Real-Time SIEM-Based Cybersecurity Framework for Threat Detection and Prevention in IoMT Environments

25-26J-70

Project Proposal Report

Basheer MS – IT22031570

(Gunasekara A.G.M.K, Ukasha MMM, Firaz MMN)

B.Sc. (Hons) in Information Technology

Specializing in Cyber Security

Department of Information Technology

Sri Lanka Institute of Information Technology

Sri Lanka

September 2025

Real-Time SIEM-Based Cybersecurity Framework for Threat Detection and Prevention in IoMT Environments

25-26J-283

Project Proposal Report

Basheer MS – IT22031570

Supervisor: Mr. Kanishka Yapa
Co-supervisor: Mr. Deemantha Siriwardhana

B.Sc. (Hons) in Information Technology
Specializing in Cyber Security

Department of Information Technology

Sri Lanka Institute of Information Technology
Sri Lanka

September 2025

DECLARATION

We declare that this is our 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.

Name	Student ID	Signature
Basheer M. S.	IT22031570	

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

Name of supervisor: Kanishka Yapa

Name of co-supervisor: Deemantha Siriwardhana

KY 27/08/2025

Signature of the supervisor: **Date**

(Kanishka Yapa)

Abstract

The Adaptive Incident Correlation Engine (AICE) forms the decision-making hub of the proposed IoMT cybersecurity framework, bridging raw AI threat predictions and real-time automated responses. Unlike traditional SIEM integrations, AICE correlates heterogeneous alerts, performs contextual analysis, assigns severity levels, and prioritizes mitigation workflows. By integrating outputs from AI models, Intrusion Detection Systems (IDS) such as Suricata and Zeek, and medical-specific logs, the engine reduces false positives and improves incident prioritization. Core innovations include anomaly correlation, severity scoring, and visualization dashboards for SOC teams, ensuring rapid and accurate decision-making within resource-constrained Sri Lankan healthcare environments. This component transforms isolated alerts into actionable intelligence, enabling hospitals to enhance cyber resilience, regulatory compliance, and patient safety.

Keywords:

- Adaptive Correlation
- Incident Severity Scoring
- IDS + AI Fusion
- IoMT Cybersecurity
- SOC Dashboards
- Healthcare Compliance (HIPAA, Local laws)

TABLE OF CONTENT

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

- **AICE** – Adaptive Incident Correlation Engine
- **IoMT** – Internet of Medical Things
- **SIEM** – Security Information and Event Management
- **IDS** – Intrusion Detection System
- **SOC** – Security Operations Center
- **PHI** – Protected Health Information
- **HIPAA** – Health Insurance Portability and Accountability Act
- **XAI** – Explainable Artificial Intelligence
- **ELK** – Elasticsearch, Logstash, Kibana
- **DoS/DDoS** – Denial of Service / Distributed Denial of Service

1 INTRODUCTION

The Internet of Medical Things (IoMT) has revolutionized healthcare delivery by connecting medical devices such as infusion pumps, ventilators, imaging systems, and wearable sensors into secure digital ecosystems [1], [2]. These devices generate heterogeneous streams of sensitive data that support real-time monitoring, diagnostics, and patient safety. However, the rapid adoption of IoMT has exposed hospitals to a growing range of cyber threats, including denial-of-service (DoS) attacks, ransomware, data tampering, and device hijacking [3], [4]. In healthcare, where disruptions directly affect patient outcomes, timely and accurate incident detection and response are essential.

Traditional solutions, including Intrusion Detection Systems (IDS) and SIEM platforms, are effective at generating alerts but often overwhelm analysts with false positives and uncorrelated anomalies [5], [6]. Alerts from IDS or AI engines typically remain isolated, leaving security teams with fragmented threat visibility. Furthermore, most tools lack severity prioritization, clinical context awareness, and transparent decision-making mechanisms, which results in delayed responses and regulatory risks [7], [8].

To address these challenges, the Adaptive Incident Correlation Engine (AICE) has been proposed as the decision-making hub of IoMT cybersecurity pipelines. AICE fuses alerts from AI-based anomaly detectors, IDS logs, and device-specific events, correlating them into coherent incidents [9]. By assigning severity scores based on device criticality, patient impact, and attack type, the system reduces analyst fatigue and improves prioritization [10], [11]. It further provides decision dashboards, contextual audit trails, and automated compliance tagging to align with regulations such as HIPAA and Sri Lanka's data protection laws [12], [13].

In resource-limited healthcare environments like Sri Lanka, AICE offers a tailored, cost-effective, and resilient solution. By transforming raw alerts into actionable intelligence, AICE enhances patient safety, strengthens compliance, and builds a proactive cybersecurity posture in IoMT ecosystems [14], [15].

1.1 Background Study

Adaptive correlation technology has rapidly evolved, providing intelligence-driven platforms for multi-source alert fusion, contextual decision support, and compliance reporting in complex environments. Tools such as Suricata, Zeek, and AI-based anomaly detectors enable integration of diverse data sources, real-time pattern analysis, and event correlation across IoMT devices [1], [2]. In healthcare, correlation engines are critical for unifying heterogeneous alerts from endpoints, medical devices, and applications to detect coordinated intrusions or misuse of patient data [3], [4]. For example, recent studies demonstrate that ensemble IDS models combined with AI significantly improve attack detection performance, but the correlation step that transforms raw alerts into actionable intelligence remains underdeveloped [2], [5]. Commercial platforms extend SIEM functionality with large-scale correlation, but open-source or adaptive solutions remain more cost-effective for resource-constrained hospitals [6].

However, conventional detection systems in healthcare still face significant challenges. There is a lack of adaptive mechanisms to correlate alerts from multiple layers—such as device anomalies, IDS logs, and AI predictions—which results in fragmented views of threats [7], [8]. Many tools produce high volumes of false positives without contextual prioritization, creating analyst fatigue and delayed responses [9]. Furthermore, most systems lack clinical context: they fail to differentiate between alerts from life-critical devices (e.g., ventilators) and those from non-critical assets (e.g., administrative PCs) [10]. This absence of severity scoring reduces decision-making efficiency in high-risk environments [11].

Recent progress highlights the use of explainable AI and interpretable feature reduction to improve transparency and trust in IoMT detection models [12]. Yet, their integration into correlation engines is limited, leaving SOC teams with black-box outputs rather than clear, contextual decisions. In Sri Lanka and other resource-limited healthcare systems, this challenge is amplified by constrained IT staffing, unstable connectivity, and diverse staff language needs [6], [13]. Effective correlation engines must therefore emphasize adaptive severity assignment, transparent outputs, and resilience to infrastructure limitations.

Despite advances in IDS, anomaly prediction, and ensemble detection, existing solutions often fail to unify anomalies into coherent incident narratives [1], [4], [7]. Attack campaigns targeting IoMT frequently span multiple devices and timeframes; without correlation, these patterns remain undetected [8], [9]. Moreover, audit trails for compliance (HIPAA and local data protection laws) are inconsistently applied, with tagging often requiring manual intervention [10], [11]. In environments like Sri Lanka, unreliable connectivity and limited resources further degrade traditional monitoring systems [6], [13].

In summary, while IDS and AI-based tools provide strong anomaly detection [1], [2], [5], the adaptive correlation of alerts, severity-aware scoring, explainability, and compliance tagging [9], [12], [14] remain critical gaps. These gaps highlight the need for an Adaptive Incident Correlation Engine (AICE) specifically tailored for IoMT security in Sri Lankan hospitals.

1.2 Research Gap

Although significant progress has been made in IoMT security, major gaps remain that hinder effective cybersecurity in healthcare environments. Current IDS and AI solutions excel at detecting anomalies but rarely integrate alerts from diverse sources into a unified decision-making framework [1], [3]. The lack of multi-source alert fusion results in fragmented visibility, making it difficult for analysts to identify coordinated or persistent threats [7].

Another limitation is the absence of severity scoring and contextual prioritization. Existing systems often label threats as anomalous or benign, but do not evaluate their impact relative to device criticality or patient safety [10]. For instance, a low-severity attack on a ventilator may pose a greater risk than a higher-scored anomaly on a non-clinical workstation, but traditional tools fail to capture such distinctions [11].

False positives and analyst fatigue also remain critical concerns. Studies highlight that IDS systems generate large volumes of redundant alerts, overwhelming SOC teams and delaying timely responses [5], [9]. Without correlation and reduction mechanisms, healthcare staff cannot focus on the most pressing incidents.

Furthermore, explainability is underutilized in decision-making. While recent approaches demonstrate interpretable models for anomaly detection [12], most correlation engines still operate as black-box systems, limiting analyst trust. Additionally, compliance mechanisms such as HIPAA tagging and audit logs are inconsistently integrated, often requiring manual effort [13].

In resource-limited environments such as Sri Lanka, where connectivity is unstable and IT resources are scarce, these limitations are amplified [6], [14]. Therefore, the need for an Adaptive Incident Correlation Engine (AICE) that unifies alerts, assigns contextual severity, reduces false positives, ensures compliance, and provides explainable, real-time dashboards remains pressing [15].

1.3 Research Problem

Increased integration of IOMT units on Sri Lanka's health exposes the system to cyber hazards, but the existing Siem integration is unable to provide localized, flexible and obedient alerts. Standard Gords lacks multilingual support, causing communication barriers between different employees [5]. They also produce generic notifications without urgent adjustment or references of compliance, causing fatigue and regulatory risk [1], [9], [12]. Connection addiction prevents offline scenarios, important in rural areas [3], [15].

It is especially acute in Sri Lanka, where hospitals work with limited resources, heritage systems and multilingual teams, which require effective SIM integration to handle real -time threats while ensuring patient protection and safety.

The research problem is: How to design a Siem integration that further normalizes the notice from (i) IOMT sources, (ii) provides multilingual, urgent information, (iii) automatic automatic tagging, and (iv) supports offline monitoring? Solving this will increase cyber flexibility in the Srilanka IOMT environment.

2 OBJECTIVES

The primary objective of the **Adaptive Incident Correlation Engine (AICE)** is to act as the decision-making hub of the IoMT cybersecurity framework by correlating heterogeneous alerts, prioritizing incidents based on severity, and providing actionable intelligence to healthcare security teams. This ensures faster and more reliable responses to threats, while reducing analyst fatigue and maintaining compliance with healthcare regulations.

The following sub-objectives support this goal:

- **Integration of Multi-Source Alerts**
  Collect and unify alerts from diverse sources, including IDS systems (Suricata, Zeek), AI anomaly detection engines, and device-specific logs. By consolidating heterogeneous alerts into a single pipeline, AICE minimizes fragmentation and improves visibility across the IoMT ecosystem.
- **Adaptive Correlation and Pattern Recognition**
  Develop correlation algorithms that identify relationships between anomalies occurring across multiple devices, network segments, and time intervals. This allows early detection of coordinated or stealthy attack campaigns that single-source monitoring tools may miss.
- **Severity Scoring and Prioritization**
  Implement a dynamic severity scoring mechanism that evaluates incidents based on device criticality, patient impact, and attack type. This ensures that life-critical devices such as ventilators or infusion pumps receive higher response priority compared to administrative or non-critical systems.
- **Decision Support Dashboards**
  Design intuitive dashboards for SOC analysts to visualize correlated incidents, severity levels, and recommended response actions. Multilingual support (Sinhala/Tamil/English) will be provided to improve accessibility in Sri Lankan healthcare environments.
- **False Positive Reduction**
  Reduce redundant and low-priority alerts through intelligent correlation and contextual filtering. This directly addresses the problem of alert fatigue, enabling analysts to focus on incidents that truly matter.
- **Regulatory Compliance and Auditability**
  Automate compliance tagging for frameworks such as HIPAA and Sri Lanka's data protection laws. Immutable audit logs will be maintained to support forensic analysis and regulatory reporting requirements.
- **Scalability and Resource Awareness**
  Ensure the system can scale to monitor 500–1000 IoMT devices while remaining efficient in resource-constrained hospital environments. Local fallback mechanisms will be included to handle unreliable internet connectivity.

.

3 METHODOLOGY

The development of the Adaptive Incident Correlation Engine (AICE) follows a structured methodology to ensure systematic integration, correlation, and decision-support in IoMT cybersecurity environments. The process combines IDS alerts, AI predictions, and medical device logs into a unified decision-making pipeline.

Step 1: Data Collection and Pre-Processing

Install and configure IDS tools (Suricata, Zeek) on IoMT gateways and hospital networks. Simultaneously, AI-based anomaly detection models (ML/DL frameworks) will monitor device behaviors. Raw data and alerts from these sources will be pre-processed and normalized into a standardized format using Python scripts and Logstash pipelines to handle heterogeneity across HL7, MQTT, and DICOM protocols.

Step 2: Multi-Source Alert Fusion

Design a data integration layer to combine IDS alerts, AI anomaly scores, and device-specific logs into a unified alert repository stored in Elasticsearch. This step ensures that all alerts are synchronized with timestamps, device identifiers, and severity attributes for efficient correlation.

Step 3: Correlation Engine Development

Develop correlation rules and algorithms that identify temporal, spatial, and logical relationships among alerts. This includes detecting attack patterns such as lateral movement, repeated probing, or distributed DoS attempts. Machine learning-based clustering may also be applied to group related anomalies.

Step 4: Severity Scoring and Prioritization

Implement a scoring model that dynamically assigns severity based on three factors:

- Device criticality (life-support vs. diagnostic vs. administrative)
- Patient impact (direct vs. indirect risk)
- Attack characteristics (malware, DoS, unauthorized access)
  This scoring system ensures that the most critical incidents are escalated first to healthcare SOC teams.

Step 5: Decision-Support Dashboards

Build an interactive visualization dashboard (using Kibana/Grafana) that displays correlated incidents, severity levels, and recommended responses. The dashboard will support multilingual notifications (Sinhala, Tamil, English) and provide compliance views for regulatory auditing.

Step 6: Compliance and Audit Logging

Automate tagging of incidents with references to HIPAA and Sri Lanka's Data Protection Act. Maintain immutable audit trails for forensic analysis, ensuring accountability and transparency.

Step 7: Testing and Validation

Conduct rigorous testing using IoMT attack datasets such as CICIoMT2024 and simulated hospital network environments. Evaluate system accuracy (precision, recall), false positive reduction, severity assignment reliability, and response times. Feedback will be gathered from healthcare security professionals to refine system usability.

Step 8: Deployment and Optimization

Deploy AICE in a controlled hospital test environment. Optimize performance to ensure scalability (up to 1000 IoMT devices), resilience in unstable connectivity scenarios, and compatibility with legacy systems.

3.1 System Overview Diagram

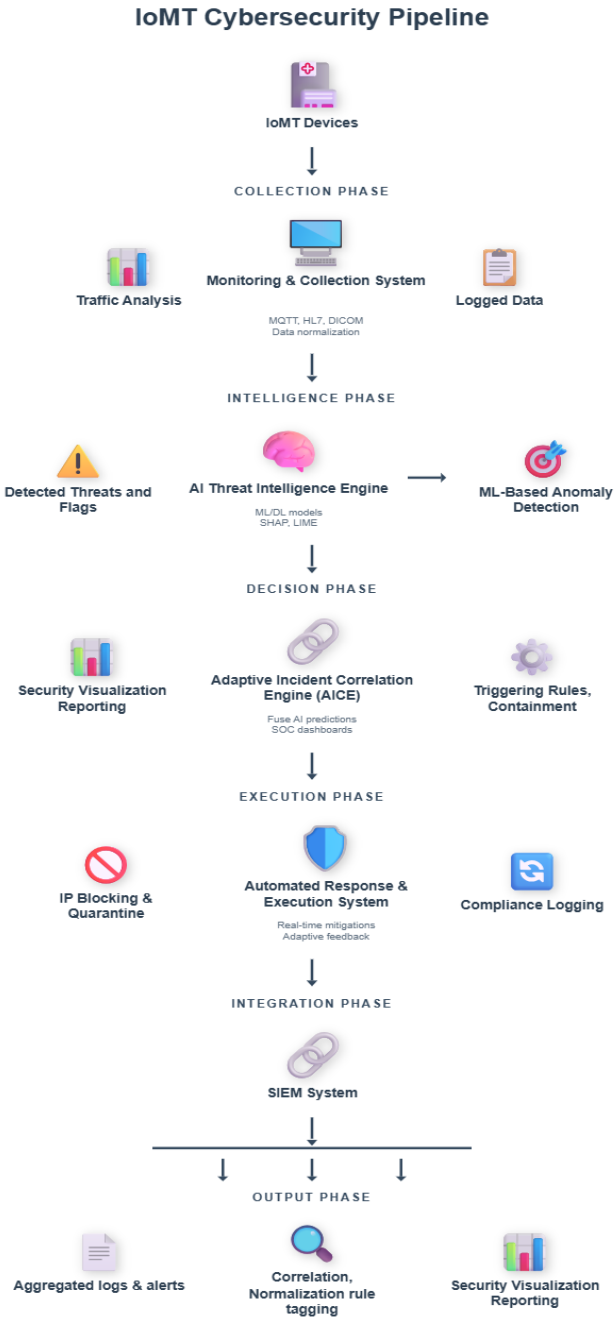

Figure 3.2 - System overview diagram

Above figure 3.1 shows the representation of proposed system diagram with all functions. The specific functions that are discussed in this document are shown in figure 3.2 component diagram below.

3.2 Component Overview Diagram

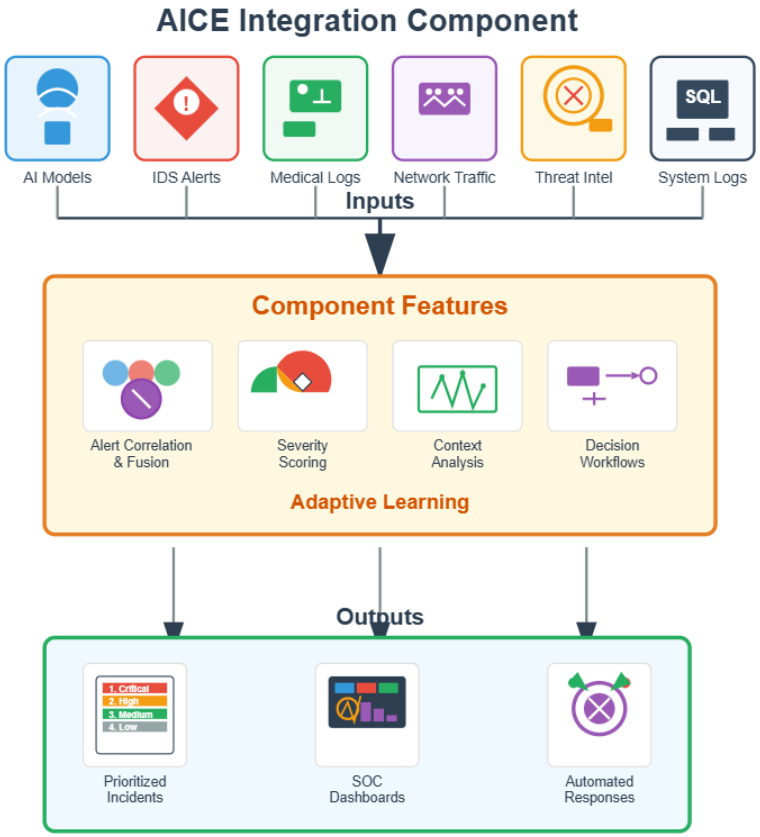

Figure 3.3 - Component overview diagram

4 TECHNOLOGIES TO BE USED

- **IDS Tools:** Suricata, Zeek – for detecting suspicious network activity in IoMT devices.

- **AI Models:** ML/DL frameworks (TensorFlow, PyTorch) – for anomaly detection and threat prediction.

- **Correlation Engine:** Python scripts – to fuse AI and IDS alerts and identify related incidents.

- **Databases:** Elasticsearch – for storing correlated alerts and incident logs.

- **Visualization:** Kibana / Grafana dashboards – to display incidents, severity scores, and recommended actions.

- **Compliance Layer:** Automated tagging – to ensure HIPAA and Sri Lankan data protection law adherence.

5 SYSTEM REQUIREMENTS

5.1 Functional requirements

- Real-time fusion of AI anomaly scores with IDS alerts.
- Correlation of related incidents.
- Dynamic severity scoring.
- Decision dashboards with escalation workflows.
- Immutable audit logs for compliance.

5.2 Non-functional requirements

- **Accuracy:** Reduce false positives by ≥30%.
- **Scalability:** Support 500–1000 IoMT devices.
- **Availability:** 99.9% uptime.
- **Usability:** Dashboards with Sinhala, Tamil, and English support.
- **Security:** Data encryption and role-based access control.

6 USE CASE SCENARIO

Table 6-1 - Use Case Scenario

Use case Name	**Critical Device Correlation & Escalation**
Actor	SOC Analyst, AICE Component
Goal	Correlate multiple suspicious events on an ICU ventilator, assign severity, and escalate for response
Trigger	AI detects anomalous traffic on IoMT device + IDS reports potential intrusion
Preconditions	- IoMT devices are monitored and generating logs - IDS and AI anomaly detection systems are active - AICE is running and integrated with SOC dashboards
Postconditions	- Incident is logged with severity and contextual details - Recommended mitigation actions are visible on dashboard - SOC team initiates response - Audit trail is maintained for compliance
Trigger	AI Detection Engine identifies anomalous behavior on ventilator system with potential patient safety impact
Basic flow	1. AI detects anomalous traffic and generates a threat alert. 2. IDS reports suspicious activity on the same device. 3. AICE ingests alerts from AI and IDS. 4. Engine performs correlation to identify related incidents. 5. Contextual analysis maps the incident to the ICU ventilator and patient-critical systems. 6. Severity score is calculated and assigned (e.g., Critical). 7. SOC dashboard highlights the correlated incident and recommends mitigation steps. 8. Escalation is triggered to SOC team or automated response system.

7 WORK BREAKDOWN STRUCTURE

The Work Breakdown Structure (WBS) provides a clear roadmap for the development and implementation of the Adaptive Incident Correlation Engine (AICE). It divides the project into distinct phases, ensuring systematic progress from requirement analysis to deployment. Each phase focuses on specific objectives, allowing for iterative testing, validation, and integration. The structured approach helps manage resources efficiently, maintain timelines, and ensure that all functional and non-functional requirements are addressed.

Figure 7.1 - Work breakdown distributed diagram

8 GHANT CHART

The project timeline spans 12 months with clearly defined phases, milestones, and dependencies. The schedule accounts for iterative development, comprehensive testing, and stakeholder feedback incorporation while maintaining realistic timeframes for each deliverable.

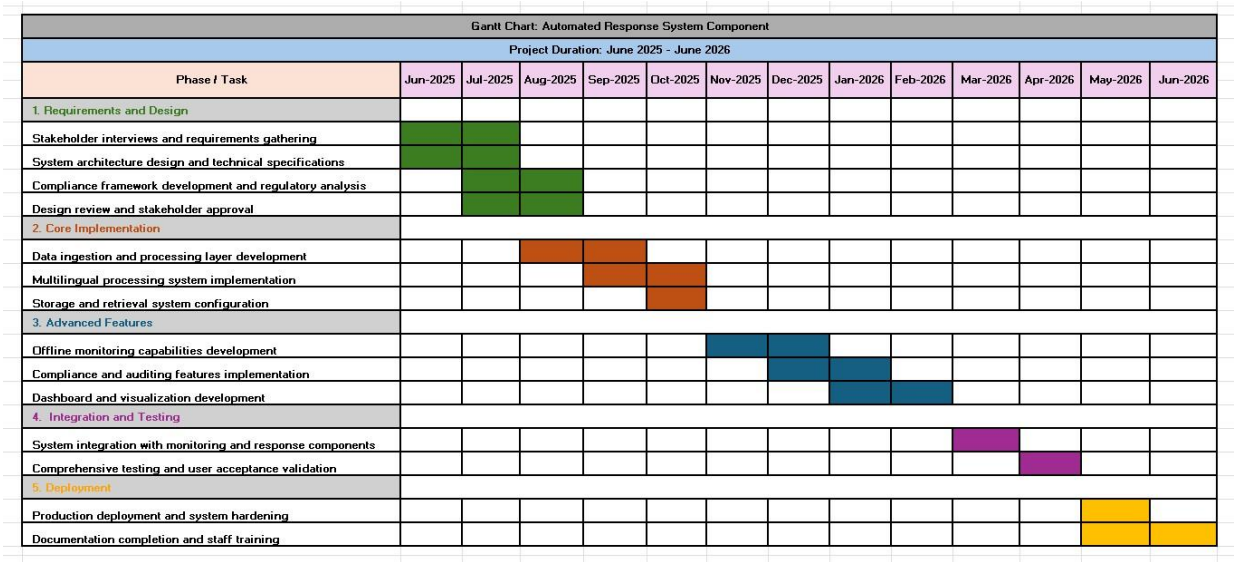

Figure 8.1 - Project gantt chart showing the timeline

9 BUDGET AND BUDGET JUSTIFICATION

Table 9-1 - Budget and justification

Category	Item Description	Estimated Unit Cost (LKR)
Hardware (Test Lab)	Server Computer	20000
	IoMT Simulation	10000
Software	ELK/Grafana Setup	5000
Software	IDS Tools (Open Source)	0
Development	AI/ML libraries	20000
Contingency	Miscellaneous & Contingency	10000
Total		65000

10 REFERENCES

1. Shaikh, J. A., et al. (2024). *RCLNet: an effective anomaly-based intrusion detection for Internet of Medical Things.* PMC. PubMed Central
2. Alalhareth, M., et al. (2024). *Enhancing the Internet of Medical Things (IoMT) Security via Ensemble & Meta-Learning Techniques.* PMC. PubMed Central
3. Hernandez-Jaimes, M. L., et al. (2023). *Artificial intelligence for IoMT security: A review of intrusion detection schemes.* ScienceDirect. ScienceDirect
4. Balhareth, G., et al. (2024). *Optimized Intrusion Detection for IoMT Networks with Tree-based Models.* MDPI Sensors. MDPI
5. Wang, C., et al. (2023). *Anomaly prediction of CT equipment based on IoMT data.* BMC Medical Informatics and Decision Making. BioMed Central
6. Elsayed, N., Dzamesi, L., & Ozer, M. (2025). *Extreme Learning Machine Based System for DDoS Attacks Detection on IoMT Devices.* arXiv. arXiv
7. Naghib, A., et al. (2025). *A comprehensive and systematic literature review on intrusion detection systems in IoMT.* Springer Link. SpringerLink
8. Doménech, J., et al. (2024). *Ensuring patient safety in IoMT: A systematic literature review.* ScienceDirect. ScienceDirect
9. Sheeraz, M., et al. (2024). *Revolutionizing SIEM Security: An Innovative Correlation Engine.* MDPI Sensors. MDPI
10. Lipsa, S., Dash, R. K., & Ivković, N. (2025). *Interpretable dimensional reduction technique with an explainable model for detecting attacks in IoMT devices.* Scientific Reports. Nature
11. Si-Ahmed, A., Al-Garadi, M. A., & Boustia, N. (2022). *Survey of Machine Learning Based Intrusion Detection Methods for Internet of Medical Things.* arXiv. arXiv
12. Almotiri, S. H., et al. (2025). *AI-driven IoMT security framework for advanced malware and ransomware detection in SDN.* Journal of Cloud Computing. SpringerOpen
13. Pratt, P., Chandekar, P., Mehta, M., Chandan, S. (2025). *Enhanced Anomaly Detection in IoMT Networks using Ensemble AI Models on the CICIoMT2024 Dataset.* arXiv. arXiv
14. Gupta, D., Gupta, M., Bhatt, S., & Tosun, A. S. (2021). *Detecting Anomalous User Behavior in Remote Patient Monitoring.* arXiv. arXiv
15. Osman, R. A., et al. (2025). *Energy-efficient communication between IoMT devices and emergency vehicles using 1D-CNN and optimization techniques.* PMC. PubMed Central