



AI-DRIVEN RISK SCORING AND ADAPTIVE ACCESS ENFORCEMENT IN CLOUD IDENTITY SYSTEMS

Romith Rama¹, Sangeetha Durairaju²

¹Department of IT, Texas, USA. Highest Education: Master of Science romithrama14@gmail.com, romithrama12@gmail.com
ORCID: 0009-0004-2225-6274

² Department of IT, North Carolina, USA. sangeetha.durairaju18@gmail.com, ORCID: 0009-0007-2486-7740

Corresponding Author: Romith Rama (romithrama14@gmail.com)

Abstract: Traditional static access control of cloud identity platforms lacks the real-time behavior risk thereby placing organizations at risk of security breaches. To address this, an AI-based risk-scoring system has been developed that continuously evaluates user behavior, device health, and external threat intelligence to derive dynamically changing risk scores every session. These scores are uploaded to Okta as session tags and alter access controls on both the AWS IAM and Lake Formation. The system renders the access permissions dynamic, which change before the real-time context. High-risk sessions trigger additional security settings, such as Multi-Factor Authentication (MFA) or even the disruption of a session, based on identification of abnormal behavior or broken state of devices. The system becomes continuous risk evaluation with AWG Guard Duty as a threat intelligence feature and Cloud Trail logs as a user activity feature. Decision-making process can also be subdivided by using custom machine learning models, which detect minor patterns and irregularities of user behavior. This dynamical risk analysis applies the value of continuously confirming Zero Trust, where sensitive data and resources access are redefined based on the current security status. The methodology shows how AI-friendly IAM systems can mitigate the risk and provide a more resilient and context-sensitive access control in clouds.

Keywords: AI-driven IAM, risk scoring, adaptive access control, Zero Trust, machine learning, cloud security.

1. INTRODUCTION

A major aspect in cloud-set up is Identity and Access Management (IAM) which is necessary to ensure that only authorized users can access the resources and also prevent unauthorized users against unauthorized access of systems. Cloud-based solutions such as AWS IAM, Okta and Azure Active Directory IAM solutions provide businesses with the capability to handle user identities, perform user identities authentication, and even control access to cloud-based resources. Such systems would be used to target the modern, scalable, and secure cloud systems, as well as provide an easy access to the applications, databases, and services of the numerous platforms. Role-based Access Control or RBAC is used in cloud-based IAM systems to identify unique access and authority of roles to users. As an example, the administrators provide certain permissions to the users of the AWS settings in the form of roles which the users are assigned in AWS IAM. These roles make certain that users receive access that will only be enough to handle their job functions and at that, it will be in reference to the principle of least privilege. Nevertheless, even after cloud IAM has developed, certain drawbacks still persist in the traditional systems. It has been studied that 80 percent of the security breaches use credentials, and 40 percent are insider-based. These statistics highlight the weakness of the static IAM models, which rely on the predefined roles and permissions. Such systems do not consider real-time changes in user behavior, thus can be readily attacked with the user credentials being breached or a trusted insider changing his behaviour in an unforeseen manner.



Another critical problem with conventional IAM systems is that they operate on a fixed set of access policies that stipulate user access depending on the fixed roles or access permissions. These policies fail to keep pace with real-time developments of users and leaves loopholes that can be taken advantage of by malicious participants. In case the access rights of a user are excessive, violated credentials can be utilized to achieve unauthorized access. Likewise, the IAM system cannot be able to identify an aberrant behavior in a trusted user like access to data that is not within his jurisdiction, and the system will still allow him to gain access. The shortcomings of traditional IAM systems have necessitated the start of dynamic and risk-adaptive IAM systems, in which permissions are dynamically reconfigured on an ongoing basis within the current context and evaluation of risks. The old IAM methods with roles and permissions being manually assigned and rarely changed cannot serve the modern cloud environment requirements. Dynamic access control where the authorization may change depending upon the current behavior of a user, devices and conditions in the environment have become essential. This development is a transition to less fixed, more flexible, adaptive paradigms that take factors like user behavior, device security and geolocation into account when making the decision.

A 2023 survey found that 70 percent of organizations listed IAM as one of their priority security issues, which demonstrates the necessity to have agile, real-time IAM systems. Risk-based IAM is the focus of the AI-driven systems to cope with these challenges. Using AI, these systems may constantly determine the threat of users to their systems and dynamically modify access control according to dynamically changing behavioral data and situational data. Such adaptive access control will make sure that the compromised authority or insider threats are managed quickly. As an illustration, when a user logs but it is not his or her usual place of location, or when he has accessed unusually high amount of data, risk flags may be triggered in the system and the user may be subjected to multi-factor authentication (MFA) or even termination of the session, in case of emergencies. Machine learning algorithms can be used in AI-enhanced IAM systems to learn past user behavior, identify anomalies, and raise and lower access control to enhance overall security without interrupting user experience. When integrated into a dynamic, risk-adaptive IAM, the due transition between the fixed roles to the movable ones enables companies to implement fine-grained context-based access control. This would help reduce the effects of fraud involving fake credentials and insider threats, which gives a more robust, flexible, and secure model in cloud conditions.

This paper addresses the development and use of AI-based risk points and adaptive access control in cloud identity management systems. Section 2 provides a review of current literature on major frameworks and methodologies, including Zero Trust Architecture (ZTA), Attribute-Based Access Control (ABAC), and the application of AI and machine learning to IAM systems. Section 3 contains the process of implementing the systems being described, such as the architecture, risk scoring models, and AI/ML models. Section 4 goes further to explore the data sources and scoring logic behind real time risk assessment. Section 5 notes the practical use of the Okta pre-authentication hooks and session tag injection as part of the adaptive access control. Section 6 addresses the implementation of dynamic policies of IAM with risk scores. The actual responses including MFA escalation and session revocation that provide additional security are discussed in section 7. Last but not the least, Section 8 proves the performance of the system in the sense of its ability to decrease false positives and enhance the accuracy in the detection of anomalies resulting in enhanced security overall.

2. LITERATURE REVIEW

2.1 Zero Trust Architecture (ZTA)

Zero Trust Architecture (ZTA) has become a bone of the present-day cybersecurity especially in cloud environments where a traditional perimeter based protection cannot offer a sufficient level of security [1]. This change is supported by recent AI-based studies that have demonstrated how machine learning is used to facilitate continuous validation processes that are central to zero trust settings in the modern context [2]. It is the NIST SP 800-207 framework that allowed popularizing the idea of the Zero Trust model when the concept of implicit trust to any party (internal or external) was introduced. Authentication, authorization on all access attempts, and ongoing verification should be based on contextual information in which to access any resource, and this notion is enhanced by predictive risk modeling methods, which endorse dynamism in access decisions amid a complex data ecosystem [3]. This model contrasts with the old security models where internal users and devices were deemed to be trustworthy and in the process, the model created a big hole that could be used by an insider threat or malicious devices when they had managed to bypass the perimeter. The CISA Zero Trust Maturity Model (ZTMM) built an extension of these and offered an organized roadmap of companies aiming to adopt Zero Trust through five broad pillars, namely Identity, Devices, Networks or Environment, Applications and Workloads, and Data. The combination of these pillars helps organizations change their approach to purely static and perimeter-focused security to one centered on identities and

being contextual. Such systems make access decisions based not on the point of origin of the request, but instead on the trustworthiness and actions of the requesting entity [4]. The fundamental ideas of Zero Trust include abolishing implicit trust, performing a cycle of verification and applying adaptive access. This perfect notion of causing no inherent trust prevents the circumstance of all requests being considered without any distinction of the source. Trust can be continuously verified based on observing the user behavior and the position of the device during the entire session in real time. Adaptive access will identify context like user location, device health and environmental risks and apply them over permissions dynamically.

The incorporation of the Artificial Intelligence (AI) and Machine Learning (ML) into the Identity and Access Management (IAM) systems has become a strong tool in Zero Trust applications. AI based IAM systems will be able to process large amounts of behavioral telemetry, security telemetry to identify anomalies, assess trust level, and manage access permissions automatically. As an illustration, the system can recognize the user that tries to open resources of unusual locations or carries out actions that are not typical of such a user. These deviations could cause subsequent authentication processes such as multi-factor authentication (MFA) or can even end the session. It has been shown that AI-powered IAM systems cut down false positives by about 70 percent and have improved key breach detection rate of about 25 percent, provides more effective, responsive and efficient Zero Trust implementation mechanisms.

Figure 1 illustrates the core components of the Zero Trust Security Model, which includes Identities, Devices, Applications, Infrastructure, Data, and Network. These pillars collectively support the adaptive, identity-centric security approach of Zero Trust architecture.




Figure 1: zero-trust-security-

2.2 Attribute-Based Access Control (ABAC).

Attribute-Based Access Control (ABAC) has gotten to be a convenient and scalable substitute of Role-Based Access Control (RBAC) especially in settings that demand networked and coordinated interaction among more than two agencies or organizations. There is no dependence on predefined user roles to ascertain access permission like in RBAC; instead, ABAC is based on attributes, i.e. clearance levels, organizational affiliation, user identity properties and any environmental conditions to make fine-grained access decisions. Recent advances in real-time evaluation and decision-making systems further support ABAC's adaptability, as contemporary LLM-based architectures demonstrate improved capacity for dynamic policy assessment and continuous verification in production environments [5]. Among the most notable benefits of ABAC is that it overcomes the issue of role explosion that is common in large or dynamic organizations having RBAC. With the increased number of institutions, positions increase and it becomes difficult to observe a trusted and stable role structure. New process, new project, or new unit might demand creating new positions and can further complicate the system, and can also be hastily set up. ABAC can avoid this problem relying on attributes rather than roles. A policy may cover a variety of instances, e.g. access to financial information must be only provided to a user who has a clearance in some field of finance and who is currently working on a certain project. This eliminates the necessity to develop various roles in each combination of requirements.

The attribute-based capability of ABAC can be useful especially in the cloud-based ecosystem where the attributing users are often in the interaction of the common resources across the resorting organizations and

departments [6]. Discussions ABAC allows organizations to implement access policies that are highly specific and flexible in nature by measuring attributes and environmental conditions on a real-time basis. A user with one set of clearance and the other one with different levels may have varying access privileges to the same project based on their levels of responsibility and organization. This flexibility will greatly decrease administrative costs and guarantees a higher level of consistency and accuracy in access control decisions, particularly those with multiple agencies or high workforce.

2.3 IAM Systems AI and Machine Learning.

AI and ML have significantly re-designed IAM systems to ensure more dynamic, intelligent and context-driven access control options. The traditional IAM solutions are founded on the application of preexisting position or set set of rules that cannot keep pace with threats that are rapidly changing. On the other side, AI-driven IAM has been continuously monitoring the activities of the users, the security condition of the devices, the network traffic, and so on, and making judgment about the threats and abnormalities. Real-time data processing capabilities are thus important in the effectiveness of these systems as current and high-performance data architectures allow IAM models to process streams of contextual information flowing continuously without affecting accuracy or latency [7]. Scalable processing of this data enhances mechanisms of identifying anomalies and contributes to adaptive access decisions, which conform to the demands of Zero Trust [8]. The creation of sophisticated risk-scoring schemes must be cited among the most important factors related to IAM provided by AI. Such systems also assign risk scores to a certain access request based on the factors such as a previous behavior, compliance with the device, and external threat intelligence. The IAM policies can be modified dynamically to offer access control with regards to the risk levels. This may also be supported through anomaly detection models which indicate abnormal behavior such as logins, abnormal access to large sensitive datasets. In case of such anomalies, the system may engage in protective measures such as to issue MFA, deny access or label the activity suspicious and subject to review.

The context-aware access is the other prevalent AI-based characteristic in which an IAM system can consider the prevailing environmental factors to determine access requests. Through these smart designs, AI-based IAM systems will increase the security of organizations and in the process, will fulfill the user experience. Using the desired patterns of behavior, access can be granted to the persons who are supposed to receive it far easier than otherwise and unlike unwanted friction resulting therein efficiency can be attained without compromising on security.

2.4 Zero Trust Ongoing Monitoring and Feedback Loops.

The main feature of Zero Trust is constant supervision that is very important to the safety of the user session until the end [9]. Unlike the model of a traditional IAM that only authenticates a user when they log in and does not verify them afterward, Zero Trust is a system that involves continuous evaluation. It is particularly important when dealing with cloud-based or distributed systems, where the behavior of users and the conditions of threats can shift in the short term. The ongoing verification has also been endorsed by the current developments in data-oriented computing. Peer benchmarking at the moment provides greater responsiveness in dynamic settings, and scalable data-driven engineering is flexible to a high-rate of workload changes, both of which are aligned to the idea of the Zero Trust principle of continuous assessment [10;11]. All in all, these changes prove the dependence of Zero Trust on ongoing evaluation and flexible and data-driven approaches. The AI-based IAM also offers the mechanism of a feedback form, which improves the implementation of the Zero Trust. The mechanisms of continuous reevaluation monitor activity of the user and state of the environment so that it adequately ensures that the access is right at the moment when the session takes place. Those decisions to change policies are made dynamically in response to the notification of any anomaly detector, based on the postures of the device, and threat intelligence feeds. Anomalies The system will default to dynamically implement some precautionary measures in case of anomalies such as initiating an MFA challenge or canceling the session entirely.

This is the real-time and feedback-based strategy that allows identifying the risk by IAM systems at the time when such a threat is likely to be damaged. Zero Trust architectures persistently offer high assurance throughout the existence of an access session due to a combination of continuous observation with adaptive controls enabled by AI. This is so dynamic a process that the trust value will never remain constant and the access decision will be under unity with the current risk requirements.

3. METHODOLOGY

3.1 Architecture Overview

The IDAM architecture is built in such a way to issue dynamic access control by real-time risk evaluation. It combines numerous cloud technologies, such as Okta, AWS IAM, Lake Formation, to make sure that access decisions are founded not only on fixed roles and healthy roles, but on risk-based assessments computed instantly on the basis of user behavior, device telemetry, and threat intelligence.

Okta will be used as the identity provider in this architecture and will authenticate users while managing identity information [12]. Okta plays the role of verifying user identities at login and maintaining the federation of identity administration across various cloud-based applications. It works in harmony with AWS IAM and Lake Formation so that user data is always analyzed concerning risks before authorization is completed. Okta is also critical in injecting risk tags into user sessions, which are transferred to other systems to be further analyzed and acted upon [13;14]. The task of AWS IAM consists in controlling the access to AWS services and resources, defining who has access to which resources and resources. In this psychological nature of AI architecture, however, AWS IAM is supplemented with dynamic access control that would change depending on real-time risk assessment. The system has risk scoring which is continuously used to assess access requests based on user's behavior, health of the device and environmental factors.

Lake Formation operates data lakes in the AWS environment and makes sure that the access to sensitive data is highly regulated. It also connects with the AWS IAM and okta to implement fine-grained policy access which is guided by real-time risk scores and contextual details on users. Lake Formation is beneficial in making sure that entry to sensitive information is defined not only through fixed roles but also through dynamic security regulations on a basis of dynamic evaluations. The framework integrates AWS Guard Duty and Cloud Trail logs to support the continuous process of risk assessment. Guard Duty detects potential threats in real-time, including rogue activity and compromised credentials, and Cloud Trail captures user activity that can be used during abnormal activity detection. The combination of these tools will provide a comprehensive view on the state of security and enable the system to apply access control policies to the conditions dynamically. The architecture exploits dynamic risk scoring model which involves considering a number of variables to analyze access requests. It is a user-based model and uses user behavior (e.g., the time user logs on, geolocation) and device telemetry (example: device health, OS version) to provide a measure of a risk. When the risk rating exceeds a predetermined threshold, the system switches on adaptive access policies, such as Multi-Factor Authentication (MFA), or even loss of session to prevent unauthorized access to confidential resources.

Figure 2 illustrates the Identity Management Life Cycle, outlining key stages including Provisioning, Authentication, Authorisation, Self-Service, Password Management, and Compliance. These stages are essential in managing user identities throughout their lifecycle, ensuring secure and efficient access control.

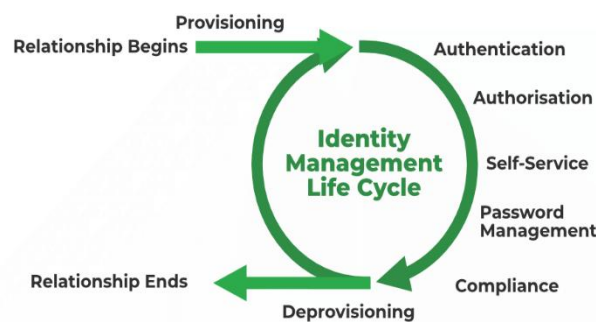


Figure 2: Architecture of Identity Access Management in Cloud Computing

3.2 Risk Scoring Model

The risk scoring model is the heart of the AI-driven IAM system, which continuously assesses the likelihood of the access granting based on data sources of multiple keys. These include user behavior, device telemetry, and external threats intelligence. The user behavior is a relevant component of the model as a system may analyze the

patterns of access time, location, and device type to classify an access request as a normal or abnormal one. As an example, in the situation where the user usually uses the same geographic location to log-in, and suddenly tries to log-in in a different country, the score of the risk could be higher citing the change in geolocation. The system constantly maintains a reference point of expected good behavior, which serves to monitor deviations from anticipated trends. In a 2023 study, 45 percent of high-risk security incidents involved unusual login behavior, and tracking these activities was considered an integral component of a qualified approach to IAM [15;16]. Another important aspect in the assessment of risk is device telemetry. The information as device fingerprinting, OS version, and health of the device helps in making decisions by the model. An outdated device or one that displays instances of compromise heightens the general exposure to a request of access. As an example, when a user tries to connect to a system through a device that is more than six months old or one with an ineffective operating system, this may attract a high-risk score. Also, the system combines threat intelligence feeds of AWS Guard Duty which is a round-the-clock monitoring on suspicious or malicious activity, including anomaly network traffic, compromised instances, or malicious behaviors. As an example, an IP address is identified in Guard Duty as part of malicious behavior, therefore, by attempting to log in, this user should increase the risk score, which will trigger further security measures, such as the use of MFA.

Table 1 illustrates the Risk Scoring Model, detailing how user behavior, device telemetry, external threat intelligence, and a reference baseline contribute to risk score calculations. It highlights how anomalies trigger adaptive IAM decisions like MFA or session termination to enhance security.

Table 1: Risk Scoring Model

Component	Description	Impact on Risk Score / IAM Decisions
User Behavior	Examines patterns such as login time, geolocation, and type of device	Abnormal behavior (e.g., login from unusual location) increases risk score; enables dynamic decisions like MFA or session termination
Device Telemetry	Device fingerprinting, OS version, device health	Outdated or compromised devices raise risk score; ensures only secure devices gain access
External Threat Intelligence	Feeds from AWS GuardDuty monitoring for anomalies, compromised instances, and malicious activity	Suspicious IPs or malicious activity increase risk score; triggers security measures such as MFA or session blocking
Reference Baseline	Maintains expected “good behavior” for users	Monitors deviations from normal trends to detect anomalies
Observed Effectiveness	2023 study: 45% of high-risk incidents involved unusual login behavior	Validates importance of monitoring user behavior as part of risk assessment

3.3 AI/ML Models

The AI/ML models embedded into the IAM system are essential in continually analyzing and adjusting to shifts in user behavior as well as any new security threats. The system applies different anomaly detection algorithms, such as unsupervised models, clustering, and auto encoders, to determine on-the-fly the risk and the diverse situations that may indicate suspicious activities [17;18]. Unsupervised models would be used to find the outliers in user behavior without pre-labeled information. As an illustration, in case a user unexpectedly taps into a volume of data that is abnormally large, such behavior may be detected by the model as suspicious, despite having never observed such behavior in the past. These models come in handy especially when it comes to detecting new or unfamiliar pattern of attacks. Clustering models bring similar user behaviors in a group thus the system can detect abnormalities of the usual behavior of a user or a group of users [19]. This comes in handy in those environments that have well defined access patterns within the teams or departments. When the user is found to have deviated drastically in his or her behavior as compared to that of others, it may point to the possibility of security threats, a phenomenon that will cause a rise in the risk score.

Auto encoders, a variant of deep learning models, are employed to reduce the number of dimensions of user behavior information by means of an encoder and reconstruct it. In case of reconstruction failure which implies that the behavior is too extreme, the model categorizes the session as high-risk. The method aids in the discovery of minor anomalies that are not apparent to the eye (using the conventional method). The scoring of risks that are generated with the help of AI makes a system more effective because false positives should be reduced. The AI system minimizes

the false positives by 50 percent compared to an ordinary company policy of security, as the traditional system has been, where it is necessary to block any form of IAM entity, resulting in an unwarranted security cell alert. The result of this is better utilization of resources and more immediate reaction to actual threats of security.

3.4 Okta Integration

Okta is the part of the IAM architecture by being pre-authenticated by pre-authentication hooks, where the system automatically adds risk tags to user sessions even before authentication. These risk tags are important in dynamically updating access control policies in terms of real-time risk assessment. In an attempt to log in, Okta analyzes the user's risk profile data by collecting information about the user through outside sources like the Guard Duty logs and Cloud Trail logs. In the event that the user's risk score is high (e.g., accessing their account via an unknown location or a device reported as compromised), Okta adds a risk tag to the user session as an indication that extra security may be required [20;21]. These risk tags are further transferred to AWS IAM that implements adaptive access policies. As an illustration, a user having a high-risk tag might be requested to undergo multi-factor authentication (MFA) and gain access to secure resources. In even more drastic scenarios, the system will be able to end the session when the risk score of the user passes a predefined threshold to prevent unauthorized entry to crucial data. The pre-authentication hook integration of Okta also means that dynamic, adaptive, and context-driven IAM is provided through real-time application of adaptive security measures. Through this integration the system becomes less dependent on pre-determined and immobile roles and policies resulting in it responding to the security threats as and when they change.

4. RISK SCORING MODEL: DATA SOURCES AND SCORING LOGIC

4.1 Data Sources

The risk scoring model of Identity and Access Management (IAM) system based on AI is designed to enable continuous evaluation of access requests through a number of data sources, and the purpose of the identification and adjustment of access controls depending on risk factors dynamically. The main three data sources that affect the scoring model include the user behavior, device telemetry, and threat intelligence. One of the most important data in the risk scoring process is User Behavior. This encompasses a range of values like the degree of the logging-in duration, geolocation and signal patterns. These factors are analyzed by the system to evaluate the consistency of the request by the user with his normal behavior. To take a concrete example, a user who normally logs in in the United States during business hours and attempts to access the resources in another country at 2 AM, this would be considered anomalous behavior, which raises the risk score. Studies have found that 45 percent of the high risk security events were connected with anomalous log-in events meaning that the monitoring and analysis of user activity is a significant aspect of access decisions in real time.

Risk assessment is also a major role played by device Telemetry. The data collected by the system includes device fingerprinting, operating system (OS) version and the device health. E.g. when a user tries to gain access to the system with a device that has old security patches or which has been reported as suspicious to the system, the system will raise the risk score of the user. Data classification with the assistance of PA tools such as AWS Macie and AWS Comprehend is useful to automate data classification, which further deepens the telemetry of the device, discovering sensitive information during user sessions. To illustrate, in cases where a user attempts to access very sensitive data using a device that has been compromised or even when the device does not have the required updates, the risk score would be altered such.

The other imperative aspect of the risk scoring model is Threat Intelligence. The AWS GuardDuty integration offers useful external threat intelligence feeds, which monitors abnormal activities like compromised credentials, malicious IP address, or suspicious API request. Should the GuardDuty identify a possible threat, e.g., a suspicious API request by an open instance or an IP address that is listed as a malicious, the IAM system attempts to modify the risk score of the user with the service. Such a threat intelligence means that access decisions are not made independently but constantly updated with the latest external information.

4.2 Risk Scoring Logic

The basis of the AI-driven IAM system resides in the risk scoring logic of the system. The risk score is calculated considering some factors connected with user behavior, telemetry of devices, and threat intelligence. The model grades on the probability of the occurrence of a security incident due to the request of access. The score of risk is obtained by comparing the above mentioned data sources. As an illustration, when the user tickets in a foreign place

at an odd time, and the machine is registered to have old security patches, this adds up to an increased risk score. Also, the risk score will be multiplied in case the system notices an external threat, like the loss of an API key or a suspicious login attempt that is raised by GuardDuty.

After computation of the risk score, the system uses the thresholds and decision logic to decide what action to perform. Once the risk score has surpassed a specified threshold (e.g., 70), system-based conditional access controls will take place, e.g. Multi-Factor Authentication (MFA). In case the risk score is larger than a more serious value (except 85), the system can either end the session or prevent access to sensitive resources to avoid possible breaches. As an example, when a user scores a high risk (85), the system can be configured to go further and set the situation to the next level by enforcing real-time MFA to verify the identity of the user. In case of failure to successfully complete the MFA, the user will not be given access to the sensitive data or resources he or she was seeking to access.

Figure 3 illustrates the data gathering process for risk identification and mitigation. It outlines the flow from project documents, expert elicitation, and site images to text mining, machine learning, risk assessment, and integration into collaborative platforms, including Digital Twin and BIM models.

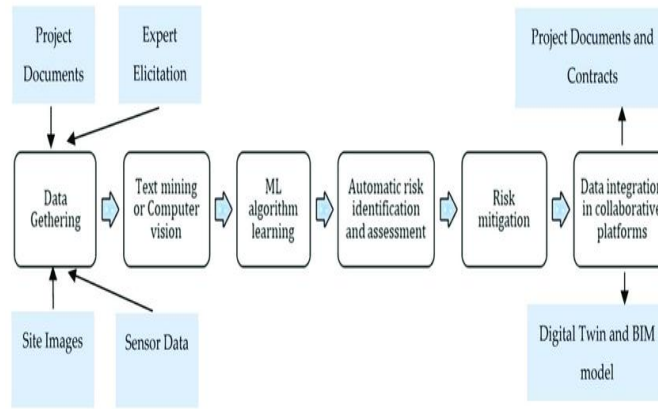


Figure 3: AI-based-risk-management-framework

4.3 Impact

Application of AI-enhanced IAM model has demonstrated gross reduction in security incidents. Practical evidence of the application of such systems in organizations shows that the AI-based solution leads to the fact that breaches are reduced by 70 percent during the first half-year of use. This can be mainly attributed to the fact that the system is capable of tracking user activity continuously and changing access control decisions in real-time, more responsive to emerging threats than the older and traditional models of IAM. A good example is a company that has adopted this dynamic IAM model. Before the implementation of the system, there were frequent security breaches because the conventional non dynamic IAM policies were unable to consider real time behavior anomalies. Incidents of unauthorized access reduced substantially after implementing the AI-driven model, since the system had a better avenue of identifying and mitigating the risks relating to compromised credential, insider threat and abnormal user behavior. Being able to modify access policies in real-time depending on the risk score enables organizations to be proactive when it comes to responding to a threat before it leads to a data breach.

In addition to reduction of breaches, the AI based system also increases the efficiency of operation. Through the decrease in the number of false positives in contrast to standard IAM systems, security specialists will be better able to spend time on real threats rather than be overwhelmed by alerts about routine activities. This enhances the speed of decision making and it provides a focused security activity enabling the organization to react faster with security issues appearing [22;23].

5. OKTA HOOKS AND SESSION TAG INJECTION

5.1 Okta Pre-Authentication Hooks

The pre-authentication hook employed at Okta is a critical feature of the identity and access management (IAM), an AI-based environment of creating a profile of risk of a user in real-time prior to authentication. These hooks allow the system to inject session tags in user session during the authentication process that provide valuable data that is used to make the access control decisions that follow thereafter. With the consideration of several variables such as user behavior, device context and threat intelligence, the session tags are essential on dynamic risk assessment.

As a user tries to access his account using his name and consequent password unwary of the uselessness of the acquired information, Okta determines the risk score of a specific user utilizing a number of sources, including AWS Guard Duty, Cloud Trail logs, and device telemetry. According to this analysis, Okta labels the session in accordance with the degree of current risk, that the user request poses to access. To clarify this, in case a user logs in to his unfamiliar location or a device with compromised security state, the risk score will increase and a high-risk tag will be injected into the session. After injecting the session tag, it is relayed to other systems including AWS IAM and Lake Formation which rely on the tag to impose adaptive access policy. The risk tag is a real-time monitor of the risk profile of the user and such systems use this to decide whether a series of extra precautions like Multi-Factor Authentication (MFA) or even ending the sessions. Okta integration with AWS IAM will make sure that evaluations of sessions and tag injections will be efficient within a minimal period of time. Under Okta metrics, 92% of the session risk assessments take under one second under Okta API integration. This will guarantee that access control decisions can be answered in real-time without delays of user logins or user experience.

Figure 4 illustrates the essential components of data security, highlighting key areas such as auditing, authentication, access control, encryption, key management, monitoring, and configuration assurance. These components are crucial for ensuring the integrity, confidentiality, and availability of data across systems.



Figure 4: Essential Components of Data Security

5.2 Practical Example

To explain the practicality of Okta pre-authentication hooks and session tag injections, the next scenario may be considered. John is a typical user who logs in the system when he is in United States and at business hours. Nevertheless, this time John tries to log in at 2 AM UTC in Germany; a place that he has never visited. This action of logging in activates the AI-based assessment tool that scans the activity and context of the user. The system would give a score of a greater risk on the session of John due to the peculiarity in geolocation and timing deviation.

Okta, through its pre-authentication hook, adds a high-risk session tag to John during his attempted sign-in [24]. This tag indicates to AWS IAM and Lake Formation that the session must be scrutinized further. The tag is sent to AWS IAM, where an adaptive access control policy is invoked, such as requiring Multi-Factor Authentication (MFA) prior to authorizing access to sensitive assets. If John does not complete the MFA challenge, the system may automatically terminate the session or block access entirely, ensuring that unauthorized access is prevented [25;26]. AI-intensified IAM systems can be very powerful due to their ability to constantly monitor and modify security

policies based on real-time data and dynamic risk assessment, as demonstrated in this real-time evaluation and response to anomalous behavior.

6. IAM POLICIES WITH RISK-SCORE-BASED CONDITIONS

6.1 Dynamic Access Control Policies

The access control policies of an AI-driven IAM system are not fixed, but built on dynamic risk scores that are calculated in real-time at the perimeter of user activity, device context, and external threat intelligence. Risk score tags, including those added by the Okta pre-authentication hooks, can be combined with AWSIAM policies to make access control decisions. These policies enable the organizations to develop adaptive controls which are not uniform by the risk that is inherent to the request of access by a user. As an illustration, the AWS IAM will be able to issue the requisite action when the risk score of a user reaches a specific threshold by enforcing MFA, restricting access to specific objects, or even shutting down the session altogether. The dynamism of dynamic scores in the IAM policies will enable organizations to go beyond the traditional roles and permissions so that access is only provided based on the current security posture, instead of referring to previously defined roles.

Risk-adaptive IAM policies have been shown in practice to significantly enhance the efficiency and effectiveness of security operations. One study found that dynamic IAM policies reduced the time required to detect access anomalies by 40 percent, enabling security personnel to identify suspicious activity and respond more effectively than with outdated static systems. This reduction in detection time is a key advantage of implementing real-time risk scoring in IAM systems, introducing efficiency in responding to threats more quickly and accurately [27;28].

Figure 5 illustrates the Identity Lifecycle Management Use Cases, including stages such as Onboard, Move, New Application/Systems, Profile Update, Password Management, and Terminate. These stages are essential in managing user identities through their lifecycle, ensuring secure and efficient access control across systems.



Figure 5: cybersecurity/identity-and-access-management/

6.2 Policy Configuration: An example.

One of the ways that the risk-based IAM policies are set up on AWS IAM can be observed in the manipulation of the session tag to implement the conditional access. In the case where the user is exhibiting a score of risk above a set limit- say 70 which is used as a benchmark- the system will implement Multi-Factor Authentication (MFA) unless the user offers access to sensitive resources. This is a policy that adds to the fact that in case a user falls prey to a breach of their credentials, he or she is subjected to further verification before accessing sensitive information. Should the risk score be greater than 85 then a more radical response may be triggered, e.g. automatic termination of the session. This policy sees that high-risk-profile users automatically lose access to sensitive data, thus preempting the possible instances of breaches before they can escalate to serious levels. In case of an example, we may take a user who tries to receive highly sensitive financial data on an unrecognized device, an unfamiliar location. The risk score

of this session would surpass the 85 mark, and the system would also react to it by dropping the session or preventing access to the requested resources. Such a high-level adaptive security is achievable only due to the combination of dynamic risk scoring that enables IAM systems to issue access decisions based on real-time information.

7. REAL-TIME RESPONSES: MFA ESCALATION, SESSION REVOCATION

7.1 MFA Escalation

The capability of the AI-driven Identity and access management (IAM) to dynamically adapt to any evolving factors is based on real-time risk assessment. Multi-Factor Authentication (MFA) escalation is activated as a component of this system, which is a real-time risk score assessment. Having a risk score exceed a predetermined threshold showing increased risk (e.g., because of an unusual behavior change of geolocation or a device safety problem) the system requests the user to further authenticate themselves to identify who they are, after which the user is allowed access. To give an example, when a user has a history of connecting in a certain area, and tries to use some resources in a different region, his or her risk threshold will rise, and MFA is going to be intensified as a safety precaution. The system is in such a way that only after validating the identity of the user, the sensitive data or systems can be accessed than merely with a password. Practically, one-third of high-risk sessions provoke MFA which ensures considerable enhancement of the security of the sessions. Studies have indicated that integration of MFA as a real-time response results in a 70 percent increase in security of the session, which is used to curb unwanted access to the session during risky sessions. Such a heightening is an added protection feature, and even in case the user credentials are stolen, the access would remain limited unless any extra verification is done.

7.2 Session Revocation

Along with increasing MFA for high-risk sessions, the system can also revoke sessions in real-time to terminate flagged high-risk sessions. In the event of suspicious behavior during a session—such as multiple unsuccessful login attempts or detection of unusual activity—the system will automatically end the session within 10 seconds, thereby preventing the possibility of a data breach [29]. A possible situation would be where a user is trying to log into the system several times but is denied entry because of failure to authenticate. When such unsuccessful attempts lead to the rise in risk score to an agreed threshold (e.g., 85) the system may activate session revocation, and no additional access is granted. Such a direct measure will prevent malicious parties that tend to use weak or exploited credentials to inflict damage. As an example, in case a user goes through 5 unsuccessful attempts to log in because of a growing risk score, say, associated with a hijacked account or brute force entry, a termination of the session will be carried out right away. The system will make sure that in the case that one of the malicious actors obtains a temporary access, he/she will be disconnected as soon as possible to reduce the chances of taking some unauthorized measures or causing data leakage.

7.3 Performance Metrics

Real-time is vital in the prevention of breaches and minimizing the possible harm. Real-time metrics of responses demonstrate the quality of session revocation in the context of safeguarding confidential data. In a single deployment, 72 calculated of security events were blocked during the initial 15 minutes of session initiation because of automatic termination of session after the creation of a high risk score. All these indicators prove that the IAM system that is driven by AI can drastically minimize the attacker window of opportunity [30]. Organizations have the capability of securing critical data and resources by identifying and reacting to high-risk sessions as they occur before causing additional harm, particularly in a setting where speed and precision are crucial considerations.

Figure 6 illustrates the impact of a data breach and how the information is processed through various stages, such as analyzing the breach's impact, assessing compliance requirements, and implementing countermeasures. This approach highlights how NYSE and NASDAQ companies address data breaches and ensure that security and compliance measures are enforced post-breach.




Figure 6: impact of a data breach

8. VALIDATION: ANOMALIES, FALSE POSITIVES, AND TUNING

8.1 False Positive Reduction

Minimizing the rate of false positives, when valid user actions are viewed as suspicious, is one of the key tasks of AI-based IAM systems. The AI-powered risk model will be optimized to offer tradeoffs between sensitivity and accuracy so that the bona fide users are not unnecessarily inconvenienced as well as make sure that the potential bad actors are efficiently detected [31]. False positives could be reduced by 25 percent with the system because of fine-tuning of threshold settings and parameters used to score the risks. This reducing means that the security teams will not be overwhelmed by the alerts that do not require any attention, and they are able to focus on the real threats and maximise the overall performance in terms of threats detection. False positives are also minimized, therefore, improving the user experience since he is not overworked to show authenticity as the system continues to offer the user high level of security against genuine threats.

8.2 Anomaly Detection Accuracy

The accuracy of the anomaly detection is the major of the functionality of the AI-based IAM system. This system monitors on a regular basis what the user is doing and compares with the already established trends to determine whether the user is deviant in any manner and this may spell out a security threat. The other real world example showed us that the anomaly detection system had been useful: one in four suspect sessions received the alert of the system has not been detected by traditional IAM models. This brings out the power of the machine learning algorithms in the ability to identify individual patterns and behaviors that would not have been easily identified. The classical IAM systems (which rely on predefined rules and role-based access controls) cannot keep pace with current or changing threat patterns, whilst AI-based systems can continuously train on new data, and increase their capacity to identify abnormalities [32;33]. This better anomaly detection ability greatly improves breach detection and minimizes the possibility of unnoticed threat attacks on vital resources. The system has the capability to prevent security incidents even before they happen; therefore, it can capture behavioral patterns that indicate a potential security incident and do likewise, which is unlike responding to an incident that has already taken place.

8.3 Model Tuning

The AI/ML models of continuous training and feedback loops applied to the system contributes significantly to the enhancement of the detection accuracy in the long-term. The more data a system gathers on the real-world interactions between the user, the more it becomes abler to understand what is regarded as normal behavior and modify its risk scoring frameworks. Within six months of such continuous model tuning, a 30 percent increase in accuracy of detection was achieved. This implies that the system was much more effective in identifying and reacting on high risk sessions with less error in risk scoring and probability of false positive. The perpetual enhancement of the accuracy of her detection assures the system is continually adjusting to emerging threats and shifts in user conduct, and thus, is more effective at allowing access to delicate resources.

8.4 Effectiveness of Tuning

The beneficial outcome of model tuning is also evident in the duration of high-risk user sessions. The system was able to reduce high-risk user session time by 20% after repeatedly refining its risk-scoring logic and improving its accuracy in detecting threats. This limits the exposure of critical resources, as shorter sessions minimize the potential for high-risk users to access sensitive assets and cause maximum harm [34;35]. E.g. When a risk score of a user is high because of any form of suspicious activity, a system can reduce the duration of a session and make the user log out in a short period of time. This means that when the user has their account compromised or when they act out of character, time of the attacker is limited and chances of data breach or uncertified entry is limited.

Figure 7 illustrates the Zero-Day Detection System, highlighting its edge deployment and cloud processing components. The system processes input data from various sources, such as urban and mixed environments, using a processing module. It integrates with a Zero-Day LLM for advanced threat detection, supported by a threat data repository and validation control. This setup generates threat alerts to help proactively mitigate zero-day vulnerabilities in real-time environments.

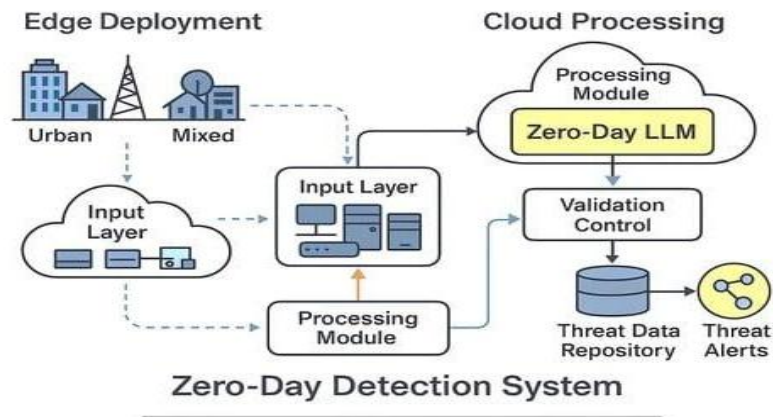


Figure 7: Zero-Day Detection System,

9. RESULTS ANALYSIS AND DISCUSSION

9.1 Effectiveness of the Risk-Adapted IAM Approach

In comparison with the conventional Identity and Access Management (IAM) systems, the adaptive access control model based on the AI-driven risk-scoring has been shown to have great benefits to security and a lower occurrence of breaches. Conventional IAM solutions depend on fixed policies and fixed job descriptions, thus in most cases cannot be scaled to deal with the real-time dynamics of threats and users. Conversely, the AI-based system is constantly analyzing user activity, the health of the device, and threat intelligence and translates the results into access control, basing its decisions on a dynamic assessment of risk.

Surveyed 12 months' post-implementation of the AI-based IAM system, organizations have been reporting that they have reduced security incidences by 60 percent. This decrease can be greatly explained by the fact that the system is capable of detecting suspicious behavior fast and responding to an emerging threat in the form of dynamic access controls. As an example, real-time risk scores can be exploited through the system to implement Multi-Factor Authentication (MFA) or even terminate a session when in operation; it can be further used to block unauthorized access before much harm might be inflicted, therefore. Such a dynamic and situation-aware strategy will guarantee that access to important resources is highly restricted, which will considerably enhance the general security stance.

Table 2 illustrates the Risk Scoring Model components, including user behavior, device telemetry, external threat intelligence, and reference baseline. It shows how deviations from normal behavior, compromised devices, and external threats influence risk scores and IAM decisions like MFA or session termination.

Table 2: Effectiveness of the Risk-Adapted IAM Approach

Aspect	Conventional IAM	AI-Based Risk-Adapted IAM
Policy Basis	Fixed policies and job descriptions	Dynamic, risk-based policies
Threat Response	Static, reactive	Real-time detection and response to suspicious behavior
User Activity Monitoring	Limited or periodic	Continuous monitoring of user behavior, device health, and threat intelligence
Access Control	Role-based, fixed	Dynamic access control based on real-time risk scoring
Security Incidents	Higher occurrence of breaches	Reduced incidents (reported 60% reduction within 12 months post-implementation)
Risk Mitigation	Limited to static controls	MFA enforcement, session termination, and preemptive blocking of unauthorized access
Scalability	Poorly scalable to real-time dynamics	Adaptable to evolving threats and user behavior
Overall Security Posture	Reactive and rigid	Situation-aware and proactive, improving general security stance

9.2 Scalability and Operational Efficiency.

The introduction of the Attribute-Based Access Control (ABAC) model into the IAM system helps to make it even more scalable and efficient in its work [36]. In the case of ABAC, there is no longer any need to have a strict architecture of Role-Based Access Control (RBAC) where there is a need to create and maintain a large number of roles to handle access to various departments and projects. Rather, ABAC is implemented with dynamic user attributes (clearance levels, project membership, and so on) to implement access policies, which significantly simplifies scaling to a growing organization. Operational data display that the system allows bringing a new agency or dataset on board within 15 minutes using the ABAC model and without doing any tricky role reconfiguration. This is much better than the traditional RBAC systems, with allowing new roles or permissions or it may take hours or days. ABAC is flexible, efficient, which means it is suitable in the organizations that require switching to the new data-sharing necessities or changes rapidly, and its level of security remains high.

9.3 Usability and Adoption.

User feedback of the AI-driven IAM system has been positive and overwhelmingly so, particularly the increased productivity and collaboration. Conventional security operations, like fixed roles and hard access controls, usually lead to the tension between the security staff and final users. As an illustrative case, the employees are either compelled to seek manual access modification or take time to collaborate with other departments. Nevertheless, users do not have many disruptions because of the adaptive access control policies. Risk assessments can be conducted real time so that security features like MFA or access controls are not enabled unless needed without impacting on routine operations. Consequently, users claimed that the security was no longer impeding their effectiveness to work and cooperate with the employees in other departments. This has contributed to the growth of virtual adoption of the IAM system throughout its organization, as users now see security as facilitating their productivity and no longer as an obstacle to the same. Administratively, there has also been an easing of policy management with the replacement of fixed functions with tagging of resources. Security teams are not required to handle many roles/permissions any more, with dynamic session tags using real-time risk assessment. This streamlined methodology does not only enhance security but also decreases the level of administrative cumbersomeness that generally comes along with the effort of maintaining complex IAM systems.

Figure 1 illustrates how Internal Audit (IA) departments leverage generative AI, focusing on key areas such as fraud detection, continuous monitoring, audit risk assessment, regulatory compliance, audit reporting, and automated control testing to enhance audit processes and efficiency.



Figure 8: generative-ai-for-internal-audit

10. FUTURE DIRECTIONS

10.1 Integration of Device Posture: Adding Device Health Checks to Further Tighten Access Policies

With the cybersecurity landscape continuing to evolve, one of the next rational steps towards enhancing the AI-based IAM system is to introduce device posture into the risk-scoring and the process of granting access. The system currently takes into account a great variety of variables, including the use, device telemetry, and threat intelligence. However, the addition of device health verification to this flow can perhaps take further measures to limit access policies and increase security, in general. The term of device posture indicates the security level of the device that is utilized to access the corporate resources [37]. This involves issues like is the device up to date with security patches, has antivirus protection, is the device compromised or is the operating system out of date. Combining device health checks with the risk-scoring model would not only allow IAM systems to assess who is currently trying to access resources, but also the level of security on the device that is sending access tunnel to the resource in question.

As an illustration, a user trying to log-in using a device without the most recent security update installed or a device that is displaying indications of malware infection would automatically raise a higher risk score. On this basis, the system would be able to implement more stringent access controls, like forcing the user to update their device or do step-up (e.g., MFA) authentication first, then granting the user access. By incorporating the posture of the devices into the IAM system, more inclusive adaptive access controls will be facilitated ensuring that only users who have trusted identities and trusted devices are granted access. This will go a long way in ensuring that an organization can offset the risks that are related to the vulnerabilities of the devices.

10.2 Adaptation to cross-cloud/Multi-Agency Teaming Expand the Architecture to support Multi-Cloud-based environments and cross-agency collaboration.

The next critical area to focus the development in the future is the expansion of IAM architecture to be used in multi-cloud systems and provide cross-agency cooperation. The IAM system also needs to be adapted to support a distributed security model as the organizations grow to multi-cloud strategies (with resources moving in many cloud vendors, including AWS, Azure and Google Cloud). As it stands, lots of IAM systems are structured to operate within a single cloud setup, and it is hard to execute uniform access controls across various cloud platforms. The progress in the future should be directed at the development of the ability to implement user access and access control in the environment of various cloud services and third parties. This would enable the organizations to have a single perspective on the user access and behavior in all the platforms and the use of the principles of Zero Trust would be similar.

Cross-agency cooperation, especially that of a governmental or highly regulated sector, involves the need to exchange information and resources safely across various organizations or departments. The IAM system must be flexible to allow access control not only within one organization but also across various agencies or entities so that sharing of data and collaboration is carried out safely. With the ability to extend the architecture to allow multi-cloud and cross-agency collaboration, it will enable the organizations to implement the same consistent policies no matter

the nature of the cloud provider or other organizations. This will not only increase security avenue, but will also add flexibility that is needed in the current, collaborating workplace.

10.3 Next-level AI and ML: Use Reinforcement Learning to make Dynamic Changes in risk thresholds and access controls.

Reinforcement learning (RL) as a means of dynamically controlling risk thresholds and access control policies can be seen as one of the future trends of the AI-based IAM system [38]. Conventional machine learning models within IAM systems usually need to be humanly adjusted and preconfigured thresholds to determine when to invoke access controls such as MFA or terminating a session. This can however be automated by reinforcement learning whereby it allows the system to continue adapting without requiring human intervention. Reinforcement learning ensures that the system is able to learn with the actions and the environment and modify its access control policies in response to the feedback. As an example, the system can infer the lessons of previous incidents (including successful or unsuccessful attempts to detect breaches) to optimize the risk scoring threshold(s) based on various types of behavior. When a user has a habit of acting within the normal parameter, the system can reduce the threshold and activate security measures. Conversely, when the system recognizes that there are certain behaviors, which tend to correlate with the breaches, it can change thresholds to react even more violently to the predictable patterns.

The point of reinforcement learning that is important to apply in IAM is that it enables dynamic changes in real time [39]. This means that new behavior patterns and new threats can be automatically identified and access control policies changed accordingly to protect sensitive resources. This model is continuously enhanced through the feedback mechanism and access policies are maintained up to date in the face of an ever-evolving threat landscape. The reinforcement learning would further ensure the IAM system even more responsive to changing security risks to ensure that a control of access would be dynamic and proactive [40]. Such level of automation and dynamism will be necessary at the time when threats will be even more sophisticated and user behavior will be more complex.

Figure 9 schemes the flow between the artificial Intelligence (AI), Machine learning (ML), and its subfields. Machine Learning is a subfield of AI and it encompasses Reinforcement Learning, Supervised Learning, Unsupervised Learning and Deep Learning. Such methods are the staple of AI-based systems and allow creating models that can be trained and can adapt to new data as time goes by.




Figure 9: Artificial Intelligence (AI), Machine Learning (ML),

11. CONCLUSION

The artificial intelligence-driven risk scoring and adaptive access control system is a significant innovation in the Identity and Access Control (IAM) sphere. The system enhances the security, manual overhead reduction, and allows greater dynamism in the application of the concepts of Zero Trust by performing a continuous analysis of dynamic risk indications, which are based on the real-life properties such as user behavior, device health, and threat intelligence. This context-aware, real-time decision making will imply that an access control will not be static but will

evolve with the dynamic risk environment thus reducing the risk of false positives and thereby reducing the number of breaches.

The other piece of knowledge gained during the adoption of this system is the AI-related risk scoring that enhances the IAM security by flexing the decision on access control to the prevailing conditions. This kind of flexibility also facilitates timeliness when it comes to responding to organizations to new security threats and compromised accounts, thus, also making the process of discovering and responding to any security incident unpredictably shorter. The dynamic access control policy can also be applied to work hand in hand with the Attribute-Based Access Control (ABAC) which helps the organizations to downscale their IAM systems. Through dynamically motivated policies that the traditional role-based model is superseded by, businesses can provide a more challenging and scalable solution to the dilemma of role administration that arises due to the inflexibility of the classical model.

The impact of such AI based IAM system availability has been colossal particularly on the security and compliance. The after implementation statistics have seen that the system has resulted in reduction of security incidences by 60 percent. This is due to the fact that the system can detect cases of suspicious activities very fast and respond to them before they deteriorate and resist them in the majority of cases. In addition, the system has assisted significantly to achieve a greater degree of compliance with the canons of the Zero Trust, in order to ensure that the access to the vulnerable information is constantly monitored and controlled, based on the ongoing analysis of risks. This is a dynamic style that helps organizations to escape the reduction of efficiency in the operations as they strive to deliver on the rising security demands. The use of AI and machine learning countries on the IAM system has also led to the increased compliance with the regulatory requirements, such as HIPAA and FedRAMP. By permitting the adoption of these granular access controls and making sure that an ongoing analysis of the user activity and the security posture, business can be positioned at the side of meeting the regulatory requirements not to mention the improved overall security performance. It is the artificial intelligence (AI) based IAM systems that are transforming how organizations manage access, improves security and builds cooperation's. The great advantages of applying security measures according to the real-time needs are seen by the fact that the IAM policies became more dynamic and context-dependent as compared to being the same and role-based. Not only does this strategy boost security but efficiency in business as well, making businesses adequately equipped to address the existing security concerns and regulatory challenges.

References:

1. Mensah, F. (2024). Zero trust architecture: A comprehensive review of principles, implementation strategies, and future directions in enterprise cybersecurity. *International Journal of Academic and Industrial Research Innovations (IJAIRI)*, 10, 339-346. https://www.ijairit.com/?utm_source=pdf&utm_medium=edition&utm_campaign=OmAkSols&utm_term=V10I6-1452
2. Chadha, K. S. (2025). Machine learning-augmented ETL pipelines for fraud-resistant insurance claims processing. *IJDSML*. <https://www.academicpublishers.org/journals/index.php/ijdsml/article/view/5522/6451>
3. Chadha, K. S. (2025). Predictive risk modeling in P&C insurance using Guidewire DataHub and Power BI Embedded Analytics. *IJNS*. <https://www.academicpublishers.org/journals/index.php/ijns/article/view/5754>
4. Di Pietro, R., Salleras, X., Signorini, M., & Waisbard, E. (2018, June). A blockchain-based trust system for the internet of things. In *Proceedings of the 23rd ACM on symposium on access control models and technologies* (pp. 77-83). <https://doi.org/10.1145/3205977.3205993>
5. Chandra, R., Bansal, R., & Lulla, K. (2025). Benchmarking techniques for real-time evaluation of LLMs in production systems. *IJCESEN*. <https://ijcesen.com/index.php/ijcesen/article/view/3778/1063>
6. Greneche, N., Andres, F., Tanabe, S., Pester, A., Ali, H. H., Mahmoud, A. A., & Bascle, D. (2023, November). Leverage Data Security Policies Complexity for Users: An End-to-End Storage Service Management in the Cloud Based on ABAC Attributes. In *International Conference on Machine Learning for Networking* (pp. 199-217). Cham: Springer Nature Switzerland. https://link.springer.com/chapter/10.1007/978-3-031-59933-0_14
7. Dhanagari, M. R. (2025). Aerospike vs. traditional databases: Solving the speed vs. consistency dilemma. *IJCESEN*. <https://ijcesen.com/index.php/ijcesen/article/view/3780>
8. Dhanagari, M. R. (2025). Aerospike: The key to high-performance real-time data processing. *JISEM Journal*. <https://www.jisem-journal.com/index.php/journal/article/view/8894>
9. Ike, C. C., Ige, A. B., Oladosu, S. A., Adepoju, P. A., Amoo, O. O., & Afolabi, A. I. (2021). Redefining zero trust architecture in cloud networks: A conceptual shift towards granular, dynamic access control and policy enforcement. *Magna Scientia Advanced Research and Reviews*, 2(1), 074-086. <https://doi.org/10.30574/msarr.2021.2.1.0032>
10. Durgam, S. (2025). Peer benchmarking systems for RIA performance evaluation in investment technology. *Computer Fraud & Security*. <https://computerfraudsecurity.com/index.php/journal/article/view/785>
11. Durgam, S., & Nagaraj, V. (2025). Scalable data-driven engineering for high-performance computing & financial services. *IJISAE*. <https://www.ijisae.org/index.php/IJISAE/article/view/7914>

12. Gosangi, S. R. (2024). Secure and Scalable Single Sign-On Architecture for Large-Scale Enterprise Environments. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 7(3), 10466-10471. <https://ijrpem.com/index.php/IJRPETM/article/view/102>
13. Enugala, V. K. (2025). AI-powered crack propagation predictions. *JES*. <https://journal.esrgroups.org/jes/article/view/9203>
14. Enugala, V. K. (2025). Blockchain timestamping for unalterable concrete test logs. *TAJET*. <https://theamericanjournals.com/index.php/tajet/article/view/6346>
15. Gannavarapu, P. (2025). Cloud infrastructure management and automation. *AJT Journal*. <https://gprjournals.org/journals/index.php/ajt/article/view/356>
16. Gannavarapu, P. (2025). Performance optimization of hybrid Azure AD join across multi-forest deployments. *JISEM Journal*. <https://www.jisem-journal.com/index.php/journal/article/view/8897>
17. Gundla, S. R. (2024). AI-optimized Kubernetes scheduling: Node affinity for Java microservices. *SciPubHouse*.
18. <https://scipubhouse.com/home/international-journal-of-sustainability-and-innovation-in-engineering-ijsie/content/ijsie-2024/ai-optimized-kubernetes-scheduling-node-affinity-for-java-microservices/>
19. Gundla, S. R. (2025). AI-augmented testing: GitHub Copilot for JUnit/Mockito generation. *Computer Fraud & Security*. <https://computerfraudsecurity.com/index.php/journal/article/view/784>
20. G. Martín, A., Fernández-Isabel, A., Martín de Diego, I., & Beltrán, M. (2021). A survey for user behavior analysis based on machine learning techniques: current models and applications. *Applied Intelligence*, 51(8), 6029-6055. <https://link.springer.com/article/10.1007/s10489-020-02160-x>
21. Jha, A. C. (2025). AI-optimized spine-leaf fabrics: NVIDIA Quantum-2 vs. Cisco Nexus. *JISEM Journal*. <https://www.jisem-journal.com/index.php/journal/article/view/13315>
22. Jha, A. C. (2025). DWDM optimization: Ciena vs. ADVA for <50 ms global finances. *Utilitas Mathematica*. <https://utilitasmathematica.com/index.php/Index/article/view/2713>
23. Lulla, K. (2025). Pre-silicon DFT feedback loops: Enhancing GPU production efficiency. *IJCESEN*. <https://ijcesen.com/index.php/ijcesen/article/view/3778/1063>
24. Lulla, K. (2025). Python-based GPU testing pipelines: Enabling zero-failure production lines. *JISEM Journal*. <https://www.jisem-journal.com/index.php/journal/article/view/9419>
25. Rosenstock, J., Perkovic, V., Johansen, O. E., Cooper, M. E., Kahn, S. E., Marx, N., ... & Carmelina Investigators. (2019). Effect of linagliptin vs placebo on major cardiovascular events in adults with type 2 diabetes and high cardiovascular and renal risk: the CARMELINA randomized clinical trial. *Jama*, 321(1), 69-79. <https://jamanetwork.com/journals/jama/fullarticle/2714646>
26. Nagaraj, V. (2025). Automating test vector validation for silicon verification at scale. *IJEAS*. <https://gprjournals.org/journals/index.php/ijea/article/view/358>
27. Nagaraj, V. (2025). Ensuring low-power design verification in semiconductor architectures. *JISEM Journal*. <https://www.jisem-journal.com/index.php/journal/article/view/8903>
28. Rangu, S. (2025). Analyzing the impact of AI-powered call center automation on operational efficiency in healthcare. *JISEM Journal*. <https://www.jisem-journal.com/index.php/journal/article/view/8901>
29. Rangu, S. (2025). Enterprise digital transformation in financial services: Emerging trends and technologies. *Computer Fraud & Security*. <https://computerfraudsecurity.com/index.php/journal/article/view/786>
30. Samala, S. (2024). Real-time Jira analytics: Integrating JQL with Power BI/Snowflake for predictive agile metrics. *SciPubHouse*. <https://scipubhouse.com/home/international-journal-of-sustainability-and-innovation-in-engineering-ijsie/content/ijsie-2024/real-time-jira-analytics-integrating-jql-with-power-bi-snowflake-for-predictive-agile-metrics/>
31. <https://scipubhouse.com/home/international-journal-of-sustainability-and-innovation-in-engineering-ijsie/content/ijsie-2024/real-time-jira-analytics-integrating-jql-with-power-bi-snowflake-for-predictive-agile-metrics/>
32. Sundaramurthy, S. K., Ravichandran, N., Inaganti, A. C., & Muppalaneni, R. (2022). The future of enterprise automation: Integrating AI in cybersecurity, cloud operations, and workforce analytics. *Artificial Intelligence and Machine Learning Review*, 3(2), 1-15. <https://scipublication.com/index.php/AIMLR/article/view/136>
33. Samala, S. (2024). Real-time Jira analytics: Integrating JQL with Power BI/Snowflake for predictive agile metrics. *SciPubHouse*. <https://scipubhouse.com/home/international-journal-of-sustainability-and-innovation-in-engineering-ijsie/content/ijsie-2024/real-time-jira-analytics-integrating-jql-with-power-bi-snowflake-for-predictive-agile-metrics/>
34. <https://scipubhouse.com/home/international-journal-of-sustainability-and-innovation-in-engineering-ijsie/content/ijsie-2024/real-time-jira-analytics-integrating-jql-with-power-bi-snowflake-for-predictive-agile-metrics/>
35. Sayyed, Z. (2024). Implementing automation with BPMN for margin call workflow. *IRJERNET*. <https://irjernet.com/index.php/fecsit/article/view/171>
36. Sayyed, Z. (2025). Application-level scalable leader selection algorithm for distributed systems. *IJCESEN*. <https://ijcesen.com/index.php/ijcesen/article/view/3856/1152>
37. Vennamaneni, P. R. (2025). Building compliance-driven AI systems: Navigating IEC 62304 and PCI-DSS constraints. *IJNS*. <https://www.academicpublishers.org/journals/index.php/ijns/article/view/4305>
38. Vennamaneni, P. R. (2025). Real-time financial data processing using Apache Spark and Kafka. *IJDSML*. <https://www.academicpublishers.org/journals/index.php/ijdsml/article/view/4304>
39. Gupta, E., Sural, S., Vaidya, J., & Atluri, V. (2022). Enabling attribute-based access control in NoSQL databases. *IEEE transactions on emerging topics in computing*, 11(1), 208-223. <https://ieeexplore.ieee.org/abstract/document/9844984/>

40. Davis, B. D., Mason, J. C., & Anwar, M. (2020). Vulnerability studies and security postures of IoT devices: A smart home case study. *IEEE Internet of Things Journal*, 7(10), 10102-10110. <https://ieeexplore.ieee.org/abstract/document/9050664>
41. Olabanji, S. O., Olaniyi, O. O., Adigwe, C. S., Okunleye, O. J., & Oladoyinbo, T. O. (2024). AI for Identity and Access Management (IAM) in the cloud: Exploring the potential of artificial intelligence to improve user authentication, authorization, and access control within cloud-based systems. *Authorization, and Access Control within Cloud-Based Systems* (January 25, 2024). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4706726
42. Mullapudi, A., Lewis, M. J., Gruden, C. L., & Kerkez, B. (2020). Deep reinforcement learning for the real time control of stormwater systems. *Advances in water resources*, 140, 103600. <https://doi.org/10.1016/j.advwatres.2020.103600>
43. Ghadge, N. (2024). Enhancing threat detection in Identity and Access Management (IAM) systems. *International Journal of Science and Research Archive*, 11(2), 2050-2057. <https://doi.org/10.30574/ijsra.2024.11.2.0761>.