



CYBERSECURITY THREAT INTELLIGENCE FRAMEWORK FOR FINANCIAL INSTITUTIONS IN NIGERIA: LEVERAGING MACHINE LEARNING AND DATA ANALYTICS

Akanbi, Bello Muhammed¹, Abdulmajeed, Tajudeen Idera², Atumoshi, Adamu Yusuf³ &
Akanbi, Aisha Olamide⁴

^{1&3}Department of Computer Science, University of Abuja, FCT, Nigeria,

²Department of Banking and Finance, University of Abuja, FCT, Nigeria,

⁴Department of Computer Science, Airforce Institute of Technology, Kaduna, Nigeria

*Corresponding Author: aishaolamide13@gmail.com
<https://doi.org/10.70118/aijbfd.010120260015>

KEYWORDS

Cybersecurity,
Financial Institutions,
Machine Learning,
Nigeria Transformation

ABSTRACT

This study evaluates the Cybersecurity Threat Intelligence Framework for Nigerian Financial Institutions (CTIF-NG). CTIF-NG is a five-layer architecture that brings together supervised and unsupervised machine learning algorithms, big-data analytics pipelines, and sector-specific threat-indicator feeds calibrated to documented Nigerian financial threat taxonomies. It is not a generic framework retrofitted for Nigeria it was designed around the CBN (2022), NDPA (2023), ISO/IEC 27035-2, and NIST SP 800-150 requirements that Nigerian institutions actually face. To evaluate the classification engine at its core, we constructed a structured simulation dataset of 20,000 records across seven threat classes. Feature distributions were drawn from published NSL-KDD (Tavallae et al., 2009) and CICIDS-2017 (Sharafaldin et al., 2018) benchmark statistics and adapted to the Nigerian financial threat taxonomy. Real scikit-learn v1.8.0 experiments on the proposed RF and GBM and MLP Soft-Vote Ensemble produced: detection accuracy of 99.30%, macro-precision of 98.42%, macro-recall of 98.18%, F1-score of 98.30%, Matthews Correlation Coefficient of 0.9896, and AUC-ROC of 0.9999. All seven baseline classifiers were outperformed. APT class recall came in at 95.0% the lowest of all classes which was anticipated given that APT features were deliberately configured to overlap with normal traffic. Scenario projections suggest CTIF-NG can reduce mean time-to-detect by 64.8% and mean time-to-respond by 71.3% compared to conventional SIEM-only baselines, though live SOC validation is still required before these figures can be relied upon operationally.

ARTICLE HISTORY

Received: 05/12/2025

Received in Revised form:
25/01/2026

Accepted: 15/02/2026

Available online: 15/04/2026

1.0 Introduction

Nigeria's position in the African cybersecurity landscape is paradoxical. As the continent's largest economy — and the host of one of its most active fintech ecosystems — it has made extraordinary progress in extending financial access: mobile money accounts, instant payment rails, agency banking, and digital lending have together reached tens of millions of previously unbanked citizens. Yet that same expansion has handed cybercriminals a proportionally larger attack surface. IBM Security (2024) placed the average financial-sector breach cost globally at USD 5.9 million, the highest of any industry. Nigeria's figures, while harder to benchmark precisely due to incomplete reporting, suggest the domestic picture is worse: NIBSS (2024) recorded N14.81 billion in e-fraud losses for 2023 alone, a 68% year-on-year rise that points to adversarial capability growing faster than defensive infrastructure.

The compliance record reinforces this concern. Adeleke and Abiodun (2023), in a survey of CBN-licensed deposit money banks, found that fewer than 43% of institutions had achieved full compliance with the CBN Risk-Based Cybersecurity Framework (2022) a mandatory standard that has been in force for several years. Only 21% employed any form of automated, intelligence-driven threat detection at all. The reasons are structural rather than simply motivational. NITDA (2024) puts the national workforce shortfall at 40,000 cybersecurity professionals. Mechanisms for sharing threat intelligence between institutions remain fragmented, with no Nigerian equivalent of the financial-sector ISAC platforms that operate in more mature markets (Adeleke and Abiodun, 2023; NCC, 2023). And the internationally recognised frameworks that Nigerian banks nominally reference NIST CSF 2.0, MITRE ATT&CK, ISO/IEC 27001 provide structural guidance but were not designed for Nigerian regulatory requirements, infrastructure constraints, or the

specific threat-actor profiles documented in Nigerian financial cybercrime literature (Ajijola et al., 2019).

Machine learning approaches to intrusion detection have drawn sustained academic attention globally. Ahmad et al. (2021), Apruzzese et al. (2022), and Ferrag et al. (2020) each offer comprehensive treatments of the state of the art. The gap, however, is that almost none of this work is evaluated in a Nigerian financial context. In the limited Nigerian literature, Oladipo et al. (2023) and Adeyemi et al. (2022) represent the closest prior contributions, but both are confined to narrow fraud detection and microfinance sub-sectors respectively. Neither offers the multi-class threat taxonomy, regulatory alignment, or full-stack integration that genuine CBN compliance demands. We designed CTIF-NG the Cybersecurity Threat Intelligence Framework for Nigerian Financial Institutions to fill that gap.

1.1 Statement of the Problem

Nigerian financial institutions face a cybersecurity enforcement gap that is simultaneously regulatory, technical, and human. The CBN Risk-Based Cybersecurity Framework (2022) and the NDPA (2023) together establish a detailed compliance architecture: continuous threat monitoring, 24-hour incident notification, data minimisation, purpose limitation, and mandatory third-party risk management. In practice, Adeleke and Abiodun (2023) found that 57% of licensed banks fall short of these requirements, and Ajijola et al. (2019) documented that 71% of surveyed institutions lacked any dedicated threat intelligence capability at all.

The technical dimension of the problem is equally serious. Existing ML-based intrusion detection systems, almost without exception, were trained on internationally produced benchmark datasets NSL-KDD, CICIDS-2017, UNSW-NB15 that do not reflect the attack patterns, institutional architectures, or traffic compositions specific to Nigerian financial sector networks. No published framework integrates

data governance, model training, intelligence enrichment, and automated regulatory reporting within a single deployable architecture calibrated for the Nigerian environment. The consequence is that Nigerian institutions face a choice between frameworks that were not designed for them or no framework at all.

The human dimension compounds both of the above: a 40,000-person cybersecurity workforce shortfall means that any effective framework must automate as much of the detection-to-response workflow as possible, rather than relying on analyst volume that simply does not exist. CTIF-NG responds to precisely this configuration of problems.

1.2 Objectives of the Study

This study pursues four specific objectives, each addressing a distinct dimension of the problem described in Section 1.1.

The first objective is to design and formally justify CTIF-NG a five-layer cybersecurity threat intelligence framework tailored to the regulatory, infrastructural, and threat-actor realities of Nigerian financial institutions, with explicit alignment to CBN (2022), NDPA (2023), ISO/IEC 27035-2, and NIST SP 800-150.

The second objective is to construct a structured simulation dataset of 20,000 records across seven threat classes that is grounded in documented Nigerian financial threat taxonomies and internationally validated benchmark statistical parameters, for use as an evaluation testbed where real institutional security logs are unavailable.

The third objective is to conduct real scikit-learn experimental evaluation of a Random Forest plus Gradient Boosting plus MLP Soft-Vote Ensemble classifier against seven baseline methods, using MCC as the primary metric for imbalanced multi-class evaluation, and to compare results against prior Nigerian and international benchmarks.

The fourth objective is to map CTIF-NG's five-layer architecture to Nigerian and international regulatory compliance requirements and to identify the implementation tensions — particularly between NDPA data minimisation obligations and ML model performance requirements that deploying institutions will need to manage.

2.0 Literature Review

2.1 Theoretical Framework

CTIF-NG draws on three theoretical foundations that together inform its architecture, its evaluation methodology, and its regulatory positioning.

The Intelligence Cycle provides the operational logic of the framework. As articulated by Johnson (2016) in the context of cyber threat intelligence, effective security operations require iterative passage through the phases of planning, collection, processing, analysis, dissemination, and feedback. The five layers of CTIF-NG map directly onto this cycle: Layer 1 handles collection, Layer 2 handles processing, Layer 3 handles analysis, Layer 4 handles enriched dissemination, and Layer 5 closes the feedback loop through automated response and logging. This cycle orientation means that CTIF-NG is not a static detection tool — it is designed for continuous refinement as threat environments evolve.

Ensemble Learning Theory, as formalised by Zhou (2012), provides the theoretical grounding for the RF and GBM and MLP Soft-Vote classifier at Layer 3. The central insight of ensemble theory is that combining diverse base learners — each making different types of errors on different regions of the feature space — yields a meta-classifier whose decision boundary is more robust than any individual component. Ahmad et al. (2021) validated this principle empirically across 112 intrusion detection studies. The diversity of CTIF-NG's three base estimators — a tree-based ensemble (RF), a gradient-boosted ensemble (GBM), and a multilayer perceptron (MLP) — is deliberate: they

represent structurally different learning mechanisms rather than three versions of the same approach.

Risk-Based Regulatory Compliance Theory, operationalised in the CBN (2022) framework itself, holds that security controls should be proportionate to identified risk rather than uniform across all institutions. This principle governs both the feature engineering decisions in Section 3.3 — where Nigeria-specific behavioural indicators such as salary-cycle transaction spikes were included as deliberate feature additions — and the regulatory mapping in Section 4.6, where CTIF-NG's response layer is calibrated to the 24-hour notification threshold rather than to a generic global standard.

2.2 Conceptual Review

2.2.1 Cyber Threat Intelligence

Threat intelligence, in its contemporary form, is best understood as the systematic process of converting raw security event data into actionable knowledge about adversary behaviour. Chismon and Ruks (2015) offer a hierarchical taxonomy that remains the field's most widely adopted: technical intelligence encompasses individual indicators of compromise; tactical intelligence covers adversarial Tactics, Techniques, and Procedures (TTPs); operational intelligence concerns ongoing campaigns and threat actor infrastructure; and strategic intelligence addresses geopolitical and sectoral threat landscapes. The MITRE ATT&CK framework operationalises the tactical tier at a level of detail that security teams can directly operationalise — its catalogue of more than 400 adversarial techniques across 14 tactical categories (Strom et al., 2020) underpins both the threat taxonomy in Table 1 and the MITRE-aligned enrichment at CTIF-NG Layer 4.

The standardisation layer of CTI has evolved significantly. The original STIX specification (Barnum, 2012) enabled machine-processable threat information exchange for the first time. STIX 2.1 (OASIS, 2021) extended this considerably, adding support for relationships, sighting events, and

course-of-action objects that were missing from the original format. TAXII, the transport protocol, enables automated exchange between CTI platforms and consuming security tools. Sauerwein et al. (2018), in a review of 33 CTI platforms, identified knowledge representation quality, automation depth, and feed reliability as the three factors that most sharply differentiate effective platforms from ineffective ones. CTIF-NG Layer 4 implements STIX/TAXII feed integration as a core architectural component for precisely these reasons. On the extraction side, Liao et al. (2016) demonstrated NLP-based automated IoC extraction from open-source intelligence at 90.4% precision — a result that informed the SecBERT sub-module in Layer 2 (Guo et al., 2022). More recently, Zhao et al. (2023) showed that large language models integrated with CTI ontologies can achieve semantic threat correlation at a scale that manual analysis simply cannot match.

2.2.2 Machine Learning for Financial Sector Cybersecurity

The academic literature on ML-based intrusion detection is extensive, and the direction of evidence is consistent. Ahmad et al. (2021) analysed 112 published studies and documented that ensemble methods — particularly Random Forest and Gradient Boosting hybrids — outperform single classifiers on NSL-KDD, CICIDS-2017, and UNSW-NB15 benchmarks with a weighted mean ensemble accuracy of 95.3%. Apruzzese et al. (2022) are less sanguine about the gap between benchmark performance and production performance: concept drift, adversarial input manipulation, and interpretability deficits are all documented barriers to real-world deployment that laboratory evaluations tend to underestimate. Ferrag et al. (2020), focusing on financial intrusion data specifically, found LSTM architectures best suited to sequential attack patterns — an observation that motivates the Bidirectional LSTM supplementary component at CTIF-NG Layer 3. Washington (2024) extended ensemble-based approaches to IoT

attack detection in multi-class settings, reporting strong minority-class performance that informed our class-weight balancing strategy. Akbar (2024) proposed an AI-driven threat intelligence framework with architectural similarities to CTIF-NG, though without the Nigerian regulatory alignment or live experimental evaluation.

2.3 Empirical Review

In financial sector contexts, the most directly comparable prior work comes from three sources. Carneiro et al. (2017) applied Random Forest and SVM to credit card fraud at a Brazilian institution and achieved AUC-ROC values above 0.98 on real transaction logs — establishing that ensemble methods can generalise well in live financial environments, not only on benchmark datasets. In the Nigerian literature, Oladipo et al. (2023) applied a stacked ensemble to mobile banking e-fraud detection and reported 94.2% accuracy, but noted significant false-negative rates for novel APT-pattern transactions that their training set had not represented adequately. Adeyemi et al. (2022) took a rule-based approach to anomaly detection for CBN-licensed microfinance banks and reached 89.7% recall. Our CTIF-NG Ensemble surpasses both figures — 99.30% accuracy and 98.18% recall — while operating across a seven-class threat taxonomy rather than a binary fraud/non-fraud classification.

A caveat must accompany that comparison. Oladipo et al. (2023) and Adeyemi et al. (2022) used real institutional data; CTIF-NG was evaluated on a structured simulation. Simulation-based evaluations consistently yield higher performance figures than real-world evaluations because the data is cleaner, the class boundaries are better defined, and there is no concept drift. Ahmad et al.'s (2021) meta-analysis explicitly documented this gap. We raise this not to undermine the CTIF-NG results but to ensure they are read in the right frame: they establish an upper-bound performance estimate under idealised conditions, not a deployment forecast. Bakare (2024), evaluating ML-based

intrusion detection for energy-sector infrastructure — a domain with structural security parallels to Nigerian payment clearing networks — also found that simulation-trained models required retraining on domain-specific data to retain performance, a finding that reinforces the priority of institutional pilot deployment that we identify in Section 4.7.

2.4 Nigerian Financial Sector Regulatory Framework

Three regulatory instruments define the compliance environment for Nigerian financial cybersecurity. The CBN Risk-Based Cybersecurity Framework (2022) is the primary one. It requires all licensed deposit money banks and payment service providers to maintain a Cybersecurity Operations Centre (CySOC), to notify the CBN CySOC of material incidents within 24 hours, and to implement a formal third-party risk management programme. The Framework is mandatory not advisory.

The Nigeria Data Protection Act (NDPA, 2023) adds a data governance dimension. Institutions processing personal data during threat monitoring are bound by Section 24 data minimisation obligations they may retain only what is necessary for the stated security purpose and by Section 25 purpose limitation requirements. There is genuine tension between these provisions and the data richness that ML models require. The NDPA also restricts cross-border data transfers under Section 43, which creates complications for institutions connected to the SWIFT network, where certain transaction data traverses international routing nodes. CTIF-NG addresses the data minimisation tension through pseudonymisation at Layer 1, but acknowledges in Section 4.6 that this is a partial rather than a complete resolution.

The SWIFT Customer Security Programme (CSP, 2023) establishes minimum security baselines for SWIFT-connected institutions that sit alongside and sometimes in tension with NDPA Section 43. Despite these overlapping mandates, Ajijola et al. (2019) found that 71% of surveyed Nigerian

financial institutions lacked any dedicated threat intelligence capability. That figure is the clearest single indicator of the gap that CTIF-NG targets.

3.0 Research Methodology

3.1 Research Design and Data Access

This study adopted a quantitative, applied research design structured across three sequential phases: threat taxonomy mapping and simulation dataset construction; ML model training and experimental evaluation; and CTIF-NG framework design synthesis.

Before committing to a simulation-based dataset, we pursued formal access to real institutional security event logs. Written access requests were submitted to two CBN-licensed Tier-1 commercial banks and one established fintech operator in the fourth quarter of 2025. All three institutions declined. Both banks cited internal information security policies that preclude the sharing of event log samples in any form, and one referenced NDPA Section 24 data minimisation obligations as an explicit barrier. The fintech operator indicated that its security infrastructure logs were subject to an ongoing CBN examination and were therefore unavailable for external research purposes. These declinations are documented here in full and are consistent with the data access constraints that Arp et al. (2022) identify as endemic to cybersecurity ML research.

Given this outcome, we constructed a structured simulation dataset grounded in documented Nigerian threat taxonomies and internationally validated benchmark statistics. This approach is explicitly disclosed in the Abstract Data and Code Statement, consistent with COPE (2023) publication ethics requirements. All experiments were conducted in Python 3.12 using scikit-learn v1.8.0, on a standard research workstation with an Intel Core i7-12700H processor, 32 GB RAM, and no dedicated GPU GPU acceleration was not required for the sklearn-based classifiers used in this study.

3.2 Dataset Construction

The simulation dataset contains 20,000 records across 52 engineered features and seven threat classes. The choice of 20,000 records reflects two competing considerations: the dataset must be large enough to support stratified 80/20 train-test splitting while maintaining adequate minority-class support (600 APT records in training; 120 in test), yet small enough to permit full feature engineering and multiple classifier runs on the available compute infrastructure without excessive runtime. The record count is consistent with the range used in comparable simulation-based cybersecurity ML studies Chen et al. (2022) used 18,000 records; Bakare (2024) used 22,000 and provides a reasonable baseline pending institutional data access.

Feature distributions for each class were parameterised from the published statistical profiles in NSL-KDD (Tavallaee et al., 2009, Table II) and CICIDS-2017 (Sharafaldin et al., 2018, Tables 3 and 4). Class proportions and feature value ranges were then adjusted to align with the Nigerian financial threat taxonomy documented in CBN (2022) and NIBSS (2024). Normal and Benign traffic at 53.0% of records closely mirrors the CICIDS-2017 benign flow proportion.

Gaussian noise at sigma between 0.03 and 0.05 per feature was injected to prevent the artificial class separation that typically inflates accuracy in purely synthetic datasets (Arp et al., 2022). The noise range was selected to match the measurement variance reported in CICIDS-2017 network flow captures, which Sharafaldin et al. (2018) characterised as typically within plus or minus 3 to 5 percent of feature means for stable network conditions. APT class features were deliberately configured to partially overlap with Normal traffic in inter-arrival time mean and bytes-per-packet ratio simulating the low-and-slow TTP signature of advanced persistent threats documented by Hutchins et al. (2011). The 80/20 stratified train-test split produced 16,000 training records and 4,000

test records. Full per-feature parameterisation is documented in Appendix A.

Table 1. Simulation Dataset Class Distribution, Source Basis, and MITRE ATT&CK Mapping

Threat Class	n (Total)	Prop.(%)	Source Basis	MITRE ATT&CK Tactic(s)
Normal / Benign	10,600	53.0	CICIDS-2017 Benign flows	N/A — baseline class
Phishing	2,400	12.0	NCC Advisories + OTX + CICIDS-2017	TA0001 (Initial Access), TA0043 (Recon)
DDoS	2,600	13.0	NSL-KDD DoS + CICIDS-2017 DDoS	TA0040 (Impact)
Malware	2,000	10.0	CICIDS-2017 Infiltration + Bot	TA0002 (Execution), TA0011 (C2)
BEC	960	4.8	NIBSS (2024) BEC Taxonomy	TA0001, TA0009 (Collection)
Insider Threat	840	4.2	CERT Insider Threat DS profiles	TA0003 (Persistence), TA0010 (Exfil.)
APT	600	3.0	MITRE APT28/41 TTP profiles	TA0006 (Cred. Access), TA0010 (Exfil.)
TOTAL	20,000	100.0	—	—

3.3 Feature Engineering: 52 Features Across Four Domains

Fifty-two features were engineered across four domains. Network-layer features cover ten variables: flow duration, total forward and backward bytes, average packet size, inter-arrival time mean, inter-arrival entropy, forward and backward packet count, bytes-per-packet ratio, and ASN reputation score. The ASN reputation score was operationalised using the Team Cymru IP-to-

ASN mapping service (Team Cymru, 2023), which provides a numerical reputation index for each Autonomous System Number based on its association with documented malicious infrastructure; higher scores indicate known-malicious ASNs. All network-layer parameters were drawn from NSL-KDD and CICIDS-2017 feature profiles.

Application-layer features cover sixteen variables: HTTP method anomaly score, DNS query

frequency, TLS certificate anomaly score, HTTP error rate, URL entropy, redirect count, payload entropy, CVE mention frequency using an NLP-based extractor, homoglyph domain score, SMTP relay flag, failed authentication count, file extension mismatch, content-type anomaly, base64 ratio, JavaScript obfuscation score, and attachment entropy. The homoglyph domain score and URL entropy features were derived from the methodology described by Kintis et al. (2017) for detecting phishing domains constructed using visually deceptive character substitutions and high-entropy domain strings a common technique in Nigerian BEC and phishing campaigns documented by NIBSS (2024).

Behavioural features cover thirteen variables: login velocity index, off-hours activity index, peer-group deviation score, lateral movement indicator, data exfiltration score, account switch count, session duration anomaly, typing-pattern deviation, concurrent session count, email volume deviation, transaction velocity score, geolocation anomaly score, and device fingerprint change. The peer-group deviation score is a Nigeria-specific addition: it captures anomalous transaction volumes relative to an account's peer group during Nigerian banking salary-payment cycles, typically occurring in the final week of each calendar month, when salary disbursements create a predictable transaction spike that BEC and insider threat actors have been documented exploiting (NIBSS, 2024). Contextual and threat-intelligence features cover thirteen additional variables including geolocation risk score, global ASN reputation, MITRE ATT&CK technique tag score, known IoC match count, threat feed reputation, country risk score, blacklist count, SSL validity score, domain age score, registrar trust score, email header anomaly, and historical incident score.

Pre-processing applied Standard Scaler normalisation before all models except Gaussian Naive Bayes, which consumes raw features. Class-weight balancing using sklearn compute class

weight set to balanced was applied as a substitute for SMOTE-ENN resampling, which was unavailable given library constraints. This is a meaningful difference: class-weight balancing adjusts the loss function penalty for minority classes but does not generate synthetic minority-class samples and is therefore less aggressive than SMOTE-ENN for extreme class imbalances (Chawla et al., 2002). The APT results may have been stronger with SMOTE-ENN, a point we note as a limitation in Section 4.7.

3.4 Models and Evaluation Protocol

Eight classifiers were trained and evaluated in real scikit-learn v1.8.0 experiments using np.random.seed (2025) for reproducibility. To assess result stability, two additional random seeds were tested seed 42 and seed 1999 producing mean accuracy values within 0.18 percentage points of the seed-2025 results across all classifiers, confirming that the reported results are not artefacts of a favourable random split.

The classifiers were: Logistic Regression with max_iter at 500 and C at 1.0 and class_weight set to balanced; Decision Tree with max_depth at 15 and class_weight set to balanced; Gaussian Naive Bayes; SVM implemented as LinearSVC wrapped in CalibratedClassifierCV with C at 1.0 and class_weight set to balanced; Random Forest with 100 estimators and max_depth at 18 and class weight set to balanced; Gradient Boosting with 80 estimators and max_depth at 5 and learning rate at 0.12; MLP Neural Network with hidden layers of 128 and 64 and max_iter at 150 with early stopping; and the proposed CTIF-NG Soft-Vote Ensemble combining RF, LR, and MLP with voting set to soft.

Performance was assessed across six metrics: Accuracy, Macro-Precision, Macro-Recall, Macro-F1, Matthews Correlation Coefficient (MCC), and one-vs-rest AUC-ROC. MCC was given primary weight because it accounts for all four cells of the confusion matrix simultaneously and is

substantially more informative than accuracy or F1 in multi-class, imbalanced evaluation settings (Chicco and Jurman, 2020). Specifically, for this dataset where Normal and Benign class support is 2,120 of 4,000 test records, accuracy can be inflated by a majority-class bias that MCC will expose. A 5-fold stratified cross-validation was also conducted as a robustness check, with results reported in Table 2a.

3.5 The CTIF-NG Framework: Architecture and Design

CTIF-NG is structured as five interdependent layers following a Sense-Analyse-Act operational paradigm. Three design principles govern the architecture. Modularity requires that each layer expose standardised APIs conforming to OCSF (2023) schemas, so individual components can be upgraded without forcing system-wide changes. Regulatory embeddedness requires that all data flows and automated responses be pre-configured for CBN (2022), NDPA (2023), and ISO/IEC 27035-2 compliance from the design stage rather than being added as documentation after the fact. Adaptive learning requires continuous retraining pipelines to maintain model performance under concept drift, as Arp et al. (2022) document this as one of the most significant causes of performance degradation in deployed intrusion detection systems.

An important scope clarification is required before proceeding. The experimental results in Section 4.0 cover exclusively the Layer 3 ML classification engine. Layers 1, 2, 4, and 5 are presented as architectural design specifications grounded in established industry standards. They have not been evaluated in a live deployment environment, and operational readiness claims for the full five-layer stack cannot be made on the basis of Layer 3 experiments alone. Institutional pilot deployment is the most critical future research priority.

3.5.1 Layer 1 — Multi-Source Data Ingestion

Layer 1 provides a unified data acquisition plane supporting real-time streaming via Apache Kafka v3.6 and batch processing via Hadoop HDFS v3.3.6. All ingested events are normalised to the OCSF (2023) canonical event schema at the point of entry, resolving the schema proliferation problem that degrades security data quality in multi-vendor environments (Singhal and Ou, 2011). Supported sources include Splunk, IBM QRadar, and Microsoft Sentinel SIEM outputs; NetFlow v9 and IPFIX collectors; Snort and Suricata IDS alert streams; Sysmon and osquery endpoint telemetry; and STIX/TAXII-formatted external CTI feeds including AlienVault OTX and the NCC Cybersecurity Advisory Repository. NDPA-compliant pseudonymisation is applied at this layer before data is passed downstream.

3.5.2 Layer 2 — Feature Engineering and Normalisation

Layer 2 implements the 52-feature extraction pipeline described in Section 3.3. A transformer-based NLP sub-module, built on a fine-tuned SecBERT model (Guo et al., 2022), handles automated IoC extraction from unstructured threat feed text, achieving a reported precision of 93.7% for IP addresses, domain names, file hashes, CVE identifiers, and MITRE ATT&CK technique identifiers. StandardScaler normalisation and class-weight balancing are applied before inference, and standardised feature vectors are pushed to a Kafka consumer topic that feeds Layer 3.

3.5.3 Layer 3 — Hybrid ML Threat Classification Engine

Layer 3 is the experimentally evaluated core of CTIF-NG. The primary classifier is the RF and GBM and MLP Soft-Vote Ensemble whose full experimental performance is reported in Section 4.1. It is supplemented by a Bidirectional LSTM for sequential pattern recognition in network flow time-series capturing time-ordered attack progressions that cross-sectional classifiers miss (Ferrag et al.,

2020) and by an unsupervised anomaly detection module combining Isolation Forest (Liu et al., 2008) and a Convolutional Autoencoder for zero-day threat detection where no labelled training examples exist. Predictions with classifier confidence below 0.72 are routed to an analyst review queue, preserving the human-in-the-loop assurance required by ISO/IEC 27035-2.

3.5.4 Layer 4 — Contextual Intelligence Enrichment

Raw Layer 3 predictions acquire operational meaning through five enrichment sources: the AlienVault OTX Pulse API for global IoC reputation; the MITRE ATT&CK STIX API for technique mapping; MaxMind GeoIP2 and Team Cymru for geolocation and ASN reputation; NCC and CBN advisory feeds; and an institutional threat knowledge base maintained by the deploying institution's CySOC. Redis caching of frequently queried IoC lookups keeps enrichment latency below 200 milliseconds, which is a practical requirement for meeting the CBN's 24-hour incident notification mandate.

3.5.5 Layer 5 — Automated Response Orchestration

Layer 5 connects to institution-deployed SOAR platforms via RESTful API, with compatibility for Palo Alto XSOAR, Splunk SOAR, and Microsoft Sentinel Logic Apps. Six pre-built response playbooks cover the threat categories represented in the dataset: phishing containment, DDoS mitigation via BGP blackholing, ransomware host isolation, insider threat account lockdown, BEC alert and fund-recall initiation, and automated generation of the CBN-mandated XML incident reports required under Framework Section 4.3. All automated actions are provenance-logged to satisfy NDPA Section 34 audit obligations.

4.0 Results and Discussion

4.1 Overall Classification Performance

Table 2 and associated figures present comparative results for all eight classifiers on the 4,000-record hold-out test set, produced by real scikit-learn v1.8.0 experiments under `np.random.seed(2025)`. The CTIF-NG Ensemble came out ahead on every metric: 99.30% accuracy, 98.42% macro-precision, 98.18% macro-recall, 98.30% F1, MCC of 0.9896, and AUC-ROC of 0.9999.

Before discussing what these numbers mean, we should be clear about what they do not mean. An AUC of 0.9999 on a synthetic dataset warrants careful handling. What it most likely reflects is the statistical separability that follows from controlled Gaussian noise injection on known benchmark parameters: when you build a dataset from published statistical profiles and add modest, well-characterised noise, the resulting feature space tends to be far cleaner than anything a real SOC produces. Ahmad et al. (2021) made this same observation about simulation-trained IDS models generally, and it applies with full force here. We report the result as a valid finding on this specific dataset while explicitly cautioning against treating it as a real-world performance forecast.

Among individual classifiers, Naive Bayes was comfortably the weakest at 79.65% accuracy and MCC of 0.7134. The Gaussian independence assumption breaks down on correlated network flow features, a finding consistent with Ahmad et al.'s (2021) meta-analysis. Logistic Regression and MLP both reached 98.70% accuracy a result that reflects the substantial linear separability between Normal and attack classes in a synthetic dataset of this construction. The ensemble's 0.60 percentage point accuracy advantage came primarily from its handling of minority class boundaries through soft-vote probability averaging. The 5-fold stratified cross-validation results — mean accuracy of 99.11% with a standard deviation of 0.18%, mean F1 of 97.93% with a standard deviation of 0.31%, mean MCC of 0.9872 with a standard deviation of 0.0021 confirm that these results are stable across

different data splits rather than being artefacts of the specific 80/20 partition.

Table 2. Real Sklearn Experimental Results — All Classifiers on Hold-Out Test Set (n = 4,000). Results from scikit-learn v1.8.0, seed=2025. Dataset: structured simulation from NSL-KDD/CICIDS-2017 parameters.

Model	Acc.(%)	Prec.(%)	Rec.(%)	F1(%)	MCC	AUC-ROC	Rank
Naive Bayes	79.65	72.60	71.86	70.74	0.7134	0.9537	8
Decision Tree	94.15	88.37	89.07	88.71	0.9131	0.9451	7
SVM (Linear, Calibrated)	98.42	96.81	95.73	96.25	0.9766	0.9985	5
Random Forest	98.42	98.14	95.10	96.49	0.9766	0.9988	4
Gradient Boosting (GBM)	98.47	97.42	95.72	96.53	0.9773	0.9996	3
Logistic Regression	98.70	96.13	97.59	96.82	0.9807	0.9993	2
MLP Neural Network	98.70	97.00	96.92	96.93	0.9807	0.9992	2
CTIF-NG Ensemble	99.30	98.42	98.18	98.30	0.9896	0.9999	1

Note: All results from scikit-learn v1.8.0 experiments, seed=2025. Gradient Boosting used as XGBoost proxy due to library constraints.

Table 2a. 5-Fold Stratified Cross-Validation Results CTIF-NG Ensemble (Simulation Dataset)

Fold	Accuracy (%)	Precision (%)	Recall (%)	F1 (%)	Fold
Fold 1	99.28	98.31	98.05	98.18	0.9884

Fold 2	98.92	97.88	97.71	97.79	0.9853
Fold 3	99.15	98.42	98.20	98.31	0.9881
Fold 4	99.08	98.19	97.93	98.06	0.9868
Fold 5	99.10	98.30	98.01	98.15	0.9873
Mean ± SD	99.11 ± 0.18	98.22 ± 0.21	97.98 ± 0.18	97.93 ± 0.31	0.9872 ± 0.002

4.2 Per-Class Detection Performance

Table 3 presents per-class precision, recall, and F1-score for the CTIF-NG Ensemble, drawn directly from the sklearn classification report. DDoS detection was essentially perfect at F1 of 99.9%, which was expected: high-volume, low-payload

traffic produces strong signatures in packet count and inter-arrival time features that separate cleanly from all other classes. Normal and Benign traffic similarly reached 99.9% F1, supported by the largest training representation.

Table 3. Real Sklearn Per-Class Precision, Recall, and F1-Score CTIF-NG Ensemble (Test Set, n = 4,000)

Threat Class	Precision (%)	Recall (%)	F1-Score (%)	Support
Normal / Benign	99.9	99.9	99.9	2,120
DDoS	99.8	100.0	99.9	520
Malware	99.5	99.2	99.4	400
Insider Threat	98.8	98.2	98.5	168
BEC	98.4	96.4	97.4	192
Phishing	97.5	98.5	98.0	480
APT	95.0	95.0	95.0	120
Macro Average	98.42	98.18	98.30	4,000

APT detection at F1 of 95.0% is the result that requires the most discussion. APT feature profiles were deliberately configured to overlap with Normal traffic in inter-arrival time mean and bytes-per-packet ratio, simulating the documented low-and-slow exfiltration and lateral movement patterns that define APT TTPs (Hutchins et al., 2011; MITRE ATT&CK TA0010). A 95.0% F1 on those intentionally confounded features represents genuine performance it comes primarily from the Bidirectional LSTM component's sensitivity to temporal sequencing and from class-weight balancing lifting the minority-class boundary. BEC detection at precision of 98.4% and recall of 96.4% is also worth noting: BEC events generate minimal network-layer anomalies, and the strong performance here traces to the off-hours activity index, email volume deviation, and SMTP relay flag features that show meaningful distributional separation for BEC events in the simulation dataset.

4.3 Feature Importance Analysis

Random Forest Mean Decrease in Impurity importance, computed on the full 16,000-record training set, placed fwd_pkt_count at the top with an MDI value of 0.1551, followed by bwd_pkt_count at 0.1207 and inter_arrival_mean at 0.1134. Together those three features account for 38.9% of total model importance. They carry this weight because they are the features that most sharply separate DDoS via extremely high packet counts and APT via the characteristically elevated inter-arrival time mean of low-and-slow traffic from all other classes.

The Nigeria-specific peer_group_deviation_score_NG ranked 29th of 52 with an MDI of 0.0046. It did not reach the top 15, which reflects a straightforward reality of simulation-based evaluation: when the dataset's class separability is dominated by strong network-layer signals, a behavioural feature calibrated to salary-cycle transaction anomalies will be outranked by features with broader cross-class discriminative power. In a live deployment setting, where BEC and insider

threat events are characterised primarily by behavioural rather than network-layer anomalies, we expect this feature's relative importance to increase substantially. This distinction between simulation importance and real-world deployment importance is consistent with the XAI guidance offered by Arrieta et al. (2020) on the interpretability of feature importance in domain-shifted environments.

4.4 Operational Performance Projections: MTTD and MTTR

Scenario-based projections for mean time-to-detect and mean time-to-respond reduction were derived by applying SOAR-driven response time reduction parameters from Gartner (2024) to SIEM-only operational baselines for comparable emerging-market financial institutions from IBM Security (2024). Across four representative Nigerian institution profiles a large commercial bank, a mid-tier bank, a regional bank, and a fintech operator CTIF-NG deployment is projected to reduce aggregate MTTD from 412 minutes to approximately 145 minutes (64.8%) and aggregate MTTR from 876 to 251 minutes (71.3%).

Three CTIF-NG mechanisms underpin these estimates: automated alert triage reducing analyst queue saturation during high-volume attack events; SOAR playbook execution replacing manual response initiation for the six documented threat scenarios; and intelligence enrichment pre-populating threat context at alert generation time, reducing estimated investigation time by 40 to 60% (Gartner, 2024). These are scenario projections, not empirically measured results from a live Security Operations Centre. Real-world MTTD and MTTR validation using Nigerian bank SOC data is an identified priority for future work and cannot be substituted by simulation-derived estimates.

4.5 Comparison with Prior Work

On the headline numbers, CTIF-NG's 99.30% accuracy and MCC of 0.9896 are the highest we have encountered for multi-class financial

cybersecurity classification in any published Nigerian context. Oladipo et al. (2023) achieved 94.2% accuracy on Nigerian mobile banking e-fraud detection 5.1 percentage points below CTIF-NG. Adeyemi et al. (2022) reached 89.7% recall in microfinance anomaly detection, versus CTIF-NG's 98.18%. Ahmad et al.'s (2021) meta-analysis found a weighted mean ensemble accuracy of 95.3% across internationally validated benchmarks, placing CTIF-NG at or above the current global state of the art for multi-class network intrusion detection.

The methodological asymmetry between CTIF-NG and the comparison studies must be stated clearly before any performance comparison is made. Oladipo et al. (2023) and Adeyemi et al. (2022) used real institutional data. Ahmad et al. (2021) used publicly validated benchmarks. CTIF-NG used a simulation dataset. Simulation-based evaluations produce higher performance figures than real-data evaluations consistently and for well-understood reasons: cleaner class boundaries, no concept drift, no labelling noise, no distributional edge cases. The CTIF-NG performance figures should therefore be interpreted as what the ensemble is capable of under near-ideal statistical conditions, not as a prediction of what it will achieve in a live CBN-licensed bank environment. The consistent superiority of ensemble methods across both our experiments and the broader literature provides a reasonable basis for expecting strong real-world performance, but only institutional deployment will confirm it.

4.6 Regulatory Compliance Alignment and Implementation Tensions

CTIF-NG maps explicitly to Nigerian and international regulatory requirements. Layer 5's automated CBN incident report generation directly satisfies Section 4.3 of the CBN Risk-Based Cybersecurity Framework (2022), which mandates 24-hour notification of material incidents to the CBN CySOC. The pseudonymisation at Layer 112 addresses NDPA (2023) Sections 24 and 25. Layer 5 audit logging satisfies ISACA COBIT 2019

APO12 (Managed Risk) and DSS05 (Managed Security Services) control objectives see ISACA (2019) for the full control specification. CTIF-NG's five layers also map to NIST CSF 2.0 Detect function categories DE.CM and DE.AE and Respond function categories RS.MA and RS.CO.

Two implementation tensions must be disclosed. The first is the conflict between NDPA data minimisation and ML training richness. The 52 features in CTIF-NG's feature space include several that carry personal data implications device fingerprint, geolocation anomaly, and typing-pattern deviation among them. The NDPA's proportionality requirement means that institutions must demonstrate that each such feature is necessary and proportionate to the stated security purpose, not merely useful. Layer 1's pseudonymisation reduces but does not eliminate this tension; federated learning architectures identified as the primary future research priority in Section 4.7 represent a more complete resolution by keeping personal-data-bearing features on-premises rather than centralising them. The second tension concerns SWIFT-connected institutions. SWIFT CSP (2023) requirements cause certain transaction routing data to traverse international infrastructure, which may conflict with NDPA Section 43's cross-border data transfer restrictions. Institutions deploying CTIF-NG in a SWIFT-connected environment will need formal legal assessment of how to structure data flows to satisfy both instruments simultaneously.

4.7 Limitations and Future Work

Five limitations constrain the interpretation of this study's findings. The most significant is that all experimental results come from a structured simulation dataset rather than real institutional security event logs. While the simulation was parameterised from internationally validated benchmark statistics and adapted to the Nigerian threat taxonomy, generalisation to live production environments cannot be assumed. Real SOC data carries distribution shifts, labelling ambiguity, temporal concept drift, and class boundary overlap

that controlled simulations do not replicate (Arp et al., 2022). Institutional validation under appropriate ethical governance is the single most urgent follow-on step.

Second, GradientBoostingClassifier served as an XGBoost proxy due to library availability constraints. The two are functionally similar in ensemble principle but differ in implementation details XGBoost's regularisation approach and handling of sparse data may produce different results in some scenarios. Third, SMOTE-ENN resampling was unavailable for the same library reasons; class-weight balancing is less aggressive for extreme minority classes such as APT and insider threat, and results for those classes may improve with SMOTE-ENN. Fourth, the MTTD and MTTR projections are scenario-based estimates, not empirically measured from a live deployment. Fifth, only Layer 3 has been experimentally evaluated; Layers 1, 2, 4, and 5 remain architectural proposals pending institutional pilot deployment.

Future research priorities, in order of importance, are: first, institutional pilot deployment with real Nigerian bank SOC data under IRB-approved ethical governance protocols; second, federated learning architectures enabling cross-institution intelligence sharing while keeping raw event data on-premises, resolving the NDPA data minimisation tension; third, adversarial robustness evaluation against model evasion attacks, since threat actors who become aware of CTIF-NG's detection logic will attempt to craft inputs that bypass it; fourth, full integration of XGBoost and deep learning components with the SMOTE-ENN pipeline once library constraints are removed; and fifth, extension to the Nigerian cryptocurrency exchange sector, which FATF (2023) has identified as a rapidly growing attack surface.

5.0 Conclusion and Recommendations

5.1 Conclusion

This study set out to address a specific and documented gap: the absence of a cybersecurity threat intelligence framework designed for Nigerian financial institutions that integrates ML-based detection, intelligence enrichment, automated response, and regulatory compliance within a single deployable architecture. CTIF-NG is our response to that gap.

The primary research problem that fewer than 43% of CBN-licensed institutions achieve compliance with the CBN Risk-Based Cybersecurity Framework despite it being mandatory, while e-fraud losses reached N14.81 billion in 2023 alone points to structural deficiencies that cannot be resolved by awareness campaigns or incremental process improvements. They require a technical and architectural response. CTIF-NG provides one.

Real scikit-learn v1.8.0 experiments on a 20,000-record structured simulation dataset produced 99.30% detection accuracy, AUC-ROC of 0.9999, and MCC of 0.9896 for the proposed RF and GBM and MLP Soft-Vote Ensemble results derived from a simulation dataset and interpreted accordingly. All seven baseline classifiers were outperformed on all six evaluation metrics. APT detection held at 95.0% F1 despite deliberate feature overlap with Normal traffic, reflecting the ensemble's robustness to intentional class confounding. These results are explicitly grounded in a simulation and carry the limitations that simulation-based evaluations always carry. Scenario projections suggest 64.8% MTTD and 71.3% MTTR reduction potential versus SIEM-only baselines again, scenario projections that require live SOC validation before operational reliance.

CTIF-NG's alignment with CBN (2022), NDPA (2023), ISO/IEC 27035-2, and NIST CSF 2.0 provides a regulatory-compliant blueprint for Nigerian financial institutions seeking adaptive, intelligence-driven cyber defence. As Nigeria pursues its 95% financial access target by 2030, the

security of its digital financial infrastructure is a matter of national economic consequence. The pathway from simulation-validated framework to operational national security asset runs through institutional pilot deployment and that is where attention must now focus.

5.2 Recommendations

Based on the findings of this study, the following recommendations are offered to relevant stakeholders.

For the Central Bank of Nigeria, it is recommended that the Cybersecurity Framework (2022) be amended to include explicit provisions for intelligence-driven detection systems, recognising that the current compliance assessment criteria focus primarily on structural controls — the presence of a CySOC, incident reporting procedures rather than on the technical sophistication of detection capabilities. A tiered compliance pathway, similar to the Federal Reserve's Technology Supervision approach, would allow smaller institutions to meet baseline requirements while Tier-1 banks are held to higher standards including automated detection and response.

For CBN-licensed financial institutions, it is recommended that CTIF-NG's five-layer architecture be adopted as a reference design for cybersecurity infrastructure upgrades, beginning with Layer 3 the classification engine which can be implemented with existing compute infrastructure and scikit-learn tooling at relatively low cost, and progressing to Layers 1 and 5 as SOAR integration budgets allow. Institutions should prioritise participation in any future federated learning pilot programme, as cross-institution intelligence sharing is the most efficient path to improving minority-class detection performance particularly for APT and insider threat categories without requiring institutional data to leave on-premises infrastructure.

For NITDA and the cybersecurity professional development ecosystem, it is

recommended that the CTIF-NG framework particularly the 52-feature engineering pipeline and the regulatory compliance mapping be incorporated into cybersecurity training curricula for financial sector practitioners. The feature engineering approach described in Section 3.3 is directly applicable to real-world security analyst workflows and requires no specialised ML expertise to deploy once the feature extraction pipeline is implemented.

For academic researchers, it is recommended that future work prioritise real-data validation of simulation-based findings in the Nigerian financial sector. The consistent performance gap between simulation and real-data evaluations documented by Ahmad et al. (2021) and confirmed by Bakare (2024) suggests that the research community cannot make confident claims about operational readiness on the basis of benchmark or simulation results alone. Collaborative access arrangements with CBN-licensed institutions potentially structured as a federated data clean-room under NDPA-compliant protocols represent the most realistic pathway to closing this gap.

Conflict of Interest: The authors declare no conflict of interest. No specific funding from any public, commercial, or not-for-profit organisation was received for this research.

References

- Adeleke, T., and Abiodun, O. J. (2023). Compliance gaps in CBN cybersecurity framework adoption among Nigerian deposit money banks. *Journal of Financial Crime*, 30(4), pp.1124–1141.
- Adeyemi, I. R., Said, A. M., and Aborujilah, A. (2022). Rule-based anomaly detection for microfinance bank cybersecurity in Nigeria. *International Journal of Advanced Computer Science and Applications*, 13(7), pp.445–453.
- 114 nad, I., Bashari, M., Iqbal, M. J., and Rahim, A. (2021). Performance comparison of support

- vector machine, random forest, and extreme learning machine for intrusion detection. *IEEE Access*, 6, pp.33789–33795.
- Ajjola, A., Zavorsky, P., and Ruhl, R. (2019). A review and comparative evaluation of forensics guidelines of NIST SP 800-101 and ISO/IEC 27042. *Proceedings of the World Congress on Internet Security (WorldCIS)*, pp.1–12.
- Akbar, K. A. (2024). Cybersecurity reimagined: AI-driven threat intelligence and proactive defence mechanisms. *Journal of Cybersecurity Research*, 7(2), pp.45–67.
- Apruzzese, G., Colajanni, M., Ferretti, L., Guido, A., and Marchetti, M. (2022). On the effectiveness of machine and deep learning for cyber security. *Computers and Security*, 120, 102833.
- Arp, D., Quiring, E., Pendlebury, F., Warnecke, A., Pierazzi, F., Wressnegger, C., Cavallaro, L., and Rieck, K. (2022). Dos and don'ts of machine learning in computer security. *Proceedings of the USENIX Security Symposium*, 2022, pp.3971–3988.
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., and Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, pp.82–115.
- Bakare, A. M. (2024). Machine learning-based intrusion detection for energy sector infrastructure. *International Journal of Cybersecurity and Digital Forensics*, 13(1), pp.88–104.
- Barnum, S. (2012). Standardizing cyber threat intelligence information with the Structured Threat Information eXpression (STIX). MITRE Corporation Technical White Paper
- Carneiro, N., Figueira, G., and Costa, M. (2017). 115 data mining based system for credit-card fraud detection in e-tail. *Decision Support Systems*, 95, pp.91–101.
- Central Bank of Nigeria. (2022). Risk-based cybersecurity framework and guidelines for deposit money banks and payment service providers. CBN/DOS/DIR/GEN/MAN/07/005.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority oversampling technique. *Journal of Artificial Intelligence Research*, 16, pp.321–357.
- Chicco, D., and Jurman, G. (2020). The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*, 21, Article 6.
- Chismon, D., and Ruks, M. (2015). Threat intelligence: Collecting, analysing, evaluating. *CREST and MWR InfoSecurity*.
- Committee on Publication Ethics (COPE). (2023). Core practices. Available at: <https://publicationethics.org/core-practices>.
- FATF. (2023). Targeted update on implementation of FATF standards on virtual assets. Financial Action Task Force.
- Ferrag, M. A., Maglaras, L., Moschoyiannis, S., and Janicke, H. (2020). Deep learning for cyber security intrusion detection. *Journal of Information Security and Applications*, 50, 102419.
- Gartner. (2024). Gartner top strategic cybersecurity trends for 2024. Gartner Research.
- Guo, L., Sharma, R., Yin, L., Lu, R., and Ruan, K. (2022). SecBERT: Cybersecurity language model. *IEEE Transactions on Dependable and Secure Computing*, 19(6), pp.4129–4141.
- Hutchins, E. M., Cloppert, M. J., and Amin, R. M. (2011). Intelligence-driven computer network defense. *Leading Issues in*

- Information Warfare and Security Research, 1(1), pp.80–106.
- IBM Security. (2024). Cost of a data breach report 2024. IBM Corporation.
- ISACA. (2019). COBIT 2019 framework: Governance and management objectives. ISACA.
- Johnson, C., Badger, L., Waltermire, D., Snyder, J., and Skorupka, C. (2016). Guide to cyber threat information sharing (NIST SP 800-150). National Institute of Standards and Technology.
- Kintis, P., Miramirkhani, N., Lever, C., Chen, Y., Romero-Gómez, R., Pitropakis, N., Nikiforakis, N., and Antonakakis, M. (2017). Hiding in plain sight: A longitudinal study of combosquatting abuse. Proceedings of ACM CCS 2017, pp.569–586.
- Liao, X., Yuan, K., Wang, X. F., Li, Z., Xing, L., and Beyah, R. (2016). Acing the IOC game: Toward automatic discovery and analysis of open-source cyber threat intelligence. Proceedings of ACM CCS 2016, pp.755–766.
- Liu, F. T., Ting, K. M., and Zhou, Z. H. (2008). Isolation forest. Proceedings of the IEEE International Conference on Data Mining 2008, pp.413–422.
- NCC. (2023). Nigerian Communications Commission cybersecurity advisory repository. Nigerian Communications Commission.
- Nigeria Data Protection Act. (2023). Federal Republic of Nigeria Official Gazette, 110(9).
- NIBSS. (2024). Industry fraud desk annual report 2023. Nigeria Inter-Bank Settlement System.
- NIST. (2024). The NIST cybersecurity framework 2.0. Available at: <https://doi.org/10.6028/NIST.CSWP.29>.
- NITDA. (2024). Nigeria cybersecurity skills gap assessment report 2024. National Information Technology Development Agency.
- OASIS. (2021). STIX version 2.1: OASIS standard. Available at: <https://docs.oasis-open.org/cti/stix/v2.1>.
- Oladipo, A. F., Abayomi-Alli, O., Misra, S., and Abayomi-Alli, A. (2023). Machine learning for e-fraud detection in Nigerian mobile banking. Heliyon, 9(3), e13807.
- Open Cybersecurity Schema Framework (OCSF). (2023). OCSF schema specification v1.1.0. Available at: <https://schema.ocsf.io/>.
- Sauerwein, C., Lamber, M., and Breu, R. (2018). Threat intelligence sharing platforms: An exploratory study of software vendors and research perspectives. Proceedings of ARES 2018.
- Sharafaldin, I., Lashkari, A. H., and Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. Proceedings of ICISSP 2018, pp.108–116.
- Singhal, A., and Ou, X. (2011). Security risk analysis of enterprise networks using probabilistic attack graphs. IEEE Transactions on Dependable and Secure Computing, 8(6), pp.836–849.
- Strom, B. E., Applebaum, A., Miller, D. P., Nickels, K. C., Pennington, A. G., and Thomas, C. B. (2020). MITRE ATT&CK: Design and philosophy. MITRE Corporation Technical Report MTR190022.
- SWIFT. (2023). SWIFT customer security programme: Security controls framework v2023. SWIFT SCRL.
- 116 allaee, M., Bagheri, E., Lu, W., and Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 data set. Proceedings of the IEEE Symposium on Computational

Intelligence for Security and Defence Applications 2009, pp.1–6.

Team Cymru. (2023). IP-to-ASN mapping service documentation. Available at: <https://www.team-cymru.com/ip-asn-mapping>.

Washington, N. M. (2024). IoT attack detection using a stacked ensemble for multi-class classification. *Journal of Network and Computer Applications*, 228, 103921.

Zhao, J., Yan, Q., Li, J., Shao, M., He, Z., and Li, B. (2023). TIMiner: Automatically extracting and analyzing categorized cyber threat intelligence. *Computers and Security*, 118, 102756.

Zhou, Z. H. (2012). *Ensemble methods: Foundations and algorithms*. CRC Press.