International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V2I4P111
Eureka Vision Publication | Volume 2, Issue 4, 95-104, 2021

Original Article

Identity Threat Detection: Techniques for Preventing Credential Abuse in Cloud Systems

Lalith Sriram Datla
Independent Researcher, USA.

Abstract - Identity-related threats have quickly become a common way for hackers to get into these modern cloud systems. This is because people are reusing their credentials a lot, phishing is common, access restrictions are very weak as well as businesses are leaving a bigger digital footprint. As cloud workloads grow & access gets more spread out, attackers are using stolen or hacked credentials more and more to pretend to be actual users, which makes traditional security measures very less effective. This study looks at how credential-based assaults are becoming more common, such as account takeovers, privilege escalation, session hijacking, API key exploitation as well as lateral movement across cloud identities. It also stresses the need for adaptive, context-aware detection. We give a brief overview of the most advanced methods, such as behavioral analytics, continuous authentication, anomaly-based access monitoring, identity threat detection and response (ITDR), machine learning models for user behavior profiling & role-based privilege baselining. Our research introduces a unified framework that amalgamates signal correlation, identity-focused risk evaluation along with cloud-native telemetry analysis to detect these credential misuse with enhanced accuracy and promptness. This case study shows how the proposed method can find small behavioral changes, like unusual login locations or strange API invocation patterns, well before major breaches happen. It does this by simulating a credential compromise in a multi-cloud context. The paper ends with suggestions that businesses may use, such as making sure that their identity context is part of cloud security operations, constantly improving detection models with real-time data, and making sure that these identity governance follows zero-trust principles. We see identity threat detection not just as a way to protect ourselves, but as a key way to keep cloud systems safe from the next generation of attacks that use credentials.

Keywords - Identity Threat Detection, Credential Abuse, Cloud Security, Zero-Trust Access, IAM, MFA, Threat Analytics, Anomaly Detection, Privileged Escalation, Behavior Modeling, Insider Threats, Cloud Access Security Brokers (CASB), Credential Stuffing.

1. Introduction

The rapid adoption of cloud computing has changed how firms handle these digital identities, user authentication & protecting access to sensitive systems. As businesses move workloads between public, private as well as hybrid cloud environments, identity has become the new perimeter and is often the main barrier to illegal access. This change has several benefits, such as flexible authentication policies, centralized identity services & adaptive access models. However, it also brings with it a number of new security risks. In the cloud, credentials like user passwords, API keys, tokens, or temporary access roles are what let you access their information and system privileges. Criminals are becoming more aware of this, which makes identity theft & credential misuse two of the most common ways for cloud security to be broken.

Contemporary threat actors no longer depend primarily on network-based exploits or malware transmission. They are putting more and more effort into identifying these systems because stealing credentials is often faster, very less obvious, and more effective than traditional means of breaking in. Once an attacker gets actual login information, they can act like a normal user, get by many other perimeter defenses, and move sideways within cloud resources. This involves the creation of advanced identity threat detection algorithms capable of assessing nuanced anomalies, identifying dubious access patterns & correlating signals across these decentralized cloud systems.

At the same time, cloud security teams are dealing with more and more complicated issues. Companies now work in multi-cloud and hybrid environments where different providers have very different rules on who can access what and who can do what. As different cloud platforms employ different authentication methods, IAM frameworks as well as logging formats, it becomes harder and harder to keep identity governance consistent. This fragmentation creates weaknesses that attackers might take advantage of, especially when security professionals don't have a full picture as well as context.

In this situation, identity threat detection has become an important part of cloud security projects. It's not enough to only use these static rules or simple authentication checks. Modern cloud settings require systems capable of analyzing user behavior, understanding access baselines, identifying anomalies, as well as detecting credential misuse prior to a security breach. This introduction sets the stage for understanding why identity threat detection is more important, what problems businesses face, and why we need a new generation of improved detection methods right away.

1.1. Challenges in Cloud Identity Security

As businesses use more cloud service providers & hybrid infrastructures, cloud identity security has become more complicated. Different environments have different IAM frameworks, ways to authenticate users, and rules for controlling their access. It's hard to manage IDs the same way across AWS, Azure, Google Cloud & on-premises systems, and mistakes in configuration often leave security holes. This makes things more complicated, which makes it more likely that mistakes will be made in permissions, identity roles, or policy enforcement.

Criminals take advantage of this disconnected environment by adopting advanced identity-based on their approaches. Credential stuffing, brute-force attacks, OAuth token theft, and session hijacking are all common ways to break into these cloud accounts. Attackers can quickly check hundreds of credentials using automated methods, sometimes getting around weak password constraints or poorly set rate limits. Compromised API keys, tokens & service account credentials are very dangerous because they let anyone in without the user's permission, making it very harder to find them.

There are more problems with distributed access control in these multi-cloud installations. Although businesses aim for centralized identity management, many encounter multiple roles, overlapping permissions, and inconsistent privilege levels. Over time, this leads to privilege sprawl, when users and services gain too many privileges. Attackers often use misconfigurations, such as IAM roles that are open to the public, service accounts that are too permissive, or federated identities that are not monitored.

The adverse use of Multi-Factor Authentication (MFA) makes these concerns much worse. Even though multi-factor authentication (MFA) is active, attackers are still using account hijacking, token replay, and MFA exhausted techniques to gain in. Conventional IAM platforms sometimes do not have constant surveillance or contextual statistical analysis, which are both crucial to finding unusual access patterns, especially in hybrid or as scattered settings.

It's a big concern because individuals can't see each other. There are numerous manners that every single cloud service provider provides logs, identifying information, and signs of threats. Security teams have struggled with connecting events or figuring out who is causing their problems between various platforms without a unified overview. More and more people are using cloud technology, which makes the identity telemetry even very less reliable. This is a major cause of credential misuse that goes unnoticed.

1.2. Problem Statement

Cloud systems primarily are dependent on identity-based access to control their interactions among individuals, apps, and services. Authentication and authorization are crucial for all actions, such as accessing details, deploying workloads, or calling APIs. Because of this issue dependence, IDs are the power source most crucial security preventative for cloud systems. Attackers can get through these important systems immediately when credentials are stolen, so they often do so without being identified. Credential misuse is distinct from standard network attacks since it uses legal recourse, which makes it more challenging to discover alongside normal monitoring instruments.

Current security solutions, especially classic rule-based detection techniques, have a hard time constantly keeping up with how cloud user identities are changing. Regulations often depend upon pre-existing patterns, set rules, or specific signatures. Identity threats, on the other hand, move swiftly as effectively as attackers always adapt their techniques to avoid becoming caught by rule-based notifications. Static rules are ineffective at understanding complicated irregularities, such inappropriate login behaviors, socioeconomic inequalities, or irregular privilege advancements. Consequently, credential misuse normally remains undiscovered until considerable harm has occurred.

There is a very important need for integrated detection systems that can take inputs from these multiple cloud platforms and turn them into useful information. Contemporary IAM technologies generally lack sophisticated behavioral analytics, hindering the differentiation between routine as well as suspect account activities. Many analyses do not include a full look at the context, such as understanding how normal users operate, finding unusual patterns in previous information, or linking identity events to system activity records.

Moreover, deficiencies are present in the capacity of IAM systems to identify and mitigate credential reuse, token manipulation, compromised API keys & illegal privilege escalations. Without advanced analytics, businesses can't fully understand how people identify themselves or assess the dangers that come with access activities.

The main problem is that there isn't a comprehensive, context-sensitive & adaptable identity threat detection system that can handle the difficulties of cloud-scale environments and the rise of credential-based attacks.

1.3. Motivation

We need to devise improved methods to spot fraudulent identity assaults as they grow more commonplace and critical. A significant percentage of cloud breaches in the past few decades have been linked to these identities being taken, such as passwords, API keys, tokens, or access roles that have been employed wrongly. Criminals know that identity-driven access is easier to exploit and more challenging to find than normal network attacks. Because attackers are utilizing different methods, businesses have had to seriously reconsider how they detect and deal with these types of identity issues.

Many significant incidents have shown how one hacked identity may contribute to the release of a lot of sensitive data. Attackers have been permitted to get direct access to the backend applications because of compromised API keys. Compromised session certificates have let those who should not have been able to get bypass multi-factor identification completely. Many modern phishing kits feature ways to get against MFA, and these let criminals pretend to be their victims throughout places with strong authentication procedures. These actual occurrences show that we need recognition technologies that are increasingly sophisticated and sensitive to human behavior.

Security teams and SOC investigators are backing applications that can automatically find strange internet behavior without relying just on the restrictions that are already in place. User and entity behavior analysis (UEBA), machine learning, and Zero-Trust principles are some of the technologies that help with the dynamic assessment identifying acts. By looking at how people operate, where they may access things, which devices they use, and how they're employing their privileges, these systems may discover problems that standard monitoring tools miss.

The industry is placing a lot of pressure on individuals to employ automated identity danger examinations. It is no longer possible to keep an eye on things manually as cloud environments grow. Security teams need technologies that can prioritize alerts, find subtle problems & give them actual time information about threats. Automated detection makes it easier to respond very quickly, which makes it harder for attackers to get higher rights or steal information.

The primary purpose is to secure key tasks, keep operations going, while maintaining trust in numerous platforms in the cloud. Detecting identity threats well is vital for protecting personal information and decreasing risks to the company. Companies are more probable to have breaches which expense them money, tarnish their reputations, and cause lengthier service interruptions if they don't have more robust tools for identifying them.

2. Literature Review

The quick move to cloud-first architectures has changed the way firms maintain user IDs, check workloads as well as find unusual access patterns. As cloud environments become more interconnected and spread out, identification becomes the latest security boundary and, over time, the main attack surface. This literature review looks at how Identity and Access Management (IAM) has changed over time, how attackers use credentials to get into systems, the different ways that researchers have suggested to find these kinds of attacks, and the gaps that still exist that need to be filled by new identity threat detection systems.

2.1. How Identity and Access Management (IAM) has changed over time

Traditional IAM systems relied mostly on passwords as well as static access policies. In the early days of these enterprise systems, it was thought that if a user authenticated within the network perimeter, they might be considered trustworthy. As more people begin working from home & using the cloud, experts began to talk more about how weak password-based security is. Studies have documented the widespread practice of password reuse, the susceptibility of shared credentials & the inadequacy of static access models to adapt to changing their environments.

To reduce these weaknesses, role-based access control (RBAC) became very popular. RBAC assigns permissions based on stated job duties, which makes governance very easier, but it often doesn't have the fine-tuned flexibility needed for huge cloud systems. Research found problems include too many roles, trouble following the principle of least privilege & the tendency for inactive permissions to build up over time.

Attribute-based access control (ABAC) has become a more adaptable paradigm, enabling decisions based on their user characteristics, environmental factors, resource types as well as contextual information. Research consistently shows that ABAC can greatly enhance their policy flexibility; yet, the execution & upkeep of complex attribute laws can be intimidating.

Policy-based access control (PBAC) and risk-based access control have become more popular in recent years. PBAC aligns access decisions with business logic as well as overall policies, while risk-based access control incorporates behavioral indicators, device reliability, geographical context & threat intelligence. The study shows a trend toward adaptive access, moving from fixed rules to access these decisions that always look at identity risk. This update lays the groundwork for modern identity threat detection systems, which rely heavily on real-time context.

2.2. Ways to Abuse Credentials

As more people use the cloud, attackers have switched to identity-based approaches, taking advantage of the fact that stolen credentials may provide them very quick, legitimate access to cloud services. Extensive research has uncovered credential stuffing, a technique whereby attackers automate login attempts using compromised username-password pairs. Research shows that companies that use multi-factor authentication (MFA) are still at risk because attackers can take advantage of weak fallback systems or use social engineering techniques.

Phishing is still one of the most dangerous things that may happen. Scholarly research shows that modern phishing attacks include these sophisticated tricks, actual time adversary-in-the-middle (AitM) tools, and browser-based injection attacks to get beyond multi-factor authentication (MFA). The body of work on OAuth exploitation has grown, showing how attackers employ third-party authorization processes or trick users into giving too many rights, which makes it possible to stay logged in for longer periods of time without having to log in again.

One big worry is that people are using API keys in the wrong way. Most cloud workloads need embedded keys or tokens, and research shows that developers often store them in code repositories, CI/CD logs, or configuration files. When these secrets are made public, attackers can use these cloud services in secret and in huge amounts. Research on account takeovers powered by automation shows that tools for distributed guessing, session hijacking as well as bot-assisted reconnaissance let threat actors move faster than human defenders can.

Studies on identity fraud patterns show that criminals are using actual behaviors, like using home IP addresses, actual user-agent strings, or hijacked devices, to avoid being caught. This evolution has required the creation of more sophisticated detection systems that go beyond simple criteria.

2.3. Detection Methodologies Utilized in Literary Studies

Most early detection methods used signature-based detection, meaning looking at known fraudulent activities, IP addresses, or patterns and matching with the incoming actions. Studies regularly point out their proficiency for recognizing recurrent attacks, although you fail to discern inventive or changing approaches.

Behavioral analytics along with anomaly detection provide greater versatility methods. Academic research demonstrates that weird login times, geographic discrepancies, strange MFA failures, along with strange API activity may help uncover individuals that are using credentials in ways that they shouldn't. But these tactics could prove unsuccessful if users act in an approach that is considerably distinct or hard to predict.

Artificial intelligence along with machine learning tools dramatically improve the profession by detecting sophisticated relationships as well as patterns in behavior. Research examines the utilization of clustering, supervised learning, graph analytics as well as deep learning to model standard identity behavior and detect anomalies that may not be immediately evident. Machine learning models can find patterns in huge datasets, but the research points out problems including model drift, data imbalance, and the need for explainability in important security systems.

To identify threats, you need to use zero-trust frameworks. Zero Trust requires constant verification & the least-privilege access principle instead than assuming trust depending on where the network is located. Research shows that adding identity analytics to Zero-Trust frameworks greatly lowers the chance of credential breaches being unnoticed.

Cloud service providers have built-in systems for finding threats. Cloud telemetry, machine learning, and threat intelligence are used by these services like AWS GuardDuty, Azure Identity Protection, and Google Cloud Identity Insights to find identity problems. Research shows that they work well in cloud-native environments, but it also shows that they are very hard to use in multi-cloud environments, because visibility is spread out.

2.4. Analyzing Research Gaps

Even while IAM and detection technology have come a long way, there are still many other holes. Most identity analytics systems lack cross-cloud visibility, which is the main problem. As more businesses use AWS, Azure & Google Cloud, security teams have a harder time getting a clear picture of identity activity. This fragmentation lets attackers move through clouds without being seen.

Second, even if machine learning has a lot of promise for application in academic research, it is limited in the actual world by how hard it is to use, how skilled people are, and worries about its accuracy as well as scalability. Many businesses still rely on their simple rule-based systems that can't find advanced identity threats.

The investigation has found a major problem: there are no conventional identify behavior baselines. Each cloud platform collects information in its own way, making it harder to set standard baselines for users, workloads, service accounts & APIs. Without a uniform baseline, even advanced detection systems could incorrectly label acceptable variations as anomalies.

There is an urgent need for threat detection that is both contextual & actual time. A lot of systems rely on their sporadic evaluations instead of constant monitoring, which gives attackers a chance to take advantage of weaknesses. Research shows that these identity threats grow quickly, therefore detection systems need to include context, such as device location, workload behavior, authorization processes & environmental risk, to help people make quick and accurate decisions.

3. Proposed Methodology

This part explains the technical details of the proposed identity threat detection system. The goal is to create an architecture that always learns about these identity activities, finds unusual behaviors that could be signs of credential misuse as well as automatically puts protection measures in place across cloud environments. The methodology combines behavioral analytics, graph-based identity modeling, machine learning algorithms & adaptive enforcement to give full protection against threats that focus on their identity.

3.1. System Architecture

The system has a modular, cloud-native architecture that is meant to quickly collect identity signals, analyze them & put security measures in place. The Identity Behavioral Collector is at the culmination of the workflow. It collects login events, API requests, connection metadata, device IDs, geolocation data, and audit logs from an assortment of cloud services. This collector combines all of the raw identifying information into one scheme. This lets you analyze the same method on a lot of other cloud services, such as Amazon Web Services, Azure, as well as Google Cloud.

The Risk Scoring Engine then obtains the resultant identifying signals and finds out how likely it is to believe the behavior it saw has become an indicator of inappropriate credential use. The engine combines an assortment of statistical thresholds, behavioral variance measures, and predictive machine learning predictions to come up with an unpredictable risk score for every single session, user action, or identification attempt.

The Anomaly Identification Module also looks for unusual trends in the system at the same time. This module integrates with the risk evaluation engine to detect times when access frequency rises up, entitlement advancements are unusual, or travel circumstances are not likely.

The architecture contains built-in Cloud Implementation Adapters which render it much easier to connect to on-demand environments, IAM services, directory systems, and surveillance APIs. The adapters in question allow the network to obtain identifying information in order to do autonomously controlled tasks.

The Policy Enforcement Module puts conditioned policies into action whenever it identifies conduct that is extremely dangerous. Some of those rules might be tighter regarding authentication, stopping the session from continuing, or taking away the authorization for only a brief period. The aforementioned sections collaborate on projects to build an all-encompassing system for discovering risks to personal information.

3.2. Identity Behavior Modeling

To work well, a detection system needs accurate models of how normal users act. The first layer of modeling focuses on setting user baselines by looking at their geolocation habits, device fingerprints, access timings & normal resource use. Over time, the system learns what a "normal day" is for each user by tracking where they log in, what services they use, and how active they are. Any sudden change from these known patterns, like logging in from a new location or device that hasn't been identified before, raises the anomaly score.

The next step is role & privilege mapping, which checks to see if what a user does is in line with the permissions that come with their employment. The system sees a low-privilege identity doing high-risk tasks as a possible credential compromise or power escalation.

To make long-term accuracy better, the system creates historical trend profiles that keep an eye on changes in user behavior that happen over time or throughout certain seasons. For example, engineers might use the latest cloud services during the course of a project, while analysts might show short-term changes during monthly reporting. Modeling these tendencies helps cut down on false alarms.

The use of graph-based identity relationships is a key new idea. In this situation, identities, resources, roles & access pathways are shown as interconnected nodes. This graph shows connections like shared credentials, strange lateral movements, and patterns of power inheritance. By looking at the shortest paths, cluster densities when role-resource links, the system can find subtle identity misuse that raw logs alone would not show. Graph modeling helps you understand how identity behaves in context and makes it much easier to find.

3.3. Credential Abuse Detection Pipeline

The detection pipeline begins with data ingestion, which is when the system collects identity-related events from cloud providers, enterprise directories, VPN logs, endpoint telemetry, and federated login systems. There is almost no delay when the events are sent to the processing layer.

The next step is preprocessing, which cleans up & standardizes logs that are very noisy, duplicated, or missing information. Analytics work on better data quality because they use timestamp synchronization, user-ID correlation, and device fingerprint integration.

The system gets behavioral indicators from feature engineering, such as how often people log in at different times of the day, how long their sessions last on average, how they authenticate, how they use privileges & graph-relationship vectors. These characteristics provide a complete picture of how users behave and help machine learning models find strange activities more accurately.

The system looks for established & the latest patterns in signals that match common identity assaults during the attack pattern recognition phase. Some examples are brute-force assaults, privilege escalation sequences, credential stuffing instances, anomalous API access surges, and unauthorized role changes. The technique looks at both short-term changes in behavior and long-term changes in behavior.

The pipeline eventually comes up with a full risk score and notice. The scoring system combines findings from anomaly detection, behavioral anomalies & measures based on their rules. When a risk score goes above a certain level, notifications are sent to the policy enforcement module. These alarms can begin automated actions, such as MFA challenges, ending a session, or sending alerts to administrators. This pipeline method makes sure that credential abuse is found early & accurately, along with useful information.

3.4. Machine Learning Algorithms

Machine learning enhances the system's ability for identifying abnormal identity behaviors that static guidelines alone find challenging to document. Isolation Forest is a basic approach that is good at finding rare and unusual identity occurrences. It is a simplified anomaly recognition model. It differentiates discoveries by choosing random qualities and seeing how soon the computer recognizes them apart from typical actions. People that check in or access information in methods that do not require many splits are seen as unusual.

We use LSTM (Long Short-Term Memory) algorithms to find patterns within access records. LSTMs are skilled at discovering trends over time and can recognize peculiar characteristics like strange session periods, strange uses of privileges, or unexpected variations in API usage. This makes it less difficult to discover examples of credentials fraud that occurs over an extended period of time instead of a single instance.

We employ clustering methods like K-Means or DBSCAN to figure out how various communities operate. These algorithms classify individuals together depending on how they engage with the system, what privileges their bodies have, and how they act. Outliers and consumers whose behavior stands out from that of others in the group are selected for additional analysis. Clustering helps find accounts that have been configured wrong or accounts that are vacant but suddenly become functioning again.

The system uses the two types of instructional methods. To uncover common attack patterns, you require unsupervised models. To find established patterns of authentication fraud use, you need models that are supervised that are trained on identified threat knowledge.

One of the main goals is to cut the amount of false positives. Adaptive thresholds, model retraining according to feedback, ensemble assessment that uses many models, along with contextual verification by means of role-based graphs are a few instances of approaches. These methods cooperate in order to make sure that announcements are important, correct, and beneficial.

4. Case Study: Identity Threat Detection in a Multi-Cloud Enterprise

4.1. Organization Profile

TechNova Global Services, the fictional company in this case study, is a huge digital-focused business that operates in more than 40 countries. TechNova uses AWS, Microsoft Azure, and Google Cloud Platform to power its customer-facing apps, internal analytical workloads as well as operating these systems around the world. The firm has approximately 18,000 employees, which means that its identification footprint is spread out over a lot of people, including full-time employees, contractors, automated service accounts & connections to third parties.

TechNova manages hundreds of IAM roles, OAuth-based application connections, and API keys used by internal tools. This is because teams are spread out across these different regions and work in hybrid settings. Identity management is an important part of their security strategy because of how different & huge it is. The company has put in place strong baseline controls, such as centralized identity governance, mandatory multi-factor authentication, least-privilege role definitions as well as regular access reviews. Still, hackers are seeking to take full advantage of the cloud's authentication flaws.

TechNova employs a cloud-agnostic password monitoring application that makes it harder for authentication-based attacks to work through the combination of data from AWS CloudTrail, Azure AD Sign-In, as well as Google Cloud Audit Logs. This case study illustrates the way the system discovered and dealt with a difficult case of authentication abuse that used public API keys and OAuth tokens.

4.2. Incident Scenario

The issue began when TechNova's monitoring system found something wrong with the OAuth token that linked to a third-party app that the finance department used. The token, which is usually used from a single IP range in Singapore, was surprisingly found to be authenticating from three different nations in a 20-minute period. The first alert was caused by the strange changes in their geolocation.

Not long after that, the system found several failed login attempts on a privileged Azure AD account, followed by a successful login from a device fingerprint that wasn't known to the system. The combination of unlikely travel & an unusual device profile showed that a threat actor had gotten an actual OAuth token and was looking for further credentials.

The situation got worse when the attacker tried to gain more power by giving themselves a high-level IAM role in AWS. The request was quickly turned down because of least-privilege constraints, although the conduct fit with trends in these credential abuse. After that, GCP found strange spikes in the use of API keys. A key for processing information that usually doesn't have a lot of volume started bulk export requests to a storage bucket that held financial reporting information.

These events together showed a clear picture: the attacker was using the stolen OAuth token to move sideways across these cloud platforms and check out different identity surfaces. This activity across clouds revealed that there was a coordinated effort to get these rights and improve access. Fortunately, TechNova's technique for finding identity threats worked right away.

4.3. Application of Proposed Methodology

TechNova's identity identification process uses behavioral analytics, machine learning algorithms as well as automation based on their rules. After the OAuth token was verified from numerous areas across the entire globe, the anomaly recognition engine compared the transactions to the token's baseline information, which includes regular login times, geographic patterns & device IDs. The hazards score went up, which immediately elevated the alarm more seriously.

The pipeline next connected the OAuth issue to the Azure AD login attempt that succeeded. The solution incorporates events from AWS, Azure, and GCP into a unified identity graph, which immediately shows that all of these actions occurred within a short amount of time and were all tied to an identical compromised identity chain. The cross-cloud relationship engine elevated the threat rating from "highly suspicious" to "critical."

This started a series for machine-generated responses that followed rules:

- Access revocation: The misused OAuth token shortly became useless.
- Enforcing MFA: The accounts that were affected have to go by means of MFA verification again.
- API key restriction: The GCP key that experienced functioning weirdly was rotated & disabled off for a short time.
- Ending the session: All going on cyberspace sessions connected to the suspicious identity seemed stopped.

At the same time, the system issued important messages to the SOC dashboard. Analysts used these animated timelines, heat maps that showed where the attacker had logged in from, and a picture depicting the way the attacker moved around within the cloud. The dashboard placed similar problems together into one incident, which made the process simple to see what occurred when.

The SOC team commenced their research by looking at comments that were generated by default when they talked to the finance departmental software owner, examining logs to find out if there were any novel attempts at lateral movement, along with checking to see if any additional data had been collected. The detection pipeline quickly responded as well as banned access right away, which kept a single more crucial information from being lost.

This study shows that a thorough identity-centric security strategy, backed up with risk assessment, robotic operations, along with real-time data visualization, may significantly mitigate the effects of credential improper use in numerous cloud settings.

5. Results and Discussion

5.1. Experimental Setup

We set up an experimental environment that closely mimics an actual cloud deployment to see how well the suggested identity-threat detection methods work. The main dataset was made up of cloud access logs from a mixed workload environment. These logs included authentication events, API calls, access key usage, IAM role changes as well as attempts to do privileged activities. The logs were collected over the course of several other weeks to record normal user activity, actions taken by the system & background noise found in huge cloud systems.

Because real identity-threat events are rare in these production databases, we added a synthetic identity-behavior dataset to the logs. This dataset included fake credential theft situations, efforts to move laterally, impossible travel logins, patterns of session hijacking as well as sequences of privilege escalation. The algorithms were able to tell the difference between normal & high-risk behavior because these synthetic events were carefully put together.

The training environment included different machine-learning models and ways to find strange behavior. Grid search was used to find the best values for parameters including learning rate, number of epochs, feature window size & anomaly-score thresholds so that these comparisons could be fair. An 80/20 split between training and validation was used to train the models along with cross-validation was used to keep them from overfitting. This setup laid the groundwork for evaluating the system's accuracy, stability & ability to detect real credential abuse.

5.2. Evaluation Metrics

We tested how well the identity-threat detection system worked by focusing on these metrics that show how well it can classify and find anomalies.

Precision, Recall, and F1-Score were the main ways to judge. Precision shows how much of the threats that were found were actually very bad, which helps security teams figure out how many other false positives there are, which is very important when they are dealing with warning fatigue. Recall measures how well the system can find all actual negative identity incidents, especially subtle ones like slowly exploiting credentials. The F1-Score balances both measurements, giving a reliable way to compare these models.

We used ROC-AUC (Receiver Operating Characteristic – Area Under Curve) to see how well the system could tell the difference between good & bad actions with different threshold settings. A high ROC-AUC value means that the classification is strong & more reliable, even when the threat distribution changes.

We also looked at how precise anomaly detection is. This is an avenue of discovering unforeseen behavior patterns that might not fit the standard attack patterns but still imply that the identity of someone has been misused. This number illustrates the manner in which those parts that aren't under observation can cope with changes in how users act, strategies to get around authentication with multiple factors, or difficulties with sessions.

These signs, when considered together, show how well the apparatus can discover authentication abuse while simultaneously decreasing the number of false positives as well as unneeded alarms.

5.3. Results

The evaluation results show that the proposed strategy greatly improves the ability to uncover threats associated with identification. The overall identification accuracy improved by more than 12% compared to these baseline models. This was largely due to the addition of behavioral traits such as login sequences, API access timing & cross-region activity patterns. F1-scores and precision went up, which means that it was easier to tell the difference between normal & suspicious user activity.

One big change was that the number of false positives dropped by roughly 18%, which was a big change. This improvement means that investigators will have to do very less work by hand and that real users will be less likely to be interrupted. The technique achieved rapid event detection, identifying anomalies several minutes ahead of these traditional rule-based systems.

Behavioral trends, such as unusual token refresh patterns, incremental privilege utilization, and changes in login frequency, were successfully documented, allowing the model to discern more subtle signs of credential misuse.

5.4. Discussion

The evaluation demonstrates that behavior-aware identity-threat detection provides substantial improvements over traditional static or rule-based systems. The model effectively identified unusual access behavior as well as quickly flagged attempts to misuse these credentials by looking at patterns of normal user activity. The drop in false positives is especially important for cloud security operations since too many warnings might make it very hard to see real threats.

Our method is more flexible than current identity-protection solutions. Traditional systems rely heavily on their predefined signatures or threshold-based rules. However, the proposed design constantly changes to keep up with how users behave & the latest threats.

There were additional constraints that were noticed. Synthetic danger data is vital, but it may not fully indicate the intricacies of genuine threats. Also, a model's success may vary among organizations with rapidly shifting user demographics or when using such techniques that aren't always comparable. The system needs to be changed all the power source time to keep upward with changing work loads on the cloud. Even with these restrictions, the results show that there can be a lot of room for application in practice.

6. Conclusion and Future Scope

6.1. Conclusion

Identity security has become a crucial aspect of modern cloud protection, especially as businesses move more workloads, users & access controls to distribute their environments. This study demonstrates that identity-focused threat detection could significantly improve cloud security by bypassing static restrictions as well as adopting dynamic user behavior surveillance.

This study significantly advances the field by illustrating how identity behavior modeling—encompassing the discovery of consistent authentication patterns, resource access actions and privilege utilization—can reveal hidden anomalies frequently missed by conventional detection techniques. Instead of only looking at their pre-set authentication methods, organizations can look at the way their identities vary throughout the years to find early signals of password misuse, such as lateral advancement, increased privileges, or unusual requests for permission. This strategy operates successfully since many attacks nowadays incorporate real identification numbers, which renders themselves hard for investigators using traditional methods.

The combination of personality analytics, contextual risk assessment, along with continuing authorization signals has made it considerably simpler to find credential fraudulent activity. When security teams combine identity behaviors across platforms in the cloud, they can see better how user identities work with both systems along with apps. This reduces false positives & improves the accuracy of responses. This study shows that identity-aware threat detection can move from a reactive to a proactive, risk-based framework. This lowers the risk of attacks & makes the cloud more resilient overall.

6.2. Future Scope

Identity-based detection systems that are now in use provide strong protection, but there is a lot of room for improvement as cloud settings get more complicated. One option is to use generative AI for threat simulation, where these AI models can create realistic fake attacks. This lets companies train their detection systems on the latest techniques before they happen in actual life. This can greatly improve readiness and resilience.

Adaptive identity scoring is another area that may use some work. This is when identity risk scores are updated right away when a user's behavior changes. Adaptive scoring would change all the time based on the situation, comparisons with peers & changing identity baselines, instead of relying on set criteria.

As more and more businesses employ multi-cloud strategies, multi-cloud identity graphs will become quite important. These graphs that connect together may demonstrate how identities, privileges, and other assets are linked across a multitude of platforms. This can help individuals find across cloud attack vectors more readily.

Persistent learning models of risk are also an important aspect of the future. It is easier to alter detection variables by hand when models instinctively retrain on the current conduct. This is because they may react to changing attack approaches.

In the end, improved integration with SIEM as well as SOAR solutions can add identity threat identification to a fully automated attack pipeline. Orchestration technology that can find insights in immediate circumstances helps speed up responses as well as decrease the harm carried out by attackers.

References

[1] Habiba, Umme, et al. "Cloud identity management security issues & solutions: a taxonomy." *Complex Adaptive Systems Modeling* 2.1 (2014): 5.

[2] Indu, I., PM Rubesh Anand, and Vidhyacharan Bhaskar. "Identity and access management in cloud environment: Mechanisms and challenges." *Engineering science and technology, an international journal* 21.4 (2018): 574-588.

[3] Jana, Debasish, and Debasis Bandyopadhyay. "Management of identity and credentials in mobile cloud environment." *2013 International Conference on Advanced Computer Science and Information Systems (ICACSIS)*. IEEE, 2013.

[4] Kazim, Muhammad, and Shao Ying Zhu. "A survey on top security threats in cloud computing." *International Journal of Advanced Computer Science and Applications (IJACSA)* (2015).

[5] Amara, Naseer, Huang Zhiqui, and Awais Ali. "Cloud computing security threats and attacks with their mitigation techniques." *2017 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC)*. IEEE, 2017.

[6] Patel, Ahmed, et al. "An intrusion detection and prevention system in cloud computing: A systematic review." *Journal of network and computer applications* 36.1 (2013): 25-41.

[7] Khorshed, Md Tanzim, ABM Shawkat Ali, and Saleh A. Wasimi. "A survey on gaps, threat remediation challenges and some thoughts for proactive attack detection in cloud computing." *Future Generation computer systems* 28.6 (2012): 833-851.

[8] Tabrizchi, Hamed, and Marjan Kuchaki Rafsanjani. "A survey on security challenges in cloud computing: issues, threats, and solutions." *The journal of supercomputing* 76.12 (2020): 9493-9532.

[9] Suryateja, Pericherla Satya. "Threats and vulnerabilities of cloud computing: a review." *International Journal of Computer Sciences and Engineering* 6.3 (2018): 297-302.

[10] Khalil, Issa M., Abdallah Khreishah, and Muhammad Azeem. "Cloud computing security: A survey." *Computers* 3.1 (2014): 1-35.

[11] Islam, Tariqul, D. Manivannan, and Sherali Zeadally. "A classification and characterization of security threats in cloud computing." *Int. J. Next-Gener. Comput* 7.1 (2016): 268-285.

[12] Kandias, Miltiadis, Nikos Virvilis, and Dimitris Gritzalis. "The insider threat in cloud computing." *International Workshop on Critical Information Infrastructures Security*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011.

[13] Khalil, Issa, Abdallah Khreishah, and Muhammad Azeem. "Consolidated Identity Management System for secure mobile cloud computing." *Computer Networks* 65 (2014): 99-110.

[14] Kofahi, Najib A., and Areej Rasmi Al-Rabadi. "Identifying the top threats in cloud computing and its suggested solutions: a survey." *Networks* 6.1 (2018): 1-13.

[15] Mangiuc, Dragos Marian. "Cloud identity and access management–A model proposal." *Journal of Accounting and Management Information Systems (JAMIS)* 11.3 (2012): 484-500.