FOR LENS-AI DRIVEN LIGHTWEIGHT ENDPOINT SECURITY FOR REAL TIME THREAT DETECTION IN SMES

25-26J-076

Project Proposal Report

Sanjula E A Y – IT22340078
Supervisor: Mr. Kanishka Yapa

Bsc (Hons) in Information Technology Specializing in Cyber Security

Department of Information Technology

Sri Lanka Institute of Information Technology

Sri Lanka

August 2025

FOR LENS-AI DRIVEN LIGHTWEIGHT ENDPOINT SECURITY FOR REAL TIME THREAT DETECTION IN SMES

25-26J-076

Project Proposal Report

Sanjula E A Y – IT22340078
Supervisor: Mr. Kanishka Yapa

Bsc (Hons) in Information Technology Specializing in Cyber Security

Department of Information Technology

Sri Lanka Institute of Information Technology

Sri Lanka

August 2025

Declaration of the Candidate & Supervisor

I declare that this is my own work and this proposal does not incorporate without acknowledgement any material previously submitted for a degree or diploma in any other university or Institute of higher learning and to the best of our knowledge and belief it does not contain any material previously published or written by another person except where the acknowledgement is made in the text.

Student Name	Student ID	Signature
Sanjula E A Y	IT22340078	

The above candidates are carrying out research for the undergraduate Dissertation under my supervision.

Signature of the supervisor:

Date : 28/08/2025

Abstract

Small and Medium-sized Enterprises (SMEs) are an easy target for cyber criminals as when starting and running a business, no one considers cyber security as a requirement for running a business. They lack the technical and financial resources to deploy advanced defensive tools. Existing enterprise-grade SOAR (Security Orchestration, Automation and Response) platforms are powerful but often too costly and complex for these smaller organizations. This project proposes the development of a **Lightweight AI-Driven Endpoint Security Solution for Real-Time Threat Detection in SMEs** where this document focus on the **Security Alert Interface with SOAR-Style Automation** which is a subcomponent of the main project. The solution will feature a user-oriented dashboard that consolidates alerts from multiple sources, applies machine learning–based prioritization, and generates actionable recommendations. To support rapid mitigation, the system will embed automated response workflows for enrichment, triage, and remediation, while also integrating seamlessly with communication channels already common in SMEs (e.g., email, Slack, Microsoft Teams). A distinguishing element of the design is an adaptive response component that learns from past user interactions, progressively reducing manual input requirements. By providing automation without excessive complexity, the project aims to balance accessibility and effectiveness, contributing practically to SME resilience.

Keywords: SME Cyber Defense, Endpoint Protection, XDR, SOAR, Automation, Artificial Intelligence

Contents

List of Tables

Table of Figures

Introduction

The trend towards rapid digitalization has increased cyber risks for SMEs due to higher digital process which increased the attack surface of SMEs. Technology has become a major part in SMEs as around 90% depend on digital technologies for their operations. From sensitive customer data to financial data are stored and transferred across multiple distributed networks. When compared to large enterprises and corporations SMEs do not possess expertise and resources to operate Security Operations Center (SOC)'s. Though SMEs might receive less cyber threats than others which can tarnish the need for a SOC, this is not always true. SMEs are becoming extremely attractive to cyber criminals as they are aware that these companies lack the safeguards and know how to respond to incidents. One incident can lead to closure of business in SMEs which highlights how critical and vulnerable these companies are towards cyber threats. SMEs are tend to use multiple technologies and systems that are not interconnected at all which result in decentralized security which can lead to "alert fatigue" [5].

SOAR platforms were created to address challenges that enterprises face. SOAR platforms has improved SOC operations by automating incident response and improving efficiency in security operations. However, most commercial offerings, such as Palo Alto's Cortex XSOAR or Splunk Phantom are resource-heavy, requires skilled analysts and is technically complex [6]. This highlights a question: How can SMEs benefit from automation and orchestration without enterprise level complexity and cost?

Furthermore, the proposed solution will house an alert interface just like an enterprise grade Security Information and Event Management (SIEM) and capabilities of Extended Detection and Response (XDR) platforms to get a holistic view of threats. SOAR, XDR and SIEM are used for distinct purposes and are bundled together to improve incident management and threat detection [8]

The proposed project aims to undertake the challenge by developing a simplified AI-driven alert interface with SOAR-style automation which is used as a component of the overall solution of building a endpoint security solution. It emphasizes ease of deployment, intuitive visualizations, and adaptive response mechanisms to ensure that SMEs can strengthen security postures without requiring specialized staff.

Background and Literature Survey

Over the past 10 years, SOAR technologies play a major and vital role in security operations through automations in large organizations. These platforms support a wide range of integrations with multiple security products and have ability to build playbook to fast track the incident response times. Such tools such as Splunk and Cortex XSOAR have many attractive and valuable features but these products are targeted to large enterprises who possess dedicated security operations center and dedicated specialists in the field [1]. These platforms require significant investments and requires dedicated personal who possess the know how to integrate with other solutions and create workflows. For smaller organizations such as SMEs this will be impractical due to limited resources and budgetary constraints.

Research shows that "Over 40% of SMEs have experienced cyberattacks", "75% would struggle to survive a ransomware attack" and "Many SMEs underestimate their risk, leading to low investment in cybersecurity" as quoted in recent research [11]. This shows how vulnerable is SMEs in the face of cyber threats. In the face of cybersecurity threats such companies are helpless despite how vital they are to a country's economy. It is suggested that SMEs should invest on cybersecurity tools and training, conduct risk assessments, create incident response plans and make security awareness program to build a cyber aware culture.

A report published in the year 2021 by European Union Agency for Cyber Security highlights the issues most SMEs face in terms of cyber security. When compared to larger firms, SMEs have smaller budgets, small IT teams and no proper centralized solutions or coordination across cyber security solutions leading to a fragmented approach [2]. Their findings highlight the need for an affordable, lightweight platforms that reduce operational overheads. If they had a SME friendly cyber security solution most threats could have been prevented completely or up until controllable level thereby reducing the overall impact.

As the world is moving towards Artificial Intelligence it create new opportunities for cyber security. Cybersecurity expertise is more affordable than before as advice can be received through a single prompt. AI also helps in the decision making and increasing efficiency. Research also highlight the concerns in SMEs operating in low poverty regions in the world which are prone to regional conflicts and financial problems. Research conducted on this subject shows that AI will bring in more opportunities to SMEs to employ cost effective solutions, proactive defense and automated incident response. These researches believe that cybersecurity will become more accessible with the development of AI as cyber security is becoming more accessible [12]. The development and evolution of Artificial Intelligence is changing the security landscape. Machine learning algorithms are vastly contributing in reducing false positives, grouping alerts and dynamically prioritizing incidents [3]. These developments have led to increased accuracy than traditional rule-based systems

Human-in-the-loop systems is more recognized in the cyber security space as complete automation may look appealing but, human judgement remains critical in specific scenarios and business critical incidents. Past studies that have been conducted on human AI interaction show

that machine learning systems that learn from user feedback can help reduce repetitive tasks while also maintaining oversight on those systems [4]. In the SME scenario most of the incident handling, decision making and user response analysis should be done by the system itself while also having a human involvement for critical decisions. As time goes the ML algorithms will learn more about the organization which will increase efficiency, accuracy and build trust.

There is research conducted to sort the issues small and medium enterprises face in implementing cyber security. One of the most prominent way is building playbooks that suite such companies. Playbooks help SMEs to reduce response times, minimize disruption and maintain compliance. Playbooks can be customermized to address the lack of dedicated security teams [10]. Playbooks are designed to help the SME leaders an IT staff to respond to cyber threats.

Although modern SOAR platforms provide many benefits, their complexity remains a significant challenge for smaller organizations. Studies show that effective implementation often requires a basic level of security knowledge and internal expertise. This includes an understanding of the organization’s security status and established incident response plans. The need for special skills, such as knowledge of scripting languages like Python for successful integration, usually drives up the total cost of ownership beyond what most budget-conscious SMEs can afford. This issue goes beyond just costs; it also affects operational readiness. Moreover, traditional SOAR platforms are typically stand-alone solutions that require time-consuming, manual integration with existing security systems. This leads to a kind of "integration debt" that smaller teams find hard to manage. A truly SME-friendly solution should be built from the ground up to address these challenges rather than being a stripped-down version of an enterprise-level product.

Research Gap

Existing SOAR solutions are built targeting larger enterprises and do not tailor such systems for SMEs. These solutions are not SME friendly due to limitations of usability and affordability. SOAR platforms have shown strong automation capabilities but, there is limited research conducted to make these systems more accessible to SMEs. ForLens aims to fix this gap by introducing SOAR architectures that can be adapted to SME organizations which have limited budgets, smaller IT teams and lower technical capacity. Researchers have also found that an overwhelming number of alerts and false positives reduces the effectiveness of security teams. Current research shows that to solve this issue and to prioritize alerts very resource heavy models and large datasets are used which is not suitable for SMEs. There is less research conducted on centralized monitoring of security tools for SMEs which result in them to choose fragmented security tools which are not integrated with each other. There is room to conduct further research on this and bridge this gap by introducing a simple, unified dashboard that is capable enough to normalize and correlate heterogeneous data in resource constrained organizations.

Another gap is seen in artificial intelligence adoption for security operations. Current research highlights the need for security solutions that is adaptive in nature. These systems can learn from user feedback to improve efficiency by automating manual and redundant tasks. It also highlights that these systems should operate under human oversight. To understand and supervise these systems skilled analysts are required which is a scarce resource in SME environment. There is room to research on more resource friendly AI algorithms which require less resources. Resource optimization can be done by building more efficient architectures and algorithms for AI and Machine Learning (ML) algorithms. An example of such optimization is visible in research on Deepseek and Chatgpt. Research shows that Deepseek consumes less amount of energy compared to Chatgpt as it offers scalable architecture while the other requires large amount of computational resources, despite both operating in transformer-based models [7].

Adaptive models are a necessity for behavior-Driven workflow adaption. In Enterprise level solutions the feedback loop which is “adaptive” is human centric. A human analyst is required to take certain decisions in contexts such as geopolitics and business risks that automations cannot. Even if SMEs procure the most cost-effective SOAR solution it still requires human expertise to give feedback and refine automations. However, how can SMEs use the feedback loop effectively remains a research gap. The proposed solution must discover a mechanism to operate with minimal human intervention and feedbacks from non-expert users. The system should be capable enough to translate subjective human statements from users to automated workflows and make subjective decisions. Research has been already conducted on reducing human intervention through self-adaptive systems which require the least human feedback [15]. However, there is not enough research done to validate it for SMEs as well as large scale enterprises.

Traditional systems are built and designed in a way that it requires expert knowledge to operate or to assess feedback [14]. There is limited research on building SOAR solutions omitting the requirement or need for expert knowledge, instead to develop a solution that provides a sense of control and transparency without demanding expertise. For example, the proposed solution can handle 95% of incidents autonomously while presenting non-technical explanations and insights with actionable steps to remaining 5% of incidents, where a user can approve in a single click. This approach will retain efficiency and autonomy while maintaining user trust.

The increasing adoption of explainable AI (XAI) has created a significant "trust gap" between AI generated decisions and understanding. Larger machine learning algorithms are not transparent and does not show how certain activities were flagged as malicious (what lead to the decision). An expert personal might have to dig deeper to analyze the artifacts in the incidents to determine what lead to the AI's decision. There needs to be more adequate research conducted to build more transparent models that are fully transparent that shows the thinking pattern of the AI. Further, XAI should generate less technical responses like instead of "unusual API call pattern" it can explain as "This file is behaving like ransomware, which could lock your business's files. Our system has automatically isolated it". It can be improved to include financial estimates if the risk was not addressed what repercussions the organization has to face and provide recommendations and educate users. While doing this the solution should be able to show how it led to taking actions, how financial values were determined and what led to recommendations. The proposed solution should be transparent in the decision making so that the SMEs understand what caused or led to the decision in fully automated scenarios.

Decentralized Security architectures are gaining attention in the market, these systems focus on moving data processing from cloud environments and data centers into local devices. This will greatly benefit to automate alerts and incident handling even during a network outage. This approach leads to faster threat detection, reduced latency and autonomous operations when the device in an unmanaged state (disconnected from network). There needs to be adequate research conducted on how to protect devices from risks inherent to decentralized approaches. The systems should be tamper resists and should be secure and self-heal. The proposed solution should be able to run the SOAR automations and incident resolution actions even at a time of network outage.

Research has been conducted of deep Reinforcement Learning (RL) which uses human feedback and preferences to create a reward function which is used to train RL agents. These RL algorithms one of the best ways to train the AI to act as a human and work on practical real-world scenarios. There are many novel ways that can be used to do this. Some research suggests that the agent should be allowed to perform actions in the environment and humans to view or watch which behavior they like; this input is used by a reward predictor that can learn from the preferences of the user to maximize the predicted reward which makes the agent get better overtime [16]. However, in critical cybersecurity solutions how can this be done? These systems are critical and cannot tolerate errors which will result in major consequences. So how this training mechanisms can be done in such systems remains as a major gap in research.

Lastly, most of research and latest technologies that use of AI is mostly accessible to larger enterprises, there is less research on solutions that focus on SMEs.

While AI-enhanced SOAR solutions exist, few attempts explicitly tailor such systems for SMEs, especially in ways that emphasize usability, affordability, and adaptive learning. Below table shows some research comparisons involving existing research gaps.:

Research Gap	Existing Enterprise SOAR Solutions	Other SME Security Tools (e.g., EDR, AV, basic SIEM)	Proposed Lightweight AI-Driven SOAR for SMEs
Complexity vs. SME Capacity	Assume mature SOCs with skilled operators; high complexity and cost.	Simpler tools, but lack orchestration or end-to-end automation.	Designed for SMEs: lightweight, lower cost, minimal expertise required.
Alert Overload & Prioritization	Provide advanced correlation but require large datasets and tuning; not SME-friendly.	Generate large numbers of alerts without intelligent prioritization.	AI-driven prioritization with efficient models adapted for limited resources.
Fragmented Tools & Integration	Integration possible but resource-intensive; SMEs struggle with implementation.	Operate in silos; weak interoperability, no unified dashboard.	Normalizes and correlates alerts from multiple sources into one dashboard.
Rigid Automation vs. Adaptive Learning	Playbooks are rigid, limited adaptation; customization requires expertise.	Automation minimal or non-existent; heavily manual processes.	Incorporates human-in-the-loop adaptive learning to refine workflows.
Usability & Analyst Experience	Dashboards designed for trained analysts; overwhelming for SME staff.	Interfaces somewhat easier but limited insights and contextual guidance.	User-friendly dashboard with clear risk scores, actions, and automation options.
Evaluation Metrics (Analyst-Centric)	Focus on technical metrics (accuracy, coverage), not operational workload or ROI.	Metrics rarely assessed; SMEs rely on incident counts or vendor reports.	Evaluation based on analyst workload reduction, MTTR, usability, and cost-benefit.

Table 1: Previous research comparison

Research Problem

The research problem is made up of the mismatch between the needs of SMEs and the capabilities of existing cybersecurity solutions. The main goal is to build a “Security Alert Interface” that offer the advanced capabilities of SOAR while also maintaining minimum to no interaction from engineers. The problem can be broken down into 3 main problems.

The usability gap remains a significant problem. SOAR platforms are made for large enterprises that have dedicated SOCs and expertise teams. SMEs do not have specialized teams and know how to manage such platforms [13]. The problem is how can SOAR be simplified with actionable and non-technical guidance for users with little to no security background. To do the translation from technical concepts to business context a new approach is needed to make use of XAI to solve the problem. To explain further if a business user is given a SOAR platform like XSOAR offered by Palo alto , the user would not be able to do anything as the user has no technical knowledge.

The proposed solution heavily relies on automations for threat handling and incident resolutions with minimal human intervention. If near full automation is to be employed the user should have adequate trust with the AI to handle incidents resolutions and playbook management. To build trust the AI needs to be transparent enough to show its internal logic and decision points involving critical actions. The problem is how this trust be created and how transparent the AI should be to build the trust between non-technical user and system automation and how user can prevent false positives that could affect business operations.

The next problem statement is SMEs operate in “high noise” environments as they use unmanaged devices and legacy tools. Traditional SOAR platforms require expertise and is costly to integrate as they are complex in nature. The SOAR platform needs to be user friendly in a way that it does not require help from an expert engineer to perform the integrations. The research problem involves how a solution can be build which is simpler to integrate and also “self-tuning” as it requires to adapt to data, responses and alerts from multiple integrated sources. Automation and playbooks execution will be different from one source to another.

There are plenty of research conducted in using open source solutions for SOAR and SIEM solutions. One of the most loved and famous solution being Wazuh. Wazuh remains at the top when comparing similar solutions in the market as it has the most security features and the best performance. However, research have raised concerns that these open source tools are no where near to enterprise grade solutions [13]. This research will try the best to utilize these existing open source tools to create a better version that is in par with enterprise grade solutions and are more SME friendly.

The ultimate research problem is how can the above 3 gaps be solved from the proposed solution that will handle the backend complexity and provide user with a readable actionable dashboard.

Objectives

Main Objective

To design and implement a lightweight, AI-driven endpoint security solution that is tailored for SMEs to detect and respond to threats in real time. The solution will address pain points that SMEs face with advanced features and novelties that are not found in enterprise grade solutions. Thus, uplifting SME security beyond enterprise grade protection. The proposed 'Security Alert Interface' for the novel endpoint protection solution will bring in enterprise grade capabilities with additional novel capabilities.

A very thorough market and academic research shows development gaps in existing solutions that expect companies to have technical security teams and higher budgets which indicate that existing platforms are not SME friendly. Once these pain points are addressed SMEs can focus on growth and not worry about threats and risks that will slow them down.

Specific Objectives

1. To build a dashboard interface that unifies and normalizes security alerts from endpoints, firewalls, cloud services and other sources. It will act as a single pane of glass bring security alerts from multiple sources together into an intuitive dashboard.

2. To build a machine learning model to prioritize alerts based on severity and relevance. The model will be able to correlate these alerts across multiple sources and past data to understand the context of threats and impact level on the organization.

3. To create automated response workflows for enrichment, triage, and remediation. The Machine learning module will learn from threats to business as well as analyze market attack trends to create custom workflows by itself.

4. Design an adaptive response learning mechanism to reduce repetitive manual intervention.

5. Evaluate the system using SME-relevant metrics, such as alert reduction, usability, and mean-time-to-respond (MTTR).

6. Create a model that can read user responses and understand business contexts to resolve and act upon incidents and alerts without intervention from an analyst.

Methodology

This investigation utilized a qualitative research design, organized around a multi-step methodology to guarantee a thorough and justifiable comprehension of the SME cybersecurity environment and the creation of an innovative solution. The strategy was intended to initially assess the issue area, subsequently to consolidate findings into a conceptual framework, and ultimately to delineate a distinct route for assessment.

As the starting stage a thorough examination of academic publications, industry reports, and market analyses to pinpoint the shortcomings of conventional enterprise-level security frameworks when implemented in the SME sector. A comprehensive study was done to identify limitations of enterprise grade solutions and how they are not feasible for SME environments. Key market dynamics were explored to grasp the significance and expansion of the SOAR, SIEM and XDR markets as they are closely related and work similar to the proposed solution. Studies shows that SMEs face budget constraints and lack expert staff to operate cyber security solutions which remains as the core challenge [9]. Documents involving EDR and XAI was also analyzed to understand how these technologies are used for various security functions. It was found that these solutions are primarily intended for security analysts, a position that is typically absent in most SMEs.

In order to better understand the practical challenges and adaptive strategies that SMEs currently use, a theoretical case analysis was carried out after the literature review. To do this, the operational realities of a typical SME—which are characterized by heterogeneous, high-noise environments with a mix of legacy and unmanaged devices—were analyzed, and a crucial component of this phase was modeling how a non-technical user would interact with a security system. Based on research on behavior-driven security and adaptive learning, the study found a critical gap that current enterprise solutions are unable to fill: a system that can learn from subjective user judgments and "low-fidelity human input’.

The third phase involved a combination of operational studies and literature reviews to create a conceptual framework. This conceptual framework which is not yet finalized involves designing a novel solution with “human in the loop” interface with XAI, creating an “autonomous tuning engine” and developing an “adaptive playbook generator”. The system will be developed by a human to think and act autonomously with a transparent override mechanism for critical decisions and displays the thinking patterns of AI. This will help bridge the trust issue that is highlighted in the research documents.

In order to ensure that the proposed solution fulfils its essential commitment a concluding evaluation phase will be conducted. In this phase test data and metrics similar to SME environments will be used to test out the effectiveness of the solution. An array of test scenarios involving alert resolution, user-friendliness and Mean Time to Respond (MTTR) will be monitored. This will ensure the solution matches the benchmarks or standards and solves the bottlenecks in SMEs as stated in the proposal document.

System Overview

The proposed system consists of:

- Alert Collector Module: Collects and visualize the alerts generated by the AI-Powered behavioral detection engine. This particular engine will be developed separately and wont be part of this proposal scope. However, this module will be integrated with the engine to gather alert data for visualization purposes.

- Normalization & Correlation Engine: Converts heterogeneous data into a common schema (e.g., STIX/TAXII).

- AI Prioritization Engine: Employs supervised learning algorithms (Random Forest, Gradient Boosting) to rank alerts.

- SOAR Workflow Module: Automates responses such as blocking IPs, quarantining files, or disabling compromised accounts.

- Dashboard Interface: Provides intuitive visualizations for SMEs with minimal technical overhead. This dashboard is capable of generating business context dashboards with explainable AI which explains about incidents in a tone which the management understand. The dashboard will have an AI where the user can use a generative AI to get any information they want.

- Adaptive Response Component: Learns from prior user responses to refine future automation. The solution can read user input and understand the context to respond and refine incident resolutions.

Approach

- Datasets needs to be collected to simulate the capabilities. Use benchmark datasets (CICIDS 2017, MITRE ATT&CK–based scenarios) alongside simulated SME log environments.

- Tools: Python, Elastic Stack, TensorFlow, and lightweight SOAR frameworks (e.g., Shuffle, StackStorm).

- Evaluation Metrics: Alert reduction percentage, false positive/negative ratio, usability testing feedback, MTTR improvement.

- The system will be integrated with AI-Powered Behavioral Detection engine which will be developed as subcomponent of the main research topic.

Project Requirements

Functional:

- Collect and normalize alerts.
- Prioritize and visualize threats.
- Automated enrichment and remediation.
- Enable integration with SME collaboration tools.

Non-Functional:

- Lightweight deployment (runs on modest hardware).
- Usable by non-expert staff.
- Scalable to growing SMEs.
- Secure and compliant with data protection standards.

Expected Test Cases

1. Validate normalization accuracy across heterogeneous alert sources.
2. Assess prioritization effectiveness using labeled datasets.
3. Test automated workflows in simulated incident scenarios.
4. Measure adaptive learning improvement after repeated user interactions.

Use Cases

SOAR platforms greatly contribute to automation and orchestration. However, these platforms are made for enterprises who possess the investment capabilities to maintain the platforms. This restricts SMEs to experience state of the art security thereby exposing to multiple threats and risks. To overcome the limitations of resources and restriction of budgets this project aims to propose a security alert interface with SOAR-style automation. This not only addresses concerns of SMEs but also enhances the SOAR far beyond what is offered to enterprises in certain functionalities.

Use Case ID	Name
U01	Receive and Normalize Alerts
Description	The interface unifies security alerts from various sources (endpoints, firewalls, etc.) into a single, normalized, and easy-to-read dashboard view, reducing alert fatigue for the user.
Application	The central dashboard and its data ingestion engine.
Primary Actor	The Security Alert Interface.
Pre-condition	The dashboard is configured to receive alert data from at least one security tool (e.g., firewall, endpoint agent).
Trigger	A security event occurs on an endpoint or network device, generating an alert.
Basic Flows	• **Steps 1:** The alert ingestion engine receives raw data from the external source. • **Steps 2:** The engine normalizes the data, standardizing fields like source, severity, and event type across all alerts. • **Steps 3:** The normalized alert is sent to the central processing module for prioritization. • **Steps 4:** The alert appears on the user's unified dashboard.
Alternative Flows	• **Steps 2a:** If the ingestion fails, an error is logged to the system health dashboard. • **Steps 3a:** If the alert data is too corrupt to be normalized, the original raw data is stored

	in a log for later review and a system-level alert is created.
Post-condition	A single, unified list of security events is available on the dashboard, ready for review and action.

Table 2: Use Case 1

Use Case ID	Name
U02	User-Approved Threat Containment
Description	The interface presents a high-priority threat with a simplified explanation and a single, clear button for the user to approve a complex, automated response workflow.
Application	The dashboard's user interaction module.
Primary Actor	The SME Administrator or a non-technical user.
Pre-condition	A high-fidelity threat is detected that requires human confirmation before an autonomous response can be executed.
Trigger	A high-priority alert appears on the dashboard with a "Pending Approval" status.
Basic Flows	• **Steps 1:** The dashboard displays a critical alert with a simplified, business-contextual explanation of the threat. • **Steps 2:** It presents a clear, one-click "Approve Remediation" button to the user. • **Steps 3:** The user reviews the explanation and clicks the button to approve the action. • **Steps 4:** The SOAR engine executes a complex, predefined workflow (e.g., endpoint isolation, process termination, file quarantine). • **Steps 5:** The dashboard is updated with a log of the successful remediation, and the incident is closed.
Alternative Flows	• **Steps 3a:** The user declines the action, and the system prompts them to provide a reason

	or mark it as a false positive, feeding this data back into the learning model. • **Steps 4a:** If the automated workflow fails, the dashboard displays an error and provides the administrator with manual remediation instructions.
Post-condition	A high-priority threat is contained, and the user has a clear record of the decision and the automated actions taken.

Table 3: Use Case 2

Use Case ID	Name
U03	Prioritize Alerts with Business Context
Description	The dashboard's prioritization model automatically ranks alerts based on their potential business impact rather than just technical severity, providing a clear, actionable view for a non-technical user.
Application	The dashboard's alert prioritization and visualization module.
Primary Actor	The Security Alert Interface.
Pre-condition	The system has access to a simple asset inventory that labels key business assets (e.g., "Critical" servers, "Accounting" data) and user roles.
Trigger	A new alert is generated from an integrated security source.
Basic Flows	• **Steps 1:** The system receives a new alert from a security tool. • **Steps 2:** The system's prioritization model correlates the alert with the business's asset inventory and user data. • **Steps 3:** The dashboard displays the alert with a descriptive, non-technical priority label, such as "High Priority: A threat to your accounting data."
Alternative Flows	• **Steps 2a:** If the system cannot correlate the alert with a known asset, it defaults to a standard technical severity score and prompts

	the administrator to manually add business context for future alerts.
Post-condition	The user can easily understand the real-world implications of each alert and focus on the most critical threats to their business.

Table 4: Use Case 3

Test Cases

Test Case ID	Objective	Pre-conditions	Steps	Expected Result	Pass Criteria
TC-SAI-01	Verify unified alert ingestion and dashboard normalization.	An endpoint security agent and a firewall are configured to send alerts to the system's ingestion engine.[41]	1. Trigger a malware detection event on the endpoint. 2. Trigger a denied connection event from the firewall.	The dashboard displays both events in a single, unified view, with normalized fields for source, severity, and event type.	The dashboard shows two distinct alerts with consistent data fields, and the total alert count is accurate and reflects both events.
TC-SAI-02	Validate automated alert prioritization based on business context.	The system's asset inventory tags a server as "Accounting Database" with a high criticality score.	1. Trigger a low-level vulnerability scan alert from an unprivileged user on the "Accounting Database" server.	The system elevates the alert's priority from low to high due to the critical nature of the targeted asset, and the dashboard displays a clear, business-focused label.	The alert is visually marked as "High Priority" on the dashboard, and the description mentions "Accounting Database" or a similar business-contextual term.
TC-SAI-03	Verify autonomous workflow execution in response to a high-fidelity alert.	An automation workflow is configured to automatically quarantine a host when a ransomware detection alert is received.	1. A known ransomware signature is detected on an endpoint by an integrated security tool.	The system automatically triggers the "quarantine host" action on the endpoint without requiring manual user approval, and a corresponding log is created in the dashboard.	The endpoint is isolated from the network within seconds of the alert, and the dashboard log shows "Automated Quarantine Action Taken."
TC-SAI-04	Validate the feedback-driven learning mechanism for false positives.	The system has flagged a benign application activity as suspicious.	1. The user views the alert on the dashboard. 2. The user selects "Dismiss as False Positive."	The system acknowledges the user's feedback and refines its detection logic for that specific activity.	The benign activity is no longer flagged in subsequent occurrences, demonstrating that the system learned from user

					input to reduce false positives.
TC-SAI-05	Verify the translation of a technical alert into a human-friendly narrative.	A security alert is triggered by a technical event (e.g., "unusual API call pattern"). The XAI component is enabled.	1. A suspicious process is detected by the system.	The dashboard displays a simplified explanation of the event in a human-friendly narrative, such as, "This process is behaving like ransomware and could lock your files."	The alert title and description on the dashboard are non-technical and clearly articulate the potential business impact of the threat.
TC-SAI-06	Verify the generation of an automated compliance report.	The system is configured with a rule to log all critical security events and their remediation for a monthly compliance audit.	1. Several critical alerts are handled by the system's automated workflows. 2. The user generates the monthly compliance report via the dashboard.	The system compiles a detailed report of all critical alerts, the remediation steps taken, and relevant timestamps and user actions.	The generated report is a clear, standardized document that accurately summarizes the security activity and can be used for compliance verification.
TC-SAI-07	Validate the creation of an adaptive playbook for a novel threat.	A new, previously unseen threat (e.g., a zero-day exploit) is detected on an endpoint. The system has no predefined playbook for this threat.	1. The detection engine flags a novel threat.	The system's AI automatically generates a new incident response playbook, outlining a series of logical steps to contain and investigate the threat based on its characteristics.	A new playbook is created and made available in the system, with a step-by-step plan that is tailored to the specific, novel threat and its detected behaviors.
TC-SAI-08	Verify the system’s ability to provide automated user guidance.	A rule is set to automatically send a training message to an employee who triggers a low-severity phishing alert.	1. An employee clicks a suspicious link in a simulated phishing email.	The system does not create an alert for the administrator but sends a direct, non-technical message to the employee explaining the risk	The employee receives a concise, educational message, and the administrator's dashboard remains free of

				and providing a best practice.[4]	unnecessary low-level alerts.
TC-SAI-09	Validate the autonomous containment and notification for remote endpoints.	An endpoint agent is installed on a laptop. The laptop is disconnected from the main network.	1. A simulated attack occurs on the laptop while offline. 2. The laptop's agent autonomously contains the threat (e.g., quarantines a file).	The agent takes autonomous action to contain the threat while offline. Once reconnected to the network, the dashboard displays a full report of the incident and the actions taken.	The dashboard shows a complete, timestamped record of the offline event and the autonomous remediation without any data loss.
TC-SAI-10	Verify the prioritization and visualization of alerts based on user role.	The system has two users: an "Admin" with full privileges and an "Employee" with limited access. A high-priority alert is generated for a user's device.	1. The system detects a lateral movement attempt from the employee's device to a critical server.	The dashboard's threat prioritization model ranks this alert as "Critical" for the Admin's view, while a simplified, less-technical version is presented to the employee's view (if applicable) for their awareness.	The Admin's dashboard shows a "Critical" alert with a technical description, while the employee's dashboard (if configured) shows a low-priority, informative alert about their device without technical details.

Table 5: Test Cases Table

Wireframes

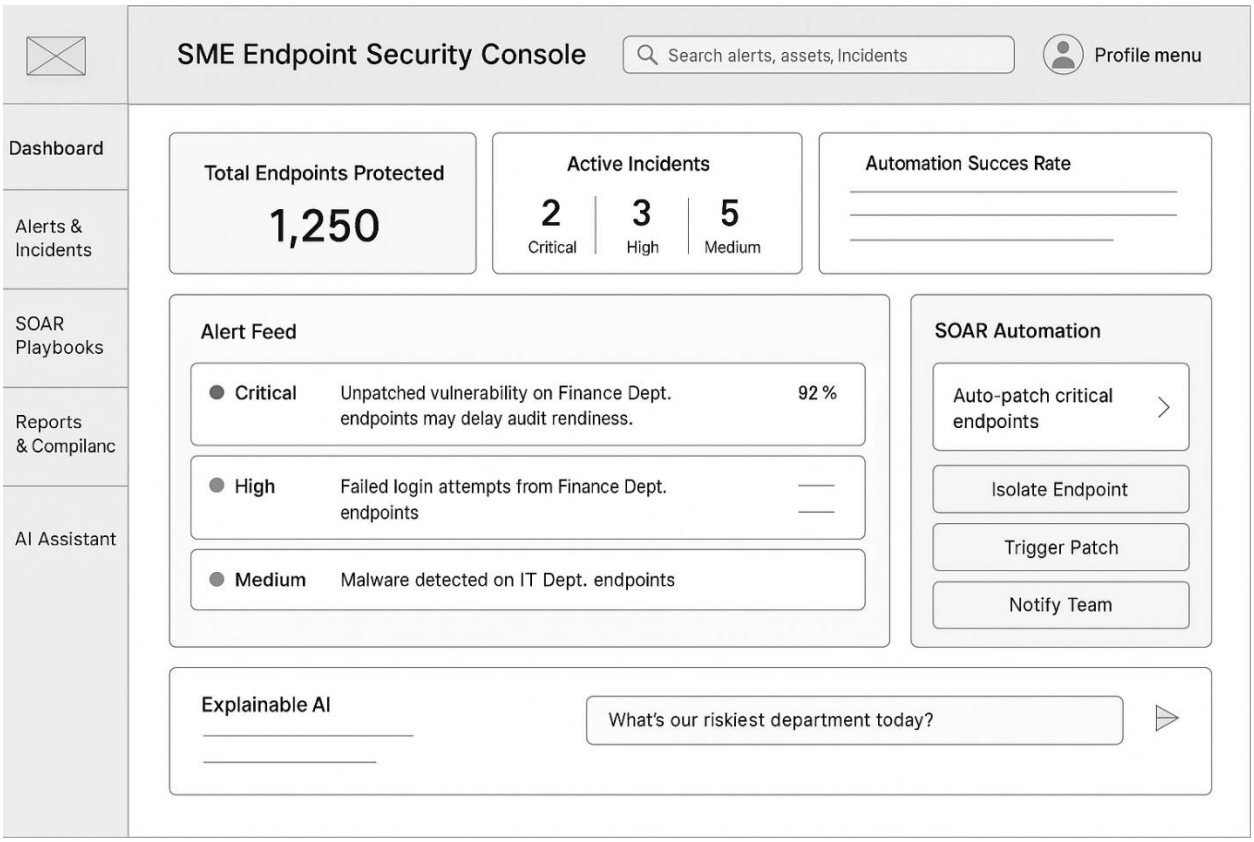

Figure 1: Alert Dashboard

Figure 1 shows the initial design of the dashboard. The essential details that are needed will be displayed here. This will involve the total number of protected endpoints, incidents and other useful essential information. The dashboard appearance will be highly customizable with drag and drop widgets and AI based chart / widget creation that does not need any code, script or query knowledge. When clicked on the widgets, it will expand into more granular levels to view more information.

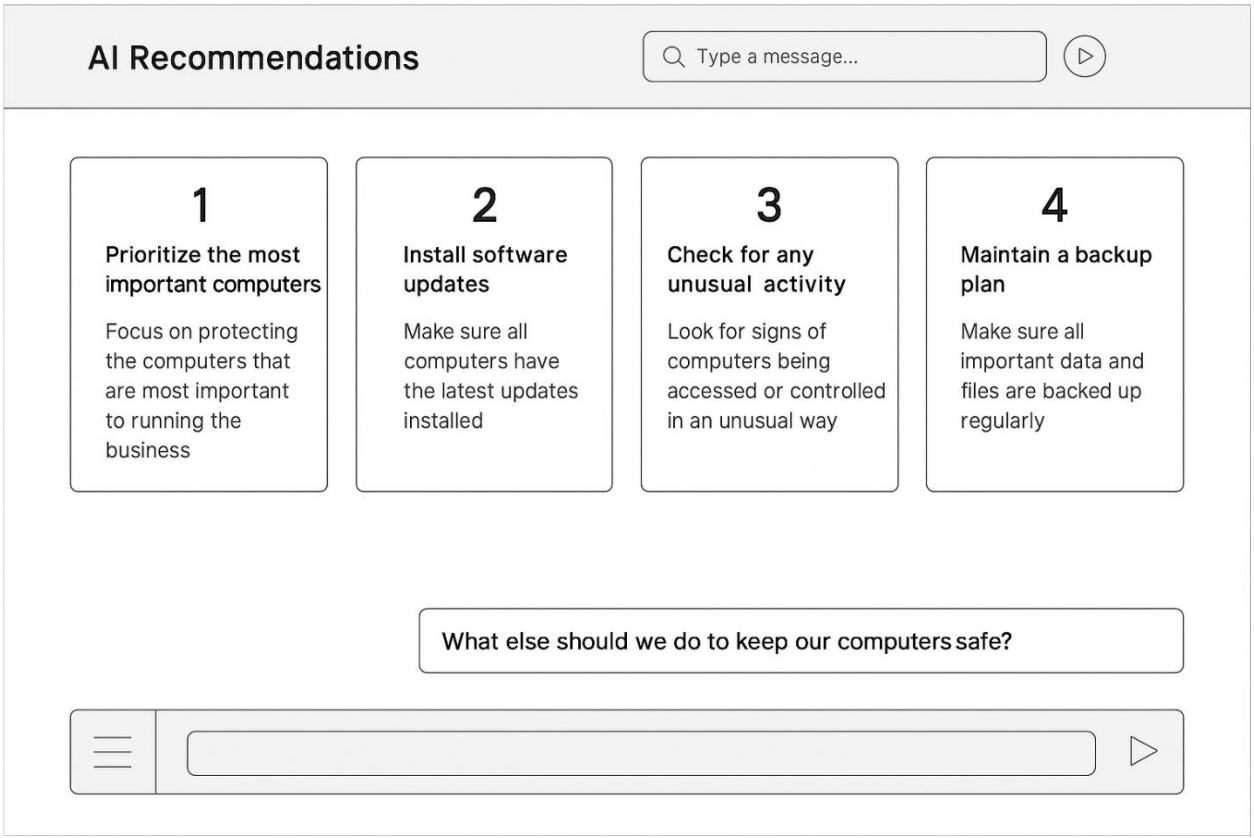

Figure 2: XAI Dashboard

To solve the scarcity of expertise analysts, the dashboard will be equipped with AI that is capable to provide the user with step by step instructions on what to do if any user actions are required. The AI will also respond to questions raised by the user to give understandable insights to non-technical people. Moreover, the system will have an inbuild audit feature which shows the thinking patterns of the AI, so that the user can see what lead the AI to take certain decisions. Through the dashboard itself the user will be able to instruct the AI on wrong actions and approaches so that the AI can refine and improve.

High level Conceptual Architecture Diagram

Below diagram shows the high-level architecture of the proposed solution. At this stage the diagram is not finalized as it is an initial draft. I have also highlighted why the current diagram might change as the project progresses.

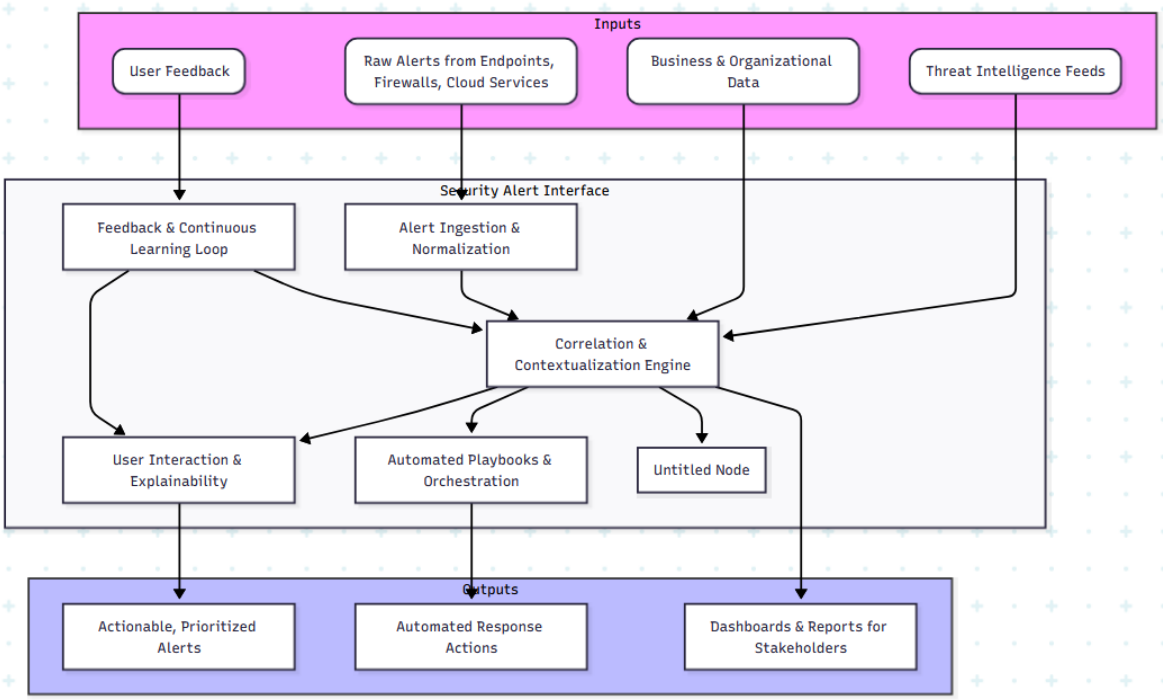

Figure 3: Architecture Diagram

Inputs – this involves the sources from which the security alert interface module receives the data form.

Outputs – describes the output the system generates and displayed after the data received is processed by the system.

Alert Ingestion & Normalization – This module collects raw alerts from multiple sources. However, there is a possibility that this might change to receiving processed alerts as the system will integrate with AI-powered Behavioral Detection Engine (another subcomponent of the main project) which will also do some processing. It is not properly defined yet whether the processing is entirely done in the AI-powered Behavioral Detection Engine or whether some processing will be passed to the alert interface module. Aspects such as efficiency and cost will be considered to decide on processing raw alerts which will be decided when the project progresses.

Correlation & Contextualization Engine - In order to identify trends and linkages between occurrences, warnings are now connected across systems. A brute force assault can be indicated, for instance, by several unsuccessful login attempts followed by a privilege escalation attempt. By adding external threat intelligence and internal organizational data to alarms, this module creates a more comprehensive picture of every possible occurrence. It will be decided how this processing will be balanced with the processing done in the AI-powered Behavioral Detection Engine as some processing will be have to done in both modules.

Automated Playbooks and Orchestration – This module is capable or creating playbooks both manually and dynamically generated workflows. Unlike traditional SOAR systems that highly depend on manually created playbooks this system will be designed to automatically create playbooks and test them out in test environments based on attack patterns, human feedback and contextual data. The user can also enter an objective of a playbook to the AI where automatically a playbook will be crafted to match the user's requirements.

User Interaction and Explainability - Trust among analysts depends on transparency. This module explains in detail why certain response actions are advised and why alerts are given priority. It ensures informed decision-making by presenting technical security findings in an understandable manner for analysts and business stakeholders through the use of natural language summaries.

Work Breakdown Structure

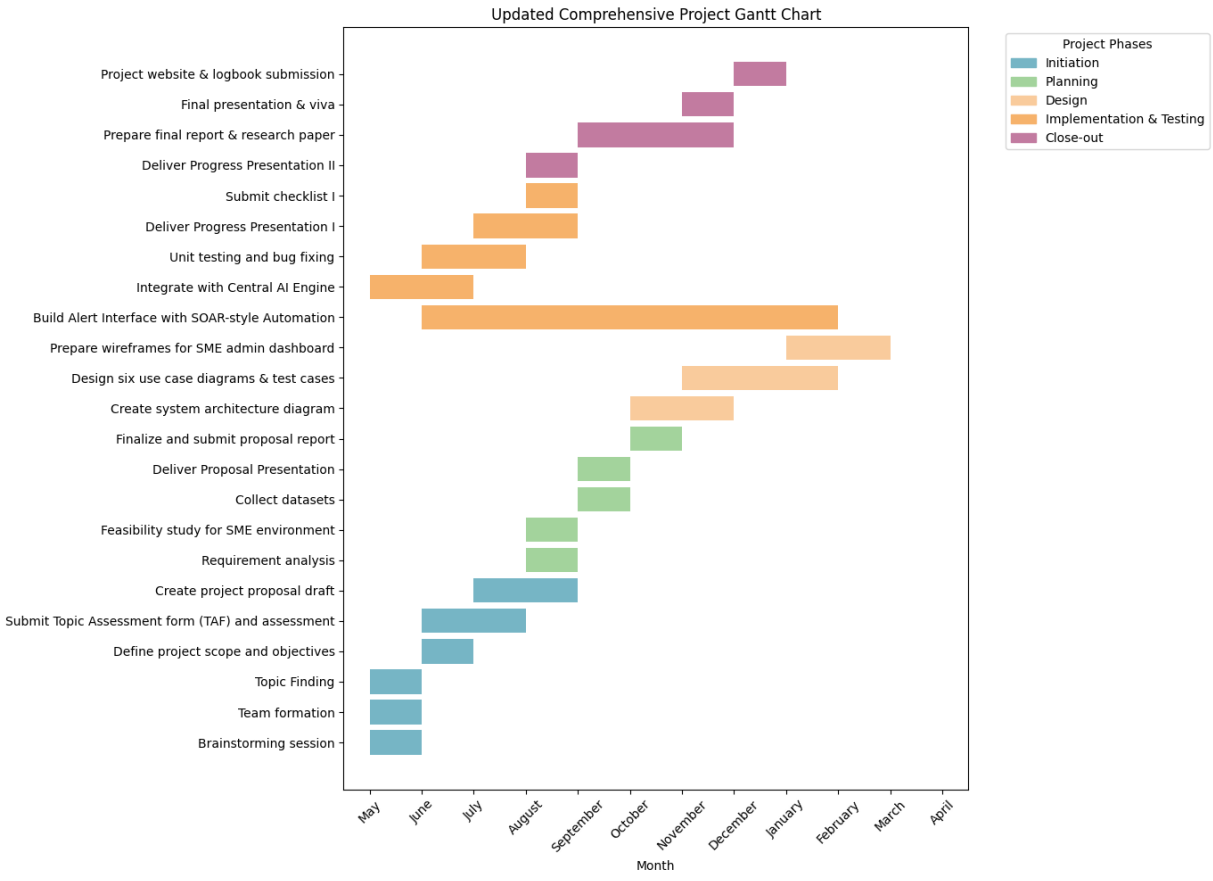

Figure 4: Gantt chart

The project commences with the Initiation phase, which entails brainstorming meetings to generate ideas, assembling the project team, and confirming the topic. This phase also includes outlining the project's scope and goals, submitting the Topic Assessment Form (TAF), and crafting the preliminary project proposal. Once the groundwork is established, the Planning phase centers on evaluating system requirements, performing a feasibility study suited to SME environments, gathering pertinent datasets, and presenting the proposal. This phase wraps up with the completion and submission of the proposal report.

The Design phase encompasses developing the system architecture diagram, creating use case diagrams and test cases, and designing wireframes for the SME admin dashboard. These design components provide clarity and organization for the development process. The Implementation & Testing phase focuses on constructing the primary component: the Alert Interface with SOAR-style Automation. This module is tasked with receiving alerts, categorizing them, and executing automated incident response measures. It is followed by the integration with a central AI engine (AI powered behavioral threat detection engine) , unit testing, bug resolution, and progress presentations to monitor development milestones.

Ultimately, the Close-out phase finalizes the project with a second progress presentation, preparation of the final report and research paper, and the concluding presentation and viva. The project ends with the submission of the project website and logbook, ensuring that all documentation and deliverables are filed and available.

Commercialization

Creating an all-encompassing, proprietary security platform from scratch is exceedingly costly and requires significant time investment. A more tactical strategy is to utilize the extensive range of open-source tools that already encompass many essential SOAR functionalities. Platforms such as Shuffle, TheHive, and Wazuh present no-cost, customizable elements for case management, threat intelligence, and host-based intrusion detection.

The suggested approach would utilize a carefully selected open-source foundation, layered with a proprietary managed service on top. This mixed model would offer the "plug-and-play" convenience and user-friendliness that SMEs require, while simplifying the inherent complexities of the open-source elements. This methodology is not only more economical but also more scalable than a unified SOAR, enabling the solution to provide enterprise-level capabilities at a significantly lower cost. The managed service layer would also deliver the specialized support and ongoing monitoring that SMEs often lack internally, addressing the technical skills gap and offering a comprehensive answer [29].

For a security solution to thrive in the SME sector, its go-to-market plan must tackle the TCO dilemma head-on. Conventional SaaS frameworks that bill based on per-user or per-feature can swiftly become too expensive for an expanding company with limited finances. The optimal strategy is the freemium model. By providing a basic, no-cost version of the product, a freemium

model reduces the entry barriers, facilitating quick user growth and brand visibility without a substantial initial monetary investment [28].

The complimentary tier would offer essential, high-value features, including minimal false positive adjustments, Explainable AI notifications, and a basic level of self-sufficient threat management. This enables the SME to "test drive" the service, experiencing its benefits directly and fostering confidence in the platform. An effective freemium model can result in conversion rates ranging from 5% to 10% for users transitioning from free to paid plans. The paid tiers would then provide access to more sophisticated features, such as enhanced automation, wider integrations, and a complete range of managed services. This tiered strategy, along with an emphasis on user-friendliness and a hybrid open-source framework, establishes a compelling value offering that aligns seamlessly with the requirements and limitations of the SME sector [17].

Target Audience

- Small and Medium Enterprises (SMEs) lacking dedicated cybersecurity teams
- Managed Security Service Providers (MSSPs) looking for cost-effective insider threat detection solutions
- IT Administrators and SOC Analysts in SME environments
- Educational institutions, healthcare centers, and financial organizations with limited budgets but high data sensitivity
- Government agencies and NGOs seeking lightweight endpoint monitoring for compliance

Budget and Budget Justification

Item	Quantity	Unit Cost (USD)	Total Cost (USD)	Justification
Cloud / Compute Resources (Google Colab Pro / AWS Free Tier upgrade)	8 months	$15 / month	$120	The entire solution will be cloud hosted as it requires computational resources to host the AI modules and data processing.
Threat Intelligence Feeds (Open-Source Integration)	–	Free	$0	Using open-source feeds such as AlienVault OTX, AbuseIPDB, and MISP for enrichment. No licensing cost.
Data Storage (Google Drive / OneDrive student subscription)	200 GB	Included / $20	$20	Storage for large log datasets and model checkpoints.
Development Tools (VS Code, Python, Jupyter)	–	Free	$0	Opensource tools (Wazuh, Shuffle and Tracecat).
API Services (Optional for TI enrichment, e.g., VirusTotal API keys)	2 keys	$50	$100	For querying file hashes, domains, and IPs to enhance detection confidence.

Software tools & licenses	–	–	$100	This is a buffer amount allocated in case we need to get licenses to purchase certain products for comparison and testing
Documentation & Reporting (Printing, Binding, reports and presentations.)	–	$20	$20	Preparation of final report, printing diagrams, tools and service subscription for resources like reference materials and journal access.

Table 6: Budget & Budget Justification

The proposed budget is compiled to be more cost effective and accessible for the project. More priority is given to free and open source tools wherever possible. The project team will also try to bring down the costs even further by exploring plans for student subscriptions for cloud services which are more budget friendly. For resources freely, available documents will be obtained from online or campus library if available. In case more indebt research of findings are needed it will be paid for. Wherever possible budget friendly or free options will be utilized to make the proposed solution.

References

[1] S. R. Pulyala, A. G. Desetty, and V. D. Jangampet, "The Impact of Security Orchestration, Automation, and Response (SOAR) on Security Operations Center (SOC) Efficiency: A Comprehensive Analysis," *Turkish Journal of Computer and Mathematics Education*, vol. 10, no. 3, pp. 1545-1549, 2019. [Online]. Available: https://turcomat.org/index.php/turkbilmat/article/view/14323/10355

[2] European Union Agency For Cybersecurity, June 2021, "Cybersecurity for SMEs – Challenges and Recommendations. [Online]. Available: https://www.enisa.europa.eu/publications/enisa-report-cybersecurity-for-smes

[3] K Shaukat et al. "A survey on machine learning techniques for cyber security in the last decade," November 2020. [Online]. Available: https://www.researchgate.net/publication/346487463_A_Survey_on_Machine_Learning_Techniques_for_Cyber_Security_in_the_Last_Decade

[4] S. Amershi et al. "Guidelines for human-AI interaction," May 2019. [Online]. Available: https://www.microsoft.com/en-us/research/wp-content/uploads/2019/01/Guidelines-for-Human-AI-Interaction-camera-ready.pdf

[5] M. Ghazizadeh et al. "Extending the Technology Acceptance Model to assess automation," October 2011. [Online]. Available: https://www.researchgate.net/publication/220579437_Extending_the_Technology_Acceptance_Model_to_assess_automation

[6] A. Horneman, J. Ray. "Benefits and Challenges of SOAR Platforms," March 2021. [Online]. Available: https://www.sei.cmu.edu/blog/benefits-and-challenges-of-soar-platforms/

[7] N. Chowdhury, I. Ahmed, A. Haque. “DeepSeek vs. ChatGPT: A Comparative Analysis of

Performance, Efficiency, and Ethical AI Considerations,” February 2025. [Online]. Available: https://www.researchgate.net/publication/388902239_DeepSeek_vs_ChatGPT_A_Comparative_Analysis_of_Performance_Efficiency_and_Ethical_AI_Considerations

[8] Palo Alto Networks, "What is SOAR?" Palo Alto Networks, [Online]. Available: https://www.paloaltonetworks.com/cyberpedia/what-is-soar. [Accessed: Aug. 25, 2025].

[9] Sirp, "3 Common Challenges to Avoid While Implementing SOAR," Sirp, [Online]. Available: https://sirp.io/blog/3-common-challenges-to-avoid-while-implementing-soar/. [Accessed: Aug. 25, 2025].

[10] L. A. Odozor, L. Omini, S. N. Berko, U. Precious, Y. Nitzan, and K. Idowu, "An Incident Response Playbook Guide for Small and Medium Enterprises (SMEs)," International Journal of Scientific and Management Research, vol. 8, no. 7, pp. 145-157, July 2025. doi: 10.37502/IJSMR.2025.8712.

[11] M. F. Arroyabe, C. F. A. Arranz, I. F. de Arroyabe, and J. C. F. de Arroyabe, "Revealing the realities of cybercrime in small and medium enterprises: Understanding fear and taxonomic perspectives," Comput. Secur., vol. 141, May 2024, Art. no. 103826, doi: 10.1016/j.cose.2024.103826.

[12] O. F. Ayepeku and S. O. Olofinlade, “Challenges and Opportunities of AI-Driven Cybersecurity for Small and Medium Enterprises (SMEs) Towards Poverty Reduction in Nigeria,” Scientific and Practical Cyber Security Journal (SPCSJ), vol. 8, no. 3, pp. 74–83, 2023.

[13] M. Awan, A. Alam, and M. Kamran, “Cybersecurity Challenges in Small and Medium Enterprises: A Scoping Review,” Journal of Cyber Security and Risk Auditing, vol. 2025, no. 3, pp. 89–102, Jul. 2025. [Online]. Available: https://doi.org/10.63180/jcsra.thestap.2025.3.7

[14] H. Aver, "Commercial vs. open-source SIEM: pros and cons," Kaspersky Official Blog. [Online]. Available: https://www.kaspersky.com/blog/open-source-siem-hidden-costs/53589/. [Accessed: Aug. 26, 2025]

[15] Y. Brun, G. Di Marzo Serugendo, C. Gacek, H. Giese, H. Kienle, M. Litoiu, H. Müller, M. Pezzè, and M. Shaw, “Engineering Self-Adaptive Systems through Feedback Loops,” in Proc. Int. Workshop on Software Engineering for Self-Adaptive Systems, LNCS 5525, Springer, pp. 48–70, 2009. doi: 10.1007/978-3-642-02161-9_3

[16] P. F. Christiano, J. Leike, T. B. Brown, M. Martic, S. Legg, and D. Amodei, “Deep reinforcement learning from human preferences,” arXiv preprint arXiv:1706.03741v4, Feb. 2023. [Online]. Available: https://arxiv.org/abs/1706.03741v4

[17] Maxio. “Freemium Model,” Maxio. [Online]. Available: https://www.maxio.com/blog/freemium-model. [Accessed: Aug. 28, 2025].

[18] Stripe. “Software Pricing Models and Strategies for SaaS Businesses,” Stripe. [Online]. Available: https://stripe.com/resources/more/software-pricing-models-and-strategies-for-saas-businesses. [Accessed: Aug. 28, 2025].

[19] CyberGL. “Cybersecurity MSP Guide,” CyberGL Blog. [Online]. Available: https://cybergl.com/blog/cybersecurity-msp-guide/. [Accessed: Aug. 29, 2025].

List of Abbreviations

- AI: Artificial Intelligence
- EDR: Endpoint Detection and Response
- ENISA: European Union Agency for Cyber Security
- GTM: Go-to-Market
- HIDS: Host-based Intrusion Detection System
- IDE: Integrated Development Environment
- IDS: Intrusion Detection System
- IOC: Indicator of Compromise
- IP: Internet Protocol
- MDR: Managed Detection and Response
- ML: Machine Learning
- MSPs: Managed Security Service Providers
- MTTR: Mean Time to Respond
- NGAV: Next-Generation Antivirus
- R&D: Research and Development
- RL: Reinforcement Learning
- SaaS: Software as a Service
- SIEM: Security Information and Event Management
- SME: Small and Medium-sized Enterprises
- SOC: Security Operations Center
- SOAR: Security Orchestration, Automation, and Response
- WBS: Work Breakdown Structure
- XAI: Explainable AI
- XDR: Extended Detection and Response