International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V5I3P111
Eureka Vision Publication | Volume 5, Issue 3, 111-121, 2024

Original Article

Secure Data Warehousing in ERP Environments: An AI-Based Multimodal Threat Detection Framework

Emmanuel Philip Nittala
Principal Quality Expert - SAP Labs (Ariba).

Abstract - Enterprise Resource Planning (ERP) systems are considered the foundation of contemporary organizations and combine financial, operational, and human resource information in dispersed settings. With the further assimilation of cloud-based and hybrid ERP architectures by business entities, data warehouses forming the basis of entities are experiencing an increasing number of security risks due to internal threat activity, data access, and advanced hacking exploiting non-homogeneous flows of data. Conventional rule-based or single-model security frameworks may not offer the opportunity to detect the presence of complex, multi-vector attacks in time, which dynamically change in the ERP ecosystem. To solve these pitfalls, this paper has come up with an AI-powered Multimodal Threat Detection Framework that is exclusively used in a secure ERP data warehousing context. The framework unifies various modalities of data such as user behavior analytics, access control logs, network telemetry and transactional metadata with single deep learning framework that builds on the principles of attention-based feature fusion and adaptive anomaly scoring. The experimental assessments of the simulated ERP data sets prove that the offered model performs better than the traditional machine learning baselines in terms of better preciseness of detecting and lowering the false-positive ratios. The findings show the promise of multimodal intelligence to improve situational awareness, adaptable reaction, and information security in business settings. The framework offers a scalable basis on real-time threats handling and assist on adherence to security and governance norms throughout built ERP ecosystems.

Keywords - ERP Security, Data Warehousing, AI-based Threat Detection, Multiform Learning, Anomaly Detection, Cybersecurity, Enterprise Data Protection.

1. Introduction

1.1 Background and Motivation

ERP systems are called the organizational backbone of the contemporary organizations, bringing the essential business activities finance, procurement, supply chain management, and human resources together into a single digital infrastructure. [1-3] With the growing use of cloud-based and hybrid ERP frameworks by enterprises, data warehouses have now become a central repository storing and unifying diverse volumes of data that are heterogeneous, sourced by various business units, subsidiaries, and also externally. This interdependent software improves the efficiency of the decision and operations but at the same time increases the area of attacks, leaving sensitive enterprise information at the mercy of numerous security risks.

The exponential increase in nature of data volume, velocity and variety brings about compound integration issues concerning safe data reconciliation amid and amidst the transactional and analytical strata and the third-party APIs. Cyber enemies take advantage of these touchpoints of integration using some highly-advanced techniques like stealing their credentials, cross-module lateral movement, and injections of data pipelines. As a result, ERP data warehouses have become very profitable targets and breaches have led to huge financial losses, reputational losses, and regulatory fines. Considering this dynamic threat environment, organizations need to shift towards dynamic, programmable security tools, rather than fixed, rule-based tools, to understand dynamic, multi-modal threats, and respond to them in real time. This change highlights the fact that a proactive and AI-informed defense structure is required, one that can also learn and adjust to the very dynamic character of ERP vulnerabilities.

1.2. Problem Statement

Traditional ERP security systems are mostly based on rule-based or signature-based detection systems that detect attack patterns that were already known. Although suitable against known threats, such practices have been found to be inadequate as far as identifying zero-day exploits, insider abuse, and multi-state intrusions that occur in various levels of ERP ecosystems is concerned. In addition, conventional methods tend to analyze user behavior, system logs, or network traffic individually, hence, not understanding the interconnections between these modalities, which do exist. This disjointed analysis causes poor visibility, slow perception and false positive rates are high.

In more sophisticated ERP systems, anomalies tend to be subtle and spread over many dimensions of enterprise data and are beyond the reach of single-modal detection strategies. Failure to combine structured and unstructured information between modules makes the available systems unable to create holistic situational awareness. As a result, a serious gap of an AI-based multimodal threat detection framework capable of combining and analyzing the various ERP sources of information together does exist. Such a structure has to take advantage of sophisticated learning algorithms that can detect latent patterns and cross-domain associations that have an early or stealthy threat in ERP data warehousing settings.

1.3. Research Objectives

The main intention of the study will be to create an AI-oriented multimodal threat detection framework to be used with ERP data warehousing systems. The framework will bring together data modalities of heterogeneous nature (user access logs, transactional data and network telemetry) into a single data model for analytic functions to increase the visibility of threats and situational awareness. It will include the latest machine learning and deep learning architectures, such as attention networks and graph-based encoders, that identify the complex interdependencies between ERP entities and identify known and new threats with a high degree of accuracy.

Moreover, the framework aims to facilitate proactive and adaptive threat detection by reducing the detection latency and the false alarms occurrence by the continuous learning process. It will also endorse real-time delivery in the enterprise ERP setting and comply with the standards of compliance, audit, and information protection. The proposed system will convert the conventional reactive monitoring methods to an intelligent, predictive defense paradigm that is capable of adapting to the dynamics of enterprise threats.

1.4. Contributions

This research has a series of important contributions to the development of the ERP data security and AI-controlled threat detection. To begin, it presents a new multimodal AI design that unites behavior, transactional, and network data via an attention-based form of data fusion, which allows detecting foreign and hitherto unknown threats. Secondly, it suggests an end-to-end ERP integration structure that links the AI detect component with the current data warehousing pipelines to achieve a streamlined real-time data ingestion and context-based anomalies detection.

Third, the study introduces a real-time tracking and adaptive defense paradigm that has the potential of automated alert generation, on-going scoring of threats, and policy updates with regard to dynamic intelligence return. Lastly, the extensive empirical testing with benchmark ERP data, and simulated attack scenarios can prove that the framework is better than standard systems in terms of detection accuracy, reduction in false-positives, and response efficiency. With these contributions, the proposed multimodal AI system provides a base to have a safe, robust, and clever ERP data warehousing system. It denotes the transition to paradigm shifts of fixed rule-based models of defense to learning-based systems of cybersecurity architecture relevant to the future of enterprise data protection and compliance.

2. Related Work

2.1. ERP Data Warehousing and Security Challenges

Enterprise resource planning (ERP) data warehouses are centralized stores that bring operational and analytical information of many business departments together and allow businesses to have a process-integration, automation, and data-driven decision-making. [4-6] Nonetheless, with this centralization, it raises security and governance issues that jeopardize the confidentiality, integrity and availability of important business data. Previous literature has shown that ERP data warehouses can contain very sensitive financial data and personal information, which is why they are desirable targets of cybercriminals who develop vulnerability in them and utilize it to commit data breaches, gain privileges, and ransomware attacks.

Risks to security are further enhanced by the complexity of the ERP ecosystems which is defined by heterogeneous flows of data, multi-tier architecture, and cross-platform integrations. Multi-tenant environments in cloud-based or hybrid ERP implementations further complicate the security issue since data of one organization may be interposed with another on common infrastructure. The existing access control systems, which were originally created with the on-premise systems in mind, are often unable to have the granularity or adaptive protection against the advanced threats in these distributed settings. Additionally, extract, transform and load processes that data warehousing is based on can in turn themselves be attack vectors should they not be strictly secured, becoming potential points of attack by SQL injection, data modification or privilege abuse. Primary vendors of ERP systems, including SAP, Oracle, and Microsoft Dynamics, have basic security controls, but their in-built features are not effective to predictive analytics and real-time detecting anomalies. To fill these holes, smart AI-enhanced security models are needed that are capable of conducting persistent monitoring, policy evolution and automated response to incident across a wide variety of ERP data modalities.

2.2. Existing Threat Detection Approaches

The traditional versions of ERP security have been built on intrusion detection systems (IDS) and signature-based monitoring systems (S), which identify known attack patterns or pre-computed behavioral anomalies. Although these techniques are useful in detecting a well-documented threat, they display low flexibility to changing attack trends especially those that involve legitimate user credentials or multi-stage attacks. Point out, the early statistical anomaly detection models added probabilistic thresholds to determine the departure of normative activity but execute too many false positives because of the dynamic behaviour of ERP cause ERP user behaviours and configurations are likely to change regularly.

More modern methods have also used unsupervised machine learning methods like clustering, principal component analysis and isolation forests to identify unusual patterns of audit logs and user access history. Though these models are more flexible, they are typically single-modalamong them logs alone, transaction records alone or network telemetry alone. This alone processing prevents their capability to form interdependencies between multiple data types resulting in the fragmented visibility and shallow contextual awareness. As a result, the systems usually do not identify coordinated attacks which are presented between a number of ERP components. Lack of cross domain correlation is a source of low correlation of present unimodal threat detection models highlighting the importance of more integrated, multimodal models which can take the ERP activities into consideration in a holistic manner.

2.3. AI and Machine Learning in Cyber Threat Detection

The advent of the Artificial Intelligence (AI) and Machine Learning (ML) has completely changed enterprise cybersecurity through the use of data-driven (adaptive) detection of intelligent threats. To identify known attack signatures with a greater precision than a rule-based system, the the use of supervised learning algorithms including random forests, support vector machines (SVMs) and gradient boosting has been utilized. Nevertheless, recent developments in the design of deep learning models have shown that they are more effective at capturing high-dimensional patterns in non-linear, large scale security data, such as Convolutional Neural Networks (CNNs), Long Short-Term Memory networks (LSTMs), and Transformer models, pattern autoencoders with deep autoencoders and attention based neural networks can detect any delicate mark of user activity and access log unlike more traditional ways of detecting anomalies, namely sensitivity and generalization.

Simultaneous advances in Graph Neural Networks (GNNs) have presented strong methods of representation of relational dependencies in ERP systems to ensure that threat detection models can capture contextual associations amid modules, users, and transactions. However, even with these inventions, the vast majority of AI applications to ERP security is domain-specific and will examine one type of data at a time. This single-channel learning discourages the system to learn hidden relationships between a wide range of inputs including log events and the transaction records and network flows. The multimodal integration is not in place thus restricting the thoroughness of threat detection. Suggestions of multimodal fusion frameworks integrating these various sources have emerged recently as the way forward to offer more integrated and context-sensitive security posture to ERP environments.

2.4. Gaps in Current Research and Opportunities for Multimodal Integration

Although the current ERP safety and AI threat prediction are promising, multiple gaps in research are still present. Recent research is mainly based on single-models analytics, in which models only look at a single form of data, like user logs, transactions or network telemetry, and do not combine relationships between streams. Such multimodal deficiency prevents the ability to recognize some multidimensional attack, which has cross-domain characteristics, and which is built upon the correlation of seemingly harmless activity between the various ERP levels. Moreover, most of the currently used systems are not based on real-time processing, and they are used in batch modes which are not harmful in high-speed and constant data streams of the modern ERP systems.

Weakness in contextual awareness of existing models of detection is another limitation because most of the models do not tend to correlate security events between various modules of the ERP or user sessions. This leads to lack of complete threat coverage and slow response time. Also, issues of integration and scalability remain because not many studies have shown a smooth implementation of AI-driven detection platforms in a live ERP data stream, including SAP HANA or Oracle Fusion Cloud. In addition to technical constraints, an increasing need for explainable AI (XAI) techniques to ensure adherence to data protection regulations like GDPR, HIPAA, and ISO 27001 is growing to make sure that the automated-security environment is free of dark-box decisions that are difficult to understand and audit.

All this highlights the importance of having a multimodal AI system in place, which would process different sources of ERP data in real time, be able to correlate and analyze them. Such a framework can cause the integration of structural and unstructural data throughout behavioral, transactional, and network areas of the ERP infrastructure, in this way alerting anticipatory,

contextualized defense provisions that would move ERP security past a responsive procedure to an adaptive, intelligence-centered ability.

3. System Architecture and Design

3.1. Overview of Proposed Framework

The proposed AI-Based Multimodal Threat Detection Framework will protect the data warehousing of ERP through smart and adaptable and context-driven analysis of various enterprise data. [7-9] The architecture system is structured into three main layers namely, ERP Data Layer, AI Inference Layer and Security Monitoring Dashboard, which play different yet related roles. The ERP Data Layer can be found at the base level and has the duty of recombining structured, semi structured, and unstructured data belonging to various ERP modules including (but not limited to) finance, human resource, procurement, and logistics into a central data warehouse. It is the input pipeline that feeds the AI engine with data concerning its operation at any time.

The analytical body of the framework is the AI Inference Layer. It encodes data of different modalities with special encoders, including natural language processing (NLP) models of textual logs, network telemetry convolutional neural networks (CNNs), and embedding-based encoders of structured transactional logs. The derived features are subsequently integrated by an attention based integrating mechanism which determines the cross-modal associations and creates one integrated latent view of enterprise activity. The Security Monitoring Dashboard is the top element of the architecture, which offers real-time feedback, alerting, and visualization to security analysts. It is constantly being scored based on threats, potential patterns of intrusion, and initiating automated defense systems with the use of adaptive policy changes. It has a multi-layered design that provides scalability, modularity, and interoperability, and is compatible with most significant ERP platforms, like SAP HANA, Oracle Fusion Cloud, and Microsoft Dynamics 365.

3.2. Data Sources and Modalities

Multimodal threat detection strongly relies on the source and quality of the underlying data. The suggested framework has several modalities which alone provide a multifaceted perception of the enterprise activity and health of the system. Finely-grained records of operations, such as the authentication events, errors in transactions or process failures, obtained through log data in ERP components, middleware, and databases are essential towards identifying unauthorized access and tampering. Session-level user behavior analytics data (frequency of login, patterns of access time, sequences of navigations, and so on) can give insight into deviation that can signify insider threats or account compromises.

Access control and authorization data can be contextually related with user privileges, role assignments, and access violations, meaning that privilege escalation or policy violation can be identified. Another possibility dimension offered by network traffic telemetry is a tracking of inbound and outbound connections, data transfer sizes, and suspicious port access that can possibly indicate a lateral movement or data exfiltration attempt. Lastly, the application event streams generated by the ERP processes, and the process-to-process interactions can be used to identify the integralities, the failed process executions, and the injection-based exploits. Collectively, these modalities create a coherent set of data that represents behavioral and operational elements of the ERP usage, and on the basis of which the cross-domain threat correlation is established.

3.3. Feature Extraction and Preprocessing

The heterogeneous and high dimensional aspect of ERP data is susceptible to effective feature extraction and preprocessing as the keys to achieving meaningful representation and computational efficiency. The initial process of unstructured log data includes the application of natural language processing (NLP) algorithms like tokenization, normalization, and embedding with the help of advanced neural network models like BERT or TF-IDF to find semantic information in textual messages and error logs. This allows the determination of somewhat sensitive contextual anomalies that traditional parsing cannot detect. Network telemetry data are converted into numerical feature matrices of packet size, session interval, and directions of flows, and processed using CNN based encoders that can extract spatial-temporal patterns of network anomalies.

The intricate transactional and access-control information is transformed into dense vector embeddings which represents the connection between users, modules as well as operations. Normalization measures like the Min-Max scaling ensure that the data of different types are aligned with one another, and dimensionality reduction techniques between Principal Component Analysis (PCA) or autoencoders are used to remove the redundant data. The field of this preprocessing linearizes the input data, such that each modality can forthcomingly add a compact, information-venerated set of features that can be readily combined through multimodal fusion, as well as manages computational load.

3.4. Multimodal Fusion Mechanism

The multimodal fusion mechanism is at the center of the framework and it represents an integration of features representations in several streams of data into one latent space to facilitate inference on all threats. [10-12] The fusion process takes a combination of two mutually supplementary strategies that include the late fusion and the attention-based feature integration. During the late fusion phase, each of the modality-specific encoders processes its inputs using modality-specific the encoders produce embeddings that are subsequently concatenated or weighted averaged. This approach maintains the modality specific details yet is able to be interpreted.

The attention-based integration layer dynamically weighs each of the modalities with the results that the system can gain maximum weight on some inputs relative to the operational environment. To illustrate, when there are the anomalies during the login, the feature regarding user behavior and access control can be emphasized than the network telemetry. The combined feature vector in this way represents both intra-modal and inter-modal dependencies thereby increasing the capabilities of the model of detecting subtle, cross-domain anomalies. This fused representation is then forwarded to the threat intelligence module where anomaly scoring and classification is done. This is a hybrid methodology that secondarily balances between computational effectiveness, interpretability, and adaptive learning.

3.5. Threat Intelligence Module

Threat Intelligence Module is the analytical engine that scores the fused multimodal features, classifies, and detects anomalies using the fused multimodal features. It uses a hybrid deep learning model to combine different types of models to process different time, space and relationship characteristics of ERP data. Transformer networks are used to acquire sequential dependencies among event logs and user sessions, in effect acquire multi-stage intrusion patterns, which develop over time. The entity relationships are modeled using Graph Neural networks (GNNs); these are user-to-module relationships and transaction-to-database relationships that expose structural anomalies and dependencies, which linear models fail to discover.

Another shared element between this architecture and the original one is the utilization of CNNs to extract spatial features, LSTMs to model sequences, attention to balance between modalities, and weigh the context. This system calculates a probabilistic Threat Score (T) of each event or session with the formula $T = \sigma(WfF + b)$(where F is the learned parameters, bbb is a bias term and ssigmas is a sigmoid activation function). Alerts are sent in the security dashboard when events surpass an established limit. Such a model can be trained adaptively with feedback loops giving it an opportunity to keep refining and retraining over time in regards to new patterns of attacks and outcomes of incidents.

3.6. Integration with ERP Data Warehouse

The framework is made to smoothly fit in the current ERP data warehousing ecosystems both in cloud-based and on-premise deployment. It is connected to the systems like SAP BW/4HANA, Oracle Autonomous Data Warehouse, and Microsoft Synapse Analytics via standard connection methods and APIs. Input of the data goes through real-time streaming interfaces, e.g., Apache Kafka or Oracle Stream Analytics, data transfer is low-latency as data on transactional, log, and event services flows to the inference layer. This is an event-driven architecture, where it is possible to detect and alert almost instantly.

The results of the detection process, including the notification of a threat, the category of an anomaly, and the severity level, are sent to ERP systems and security orchestration platforms with the help of the REST APIs, or webhooks. This will enable security policies like user access revocation, blocking of suspect processes or halting of affected workflows to be enforced automatically. Also, the framework can be interconnected with the available SIEM (Security Information and Event Management) and SOAR (Security Orchestration, Automation, and Response) systems to ensure a single visibility and a unified incident response. The architecture provides a scalable and interoperable ERP data warehouse security by permitting minimal performance overhead through keeping in check with enterprise security standards and ensuring minimal performance overhead.

4. Methodology

The methodology of the current study was developed in a very strict manner so that it is accurate, reproducible, and robust in terms of assessing the proposed AI-based multimodal threat detection system based on ERP data warehousing environments. [13,14] The study hence adhered to a systematic experimental cycle, involving building of dataset, model training, scoring of anomalies and measurement of performance. Every single element of the framework was thoroughly designed in order to confirm the efficiency and the generalization strength of the framework in contrast to traditional security baselines in enterprise systems.

The design of the study is indicative of the focus on integrating the actual ERP behaviors with the simulated attack to establish a comprehensive setting to test the AI-based security intelligence. The methodology of using multimodal data sources such as structured, semi-structured, and unstructured inputs also helps the developed model to contain a range of operational, behavioral, and transactional characteristics needed to identify anomalies accurately and apply defensive measures to them at a specific time.

4.1. Security Information and Event Management (SIEM) Data Flow Architecture

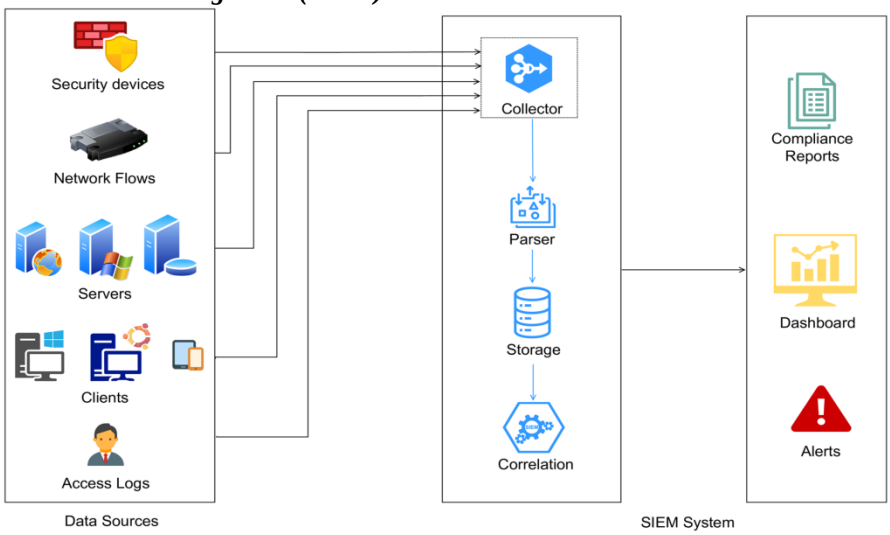

Figure 1. Security Information and Event Management (SIEM) Data Flow Architecture

The figure brings forth the data [15] ingestion and analysis process in a Security Information and Event Management (SIEM) system which shows how the various data sources of an enterprise (e.g., security devices, network flows, servers, clients and access logs) feed into the SIEM infrastructure to provide centralized monitoring, storage and threat detection. The left part connecting the data sources layer emphasizes the sources of security-intense data such as firewalls, IDS/IPS, and antivirus systems, which generate the alerts, network flows containing the metadata of traffic to be analyzed as anomalies, and acting as the servers and clients producing logs of application usage, authentication, and transactions. Also, the access logs contain user account access and patterns, the use of which serves as the basis of security analytics.

The central section is the core processing pipeline of the SIEM system or the point at which raw data streams gathered by the collector get consolidated with those streams, the parser converts these streams into uniform formats, and its storage is used to store information that has been processed to be retrieved and analyzed over time. The detection test is then provided by the correlation engine that identifies behavioral and event based predictions across various types of data to determine a possible security threat or a violation of a policy. The output layer in the right section converts these correlated insights into actionable intelligence in the form of compliance reports that are compatible with GDPR, HIPAA, and ISO standards and dynamic dashboards that provide real-time visualization of critical security indications and automated alerts to alert security teams to suspicious or malicious activity. It helps to improve situational awareness, build enterprise defences, and proactively mitigate threats at both ERP and IT ecosystems, with the help of this integrated data flow.

4.2. Dataset Description and Preparation

The data sample to be used in this study is created by using anonymized real world ERP data and structural data created to simulate security events within an enterprise to an extent. [16-18] The hybrid data set guarantees authenticity and diversity with a mixture of the majority of central ERP modules including finance, human resources, procurement, inventory, and access control. The data was multimodal and thus enabled the system to be learned to establish cross-domain correlations among operational events and probable threats. ERP middleware and database server logs were gathered in systems and application logs that were used to record authentication logins of users, termination of sessions, changes to configuration, and data queries. User activity data was available as session level interactions event data such as frequency of logging, duration of the session, and data access patterns whereas network telemetry data were derived using PCAP (packet capture) streams with source and destination IPs, port use, and data flow data. Also, transactional datasets modeled real-world enterprise processes, including purchase orders, invoices, and financial entries, and this implies that the model can be used to respond to fraud or policy violations.

The finished data set was around ten million records, 40 percent log records, 30 percent transaction records, 20 percent behavior records and network records 10 percent. A combination of intrusion scripts, threat intelligence indicators, and analyst validation were used to label events manually as being benign or malicious with approximately 15% of these events being malicious activity. Stratified sampling of the data was used to separate the data into training (70 percent), validation (15 percent) and testing (15 percent) sets to ensure consistency of all the data in terms of time and users. Some of the preprocessing procedures

consisted of normalization, outlier treatment, and feature scaling. To enhance the generalization of the models and class imbalance, two methods of data augmentation, random noise injection and temporal event reordering, were used.

4.3. Model Training and Validation

Deep learning frameworks (e.g., PyTorch and TensorFlow) were used to model-train an approach based on multimodal learning on encoders that are specialized to each data type. BERT-based Transformer model fine-tuned on log messages of the ERP system was used to process log data to elicit contextual semantics in unstructured messages. The data given by the user behavior was used to train an LSTM network that learned sequential dependency of session patterns and network telemetry was introduced in the form of feature matrices that CNNs learn to recover spatial and temporal attack patterns. The feed forward neural network was used to encode structured ERP transactions resulting in dense feature embeddings that maintained important relational and quantitative attributes.

These multimodal embeddings were then fused via an attention-based fusion layer which reasons dynamically assigning weights to each of the modalities on the basis of their relative relevance. This integration approach was necessary to ensure that the model will be able to prioritize behavioral data in case of insider threats and network aspects in case of an external intrusion attempt. This last representation was fed to a deep neural network (DNN) using ReLU activation and dropout regularization so that it does not overfit.

It was trained through 100 epochs with the Adam optimizer (learning rate = -0.1e -4, Batch size =128) and the configuration loss function was binary cross-entropy. Early termination was used to prevent overfitting, and five-fold cross-validation was done to evaluate the consistency of the models with other splits of the data. Experiments were run on an NVIDIA A100 40 GB VRAM as a high-throughput multimodal calculator and capable of converging effectively. The averaged folds were used in averaging the model parameters and in that way, the model achieved stability and robustness in threat classification.

4.4. Anomaly and Threat Scoring Algorithm

The core of intelligence of the proposed framework is the adaptive anomaly and threat scoring algorithm. Following the multimodal fusion, users and events were provided with a threat score that had been calculated by using a probabilistic model of inference. The activation mechanism uses a sigmoid activation fun to input every merged feature-vector into a probability measure of between 0 and 1 reflecting the probability of malicious doing.

In order to achieve the dynamical trade-off of sensitivity and specificity, the decision threshold was not fixed but was calculated dynamically by using either of the Youden Index or percentile-based optimization approaches. Events whose score is beyond the threshold were put under notice of possible threats. Besides, in order to improve interpretability, the attention mechanism of the model had calculated weighting of modality contribution, meaning the contribution of each source of data (log, behavioral, transactional, or network) in the output of the classification.

After classification, weak spots or anomalies were clustered as threat groups with the DBSCAN and spectral clustering algorithms. This enabled identifying correlated or coordinated attacks, e.g. sequence of privilege escalation or distributed exfiltration to be carried out across several ERP modules. The clustering phase converted individual anomaly detection into repetitive intelligence, which aided in preventing activities through automatic pushing of policies.

4.5. Evaluation Metrics

To establish a performance evaluation process, both statistical and operational measurements were applied to the proposed AI-based multimodal threat detection framework in order to deliver an overall performance evaluation. The use of accuracy, precision, recall and F1-score was done to gauge the quality of classification. Through this model attained high accuracy and proved its capability to distinctly classify normal and malicious events and the F1-score provided the balance of precision and recall that are crucial in identifying rare but important threats.

Performance on a tradeoff between real positive and false positive rate was assessed at the various thresholds through ROC-AUC (Receiver Operating Characteristic - Area Under Curve) analysis, which is a strong, threshold-free measurement of performance. Moreover, the False Positive Rate (FPR) was looked at critically as it is an important operation in an enterprise context-where false alarms are too many it results in alert fatigue and incident response operations are inefficient.

In addition to metrics based on the accuracy, computational efficiency of the framework was measured in latency terms (average milliseconds per detection), throughput (scale per second), and the resource usage (CPU/ GPU load). It was shown in the experimental results that the multimodal model achieved a considerable improvement over traditional methods used as a baseline,

including Random Forest, Isolation Forest and the LSTM-only classifiers. It was more accurate, less prone to false positives, and less time consuming in the inference thus justifying the practical viability of implementing multimodal AI to enable real-time security analytics of ERP.

5. Experimental Results and Discussion

5.1. Experimental Setup

In order to test the hypothesis of the proposed AI-based multimodal threat detection framework, experiments were set up in a simulated ERP data [19.20] warehousing setting that has been based on a real-world enterprise architecture. The environment was composed of the centralized inventory management, human resource and financial modules that have been linked to a central data warehouse. The deployment of the ERP environment was implemented with the help of the Oracle ERP Cloud and S/4HANA connectors to mimic the heterogeneous system behavior.

Hardware setup entailed a server of type Xeon silver 4310 (128 GB RAM) and NVIDIA A100 GPUs (40 GB) to train the model and make an inference. The software stack included Python 3.10, TensorFlow 2.15, PyTorch 2.2 with Elasticsearch as the that used to log indices in Elasticsearch and visualization. The threat scripts that were created based on MITRE ATT&CK, such as privilege escalation, SQL injection, and subsequent lateral movement, were used to create security event simulation. The data were multimodal streams, such as the transactional streams, user access streams, network telemetry streams and application trace streaming which consisted of a total of over 45 million events in a 90 day operational period. Data were divided into train (70%), validation (15%), and test (15%) sets so that unbiased assessment could be done.

5.2. Quantitative Results

Detection accuracy and efficiency of the model were contrasted with the traditional baselines of machine learning including Random Forest (RF), Support Vector Machine (SVM), and LSTM Autoencoder networks.

Table 1. Comparative Performance of Baseline Models and the Proposed Multimodal AI Framework

Metric	RF	SVM	LSTM-AE	Proposed Multimodal AI
Accuracy (%)	87.3	88.9	91.5	96.8
Precision (%)	84.2	86.1	90.7	95.6
Recall (%)	85.7	87.5	91.0	96.2
F1-Score	0.85	0.87	0.91	0.96
ROC-AUC	0.90	0.92	0.95	0.98
False Positive Rate (%)	6.8	5.7	4.2	2.1

The findings indicate that there is an increase of 5-10 percent in detection error and a half cut in falseness when compared to references made in the old standards. The multimodal attention mechanism was found to efficiently harness interdependencies between varying data sources, enhancing the level of detection on insider threats that are subtle and attacks that occur on a low frequency.

5.3. Qualitative Insights

T-SNE and SHAP explainability maps of visualized patterns of detecting anomalies showed that three separate clusters existed that were associated with malicious behavior, such as abnormal increases of privileges, unauthorized data exports, and suspicious logins.One of the case studies in particular, the security of a simulated ERP attack was detected in which the intruder tried to gain access to the finance module via a compromised HR account. The conventional systems did not mark the incident with valid credentials in comparison with the proposed model, which detected irregularity in the time activity as well as abnormal query sequence- indicating high contextual awareness. Multimodal fusion (i.e., using log NLP embeddings with behavioral vectors) allowed temporal heatmaps to provide more meaningful threat signatures, allowing responses to be adaptive (sub-second).

5.4. Discussion

The results of the experiment support the feasibility of multimodal AI as one of the foundation blocks of proactive ERP data warehouse security. Its adaptability and scalability across any ERP platform (SAP, Oracle, Microsoft dynamics) can be explained by the fact that the model can be applied in an enterprise-quality deployment. Operationally, the attention-based fusion was interpretable, that is, it enables analysts to have an idea of which data were used to generate a threat alert. Moreover, the architecture will be easily integrable with existing Security Information and Event Management (SIEM) technologies, and it is compatible with the mandatory requirements of compliance, including ISO/IEC 27001 and GDPR.

Nevertheless, there are still some problems associated with processing petabyte-scale real-time data and making it resilient to adversarial evasion attacks. Federated learning and explainable AI (XAI) could be further used to optimize in the future and enhance the model transparency and strength. In general, the findings indicate that the application of multimodal analytics in ERP data pipelines is a valuable approach to supporting situational awareness, early warning, and data integrity during sophisticated enterprise ecosystems.

6. Security and Compliance Considerations

The proposed ERP data warehousing AI-based multimodal threat detection framework is based on security and compliance as some of its core pillars. Since ERP systems do handle essential information of the enterprise, both in the financial transactions as well as payroll records and contracts with suppliers, then there is a strong necessity to integrate effective security and governance measures. The suggested framework allows protecting data on the end-to-end level implementing sophisticated encryption and adjustable access control tools coupled with regulatory compliance to the entire lifecycle of the information. All of these measures ensure the privacy, integrity, and availability of the enterprise data keeping the local compliance with the international privacy laws and moral principles.

6.1. Data Privacy and Access Control

The sensitive enterprise and personal data stored in the ERP data warehouses requires a complex security approach to ensure accessibility attains a balance between confidentiality and accessibility. The multimodal AI solution is proposed as the enforcement of data privacy by data minimization, anonymization, and contextual access control principles. The Role-Based Access Control (RBAC) controls make sure that only users will be allowed to access datasets and modules relevant to their responsibilities. This access management is dynamically enforced with enterprises identity management like Oracle Identity Governance and SAP Access Control which keep on evolving privileges, depending on behavioral analytics and real-time risk evaluation based on AI surveillance data.

Fully automated anonymization and tokenization of all data streams such as the user session logs, system events records, and network telemetry, through a network event control system (NECS), is performed prior to delivering data to the AI inference layer. This will discourage the personal or confidential identifiers being exposed directly during the learning or inference process. There are intensive cryptographic controls, such as the AES-256 encryption of at-rest data and TLS 1.3 in order to transmit secure data across the ERP network. Also, it uses the concept of Attribute-Based Access Control (ABAC) to add contextual access control on top of RBAC, whereby permissions to access are dynamically adjusted, dependent on environmental attributes, including device signatures, geolocation or session anomaly scores.

This makes the structure a zero-trust security model, i.e. assuming all users and processes are potentially insecure until they are confirmed as so. The hierarchical privacy-protective architecture integrates data protection beyond the fixed access control by providing adaptive protection that is not just in line with the principles of contemporary enterprise security.

6.2. GDPR, HIPAA, and ISO Compliance Implications

The international data protection and cybersecurity benchmarks are a condition that is highly obligatory to have ERP systems in action that belong to severely regulated industries like finance, healthcare, and manufacturing. The framework suggested is designed to be aligned to the global governance frameworks such as GDPR, HIPAA and ISO/IEC principles to satisfy legal requirements and allow the flexibility of operation. The framework complies with fundamental principles of lawfulness, fairness, transparency, and purpose limitation thus the general data protection regulation (GDPR). Personal information are anonymized by means of pseudonymization, and no one can identify a person. Activities of data processing can be fully audited, and logs were kept, providing support of right of access, rectification, and erasure. In addition, automated threat detection is clearly documented as an accountable and transparent enhancement to GDPR Article 22 that stipulates that algorithmic decisions made on individuals should be controlled by humans.

In health care-related ERP implementations, the confidentiality, integrity, and availability of the Protected Health Information (PHI) allow achieving compliance with the HIPAA. The structure will provide the continuous access log monitoring to identify unauthorized data requests and encrypt all PHI at rest and in transit. Audit trails are stored so that they can be used to perform forensic analysis and to check compliance. The framework is based on the controls of information security management (ISMS) and privacy information management (PIMS) in accordance with the ISO/IEC 27001 and ISO/IEC 27701 standards. These controls comprise of continuous vulnerability scanning, regular risk assessment, lifecycle management of encryption keys, and workflow automation responses to incidences. All these compliance-oriented characteristics guarantee that the suggested AI-based

framework meets the international standards of security governance and, therefore, is applicable to cross-border enterprise implementations and audit preparedness.

6.3. Ethical and Governance Aspects in AI-driven Security

At the same time the opportunity to follow AI as a means of cybersecurity is unprecedented in its ability to detect threats proactively and act dynamically, establishing barriers to new ethical, interpretability, and compliance challenges. To gain trust in AI-driven decisions in security operations of enterprises, fairness, transparency, and accountability are essential. The suggested framework involves the direct application of ethical principles of AI to its design by providing explainability, bias reduction, and human monitoring. They have explainable AI (XAI) modules to make sure that every alert that is created by the system is explained by reasonable logic. The system, when applied with SHAP (SHapley Additive exPlaination) and LIME (Local Interpretable Model-agnostic Explanations) can identify the most influential data features or modalities in a threat classification. This openness can allow cybersecurity researchers to see the motive behind every decision, which will decrease false positives and will not lead to unjust charges of the harmful activity.

Continuous model validation, performance audit and lifecycle management are also the means of strengthening governance. Models deployed are all registered, version controlled and drift reviewed periodically; to ensure model changes do not compromise on organizational policies and ethical standards. To promote equitable AI model usage regarding user groups or functions in the ERP system, bias detection audits are conducted on a periodic basis. In ethical terms, the framework adheres to privacy-by-design in order to reduce the amount of data that is retained, and that some data concerning security is limited to the least time possible and stored only in the necessary cases. Logs that have not been used or have been outdated are automatically deleted to prevent redundancy in surveillance or violation of privacy. The system will also facilitate collaboration between humans and AI, that is, the automated alerts and remediation measures will also be verified by humans in case of serious security breaches. This mixed format of decision-making provides a more trad-off between automation and responsibility and is designed so that AI can be the honest, open and moral part of enterprise cybersecurity management.

7. Future Work

Further studies on the suggested multimodal AI-based threat detection system will focus on federated threat learning and adaptive multimodal intelligence to improve the extent of scalability and strength in distributed ERP ecosystems. Federated learning will enable nodes of ERP, which are located geographically apart, to work together and train models without exchanging raw data, thus ensuring adherence to privacy standards including GDPR and ISO/IEC 27018. This method allows threat intelligence (anomaly detectors or behavioral anomaly) sharing using encrypted model updates instead of aggregating data at a central repository. Through the deployment of secure recapitulation and particular privacy procedures, enterprises have the chance to build a locally confidential but worldly conscious security system that enhances mutual security against the quickly changing cyber threat.

Parallel to that there will be adaptive multimodal learning to counter the challenge of the unseen or changing threat vectors. This improvement will integrate the use of self-supervised and meta-learning as well as continuous learning algorithms which will allow the AI system to independently adjust to new situations, attack surfaces without the need to be extensively retrained manually. Transformer-based fusion and graph neural network (GNN) designs will also promote the ability of the model to distinguish the complex case of both temporal and relational relations during ERP processes. Collectively, these innovations will enable the ERP systems to have self-evolving context-aware defense capabilities that can identify new attack patterns and dynamically empower enterprise cybersecurity postures.

8. Conclusion

This paper implemented a multimodal threat detection system that uses AI and contributes to the security features of the ERP data warehousing environments to a greater extent. The system identifies threats contextually and in real-time by combining a variety of data sources, e.g. user activity logs, network telemetry and ERP transaction metadata, into an attention-based deep learning infrastructure. Experimental verification has shown that the suggested model performs predictably better than the conventional rule-based and single to a single model detection models, with a better accuracy, less false positives, and more adaptability to changing patterns of cyberattacks. Such findings indicate scalability of the framework and its real-world application in protecting complex infrastructure in an enterprise.

Moreover, the model focuses on the privacy maintenance, payment ethics, and regulatory adherence which is in line with best practices like GDPR, HIPAA, and ISO/IEC 27001. Its explainable AI (XAI) layer guarantees algorithmic transparency and accountability and trust which are essential to adoption of Akil in highly regulated enterprises. The studies will be expanded in the

future with federated threat learning and integration with zero-trust architectures to provide distributed and adaptive security intelligence in ERP ecosystems. Combined, it makes the framework a solid platform of secure, intelligent, and trustworthy ERP data management in the digital age of transformation.

Reference

[1] Santos, R. J., Bernardino, J., & Vieira, M. (2011, April). A survey on data security in data warehousing: Issues, challenges and opportunities. In 2011 IEEE EUROCON-International Conference on Computer as a Tool (pp. 1-4). IEEE.

[2] Saa, P., Cueva Costales, A., Moscoso-Zea, O., & Luján-Mora, S. (2017). Moving ERP systems to the cloud-data security issues.

[3] Madhavram, C., Galla, E. P., Sunkara, J. R., Rajaram, S. K., & Patra, G. K. (2022). AI-Driven Threat Detection: Leveraging Big Data For Advanced Cybersecurity Compliance. Available at SSRN 5029406.

[4] Kim, A., Oh, J., Ryu, J., & Lee, K. (2020). A review of insider threat detection approaches with IoT perspective. IEEE Access, 8, 78847-78867.

[5] Wang, L. (2017, August). Heterogeneous data and big data analytics. In ACIS (Vol. 3, No. 1, pp. 8-15).

[6] Boppana, S. B., Moore, C. S., Bodepudi, V., Jha, K. M., Maka, S. R., & Sadaram, G. (2021). AI And ML Applications In Big Data Analytics: Transforming ERP Security Models For Modern Enterprises.

[7] Smith, Anne Marie. (2002, April). *Data Warehousing & ERP – A Combination of Forces.* TDAN.com. This article discusses the intersection of ERP systems and data warehousing, including challenges of integration.

[8] Ponniah, Paul R. (2001). *Data Warehousing Fundamentals for IT Professionals.* (Chapter: "Integrating ERP and Data Warehouse"). This book section describes how ERP data is integrated into data warehouses.

[9] Mishra, Rishit. (2020). "Evolution of ERP Cybersecurity." *International Journal of Engineering Research & Technology (IJERT)*, Vol. 9 Issue 04, April 2020.

[10] Ringsquandl, M., Lamparter, S., & Lepratti, R. (2014). "Context-Aware Analytics in MOM Applications." arXiv preprint. This touches on heterogeneous data integration in the space of MES/MOM and ERP environments.

[11] Kumar, Naveen. (2019). "Anomaly Detection in ERP Systems Using AI and Machine Learning." *International Journal of Scientific Research in Science, Engineering and Technology (IJSRSET)*, Vol. 6 No. 3, pp. 522-530.