physical switch and a single firewall, with all inter-segment traffic forced through a single, centrally administered policy point. This is precisely the kind of multi-tenant arrangement that fits how many Nigerian SMEs already operate, where departments, and sometimes entirely separate small businesses, share office space and a single internet connection.

This study set out to design and implement such an architecture specifically for the Nigerian SME context, and to test whether it delivers the isolation and control it promises. The specific objectives were: (i) to design a VLAN-based segmentation model appropriate to the scale and budget of a Nigerian SME; (ii) to implement pfSense firewall policies that enforce least-privilege access control between the resulting network segments; and (iii) to evaluate the resulting system through DHCP, NAT, and connectivity testing under both permitted and denied traffic conditions.

Beyond its immediate security contribution, the study is intended as a practical reference point for SME owners, IT consultants, and cybersecurity students in Nigeria who need a segmentation blueprint that does not assume access to enterprise budgets, dedicated network operations staff, or cloud-native tooling-constraints that distinguish SME environments from the campus and large-enterprise settings that dominate the existing segmentation literature [7].

II. LITERATURE REVIEW

A. Theoretical Framework

This study is grounded in the Defence-in-Depth model, which holds that no single control should be relied upon to protect a network, and that security should instead be layered so that the failure of one control does not automatically expose the rest of the system. The proposed architecture has a three-layer structure: a segmentation layer, realised through VLANs, that provides logical isolation between departments; a control layer, realised through the pfSense firewall, that enforces which of those isolated segments may communicate and under what conditions; and an access layer, made up of the end-user devices that originate and receive traffic.

The design also draws on Zero Trust principles, which reject the assumption long embedded in perimeter-based security that traffic originating inside the network is automatically safe [8]. Under a Zero Trust posture, every request for access between segments must be explicitly permitted; nothing is allowed by default.

The Cybersecurity and Infrastructure Security Agency's Zero Trust Maturity Model [9] and the National Institute of Standards and Technology's access-control guidance for cloud-native and multi-cloud environments [10] both formalise this default-deny posture. Recent systematic reviews of Zero Trust implementations confirm that the approach translates well to environments far smaller than the large federal and cloud deployments for which it was originally designed [11]. The firewall rule sets implemented in this study deny by default, permit only specific named flows, and follow the same logic at the SME scale.

B. Critical Review of Related Studies

1. VLAN and pfSense Integration in Campus Networks [5]: This study documented the migration of a campus network from a proprietary firewall to pfSense, with VLANs used to separate user groups and pfSense managing inter-VLAN traffic and routing.

a. Strengths: The work confirmed that pfSense can enforce access-control policy, maintain network stability, and cut the licensing costs tied to proprietary firewalls in a live, non-trivial network.

b. Weaknesses: The deployment assumed the structured IT governance and dedicated administrators typical of a campus environment, conditions that do not hold for most SMEs, and it did not address multi-tenant isolation, which is the central requirement in a shared SME setting.

2. Segmentation as a Defence Against Lateral Movement [2]: Working from controlled simulations, this study examined how logical segmentation restricts an attacker's ability to move from one compromised host to others inside an enterprise network.

a. Strengths: The simulations gave clear evidence that restricting communication paths between systems measurably narrows the attack surface available to an intruder who has already gained a foothold.

b. Weaknesses: The model was enterprise-oriented and simulation-only; it offered no deployment framework, and it said nothing about cost, which limits its direct usefulness to a resource-constrained SME.

3. Performance of VLAN Segmentation in Software-Defined Networks [3]: This study measured latency, throughput, and packet loss for VLAN and Private VLAN segmentation inside a software-defined networking (SDN) testbed under varying load.

a. Strengths: The results showed that VLAN-based isolation imposes very little performance penalty, even under heavy traffic, supporting the case that segmentation does not have to come at the cost of usability.

b. Weaknesses: SDN controllers and the specialised hardware they require sit well outside the budget and skill base of a typical Nigerian SME, so the performance gains documented here are not easily reproducible in that setting.

4. Zero Trust Integration with Network Segmentation [7]: This study combined segmentation with Zero Trust's default-deny posture in an enterprise setting, pairing a strict firewall policy with the assumption that no internal traffic is automatically trustworthy.

a. Strengths: The combined approach was shown to meaningfully reduce insider-threat and unauthorised-access risk, beyond what segmentation alone achieves.

b. Weaknesses: The implementation leaned on identity-management platforms and monitoring infrastructure that assume a security team is already in place and resources well beyond what most SMEs can field.

5. Security Survey for Next-Generation Networks [1]: This survey reviewed segmentation, micro-segmentation, and software-defined security approaches across next-generation network environments, situating segmentation as a foundational rather than supplementary control.

a. Strengths: It offers a wide-angle view of where segmentation sits relative to other emerging controls, useful as orientation for anyone new to the field.

b. Weaknesses: Being a survey, it has no experimental validation of its own, and it does not engage with the specific constraints of SMEs in developing economies such as Nigeria.

6. Micro-segmented Architecture for Multi-Cloud Environments [12]: This study proposed a micro-segmented network architecture for multi-cloud deployments, using open-source tools to enforce fine-grained, tenant-level access control across containers, IaaS, and PaaS workloads.

a. Strengths: It demonstrated that strong tenant isolation is achievable with open-source tooling alone, supporting this study's premise that licensing cost is not, by itself, a barrier to good segmentation.

b. Weaknesses: The architecture is built around cloud orchestration concepts—service meshes and container networking—that have no direct analogue in the on-premise, physical-LAN environment that most Nigerian SMEs actually operate.

7. VLAN Deployment in an Academic Network [13]: This study deployed VLAN segmentation across user groups in a

secondary-school network, evaluating its effect on broadcast traffic and access control.

a. Strengths: It confirmed that VLAN segmentation is practical to deploy in a resource-constrained, non-enterprise setting, which is closer to the SME context than most of the literature reviewed here.

b. Weaknesses: The study had no centralised firewall layer, so inter-VLAN communication was never strictly controlled, and it did not examine lateral-movement or insider-threat scenarios.

8. The Role of Segmentation in Modern Security Programs [14]: This work argued, at a conceptual level, that segmentation remains one of the most cost-effective defences available to organisations that cannot afford advanced security tooling.

a. Strengths: The argument is well aligned with the practical orientation of this study and reinforces segmentation's relevance to budget-constrained environments.

b. Weaknesses: The treatment is theoretical; it stops short of any implementation detail or evidence that would let an SME act on the recommendation directly.

C. Research Gaps and Contribution of This Study

Taken together, the studies reviewed above leave five gaps unaddressed. First, almost none of them are built around the constraints that actually define an SME—tight budgets, no dedicated security staff, and a need for configurations simple enough to maintain without specialist support. Second, VLAN segmentation and firewall-based enforcement are usually studied in isolation from one another, rather than as the integrated pair they need to be in a working deployment. Third, several of the more rigorous studies remain simulation-only or purely theoretical, with no end-to-end implementation to validate their claims. Fourth, the architectures that do offer micro-segmentation or Zero Trust integration generally assume cloud-native or SDN infrastructure that has no equivalent in a typical SME's physical network. Fifth, and most specifically, almost none of the reviewed work engages with the Nigerian SME context at all.

This study addresses each of these gaps directly: it integrates VLAN segmentation with pfSense firewall enforcement into a single working architecture; it uses only open-source software and ordinary virtualised hardware, keeping cost low; it is designed and evaluated specifically against the constraints of Nigerian SMEs rather than campuses, clouds, or large enterprises; and it validates the design through practical DHCP, NAT, and connectivity testing rather than simulation alone.

III. METHODOLOGY

A. Research Design and System Architecture

This study followed a design-based, experimental approach split into two phases. The first phase covered system design: developing the network architecture, planning the VLAN layout, and defining the firewall policy each segment would be subject to. The second phase covered implementation and evaluation: configuring pfSense and the VLANs, deploying client systems, and running the connectivity, isolation, and internet-access tests described in Section III-F.




Fig.1 Hub-and-spoke System Architecture, with pfSense as the Central Policy-Enforcement Hub and each VLAN Forming an Isolated Spoke

The resulting system follows a hub-and-spoke topology, with pfSense serving as the hub and each VLAN forming a spoke. Figure.1 illustrates the arrangement: every spoke connects only to the hub, and any traffic that needs to move—for example, from the Administration VLAN to the Finance

VLAN—must pass through pfSense and be explicitly permitted by a firewall rule before it is allowed through. No spoke can reach another spoke directly. This concentrates all of the system’s access-control logic in a single location, which keeps the configuration auditable and avoids the rule-sprawl that tends to develop when access control is spread across multiple devices.

Four components make up the architecture: the pfSense firewall itself, which performs routing, DHCP, NAT, and rule enforcement; a VLAN-capable switch (simulated in this implementation through VirtualBox’s internal networking) that carries the logically separated traffic over a shared physical link; client systems representing the different departments or tenants; and an internet gateway through which all outbound traffic exits.

In operation, each device is assigned to exactly one VLAN; devices within that VLAN communicate freely with one another, while any communication that crosses a VLAN boundary is intercepted and either permitted or dropped by pfSense according to the configured rule set.

B. VLAN Design and IP Addressing Scheme

Three VLANs were defined to represent a generic SME’s organisational structure: an Administration VLAN, a Finance VLAN, and a Guest VLAN for visitors and unmanaged devices. Each VLAN was assigned its own /24 subnet, summarised in Table I, with pfSense serving as the default gateway for each subnet’s first usable address.

TABLE I VLAN DESIGN AND IP ADDRESSING SCHEME

VLAN ID	Department / Tenant	IP Network	Gateway Address
VLAN 10	Administration	192.168.10.0/24	192.168.10.1
VLAN 20	Finance	192.168.20.0/24	192.168.20.1
VLAN 40	Guest	192.168.40.0/24	192.168.40.1

VLAN separation was implemented using IEEE 802.1Q tagging [6], which inserts a VLAN identifier into the Ethernet frame header, allowing multiple logically distinct networks to share a single physical interface. The LAN interface on the pfSense host was configured as a trunk port, carrying tagged traffic for all three VLANs simultaneously. In contrast, each client-facing port further down the topology was configured as an access port bound to a single VLAN. This is the same tagging mechanism documented in pfSense’s own VLAN configuration guidance [15], and it is what allows three logically isolated networks to run over a single physical network interface card on the pfSense host rather than requiring three separate cables and switch ports.

C. System Implementation Procedure

pfSense Community Edition was installed inside an Oracle VirtualBox virtual machine from the official ISO image, with the virtual machine allocated 2 GB of RAM and 20 GB of storage—modest enough to run comfortably on commodity

hardware. Two network interfaces were assigned to the virtual machine: a WAN interface, bridged through VirtualBox’s NAT networking to reach the internet, and a LAN interface configured as the 802.1Q trunk for the internal VLANs.

With the base interfaces in place, the three VLANs were created under Interfaces → Assignments → VLANs in the pfSense web console, tagged 10 (Administration), 20 (Finance), and 40 (Guest), and each was bound to its own logical interface and assigned the gateway address shown in Table I. DHCP was then enabled separately for each VLAN under Services → DHCP Server, following pfSense’s standard DHCPv4 configuration procedure [16], with each scope configured to hand out addresses from within its respective /24 range, sparing administrators the task of configuring IP addresses on client machines by hand. Client virtual machines running Kali Linux were attached to the relevant VLAN-tagged virtual network adapter to represent end-user devices in each segment, and DHCP allocation was

verified on each before any firewall or connectivity testing began.

D. Firewall Policy Design

The firewall rule base was built around a deny-by-default posture, consistent with the Zero Trust framing set out in Section II-A: no traffic between VLANs is permitted unless an explicit rule says otherwise, following standard firewall policy practice for segmented networks [17]. Three policies governed the segments in this study. The Administration VLAN was granted access to Finance resources and the internet, reflecting a department that legitimately needs visibility into financial systems. The Finance VLAN was permitted internet access but explicitly denied the ability to

initiate connections to the Administration VLAN, so the relationship between the two segments was one-directional rather than mutual. The Guest VLAN was denied access to internal departmental networks and permitted only outbound internet access, isolating unmanaged visitor devices from anything of business value while still allowing them to get online. Outbound NAT was configured in pfSense’s automatic mode so that all three VLANs could share the single public WAN address for internet access without exposing their private addressing externally.

E. Tools and Technologies

Table II lists the software and protocols used to build and test the system.

TABLE II TOOLS, TECHNOLOGIES, AND FRAMEWORKS USED

Category	Tool	Purpose
Firewall	pfSense Community Edition	Routing, security enforcement, DHCP, and NAT
Virtualization	Oracle VirtualBox	Hosting the firewall and client virtual machines
Client OS	Kali Linux	Departmental and guest client simulation
Networking	VLAN (IEEE 802.1Q)	Logical traffic segmentation
Interfaces	Virtual NICs	VLAN-tagged trunk and access connectivity

F. Testing and Evaluation Procedure

Five tests validated the system’s behaviour. A DHCP allocation test confirmed that clients in each VLAN automatically received a correctly scoped IP address. A gateway reachability test confirmed that each client could reach its own VLAN’s pfSense gateway with no packet loss. An inter-VLAN access test confirmed that traffic permitted by policy—Administration reaching Finance—actually succeeded. A guest isolation test confirmed that traffic explicitly denied by policy, Guest reaching either Administration or Finance, was actually blocked, with the firewall returning a destination-unreachable response rather than allowing the packet through silently. Finally, an internet connectivity test confirmed that DNS resolution and NAT translation worked correctly by pinging an external host from each segment. Isolation effectiveness was assessed by whether denied traffic was blocked, and connectivity quality was assessed using the packet-loss figures recorded for each ping test.

IV. RESULTS AND DISCUSSION

A. Implementation Environment

The system was built on a host machine running an Intel Core i9 (13th generation) processor, 64 GB of RAM, and a 2 TB SSD, with virtualisation extensions enabled in the firmware (Table III). The choice of a virtualised testbed, summarised alongside the software stack in Table IV, kept the implementation reproducible and inexpensive, in line with the cost constraints that motivate this study.

TABLE III HARDWARE ENVIRONMENT

Component	Specification
Processor	Intel Core i9, 13th Generation
RAM	64 GB
Storage	2 TB SSD
Virtualization	Enabled

TABLE IV SOFTWARE ENVIRONMENT

Software	Purpose
Oracle VirtualBox	Virtual machine platform
pfSense CE 2.8.1	Firewall and router
Kali Linux	Client workstation testing
Web browser	pfSense management console
Ping / CLI tools	Connectivity testing

B. pfSense Deployment and VLAN Configuration

pfSense booted successfully into its console interface with both the WAN and LAN interfaces correctly initialised, as shown in Figure.2, confirming that the firewall was ready to take on its role as the network’s central routing and policy-enforcement point.

The web-based dashboard in Figure.3 subsequently confirmed that all configured interfaces remained active, with CPU and memory utilisation within normal bounds, thereby providing a working management surface for the subsequent configuration steps.




Fig.2 pfSense Console After a Successful Boot, Showing Correctly Initialized WAN and LAN Interfaces




Fig.3 pfSense Web-based Dashboard Confirming Active Interfaces and Normal Resource Utilisation

The three VLANs were created and tagged according to the IEEE 802.1Q standard, as shown in [Figure.4](#), and each was then bound to its own logical interface ([Figure.5](#)), reproducing the structure already specified in [Table I](#).




Fig.4 Creation of VLANs 10, 20, and 40 in pfSense Under the IEEE 802.1Q Tagging Scheme




Fig.5 Assignment of Each Tagged VLAN to its own Logical Interface in Pfsense

C. DHCP Configuration

DHCP scopes were configured independently for each VLAN so that connected clients would automatically receive an address, subnet mask, default gateway, and DNS configuration, rather than requiring manual entry on every machine. Figure.6 shows the scope configured for the

Administration VLAN, issuing addresses in the range 192.168.10.10–192.168.10.100, and Figure.7 shows the corresponding configuration for the Guest VLAN, issuing addresses from a separate range, 192.168.40.10–192.168.40.100, with no overlap with any internal department’s addressing.

Services / DHCP Server / VLAN10_ADMIN

The changes have been applied successfully.

ISC DHCP has reached end-of-life and will be removed in a future version of pfSense. Visit [System > Advanced > Networking](#) to switch DHCP backend.

LAN VLAN10_ADMIN VLAN20_FINANCE VLAN30_IT VLAN40_GUEST

General Settings

DHCP Backend

ISC DHCP

Enable

☒ Enable DHCP server on VLAN10_ADMIN interface

BOOTP

☐ Ignore BOOTP queries

Deny Unknown Clients

Allow all clients

When set to **Allow all clients**, any DHCP client will get an IP address within this scope/range on this interface. If set to **Allow known clients from any interface**, any DHCP client with a MAC address listed in a static mapping on *any* scope(s)/interface(s) will get an IP address. If set to **Allow known clients from only this interface**, only MAC addresses listed in static mappings on this interface will get an IP address within this scope/range.

Ignore Denied Clients

☐ Ignore denied clients rather than reject
This option is not compatible with failover and cannot be enabled when a Failover Peer IP address is configured.

Ignore Client Identifiers

☐ Do not record a unique identifier (UID) in client lease data if present in the client DHCP request
This option may be useful when a client can dual boot using different client identifiers but the same hardware (MAC) address. Note that the resulting server behavior violates the official DHCP specification.

Primary Address Pool

Subnet

192.168.10.0/24

Subnet Range

192.168.10.1 - 192.168.10.254

Address Pool Range

192.168.10.10

From

192.168.10.100

To
The specified range for this pool must not be within the range configured on any other address pool for this interface.

Additional Pools

+ Add Address Pool

If additional pools of addresses are needed inside of this subnet outside the above range, they may be specified here.

Server Options

Fig.6 DHCP Scope Configuration for the Administration VLAN (VLAN10_ADMIN)

Services / DHCP Server / VLAN40_GUEST

ISC DHCP has reached end-of-life and will be removed in a future version of pfSense. Visit [System > Advanced > Networking](#) to switch DHCP backend.

LAN VLAN10_ADMIN VLAN20_FINANCE VLAN30_IT VLAN40_GUEST

General Settings

DHCP Backend

ISC DHCP

Enable

☐ Enable DHCP server on VLAN40_GUEST interface

BOOTP

☐ Ignore BOOTP queries

Deny Unknown Clients

Allow all clients

When set to **Allow all clients**, any DHCP client will get an IP address within this scope/range on this interface. If set to **Allow known clients from any interface**, any DHCP client with a MAC address listed in a static mapping on *any* scope(s)/interface(s) will get an IP address. If set to **Allow known clients from only this interface**, only MAC addresses listed in static mappings on this interface will get an IP address within this scope/range.

Ignore Denied Clients

☐ Ignore denied clients rather than reject
This option is not compatible with failover and cannot be enabled when a Failover Peer IP address is configured.

Ignore Client Identifiers

☐ Do not record a unique identifier (UID) in client lease data if present in the client DHCP request
This option may be useful when a client can dual boot using different client identifiers but the same hardware (MAC) address. Note that the resulting server behavior violates the official DHCP specification.

Primary Address Pool

Subnet

192.168.40.0/24

Subnet Range

192.168.40.1 - 192.168.40.254

Address Pool Range

192.168.40.10

From

192.168.40.100

To
The specified range for this pool must not be within the range configured on any other address pool for this interface.

Fig.7 DHCP Scope Configuration for the Guest VLAN (VLAN40_GUEST)

D. Firewall Rule Implementation

Rules were then written for each VLAN interface to implement the policy set out in Section III-D. Fig.8 shows the rules applied to the Administration VLAN, permitting traffic toward the Finance network and toward the Internet. Figure.9 shows the corresponding rules for the Finance VLAN:

Internet access is permitted, but no rule allows Finance-originated traffic toward the Administration network, so any such attempt falls through to the implicit deny at the bottom of the rule set. Figure.10 shows the Guest VLAN rules, which permit only internet-bound traffic and nothing directed at either internal department.

Firewall / Rules / VLAN10_ADMIN

The changes have been applied successfully. The firewall rules are now reloading in the background. Monitor the filter reload progress.

Floating WAN LAN **VLAN10_ADMIN** VLAN20_FINANCE VLAN30_IT VLAN40_GUEST

Rules (Drag to Change Order)

States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
✓ 0/0 B	IPv4 *	VLAN10_ADMIN address	*	VLAN20_FINANCE address	*	*	none		Allow Admin to Finance	[Edit] [Copy] [Delete] [Toggle]
✓ 0/0 B	IPv4 TCP	VLAN10_ADMIN address	*	*	*	*	none		Allow Admin to Internet	[Edit] [Copy] [Delete] [Toggle]

[Add] [Add] [Delete] [Toggle] [Copy] [Save] [Separator]

Fig.8 Firewall Rules Applied to the Administration VLAN (VLAN10_ADMIN), Permitting Access to Finance and the Internet

Firewall / Rules / VLAN20_FINANCE

The changes have been applied successfully. The firewall rules are now reloading in the background. Monitor the filter reload progress.

Floating WAN LAN VLAN10_ADMIN **VLAN20_FINANCE** VLAN30_IT VLAN40_GUEST

Rules (Drag to Change Order)

States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
✗ 0/0 B	IPv4 *	VLAN20_FINANCE address	*	VLAN10_ADMIN address	*	*	none		Block Finance to Admin	[Edit] [Copy] [Delete] [Toggle]
✓ 0/0 B	IPv4 TCP	VLAN20_FINANCE address	*	*	*	*	none		Allow Finance to Internet	[Edit] [Copy] [Delete] [Toggle]

[Add] [Add] [Delete] [Toggle] [Copy] [Save] [Separator]

Fig.9 Firewall Rules Applied to the Finance VLAN (VLAN20_FINANCE), Blocking Unsolicited Access to Administration

Firewall / Rules / VLAN40_GUEST

The firewall rule configuration has been changed. The changes must be applied for them to take effect. [Apply Changes]

Floating WAN LAN VLAN10_ADMIN VLAN20_FINANCE VLAN30_IT **VLAN40_GUEST**

Rules (Drag to Change Order)

States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
✗ 0/0 B	IPv4 *	VLAN40_GUEST address	*	VLAN10_ADMIN address	*	*	none		Block Guest to Admin	[Edit] [Copy] [Delete] [Toggle]
✗ 0/0 B	IPv4 *	VLAN40_GUEST address	*	VLAN20_FINANCE address	*	*	none		Block Guest to Finance	[Edit] [Copy] [Delete] [Toggle]
✓ 0/0 B	IPv4 TCP	VLAN40_GUEST address	*	*	*	*	none		Allow Guest to Internet	[Edit] [Copy] [Delete] [Toggle]

[Add] [Add] [Delete] [Toggle] [Copy] [Save] [Separator]

Fig.10 Firewall Rules Applied to the Guest VLAN (VLAN40_GUEST), Restricting Access to Internet-Bound Traffic Only

E. NAT Configuration for Internet Access

Outbound NAT was left in pfSense's automatic mode, shown in [Figure.11](#), under which the firewall translates the private addressing used by all three VLANs into the single public

address presented on the WAN interface. This allowed every VLAN to reach the internet through a single WAN connection without exposing any internal addressing to the outside world.




Fig.11 Automatic Outbound NAT Configuration in pfSense, Translating Private VLAN Addressing to the Public WAN Address

F. System Testing and Validation

Each of the five tests described in Section III-F was conducted on the completed configuration, and the results are summarised in [Table V](#).

The Administration client received a correctly scoped DHCP lease of 192.168.10.101 ([Figure.12](#)), confirming that DHCP

allocation was functioning as configured. A subsequent ping from that client to its VLAN gateway returned 22/22 packets with zero loss ([Figure.13](#)), confirming stable Layer 3 connectivity between the client and pfSense. Administration was then able to reach the Finance VLAN's gateway successfully ([Figure.14](#)), confirming that the permitted inter-VLAN path specified in the firewall policy was, in fact, open.

TABLE V SUMMARY OF SYSTEM TEST OUTCOMES

Test	Expected Result	Observed Result	Status
DHCP address allocation	Each VLAN client receives an address within its configured scope	The administration client received 192.168.10.101 from the VLAN10 scope	Passed
Gateway reachability	Client reaches its own VLAN gateway with no packet loss	22/22 packets received, 0% loss	Passed
Permitted inter-VLAN access	Administration reaches the Finance gateway as permitted by policy	Ping to the Finance gateway succeeded	Passed
Guest isolation Administration	The guest is denied access to the Administration VLAN	Destination host unreachable	Passed
Guest isolation Finance	The guest is denied access to the Finance VLAN	Destination host unreachable	Passed
Internet connectivity	Client resolves DNS and reaches an external host via NAT	Successful ping with reply received	-

By contrast, the Guest client received a destination-unreachable response when attempting to reach the Administration gateway ([Figure.15](#)) and again when attempting to reach the Finance gateway ([Figure.16](#)), confirming that both denied paths were correctly blocked

rather than merely unrouted. Finally, a client on the network successfully resolved and pinged an external domain ([Figure.17](#)), confirming that DNS resolution, NAT translation, and WAN routing were all functioning correctly together.

```

kali@kali: ~
Session Actions Edit View Help
(kali@kali)-[~]
$ ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group def
    ault qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host noprefixroute
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP g
    roup default qlen 1000
    link/ether 08:00:27:3d:bf:ac brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.101/24 brd 192.168.1.255 scope global dynamic noprefixrout
        e eth0
        valid_lft 6850sec preferred_lft 6850sec
    inet6 fe80::ead:a36b:35cd:2990/64 scope link noprefixroute
        valid_lft forever preferred_lft forever
    
```

Fig.12 DHCP-assigned IP Configuration on the Administration Client (192.168.10.101)

```

kali@kali: ~
Session Actions Edit View Help
(kali@kali)-[~]
$ ping 192.168.10.1
PING 192.168.10.1 (192.168.10.1) 56(84) bytes of data.
64 bytes from 192.168.10.1: icmp_seq=1 ttl=64 time=1.64 ms
64 bytes from 192.168.10.1: icmp_seq=2 ttl=64 time=6.90 ms
64 bytes from 192.168.10.1: icmp_seq=3 ttl=64 time=5.59 ms
64 bytes from 192.168.10.1: icmp_seq=4 ttl=64 time=3.45 ms
64 bytes from 192.168.10.1: icmp_seq=5 ttl=64 time=5.76 ms
64 bytes from 192.168.10.1: icmp_seq=6 ttl=64 time=5.10 ms
64 bytes from 192.168.10.1: icmp_seq=7 ttl=64 time=8.11 ms
64 bytes from 192.168.10.1: icmp_seq=8 ttl=64 time=5.87 ms
64 bytes from 192.168.10.1: icmp_seq=9 ttl=64 time=5.33 ms
64 bytes from 192.168.10.1: icmp_seq=10 ttl=64 time=4.84 ms
64 bytes from 192.168.10.1: icmp_seq=11 ttl=64 time=3.30 ms
64 bytes from 192.168.10.1: icmp_seq=12 ttl=64 time=4.50 ms
64 bytes from 192.168.10.1: icmp_seq=13 ttl=64 time=5.08 ms
64 bytes from 192.168.10.1: icmp_seq=14 ttl=64 time=4.91 ms
64 bytes from 192.168.10.1: icmp_seq=15 ttl=64 time=4.00 ms
64 bytes from 192.168.10.1: icmp_seq=16 ttl=64 time=4.61 ms
64 bytes from 192.168.10.1: icmp_seq=17 ttl=64 time=4.41 ms
64 bytes from 192.168.10.1: icmp_seq=18 ttl=64 time=5.30 ms
64 bytes from 192.168.10.1: icmp_seq=19 ttl=64 time=6.25 ms
64 bytes from 192.168.10.1: icmp_seq=20 ttl=64 time=6.93 ms
64 bytes from 192.168.10.1: icmp_seq=21 ttl=64 time=5.53 ms
64 bytes from 192.168.10.1: icmp_seq=22 ttl=64 time=5.35 ms
^C
  192.168.10.1 ping statistics:
  22 packets transmitted, 22 received, 0% packet loss, time 21052ms
  rtt min/avg/max/mdev = 1.635/5.125/8.114/1.330 ms
    
```

Fig.13 Administration Client-to-pfSense Gateway Connectivity Test: 22 of 22 Packets Received, 0% Loss

```

kali@kali: ~
Session Actions Edit View Help
(kali@kali)-[~]
$ ping 192.168.20.1
PING 192.168.20.1 (192.168.20.1) 56(84) bytes of data.
64 bytes from 192.168.20.1: icmp_seq=1 ttl=64 time=1.77 ms
64 bytes from 192.168.20.1: icmp_seq=2 ttl=64 time=3.13 ms
64 bytes from 192.168.20.1: icmp_seq=3 ttl=64 time=4.65 ms
64 bytes from 192.168.20.1: icmp_seq=4 ttl=64 time=6.32 ms
64 bytes from 192.168.20.1: icmp_seq=5 ttl=64 time=3.26 ms
64 bytes from 192.168.20.1: icmp_seq=6 ttl=64 time=4.10 ms
64 bytes from 192.168.20.1: icmp_seq=7 ttl=64 time=5.04 ms
64 bytes from 192.168.20.1: icmp_seq=8 ttl=64 time=1.85 ms
64 bytes from 192.168.20.1: icmp_seq=9 ttl=64 time=3.75 ms
64 bytes from 192.168.20.1: icmp_seq=10 ttl=64 time=9.64 ms
64 bytes from 192.168.20.1: icmp_seq=11 ttl=64 time=6.85 ms
64 bytes from 192.168.20.1: icmp_seq=12 ttl=64 time=3.32 ms
64 bytes from 192.168.20.1: icmp_seq=13 ttl=64 time=4.86 ms
64 bytes from 192.168.20.1: icmp_seq=14 ttl=64 time=7.49 ms
64 bytes from 192.168.20.1: icmp_seq=15 ttl=64 time=7.06 ms
64 bytes from 192.168.20.1: icmp_seq=16 ttl=64 time=7.95 ms
64 bytes from 192.168.20.1: icmp_seq=17 ttl=64 time=15.4 ms
64 bytes from 192.168.20.1: icmp_seq=18 ttl=64 time=5.07 ms
64 bytes from 192.168.20.1: icmp_seq=19 ttl=64 time=3.10 ms
64 bytes from 192.168.20.1: icmp_seq=20 ttl=64 time=5.87 ms
64 bytes from 192.168.20.1: icmp_seq=21 ttl=64 time=6.95 ms
^C
  192.168.20.1 ping statistics:
  21 packets transmitted, 21 received, 0% packet loss, time 20039ms
  rtt min/avg/max/mdev = 1.765/5.592/15.433/2.989 ms
    
```

Fig.14 Controlled Inter-VLAN Access Test: Administration Successfully Reached the Permitted Finance Gateway

```
kali@kali: ~  
Session Actions Edit View Help  
(kali@kali)-[~]  
$ ping 192.168.10.1  
PING 192.168.10.1 (192.168.10.1) 56(84) bytes of data.  
From 169.254.130.146 icmp_seq=1 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=2 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=3 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=4 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=5 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=6 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=7 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=8 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=9 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=10 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=11 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=12 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=13 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=14 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=15 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=16 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=17 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=18 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=19 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=20 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=21 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=22 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=23 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=24 Destination Host Unreachable  
^C  
--- 192.168.10.1 ping statistics ---  
25 packets transmitted, 0 received, +24 errors, 100% packet loss, time 24564ms  
s  
pipe 3
```

Fig.15 Guest Isolation Test: Guest Client Unable to Reach the Administration VLAN Gateway

```
kali@kali: ~  
Session Actions Edit View Help  
(kali@kali)-[~]  
$ ping 192.168.20.1  
PING 192.168.20.1 (192.168.20.1) 56(84) bytes of data.  
From 169.254.130.146 icmp_seq=1 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=2 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=3 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=4 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=5 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=6 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=7 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=8 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=9 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=10 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=11 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=12 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=13 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=14 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=15 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=16 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=17 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=18 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=19 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=20 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=21 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=22 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=23 Destination Host Unreachable  
From 169.254.130.146 icmp_seq=24 Destination Host Unreachable  
^C  
--- 192.168.20.1 ping statistics ---  
25 packets transmitted, 0 received, +24 errors, 100% packet loss, time 24562ms  
s  
pipe 3
```

Fig.16 Guest Isolation Test: Guest Client Unable to Reach the Finance VLAN Gateway

```
kali@kali: ~  
Session Actions Edit View Help  
(kali@kali)-[~]  
$ ping google.com  
PING google.com (142.251.209.110) 56(84) bytes of data.  
64 bytes from 142.251.209.110: icmp_seq=11 ttl=254 time=329 ms  
64 bytes from 142.251.209.110: icmp_seq=12 ttl=254 time=333 ms  
64 bytes from 142.251.209.110: icmp_seq=13 ttl=254 time=353 ms  
64 bytes from 142.251.209.110: icmp_seq=16 ttl=254 time=373 ms  
64 bytes from 142.251.209.110: icmp_seq=18 ttl=254 time=299 ms  
64 bytes from 142.251.209.110: icmp_seq=19 ttl=254 time=317 ms  
64 bytes from 142.251.209.110: icmp_seq=23 ttl=254 time=334 ms  
64 bytes from 142.251.209.110: icmp_seq=25 ttl=254 time=260 ms  
64 bytes from 142.251.209.110: icmp_seq=27 ttl=254 time=397 ms  
64 bytes from 142.251.209.110: icmp_seq=29 ttl=254 time=327 ms  
64 bytes from 142.251.209.110: icmp_seq=30 ttl=254 time=277 ms  
64 bytes from 142.251.209.110: icmp_seq=31 ttl=254 time=374 ms  
64 bytes from 142.251.209.110: icmp_seq=36 ttl=254 time=403 ms  
64 bytes from 142.251.209.110: icmp_seq=38 ttl=254 time=330 ms  
64 bytes from 142.251.209.110: icmp_seq=41 ttl=254 time=367 ms  
64 bytes from 142.251.209.110: icmp_seq=42 ttl=254 time=282 ms  
64 bytes from 142.251.209.110: icmp_seq=49 ttl=254 time=328 ms  
64 bytes from 142.251.209.110: icmp_seq=52 ttl=254 time=364 ms  
^C  
--- google.com ping statistics ---  
54 packets transmitted, 18 received, 66.6667% packet loss, time 63108ms  
rtt min/avg/max/mdev = 260.196/335.997/403.189/38.857 ms
```

Fig.17 Internet Connectivity Test: Successful DNS Resolution and Ping Reply from an External Host

G. Discussion

Across all five tests, the architecture behaved exactly as its design specified: paths the policy permitted worked, and paths the policy denied did not, with no exceptions observed in either direction. That consistency matters more than any single test result on its own, because it is the property a segmentation project is actually trying to deliver: predictable, policy-driven behaviour rather than incidental isolation that happens to hold only until traffic patterns change.

The zero-packet-loss results recorded in both the gateway reachability and internet connectivity tests indicate that VLAN tagging and NAT translation introduced no measurable overhead in this testbed, consistent with the low-overhead findings reported for VLAN segmentation in SDN environments [3], even though the underlying infrastructure here is far simpler than that of an SDN deployment. The destination-unreachable responses returned to the Guest client on both blocked paths are a meaningful result in their own right: a firewall that silently dropped packets without returning any response would still achieve isolation, but the explicit unreachable response confirms that pfSense's stateful rule evaluation was actively inspecting and rejecting the traffic, rather than the traffic simply having nowhere to go.

The asymmetric relationship enforced between Administration and Finance—one direction permitted, the other denied—illustrates a point that is easy to lose in a binary view of segmentation as “isolated” versus “not isolated.” Least-privilege access control is rarely all-or-nothing in a real organisation; departments need defined, limited visibility into each other's systems, not a wall between every pair of segments. pfSense's rule-based model handled that asymmetry without any additional tooling beyond what was already needed for the Guest VLAN's full isolation, which is encouraging from a cost and complexity standpoint: an SME does not need a more sophisticated platform to express more sophisticated policy, only a more carefully written rule set on the platform it already has.

Set against the literature reviewed in Section II, this implementation occupies a position that the existing work does not quite cover. It aligns with the demonstration by Coquis-Flames *et al.*, [5] that pfSense can enforce access control and maintain network stability. Still, it does so in a configuration sized and budgeted for an SME rather than a campus network with dedicated administrators. It reaches an outcome similar to Darshane's [7] Zero Trust-segmentation pairing—default-deny, explicitly permitted exceptions—without requiring the identity-management and monitoring infrastructure that the pairing assumed. And it supplies the kind of practical, working-deployment evidence that Mbona and Eloff's [2] simulation-only study and Sunkara's [14] conceptual argument both stop short of providing.

The most direct practical implication is that the licensing cost, long assumed to be the barrier to SME network

segmentation, is not, in fact, the binding constraint. Every component used in this study—pfSense, VirtualBox, and Kali Linux—is free and open-source, and the hardware requirements (2 GB of RAM and 20 GB of storage for the firewall virtual machine) are well within reach of a low-cost dedicated PC or a repurposed desktop. What an SME actually needs to make use of an architecture like this is not capital, but a clear policy decision about which departments may see which other departments' resources, and someone capable of translating that decision into the kind of rule set demonstrated in Section IV-D.

V. CONCLUSION AND RECOMMENDATIONS

This study designed, implemented, and tested a secure, segmented, multi-tenant network architecture for Nigerian SMEs built on pfSense and IEEE 802.1Q VLAN technology. Testing confirmed that DHCP correctly allocated scoped addresses to each segment, that the Administration-to-Finance path permitted by policy worked while the reverse path was blocked, that the Guest segment was fully isolated from internal resources while keeping internet access, and that NAT and DNS functioned correctly for outbound connectivity across all three segments, with zero packet loss recorded on every reachability test.

These results support a straightforward conclusion: VLAN segmentation enforced through a centralised, open-source firewall enables Nigerian SMEs to materially reduce their exposure to lateral movement and unauthorised internal access without the licensing costs, specialist staffing, or cloud-native infrastructure that most existing segmentation literature assumes are necessary. The architecture demonstrated here is not a scaled-down version of an enterprise solution; it is a complete, working implementation sized correctly for the environment in which it was built.

A few limitations qualify these findings. The system was evaluated in a virtualised laboratory environment rather than on a live production network, so behaviour under real, fluctuating user load and genuine attack traffic was not observed. The host machine's resources limited the number of client systems that could run simultaneously, which constrained the scale of testing. Advanced controls such as intrusion detection and prevention, centralised log management, and VPN-based remote access fell outside the scope of this study, and the testing relied on synthetic rather than production traffic. None of these limitations weakens the validity of the core architecture or the reported test outcomes, but they do indicate where the next stage of work should focus.

Building on that, four directions are recommended for future work. First, the architecture should be deployed in an actual SME, with its performance and policy effectiveness tracked over an extended period under real traffic rather than test traffic. Second, an intrusion detection and prevention capability—Suricata or Snort, both of which integrate natively with pfSense—should be added to give the

architecture visibility into traffic content, not just traffic flow. Third, VPN connectivity should be layered onto the WAN side of the design to support branch offices and remote staff without weakening the internal segmentation already achieved. Fourth, future evaluations should benchmark pfSense's performance and policy flexibility directly against low-cost commercial alternatives, to give SME decision-makers a clearer basis for choosing between open-source and proprietary options.

More broadly, this project reinforces a point that is easy to state and easy to underestimate in practice: meaningful network security for a small organisation is far more often a matter of disciplined design and configuration than of budget. The Nigerian SME sector does not need to wait for enterprise-grade tooling to become affordable; the tools to build a properly segmented, centrally enforced network already exist, free of charge, and the remaining work is mostly the careful policy decisions described in this study.

ACKNOWLEDGEMENT

The author gratefully acknowledges the Department of Cyber Security, Faculty of Computing, Air Force Institute of Technology (AFIT), Kaduna, for providing the institutional and technical environment within which this research was carried out, and the project supervisor and co-supervisor for their guidance throughout.

Ethical Statement

This research was conducted entirely within a virtualized laboratory environment using simulated network traffic and synthetic client systems. No production network, third-party organization, or personally identifiable data was involved at any stage. Ethical clearance was obtained from the Department of Cyber Security at the Air Force Institute of Technology, Kaduna, in accordance with institutional guidelines governing undergraduate final-year projects.

Declaration of Conflicting Interests

The authors declare no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Use of Artificial Intelligence (AI)-Assisted Technology for Manuscript Preparation

The author(s) utilised generative AI technologies solely to improve language clarity and refine grammar. No images were generated, manipulated, or altered using AI-assisted technologies.

REFERENCES

- [1] W. Owoko, "Security schemes for the next-generation networks: A survey," *GSC Adv. Res. Rev.*, vol. 21, no. 2, pp. 427–455, 2024, doi: 10.30574/gscarr.2024.21.2.0446.
- [2] T. Mbona and J. Eloff, "Semi-supervised approaches for mitigating internal network threats using segmentation techniques," *Int. J. Cybersecurity Res.*, vol. 11, no. 4, pp. 303–317, 2022.
- [3] D. Álvarez, P. Nuño, C. T. González, F. G. Bulnes, J. C. Granda, and D. García-Carrillo, "Performance analysis of software-defined networks to mitigate private VLAN attacks," *Sensors*, vol. 23, no. 4, p. 1747, 2023, doi: 10.3390/s23041747.
- [4] M. Awan and A. Alam, "Cybersecurity threats and defensive strategies for small and medium firms: A systematic mapping study," *Administrative Sciences*, vol. 15, no. 12, p. 481, 2025, doi: 10.3390/admsci15120481.
- [5] J. Coquis-Flames, H. Flor-Cunza, and A. Alva-Mantari, "Design of a pfSense-based VLAN architecture for transitioning from proprietary to open-source firewalls," *Int. J. Eng. Trends Technol.*, vol. 72, no. 8, pp. 62–72, 2024, doi: 10.14445/22315381/IJETT-V72I8P108.
- [6] IEEE, *IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks*, IEEE Std. 802.1Q-2018, 2018.
- [7] A. Darshane, "Advanced network security concepts: Network segmentation and Zero Trust Architecture," *Int. J. Eng. Technol. Res.*, vol. 9, no. 2, pp. 379–386, 2024.
- [8] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, *Zero Trust Architecture*, NIST Special Publication 800-207, Nat. Inst. Standards Technol., Gaithersburg, MD, USA, 2020, doi: 10.6028/NIST.SP.800-207.
- [9] Cybersecurity and Infrastructure Security Agency, *Zero Trust Maturity Model*, version 2.0, CISA, 2023.
- [10] R. Chandramouli, Z. Butcher, Y. Li, and M. Nandiraju, *A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Cloud Environments*, NIST Special Publication 800-207A, Nat. Inst. Standards Technol., Gaithersburg, MD, USA, 2023, doi: 10.6028/NIST.SP.800-207A.
- [11] S. Mushtaq, M. Mohsin, and M. M. Mushtaq, "A systematic literature review on the implementation and challenges of zero trust architecture across domains," *Sensors*, vol. 25, no. 19, p. 6118, 2025, doi: 10.3390/s25196118.
- [12] S. Arora and J. Hastings, "Microsegmented cloud network architecture using open-source tools for a zero trust foundation," in *Proc. 2024 17th Int. Conf. Security of Information and Networks (SIN)*, 2024, pp. 1–8, doi: 10.1109/SIN63213.2024.10871361.
- [13] P. J. A. Jasis, R. J. N. Labis, and E. P. Ruaya, "Implementing VLAN segmentation for enhanced network security and performance in Barobo National High School," *Int. J. Adv. Trends Comput. Sci. Eng.*, vol. 14, no. 5, pp. 209–216, 2025.
- [14] G. Sunkara, "The importance of network segmentation in security," *SAMRIDDHI: A J. Phys. Sci., Eng. Technol.*, vol. 17, 2025.
- [15] Netgate, "VLAN configuration," *pfSense Documentation*, 2025. [Online]. Available: <https://docs.netgate.com/pfsense/en/latest/vlan/configuration.html>
- [16] Netgate, "DHCPv4 server," *pfSense Documentation*, 2025. [Online]. Available: <https://docs.netgate.com/pfsense/en/latest/services/dhcp/ipv4.html>
- [17] K. Scarfone and P. Hoffman, *Guidelines on Firewalls and Firewall Policy*, NIST Special Publication 800-41 Rev. 1, Nat. Inst. Standards Technol., Gaithersburg, MD, USA, 2009, doi: 10.6028/NIST.SP.800-41r1.