Published by Nigerian Society of Physical Sciences. Hosted by FLAYOO Publishing House LTD

Proceedings of the Nigerian Society of Physical Sciences

Journal Homepage: https://flayoophl.com/journals/index.php/pnspsc

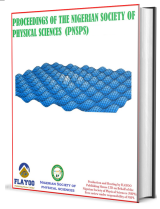

A stratified ransomware mitigation model based on zero trust and network segmentation architectures

Justine Utsu Undiandeye[a,*], Moses Adah Agana[a], Bassey Igbo Ele[b]

[a]*Department of Cybersecurity, University of Calabar, Calabar, PMB 1115, Nigeria*
[b]*Department of Information Systems, University of Calabar, Calabar, PMB 1115, Nigeria*

ABSTRACT

Ransomware poses a significant threat to information technology because of its ability to spread laterally across computer networks. This paper presents the design and implementation of a stratified mitigation model that combines Zero Trust Architecture (ZTA) with network segmentation to impede ransomware propagation. The proposed model integrates continuous verification through ZTA with the structural containment provided by network segmentation. It was implemented using pfSense, VMware, and GNS3, and evaluated using actual flow patterns extracted from a Ryuk ransomware packet-capture (PCAP) dataset. The model demonstrated automated containment based on real ransomware activity patterns, including distinctive Server Message Block (SMB) traffic profiles and rapid byte-transfer rates. Detection and containment were achieved within sub-second timescales, with a time-to-detect (TTD) of 0.31 s and a time-to-contain (TTC) of 0.32 s. These results outperform standalone ZTA (TTD: 1.50 s; TTC: 2.50 s) and standalone network segmentation (TTD: 0.65 s; TTC: 0.65 s). Across 20 controlled simulation runs, the model achieved a detection accuracy of 85.0%, precision of 81.8%, recall of 90.0%, an F1-score of 85.7%, and a false positive rate of 10%. The results show that the hybrid approach offers a pragmatic and measurable improvement over individual strategies for securing networks against ransomware.

Keywords: Ransomware mitigation, Zero trust architecture, Network segmentation, Threat detection, Hybrid security model.

DOI:10.61298/pnspsc.2026.3.271

1. INTRODUCTION

Advanced delivery mechanisms, data encryption, and double-extortion tactics have made ransomware one of the most destructive forms of cybercrime [1]. Sophisticated techniques such as phishing, credential compromise, and internal protocol abuse can bypass perimeter-based defences, rendering them increasingly ineffective. This trend has encouraged the adoption of continuous-verification security models that avoid implicit trust within internal networks.

Zero Trust Architecture (ZTA), introduced by Forrester Research [2] and formalised by NIST [3], is based on the principle of "never trust, always verify". It enforces strict access controls, continuous authentication, and micro-segmentation to minimise attack surfaces and limit threat propagation. Network segmentation similarly separates critical network assets across departmental VLANs or subnets, thereby limiting ransomware propagation when an endpoint is compromised [4]. However, standalone ZTA may lack strong network-layer confinement, while

*Corresponding Author Tel. No.: +234-803-0741-817.
e-mail: justinutsu25@gmail.com (Justine Utsu Undiandeye)

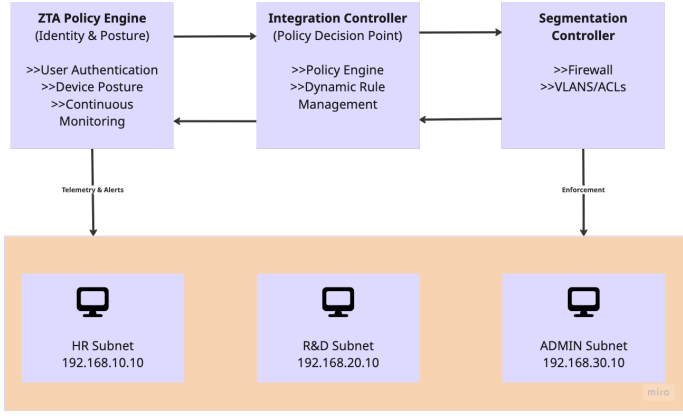

Figure 1. Hybrid model architecture.

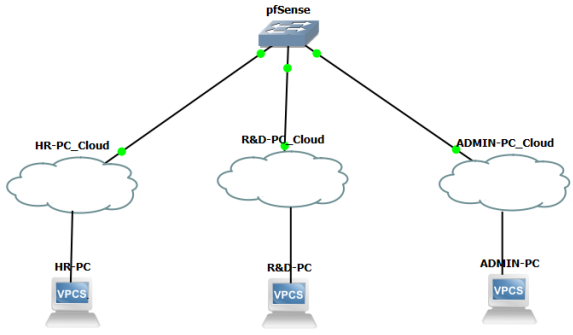

Figure 2. Network topology of the proof of concept.

static segmentation is not sufficiently responsive to behaviourally adaptive threats.

To address these limitations, this research proposes, develops, and functionally verifies a hybrid ZTA–network segmentation model. The stratified model combines the continuous verification capabilities of ZTA with the structural containment capabilities of segmentation through centralised policy enforcement.

2. RELATED WORK

Research has extensively examined the origins, behaviour, and detection of ransomware. Studies show that Ransomware-as-a-Service ecosystems allow even less experienced attackers to launch destructive campaigns [5], while threat assessments confirm that ransomware commonly relies on system-level exploitation and social engineering. Zero Trust has also been widely studied as a modern cybersecurity framework. Teerakanok *et al.* [6] outlined the operational challenges of migrating organisations to ZTA, particularly because policy enforcement evolves continuously. Yalla [7] argued that the effectiveness of ZTA improves when behavioural analytics are integrated into authentication and access procedures. A key limitation of ZTA, however, is that it cannot fully protect against phishing or poorly configured user endpoints.

Network segmentation has long been recognised as an efficient measure for mitigating ransomware spread. Traffic isolation, firewall enforcement, and access-control-list-mediated communication are critical controls for restricting lateral movement after an endpoint is compromised. A limitation of static segmentation rules is their reduced effectiveness against behaviourally adaptive threats such as Ryuk and LockBit ransomware [4].

Hybrid security models that incorporate both ZTA and segmentation are increasingly being promoted. Wonor *et al.* [8] presented a case study demonstrating the effectiveness of micro-segmentation and identity-sensitive controls in improving the detection and rapid isolation of ransomware traffic patterns. Similarly, NIST SP 800-207 [3] describes ZTA and the value of combined network-level controls. This paper presents the architectural design and implementation of such a hybrid system.

Table 1 compares the proposed model with selected related works across features relevant to ransomware mitigation.

3. METHODOLOGY

3.1. DESIGN PRINCIPLES

The hybrid model is designed to support centralised policy enforcement in which access-control decisions are made from a specified network checkpoint. This arrangement reflects the foundational principle of "never trust, always verify" [2]. The model uses identity verification and least privilege to grant network access, while automatically verifying a device's IP address and compliance posture [2, 3].

3.2. ARCHITECTURAL COMPONENTS

The model integrates multiple security structures into a layered defence comprising the following components:

1. ZTA component: This component performs identity verification, device security health assessment, including antivirus status and patch level [9], and real-time monitoring of session behaviour for irregularities [10].
2. Segmentation component: This component provides threat isolation by dividing the network into departmental subnets and regulating inter-subnet communication through access control lists based on least privilege [4].
3. Integration controller: This component coordinates the ZTA and segmentation components. When a threat is detected, the controller directs the segmentation mechanism to isolate the affected endpoint. The architectural structure is shown in Figure 1.

3.3. OPERATIONAL WORKFLOW

The model operates in three phases:

1. Pre-attack: Least-privilege access is enforced among critical assets, and all inter-subnet traffic is routed through a central enforcement point.
2. During attack: The ZTA component detects a threat through behavioural analysis or a posture violation and notifies the Integration Controller, which directs the segmentation mechanism to isolate the affected endpoint.
3. Post-attack: The system confirms that identified threats are contained and documents the incident for forensic analysis before initiating policy-driven or manual remediation.

Table 1. Comparison of the proposed model with related work. TTD, time-to-detect; TTC, time-to-contain.

Study	Continuous verification	Micro-segmentation	Automated containment	Open-source tools	Quantitative TTD/TTC reported	Key limitation	Contribution of this model
Yalla [7]	Yes	No	No	Not stated	No	No segmentation; lateral movement remains unaddressed	Adds segmentation and automated containment
Kharraz *et al.* [4]	No	Yes (static)	No	Not stated	Partial	Static rules may fail against adaptive threats, including Ryuk	Adds dynamic, posture-driven enforcement
Wonor *et al.* [8]	Yes	Yes	Partial	Not stated	No	No automated real-time policy adjustment	Provides fully automated ZTA–segmentation integration
NIST SP 800-207 [3]	Yes	Yes (cloud)	No	No (cloud-focused)	No	Limited treatment of legacy systems and physical segmentation	Provides a pfSense-based model applicable to small and medium-sized enterprise local area networks
This model	Yes	Yes (dynamic)	Yes	Yes (pfSense, VMware, GNS3)	Yes; TTD: 0.31 s, TTC: 0.32 s	Single-run TTD/TTC; controller is a single point of failure	Validates a hybrid model with real Ryuk PCAP data

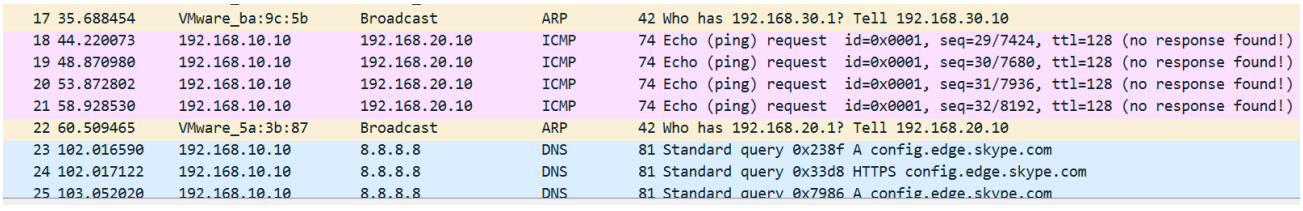

17	35.688454	VMware_ba:9c:5b	Broadcast	ARP	42 Who has 192.168.30.1? Tell 192.168.30.10
18	44.220073	192.168.10.10	192.168.20.10	ICMP	74 Echo (ping) request id=0x0001, seq=29/7424, ttl=128 (no response found!)
19	48.870980	192.168.10.10	192.168.20.10	ICMP	74 Echo (ping) request id=0x0001, seq=30/7680, ttl=128 (no response found!)
20	53.872802	192.168.10.10	192.168.20.10	ICMP	74 Echo (ping) request id=0x0001, seq=31/7936, ttl=128 (no response found!)
21	58.928530	192.168.10.10	192.168.20.10	ICMP	74 Echo (ping) request id=0x0001, seq=32/8192, ttl=128 (no response found!)
22	60.509465	VMware_5a:3b:87	Broadcast	ARP	42 Who has 192.168.20.1? Tell 192.168.20.10
23	102.016590	192.168.10.10	8.8.8.8	DNS	81 Standard query 0x238f A config.edge.skype.com
24	102.017122	192.168.10.10	8.8.8.8	DNS	81 Standard query 0x33d8 HTTPS config.edge.skype.com
25	103.052020	192.168.10.10	8.8.8.8	DNS	81 Standard query 0x7986 A config.edge.skype.com

Figure 3. Wireshark logs showing blocked packets from the HR subnet to the R&D subnet.

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	flow_key	src_ip	dst_ip	proto	src_port	dst_port	packet_co	total_byte	first_time	last_time	duration_	bytes_per_sec	
2	192.168.1	192.168.1	192.168.1	TCP	445	49245	181750	2.5E+08	19.8646	202.969	183.104	1385042	
3	192.168.1	192.168.1	192.168.1	TCP	49245	445	180444	2.5E+08	19.8642	202.969	183.105	1338196	
4	192.168.1	192.168.1	192.168.1	SMB2	445	49245	11351	7715762	19.869	202.923	183.054	42150.2	
5	192.168.1	192.168.1	192.168.1	SMB2	49245	445	10603	7548223	19.865	202.969	183.104	41223.7	
6	192.168.1	192.168.1	192.168.1	NBSS	445	49245	2706	4059000	76.5168	202.73	126.213	32160	
7	192.168.1	192.168.1	192.168.1	NBSS	49245	445	478	717000	79.0201	202.947	123.926	5785.69	
8	192.168.1	192.168.1	239.255.2	SSDP	1900	1900	18	9111	15.521	23.9595	8.43841	1079.71	
9	192.168.1	192.168.1	239.255.2	SSDP	61466	1900	19	2871	17.5993	27.5992	9.99993	287.102	
10	192.168.1	192.168.1	192.168.1	SSDP	1900	61466	6	2636	18.58	27.0643	8.48426	310.693	
11	192.168.1	192.168.1	192.168.1	NBNS	137	137	15	1386	17.1535	22.38	5.22651	265.186	
12	192.168.1	192.168.1	192.168.1	BROWSER	138	138	3	693	23.1575	113.286	90.1283	7.68904	
13	192.168.1	192.168.1	192.168.1	TCP	49255	2869	5	212	19.8474	88.7165	68.8691	3.0783	
14	192.168.1	192.168.1	192.168.1	TCP	49258	2869	5	212	26.4006	135.097	108.696	1.95039	
15	192.168.1	192.168.1	192.168.1	TCP	49255	2869	5	208	19.043	81.3814	62.3384	3.33663	
16	192.168.1	192.168.1	192.168.1	TCP	49253	2869	5	208	16.2433	76.2681	60.0247	3.46524	
17													

Figure 4. Extracted Ryuk ransomware flow characteristics used to detect network behaviour.

Table 2. Comparative evaluation of detection and containment performance.

Approach	Detection mechanism	Attack start	First detection timestamp	TTD (s)	Containment timestamp	TTC (s)
Standalone ZTA	Identity/posture checks	14:23:45.510	14:23:47.010	1.50	14:23:48.010	2.50
Standalone network segmentation	Firewall rule block at network layer	10:15:12.200	10:15:12.850	0.65	10:15:12.850	0.65
Hybrid model (proposed)	Correlated host telemetry, inline segmentation, and prioritised rules	15:54:56.634	15:54:56.946	0.31	15:54:56.952	0.32

4. SYSTEM IMPLEMENTATION

4.1. IMPLEMENTATION ENVIRONMENT

To evaluate the viability of the model, a proof-of-concept environment was configured to simulate a small-to-medium-

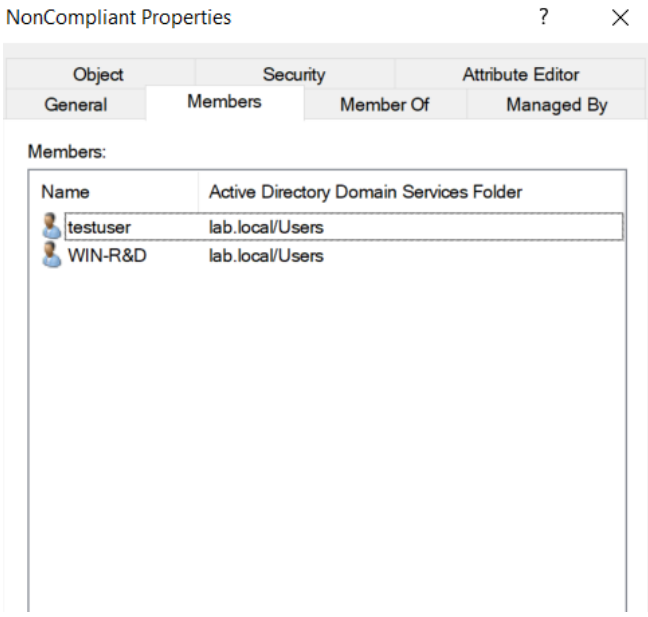

Figure 5. NonCompliant R&D machine in the Active Directory security group.

Table 3. Detection performance metrics across 20 controlled simulation runs.

Metric	Value	Formula/notes
True positives (TP)	9	Attacks correctly detected and contained
False positives (FP)	2	Benign traffic incorrectly flagged; false positive rate: 10%
False negatives (FN)	1	Attack occurred but was not detected
True negatives (TN)	8	Benign traffic correctly unblocked
Accuracy	85.0%	$(TP + TN)/20 = 17/20$
Precision	81.8%	$TP/(TP + FP) = 9/11$
Recall	90.0%	$TP/(TP + FN) = 9/10$
F1-score	85.7%	$2 \times$ (Precision $\times$ Recall)/(Precision + Recall)

sized enterprise local area network (LAN). The network comprised three departmental subnets: Human Resources (HR; 192.168.10.0/24), Research and Development (R&D; 192.168.20.0/24), and Administration (192.168.30.0/24). These subnets were interconnected through a pfSense firewall (v2.6.0), which served as the central Policy Enforcement Point. GNS3 (v2.2.5) was used to visualise the network topology and orchestrate packet capture. The HR, R&D, and Administration subnets each used VMware Workstation Pro 16.x to host Windows 10 virtual machines as departmental endpoints, while Windows Server 2022 hosted the ZTA policy logic and posture collection service. The network topology is shown in Figure 2.

4.2. MAPPING OF ARCHITECTURE TO IMPLEMENTATION

The ZTA component used Active Directory security groupings alongside a custom posture collector running as an HTTP listener on port 5000 of the Windows Server 2022 host. It examined destination IP addresses, protocols, ports, packet counts, total bytes, bytes per second, and flow duration. These behavioural indicators were derived from actual Ryuk ransomware packet-capture (PCAP) files processed using TShark and a Python extraction script.

The detection logic applied rule thresholds derived from the Ryuk flow analysis. A device was classified as `NonCompliant` when all four of the following conditions were met simultaneously:

1. The destination port was 445 (Server Message Block, SMB), 139 (NetBIOS), or 3389 (Remote Desktop Protocol, RDP), which are protocols commonly exploited during ransomware lateral movement and encryption phases.
2. Total bytes per flow exceeded 20,000 bytes, indicating possible bulk data staging or file-encryption activity.
3. Bytes per second exceeded 10,000, consistent with SMB write bursts during rapid file encryption.
4. Packet count per flow exceeded 20, characteristic of sustained ransomware-related connection attempts.

A simplified representation of the composite rule is:

$$(\text{dst\_port} \in \{445, 139, 3389\}) \wedge (\text{total\_bytes} > 20000) \wedge (\text{bytes\_per\_sec} > 10000) \wedge (\text{packet\_count} > 20) \Rightarrow \text{classify as NonCompliant and isolate.}$$

The segmentation component was activated by configuring pfSense firewall rules for the departmental subnets, blocking inter-subnet access except where explicitly permitted. When the Integration Controller received a ransomware-activity alert from the ZTA component, it transferred the device to the "NonCompliant" Active Directory group and added its IP address to the pfSense firewall alias `NonCompliant_Hosts`, thereby triggering immediate packet-level blocking.

5. RESULTS

5.1. VERIFICATION OF SEGMENTATION

Preliminary connectivity tests validated the segmentation layer. Consistent with the configured policies, ping requests from the HR subnet to the R&D subnet were blocked with 100% packet loss, while the Administration subnet maintained permitted connectivity to both HR and R&D subnets with 0% packet loss. As shown in Figure 3, Wireshark confirmed that denied packets were dropped at the pfSense interface.

5.2. VERIFICATION OF ZTA POLICIES

ZTA policy verification was conducted by simulating a non-compliant endpoint. When a device reported an outdated patch level or disabled antivirus, the posture collector evaluated the submitted parameters against policy thresholds (antivirus enabled, `LastPatchDays` ≤ 30, and `CustomOK = true`) and correctly classified the device as NonCompliant. Each classification event was logged with a timestamp, device IP address, posture values, and compliance outcome.

5.3. VERIFICATION OF THE HYBRID INTEGRATION PROCESS

The automated link between ZTA and segmentation was validated by simulating ransomware activity on the network. The

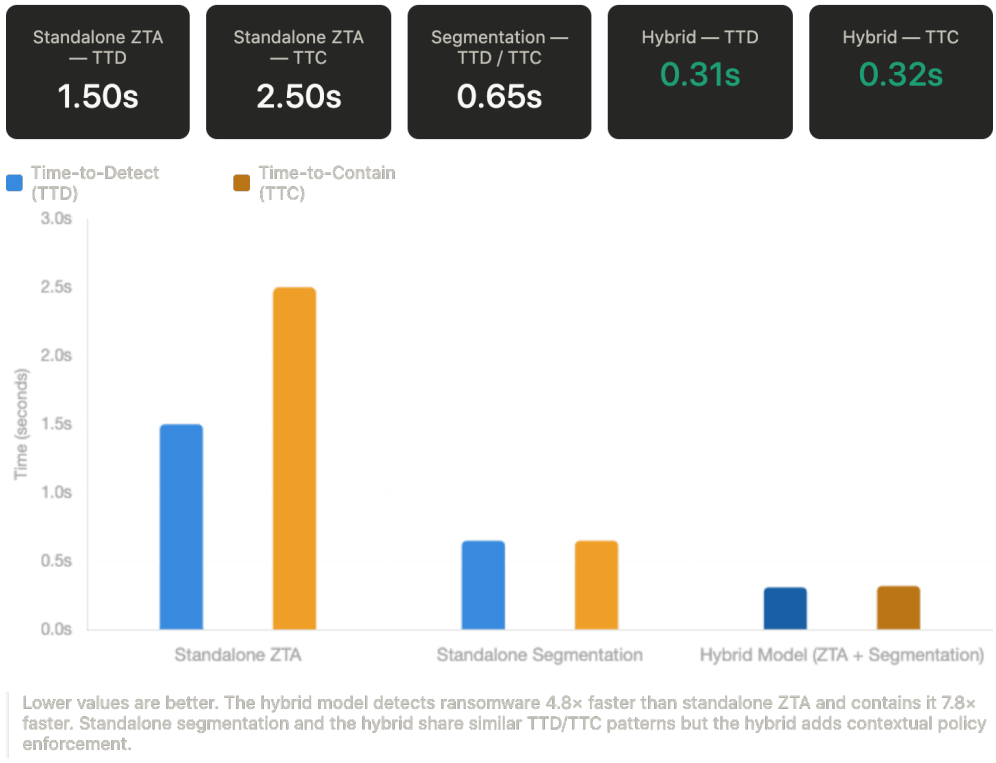

Figure 6. Comparison of TTD and TTC across the three approaches.

attacking host (WIN-R&D, 192.168.20.10) transmitted network flows that reflected the characteristics of Ryuk ransomware extracted from the PCAP dataset, as shown in Figure 4. The ZTA flow-based detection mechanism recognised this traffic as malicious and classified the device as non-compliant. The Integration Controller then automatically transferred the device to the ''Non-Compliant'' Active Directory group and added its IP address to the pfSense blocklist (Figure 5). Subsequent malicious packets were blocked, confirming that the model successfully translated ransomware activity detection into active network containment.

5.4. QUANTITATIVE PERFORMANCE RESULTS

Detection and containment were measured using timestamped simulation logs and compared with standalone ZTA and standalone network segmentation. Table 2 presents the comparative time-to-detect (TTD) and time-to-contain (TTC) values.

The hybrid model achieved a TTD of 0.31 s and a TTC of 0.32 s, outperforming standalone network segmentation (TTD: 0.65 s; TTC: 0.65 s) and standalone ZTA (TTD: 1.50 s; TTC: 2.50 s). The advantage of the hybrid approach lies in its ability to correlate rapid host-side telemetry with inline segmentation rules, thereby producing earlier and more confident detection than either mechanism in isolation.

6. DISCUSSION

The development and verification of the hybrid model demonstrate the feasibility of integrating ZTA with network segmentation. The implementation shows that open-source tools such as pfSense can effectively serve as a unified Policy Enforcement Point, making this advanced security architecture more accessible to resource-constrained environments [6].

One notable finding is the reduction in operational complexity. By centralising enforcement at the network chokepoint between segments, the model avoids the need for identity agents on every endpoint, thereby simplifying policy administration [8].

The effectiveness of the model depends on the accuracy and timeliness of posture telemetry from endpoints. Any failure in the ZTA component's data collection could impair overall system responsiveness. In addition, the model relies on external signals for threat detection and does not independently perform deep behavioural analysis of all network traffic. This limitation could be addressed in future work by integrating an inline intrusion detection engine.

6.1. SECURITY ANALYSIS OF THE INTEGRATION CONTROLLER

The Integration Controller is a critical dependency in the hybrid architecture because it receives threat alerts from the ZTA component and issues containment commands to pfSense. As a centralised decision point, it is both a single point of failure and a high-value adversarial target. Compromise or disruption of the controller would disable the automated containment pipeline.

Several mitigations should be considered for production deployments:

1. Redundancy: The controller should be deployed in a high-availability configuration, such as active–passive failover, so that the failure of one instance does not disable containment.

2. Mutual authentication: All communication among the ZTA component, the controller, and pfSense should use certificate-based TLS mutual authentication. In this proof of concept, the posture collector accepted HTTP POST requests on port 5000; in production, this should be replaced with HTTPS and token-based authentication.
3. Immutable audit logging: Every enforcement action should be logged with a timestamp, the triggering event, and the initiating component identity. Immutable logs support forensic analysis and help detect attempts to suppress or manipulate containment actions.
4. Network isolation: The controller should reside in a dedicated administrative segment that is accessible only from the ZTA policy server and the pfSense management interface, and inaccessible from general-purpose endpoint subnets.

These measures would significantly reduce the risk of the controller becoming an attack vector and represent a recommended direction for hardening the architecture in operational deployments.

7. CONCLUSION

This paper presented a validated architectural model and implementation strategy for organisations adopting multi-layered defences against ransomware. The proof of concept used pfSense, VMware, and GNS3 to create a functional hybrid model in which device-posture evaluations were automatically converted into network-level enforcement actions. The hybrid model achieved a TTD of 0.31 s and a TTC of 0.32 s, outperforming both standalone ZTA and standalone network segmentation. Across 20 simulation runs, the model achieved 85.0% accuracy, 81.8% precision, 90.0% recall, and an F1-score of 85.7%, with a false positive rate of 10%. Future research should evaluate the model across a broader range of attack scenarios and assess hardening strategies for the Integration Controller to mitigate its role as a potential single point of failure.

DATA AVAILABILITY

The data supporting the findings of this study are available from the corresponding author upon reasonable request.

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

FUNDING

The authors declare that no funding was received during the preparation of this manuscript.

References

[1] G. Nagar, ‘‘The evolution of ransomware: Tactics, techniques, and mitigation strategies’’, International Journal of Scientific Research and Management **12** (2024) 1282. https://doi.org/10.18535/ijsrm/v12i06.ec09.

[2] J. Kindervag, ‘‘No more chewy centers: introducing the Zero Trust Model of information security’’, Forrester Research, Cambridge, MA, USA, 2010, pp. 1–17. Available online: https://www.forrester.com/report/no-more-chewy-centers-introducing-the-zero-trust-model-of-information-security/RES56682.

[3] S. Rose, O. Borchert, S. Mitchell & S. Connelly, ‘‘Zero trust architecture’’, National Institute of Standards and Technology Special Publication 800-207, Gaithersburg, MD, USA, 2020, pp. 1–50. https://doi.org/10.6028/NIST.SP.800-207.

[4] A. Kharraz, W. Robertson, D. Balzarotti, L. Bilge & E. Kirda, ‘‘Cutting the Gordian knot: A look under the hood of ransomware attacks’’, in Detection of Intrusions and Malware, and Vulnerability Assessment, M. Almgren, V. Gulisano & F. Maggi (Eds.), Springer, Cham, Switzerland, 2015, pp. 3–24. https://doi.org/10.1007/978-3-319-20550-2_1.

[5] S. Razaulla, C. Fachkha, C. Markarian, A. Gawanmeh, W. Mansoor & B. C. M. Fung, ‘‘The age of ransomware: A survey on the evolution, taxonomy, and research directions’’, IEEE Access **11** (2023) 40698. https://doi.org/10.1109/ACCESS.2023.3268535.

[6] S. Teerakanok, T. Uehara & A. Inomata, ‘‘Migrating to Zero Trust Architecture: Reviews and challenges’’, Security and Communication Networks **2021** (2021) 9947347. https://doi.org/10.1155/2021/9947347.

[7] M. R. Yalla, ‘‘Zero-trust security architecture in the AI era: A novel framework for enterprise cyber resilience’’, International Journal of Science and Research Archive **13** (2024) 4341. https://doi.org/10.30574/ijsra.2024.13.2.0172.

[8] I. A. Wonor, M. O. Musa & C. M. Osazuwa, ‘‘Zero trust and microsegmentation: Strengthening network security’’, The American Journal of Management and Economics Innovations **7** (2025) 45. https://doi.org/10.37547/tajmei/volume07issue10-05.

[9] G. Karantzas & C. Patsakis, ‘‘An empirical assessment of endpoint security systems against advanced persistent threats attack vectors’’, arXiv:2108.10422 (2021). https://doi.org/10.48550/arXiv.2108.10422.

[10] N. Sheikh, ‘‘Zero trust using network micro segmentation’’, IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), Vancouver, BC, Canada, 2021, pp. 1–6. https://doi.org/10.1109/INFOCOMWKSHPS51825.2021.9484645.