



Information Security Copilot for Cyber Threat Monitoring and Risk Management in Small and Medium-Sized Businesses

Yashowardhan Thakare¹, Rutuja Ghongade², Payal Pawar³,

Radha Mankar⁴, Anisha Kaware⁵, Dr V B Gadicha⁶

B.Tech Student, Department of Computer Science and Engineering, P. R. Pote Patil College of Engineering and Management, Amravati, Maharashtra, India¹²³⁴⁵

Professor and Head, Department of Computer Science and Engineering, P. R. Pote Patil College of Engineering and Management, Amravati, Maharashtra, India⁶

Abstract: Small and Medium Businesses (SMBs) increasingly depend on cloud platforms, email, CRM systems, and other digital services, while limited budgets and cybersecurity expertise increase their exposure to threats. This paper proposes an **AI Security Copilot** for continuous security monitoring in SMB environments. The system collects security data through authorized APIs, analyzes activities and anomalies using AI/ML techniques, prioritizes risks, and provides actionable security recommendations through a centralized dashboard. The proposed approach aims to make cybersecurity monitoring more **accessible, manageable, and effective** for resource-constrained

Keywords: AI Security Copilot, Cybersecurity, SMB Security, Threat Detection, Anomaly Detection, Risk Assessment, Security Monitoring.

I. INTRODUCTION

The rapid adoption of digital technologies has transformed how **Small and Medium Businesses (SMBs)** manage communication, customer relationships, financial activities, cloud resources, and internal operations. Cloud platforms, email systems, CRM applications, and collaboration tools improve business efficiency but also increase the number of systems that need to be secured. The growing dependence on interconnected digital services therefore creates a broader cybersecurity attack surface.

Unlike large enterprises, many SMBs have limited budgets, technical expertise, and dedicated cybersecurity personnel [1][5]. This can make continuous monitoring and effective risk management difficult [1][6]. Security concerns such as weak authentication, unauthorized access, suspicious login activities, phishing, and misconfigured cloud resources require timely detection and response [1][6]. Existing literature demonstrates the growing role of machine learning, AI-assisted security techniques, and security-operation approaches in identifying and responding to cyber threats [2][6][9].

AI and Machine Learning (ML) provide opportunities to analyze security data, detect anomalies, classify threats, and prioritize security events [2][4]. Recent research has also explored the use of large language models for cyber-threat detection and strategic cybersecurity reasoning [7][9]. AI-assisted security platforms such as Microsoft Security Copilot and Google Cloud Security AI Workbench further demonstrate the potential of AI in supporting cybersecurity operations. However, SMBs require solutions that are not only capable of detection but are also **affordable, easy to integrate, and capable of providing understandable recommendations**.

This research proposes an **AI Security Copilot** for SMB environments that acts as a supportive cybersecurity assistant. The system collects relevant security information from authorized business services, analyzes activities and potential vulnerabilities, prioritizes risks based on severity and business impact, and provides actionable recommendations through



centralized dashboard. The proposed approach aims to simplify continuous cybersecurity monitoring and decision-making for SMBs while complementing existing security controls [1][6].

II. OBJECTIVES

The proposed research is aligned with the following objectives associated with an AI-based cybersecurity assistance system for Small and Medium Businesses.

- (1) To design an AI-powered cybersecurity assistant for continuous monitoring of SMB digital environments.
- (2) To collect and analyze security logs, user activities, authentication events, and configuration information from connected business platforms through authorized interfaces.
- (3) To identify potential vulnerabilities and suspicious activities such as weak authentication, phishing indicators, unusual login behavior, and cloud misconfigurations.
- (4) To prioritize detected risks using severity, likelihood, and potential business impact.
- (5) To generate clear security recommendations, alerts, and reports through an interactive dashboard.
- (6) To evaluate the system using threat detection accuracy, false positive rate, response time, and risk prioritization efficiency.

III. LITERATURE REVIEW

Existing research shows that cybersecurity remains a significant challenge for SMBs due to limited financial resources, technical expertise, and cybersecurity personnel. Kandpal et al. [1] highlight organizational and resource-related barriers to effective security adoption, while research on resource-constrained environments shows that computational and deployment limits affect the practicality of AI-based security solutions [5].

Machine learning has been widely investigated for automated intrusion detection. Ahmed and Hamad [2] review machine learning techniques for network-based intrusion detection systems and demonstrate the applicability of ML approaches to identifying malicious network activity. Similarly, Soliman et al. [3] examine machine learning techniques for host-based intrusion detection, highlighting the use of behavioral and system-level information for detecting suspicious activities. These approaches demonstrate the potential of automated detection; however, network- and host-based techniques generally focus on specific security layers rather than providing a unified view of security events. Deep learning techniques have also been investigated for detecting anomalies in security and system logs, providing a basis for automated analysis of large volumes of event data [4]. Security logs provide another important source of information for automated threat analysis. Pithode and Patheja [4] investigate deep learning techniques for log anomaly detection, demonstrating how learning-based methods can be applied to identify abnormal patterns within log data. However, effective log-based analysis can depend on the availability and quality of suitable security data, as well as appropriate preprocessing and computational resources [4][5].

Security monitoring and operational response are also important components of cybersecurity. Perera et al. [6] discuss next-generation Security Operations Center (SOC) approaches that emphasize centralized monitoring, analysis, and response to security events. Such approaches provide a foundation for managing security information from multiple sources, but effective interpretation and response can still require specialized cybersecurity knowledge, which may be difficult for resource-constrained SMBs [1][6].

Recent research has also explored the use of large language models (LLMs) in cybersecurity. Ferrag et al. [7] investigate a lightweight BERT-based approach for cyber-threat detection, demonstrating the potential of language-model-based techniques for security analysis. Jin et al. [9] explore the use of LLMs for strategic cybersecurity reasoning, indicating the potential of language models to transform technical security information into more meaningful cybersecurity insights. However, AI-based cybersecurity systems still require appropriate validation and contextual interpretation before their outputs can be relied upon for security decisions [7][9]. Recent research has expanded cybersecurity analysis toward large language models, including LLM-based cyber-threat detection and security reasoning [7][9]. AI and machine-learning techniques have also been explored for Cyber Threat Intelligence (CTI), supporting proactive identification and analysis of potential cyber threats [8].

AI and machine learning have additionally been investigated for Cyber Threat Intelligence (CTI). Simran et al. [8] examine machine-learning-based approaches for CTI and demonstrate the potential of AI techniques for supporting proactive cyber-threat analysis. Integrating such threat intelligence with security events, risk prioritization, and actionable recommendations remains an important challenge when designing an integrated security-assistance system [6][8].



Existing research provides strong foundations for network intrusion detection [2], host-based detection [3], log anomaly detection [4], security operations [6], AI-based threat detection [7], CTI [8], and LLM-based cybersecurity reasoning [9]. However, these capabilities are often studied separately. SMBs need a unified, lightweight solution that can collect security data, detect anomalies, prioritize risks, and provide clear, actionable recommendations. The proposed AI Security Copilot addresses this gap through a centralized, SMB-focused security-assistance workflow [1][5][6].

TABLE 01: COMPARISON OF EXISTING SYSTEMS

Ref.	Author(s) & Year	Method / Technique	Application	Key Findings	Limitations / Gap
I	Kandpal et al. (2023)	SMB cybersecurity analysis	SME/SMB security	Identifies cybersecurity implementation challenges in SMEs	Does not provide an integrated AI-based security assistant
II	Ahmed & Hamad (2021)	ML-based NIDS	Network security	Reviews ML techniques for network intrusion detection	No automated monitoring; need AI-driven Focuses on network-level detection
III	Soliman et al. (2021)	ML-based HIDS	Host security	Reviews ML approaches for host intrusion detection	Limited cross-platform security integration
IV	Pithode & Patheja (2022)	Deep learning	Log anomaly detection	Demonstrates the use of DL to identify anomalous log behavior	Requires suitable data and model development
V	Kornaros (2022)	Hardware-assisted ML	Resource-constrained environments	Examines ML security under computational constraints	Focused on IoT/resource-constrained environments
VI	Perera et al. (2021)	Next-generation SOC architecture	Security monitoring	Highlights centralized monitoring	Conventional SOC approaches require expertise

The reviewed literature indicates that advanced threat detection, centralized monitoring, AI assistance, and cybersecurity frameworks are already available. However, these capabilities are often distributed across specialized tools, enterprise ecosystems, or individual platforms. SMBs may therefore face fragmented security visibility, higher management effort, and difficulty interpreting alerts. The identified research gap is the need for a lightweight AI-assisted layer that can collect security information from multiple authorized services, prioritize risks, and provide practical recommendations in language that can be understood by users without deep cybersecurity expertise.

IV. LIMITATIONS OF EXISTING SYSTEMS

Existing cybersecurity solutions and research approaches provide effective capabilities for individual security tasks; however, several limitations remain when considering the requirements of small and medium-sized businesses (SMBs). First, many cybersecurity approaches focus on a specific security layer, such as network-based intrusion detection, host-based intrusion detection, or log anomaly detection, rather than providing a unified view of security events across multiple sources.

Second, machine-learning-based security systems generally depend on appropriate datasets, feature engineering, computational resources, and model maintenance. These requirements can create practical challenges for resource-constrained organizations that may not have dedicated cybersecurity teams or sufficient infrastructure.

Third, conventional security monitoring and Security Operations Center (SOC) approaches can provide centralized monitoring and analysis, but their effective operation may require specialized cybersecurity expertise to interpret alerts and determine appropriate responses. For SMBs, simply generating security alerts may therefore be insufficient if users cannot easily understand the severity and recommended action. Fourth, AI-based cybersecurity systems and recent large



language model approaches demonstrate potential for automated threat detection and cybersecurity reasoning, but their outputs still require appropriate validation, contextual interpretation, and human oversight before being used for security decisions.

Finally, AI and machine-learning approaches for Cyber Threat Intelligence can assist with identifying and analyzing threats, but integrating threat intelligence, security events, risk prioritization, and actionable recommendations into a single workflow remains a challenge. These limitations indicate the need for a lightweight and unified security-assistance layer that can combine information from multiple security sources, prioritize risks, and present understandable recommendations for SMB users

V. PROPOSED AI SECURITY COPILOT

The proposed AI Security Copilot is designed as a layered cybersecurity assistance system for Small and Medium Businesses. It does not assume that a single source contains all relevant security information. Instead, the system collects selected data from authorized business services and processes it through a common security analysis workflow. The main output is a prioritized set of risks and recommendations that can be viewed by administrators through a dashboard.

A. Data Collection Layer

The data collection layer is responsible for securely obtaining security-relevant information from connected services. Possible sources include authentication logs, user activity records, security alerts, email-related events, configuration information, endpoint signals, and cloud security settings. Integration is proposed through authorized APIs and service-specific connectors. Only information required for security analysis should be collected, and access should follow least-privilege principles.

B. AI Security Analysis Engine

The AI Security Analysis Engine processes collected data to identify suspicious patterns and potential vulnerabilities. The proposed analysis can combine rule-based checks with machine learning techniques. Rule-based analysis is useful for known conditions such as missing MFA or risky configuration settings. Machine learning or anomaly detection can be used to identify unusual login behavior, abnormal activity sequences, or deviations from expected patterns.

C. Risk Prioritization

Not every security event requires the same level of attention. The proposed system assigns a risk priority by considering threat severity, likelihood, asset importance, exposure, and potential business impact. A conceptual risk score is expressed as: $\text{Risk Score} = f(\text{Severity}, \text{Likelihood}, \text{Business Impact}, \text{Exposure})$. The score is then mapped to Critical, High, Medium, and Low categories so that SMB users can focus first on issues that may create the greatest potential harm.

D. Recommendation and Alert Module

For each prioritized risk, the system generates an actionable recommendation. Recommendations identify the issue, explain the potential impact, suggest a corrective action, and indicate the priority. Examples include enabling MFA, reviewing suspicious login activity, changing weak credentials, reducing excessive permissions, correcting a cloud configuration, or reviewing a phishing-related alert.

E. User Dashboard

The dashboard provides a consolidated view of the SMB security posture. Proposed components include an overall risk score, counts of open issues by severity, recent suspicious events, recommended actions, trend indicators, and security status summaries. The design emphasizes clarity and prioritization so that users can understand what requires attention without reviewing raw security logs from every connected platform.




Fig 1: AI Security Copilot for Small and Medium Businesses (SMBs)

VI. RESEARCH METHODOLOGY

The research follows an analytical and experimental design. The analytical component defines the architecture, security workflow, data requirements, and risk prioritization logic. The experimental component evaluates the proposed analysis using simulated and publicly available cybersecurity data where appropriate. The methodology is organized into data



collection, preprocessing, security analysis, risk scoring, recommendation generation, dashboard presentation, and evaluation.

- (1) Connect authorized business services and obtain selected security-related data.
- (2) Validate, normalize, and preprocess the collected records.
- (3) Apply rule-based checks and anomaly or machine learning analysis.
- (4) Detect vulnerabilities, suspicious events, and potential threats.
- (5) Assign risk scores using severity, likelihood, exposure, and business impact.
- (6) Generate prioritized recommendations and alerts.
- (7) Present results through a dashboard and maintain an auditable history of actions.
- (8) Evaluate the system using defined performance metrics.

A. Data Preparation and Feature Engineering

Security data collected from different sources can have different formats. Therefore, preprocessing includes timestamp normalization, removal of duplicate records, handling missing values, categorical encoding where required, and extraction of security-relevant features. Possible features include login time, location change, failed login count, device or IP changes, authentication status, user privilege level, alert type, configuration state, and event frequency. Feature selection should focus on variables that contribute meaningfully to detection and risk prioritization.

B. Model and Analysis Strategy

The proposed work uses a hybrid strategy. Deterministic rules are applied to known security requirements and configuration checks. Machine learning is proposed for patterns that are difficult to describe using fixed rules. Depending on the available dataset, supervised classification can be evaluated for labeled threat data, while anomaly detection can be evaluated for previously unseen or unusual activity. The final model selection should be based on validation results, computational requirements, explainability, and suitability for SMB deployment rather than accuracy alone.

C. Technologies and Experimental Environment

- (1) Programming language: Python for data processing, analysis, and backend services.
- (2) API integration: REST APIs and authorized service connectors.
- (3) Machine learning: scikit-learn and related Python libraries for classification and anomaly detection.
- (4) Backend services: Python-based web APIs for analysis and dashboard integration.
- (5) Database: a relational or document-oriented database for security events, users, recommendations, and audit information.
- (6) Deployment environment: Windows or Linux for development and testing, with Linux-based environments suitable for server deployment.

VII. EXPERIMENTAL SETUP AND PERFORMANCE EVALUATION

The proposed AI Security Copilot is evaluated using simulated and real-world cybersecurity datasets, including security logs, user activity records, and network traffic data. Different cyberattack scenarios, such as phishing attempts, unauthorized access, repeated failed authentication, suspicious login activities, missing MFA, abnormal access patterns, and cloud misconfigurations, are tested to assess the system's ability to detect threats and generate appropriate security recommendations. The experimental environment may include a Windows or Linux development system, Python, a REST API service, a database, and a browser-based dashboard.

Scenario	Security Event	Evaluation Goal
S1	Repeated failed logins attempts	Identify potential brute-force attacks, credential misuse, or unauthorized access attempts.
S2	Anomalous login behavior	Detect unusual authentication patterns based on location, device, access time, or other contextual indicators.
S3	Missing MFA	Identify user accounts that lack MFA and require stronger authentication controls.
S4	Phishing-related security event	Detect and prioritize suspicious emails, links, or user-reported phishing indicators.
S5	Cloud configuration misconfiguration	Identify potentially insecure cloud configuration conditions that may expose organizational resources or data.



S6	High-risk user activity	Prioritize potentially harmful user activities based on security severity, contextual risk, and potential business impact.
----	-------------------------	--

TABLE 02: TEST SCENARIOS

A. Performance Evaluation Metrics

- (1) Accuracy: proportion of correctly classified events among all evaluated events.
- (2) Precision: proportion of detected threat events that are actually relevant threats.
- (3) Recall: proportion of actual threats successfully detected.
- (4) F1-score: harmonic mean of precision and recall.
- (5) False Positive Rate: proportion of benign events incorrectly classified as threats.
- (6) Response Time: time required to process an event and produce a result.
- (7) Risk Prioritization Efficiency: ability of the system to rank more important issues above lower-risk issues.

B. Evaluation Equations

$$\text{Accuracy} = (\text{TP} + \text{TN}) / (\text{TP} + \text{TN} + \text{FP} + \text{FN})$$

$$\text{Precision} = \text{TP} / (\text{TP} + \text{FP})$$

$$\text{Recall} = \text{TP} / (\text{TP} + \text{FN})$$

$$\text{F1-score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$$

$$\text{False Positive Rate} = \text{FP} / (\text{FP} + \text{TN})$$

VIII. LIMITATIONS AND RECOMMENDATIONS

The proposed system has limitations related to the availability of real-time security data, API permissions, cross-platform integration, data quality, false positives, false negatives, privacy, and model maintenance. Machine learning performance depends on the quality and representativeness of the available security data. Therefore, the system should be periodically validated and improved before production-scale claims are made.

- (1) Use least-privilege access and secure credential storage for all connected APIs.
- (2) Apply encrypted communication, audit logging, and role-based access control.
- (3) Minimize the collection of sensitive raw data and define retention rules.
- (4) Combine rule-based checks with machine learning to improve explainability and coverage.
- (5) Evaluate the prototype using representative datasets and realistic SMB security scenarios.

IX. CONCLUSION

This paper proposes an AI Security Copilot designed to support Small and Medium Businesses in monitoring, analysing, and responding to cybersecurity threats. The proposed architecture integrates authorized data collection, security analysis, risk prioritization, recommendation generation, alert management, and centralized visualization into a unified assistance layer. The literature review highlights that existing cybersecurity frameworks, intrusion detection techniques, AI-based security solutions, and platform-specific monitoring tools offer valuable capabilities, but their fragmented nature and operational complexity can create challenges for resource-constrained organizations. The proposed approach aims to address this gap by transforming diverse security information into prioritized, explainable, and actionable guidance. Future work will involve prototype implementation, integration with selected business platforms through authorized APIs, evaluation of appropriate machine learning and anomaly detection techniques, enhancement of recommendation explainability, and controlled experimentation using representative security datasets and test scenarios. The effectiveness, reliability, and practical value of the proposed Security Copilot will ultimately be established through quantitative evaluation and measured experimental results.

REFERENCES

- [1]. S. Kandpal et al., "Cyber Security Implementation Issues in SMEs and Their Potential Solutions," in *Proc. ICCCNT*, 2023, pp. 1–5, doi: 10.1109/ICCCNT56998.2023.10307363.



- [2]. L. A. H. Ahmed and Y. A. M. Hamad, "Machine Learning Techniques for Network-based Intrusion Detection System," in *Proc. NCCC*, 2021, pp. 1–7, doi: 10.1109/NCCC49330.2021.9428827.
- [3]. K. Soliman et al., "Survey of Machine Learning HIDS Techniques," in *Proc. ICCES*, 2021, pp. 1–5, doi: 10.1109/ICCES54031.2021.9686138.
- [4]. K. Pithode and P. S. Patheja, "A Study on Log Anomaly Detection using Deep Learning Techniques," in *Proc. ICAAIC*, 2022, pp. 290–298, doi: 10.1109/ICAAIC53929.2022.9793238.
- [5]. G. Kornaros, "Hardware-Assisted Machine Learning in Resource-Constrained IoT Environments for Security," *IEEE Access*, vol. 10, pp. 58603–58622, 2022, doi: 10.1109/ACCESS.2022.3179047.
- [6]. M. Conti et al., "Internet of Things Security and Forensics: Challenges and Opportunities," *Future Generation Computer Systems*, vol. 78, pp. 544–546, 2018, doi: 10.1016/j.future.2017.07.060.
- [7]. M. A. Ferrag et al., "Revolutionizing Cyber Threat Detection With Large Language Models," *IEEE Access*, vol. 12, pp. 23733–23750, 2024, doi: 10.1109/ACCESS.2024.3363469.
- [8]. Simran et al., "The AI Shield and Red AI Framework: Machine Learning Solutions for Cyber Threat Intelligence," in *Proc. ISCS*, 2024, pp. 1–6, doi: 10.1109/ISCS61804.2024.10581195.
- [9]. J. Jin et al., "Crimson: Empowering Strategic Reasoning in Cybersecurity through Large Language Models," in *Proc. ICCBD+AI*, 2024, pp. 18–24, doi: 10.1109/ICCBD-AI65562.2024.00011.