© NOV 2025 | IRE Journals | Volume 9 Issue 5 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I5-1712272

Bridging The Cybersecurity Divide: A Systematic Review of AI And Lightweight Monitoring For Resource-Constrained Institutions

MUGERWA JOSEPH[1], ODO FRANCIS IKECHUKWU[2], KITUMBA DAVID[3]
[1] *Department of Computer Science, Babcock University, Ilishan Remo, Ogun State, Nigeria*
[2] *Department of Computer Science, Babcock University, Ilishan Remo, Ogun State, Nigeria*
[1,3] *Department of Computing and Informatics, Bugema University, Kampala, Uganda*

Abstract: ***The digital transformation has exacerbated a critical cybersecurity divide, systematically excluding low-resource institutions for example schools, NGOs, and SMEs from modern network protection due to profound constraints in budget, hardware, and expertise. While research in network monitoring has advanced, featuring sophisticated AI for anomaly detection and cloud-based orchestration, its applicability to constrained environments remains critically underexplored. This study conducts a systematic literature review of 150 peer-reviewed works (2018-2025) to map the state-of-the-art and evaluate its alignment with the needs of low-resource settings. Our analysis, structured across four thematic areas namely AI Anomaly Detection, Lightweight Monitoring, Cloud Integration, and Configuration Management reveals a pervasive convergence gap. We identify that research excels in technological silos but fails at their intersection: AI models are computationally prohibitive, lightweight solutions lack adaptive intelligence, cloud frameworks assume stable connectivity, and configuration tools operate in isolation. This siloed progression creates a landscape where solutions are either too bulky, too simple, too cloud-dependent, or too complex for practical deployment in target institutions. The primary contribution of this review is the articulation of four core requirements for a new holistic paradigm: (1) Resource-Aware AI, (2) Progressive Intelligence, (3) Intermittent-Cloud Resilience, and (4) a Unified Management Plane. In direct response, we propose the development of a lightweight, AI-cloud-integrated platform as a direct path forward. This review synthesizes a fragmented field and provides a definitive research agenda aimed at bridging the cybersecurity divide through architectural convergence.***

Index Terms: ***Cybersecurity Divide, Resource-Constrained Networks, Lightweight Network Monitoring, AI Anomaly Detection, Cloud-Edge Integration***

I. INTRODUCTION

The digital transformation has irrevocably altered the operational fabric of modern society, embedding network connectivity as a critical utility for organizations of all sizes and mandates. For institutions such as schools, non-governmental organizations (NGOs), and small to medium enterprises (SMEs), this connectivity enables essential services ranging from digital learning and remote healthcare to global outreach and commerce. However, this reliance has concurrently forged a dangerous dependency, creating an expansive attack surface that is increasingly targeted by malicious actors. Within this perilous landscape, a profound and systemic inequality has emerged: a cybersecurity divide that systematically excludes low-resource institutions from the protections enjoyed by their well-funded counterparts. These organizations face a critical paradox, tasked with safeguarding sensitive data and critical digital services while being systematically disqualified from the very tools designed to provide such security due to profound constraints in budget, computational hardware, and technical expertise [1], [2].

The academic and industrial communities have responded to the escalating threat landscape with remarkable innovation, particularly in the domain of network monitoring and intrusion detection. A significant research thrust has focused on leveraging artificial intelligence (AI) and machine learning (ML) to move beyond signature-based detection towards proactive and adaptive anomaly identification. This has yielded a sophisticated arsenal of models, including deep convolutional networks, generative adversarial networks, and transformer-based architectures, which demonstrate exceptional accuracy on benchmark datasets [3], [4]. Parallel to these advancements, the paradigms of cloud computing and software-defined networking promise unprecedented levels of orchestration, scalability, and automation, potentially enabling a new era of intelligent security management [5], [6].

© NOV 2025 | IRE Journals | Volume 9 Issue 5 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I5-1712272

Yet, a critical disconnect persists between this trajectory of technological advancement and the operational realities of low-resource environments. Incumbent monitoring solutions and nascent AI-driven approaches are overwhelmingly architected for an enterprise context, presupposing abundant computational resources, persistent high-bandwidth connectivity, and dedicated IT security personnel. This fundamental architectural mismatch renders them practically unusable in constrained settings, where high computational overhead exceeds the capacity of low-cost hardware and complex configuration processes present an insurmountable barrier to adoption [7], [8]. Consequently, a chasm has formed between research potential and practical deployment. As recent surveys starkly illustrate, the adoption of AI/ML-based security tools is severely hampered by a neglect of critical cross-cutting challenges, including explainability, usability, and computational efficiency in real-world settings [9]. This leaves the most vulnerable institutions with a choice between enterprise tools that are functionally inaccessible and lightweight solutions that lack the intelligent, context-aware detection capabilities necessary to combat modern threats.

Therefore, a systematic consolidation of the current state of knowledge is urgently required to map this fragmented landscape and identify a viable path forward. While individual domains of AI for security, lightweight monitoring, and cloud integration are well-studied in isolation, their synthesis and critical evaluation through the lens of resource constraints is lacking. This systematic review seeks to address this gap by asking: *What is the state-of-the-art in network monitoring and anomaly detection?* and to *what extent are these solutions applicable to the stringent demands of low-resource institutions?* To answer this, we conduct a thorough analysis of recent literature, categorizing research into four interconnected themes: AI-driven anomaly detection, lightweight monitoring strategies, cloud integration and synchronization, and automated configuration management. Through this synthesis, this review aims to conclusively demonstrate that the central challenge is no longer a lack of innovation within technological silos, but a critical gap at their convergence. The ultimate contribution of this work is to articulate the specific requirements for a new class of systems lightweight, intelligently adaptive, and holistically integrated that can finally bridge the cybersecurity divide and empower the institutions that form the bedrock of an equitable digital society.

II. RESEARCH METHODOLOGY

To ensure a comprehensive, unbiased, and reproducible analysis of the literature, this review was conducted following the systematic literature review (SLR) methodology, guided by the principles outlined in the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) statement. The objective was to identify, evaluate, and synthesize all relevant research concerning network monitoring and anomaly detection, with a specific focus on its applicability to resource-constrained environments.

A. Search Strategy

A systematic search was performed across four major electronic databases renowned for their coverage in computer science, engineering, and technology: *IEEE Xplore, Scopus, ACM Digital Library, and Google Scholar.* The search was deliberately scoped to literature published between January 2018 and September 2025 to capture the most recent advancements in AI-driven monitoring and lightweight systems, a period of rapid evolution in these fields. The initial search query was designed to be broad yet focused, structured around key concepts and their synonyms:

Population/Problem: ("resource-constrained" OR "low-resource" OR "low-cost" OR "SME" OR "non-profit" OR "NGO" OR "Raspberry Pi" OR "IoT" OR "edge")

Intervention/Technology: ("network monitoring" OR "network management" OR "NIDS" OR "anomaly detection")

Mechanism: ("machine learning" OR "deep learning" OR "AI" OR "lightweight" OR "cloud" OR "configuration management")

These concepts were combined using Boolean operators, and the query was adapted to the specific syntax of each database. For instance, a typical query structure was: *("network monitoring" OR "NIDS") AND ("lightweight" OR "resource-constrained") AND ("AI" OR "machine learning").* The Google Scholar search was utilized primarily for forward and

backward snowballing, identifying seminal works cited by or citing the papers found in the primary databases.

B. Study Selection and Eligibility Criteria

The retrieved records underwent a multi-stage screening process to identify the most relevant studies. The inclusion and exclusion criteria were defined a priori as follows:

A. Inclusion Criteria:

1. Peer-reviewed journal articles, conference proceedings, and workshop papers.
2. Studies presenting a technical solution, framework, model, or empirical evaluation related to network monitoring, intrusion detection, or configuration management.
3. Studies explicitly discussing aspects of computational efficiency, resource constraints, cost-effectiveness, or deployment in environments with limited IT expertise.
4. Publications within the date range of January 2018 to September 2025.
5. Publications in the English language.

B. Exclusion Criteria:

1. Studies focused solely on enterprise-grade data centers or high-performance computing without consideration for constrained environments.
2. Papers where the full text was not accessible.
3. Short papers (less than 4 pages) without substantial technical detail.
4. Duplicate publications describing the same study; the most complete version was retained.

The study selection process, detailed in the PRISMA flow diagram (Figure 1), was conducted in two phases. First, titles and abstracts were screened against the eligibility criteria. Second, the full text of the remaining articles was thoroughly reviewed to make a final inclusion decision. Any uncertainties regarding a paper's eligibility were resolved through discussion between the authors until a consensus was reached.

C. Data Extraction and Synthesis

From each of the 150 included studies, data was systematically extracted into a standardized spreadsheet to facilitate thematic synthesis. The extracted data included: bibliographic information (authors, title, year, source), research methodology for example simulation, prototype and case study core technological contribution, key findings related to performance and efficiency, and explicitly stated limitations or gaps.

The synthesis was primarily narrative and thematic. Following extraction, the studies were iteratively categorized into four emergent, interconnected themes that collectively address the research problem:

1. *AI Anomaly Detection*: Focusing on the models and their computational demands.
2. *Lightweight Monitoring*: Encompassing strategies for efficient operation on constrained hardware.
3. *Cloud Integration / Sync*: Addressing architectures for management and challenges of intermittent connectivity.
4. *Configuration Management*: Covering tools and methods for automation and integrity.

The analysis within each theme involved identifying trends, comparing approaches, and critically evaluating the applicability of the proposed solutions to the target environment of low-resource institutions. The convergence and gaps between these themes form the core of the discussion in this review.

© NOV 2025 | IRE Journals | Volume 9 Issue 5 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I5-1712272

Figure 1. PRISMA Flow Diagram Illustrating the Study Selection Process

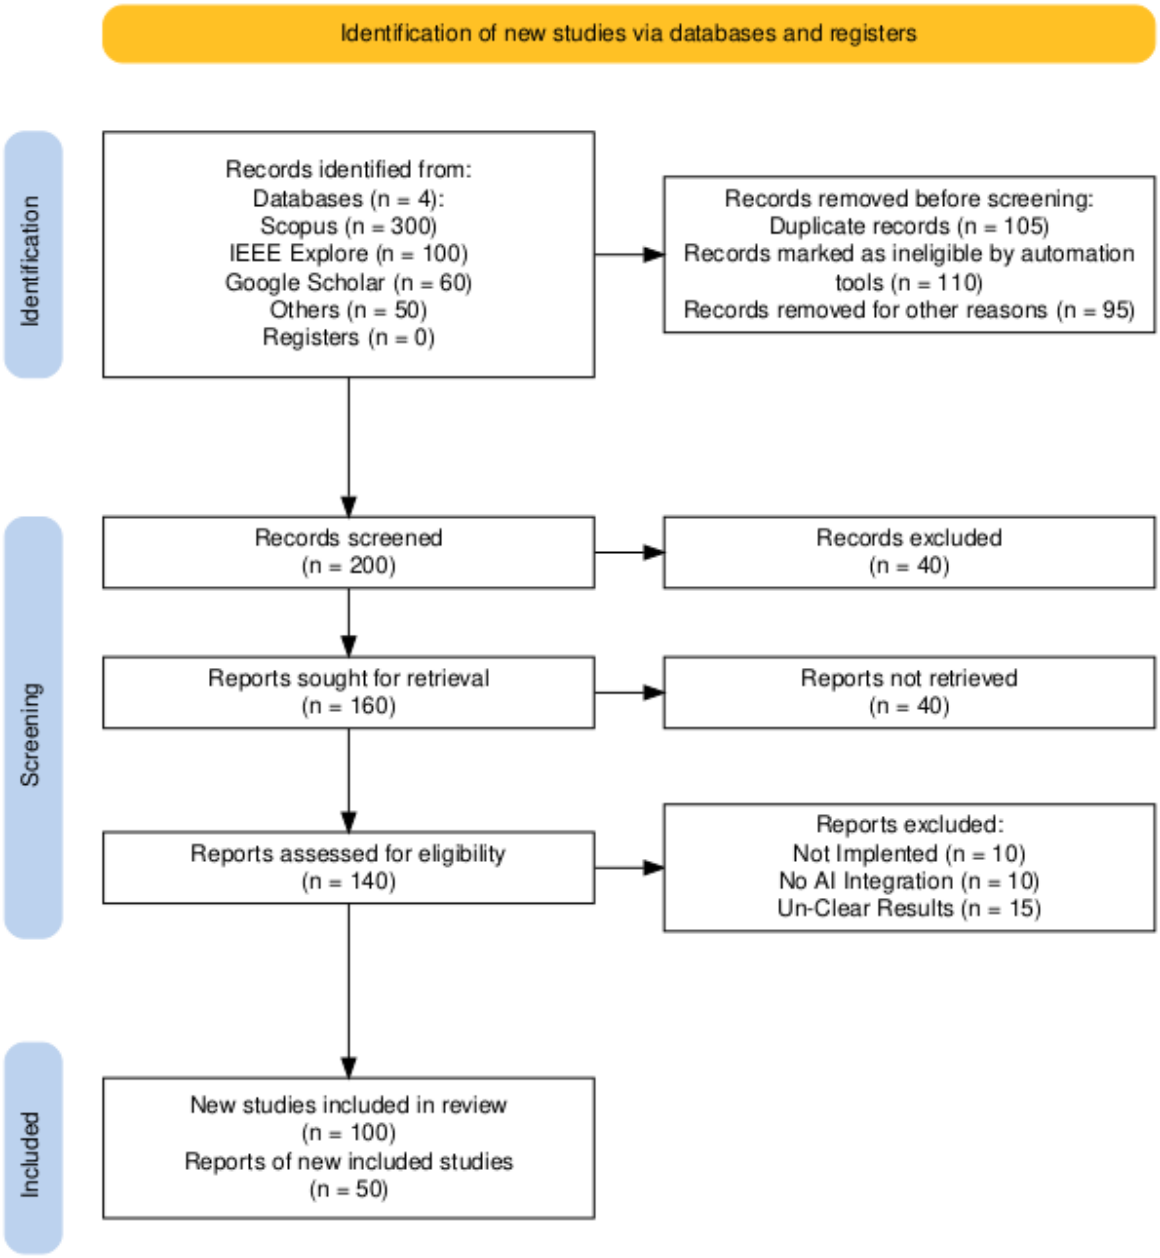

The PRISMA flow diagram outlines the process of identifying, screening, and selecting the 150 studies included in this systematic review, ensuring a transparent and methodologically sound selection process. [10]

III. THEMATIC ANALYSIS AND FINDINGS

A. The Pursuit of Accuracy: AI/ML in Anomaly Detection

The application of Artificial Intelligence and Machine Learning for network anomaly detection represents a dominant and rapidly evolving research frontier, driven by the need to identify sophisticated and novel threats that evade traditional signature-based methods. The literature reveals a clear taxonomy of approaches, each pushing the boundaries of reported performance[11]. Deep learning architectures, particularly Convolutional Neural Networks (CNNs) and Deep Neural Networks (DNNs), are extensively employed for their powerful feature extraction capabilities from high-dimensional network traffic data. Studies such as that by [4] demonstrate that CNNs can achieve near-perfect accuracy on multi-class intrusion detection tasks, significantly outperforming more complex Generative Adversarial Networks (GANs) in direct comparisons. This is complemented by hybrid models, like the CNN-GAN proposed by [12], which use generative models for data augmentation to address class imbalance and further refine detection

rates, achieving accuracies upwards of 99% on benchmark datasets like NSL-KDD.

Beyond supervised classification, the field explores more nuanced architectures. Transformer models, celebrated in natural language processing, are being adapted to capture long-range dependencies in sequential network traffic, offering a potential advantage over recurrent networks [3]. Furthermore, the paradigm of Federated Learning (FL) has emerged as a promising solution to the data privacy challenges in collaborative security. Research by [6] integrates FL with Zero-touch Service Management (ZSM) for B5G networks, demonstrating how multiple actors can collaboratively train an intrusion detection model without sharing raw data, thereby achieving a balanced accuracy while preserving privacy. The pursuit of proactive security has also led to forecasting models, where deep learning architectures like Long Short-Term Memory (LSTM) networks are used to predict future network traffic states, aiming to identify anomalies before they fully manifest [13], [14].

However, a critical synthesis of this body of work reveals a pervasive and troubling gap between technological potential and practical deploy ability, particularly in low-resource contexts. The primary focus of these studies is overwhelmingly techno-centric, with success almost exclusively measured by metrics like accuracy, precision, and recall on static, often curated, datasets. This narrow focus obscures critical operational constraints. As [9] bluntly identify in their user-centric survey of 166 academic papers, there is a systematic neglect of explainability, usability, and practical reproducibility in AI/ML-based Network Intrusion Detection Systems (NIDS). The complex "black-box" nature of deep learning models like Transformers and deep CNNs renders their decisions inscrutable to the non-expert administrators typical of low-resource institutions, undermining trust and actionable response.

Moreover, the computational cost of state-of-the-art models is frequently prohibitive. Training and inferencing with these sophisticated architectures demand significant CPU, memory, and often GPU resources, placing them far beyond the operational envelope of low-cost hardware such as Raspberry Pis. The survey by [15] on enterprise network security explicitly identifies the need for "explainable ML detection" and highlights the computational challenges of deploying advanced AI in dynamic environments. While studies like that of [16] propose conceptual frameworks for AI-based threat detection, they often remain silent on the immense implementation overhead and real-time inference latency on constrained devices. Consequently, this relentless pursuit of incremental accuracy in controlled environments has created a gap, leaving a wealth of powerful AI research largely inaccessible to the very institutions that are in most desperate need of intelligent, automated threat detection.

B. The Imperative of Efficiency: Lightweight Monitoring

In direct contrast to the resource-intensive nature of advanced AI models, a parallel and equally critical research trajectory focuses on the imperative of efficiency, developing monitoring strategies expressly designed for resource-constrained environments. This body of work addresses the fundamental architectural mismatch by prioritizing low computational overhead, minimal energy consumption, and operational feasibility on hardware such as single-board computers (SBCs), IoT devices, and within wireless sensor networks (WSNs)[17]. A significant portion of this research is dedicated to optimizing the performance of standard protocols on low-power devices. For instance, [7] provide a crucial empirical evaluation of SNMP versions on Raspberry Pi, quantifying the significant performance overhead introduced by SNMPv3's security features, a key consideration for practical deployment. Further optimizations are explored through novel data plane programming; [18] demonstrate how eBPF and XDP can be used to create efficient network monitoring applications within the kernel, consolidating primitives to reduce CPU usage by triggering analytics only upon metric anomalies.

The challenge of lightweight monitoring is particularly acute in the realm of the Internet of Things (IoT), where devices are characterized by extreme limitations in processing, memory, and power. In response, researchers have proposed specialized, flow-based Intrusion Detection Systems (IDS) that sacrifice machine learning altogether. [19] presents IoT-FIDS, a system that establishes behavioral baselines for IoT devices and uses the computationally trivial Hamming distance for anomaly detection, achieving high accuracy without

the need for model training. Similarly, studies like that of [20] implement lightweight ML-based IDS directly on constrained switch hardware, demonstrating practical deployment potential despite limitations in detecting certain attack types like MITM. The systematic review by [21] on lightweight ML/DL for IoT networks confirms this trend, finding a dominance of machine learning over deeper models and highlighting the critical role of feature engineering in achieving lightness. Beyond detection, lightweight principles extend to core network functions, [22] proposing LowPaxos, a consensus protocol for state machine replication designed for heterogeneous, low-resource settings, dynamically adapting leader election to fluctuating network conditions.

A critical synthesis of these efficiency-focused solutions, however, reveals a consistent trade-off; the achievement of lightness often comes at the cost of sophisticated intelligence. While these approaches successfully solve the problem of operational feasibility on constrained hardware, they frequently revert to simpler analytical methods. Many solutions rely on static rules, pre-defined signatures, or straightforward statistical baselines. For example, the highly efficient IoT-FIDS [19] and the characteristic functions method for CAN bus intrusion detection proposed by[23] are fundamentally rule-based, offering high speed and transparency but lacking the adaptive learning capabilities of their AI counterparts. This creates a distinct "*intelligence gap*"; a disconnect between the efficient operation of a monitoring system and its ability to understand complex, evolving network contexts and identify novel, multi-stage attacks that do not match a pre-defined rule or a simple statistical deviation. The survey by [21] further underscores this limitation, noting that the field lacks real-world testing and comprehensive reporting of computational metrics, suggesting that even the reported "lightweight" ML models may not be sufficiently intelligent or practical for broad deployment.

C. The Promise and Problem of the Cloud: Integration & Sync

The paradigm of cloud computing promises a transformative leap in network monitoring, shifting the locus of intelligence from isolated on-premises appliances to centralized, scalable platforms capable of sophisticated analytics and global threat intelligence[24]. This research stream explores frameworks that leverage the cloud for enhanced orchestration, management, and collaborative security. A prominent example is the move towards data-plane integrated security in cloud infrastructures, as demonstrated by [5] with OpenStackDP. This framework embeds threat detection and mitigation directly into the data plane of an SDN-based OpenStack cloud, achieving high accuracy with low latency by bypassing control-plane bottlenecks. This vision of centralized intelligence is further refined in industrial contexts, where [25] implement a unified service-oriented architecture using the Arrowhead Framework, enabling dynamic orchestration of IIoT monitoring and control services across the edge-cloud continuum. Such approaches exemplify the cloud's promise: to unify disparate systems, automate complex operations, and provide a holistic view of distributed network assets.

Furthermore, the cloud enables advanced collaborative learning models that are otherwise impossible. The integration of Federated Learning (FL) with Zero-touch Service Management for B5G networks, as proposed by [6], illustrates a future where cloud orchestration facilitates privacy-preserving collaboration between different network operators, collectively improving anomaly detection models without sharing sensitive raw data. Cloud platforms also serve as a central hub for aggregating and analyzing data from distributed lightweight monitors, potentially compensating for their local "intelligence gap" by providing access to more powerful, centralized AI models and global threat feeds.

However, a critical synthesis reveals a fundamental and often unstated assumption underpinning most cloud-integrated frameworks: the presence of a stable, high-bandwidth, and low-latency connection between the edge and the cloud. This assumption constitutes a critical flaw when applied to the context of low-resource institutions, which often operate in regions with intermittent, costly, or low-bandwidth internet connectivity. The very architectures designed to provide resilience and scalability, such as those described by [5] and [25], can become single points of failure or sources of operational paralysis when the network link is unreliable. In such environments, a cloud-centric model that requires constant communication for functionality or real-time

analytics breaks down, rendering the monitoring system ineffective at the moment it may be most needed.

This core problem has begun to be acknowledged in an emerging body of research focused on resilience and abstraction in challenging environments. [26] address this directly with Crane Cloud, an open-source multi-cloud abstraction layer explicitly designed for resource-constrained settings. Their work provides a crucial counterpoint, demonstrating that practical cloud integration must prioritize resilience to network issues and the ability to maintain localized service delivery during cloud outages. This highlights the need for an "offline-first" design philosophy, where the core monitoring intelligence resides locally, and the cloud serves as an optional force-multiplier for intermittent synchronization, model updates, and remote oversight, rather than a central nervous system.

D. The Missing Link: Automated Configuration Management

The integrity of a network's security posture is intrinsically linked to the configuration of its constituent devices; a single misconfigured firewall rule or an unpatched system vulnerability can render all other security measures doubtful [27]. Recognizing this, a significant body of research is dedicated to automating network configuration management to eliminate human error, enforce policy consistency, and enable rapid recovery from failures. This domain is dominated by two primary approaches: agent-based automation tools and software-defined networking (SDN). Empirical evaluations, such as the comparative analysis by [8], systematically assess tools like Ansible, Puppet, and Chef for SDN environments. Their findings highlight a key trade-off: Ansible's agentless, YAML-driven approach offers superior versatility for rapid deployment, while Puppet and Chef provide more robust state enforcement at the cost of higher complexity. This evolution from manual CLI configuration towards declarative, code-driven automation represents a fundamental shift in network operations, or NetDevOps.

Concurrently, the rise of SDN has introduced a paradigm for centralized, programmatic control over the entire network. Research in this vein focuses on leveraging this centralization to simplify and verify configuration. [28] propose an extensible automated troubleshooting framework that integrates the formal NetKAT language with monitoring tools to localize configuration errors and link failures against a specified intent. Earlier foundational work, such as that by [29], demonstrates the operational efficiency gains from automated configuration acquisition, using a unified XML model to integrate physical and virtual network data and reduce operator configuration time by approximately 40%. These approaches underscore a powerful trend: treating the network as a single, programmable entity rather than a collection of independent devices.

A critical synthesis of this field, however, reveals that automated configuration management largely exists as a standalone operational silo, disconnected from the continuous monitoring and anomaly detection processes. The tools and frameworks discussed whether Ansible for automation or SDN controllers for centralization operate in a separate domain from NIDS and performance monitoring platforms like Zabbix or Nagios. This disconnection creates a significant operational burden, particularly for the small, non-specialist teams characteristic of low-resource institutions. An administrator may receive an alert from a monitoring system indicating a performance degradation or a security anomaly but must then context-switch to a separate configuration management tool for example Ansible playbooks or SDN controller to investigate and enact a remediation, such as pushing a new access control list or rolling back a faulty configuration. This fragmentation of tools and interfaces dramatically increases cognitive load and delays incident response.

The lack of a unified management plane means that the powerful feedback loop between detection and remediation remains manual. A monitoring system might detect an anomaly stemming from a configuration drift, but it cannot autonomously trigger a configuration backup or a restorative action. Conversely, a configuration management tool can enforce policy but typically lacks the real-time intelligence to understand if a configuration change has inadvertently caused a performance or security issue. This operational silo is the "missing link" in building a truly autonomous and resilient network management system for constrained environments.

© NOV 2025 | IRE Journals | Volume 9 Issue 5 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I5-1712272

IV. DISCUSSION: THE CONVERGENCE GAP

The systematic and thematic analysis presented in the previous sections reveals a landscape of significant but profoundly siloed innovation. Individually, the domains of AI anomaly detection, lightweight monitoring, cloud integration, and configuration management have each advanced considerably, pushing the boundaries of what is technically possible within their respective scopes. However, when evaluated through the critical lens of applicability for low-resource institutions, a consistent and systemic failure pattern emerges. The central finding of this review is that the research community is excelling in technological silos but is failing at the essential point of *convergence*. The next frontier for research in this field is not a more accurate AI model, a more efficient protocol, or a more robust cloud service in isolation, but rather the development of a new *holistic architectural paradigm* that deliberately unifies these capabilities into a cohesive and accessible system.

This critical gap is synthesized in Table 1 below, which maps the strengths and limitations of each research domain, illustrating how their individual shortcomings collectively create an insurmountable barrier for low-resource institutions.

Table 1. Synthesis of Literature and Identification of the Convergence Gap

Research Domain	Strengths and Focus (From Literature)	Limitations & Gaps (As Applied to Low-Resource Contexts)
AI Anomaly Detection	High-accuracy models (CNNs, GANs, Transformers); Proactive forecasting; Federated Learning [3], [6].	Computationally intensive; poor explainability; neglects usability and deployment constraints [9].
Lightweight Monitoring	Efficient operation on SBCs/IoT; low-overhead protocols (eBPF); resource-aware algorithms [18], [19].	Lacks integrated, adaptive AI; often relies on simple rules/signatures, creating an "intelligence gap."
Cloud Integration / Sync	Centralized management; scalable analytics; service orchestration [5], [25].	Assumes stable, high-bandwidth connections; lacks efficient "offline-first" sync mechanisms for intermittent links [26]
Configuration Management	Automation of complex tasks; reduced human error; SDN-based control [8], [28].	High complexity for setup; standalone tools that are not integrated with monitoring, creating management overhead and operational silos.
Converged Solution (Proposed / Required)	A unified platform integrating resource-aware AI, lightweight operation, resilient cloud-sync, and automated configuration management.	Fills the gap at the intersection of the four domains, transforming isolated capabilities into a synergistic, practical system.

As Table 1 indicates, a solution that merely combines off-the-shelf components from these silos would inevitably inherit their collective limitations, it would be either too bulky, too simple, too dependent on the cloud, or too complex to manage. Therefore, to effectively bridge the cybersecurity divide, a future solution must be architected from first principles around the following core requirements:

1. *Resource-Aware AI Architecture:* Move beyond simply porting large models to the edge. This requires designing AI/ML techniques where computational and memory footprints are primary design constraints, not afterthoughts. This may involve leveraging efficient feature engineering, model compression, or inherently lightweight algorithms that prioritize operational feasibility over marginal gains in accuracy on benchmark datasets.
2. *Progressive Intelligence:* The system must embody a "start simple, grow smart" philosophy. It should begin its lifecycle with highly efficient, transparent baseline establishment and simple anomaly detection, autonomously enhancing its analytical models and contextual understanding as it accumulates operational data. This ensures immediate value upon deployment and continuously adaptive defense without initial configuration complexity.
3. *Intermittent-Cloud Resilience:* The cloud relationship must be redefined from "central-

brain" to "optional force-multiplier." The core system must be fully autonomous and operational with all critical functions executing locally. The cloud component should serve solely for bandwidth-efficient delta-synchronization, aggregated threat intelligence updates, and remote oversight, using robust abstraction layers that gracefully handle prolonged connectivity loss.

4. *Unified Management Plane:* Perhaps the most critical requirement for usability is the dissolution of operational silos. Monitoring, alerting, and configuration management cannot be separate tools. They must be integrated into a single interface where an anomaly alert can be directly correlated with device configuration history and where a one-click remediation (e.g., configuration rollback) is an inherent capability of the platform, not a separate manual process.

V. CONCLUSION AND FUTURE RESEARCH DIRECTIONS

This systematic literature review has provided a comprehensive mapping of the current state of research in network monitoring and anomaly detection, critically evaluated through the lens of applicability for low-resource institutions. The central contribution of this work is the identification and articulation of a critical *convergence gap*, a persistent disconnect between four key domains of innovation: high-accuracy AI, lightweight monitoring, resilient cloud integration, and automated configuration management. While each domain demonstrates significant individual advancement, our synthesis reveals that their isolated progression has inadvertently created a technological landscape where solutions are either too resource-intensive, too simplistic, too connectivity-dependent, or too operationally complex to be viable for schools, NGOs, and small enterprises. This gap perpetuates the dangerous cybersecurity divide, leaving the most vulnerable organizations systematically exposed.

Through this analysis, we have distilled the path forward into a set of four core requirements for a new class of systems designed to bridge this divide: *(1) Resource-Aware AI Architecture* that prioritizes efficiency and explainability; *(2) Progressive Intelligence* that enables systems to start simple and autonomously enhance their capabilities; *(3) Intermittent-Cloud Resilience* that ensures core functionality through an "offline-first" design with efficient synchronization; *and (4) a Unified Management Plane* that integrates monitoring, alerting, and configuration into a single, coherent interface to eliminate operational silos.

In direct response to the findings of this review, our ongoing research is developing a *Lightweight AI Cloud-Integrated Network Monitoring and Configurations Management System*, a platform architected from the ground up to embody these precise requirements. The proposed system will leverage a resource-aware adaptive baseline engine and context-aware anomaly detection, all tightly integrated with automated configuration backup and a resilient cloud-sync layer, deployed on low-cost hardware. Future work, therefore, will focus on the implementation, deployment, and rigorous empirical evaluation of this system to validate its efficacy in closing the convergence gap identified herein. By doing so, this research program aims to translate the synthesized knowledge of this review into a tangible tool, actively working to democratize enterprise-grade network security intelligence and build a more resilient and equitable digital infrastructure for all.

STATEMENTS AND DECLARATIONS

A. Competing Interests

The authors declare that they have no competing interests that could have influenced the conduct of this study.

B. Funding

This study received no external funding. All research activities were self-sponsored by the authors.

C. AI Statement

The authors declare that no generative artificial intelligence (AI) tools were used to generate the manuscript. However, AI-based language assistance (such as grammar and clarity enhancement tools) was utilized solely to improve the readability, coherence, and formatting of the text during manuscript preparation.

All substantive ideas, interpretations, and conclusions are the original work of the authors, who take full responsibility for the content and integrity of the publication.

REFERENCES

[1] W. J. Triplett, “Digital Divide in Cybersecurity,” Cybersecurity and Innovation

Technology Journal, vol. 3, no. 1, pp. 1–8, 2025, doi: 10.52889/citj.v3i1.622.

[2] L. Huey and L. Ferguson, “Another Digital Divide: Cybersecurity in Indigenous Communities,” CrimRxiv, Feb. 2022, doi: 10.21428/cb6ab371.bbf05cbc.

[3] F. M. Anis, M. Alabdullatif, S. Aljbli, M. Hammoudeh, and S. Member, “Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000. A Survey on the Applications of Deep Learning in Network Intrusion Detection Systems to Enhance Network Security”, doi: 10.1109/ACCESS.2017.DOI.

[4] N. R. Al-Milli and Y. A. Al-Khassawneh, “Intrusion Detection System using CNNs and GANs,” WSEAS Transactions on Computer Research, vol. 12, pp. 281–290, 2024, doi: 10.37394/232018.2024.12.27.

[5] P. Krishnan, K. Jain, A. Aldweesh, P. Prabu, and R. Buyya, “OpenStackDP: a scalable network security framework for SDN-based OpenStack cloud infrastructure,” Journal of Cloud Computing, vol. 12, no. 1, Dec. 2023, doi: 10.1186/s13677-023-00406-w.

[6] P. F. Saura, J. M. Bernabé Murcia, E. G. de la Calera Molina, A. M. Zarca, J. B. Bernabé, and A. F. Skarmeta Gómez, “Federated Network Intelligence Orchestration for Scalable and Automated FL-Based Anomaly Detection in B5G Networks,” Computers, Materials and Continua, vol. 80, no. 1, pp. 163–193, 2024, doi: 10.32604/cmc.2024.051307.

[7] E. Gamess and S. Hernandez, “Performance Evaluation of SNMPv1/2c/3 using Different Security Models on Raspberry Pi.” [Online]. Available: www.ijacsa.thesai.org

[8] F. S. Bruschetti, J. Guevara, M. C. Abeledo, and D. A. Priano, “An Empirical Evaluation of Automated Configuration Tools for Software-Defined Networking: A Usability and Performance Perspective,” Ingenierie des Systemes d’Information, vol. 28, no. 5, pp. 1127–1134, 2023, doi: 10.18280/isi.280502.

[9] K. Dietz et al., “The Missing Link in Network Intrusion Detection: Taking AI/ML Research Efforts to Users,” IEEE Access, vol. 12, pp. 79815–79837, 2024, doi: 10.1109/ACCESS.2024.3406939.

[10] N. R. Haddaway, M. J. Page, C. C. Pritchard, and L. A. McGuinness, “PRISMA2020: An R package and Shiny app for producing PRISMA 2020-compliant flow diagrams, with interactivity for optimised digital transparency and Open Synthesis,” Campbell Systematic Reviews, vol. 18, no. 2, p. e1230, Jun. 2022, doi: https://doi.org/10.1002/cl2.1230.

[11] C. Tselios, I. Politis, and C. Xenakis, “Improving Network, Data and Application Security for SMEs,” in ACM International Conference Proceeding Series, Association for Computing Machinery, Aug. 2022. doi: 10.1145/3538969.3544426.

[12] V. S. Rao, R. Balakrishna, Y. A. B. El-Ebiary, P. Thapar, K. A. Saravanan, and S. R. Godla, “AI Driven Anomaly Detection in Network Traffic Using Hybrid CNN-GAN,” Journal of Advances in Information Technology, vol. 15, no. 7, pp. 886–895, 2024, doi: 10.12720/jait.15.7.886-895.

[13] G. Nguyen, S. Dlugolinsky, V. Tran, and A. Lopez Garcia, “Deep learning for proactive network monitoring and security protection,” IEEE Access, vol. 8, pp. 19696–19716, 2020, doi: 10.1109/ACCESS.2020.2968718.

[14] O. Lytvyn and G. Nguyen, “Secure Federated Learning for Multi-Party Network Monitoring,” IEEE Access, vol. 12, pp. 163262–163284, 2024, doi: 10.1109/ACCESS.2024.3486810.

[15] M. Lyu, H. Habibi Gharakheili, and V. Sivaraman, “A Survey on Enterprise Network Security: Asset Behavioral Monitoring and Distributed Attack Detection,” IEEE Access, vol. 12, pp. 89363–89383, 2024, doi: 10.1109/ACCESS.2024.3419068.

[16] M. A. Paracha, S. U. Jamil, K. Shahzad, M. A. Khan, and A. Rasheed, “Leveraging AI for Network Threat Detection—A Conceptual Overview,” Electronics (Switzerland), vol. 13, no. 23, Dec. 2024, doi: 10.3390/electronics13234611.

[17] M. N. Mowla, N. Mowla, A. F. M. S. Shah, K. M. Rabie, and T. Shongwe, “Internet of Things and Wireless Sensor Networks for Smart Agriculture Applications: A Survey,” IEEE Access, vol. 11, pp. 145813–145852, 2023, doi: 10.1109/ACCESS.2023.3346299.

[18] M. Abranches, O. Michel, E. Keller, and S. Schmid, “Efficient Network Monitoring Applications in the Kernel with eBPF and XDP,” in 2021 IEEE Conference on Network Function Virtualization and Software Defined

Networks (NFV-SDN), IEEE, 2021, pp. 28–34. doi: 10.1109/NFV-SDN53031.2021.9665108.

[19] I. Mutambik, "An Efficient Flow-Based Anomaly Detection System for Enhanced Security in IoT Networks," Sensors, vol. 24, no. 22, Nov. 2024, doi: 10.3390/s24227408.

[20] G. Ingletto, N. Xiong, M. L. Ortiz, T. Markovic, A. Egeberg, and P. E. Strandberg, "ANOMALY DETECTION FOR NETWORK TRAFFIC IN A RESOURCE CONSTRAINED ENVIRONMENT," 2023. [Online]. Available: https://github.com/westermo.

[21] G. AL Mukhaini, M. Anbar, S. Manickam, T. A. Al-Amiedy, and A. Al Momani, "A systematic literature review of recent lightweight detection approaches leveraging machine and deep learning mechanisms in Internet of Things networks," Jan. 01, 2024, King Saud bin Abdulaziz University. doi: 10.1016/j.jksuci.2023.101866.

[22] A. Mwotil, T. Anderson, B. Kanagwa, T. Stavrinos, and E. Bainomugisha, "LowPaxos: State Machine Replication for Low Resource Settings," IEEE Access, vol. 12, pp. 91272–91288, 2024, doi: 10.1109/ACCESS.2024.3421582.

[23] Y. Chevalier, F. Fenzl, M. Kolomeets, R. Rieke, A. Chechulin, and C. Krauss, "CYBERATTACK DETECTION in VEHICLES USING CHARACTERISTIC FUNCTIONS, ARTIFICIAL NEURAL NETWORKS and VISUAL ANALYSIS," Informatics and Automation, vol. 20, no. 4, pp. 845–868, Aug. 2021, doi: 10.15622/IA.20.4.4.

[24] N. Cai, Z. Deng, and H. Wang, "An Intelligent Data Flow Security Strategy Model of Cloud-Network Integration," in Communications in Computer and Information Science, Springer Science and Business Media Deutschland GmbH, 2022, pp. 3–27. doi: 10.1007/978-981-19-8285-9_1.

[25] D. Hastbacka et al., "Dynamic Edge and Cloud Service Integration for Industrial IoT and Production Monitoring Applications of Industrial Cyber-Physical Systems," IEEE Trans Industr Inform, vol. 18, no. 1, pp. 498–508, Jan. 2022, doi: 10.1109/TII.2021.3071509.

[26] E. Bainomugisha and A. Mwotil, "Crane Cloud: A resilient multi-cloud service abstraction layer for resource-constrained settings," Dev Eng, vol. 7, Jan. 2022, doi: 10.1016/j.deveng.2022.100102.

[27] M. Fuentes-Garcia, J. Camacho, and G. Macia-Fernandez, "Present and Future of Network Security Monitoring," IEEE Access, vol. 9, pp. 112744–112760, 2021, doi: 10.1109/ACCESS.2021.3067106.

[28] I. Pelle and A. Gulyás, "An extensible automated failure localization framework using NetKAT, Felix, and SDN traceroute," Future Internet, vol. 11, no. 5, May 2019, doi: 10.3390/fi11050107.

[29] H. Okita, M. Yoshizawa, K. Uehara, K. Mizuno, T. Tarui, and K. Naono, "Proposal of Virtual Network Configuration Acquisition Function for Data Center Operations and Management System."