LeakAlert - Insider Threat & Anomaly Detection System

Sakshi Bansal[1], Anushka Kashyap[2], Khushi Gupta[3], Swayam Nain[4], Ms. Prerna[5]
[1,2,3,4,5] *Department of Computer Science and Engineering Meerut Institute of Engineering & Technology Meerut, India*

Abstract: Organizations nowadays are very dependent on internal computer systems and there is a need to keep track of how users behave and what they are doing in the systems. Despite the fact that the majority of security tools have been developed to prevent external attacks, a significant number of security challenges are caused by insiders who have already gained access to the systems either purposefully or through some unintentional practices. In this project, a basic system of insider-monitoring was established to track the attempts of logging in, information about the device, IP location, file access pattern and other suspicious patterns in real time. The system computes a risk score based on failed logins, access to restricted files, new device or unusual location logins among others. The system is not based on labeled data but rather a combination of rule-based checks with unsupervised machine-learning models to identify abnormal behavior. An Isolation Forest model is used to examine the patterns of logins, device changes, change of location, and file-access activity to determine the abnormal behavior of the user. A dynamic risk score is created by combining the model output with the real-time security rules, which enhances the accuracy of the detection without proving to be an unrealistic concept when implemented in real-world enterprise settings.

Index Term - Anomaly Detection, cybersecurity, insider threat, Isolation Forest, machine learning.

I. INTRODUCTION

Many people tend to perceive cybersecurity threats as those posed by external parties but the reality is that most of them are initiated in an organization. The legitimate access by employees or trusted users makes it more challenging to the traditional rule-based systems to identify harmful activities. These internal risks can accumulate over time, e.g. by attempting to get access to certain files several times, accessing unusual logins, or several unsuccessful attempts to enter the passwords. Because organizations produce thousands of login and file related events every day, it is not feasible to review all the activities manually. Hence, it is necessary to have automated monitoring systems that will constantly monitor system users and immediately detect suspicious behaviors. With the emergence of remote work, users have been able to access systems across various locations and devices, thus making it harder to identify unauthorized access according to the traditional approach.

To address these challenges, this project presents a practical insider monitoring system that combines login tracking, location checking, device identification and file access control. Instead of relying on heavy ML models, it uses simple real-time rules and a dynamic risk score. The goal is to immediately catch anything out of the ordinary and give admins clear information to act on.

I.1 Motivation and Design Objectives

Organizations in the modern world require insider threat detecting systems that operate in the background without affecting the normal operations. Some of the currently available solutions are however hard to deploy or rely on labelled datasets, which in practice are not always available in the real world.

Consequently, there is evident disparity between theoretical and applied enterprise security needs of research. The primary goal of the work is to create a light-weight, however, efficient system of insider threat monitoring that could be installed in the organizational setting with the smallest amount of setup. The system will be built to identify suspicious activity in real time, provide explainable alerts to the administrators and provide swift response measures like suspending users. The proposed approach balances the usefulness and practicality of activities, precision of detection, and interpretability by combining rule based security checking with unsupervised methods of machine learning.

II. RELATED WORK

A significant number of researchers studied insider threats in terms of analyzing the behavior of normal users and its development. An example is a study conducted in [1], which investigated real time behavioral patterns and used the methods of clustering to determine when a user was abnormal to the routine. The findings showed that a constant surveillance will aid in ensuring abnormal behaviors are noticed at a tender age. In another study [2], the authors of the article concentrated on the management of insider risk and suggested that the level of risk of a user should not be fixed but should fluctuate. Under this method, the risk score will increase or reduce depending on the behavior of the user at the moment. This theory of real time risk revising at a later stage inspired the development of a number of insider threat detection systems. The study of insider behavior has also been performed with the help of deep learning techniques. In [3], the authors examined sequence of user activities with an aim of detecting subtle changes that can signal misuse. They found that any minor differences between the normal behavior of a user can indicate a possible threat. The other method as described in [4] utilized machine learning to conduct analysis on the system logs, file access and permission histories. These features were applied to determine behavior patterns which could be an indication of risky or illegal behaviors in an organization. In these studies the similar observation is that the detection of insider threats is enhanced when multiple types of the user details, such as log in details, files access, device use and location are considered jointly. Nevertheless, most of the research techniques rely on complicated models or need labelled data, which are not always accessible and feasible in the case of ordinary organizations.

1.1 Comparison with Existing Approaches

Existing insider threat detection systems can generally be classified as rule-based, supervised machine learning, or unsupervised and hybrid approaches. Rule-based systems provide transparency but struggle to detect new or evolving attack patterns. Supervised learning methods can be accurate, yet they depend on labeled insider threat datasets that are hard to obtain due to privacy concerns and the rare nature of such incident

Unsupervised methods, including clustering and anomaly detection, address this limitation by learning normal behavior patterns without the need for labeled data. However, many earlier studies depend on computationally intensive deep learning architectures, which increase deployment complexity and limit interpretability.

The proposed system differs by integrating lightweight anomaly detection using Isolation Forest with real-time rule enforcement. This hybrid strategy achieves a balance between detection accuracy and system simplicity, making it suitable for practical enterprise use where transparency and fast response are critical.

III. PROPOSED METHOD

III.1 System Overview

Previous works were mostly based on behavior modeling or machine learning to identify insider activity including clustering or pattern recognition using deep learning [1]-[4]. On the contrary, the offered system is simpler and more realistic. Instead of relying on complicated ML pipelines, it gathers continuous data of the log-in, device fingerprints, IP location and attempted file-access with a combined rule-based and behavior based design. Our approach is inspired by the adaptive scoring logic presented in [2] and the continuous behaviour analysis presented in [1], and is a combination of lightweight anomaly rules and real-time checks. This would assist the system in real-time pointing out suspicious activity and creating a dynamic risk rating. It is to provide a solution that is simple to implement yet can still track activities of users during the day. The proposed methodology is a low-training pipeline with no complicated feature engineering by avoiding the expensive training pipelines and feature engineering. computational overhead and facilitates the integration with the existing organizational systems.

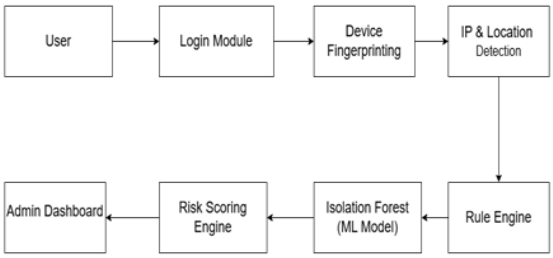

Figure 1. High-level architecture of the proposed insider threat monitoring system.

Figure 1 illustrates the overall system workflow, showing how user activities are captured, analyzed through rule-based and machine-learning modules, and visualized on the administrative dashboard. This architecture enables real-time monitoring and rapid response to suspicious behavior.

III.2 Login Security Flow

Past research has shown that any unexpected change in the device, abnormal time logs-in or the aberrant behavioral patterns can be an indication of insider abuse [1]-[3]. Developing this concept, the system verifies the attempts of logging in by analysing the device fingerprints, the geographical information, and the history of the user logging in to the system. Put together, the system is capable of detecting authentication attempts that do not comply with the patterns of use and marking them as subject to further investigation. In case a login happens on a new device or on a new location, the system processes it as a flag of a strange behaviour, as was previously seen in the behaviour-deviation approach in [3]. The failed attempts to log in on a few occasions automatically raise the risk score of the user according to the principle of adaptive. scoring highlighted in [2]. This activates live monitoring as opposed to the waiting of logs that will be analysed later.

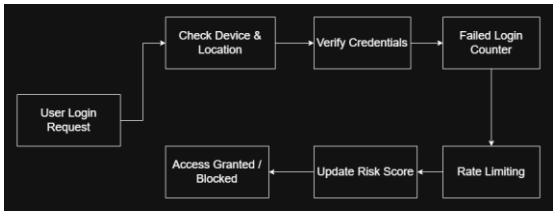

Figure 2. Login security flow with anomaly detection and rate-limiting.

Figure 2 demonstrates the workflow of the system where the authentication requests are compared with the help of the device fingerprinting, geographic location analysis, and user login history. Another major point in the workflow is the limitations in the cases of failed logs in and dynamical updating the risk score to avoid brute-force attacks and to identify suspicious logins in real time. The system applies rate-limiting of failed logins to avoid brute-force. In case the number of login failures goes over a set limit in a time frame, the subsequent attempts are blocked temporarily and recorded as high-risk events. This process prevents user accounts as well as increases anomaly detection.

III.3 Machine Learning–Based Anomaly Detection

In order to go beyond static rules during detection, the system includes an unsupervised machine-learning model to analyze the behavior. It learns the normal user behavior pattern using an isolation forest algorithm without the need to have labeled attack data.

The model is trained based on historical features of activities like frequency of login, number of failed logins, device change rate, location change.

and file-access behavior. At the time of live operation, the model assesses every new activity to define whether it is an outlier to the established patterns of behaviors. Anomalous events boost the total risk score of a user.

This will allow the system to identify finer insider threats that do not necessarily break the set rules, but reflect suspicious activity, which is consistent with the actual enterprise security practices.

III.4 File-Access Flow

Previous researchers have reported that privilege misuse and anomalous file-access patterns are the excellent predictors of insider threats [3], [4]. Conventional systems tend to examine such activities once they have taken place that may slow the response measures. To overcome this weakness, the suggested system analyses each file-access request on a real-time basis.

When a user is trying to access a restricted or sensitive file, the request is instantly evaluated on the basis of predefined rules and role-based permissions. Any illegal action is prevented and captured in the logs which can be used as evidence in proactive detection strategy as indicated in [4].

III.5 Risk Score Computation

Based on the conceptualization of adaptive scoring presented in the previous literature [2], the suggested system computes a risk score that is dynamically calculated by integrating a rule-based security breach with a machine learning-based anomaly score. Every action of the user is evaluated on the basis of the predefined security rules, whereas the Isolation Forest model offers the score on the anomaly confidence on the basis of deviation of behavior. This integrated approach has the capability of independent detection of known policy violations and also of new anomalous behavior. The score on risks is constantly updated so that real time monitoring is possible and to issue alerts

at the earliest time possible in case of any insider threats.

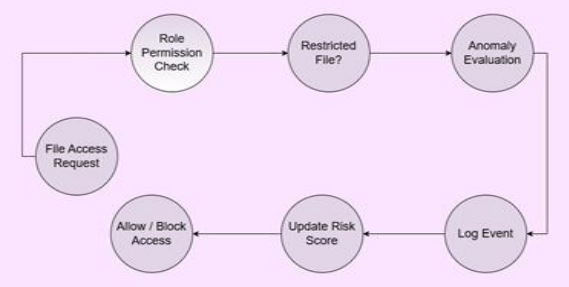

Figure 3. Real-time file access monitoring and anomaly evaluation process.

Figure 3 is a real time file access monitoring process which indicates how a user request is authenticated against role-based access policies and security rules. Attempts of accessing unauthorized or suspicious files are blocked instantly, logged, and sent to the anomaly detection module, where they are used in the computation of dynamic risk scores.

Weighted contributions of the two elements are used to compute the final risk score, whereby recurring suspicious behavior, however subtle, will result in gradual increase in risk visibility to the administrators. There are several behavioral indicators of the risk score, such as failed attempts to log in, unauthorized or blocked file-access requests, a login in a suspicious or unusual geographical location, and a login in a new or unrecognized device.

The score is higher when risky behaviours reoccur and this aids the administrators to identify what could be harmful behaviour in the early stages before it blows out of proportion into a serious event. Besides the access checks performed by the rule, the file-access events are also assessed by the anomaly detection model. Avoiding the normal working pattern of accessing sensitive files, attempting to access files that should not be accessed, and the display of unusual access frequency is considered a behavioral anomaly. These indicators are useful in detecting the misuse patterns, which are not easily identified via permission checks.

III.6 Real-Time Dashboard

Most previous solutions rely on offline logs or post-activity analysis as their major methods of operation [3], [4], which may postpone response actions. Conversely, the suggested system will have an administrative dashboard in real time, where important security data will be updated, including, but not limited to, active user sessions, suspicious log-in attempts, failed file-access attempts, dynamically-calculated user risk scores, and notices that need prompt administrative action. Confidence indicators that are generated by the machine-learning model are also displayed on the dashboard, enabling the administrators to differentiate between the rule violations and behavioral deviations.

The dashboard allows the administrators to directly suspend a user account in case any of the activities seems dangerous. This real time visibility allows making decisions faster, and shortening the time between action and detection.

IV. SYSTEM ARCHITECTURE

The recommended system will be modular-based, which is scalable and has real-time monitoring capabilities. On the application layer, user actions like login requests, file access requests are recorded and sent to the monitoring engine. Fingerprinting of the device and IP based location detection modules add contextual metadata to every event.A rule engine evaluates security policies, including authentication limits, access permissions, and rate-limiting thresholds. At the same time, selected features are forwarded to the Isolation Forest model for anomaly assessment. The outputs from both components are then combined by the risk scoring engine, which dynamically updates user risk levels.

All processed events and risk indicators are displayed on the administrative dashboard, allowing security teams to monitor system status and take immediate action. This layered design maintains a clear separation of concerns while supporting real time threat detection.

V. FEATURE EXTRACTION AND ENGINEERING

Anomaly detection depends on the choice of meaningful behavioral features to make. The offered system derives some properties based on the authentication and file-access logs such as the frequency of logins, the number of failed logins during a timeframe, the frequency of changing a device, geographical location deviation and attempts to access confidential files.

Those features are scaled and averaged during set

periods in order to filter the noise and enhance the model consistency. Sudden spikes in attempts to log in or access activity during unusual working hours are also fitted as temporal patterns. One advantage of feature engineering is that the Isolation Forest model can be trained to model both short term anomalies and long term behavioral deviations.

Forest Model Implementation Isolation Forest Model.

Isolation Forest It is an unsupervised anomaly detector algorithm, which isolates anomalies through random partitioning of data points. Contrary to distance-based techniques, it is an efficient algorithm on high-dimensional data, and requires no underlying data distribution. It builds a grouping of isolated trees, in which anomalies tend to get separated closer to the root of the tree and are more unique in nature. The algorithm scores each point with an anomaly score, which is easy to determine an outlier in the data set.

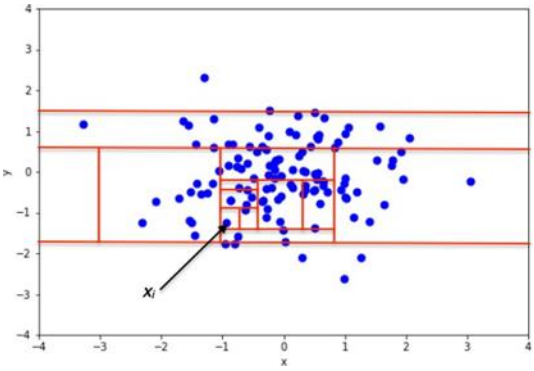

Figure 4. Isolation of a normal (non-anomalous) data point in a 2D Gaussian distribution using random partitioning.

Figure 4 illustrates the isolation process of a normal (non-anomalous) data point within a two-dimensional Gaussian distribution.

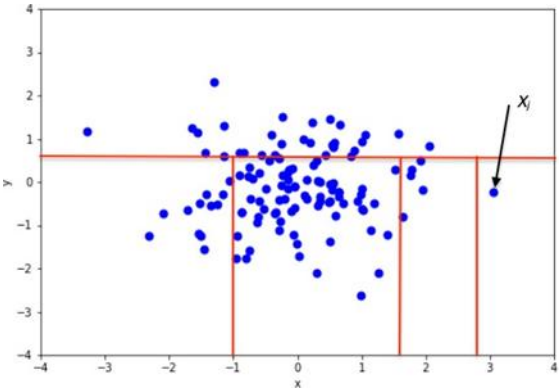

Figure 5. Isolation of an anomalous data point in a 2D Gaussian distribution, requiring fewer partitions compared to normal points.

Figure 5 shows the odd data point that is isolated in a two-dimensional Gaussian distribution. The anomalous observations are segregated by less random partitions as compared to normal data points, and it signifies that they are not associated with most samples. This idea is the main idea of the Isolation Forest algorithm in which anomalies are detected by their speed of isolation in random decision trees. This mechanism in the proposed system helps to identify abnormal logs in terms of login, device modifications, and abnormal file-access patterns. In the given system, historical records of activities used by the normal user are used to train the model. In the deployment stage, the incoming events are measured by their easy isolability in the decision trees. Higher scores of anomalies denote intensified deviations of the regular patterns of behavior.

It works together with checks based on rules, and any explicit violations as well as the subtle changes in behavior add to the detection of the threat. The design enhances a high level of resistance to insider threats which develop over time.

VI. IMPLEMENTATION DETAILS

The client-server architecture is used to implement the system. Authentication and logging of activities are handled at the server end, but the monitoring modules are asynchronous, therefore not affecting the performance. Fingerprinting of devices depends on attributes of browsers and systems and geolocation of IP addresses is acquired based on open IP mapping services. The anomaly detection model is a background service and it is periodically retrained to fit the changing user behavior. Audits, Logs and risk scores are safely stored.

The dashboard interface is a real-time display with role-based access to administrators.

VII. EXPERIMENTAL EVALUATION

To evaluate system effectiveness, simulated user behavior datasets were generated containing both normal and anomalous activities. Metrics such as anomaly detection accuracy, false-positive rate, and response latency were analyzed.

Method	Detectio n Accurac y	False Positive	Respo nse Time

		Rate	
Rule-Based Detection	Moderate	High	Fast
Isolation Forest Only	High	Moderate	Moder ate
Proposed Hybrid System	High	Low	Real-Time

Table I. Performance Comparison of Detection Approaches

Results indicatethat the hybrid approach significantly reduces false alerts compared to rule-only systems while detecting anomalies earlier than offline log analysis methods. The system maintained real-time responsiveness under increased login and file-access loads, demonstrating its suitability for enterprise environments.

7.1 Dataset Description

As the real organizational insider threat data are scarce, because of the privacy issue, a simulated dataset was developed to describe the normal and abnormal user behavior. The data contains authentications, devices, IP addresses, times and file access logs. Normal patterns of behavior were assumed on the basis of regular working hours and authorized access policies and some anomalous activities were used to simulate insider misuse cases.

7.2 Performance Evaluation Metrics

In order to measure the efficiency of the proposed insider threat detection system, several assessment metrics were taken into consideration. Detection accuracy was also the ability of the system to detect anomalous user behavior. The false-positive rate was examined to learn the rate of the user actions that are considered normal and were mistaken to be suspicious. Response latency was also considered to make sure that the system remains real-time when subjected to large volumes of the login and file-access events.

VIII.DISCUSSION AND LIMITATIONS

The suggested system passes the test of balancing both real-time monitoring and interpretability, which is critical to administrative trust. Anomaly detection quality is however determined by the quality of the historical behavior data. The model reliability may be compromised temporarily when new users come or change roles. Besides, though Isolation Forest is useful in addressing behavioral abnormalities, it fails to resolve sequential dependencies as much as recurrent neural networks. More efficient systems in the future can include timely models.

IX.CONCLUSION AND FUTURE WORK

The paper introduces a viable system of insider threat and anomaly detection that combines real-time monitoring, rulebased security enforcement, and unsupervised machine learning. The system is able to identify explicit security violations as well as subtle behavioral deviations by integrating Isolation Forest with adaptive risk scoring.

The direction of further work will include incorporating sequence-based models, scale-up behavioral attributes and assessing the performance on real organization data.

REFERENCE

[1] Anas Ali, Mubashar Husain & Peter Hans, “Real-Time Detection of Insider Threats Using Behavioral Analytics and Deep Evidential Clustering, 2025. Available: https://arxiv.org/abs/2505.15383 arXiv+1

[2] Lokesh Koli, Shubham Kalra, Rohan Thakur, Anas Saifi & Karanpreet Singh, AI-Driven IRM: Transforming Insider Risk Management with Adaptive Scoring and LLM-based Threat Detection, 2025. Available: https://arxiv. org/ abs /2505.03796 arXiv

[3] Rida Nasir, Mehreen Afzal, Rabia Latif & Waseem Iqbal, Behavioral Based Insider Threat Detection Using Deep Learning, IEEE Access, 2021.Available:10.1109/ACCESS.2021.3118297 ElsevierPure+1

[4] Bushra Bin Sarhan & Najwa Altwaijry, Insider Threat Detection Using Machine Learning Approach, Applied Sciences, 2023, 13(1), 259. Available: 10.3390/app13010259